


UKRAS20 CONFERENCE:



UK-RAS
NETWORK
ROBOTICS & AUTONOMOUS SYSTEMS

“Robots into the real world” PROCEEDINGS

UKRAS20 Conference Proceedings: 17th April 2020



EPSRC Centre for Doctoral Training in Agri-Food Robotics



UNIVERSITY OF
LINCOLN



Engineering and
Physical Sciences
Research Council

Table of Contents

Introduction

<i>UKRAS20: the 3rd UK Robotics and Autonomous Systems Conference.</i> C. Fox, T. Duckett, A. Richards	1
--	---

Research Papers

<i>Enhancing Unsupervised Natural Language Grounding through Explicit Teaching.</i> O. Roesler	3
<i>A Study Assessing the Impact of Task Duration on Performance Metrics for Multi-Robot Teams.</i> G. Miyauchi	6
<i>An Incremental Learning Approach for Physical Human-Robot Collaboration.</i> A. Buerkle	9
<i>Plastic ‘personalities’ for effective field swarms.</i> E. R. Hunt	12
<i>Reliability-Aware Multi-UAV Coverage Path Planning Using Integer Linear Programming.</i> M. Li	15
<i>Towards Intention Recognition for Human-Interacting Agricultural Robots,</i> A. Gabriel	18
<i>Feasibility Study of In-Field Phenotypic Trait Extraction for Robotic Soft-Fruit Operations.</i> R. Kirk	21
<i>Human-in-the-Loop Adaptation and Reuse of Robot Assistance Policies for Ambient Assisted Living,</i> R. Smith	24
<i>Automated Topological Mapping for Agricultural Robots.</i> W. Mandil	27
<i>Trajectory Tracking and Control of Multiple Robot Arms on a Free-Floating Spacecraft for Debris Removal.</i> A. Babu	30
<i>Improving Quadrupedal Locomotion on Granular Material Using Genetic Algorithm.</i> J. Hulas	33
<i>Analysis of two-wheeled robot morphology for a slope environment/</i> R. Woolley	35
<i>The Goods and Bads in Dyadic Co-Manipulation: Identifying Conflict-Driven Interaction Behaviours in Human-Human Collaboration.</i> I. Issak	37
<i>Does Expression of Grounded Affect in a Hexapod Robot Elicit More Prosocial Responses?</i> L. Hickton	40
<i>Optimising Soft Fin Ray Robotic Fingers using Finite Element Analysis to Reduce Object Slippage,</i> J. Emerson	43

Poster Papers

<i>A Cable-based Manipulator for Chemistry Labs.</i> L. Manes, S. Fichera, D. Marquez-Gamez, A. I. Cooper and P. Paoletti	46
<i>An Open Source Seeding Agri-Robot.</i> H. Rogers and C. Fox	48
<i>Investigating PID Control for Station Keeping ROVs.</i> K. L. Walker, A. A. Stokes, A. Kiprakis and F. Giorgio-Serchi	51
<i>Control of hydraulically-actuated manipulators with dead-band and time-delay uncertainties.</i> O. Albrecht, M. Bandala, S. Monk and C. J. Taylor	54
<i>A Recurrent Encoder-Decoder Network Architecture for Task Recognition and Motion Prediction in Human-Robot Collaboration based on Skeleton Data.</i> D. Zhang, S. McLoone and V. Ngo	56
<i>Topological Robot Localization in a Pipe Network.</i> R. Worley and S. Anderson	59
<i>The role of ‘urban living labs’ in ‘real-world testing’ robotics and autonomous systems.</i> R. Macrorie, M. Kovacic, A. Lockhart, S. Marvin and A. While	61
<i>RICA: Robocentric Indoor Crowd Analysis Dataset.</i> V. Schmuck and O. Celiktutan	63
<i>Robots Producing Their Own Hierarchies with DOSA; The Dependency-Oriented Structure Architect.</i> B. Hawker and R. K. Moore	66
<i>Development of Tools and Methods for Autonomous Fixed-wing UAV Research.</i> M. Khan, L. Alboul and J. Potts	69
<i>A review of manufacturing systems for introducing collaborative robots.</i> R. Ma, J. Chen and O. John	71
<i>Smart Suitcase Implementation Using Fuzzy Logic and Deep Learning.</i> K. T. Ly	74
<i>Game Theory For Self-Driving Cars.</i> F. Camara and C. Fox	77
<i>Deep learning for robotic strawberry harvesting.</i> X. Li, C. Fox and S. Coutts	80
<i>3D Printed Variable Infill Soft Fingers for the SIMPA Prosthetic Arm.</i> D. De Barrie, K. Goher and K. Elgeneidy	83
<i>Task Delegation and Architecture for Autonomous Excavators.</i> J. Rankin, L. Justham, Y. M. Goh and J. Morley	86
<i>Understanding human responses to errors in a collaborative human-robot selective harvesting task.</i> Z. Huang, G. Miyauchi, A. S. Gomez, Richie Bird, A. S. Kalsi, C. Jansen, Z. Liu, S. Parsons and E. Sklar	89

<i>A Point Cloud Semantic Segmentation Framework for Embedded Systems in Agricultural Robots.</i> G.Storey, L. Jiang and Q. Meng	92
<i>A Robotic Environment for Cognitive Assessment.</i> S. MacLeod, M. P. Rodriguez and M. Dragone	95
<i>Robotic ignition systems for oil fields.</i> M. Jabbar, L. Alboul and I. Ahmed Abed	97
<i>Towards the development of a deposition technology for an automated rail repair system.</i> D. De Becker, J. Dobrzanski, M. Goh and L. Justham	100
<i>An Automatic Design Tool for Soft Robotics.</i> O. Bridgewater-Smith, G. Maurizi, S. Fichera, D. Marquez, A. Cooper and P. Paoletti	103
<i>Adaptive Manipulator Control using Active Inference with Precision Learning.</i> M. Baioumy, M. Mattamala, P. Duckworth, B. Lacerda and N. Hawes	105
<i>Towards Human-Chatbot Interaction: A Virtual Assistant for the Ramp-up Process.</i> M. Zimmer, A. Al-Yacoub, P. Ferreira and N. Lohse	108
<i>Power-aware Fusion of Visual and Wheel Odometry for Mobile Platforms.</i> M. Malinowski, A. Richards and M. Woods	111
<i>Debiasing of position estimations of UWB-based TDoA indoor positioning system.</i> P. Grasso and M. S. Innocente	114
<i>The Augmented Agronomist Pipeline and Time Series Forecasting.</i> G. Onoufriou, M. Hanheide and G. Leontidis.	117
<i>BeetleBot: A Multi-Purpose AI-Driven Mobile Robot for Realistic Environments.</i> A. Nguyen, E. Tjiputra and Q. Tran	120
<i>Determining shape of strawberry crops with spherical harmonics.</i> J. Le Louedec and G. Cielniak	122
<i>Towards Gamification of the Ramp-up Process for Industry 4.0.</i> B. Foley, P. Ferreira and M. Zimmer	125
<i>Modelling variable communication signal strength for experiments with multi-robot teams.</i> T. Zhivkov and E. Sklar	128
<i>Towards bio-inspired fruit detection for agriculture.</i> Z. Hobbs, T. Duckett, S. Pearson and M. Mangan	131
<i>Towards co-creative drawing with a robot.</i> C. Jansen and E. Sklar	134
<i>Robotic Untangling of Herbs with Parallel Grippers.</i> P. Ray and M. J. Howard	137
<i>Towards Insect Inspired Visual Sensors for Robots.</i> B. Millward, S. Maddock and M. Mangan	140

<i>Vision Navigation System to Manoeuvre Unmanned Aerial Vehicle (UAV).</i> P. T. Saleh, J. Penders and L. Alboul	142
<i>Experimental Analysis of Soft Vacuum Cups for Automated Mushroom Picking.</i> H. Husain and K. Elgeneidy	145
<i>Mapping the impact of faults in a multi-robot team.</i> D. Zhang and E. Sklar	148
<i>Image Pre-processing, Triplet Learning and Route Navigation.</i> W. Smith, R. Fisher and Y. Petillot	151
<i>Unsupervised Anomaly Detection for Safe Robot Operations.</i> P. Somaiya, M. Hanheide and G. Cielniak	154
<i>Intelligent Service Robots to enter the Hospitality Industry: Job Termination or Gospel to Waiters.</i> J. Yang and E. Chew.	157

UKRAS20: the 3rd UK Robotics and Autonomous Systems Conference

Charles Fox

Tom Duckett

Arthur Richards

<https://doi.org/10.31256/Vh8Se9V>

On behalf of the Organising and Advisory Committee we take great pleasure in welcoming students, researchers and experts in Robotics to UKRAS20, the 3rd UK-RAS Conference for PhD Students and Early-Career Researchers, organised by the EPSRC UK-RAS network in collaboration with the AgriFoRwArdS and FARSCOPE Centres for Doctoral Training.

I. CONFERENCE AIMS

The conference is specifically for PhD students and early-career robotics researchers of the UK-RAS Network and will foster progress in the field of robotics research at what promises to be a very comprehensive and exciting meeting. The aim of the UKRAS20 conference is to promote quality research, networking, and community building for PhD students and practitioners at the frontier of science and technology in intelligent robots and systems, by discussing the latest advancements in this fast growing and exciting field.

II. TOPICS

This year's theme is 'Robots into the Real World', exploring how robotics can make a positive difference to societal challenges, from fundamental enabling technologies to real-world applications, such as working in challenging and extreme environments; enabling healthy / independent living; ensuring safe, efficient transport; developing next-generation manufacturing; feeding a growing population and ensuring a safe environment for the future. Presentations have been grouped into three sessions within this theme:

Artificial Intelligence and Robotics. Keynote speaker Ingmar Posner (Oxford) will discuss Robots Thinking Fast and Slow. The oral presentation topics are Enhancing Unsupervised Natural Language Grounding through Explicit Teaching [1], Enhancing Unsupervised Natural Language Grounding through Explicit Teaching [2], An incremental learning approach for physical Human-Robot Collaboration [3], Plastic 'personalities' for effective field swarm [4], and Reliability-Aware Multi-UAV Coverage Path Planning Using Integer Linear Programming [5].

Field and Service Robotics. Keynote speaker Fumiya Iida (Cambridge) will discuss Turning Soft Materials into Intelligent Machines. The oral presentation topics are Towards Intention Recognition for Human-Interacting Agricultural Robots [6], Feasibility Study of In-Field Phenotypic Trait Extraction

for Robotic Soft-Fruit Operations [7], Enabling Deep Personalisation for a Heterogeneous Ambient Assisted Living Landscape [8], Automated Topological Mapping for Agricultural Robots [9], and Trajectory Tracking and Control of Multiple Robot Arms on a Free-Floating Spacecraft for Debris Removal [10].

Novel and Emerging Robotics Technologies. Keynote speaker Adam Stokes (Edinburgh) will discuss Biologically Inspired Robotic Systems for Extreme Environments. The oral presentation topics are Biologically Inspired Robotic Systems for Extreme Environments [11], Analysis of two wheeled robot morphology for a slope environment [12], The Goods and Bads in Dyadic Co-Manipulation: Identifying Conflict-Driven Interaction Behaviours in Human-Human Collaboration [13], Expression of Grounded Affect in a Hexapod Robot [14], and Optimising Soft Fin Ray Robotic Fingers using Finite Element Analysis to Reduce Object Slippage [15].

These sessions and topics present an overview of current areas of interest across the UK robotics community in 2020.

III. SELECTION PROCESS

Accepted authors from previous UKRAS conferences were invited to review, with additional reviewers from the host institution, University of Lincoln, having one or more previous publications in TAROS, IROS or ICRA. Reviewers were free to delegate reviews to any others having publications in these venues. All submissions have been reviewed by two or three reviewers, scoring between 3 (strong accept) and -3 (strong reject), with 0 as borderline. The conference aims to be inclusive so all papers with average scores of 0 or greater have been accepted. Additional reviewers were invited where reviewer opinions differed strongly between accept and reject. Reviewers were instructed that papers should contain some novelty, such as presenting new results and/or new contextualisations of previous results, such as reviewing previous work to present one's whole PhD or research project to the UK community as a whole. Reviewers from the host institution were not assigned to papers from the host institution. All other review assignment is random, as a conference aim is for the whole UK robotics community to understand each others work in a single track event. The 15 highest scoring papers were selected for oral presentation, subject to a limit of no more than one oral presentation per author, and excluding four papers authored by the General Chair and Programme Chair ([16], [17], [18], [19]). Keynotes were invited at the discretion of the programme committee. Awards will be given to the best paper and poster presentation as selected by a committee comprising members of the UK-RAS network and senior programme committee members of UKRAS20.

Charles Fox is Programme Chair, UKRAS20, University of Lincoln, UK. Tom Duckett is General Chair, UKRAS20, University of Lincoln, UK. Arthur Richards is Technical Chair, UKRAS20, University of Bristol, and UK-RAS Network.

IV. CONFERENCE STATISTICS

67 submissions were received, of which 15 were accepted as oral presentations and 41 as poster presentations. (Acceptance rate 88%). Accepted papers are from 22 UK universities, with ratios shown in fig. 1. 44% of papers have two reviews, and 56% have three or more reviews. There were 45 participating reviewers from 28 UK universities whose affiliations are shown in fig. 2. The UKRAS20 host institution, University of Lincoln, and previous year's UKRAS19 host institution, University of Loughborough, are both highly represented in both presentations and reviewers.

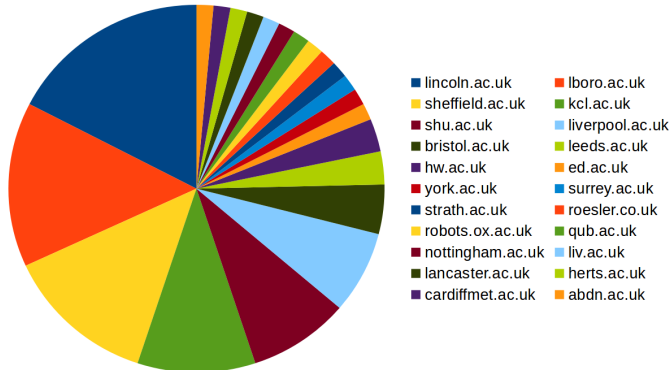


Fig. 1: Affiliations of accepted paper authors




Fig. 2: Affiliations of reviewers

V. REVIEWERS

Paul Baxter, Nicola Bellotto, Jordan Bird, Fanta Camara, Grzegorz Cielniak, Heriberto Cuayahuitl, Gautham Das, Daniel DeBarrie, Johann Dichtl, Khaled Elgeneidy, Yinfeng Fang, Diego Faria, Manuel Fernandez-Carmona, Joo Filipe Ferreira, Khaled Goher, Laura Justham, Marc Hanheide, Christos Kouppas, Weeding Li, Honghai Liu, Neils Lohse, Ahmad Lotfi, Mufti Mahmud, Luis J. Manso, Sarah Mghames, Alan Millard, Ben Mitchinson, Hector A. Montes, Harit Pandya, Simon Parsons, Martin Pearson, Riccardo Polvera, Mithun Poozhivil, David Portugal, Mini Saaj, Baris Serhan, Elizabeth Sklar, Aravinda Srinivasan, Cuebong Wong, Erfu Yang, Shing Yue, Tsvetan Zhivkov, Melanie Zimmer.

VI. LOCAL ARRANGEMENTS

Due to the 2020 COVID-19 pandemic, UKRAS20 is being held remotely as a teleconference. We ask all participants to study and respect teleconference etiquette. In particular this means that as with a physical conference, participants should devote their full day exclusively to attending the presentation and networking sessions. It is rude to perform other work on computers at the same time. Teleconferencing can deliver information and build communities as well as physical meetings if suitable etiquette is applied, and that it also reduces carbon emissions, fuel use, travel times, and attendance costs.

REFERENCES

- [1] O. Roesler, "Enhancing unsupervised natural language grounding through explicit teaching," in *Proceedings of The 3rd UK-RAS Conference*, 2020.
- [2] G. Miyauchi and E. Sklar, "A study assessing the impact of task duration on performance metrics for multi-robot teams," in *Proceedings of The 3rd UK-RAS Conference*, 2020.
- [3] A. Buerkle, A. Al-Yacoub, and P. Ferreira, "An incremental learning approach for physical human-robot collaboration," in *Proceedings of The 3rd UK-RAS Conference*, 2020.
- [4] E. Hunt, "Plastic 'personalities' for effective field swarm," in *Proceedings of The 3rd UK-RAS Conference*, 2020.
- [5] M. Li, A. Richards, and M. Sooriyabandara, "Reliability-aware multi-uav coverage path planning using integer linear programming," in *Proceedings of The 3rd UK-RAS Conference*, 2020.
- [6] A. Gabriel and P. Baxter, "Towards intention recognition for human-interacting agricultural robots," in *Proceedings of The 3rd UK-RAS Conference*, 2020.
- [7] R. Kirk, M. Mangan, and G. Cielniak, "Feasibility study of in-field phenotypic trait extraction for robotic soft-fruit operations," in *Proceedings of The 3rd UK-RAS Conference*, 2020.
- [8] R. Smith, "Enabling deep personalisation for a heterogeneous ambient assisted living landscape," in *Proceedings of The 3rd UK-RAS Conference*, 2020.
- [9] W. Mandil, K. Heiwolt, G. Cielniak, and M. Hanheide, "Automated topological mapping for agricultural robots," in *Proceedings of The 3rd UK-RAS Conference*, 2020.
- [10] A. Babu, A. Rathinam, and Z. Hao, "Trajectory tracking and control of multiple robot arms on a free-floating spacecraft for debris removal," in *Proceedings of The 3rd UK-RAS Conference*, 2020.
- [11] J. Hulas and C. Zho, "Improving quadrupedal locomotion on granular material using genetic algorithm," in *Proceedings of The 3rd UK-RAS Conference*, 2020.
- [12] R. Woolley, J. Timmis, and A. Tyrell, "Analysis of two wheeled robot morphology for aslope environment," in *Proceedings of The 3rd UK-RAS Conference*, 2020.
- [13] I. Issak and A. Kucukyilmaz, "The goods and bads in dyadic co-manipulation: Identifying conflict-driven interaction behaviours in human-human collaboration," in *Proceedings of The 3rd UK-RAS Conference*, 2020.
- [14] L. Hickton, M. Lewis, and L. Canamero, "Expression of grounded affect in a hexapod robot," in *Proceedings of The 3rd UK-RAS Conference*, 2020.
- [15] J. Emerson and K. Elgeniedy, "Optimising soft fin ray robotic fingers using finite element analysis to reduce object slippage," in *Proceedings of The 3rd UK-RAS Conference*, 2020.
- [16] Z. Hobbs, T. Duckett, S. Pearson, and M. Mangan, "Towards bio-inspired fruit detection for agriculture," in *Proceedings of The 3rd UK-RAS Conference*, 2020.
- [17] H. Rogers and C. Fox, "An open source seeding agri-robot," in *Proceedings of The 3rd UK-RAS Conference*, 2020.
- [18] X. Li, C. Fox, and S. Coutts, "Deep learning for robotic strawberry harvesting," in *Proceedings of The 3rd UK-RAS Conference*, 2020.
- [19] F. Camara and C. Fox, "Game theory for self-driving cars," in *Proceedings of The 3rd UK-RAS Conference*, 2020.

Enhancing Unsupervised Natural Language Grounding through Explicit Teaching

Oliver Roesler
Artificial Intelligence Lab
Vrije Universiteit Brussel
 Brussels, Belgium
 oliver@roesler.co.uk

<https://doi.org/10.31256/Bf9Vw8C>

Abstract—In this paper, a grounding framework is proposed that combines unsupervised and supervised grounding by extending an unsupervised grounding model with a mechanism to learn from explicit human teaching. To investigate whether explicit teaching improves the sample efficiency of the original model, both models are evaluated through an interaction experiment between a human tutor and a robot in which synonymous shape, color, and action words are grounded through geometric object characteristics, color histograms, and kinematic joint features. The results show that explicit teaching improves the sample efficiency of the unsupervised baseline model.

Index Terms—language grounding, cross-situational learning, sample efficiency, human-robot interaction

I. INTRODUCTION

The need for robots that are able to understand natural language instructions is growing due to an increasing number of service robots that are employed in human-centered environments. To this end, connections between words and percepts need to be created through grounding because language only has meaning, if it is linked to the physical world [1].

Previous studies that investigated grounding employed either unsupervised [2]–[4] or supervised [5], [6] approaches. The former have the advantage that no human tutor is required for grounding, however, they require a large number of situations to learn the correct grounding, i.e. they are less sample efficient, and are often also less accurate than supervised approaches. In comparison, the latter are often more accurate and can already learn the correct mappings from a very small number of situations, however, they do not work in the absence of a human tutor.

In this paper, both approaches are combined by extending a recently proposed unsupervised cross-situational learning based grounding framework [7], [8] to learn from explicit human teaching, if available. The main aim is to investigate whether this extension increases the model’s sample efficiency, i.e. whether it reduces the number of interactions required until the model obtains the correct mappings between words and percepts.

The rest of this paper is structured as follows: Section (II) describes the extended grounding framework. The experimental design and obtained results are described in Sections (III and IV). Finally, Section (V) concludes the paper.

II. SYSTEM OVERVIEW

The used grounding system consists of the following parts:

- 1) **3D object segmentation system**, which employs a model based 3D point cloud segmentation approach [9] to segment objects into point clouds. The shapes and colors of objects are represented through Viewpoint Feature Histogram [10] descriptors, which represent the object geometries taking into consideration the viewpoints, while ignoring scale variances, and color histograms.
- 2) **Action recording system**, which records the vertical position of the robot’s torso, the angles of the arm flex and wrist roll joints, the velocity of the robot’s base and the binary state of the gripper, i.e. open or closed, during action execution. The recorded data is then combined into an action feature vector, which represents the change of the recorded characteristics between the beginning and the end of an action.
- 3) **Percept clustering component**, which converts percepts to abstract representations through clustering to enable the CSL algorithm to use them to ground encountered words as proposed by [11]. The used clustering algorithm is DBSCAN [12] because it does not require the number of clusters to be specified in advance, which is important since it cannot be assumed that the number of percepts is known in advance. Cluster numbers were calculated prior to grounding so that they could be provided to the CSL algorithm. Shape, color, and action percepts achieved mean adjusted rand scores [13] of 0.84, 1.0, and 0.99, respectively.
- 4) **Language grounding component**, which uses an extended version of the cross-situational learning based grounding algorithm proposed by [8]. The original grounding algorithm grounds words and phrases through cluster numbers of percepts in an unsupervised manner without being able to take into account any teaching or feedback by a human tutor. Thus, in this study an extension is proposed that provides a mechanism to learn mappings from explicit teaching. The new mechanism uses similar to the original one cross-situational learning to determine the correct mappings, however, it requires the tutor to artificially create a situation where only one percept occurs twice and only one word is given

to the robot, which should be grounded through that percept (Section III). When a new mapping has been obtained through explicit teaching, it is added to the set of previously obtained mappings, which also includes mappings obtained through the unsupervised grounding algorithm during regular situations.

III. EXPERIMENTAL SETUP

During the experiment a human tutor and HSR robot [14] interact in front of a table with one or two objects on top of it. Interactions can be of two types. Either the human tutor asks the robot to perform an action on the object or the tutor tries to teach the robot the correct mapping for a shape, color, or action word. The former interactions use the following procedure.

- 1) The tutor places an object on the table and the robot determines the corresponding shape and color percepts.
- 2) The human tutor provides an instruction to the robot.
- 3) The human tutor teleoperates the robot to execute the action provided through the instruction and the robot determines a corresponding action percept.
- 4) The robot employs clustering to convert all encountered percepts to abstract representations.
- 5) Words are grounded through obtained cluster numbers by the CSL based grounding algorithm.

In contrast, situations in which the human tutor tries to teach the robot a specific mapping follow the following procedure.

- 1) The human tutor places two objects, which have either the same shape or color, on the table and the robot determines the corresponding percepts, if the tutor tries to teach the correct mapping for a shape or color word. Otherwise, to teach an action word, the human tutor places two objects on the table that have different shapes and colors and executes the same action for both of them so that only the action percept occurs twice.
- 2) The human tutor provides one single word, which refers to the percept that occurs twice.
- 3) The robot creates a corresponding mapping and adds it to the set of previously obtained mappings.

To create the 2,500 situations used in this study without having to perform 2,500 interactions, the following procedure is applied. First, a total of 125 interactions are performed to record perceptual information for all combinations of employed shapes, colors, and actions, while skipping the last two steps of the interaction procedure, i.e. steps 4 and 5. Afterwards, all possible unique sentences are obtained by creating all possible combinations of shape, color, and action words. Finally, each sentence is randomly assigned one shape, color, and action percept that correspond to the words in the sentences, leading to overall 2,500 situations.

Each sentence has the following structure: “*action* the *color* *shape*”, where *action*, *color*, and *shape* are replaced by one of their corresponding words. Each action and color can be referred to by two different words, e.g. the color green can be referred to by “green” or “greenish”, while each shape has five corresponding words, e.g. “latte”, “milk”, “milk tea”, “coffee” or “espresso” for cup.

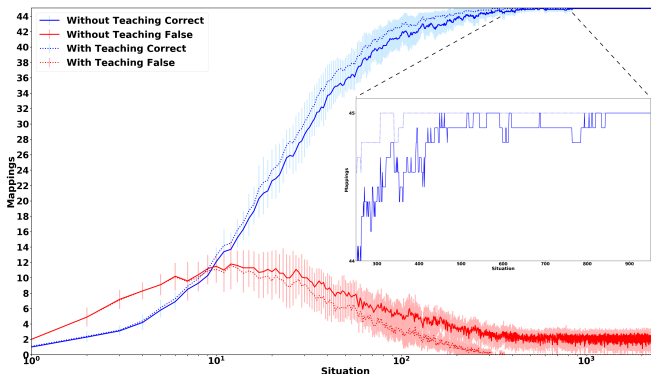


Fig. 1. Grounding results showing means and standard deviations of correct and false mappings over all 2,500 situations encountered by the robot for 10 different sequences. The dotted lines represent the results, when the tutor teaches the robot a correct mapping after on average every 9 situations, while the continuous lines represent the results, when no explicit teaching is provided.

IV. RESULTS AND DISCUSSION

The results show that teaching increases the convergence towards the correct mappings (Figure 1). If no teaching is provided, the algorithm requires about 850 situations to ground all words correctly for all 10 different sequences, while it only requires about 350 situations, when the human tutor explicitly teaches one mapping to the robot after on average every 9 situations. Before obtaining all correct mappings teaching also leads to a slightly higher number of correctly grounded words. Even after all words are correctly grounded, the number of false mappings oscillates around 2, when no teaching is provided, because the algorithm allows words to be grounded through several percepts to handle homonyms. In contrast, no false mappings are obtained in case of the extended model. If the human tutor teaches all 45 words at the beginning of the experiment, the model learns all words after 45 situations, assuming that all encountered percepts are correctly clustered. While teaching all words explicitly is possible for the small number of words used in this scenario, it would not be feasible for a much larger number of words, which illustrates the importance of the unsupervised grounding mechanism.

V. CONCLUSIONS AND FUTURE WORK

An unsupervised grounding model was extended to allow it to benefit from explicit teaching by a human tutor. The proposed model was evaluated through a human-robot interaction experiment and compared to the original model that does not allow explicit teaching. The results showed that with teaching the model grounds all words on average about 2.5 times faster than without teaching. In future work, the model will be evaluated for longer and more complex sentences that contain a larger number of words. Furthermore, the influence of wrong teaching, i.e. when the tutor on purpose or by accident provides a wrong word during teaching, will be investigated. Finally, the model will be extended to allow human feedback for already obtained groundings.

REFERENCES

- [1] S. Harnad, "The symbol grounding problem," *Physica D*, vol. 42, pp. 335–346, 1990.
- [2] C. R. Dawson, J. Wright, A. Rebguns, M. V. Escárcega, D. Fried, and P. R. Cohen, "A generative probabilistic framework for learning spatial language," in *IEEE Third Joint International Conference on Development and Learning and Epigenetic Robotics (ICDL)*, Osaka, Japan, August 2013.
- [3] O. Roesler, A. Aly, T. Taniguchi, and Y. Hayashi, "A probabilistic framework for comparing syntactic and semantic grounding of synonyms through cross-situational learning," in *ICRA-18 Workshop on Representing a Complex World: Perception, Inference, and Learning for Joint Semantic, Geometric, and Physical Understanding.*, Brisbane, Australia, May 2018.
- [4] —, "Evaluation of word representations in grounding natural language instructions through computational human-robot interaction," in *Proceedings of the 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, Daegu, South Korea, March 2019.
- [5] L. Steels and M. Loetzsch, "The grounded naming game," in *Experiments in Cultural Language Evolution*, L. Steels, Ed. Amsterdam: John Benjamins, 2012, pp. 41–59.
- [6] L. She, S. Yang, Y. Cheng, Y. Jia, J. Y. Chai, and N. Xi, "Back to the blocks world: Learning new actions through situated human-robot dialogue," in *Proceedings of the SIGDIAL 2014 Conference*, Philadelphia, U.S.A., June 2014, pp. 89–97.
- [7] O. Roesler and A. Nowé, "Simultaneous action learning and grounding through reinforcement and cross-situational learning," in *ALA 2018, Adaptive Learning Agents Workshop.*, Stockholm, Sweden, July 2018.
- [8] —, "Action learning and grounding in simulated human robot interactions," *The Knowledge Engineering Review*, vol. 34, no. E13, November 2019.
- [9] C. Craye, D. Filliat, and J.-F. Goudou, "Environment exploration for object-based visual saliency learning," in *IEEE International Conference on Robotics and Automation (ICRA)*, Stockholm, Sweden, May 2016.
- [10] R. B. Rusu, G. Bradski, R. Thibaux, and J. Hsu, "Fast 3D recognition and pose using the viewpoint feature histogram," in *Proceedings of the 2010 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Taipei, Taiwan, October 2010, pp. 2155–2162.
- [11] O. Roesler, "A cross-situational learning based framework for grounding of synonyms in human-robot interactions," in *In Proc. of the Fourth Iberian Robotics Conference (ROBOT)*, Porto, Portugal, November 2019.
- [12] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu, "A density-based algorithm for discovering clusters in large spatial databases with noise," in *Proceedings of the 2nd International Conference on Knowledge Discovery and Data Mining (KDD)*, Portland, Oregon, USA, August 1996, pp. 226–231.
- [13] L. Hubert and P. Arabie, "Comparing partitions," *Journal of Classification*, vol. 2, no. 1, pp. 193–218, December 1985.
- [14] *HSR Manual*, 2017th ed., Toyota Motor Corporation, April 2017.

A Study Assessing the Impact of Task Duration on Performance Metrics for Multi-Robot Teams

Genki Miyauchi

Dept of Automatic Control & Systems Engineering
University of Sheffield, S1 3JD, UK
g.miyauchi@sheffield.ac.uk

Elizabeth I Sklar

Lincoln Institute for Agri-food Technology
University of Lincoln, LN6 7TS, UK
esklar@lincoln.ac.uk

<https://doi.org/10.31256/Xg3Gx5E>

Abstract—The allocation of tasks to members of a team is a well-studied problem in robotics. Applying market-based mechanisms, particularly auctions, is a popular solution. We focus on evaluating the performance of the team when executing the tasks that have been allocated. The work presented here examines the impact of one such factor, namely task duration. Building on prior work, a new bidding strategy and performance metric are introduced. Experimental results are presented showing that there are statistically significant differences in both time and distance-based performance metrics when tasks have zero vs greater-than-zero duration.

I. INTRODUCTION

Assigning a set of tasks within a team of robots is known as the *multi-robot task allocation (MRTA)* problem. As the number of tasks and size of the robot team increases, the number of possible allocations rises at an exponential rate. The complexity further increases when factors are added, such as a dynamic environment that changes over time, a heterogeneous team of robots with various capabilities or tasks that have prerequisites which must be satisfied before they can be executed. Finding the optimal solution to an MRTA problem is known to be NP-hard [1], [2], so a popular family of strategies takes a *market-* or *auction-*based approach. Auctions are useful because they can distribute workload amongst team members whilst reflecting preferences of individuals [2], [3], [4]. A local optimum is determined by each robot (labelled *bidder*) and these optima are collectively reconciled by an *auctioneer*. The team can operate in real-time and respond dynamically to changes in the task landscape, such as the arrival of new tasks to be addressed while already executing tasks previously allocated.

The work presented here extends our prior work in which we demonstrated the importance of mission execution and weighing various performance metrics when comparing task allocation mechanisms [5], [4]. Here we assess the impact of *task duration*. Our contributions include: (a) a new methodology for robots to bid on tasks with varying durations; (b) a new metric that attempts to capture aspects of task duration; (c) results of experiments conducted both in simulation and on physical robots across a landscape of mission parameters; and (d) statistical analysis of results to highlight the impact on performance metrics when task duration varies.

This work was conducted while the first author was a masters student in the Dept of Informatics at King's College London, UK. The second author was a professor in the same department and project supervisor.

II. BACKGROUND

Koenig et al. [2] proposed *sequential single-item (SSI)* auctions in which several tasks are announced to team members at once; each robot responds with a bid representing a cost to the robot for executing the task, e.g. the distance the robot must travel to reach the task location. An auctioneer identifies the winner as the robot with the smallest bid. The auction repeats in *rounds* until all tasks have been allocated. SSI combines the strength of *combinatorial* [6] (bidding on bundles of tasks) and *parallel single-item (PSI)* [2] (allocating all tasks in a single round) auctions. SSI has been a popular choice for MRTA and several variants have been studied, for example tasks with temporal constraints [3], tasks with precedence constraints [7], [8], and tasks with pickup-and-delivery constraints [9]. Our prior work has focused on empirical analysis of auction mechanisms [5], comparing SSI, PSI and a baseline *Round Robin (RR)* (first come, first served) in experiments conducted on physical and simulated robots. We have defined a comprehensive set of performance metrics [10] and learned a model for selecting an appropriate mechanism given mission parameters [11].

III. APPROACH

We assess the impact of varying task durations on multi-robot team performance by executing missions on physical robots and in simulation, and then analysing differences in performance metrics. Missions are defined across a landscape of parameters [1], [12] which distinguish characteristics of tasks, robots and the environment: *single robot (SR)* vs *multi-robot (MR)*—SR tasks can be completed by one robot, whereas MR tasks require the cooperation of multiple robots; *independent (IT)* vs *constrained (CT)*—IT tasks can be executed in any order, whereas CT tasks are dependent on others due to factors such as precedence order; *static (SA)* vs *dynamic (DA)*—SA tasks are known before a mission starts and can be allocated before any execution, whereas DA tasks arrive dynamically and are allocated during execution of other tasks.

Here, we introduce a *task duration* parameter, i.e. the time it takes to execute a task: *instantaneous (ID)* vs *extended (XD)*—ID tasks take no time to execute (i.e., 0 seconds), whereas XD task length is > 0 . We compare four XD variants: *XDC*, where all tasks have the same constant length; *XDG*, where task length is chosen randomly from



(a) TurtleBot3 Burger (b) arena

Fig. 1. Experimental setup.

a Gaussian distribution; XDP , where task length is chosen randomly from a Poisson distribution; and XDR , where task length is chosen randomly from a uniform distribution¹.

We also introduce a new bidding strategy that calculates the estimated time for a robot to travel to a task location instead of using travel distance as the basis of a bid. Thus the bid takes into account the estimated travel time as well as predicted task duration. The bid value b of robot r for a new task x is calculated as: $b_r = \sum_{i=1}^{N_r} (T_{i-1,i} + E_i) + T_{N_r,x}$, where N_r is the number of uncompleted tasks robot r is assigned, T is the estimated time to travel between two task locations² and E is the predicted duration of task i .

IV. EXPERIMENTS

We conducted a series of experiments with physical and simulated Turtlebot3 robots (Figure 1a) using the *MRTeAm* [5], [10] framework built with Robot Operating System (ROS). Our experimental arena emulates an office-like area divided into rooms and corridors (Figure 1b).

For our experiments, we employed three different *auction mechanisms* (*RR*, *PSI*, *SSI*) two different *starting configurations*: *clustered* together in one portion of the arena or *distributed* around the arena; two different task constraints (*IT*, *CT*); and five task durations (*ID*, *XDC*, *XDG*, *XDP*, *XDR*). For each combination of parameters, at least 5 runs were conducted in the physical environment and 15 runs in simulation. In total, 1926 runs were completed.

Three performance metrics are analysed to assess the impact of task duration: the total *distance travelled* collectively by all the members of the robot team; the total *run time* from the start of a mission until all tasks are completed; and *service delay time*, the time from when an auctioneer awards a task until a robot begins executing the task. Service time is a new metric introduced here to reflect the time each task “waits” before a robot arrives at its location. Service time measures how quickly a task’s execution can commence or how long the task was delayed, i.e. left unattended. Arriving at a task location quickly is an important factor in many application domains, especially in emergency situations. A shorter service time is preferred because it means the robot team is able to arrive at task locations more promptly.

¹The values employed here are: XDC: length=15 seconds; XDG: $\mu = 15$ and $\sigma = 3$; XDP: $\lambda = 15$; XDR: range=(5, 35).

² T is calculated as the predicted travel distance times the average robot velocity, which was determined experimentally to be 0.12m/s.

V. RESULTS

Our analysis assesses the impact of task duration on the three performance metrics. For brevity here, we present aggregated results and make three comparisons: (a) across all five task duration values; (b) instant versus extended time; and (c) across the four extended time values. The number of runs per aggregate are: *ID*: $N = 330$, *XDC*: $N = 508$, *XDG*: $N = 469$, *XDP*: $N = 379$, and *XDR*: $N = 240$.

We looked for statistically significant results. After checking that the data within each sample is normally distributed³, we tested for differences using *analysis of variance* (*ANOVA*), with $p < 0.01$ as the threshold for statistical significance. The results of statistical testing on individual variables are shown in Figure 2, which illustrates the distributions for each sample and indicates, for each of the three types of analysis (above), which differences amongst samples are statistically significant.

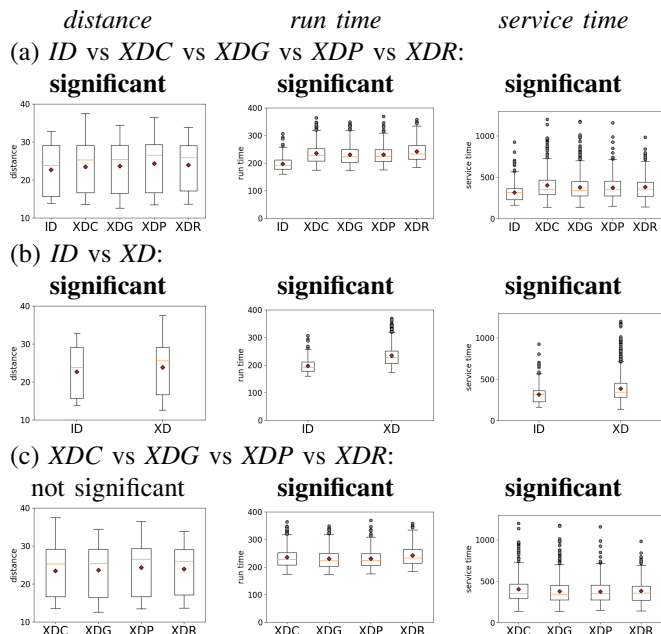


Fig. 2. Statistical differences. Statistical significance is noted where ANOVA produces $p < 0.01$.

VI. SUMMARY

We have presented a study assessing the impact of task duration on three different performance metrics collected from experiments with multi-robot teams. We have introduced a new bidding strategy which accounts for task duration and a new metric for measuring performance delay due to task duration. Statistical analysis of our results shows that there are significant differences in all performance metrics when tasks have non-zero duration (*ID* vs *XD**) and in time-related metrics when duration varies randomly according to probability distributions defined by similar parameter values.

³Using the Shapiro-Wilk test [13]

REFERENCES

- [1] B. Gerkey and M. Mataric, "A formal analysis and taxonomy of task allocation in multi-robot systems," *Intl Journal of Robotics Research*, vol. 23, no. 9, 2004.
- [2] S. Koenig, C. Tovey, M. Lagoudakis, V. Markakis, D. Kempe, P. Keskinocak, A. Kleywegt, A. Meyerson, and S. Jain, "The power of sequential single-item auctions for agent coordination," in *21st Natl Conf on Artificial Intelligence (AAAI)*, vol. 2, 2006.
- [3] E. Nunes and M. Gini, "Multi-robot auctions for allocation of tasks with temporal constraints," in *29th Conference on Artificial Intelligence (AAAI)*, 2015.
- [4] E. Schneider, "Mechanism selection for multi-robot task allocation," Ph.D. dissertation, Univ of Liverpool, 2018.
- [5] E. Schneider, E. I. Sklar, S. Parsons, and A. T. Özgelen, "Auction-based task allocation for multi-robot teams in dynamic environments," in *Towards Autonomous Robotic Systems (TAROS)*. Springer, 2015.
- [6] M. Berhault, H. Huang, P. Keskinocak, S. Koenig, W. Elmaghraby, P. Griffin, and A. Kleywegt, "Robot exploration with combinatorial auctions," in *IEEE/RSJ Intl Conf on Intelligent Robots and Systems (IROS)*, vol. 2, 2003.
- [7] M. McIntire, E. Nunes, and M. Gini, "Iterated multi-robot auctions for precedence-constrained task scheduling," in *Intl Conf on Autonomous Agents and Multiagent Systems (AAMAS)*, 2016.
- [8] E. Nunes, M. McIntire, and M. Gini, "Decentralized allocation of tasks with temporal and precedence constraints to a team of robots," in *IEEE Intl Conf on Simulation, Modeling, and Programming for Autonomous Robots (SIMPAN)*, 2016.
- [9] B. Heap and M. Pagnucco, "Repeated sequential single-cluster auctions with dynamic tasks for multi-robot task allocation with pickup and delivery," in *Multiagent System Technologies*. Springer, 2013.
- [10] E. Schneider, E. I. Sklar, and S. Parsons, "Evaluating multi-robot teamwork in parameterised environments," in *Towards Autonomous Robotic Systems (TAROS)*. Springer, 2016.
- [11] E. Schneider, E. Sklar, and S. Parsons, "Mechanism selection for multi-robot task allocation," in *Towards Autonomous Robotic Systems (TAROS)*. Springer, 2017.
- [12] D. Landén, F. Heintz, and P. Doherty, "Complex task allocation in mixed-initiative delegation: A uav case study," in *Principles and Practice of Multi-Agent Systems*. Springer, 2012.
- [13] S. S. Shapiro and M. B. Wilk, "An Analysis of Variance Test for Normality (Complete Samples)," *Biometrika*, vol. 52, no. 3/4, 1965.

An Incremental Learning Approach for Physical Human-Robot Collaboration

Achim Buerkle
Wolfson School of Engineering
Loughborough University
Loughborough, UK
a.buerkle@lboro.ac.uk

Ali Al-Yacoub
Wolfson School of Engineering
Loughborough University
Loughborough, UK
a.al-yacoub@lboro.ac.uk

Pedro Ferreira
Wolfson School of Engineering
Loughborough University
Loughborough, UK
p.ferreira@lboro.ac.uk

<https://doi.org/10.31256/It7Lm1B>

Abstract—Physical Human-Robot Collaboration requires humans and robots to perform joint tasks in a shared workspace. Since robot's characteristic strengths are to cope well with high payloads, they are utilized to assist human operators during heavy pulling or pushing activities. A widely used sensor to detect human muscle fatigue and thus, to trigger an assistance request, is an Electromyography (EMG). Many previous approaches to process EMG data are based on training Machine Learning models offline or include a large degree of manual fine tuning. However, due to recent advances in Machine Learning such as incremental learning, there is an opportunity to apply online learning which reduces programming effort and also copes well with subject specific characteristics of EMG signals. Initial results show promising potential, yet, unveil a conflict between convergence time and classification accuracy.

Keywords—EMG, Human-Robot Collaboration, Incremental Learning, Machine Learning

I. INTRODUCTION

Human-Robot Collaboration (HRC) in manufacturing aims to establish symbiotic or synergetic effects between human operators and robots [1]. This is enabled by combining the characteristic strengths of each party. Human strengths are considered to be adaptability to changes, decision making, and problem solving [1], [2]. Robot's strengths, on the other hand, are high precision, high operating speeds, and the capability of coping with high payloads [3]. Thus, in a physical collaboration, robots are able to support human operators via force amplification to handle heavy pushing and pulling activities [4]. In order to measure human muscle activity such as during the lift of heavy objects, a widely used sensor is an Electromyography (EMG) [5]. The approaches typically include pre-processing of the data, feature extraction, and a supervised learning of the model [5], [6]. However, recent advances in Machine Learning regarding incremental learning could allow to minimize the training and programming effort of such models [7]. Furthermore, the algorithm could optimize its performance over time in an online system [7]. In this work, an incremental learning approach is utilized to predict EMG data during three different states: Participants lifting light payloads, medium payloads, and heavy payloads (struggling).

II. RELATED WORK

In a Human-Robot collaborative scenario, humans and robots perform joint tasks in a shared workspace [1]. In order to communicate intentions of the human to the robot, sensors are utilized such as EMGs [5]. The EMG signals are usually acquired from a human upper-limb since they are mostly used in the given tasks [5]. The acquired data can be used to communicate movement intentions. It can also provide insights on human muscle fatigue [6]. In this case, a robot

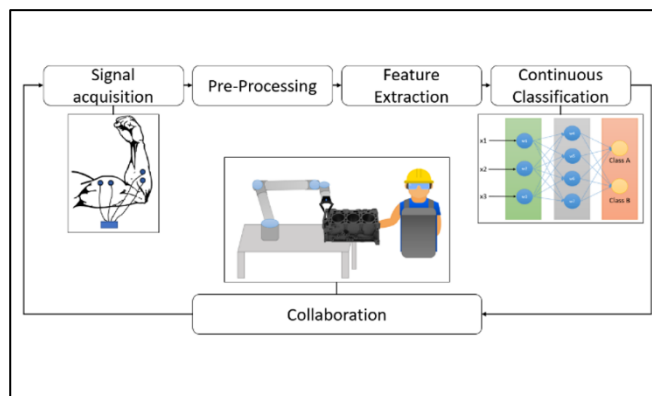


Figure 1 EMG Signal Processing for HRC (adapted) [5]

could assist a human operator during a heavy pull or push of an object or adapt its behavior to create more ergonomic working conditions for its human co-worker [6], [8]. This is intended to prevent injuries, as well as long-term health issues related to physical fatigue [8].

Figure 1 shows the general process used to integrate EMGs in Human-Robot Collaboration for a supervised, non-incremental learning approach. The first stage is EMG data acquisition. Critical attention is required during the selection of the acquisition device, the number of channels used, as well as the placement of each channel [5]. The channel acquisition device also determines the sampling rate and data transmission [9].

During the pre-processing stage, raw EMG signals are checked for baseline offset [5]. Typically, the signal is corrected by subtracting the average amplitude from each instance, however, there are also approaches based on nonlinear error-modelling [5], [6]. Raw EMG signals are susceptible to contain noise. Thus, Butterworth filters with a cut off frequency from 2Hz-20Hz are utilized [5]. The remaining features are extracted in the Feature selection and extraction stage. This is critical during EMG data processing since it has a high impact on the classification accuracy [9]. Three properties are considered as essential: class separability (minimize overlap), robustness (separability in noisy environment), and computational complexity (low complexity of features implies lower processing times) [5], [9].

The fourth stage is continuous classification of the filtered and extracted signals. There are mainly two types of prediction models. One is the use of kinematic models, the second approach is to utilize Artificial Neural Networks (ANNs) [5]. However, [9] states that Linear Discriminant Analyses (LDA) and Support Vector Machines (SVMs) are also widely used for EMG data classification. According to [5] there are few critical challenges remaining. Firstly, many

offline systems obtain high classification accuracies, yet the online performances of such systems are far from satisfactory. Secondly, there are subject-specific characteristics of EMG signals. This can even include variation of the EMG signals for the same person during different recording sessions. An opportunity to increase the performance and to lower programming and fine-tuning effort could be incremental learning. Incremental learning algorithms have the following characteristics: ability of life-long learning, ability to incrementally tune the model's performance, and no prior knowledge about the data and its properties is needed [7].

III. EXPERIMENTAL SETUP

The experimental setup aims to collect EMG data during three different stages: light payload, medium payload, and high payload, during which a participant is slightly struggling.




Figure 2 Experimental Setup

The acquired data will be fed into the classifier unlabeled. However, in order to validate the prediction results, predicted classes and the actual classification will be compared.

IV. RESULTS AND DISCUSSION

The collected data was used to train an Online Random Forest (ORF) model, that aims to classify the EMG signals into low payload, medium payload, and heavy payload. In any incremental learning approach, the most crucial property apart

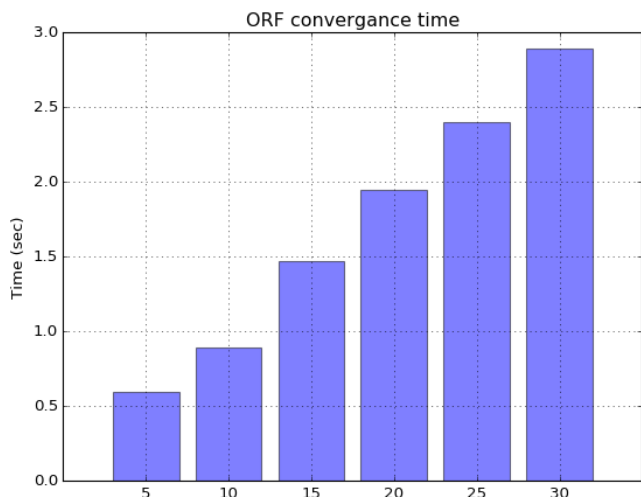


Figure 3 ORF convergence time

from accuracy is the convergence time. Since the model aims to minimize the prediction error live and immediately. In Human-Robot Collaboration, this is exceptionally important as humans and robots are physically interacting. Hence, in this validation experiment convergence time of the ORF model was measured with a different number of trees, which is illustrated in Figure 3.

As, expected, Figure 3 shows that the convergence time is directly proportional to the number of trees in the ORF model. The correspondent accuracy of the models in Figure 3 is shown in Table 1. The collected data for this experiment is ~6000 data points of EMG signals and the associated labels. Based on Figure 3 and Table 1, it can be noticed that the model must achieve a trade-off between accuracy and convergence time. The ORF model with 20 trees seems to be the most suitable model since it can converge in less than 2 seconds, and it achieves the highest detection accuracy.

Table 1 Number of Trees vs Prediction Accuracy

Number of Trees	5	10	15	20	25	30
Mean Square Error	0.26	0.19	0.250	0.14	0.15	0.17
Accuracy [%]	82.3	84.6	85.2	89.7	86.7	86.7

V. CONCLUSION AND FUTURE WORK

A novel incremental learning approach was introduced to determine physical workload from EMG data in Human-Robot Collaboration. During the online training, a conflict became clear between processing speed and accuracy. Lesser trees in the model meant faster convergence, however, it also resulted in the aforementioned lower accuracy. Overall, the accuracy could reach 89% in only two seconds. Thus, in a Human-Robot Collaborative Scenario this would allow the system to recognize a human operator struggling with the payload. The collaborative robot could then support the operator and subsequently, create a more ergonomic environment. However, prior to this technology being ready to be used in a practical application, further testing is essential. This includes the need for a larger sample size in participants and a richer variety in lifting tasks. The current setup allows to detect muscle contraction in participant's forearms and biceps. Yet, more EMG sensors placed on other muscle groups such as triceps and shoulders are expected to provide better results for predicting pushing activities. Furthermore, the system could be trained to not only detect temporary high payloads but also to recognize muscular fatigue during endurance tasks. This could help to improve human operator's posture and subsequently prevent negative long-term health effects.

Nevertheless, early results of this incremental learning approach demonstrate a reduced manual fine-tuning effort and it coping well with subject specific characteristics in the data. This offers the potential to be applied for additional human sensor technologies and subsequent data classifications. Ultimately, this could help to make Human-Robot Collaboration safer and more efficient.

VI. REFERENCES

- [1] L. Wang *et al.*, “Symbiotic human-robot collaborative assembly,” *CIRP Ann.*, 2019.
- [2] V. Villani, F. Pini, F. Leali, and C. Secchi, “Survey on human-robot collaboration in industrial settings: Safety, intuitive interfaces and applications,” *Mechatronics*, no. June 2017, pp. 1–19, 2018.
- [3] J. Krüger, T. K. Lien, and A. Verl, “Cooperation of human and machines in assembly lines,” *CIRP Ann. - Manuf. Technol.*, vol. 58, no. 2, pp. 628–646, 2009.
- [4] J. Schmidler and K. Bengler, “Fast or Accurate? – Performance Measurements for Physical Human-robot Collaborations,” *Procedia Manuf.*, vol. 3, no. Ahfe, pp. 1387–1394, 2015.
- [5] L. Bi, A. Feleke, and C. Guan, “A review on EMG-based motor intention prediction of continuous human upper limb motion for human-robot collaboration,” *Biomed. Signal Process. Control*, vol. 51, pp. 113–127, 2019.
- [6] L. Peternel, C. Fang, N. Tsagarakis, and A. Ajoudani, “A selective muscle fatigue management approach to ergonomic human-robot co-manipulation,” *Robot. Comput. Integr. Manuf.*, vol. 58, no. January, pp. 69–79, 2019.
- [7] A. Bouchachia, B. Gabrys, and Z. Sahel, “Overview of some incremental learning algorithms,” *IEEE Int. Conf. Fuzzy Syst.*, no. August, 2007.
- [8] L. Peternel, N. Tsagarakis, D. Caldwell, and A. Ajoudani, “Robot adaptation to human physical fatigue in human-robot co-manipulation,” *Auton. Robots*, vol. 42, no. 5, pp. 1011–1021, 2018.
- [9] M. Hakonen, H. Piitulainen, and A. Visala, “Current state of digital signal processing in myoelectric interfaces and related applications,” *Biomed. Signal Process. Control*, vol. 18, pp. 334–359, 2015.

Plastic ‘personalities’ for effective field swarms

Edmund R. Hunt

Department of Engineering Mathematics

University of Bristol

Bristol, U.K.

edmund.hunt@bristol.ac.uk

<https://doi.org/10.31256/Xa2Lf8K>

Abstract—Most studies on real-world multi-robot systems have been performed in controlled laboratory environments, whereas the real world is unpredictable and sometimes hazardous. I have recently suggested that the natural phenomenon of *phenotypic plasticity* provides a useful bioinspiration framework for making such systems more resilient in field conditions [1]. Phenotypic plasticity occurs when a single genotype produces a range of phenotypes (observable traits) in response to different environmental conditions. Consistent individual behavioural differences can result from such plasticity, and have been described as ‘personalities’. At the same time, in social animals, individual heterogeneity is increasingly recognised as functional for the group. We can exploit this functional heterogeneity as engineers trying to design field robot systems, and phenotypic plasticity can provide meaningful diversity ‘for free’, based on the local experience of agents. Personality axes such as bold–shy or social–asocial can be represented as single variables, with the advantage of being transparent and intuitive for human users, and predictable in their effects. For example, in a dangerous environment, robots may become more ‘shy’ and ‘social’ to stay closer together and out of harm’s way.

Index Terms—phenotypic plasticity, reaction norms, resilience, personality, functional heterogeneity, field robotics

I. INTRODUCTION

There is increasing recognition in biology that consistent individual differences in behaviour (‘personality’) among group members can be important for group function in local ecologies [2]. Examples of significant personality axes include: risk-taking behaviour (boldness—shyness), exploratory behaviour (neophilic—neophobic), activity levels (active—inactive), sociability (social—asocial), and aggression (aggressive—non-aggressive) [3]. Variation can also be seen in reaction threshold-type behaviours, for example the acceptability of options in a decision task [4]; one might term this ‘choosiness’ or ‘pickiness’. Despite the definition of personality relating to *consistent* differences, personality is also recognised as being somewhat variable – plastic – over time.

The term *developmental reaction norm* (DRN) describes the range of phenotypes generated by a given genotype (artificial agent controller) in response to experienced environmental cues [5]. There are at least five attributes to DRNs: amount of plasticity (large/small); pattern of response (e.g. monotonic increase/decrease or more complex reaction curves); rapidity

of response; reversibility of response; and competence (possibility) of the developmental system to respond at a certain stage in an organism’s (agent’s) lifetime [5]. One can refer to *behavioural reaction norms* (BRNs) if behaviour is the focus, as is the case here (plasticity can also be observed in physiology and morphology). BRNs can be a useful framework for integrating the notions of animal personality and individual plasticity [6]. In biology, the various attributes of developmental reaction norms are, in principle, subject to natural selection [5], [6]. I suggest that engineers attempting to deploy collectives of robots into real-world field conditions can undertake pre-deployment artificial evolution of DRNs. This should establish an adaptive DRN profile [1].

Examples of personality variation can be readily found at the level of the individual or the whole group, which gives rise to the notion of collective personalities [7]. Behavioural plasticity allows organisms to make relatively rapid adjustments in their function to adapt to changing environmental conditions, and can be seen as personality adaptation [6]. This is true for individuals, and also for whole groups: for instance, in leaf-cutter ants the whole colony can become more threat-responsive and aggressive in response to disturbance [8].

II. PERSONALITY PLASTICITY AS COMPLEMENTARY TO EXISTING APPROACHES

There are already several examples in the swarm robotics literature in which individual robots, though identically programmed with the same controller, end up behaving differently according to their experience of the environment. These are briefly:

- Off-line (pre-deployment) evolutionary optimisation to identify environmentally contingent behaviours that are adaptive at the group level (e.g. [9]); though effectiveness is tuned to the particular simulated environment.
- With sufficient computing power, one can undertake on-line (on-deployment) evolutionary optimisation (e.g. [10]). However, evolutionary approaches (off- or on-line) could struggle in the field owing to unanticipated circumstances or because of the ‘reality gap’ between the world and (inner) simulation (e.g. [11]).
- On-deployment learning: typically this employs (evolved) neural networks (e.g. [12]). Yet, neural network-based approaches can have difficulty in scaling to more complex problems [13] and be less human-readable.

EH acknowledges support from the Royal Academy of Engineering and the Office of the Chief Science Adviser for National Security under the UK Intelligence Community Postdoctoral Fellowship Programme.

I suggest ‘personality’ adaptation can be an effective, minimal bio-inspired approach to learning, suitable for unpredictable real-world environments and potentially complementary to the above approaches. The notion of personality maps readily to adaptive threshold-based behaviours: for example, the likelihood of switching behaviours in probabilistic finite state machines (e.g. [14]). Animal personality research can indicate simple behavioural mechanisms (‘interaction rules’) that adaptively shape personality. For example, the frequency of social interactions can relate to boldness changes [15]. Such rules can be transposed into embodied and/or virtual artificial agents working in collectives. Principle advantages of this approach could include:

- Various major personality axes could be relevant to groups of field robots, and these can each be represented as a single variable.
- Personality variables can adapt quickly and predictably to changing environments.
- The notion of personality is intuitive for human users.
- Individual experience of the environment will lead to a certain amount of group-level diversity in personality variables, the advantages of which I outline in Section III.

Plasticity occurs in response to local environmental cues, so one must also consider the relevant environmental features (physical and social) that will elicit change – and how they will be sensed. For example, one may wish to use agent density as a proxy for group size. This could be sensed via frequency of physical interactions, which in turn could be correlated with risk-taking behaviour [15], if one prefers individuals in larger groups to take on a higher risk appetite.

III. ADVANTAGES OF BEHAVIOURAL DIVERSITY

A group of identically-programmed (homogeneous) agents deployed into a variable (heterogeneous) environment will each experience different conditions, to a greater or lesser extent. Each may benefit from individual adaptation; or indeed, individual differences may occur which have little significance when considered in relation to the *direct* ‘fitness’ of a single individual. Nevertheless, at the level of the collective, such heterogeneity may make an important indirect contribution to the fitness of the swarm, as part of an adaptive collective phenotype [16]. As such, the group-level distribution of personality traits may be self-organized through interactions with the environment and others to favour a certain ecologically relevant pattern [15], as I go on to illustrate.

A. Animal collectives

1) *Diversity for decision-making*: Autonomous collectives – whether in biology, robotics or elsewhere – need to be capable of making collective decisions. Diversity of reaction thresholds or option assessment behaviour, as seen in ants, can help this process [4], [17]. As an example of this, a proportion of individuals with high acceptance thresholds may reject medium-quality options, and thus through continued exploration go on to discover higher-quality possibilities.

2) *Diversity for homeostasis*: In biological systems phenotypic diversity can also promote positive collective success: for example, in honeybees diversity in reaction thresholds for their cooling behaviour promotes stability in nest thermoregulation [18]. Although this example is driven by corresponding genetic heterogeneity, it could equally be designed in an artificial agent context to result from phenotypic plasticity.

In *Stegodyphus* social spiders the group-level distribution of boldness is important for their collective predation performance [19]. Thus, one expects some form of mechanism to maintain a suitable collective boldness phenotype, and indeed there is evidence for a link between social interactions and boldness change to achieve this via self-organization [15].

3) *Diversity as a shield against adversity*: Inbreeding in agriculture is observed as a cause of disease vulnerability, because a single pathogen virulent against one individual can quickly spread across the whole population; conversely, diversity can help resistance [20].

Similarly, robustness is frequently claimed for swarm robot systems, but if a homogeneous controller results in homogeneous behaviour it may be liable to systematic failure if the swarm encounters unexpected environmental conditions or faulty or malicious agents.

4) *Diversity for foraging and search*: Variation in individual behaviour can also be important for effective foraging and search, as demonstrated for example by Fricke et al. in immune-system-inspired search algorithms [21]. If search targets have heterogeneous configurations (for example, sometimes low density, other times high density) a collective of agents will be more effective if individuals behave differently.

B. Human teams

There is evidence that teams with more diverse personality types are more effective: for example, Mohammed & Angell found that higher variability in extraversion results in higher task performance [22]. In this context I suggest that there are opportunities to enhance collective intelligence in human-AI interaction. Virtual agents and robots could benefit from suitably plastic personalities to adapt and complement the shortcomings or absence of relevant behavioural types in their human teams. This could be done as virtual team members or as adaptive social assistance robots [23], working as facilitators, contributing to a successful ‘hybrid’ team phenotype.

IV. CONCLUSIONS

As we deploy robots into real-world field conditions, equipping them with smart behavioural reaction norms – *plastic personalities* – could help them to work effectively in unpredictable conditions: individually, and as a team. Advantages could include simplicity (single personality variables), transparency, intuitiveness, predictability, and automatic diversity in multi-robot teams from local experience. Future work will focus on demonstration of these concepts in simulation, before experiments on real world platforms such as unmanned ground vehicles (UGVs) undertaking tasks such as search and decision-making in unpredictable environments.

REFERENCES

- [1] E. R. Hunt, "Phenotypic plasticity provides a bioinspiration framework for minimal field swarm robotics," *Frontiers in Robotics and AI*, vol. 7, p. 23, 2020. [Online]. Available: <https://www.frontiersin.org/articles/10.3389/frobt.2020.00023/abstract>
- [2] S. R. Dall, A. M. Bell, D. I. Bolnick, and F. L. Ratnieks, "An evolutionary ecology of individual differences," *Ecology letters*, vol. 15, no. 10, pp. 1189–1198, 2012. [Online]. Available: <https://doi.org/10.1111/j.1461-0248.2012.01846.x>
- [3] D. Réale, S. M. Reader, D. Sol, P. T. McDougall, and N. J. Dingemanse, "Integrating animal temperament within ecology and evolution," *Biological Reviews*, vol. 82, no. 2, pp. 291–318, 2007. [Online]. Available: <https://doi.org/10.1111/j.1469-185X.2007.00010.x>
- [4] N. Masuda, T. A. O'Shea-Wheller, C. Doran, and N. R. Franks, "Computational model of collective nest selection by ants with heterogeneous acceptance thresholds," *Royal Society Open Science*, vol. 2, no. 6, p. 140533, 2015. [Online]. Available: <https://doi.org/10.1098/rsos.140533>
- [5] C. D. Schlichting and M. Pigliucci, *Phenotypic evolution: a reaction norm perspective*. Sunderland, Massachusetts: Sinauer Associates Incorporated, 1998.
- [6] N. J. Dingemanse, A. J. Kazem, D. Réale, and J. Wright, "Behavioural reaction norms: animal personality meets individual plasticity," *Trends in Ecology & Evolution*, vol. 25, no. 2, pp. 81–89, 2010. [Online]. Available: <https://doi.org/10.1016/j.tree.2009.07.013>
- [7] J. M. Jandt, S. Bengston, N. Pinter-Wollman, J. N. Pruitt, N. E. Raine, A. Dornhaus, and A. Sih, "Behavioural syndromes and social insects: personality at multiple levels," *Biological Reviews*, vol. 89, no. 1, pp. 48–67, 2014. [Online]. Available: <https://doi.org/10.1111/brv.12042>
- [8] V. C. Norman, T. Pamminer, and W. O. Hughes, "The effects of disturbance threat on leaf-cutting ant colonies: a laboratory study," *Insectes Sociaux*, vol. 64, no. 1, pp. 75–85, 2017. [Online]. Available: <https://doi.org/10.1007/s00040-016-0513-z>
- [9] E. Ferrante, A. E. Turgut, E. Duéñez-Guzmán, M. Dorigo, and T. Wenseleers, "Evolution of Self-Organized Task Specialization in Robot Swarms," *PLOS Computational Biology*, vol. 11, no. 8, p. e1004273, aug 2015. [Online]. Available: <https://doi.org/10.1371/journal.pcbi.1004273>
- [10] S. Jones, A. F. Winfield, S. Hauert, and M. Studley, "Onboard Evolution of Understandable Swarm Behaviors," *Advanced Intelligent Systems*, vol. 1, no. 6, p. 1900031, 2019. [Online]. Available: <https://doi.org/10.1002/aisy.201900031>
- [11] N. Jakobi, P. Husbands, and I. Harvey, "Noise and the reality gap: The use of simulation in evolutionary robotics," in *European Conference on Artificial Life*. Springer, 1995, pp. 704–720.
- [12] M. Hüttenrauch, A. Šošić, and G. Neumann, "Local Communication Protocols for Learning Complex Swarm Behaviors with Deep Reinforcement Learning," in *Swarm Intelligence. ANTS 2018. Lecture Notes in Computer Science, vol 11172.*, M. Dorigo, M. Birattari, C. Blum, A. L. Christensen, A. Reina, and V. Trianni, Eds. Springer, Cham, 2018, pp. 71–83. [Online]. Available: https://doi.org/10.1007/978-3-030-00533-7_6
- [13] M. Brambilla, E. Ferrante, M. Birattari, and M. Dorigo, "Swarm robotics: A review from the swarm engineering perspective," *Swarm Intelligence*, vol. 7, no. 1, pp. 1–41, 2013. [Online]. Available: <https://doi.org/10.1007/s11721-012-0075-2>
- [14] E. Castello, T. Yamamoto, F. D. Libera, W. Liu, A. F. Winfield, Y. Nakamura, and H. Ishiguro, "Adaptive foraging for simulated and real robotic swarms: the dynamical response threshold approach," *Swarm Intelligence*, vol. 10, no. 1, pp. 1–31, 2016. [Online]. Available: <https://doi.org/10.1007/s11721-015-0117-7>
- [15] E. R. Hunt, B. Mi, C. Fernandez, B. M. Wong, J. N. Pruitt, and N. Pinter-Wollman, "Social interactions shape individual and collective personality in social spiders," *Proceedings of the Royal Society B: Biological Sciences*, vol. 285, no. 1886, p. 20181366, 2018. [Online]. Available: <https://doi.org/10.1098/rspb.2018.1366>
- [16] P. Kennedy, G. Baron, B. Qiu, D. Freitak, H. Helanterä, E. R. Hunt, F. Manfredini, T. O'Shea-Wheller, S. Patalano, C. D. Pull, T. Sasaki, D. Taylor, C. D. Wyatt, and S. Sumner, "Deconstructing Superorganisms and Societies to Address Big Questions in Biology," *Trends in Ecology & Evolution*, vol. 32, no. 11, pp. 861–872, nov 2017. [Online]. Available: <https://doi.org/10.1016/j.tree.2017.08.004>
- [17] T. A. O'Shea-Wheller, N. Masuda, A. B. Sendova-Franks, and N. R. Franks, "Variability in individual assessment behaviour and its implications for collective decision-making," *Proceedings of the Royal Society B: Biological Sciences*, vol. 284, no. 1848, p. 20162237, feb 2017. [Online]. Available: <https://doi.org/10.1098/rspb.2016.2237>
- [18] J. C. Jones, M. R. Myerscough, S. Graham, and B. P. Oldroyd, "Honey Bee Nest Thermoregulation: Diversity Promotes Stability," *Science*, vol. 305, no. 5682, pp. 402–404, 2004. [Online]. Available: <https://doi.org/10.1126/science.1096340>
- [19] E. R. Hunt, B. Mi, R. Geremew, C. Fernandez, B. M. Wong, J. N. Pruitt, and N. Pinter-Wollman, "Resting networks and personality predict attack speed in social spiders," *Behavioral Ecology and Sociobiology*, vol. 73, no. 7, p. 97, 2019. [Online]. Available: <https://doi.org/10.1007/s00265-019-2715-7>
- [20] L. V. Ugelvig, D. J. Kronauer, A. Schrempf, J. Heinze, and S. Cremer, "Rapid anti-pathogen response in ant societies relies on high genetic diversity," *Proceedings of the Royal Society B: Biological Sciences*, vol. 277, no. 1695, pp. 2821–2828, sep 2010. [Online]. Available: <https://doi.org/10.1098/rspb.2010.0644>
- [21] G. M. Fricke, J. P. Hecker, J. L. Cannon, and M. E. Moses. [Online]. Available: <https://doi.org/10.1017/S0263574716000382>
- [22] S. Mohammed and L. C. Angell, "Personality Heterogeneity in Teams: Which Differences Make a Difference for Team Performance?" *Small Group Research*, vol. 34, no. 6, pp. 651–677, 2003. [Online]. Available: <https://doi.org/10.1177/1046496403257228>
- [23] K. Winkle, S. Lemaignan, P. Caleb-Solly, U. Leonards, A. Turton, and P. Bremner, "Effective Persuasion Strategies for Socially Assistive Robots," in *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 2019, pp. 277–285. [Online]. Available: <https://doi.org/10.1109/HRI.2019.8673313>

Reliability-Aware Multi-UAV Coverage Path Planning Using Integer Linear Programming

Mickey Li
Bristol Robotics Laboratory
University of Bristol
Bristol, UK
mickey.li@bristol.ac.uk

Arthur Richards
Bristol Robotics Laboratory
University of Bristol
Bristol, UK
arthur.richards@bristol.ac.uk

Mahesh Sooriyabandara
Bristol Research and Innovation Laboratory
Toshiba Research Europe Ltd
Bristol, UK
mahesh@toshiba-trel.com

<https://doi.org/10.31256/Cy5Ej9K>

Abstract—Multi-Agent Systems have the potential advantage of graceful degradation over Single-Agent systems. This is a desirable trait in applications that require area coverage with failure-prone UAVs. This paper uses methods informed by Reliability Engineering to study this formally by using a new stochastic framework which evaluates a strategy's probability of task completion. Based on this analysis, an integer linear programming formulation of the problem is then shown to provide almost optimal strategies at a fraction of the computational cost of brute force methods.

Index Terms—Multi-Robot Systems, Coverage Path Planning, Reliability Analysis

I. INTRODUCTION

Multi-agent systems are well known to provide greater reliability, resilience and fault tolerance to individual and system level failures compared to single-agent systems [1, 2]. This robustness property makes a multi-UAV system especially suited to the problem of area coverage and inspection, as UAVs are particularly prone to failure [3–5]. Whilst [6, 7] consider failure handling, they only assert eventual completion. This work takes a probabilistic planning approach to model failures to provide temporal guarantees of completion. Figure 1 motivates this problem - by what evaluation method is strategy 2 optimal? The aim of this paper is to present (1) an exact reliability evaluation framework, inspired by the field of Reliability Engineering [8], (2) a computationally feasible approximate method, for multi-drone path planning which takes into account the uncertainty in failure of individual drones, such that the *Probability of Mission Completion (PoC)* is maximised. The resultant plan is thus the Pareto-Optimal strategy with respect to reliability and efficiency.

II. MARKOV RELIABILITY ANALYSIS FRAMEWORK

The state of the system of n agents \mathbf{A} at a time t is $x^t = (\tau_1^t, \dots, \tau_n^t) \in \mathbb{N}^n = \mathbf{S}$ where τ_i^t is the length of time agent i has survived (i.e. the useful time worked) under the following Markov transition rule:

$$\tau_i^t = \begin{cases} \tau_i^{t-1} + 1 & \text{if agent } i \text{ survives} \\ \tau_i^{t-1} & \text{if agent } i \text{ fails or already failed} \end{cases} \quad (1)$$

This work has been funded by the Engineering and Physical Sciences Research Council (EPSRC) iCASE with Toshiba Research Europe Ltd and FARSCOPE Centre of Doctoral Training at the Bristol Robotics Laboratory




Fig. 1. Consider 2 drones covering a 1D cyclic path. Agent 1 suffers a failure. Should Agent 2 follow strategy 1 or 2 such that the greatest probability of completion is achieved? Intuitively, it must be strategy 2, but why?

A set of m spatially connected, discrete tasks in an environment is defined as $\mathbf{J} = (j_1, \dots, j_m)$. A **strategy** is then defined as a function $\pi_i(x, t) = j$, $x \in \mathbf{S}$ and $j \in \mathbf{J}$ where j is the task performed by agent i at time t given the health states of all agents x .

This framework is described with respect to **static** strategies: pre-allocated fixed paths for each agent with no dependencies on the state \mathbf{S} . The strategy matrix $T \in \mathbb{N}^{n \times m}$ can be defined where T_{ij} = the time at which agent i visits task j . T can be mapped onto $\pi_i(x, t)$, for any time t and agent i , by finding the task j such that $T_{ij} = \tau_i$. A task j is **completed** by agent i if agent i survives longer than the time at which the drone is scheduled to visit task j , i.e $\tau_i \geq T_{ij}$.

A state $x \in \mathbf{C} \subseteq \mathbf{S}$ is a **coverage completion state** iff:

$$F(x) = \min_{j \in [1..m]} E_j(x) = \min_{j \in [1..m]} \max_{i \in [1..n]} \tau_i - T_{ij} \geq 0 \quad (2)$$

as the **time since completion of task** j , $E_j(x) \geq 0$ for all tasks \mathbf{J} and therefore every task has been visited.

Given the failure probability density $f_i(\tau) = p_i(t = \tau)$ for each agent, the **probability of completion (PoC)** for a given strategy T at a particular time t' is the sum of the probabilities of surviving until each of the completion states of t' , $\mathbf{C}_T(t')$.

$$\mathbf{C}_T(t') = \{x \in \mathbf{C} | \forall \tau \in x, \tau \leq t'\} \quad (3)$$

$$PoC(T, t') = \sum_{x \in \mathbf{C}_T(t')} \prod_{\tau_i \in x} p_i(\tau_i, t') \quad (4)$$

$$p_i(\tau, t') = \begin{cases} p_i(t = \tau) = f_i(\tau), & \text{if } \tau < t' \\ p_i(t > \tau) = R_i(\tau), & \text{if } \tau \geq t' \end{cases} \quad (5)$$

Equation (5) describes the case where the agent fails before t' and where the agent survives at t' . $R_i(\tau)$ is known as **Reliability. Failure Rate** $\lambda_i = \frac{f_i(\tau)}{R_i(\tau)}$ is a system parameter.




Fig. 2. The PoC of all cyclic strategies for a path of length 12, with 2 agents, $\lambda = 0.1$. Top performing strategies highlighted. Legend describes strategies (initial task, direction) w.r.t one of the agents starting at 0 travelling right.




Fig. 3. The difference in PoC between optimal strategies reported by the computationally feasible brute force searches and the ILP. (#drones, #path)

Considering the state as ‘work done’, instead of explicit locations of agents and completion of tasks, allows the decoupling of the state analysis from the probability analysis.

Complex **non-static** strategies may be defined based on agent failures, but are difficult to formalise. In principle, the static framework can still be applied by running a unique simulation of π for every $x \in \mathbf{S}$ and labelling completion. In general this is non-trivial apart from the ‘teleportation’ strategy where task spatial constraints are not considered. Teleportation can be proven to be the optimal strategy.

The framework can be used to find the Pareto-Optimal strategy for a fixed time horizon in a brute force manner by evaluating every strategy for a given number of drones and path. Figure 2 shows the optimal strategies found via brute force search for the motivating problem in Figure 1. The strategy ‘[11,L]’ corresponds to Strategy 2 and provides a greater PoC by end time 11 than Strategy 1, ‘[5, R]’, which agrees with intuition.

III. INTEGER LINEAR PROGRAMMING APPROACH

The brute force method, whilst exact, is only computationally feasible for small scenarios (e.g. 4 agents, 20 path takes a week). An Integer Linear Programming (ILP) formulation is presented to provide an approximation of the optimal strategy, whilst being computationally tractable.

A strategy is represented by the binary variable X where $X_{ij}(t) = 1$ iff agent i visits task j at t . Instead of maximising PoC directly, the ILP minimises the maximum log-probability of any single task j being missed $\ln p(m_j)$. $p(m_j)$ is then the probability of all agents independently failing before their

scheduled visits of j .

$$\text{minimize}_X \max_j \ln \prod_i \prod_t F(t)^{X_{ij}(t)} \quad (6)$$

$$= \max_j \sum_i \sum_t X_{ij}(t) \ln F(t) \quad (7)$$

$$\text{subject to} \sum_j X_{ij}(t) = 1 \quad \forall i, t \quad (8)$$

$$\sum_i X_{ij}(t) \leq 1 \quad \forall j, t \quad (9)$$

$$\sum_i \sum_k X_{ij}(t) \geq 1 \quad \forall j \quad (10)$$

$$X_{ij}(t) \leq \sum_{j' \in \mathcal{N}(j)} X_{ij'}(t-1) \quad \forall i, j, t \quad (11)$$

$$\sum_t X_{ij}(t) \leq 1 \quad \forall i, j \quad (12)$$

Where (8) ensures every agent can only do 1 task at each time. (9) ensures every task is visited at some time and (10) ensures every task is visited at some time by some agent. (11) constrains agent movement to neighbouring tasks. Finally (12) stops an agent from visiting a task more than once.

By minimising $\max_j p(m_j)$, the ILP makes it unlikely that any given task is missed. Intuitively, this makes completion more likely, and it can be shown that it maximises a lower bound on PoC. It is not equivalent to a direct optimisation of PoC as it does not capture task inter-dependencies. Nevertheless, experiments suggest that the ILP selects strategies whose PoC comes close to the optimal PoC, found by exhaustive evaluation.

The PoC of the ILP’s chosen strategy is calculated by converting X into a static T matrix for the Markov Framework. To evaluate the accuracy, the ILP is applied to cyclical paths of various lengths and numbers of drones using the Gurobi LP solver [9]. Figure 3 shows the difference between the PoC of the strategy obtained through brute force evaluation compared to the outputted strategy of the ILP. The ILP appears to provide a very close PoC with its strategy choice, only a small deviation for longer paths. This is promising as close to exact accuracy can be obtained from a method which takes a fraction of the time (e.g. 4 agents, 20 path takes minutes).

IV. CONCLUSION

This paper proposes that for applications where knowing all tasks are completed are crucial, but the hardware is prone to failure, reliability may be favoured over efficient plans. Inspired by Reliability Engineering, a probabilistic evaluation framework is presented which reports on the probability of completion of a mission given a strategy. The Pareto-Optimal strategy can be found through brute force evaluation of all strategies, but this is computationally infeasible. An Integer Linear Programming approach is shown to provide almost optimal results for a fraction of the computation time. Future work will include exploring non-linear optimisation approaches as well as considering more complex drone failure models such as localisation and communication failures.

REFERENCES

- [1] G. Weiss, *Multiagent systems, Second Edition*. MIT Press, 2013. ISBN 9780262018890. [Online]. Available: <https://mitpress.mit.edu/books/multiagent-systems-second-edition>
- [2] H. Hamann, *Swarm Robotics: A Formal Approach*. Cham: Springer International Publishing, 2018. ISBN 978-3-319-74526-8. [Online]. Available: <http://link.springer.com/10.1007/978-3-319-74528-2>
- [3] V. Kumar and N. Michael, "Opportunities and challenges with autonomous micro aerial vehicles," *The International Journal of Robotics Research*, vol. 31, no. 11, pp. 1279–1291, 9 2012. doi: 10.1177/0278364912455954. [Online]. Available: <http://journals.sagepub.com/doi/10.1177/0278364912455954>
- [4] S. Gupte, P. I. T. Mohandas, and J. M. Conrad, "A survey of quadrotor Unmanned Aerial Vehicles," in *2012 Proceedings of IEEE Southeastcon*. IEEE, 3 2012. doi: 10.1109/SECon.2012.6196930. ISBN 978-1-4673-1375-9 pp. 1–6. [Online]. Available: <http://ieeexplore.ieee.org/document/6196930/>
- [5] S.-J. Chung, A. A. Paranjape, P. Dames, S. Shen, and V. Kumar, "A Survey on Aerial Swarm Robotics," *IEEE Transactions on Robotics*, vol. 34, no. 4, pp. 837–855, 8 2018. doi: 10.1109/TRO.2018.2857475. [Online]. Available: <https://ieeexplore.ieee.org/document/8424838/>
- [6] N. Hazon and G. A. Kaminka, "On redundancy, efficiency, and robustness in coverage for multiple robots," *Robotics and Autonomous Systems*, vol. 56, no. 12, pp. 1102–1114, 12 2008. doi: 10.1016/J.ROBOT.2008.01.006. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0921889008000122>
- [7] P. Fazli, A. Davoodi, P. Pasquier, and A. K. Mackworth, "Complete and robust cooperative robot area coverage with limited range," in *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 10 2010. doi: 10.1109/IROS.2010.5651321. ISBN 978-1-4244-6674-0 pp. 5577–5582. [Online]. Available: <http://ieeexplore.ieee.org/document/5651321/>
- [8] P. D. T. O'Connor and A. Kleyner, *Practical reliability engineering*. Wiley, 2012. ISBN 9780470979815
- [9] Gurobi Optimization LLC, "Gurobi Optimizer Reference Manual," 2020. [Online]. Available: <http://www.gurobi.com>

Towards Intention Recognition for Human-Interacting Agricultural Robots

Alexander Gabriel

Lincoln Centre for Autonomous Systems
University of Lincoln
Lincoln, United Kingdom
agabriel@lincoln.ac.uk

Paul Baxter

Lincoln Centre for Autonomous Systems
University of Lincoln
Lincoln, United Kingdom
pbaxter@lincoln.ac.uk

<https://doi.org/10.31256/Ye5Nz9W>

Abstract—Robots sharing a common working space with humans and interacting with them to accomplish some task should not only optimise task efficiency, but also consider the safety and comfort of their human collaborators. This requires the recognition of human intentions in order for the robot to anticipate behaviour and act accordingly. In this paper we propose a robot behavioural controller that incorporates both human behaviour and environment information as the basis of reasoning over the appropriate responses. Applied to Human-Robot Interaction in an agricultural context, we demonstrate in a series of simulations how this proposed method leads to the production of appropriate robot behaviour in a range of interaction scenarios. This work lays the foundation for the wider consideration of contextual intention recognition for the generation of interactive robot behaviour.

Index Terms—Agricultural Robotics, Human-Robot Interaction, Intention Recognition, Belief-Desire-Intention System

I. INTRODUCTION

Introducing robots into a human working space can increase efficiency but should not come at the cost of comfort or safety. To achieve this balance in a challenging setting like agriculture, a robot needs to understand the intentions behind their coworkers' behaviour and basic communication. Gestures form an ideal medium to maintain reliability in adverse circumstances but are limited to situations where the human has their hands free. Additional clues from the environment as well as behaviour analysis can be used to estimate their state. Our interpretation of intentions [5] sees them as the meaning [4] of, explanation [7] for, or idea [9] behind an action, plan or utterance. In our agricultural setting, workers pick berries into crates in a poly-tunnel environment. The robot is acting in a supporting role, supplying the human with empty crates, taking away full crates and staying out of the way the rest of the time. To facilitate the robot's autonomy, we created an integrated sensor data processing pipeline and Belief-Desire-Intention (BDI) [1] agent system. The general motivation for this system is that in order to 'understand' the intentions of their human interaction partner (from observable behaviour) and to generate appropriate responses, the robot should consider both the environmental context but also its own goals (or 'desires'); this supports our use of a BDI

Supported by the RASberry project (<https://rasberryproject.com>), and the CTP for Fruit Crop Research (www.ctp-fcr.org).

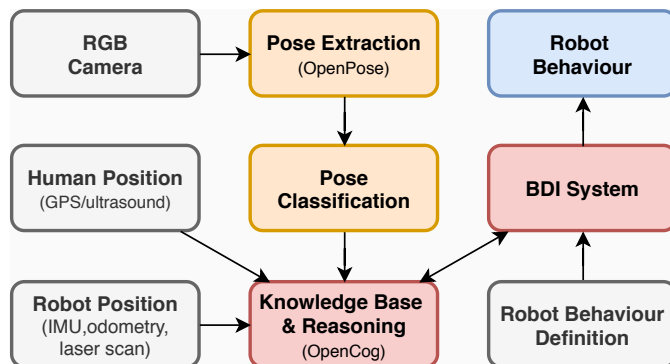


Fig. 1. HRI System Representation Grey: Datasources, Orange: Pre-Processing Red: Memory, Reasoning, and Agent Blue: Output

architecture. This work summarizes a first evaluation of the system's performance in a range of scenarios. We evaluated the system in simulation, using sensor readings as well as statistical data gathered in real-world poly-tunnels.

II. HUMAN-ROBOT INTERACTION (HRI) SYSTEM

A. Data Processing: The robot perceives its environment through a stereo RGB-D camera, a thermal camera, 2D and 3D LIDAR, as well as differential GPS and odometry. It can also receive its coworkers' location either provided by GPS or ultrasonic localization and is supplied with predefined topological and laser maps. During this simulation, the robot determines its position using simulated laser scans and odometry. The location of the human is supplied by a picker-simulation engine and abstracted using Qualitative Trajectory Calculus (QTC) [8].

As shown in Fig. 1, the video is first pre-processed using OpenPose [2] to extract joint positions. Those are further processed to extract joint angles before both are fed into a naive classifier that produces pose labels for each frame individually, based on predefined prototype poses. Series of frames are classified using voting rounds where each frame contributes a single vote towards a pose. Whichever pose first wins 10 votes, labels the round. The movement samples used in this evaluation are part of a new dataset for Action Recognition in agri-robotics. It contains samples of behaviours, such as

picking of berries and carrying of crates, and samples of gestures to communicate with the robot. The samples were recorded from 10 different subjects between morning and early afternoon in a poly-tunnel environment featuring ripe strawberry plants.

B. Belief-Desire-Intention Agent: The BDI system chooses intentions (plans to reach a goal) from its desires (abstract goals) based on its beliefs as captured in the Knowledge Base (KB)¹. This separates reasoning about which goals to achieve from managing the execution of said goals. This allows us to consider more contextual information when deciding which goal to follow and leads to an aesthetic analog to our idea of human motivation, intention and action on the robot.

In our system, plans are represented as ordered tree structures with executable actions as leafs. Actions have a set of preconditions and expected consequences, which combine to form the preconditions and expected consequences of a plan. When the agent decides which of its desires are applicable in a given situation, it searches the KB for patterns of beliefs that match its desire's preconditions. If successful, it produces a corresponding intention. When idle, the robot chooses from the possible next actions defined by its current intentions, based on utility and expected time requirements.

III. EVALUATION

In the context of interactions between a robot and a human in an agricultural context, we explore three scenarios (Fig. 2). They cover different situations in the work environment: starting to work (crate delivery), moving around (evading the human), and resupply (crate exchange). Both the delivery and exchange scenarios are initiated by the human gesturing to the robot, but require a different response. The robot can make the distinction based on prior observed human behavior (whether the person has been picking berries). The evasion scenario is triggered by human behaviour (approaching) without any conscious interaction.

¹We are using OpenCog's [6] AtomSpace and Pattern matcher for knowledge representation and reasoning.




Fig. 2. **Left:** Delivery Scenario: Meet the human, wait for 2.5 seconds, leave. **Middle:** Evasion Scenario: On human approach, move to the next waypoint. **Right:** Exchange Scenario: Meet the human, wait for 5 seconds, leave.

TABLE I
EXPERIMENT RESULTS

Scenario	Success Rate	Meeting Distance [m]		Time to Service [s]	
		μ	σ	μ	σ
Delivery	0.99	0.37	0.007	12.65	22.552
Evasion	1.00		N/A	10.06	1.706
Exchange	0.99	0.37	0.006	11.83	0.634

The recorded human behaviours in poly-tunnels (Section II), are passed as input to the proposed system (Fig. 1), acting on a simulated environment. Given 10 recorded subjects, 20 simulations per subject are performed (given a stochastic simulation), resulting in 200 simulations per scenario.

Table I shows for each scenario the mean (μ) and standard deviation (σ) of the three metrics of evaluation: success rate, meeting distance and time to service. Success rate is defined as the share of experiment runs that ended with the robot successfully interpreting the situation and performing the expected actions. Meeting distance is the distance between human and robot at which the robot decided to halt to facilitate the delivery or exchange of a crate. The variance for this metric can be interpreted as an indicator of how much reasoning affects the agent's reaction time (the robot and human don't meet in the Evasion case). Time to Service is the time between the human displaying behaviour that should trigger a change in robot behaviour, and the time at which the robot performed the expected action (delivered or exchanged a crate, or moved away from the human to the next waypoint). This time consists mainly of the time it takes to detect the behavior, the time it takes to meet the human ($\sim 4s$), and the time for the delivery (2.5s) or exchange (5s) of crates. The large variance for the Delivery scenario stems from the robot's failure to detect one subject's behavior correctly. Without these cases, the variance for the Delivery scenario is 0.298.

IV. CONCLUSION

Our evaluation shows benefits compared to a system that is unaware of its human coworkers except for their location and service calls. However, they do not directly support any claim of advantage over a simpler, reaction-based, human-aware system. The hypothesis is that by using intention recognition to anticipate the human's requirements, there will be observable benefits in both task efficiency and perceived comfort of the interaction, but this requires validation.


The service times of under 15s are short in comparison to the $110 \pm 44s$ [3] it takes for the robot to service the picker when stationed outside the field. This points at a potential increase in productivity, achievable by estimating the next time a worker requires the robot's service and arriving in the vicinity early (another source of information to facilitate anticipatory robot behaviour).

Overall, while there remain a number of outstanding challenges, this paper has outlined the fundamentals of our approach: the consideration of both observable human behaviour and the wider environmental context in supporting anticipatory interactive robot behaviour.

REFERENCES

- [1] M. Bratman et al. *Intention, plans, and practical reason*, volume 10. Harvard University Press Cambridge, MA, 1987.
- [2] Z. Cao et al. Realtime multi-person 2D pose estimation using part affinity fields. *Proceedings - 30th IEEE Conference on Computer Vision and Pattern Recognition*, 2017-Januar:1302–1310, 2017.
- [3] G. Das, G. Cielniak, et al. Discrete event simulations for scalability analysis of robotic in-field logistics in agriculture – a case study. In *IEEE International Conference on Robotics and Automation, Workshop on Robotic Vision and Action in Agriculture (WRVAA)*, 2018.
- [4] M. Fleischman and D. Roy. Why Verbs are Harder to Learn than Nouns: Initial Insights from a Computational Model of Intention Recognition in Situated Word Learning. In *Proc. of the Annual Meeting of the Cognitive Science Society*, 2005.
- [5] A. Gabriel, S. Coşar, N. Bellotto, and P. Baxter. A dataset for action recognition in the wild. In *Annual Conference Towards Autonomous Robotic Systems*, pages 362–374. Springer, 2019.
- [6] B. Goertzel, C. Pennachin, and N. Geisweiller. The opencog framework. In *Engineering General Intelligence, Part 2*, pages 3–29. Springer, 2014.
- [7] K. A. Tahboub. Intelligent human-machine interaction based on Dynamic Bayesian Networks probabilistic intention recognition. *Journal of Intelligent and Robotic Systems: Theory and Applications*, 45(1):31–52, 2006.
- [8] N. Van de Weghe et al. A qualitative trajectory calculus as a basis for representing moving objects in geographical information systems. *Control and cybernetics*, 35(1):97–119, 2006.
- [9] S. Youn and K. Oh. Intention recognition using a graph representation. *International Journal of Applied Science, Engineering and Technology*, 4(1):13–18, 2007.

Feasibility Study of In-Field Phenotypic Trait Extraction for Robotic Soft-Fruit Operations

1st Raymond Kirk 
Lincoln Centre for Autonomous Systems
University of Lincoln
Lincoln, UK
rkirk@lincoln.ac.uk

2nd Michael Mangan
Sheffield Robotics
University of Sheffield
Sheffield, UK
m.mangan@sheffield.ac.uk

3rd Grzegorz Cielniak
Lincoln Centre for Autonomous Systems
University of Lincoln
Lincoln, UK
gcielniak@lincoln.ac.uk

<https://doi.org/10.31256/Uk4Td6I>

Abstract—There are many agricultural applications that would benefit from robotic monitoring of soft-fruit, examples include harvesting and yield forecasting. Autonomous mobile robotic platforms enable digitisation of horticultural processes in-field reducing labour demand and increasing efficiency through continuous operation. It is critical for vision-based fruit detection methods to estimate traits such as size, mass and volume for quality assessment, maturity estimation and yield forecasting. Estimating these traits from a camera mounted on a mobile robot is a non-destructive/invasive approach to gathering qualitative fruit data in-field. We investigate the feasibility of using vision-based modalities for precise, cheap, and real time computation of phenotypic traits: mass and volume of strawberries from planar RGB slices and optionally point data. Our best method achieves a marginal error of 3.00cm^3 for volume estimation. The planar RGB slices can be computed manually or by using common object detection methods such as Mask R-CNN.


Index Terms—phenotyping, mobile robots, computer vision

I. INTRODUCTION

Fruit detection is an area fast gaining interest in the horticultural industry. The environmental challenges posed by the fast growing population and climate concerns are spurring new innovative approaches to fruit detection, harvesting and yield estimation using computer vision e.g [1]–[3]. Phenotypic information such as volume shown in 1 about the fruit is important for all of these approaches. For harvesting it allows to automatically grade and harvest specific berry classifications, and for yield more specific estimates such as detection of waste strawberries or estimating a total yield volume can be computed. Phenotypic information is critical for any form of quality assessment.

Our method aims to estimate mass and volume of soft-fruit from robotic platforms in-field. We present a feasibility study in lab conditions for estimating these traits from images based on the intuition that most soft-fruits are ellipsoidal in nature and symmetrical around their major-axis. Meaning the methods presented are applicable to most of the soft-fruit family. Geometrically the major axis is the longer axis of an ellipse passing through its foci or centre of gravity in the case of our planar segment; minor axis is the shorter axis directly perpendicular to the major.

This work was partially funded by the RASberry project.



(a) Actual Volume = 35.00cm^3 (b) Predicted Volume = 34.53cm^3

Fig. 1: Strawberry volume prediction, RGB image (1a), computed reconstruction surface of RGB segment (1b)

II. DATA AND METHODS

In order to evaluate our methods, we required mass and volumetric data of soft-fruit. We chose to evaluate strawberries as they are readily available and have one of the most challenging shapes in the soft-fruit family compared to blackberries, blueberries etc. their surface is not as ellipsoidal and has a more teardrop profile. We collected 20 samples of class 1 ripe strawberries. To capture the data necessary, we used a 2cm^3 precision volumetric beaker, a 5g accurate scale, a 0.01mm accurate digital caliper and an Intel Realsense D415 computer vision camera to capture RGB images and depth information, pictured in Figure 2. Each strawberry was measured in three dimensions manually through its minor, major and cross sections which are the widest, tallest and deepest lengths of the berry respectively. Then it was weighed and placed in the volumetric breaker, a control rod of a known volume was used to fully submerge the berry to get more accurate readings. Finally, the berry was placed at a set distance away from the downwards facing camera, flat on a table to simulate the conditions met in field and the RGB and depth information was captured and logged.

A. Volume Estimation

A segment (planar RGB slice) is a binary mask detailing all of the pixels that belong to an object in an image. We use



Fig. 2: Equipment used for data collection.

these segments to estimate the volume of the strawberries. The computational resources required to process these segments are very low and are a typical output of modern object detectors in this field, meaning this approach is easily integrated with existing work with negligible overhead. We present the results in Figure 3.

The three evaluated methods are ellipsoidal, surface area integration and disc summation. The ellipsoidal method computes the volume as $\frac{4}{3}\pi m_i m_a d$ where m_i is the minor axis, m_a is the major and d is the cross section length. These volume measurements are computed from both the ground truth (GT) data and measurements extracted from the depth map. The method we deem surface area integration uses the fundamental relationship in calculus that states the integral of a function f over an interval can be calculated by finding an anti-derivative F of f . For an ellipsoid the volume is the integral of the surface area with respect to the radius.

$$f(c, r) = \frac{r}{m_i} \left(\frac{\sum_{k=0}^n (c_{xk}, c_{yk})}{n} - c \right) + c \quad (1)$$

$$V = 2 \int_0^r 2a(f(c, r))dx \quad (2)$$

In Equation (2) we show the integral for computing the volume V of an irregular segment, by taking the product of dx , the height of each slice and the contour c . We scale each slice by each slice radius r in function $f(c, r)$ (1) and calculate its area $a(f(c, r))$. The function $a(x)$ is the shoelace algorithm for finding area of simple polygon (no intersection or holes) expressed as Cartesian coordinates of a segment. We use the integral range $[0, r]$ and multiply the result by 2 to only consider positive contour values.

$$v_i = \sum_{j=1}^n c_{ij} \quad v = \left(\pi \frac{v_i^2}{4} d_y \right)_{i=1}^n \quad V = \sum_{k=0}^n v_k \quad (3)$$

Finally, the disc method shown in Equation (3) estimates the volume of the segment by treating each of its rows of size d_y as a cylinder. The segment is split into sizes of d_y for each the




Fig. 3: Volume estimates using surface area integration.

volume is calculated as $\pi r^2 d_y$. This method should be more robust than the integration step in cases when the orientation estimate error is large. Since each row is treated independently, a more complete surface not dependant on axial symmetry can be reconstructed, whereas with integration the entire contour is used with a singular estimate of the cross length.

$$c'_x = \frac{z_{max}}{f_x} (c_x - p_x) \quad c'_y = \frac{z_{max}}{f_y} (c_y - p_y) \quad (4)$$

The presented methods approximate the volume in pixels (px^2). To calculate the volume in centimetres (cm^3), we simply deproject the contour c by the camera intrinsic parameters focal length f_x, f_y , principal point p_x, p_y and an estimated distance z_{max} from the camera obtained from the max value bounded by the segment. For the disc method the z_{max} value is equal to the local max at each row rather than the entire segment. The deprojection step is shown in Equation (4) and is applied prior to volume estimation.

B. Mass Estimation

We model the relationship of mass and volume as the least squares regression fit to our data, and estimate the mass from predicted volume fit.

III. RESULTS AND CONCLUSIONS

We have presented a non-invasive/destructive, inexpensive method for volume and mass estimation in-field designed for use on a robotic platform. Our results for volumetric and mass estimation of the chosen soft-fruit are presented in Table I. It's evident that this method is appropriate for calculating the volume from only two dimensional data (segments) since the median absolute error is only 3.00cm^3 for the best method, which is only 1.00cm^3 above the maximum precision of the volumetric measurements. The relatively poor results for mass estimation were due to the low precision of the equipment.

	<i>Ellipsoid GT</i>	<i>Ellipsoid Depth</i>	<i>Integration</i>	<i>Disc</i>
Volume	3.28cm ³	3.94cm ³	3.00cm³	3.22cm ³
Mass	10.19g	11.90g	9.96g	9.85g

TABLE I: Median Absolute Error of volume and mass estimation methods, bold indicates the best method.

REFERENCES

- [1] R. Kirk, G. Cielniak, and M. Mangan, "L*a*b*fruits: A rapid and robust outdoor fruit detection system combining bio-inspired features with one-stage deep learning networks," *Sensors*, vol. 20, no. 1, 2020.
- [2] "Development and field evaluation of a strawberry harvesting robot with a cable-driven gripper," *Computers and Electronics in Agriculture*, vol. 157, pp. 392 – 402, 2019.
- [3] Y. Chen, W. S. Lee, H. Gan, N. Peres, C. Fraisse, Y. Zhang, and Y. He, "Strawberry yield prediction based on a deep neural network using high-resolution aerial orthoimages," *Remote Sensing*, vol. 11, 07 2019.

Human-in-the-Loop Adaptation and Reuse of Robot Assistance Policies for Ambient Assisted Living

Ronnie Smith

Edinburgh Centre for Robotics

Edinburgh, UK

ronnie.smith@ed.ac.uk

<https://doi.org/10.31256/Fn5Eb4A>

Abstract—Personalisation and adaptation of *Ambient Assisted Living* (AAL) solutions is the subject of many existing works, which seek to embed and/or learn user preferences and needs. However, with a focus on lab-based evaluation, it has been easy to overlook the potential realities of what AAL might look like in the future: a range of heterogeneous platforms, devices, and robots will generate swathes of knowledge that must be shared, within and outside the home. As such, it is important to consider scalability and interoperability in every aspect of future AAL solutions. *Adaptivity as a Service* (AaaS) is proposed as a highly specialised service with a core function of personalising and adapting smart homes and AAL systems to the needs and wants of individuals.

Index Terms—adaptivity, personalisation, Ambient Assisted Living (AAL), robotic care, Human-in-the-Loop (HITL), Digital Twin

I. INTRODUCTION

It is expected of future AAL solutions to be *adaptive*. It has been recognised for over a decade that this adaptivity must be deeply ingrained in AAL systems at the “algorithmic, architectural, and human interface” levels [1]. They must adapt to the changing habits, evolving needs, and individual preferences of users.

A recent stakeholder participatory study on requirement gathering for future assisted living solutions found that long-term personalisation and adaptation is vital for elderly users: the system must understand the impact of the ageing process [2]. Furthermore, solutions must be goal-oriented in that they work towards delivering a desired user state. This suggests a high degree of autonomous adaptation, which is in contrast to previous works which have suggested solutions be adapted over time through iterative user surveying to ensure the solution still meets the needs of the user [3].

It has been common in the past to think of personalisation as something that is not integral to the architecture of the system, but rather as an added feature. As such, there a number of approaches which propose the use of a *Graphical User Interface* (GUI) to provide personalisation insofar as some parameters of the system are modifiable by the user (e.g. [4] [5]). However, this puts a significant burden on the user in terms of effort, effectively causing the user to adapt to the system, rather than vice versa.

This work was supported by the Engineering and Physical Sciences Research Council (grant EP/L016834/1), EPSRC Centre for Doctoral Training in Robotics and Autonomous Systems.

Other solutions may instead provide a set of predefined templates of a user profile. For instance, in an AAL system these could be “dependent, assisted, at risk, and active” [6]. Users can be manually assigned, or automatically classified, into these categories so the system can modify behaviour accordingly. However, these approaches are limited: (1) to models provided during the design phase; and (2) in that it is virtually impossible to cater for elderly individuals with a wide and diverse range of needs.

It is increasingly common to see approaches which adapt in real-time, such as the social robot ‘GrowMu’, which offers services to users based on their emotion and where the robot is in the environment [7]. The robot is trained (offline) to gravitate towards services which elicit a positive response. While useful, the approach is limited to service selection only, with little regard to adaption within those services.

In contrast, Adaptivity as a Service (AaaS) is inspired by established ‘hybrid’ approaches in Human Activity Recognition (e.g. [8]): hybrid models fuse knowledge- and data-driven sources to enable adaption to individual users and improve baseline performance and scalability. This typically means starting with initial ‘seeds’ (templates) of what a system can recognise or do (from provided knowledge), while subsequent data collection at run time enables a semi-supervised learning process to grow capability over time.

Where existing approaches focus primarily on providing one or a few dimensions of personalisation and adaptivity within AAL solutions, AaaS puts it at the core. AaaS is proposed as a highly specialised service that combines extensive user modelling, distributed learning, and knowledge transfer [9].

II. PROPOSAL

AaaS addresses three types of long-term adaptation [9]:

- 1) Adapting context-awareness itself to account for predicted physical and mental decline, based on individual’s known conditions and principles of ageing.
- 2) Adapting assistive functionality, including interaction modalities, to fit an individual’s exact needs/wants.
- 3) Adapting quickly to new users, based on experience.

AaaS is ‘distributed’ in the sense that services are delivered in-home, through ‘local’ nodes, while data is aggregated and processed centrally, at a ‘global’ level. A (simplified) conceptual architectural overview of AaaS is provided in Figure 1. As shown, AaaS sits between existing AAL components, acting

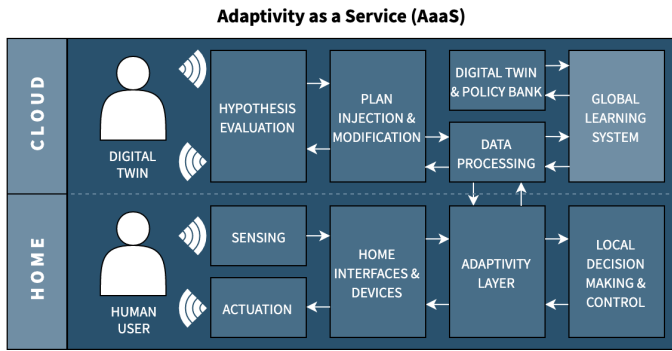


Fig. 1. Conceptual architecture overview of AaaS.

as an intermediary, similar to the concept of an ‘adaptivity layer’ proposed in [10]). The user is a single member in a population of individuals requiring personalised and adapted AAL. This population-centric view enables knowledge transfer and learning from common experience at the local level.

AaaS addresses the three types of adaption with a two-pronged approach to Human-in-the-Loop operation: (1) implicit consideration of the user in decision making, and (2) continuous learning of desired system behaviour relative to user profile for the updating of beliefs.

III. A HYBRID MODEL

It is expected that AaaS employ hybrid models to: (1) model users, and (2) enable personalisation. A *Digital Twin* (DT) is therefore proposed to provide a digital counterpart for each user, which is then evolved through a *Human-in-the-Loop* (HITL) process that enables learning while implicitly expanding the knowledge represented by the DT.

a) *Digital Twin*: The digital twin is a user model initially comprised only of explicitly provided information (e.g. from the user or their carers). Based on this information, a set of initial *Assistive Policies* (APs) are selected and assigned to the DT. Selection is based on learning from experience with similar users. Policy reuse is fundamental in improving the initial user experience and acceptance, providing a higher degree of familiarity in the first instance.

Each AP is a general plan of an assistive functionality, originally created from expert knowledge. All of these APs together form a bank of policies covering a range of scenarios. APs describe: (1) a plan of actions that can be carried out in a robot-enabled home environment; and (2) the variability in these plans. APs are not hardware-specific, and instead rely on translation from the high-level to the command-level.

Through the policy personalisation process explained below, APs evolve over time as they are fine-tuned to a given user. In terms of forecasting, this has the following implications: (1) short-term, AaaS can evaluate whether proposed actions by a local planner (e.g. from the smart home / AAL system) are sensible for the given user through hypothesis testing; and (2) long-term, it allows forecasting of user behaviour and health, which allows for monitoring of specific conditions with/without the condition being specified in the DT.

b) *Policy Personalisation*: Realising a hybrid approach, APs originally provided from knowledge engineering should be evolved and branched out into a set of source APs from a *Reinforcement Learning* (RL) and policy reuse process. Knowledge can be extracted about the impact of certain traits (e.g. the presence of a specific *Mild Cognitive Impairment* [MCI]) based on the aggregation of real-world experiences with relevant users.

APs are evolved to suit the *wants*, and *needs*, of each user. The optimal state of users’ APs are used to enhance their DT, and for policy reuse more broadly. This enables, for example, the automated selection of an AP set for new users, based on experience from similar users within the population of DTs.

IV. CONCLUSION & FUTURE WORK

Here, the fundamental principles of AaaS and its novel distributed HITL approach have been outlined in relation to existing challenges in AAL personalisation. Within AaaS, individual homes can benefit from and contribute to a wider network of adaptivity specialisation. Future research will need to focus on the best approaches to meet key goals of AaaS. Ultimately, the Digital Twin will serve as a rich source of data that accompanies a user for life, which systems that deal with personalisation independently may fail to replicate.

While the vision for AaaS is rather broad, research will focus on feasibility by addressing underpinning scientific issues, enabled by our Robotic Assisted Living Testbed (RALT)¹: a 60m², fully-furnished simulated apartment comprising a bedroom, bathroom, and combined kitchen/dining/living area.

Work will therefore focus on: (1) encoding/representation of personalisation and adaptivity policies/plans; (2) translating these for heterogeneous AAL platforms and devices; (3) merging and processing feedback to best reflect learning from similar users; and (4) management of “unhealthy” feedback, where users have moulded policy to their unhealthy wants.

The next phase will focus on adapting robot-enabled assistance for some example scenarios. The target for adaptation in the experiments is users with varying level of capability in performing cognitive and physical tasks during a scenario, and so the latent variable is user skill. The agent is the smart home, as an ensemble of *Internet of Things* and robotic devices. A scenario may represent, for example, a use case in which an immobile user wishes to have a cup of coffee and the adaptive service must learn how to meet this goal using the devices at hand, in a way that suits the user. As such, the AP for each scenario can be formulated as an RL problem, with the state and actions of the system modelled in an MDP. The transition function depends on user actions (e.g. requests, feedback, active sensing), while rewards are allocated based on successfully achieving the goal. The end product is that users can effectively build their own solutions through interactions with the system.

¹Virtual tour available at <https://ralt.hw.ac.uk/>

REFERENCES

- [1] J. Nehmer, M. Becker, A. Karshmer, and R. Lamm, "Living assistance systems: an ambient intelligence approach," in *Proceeding of the 28th international conference on Software engineering - ICSE '06*. Shanghai, China: ACM Press, 2006, p. 43. [Online]. Available: <http://portal.acm.org/citation.cfm?doid=1134285.1134293>
- [2] J. Cahill, S. McLoughlin, M. O'Connor, M. Stolberg, and S. Wetherall, "Addressing Issues of Need, Adaptability, User Acceptability and Ethics in the Participatory Design of New Technology Enabling Wellness, Independence and Dignity for Seniors Living in Residential Homes," in *Human Aspects of IT for the Aged Population. Aging, Design and User Experience*, J. Zhou and G. Salvendy, Eds. Cham: Springer International Publishing, 2017, vol. 10297, pp. 90–109.
- [3] S. Kieffer, J.-Y. L. Lawson, and B. Macq, "User-Centered Design and Fast Prototyping of an Ambient Assisted Living System for Elderly People," in *2009 Sixth International Conference on Information Technology: New Generations*. Las Vegas, NV, USA: IEEE, 2009, pp. 1220–1225.
- [4] C. Antonopoulos et al., "Robots in assisted living environments as an unobtrusive, efficient, reliable and modular solution for independent ageing: The RADIO perspective," in *International Symposium on Applied Reconfigurable Computing*. Springer, 2015, pp. 519–530.
- [5] J. Saunders, D. S. Syrdal, K. L. Koay, N. Burke, and K. Dautenhahn, "'Teach Me—Show Me'—End-User Personalization of a Smart Home and Companion Robot," *IEEE Transactions on Human-Machine Systems*, vol. 46, no. 1, pp. 27–40, Feb. 2016.
- [6] S. Stavrotheodoros, N. Kaklanis, and D. Tzovaras, "A personalized cloud-based platform for AAL support to cognitively impaired elderly people," *IFMBE Proceedings*, vol. 66, no. 2, pp. 87–91, 2018.
- [7] G. S. Martins, P. Ferreira, L. Santos, and J. Dias, "A Context-Aware Adaptability Model for Service Robots," p. 7.
- [8] L. Chen, C. Nugent, and G. Okeyo, "An ontology-based hybrid approach to activity modeling for smart homes," *IEEE Transactions on Human-Machine Systems*, vol. 44, no. 1, pp. 92–105, 2013.
- [9] R. Smith, "Adaptivity as a Service (AaaS): Enabling Deep Personalisation for a Heterogeneous Ambient Assisted Living Landscape," in *To appear in 1st ACM/IEEE Workshop on Behavioral Patterns and Interaction Modelling for Personalized Human-Robot Interaction*, Mar. 2020.
- [10] G. S. Martins, L. Santos, and J. Dias, "User-Adaptive Interaction in Social Robots: A Survey Focusing on Non-physical Interaction," *International Journal of Social Robotics*, vol. 11, no. 1, pp. 185–205, Jan. 2019. [Online]. Available: <http://link.springer.com/10.1007/s12369-018-0485-4>

Automated Topological Mapping for Agricultural Robots

Karoline Heiwolt, Willow Mandil, Grzegorz Cielniak and Marc Hanheide

Lincoln Centre for Autonomous Systems,

University of Lincoln, UK

18727004@students.lincoln.ac.uk, {wimandil, gcielniak, mhanheide}@lincoln.ac.uk

<https://doi.org/10.31256/Ze8Ex1V>

Abstract—Essential to agricultural robot deployment in farms are accurate topological maps, which are manually created in current systems. In this work we present a novel approach to automatically generate a topological map along crop rows from aerial images for the deployment of agricultural mobile robots. We evaluate our system in a digital twin of a farm environment using real-world textures and physical simulation, and also demonstrate its applicability to aerial images of a real farm.

Index Terms—Agri-robotics, Topological mapping, Mobile Robot Navigation

I. INTRODUCTION

The deployment of fully autonomous mobile platforms to real-world farms is fast approaching, aiming to solve challenges from a growing population, labour shortage, and pressure to reduce environmental impact [1]. The deployment of fully autonomous mobile platforms to real-world farms requires solving a range of technical challenges. First and foremost safe and precise navigation across the farm environment. In this work, we present a novel approach to automatic topological map creation from aerial views of a field to guide the mobile robots along crop rows.

Thanks to recent advances in mobile robotics, manipulation and computer vision, modern agricultural robots can be deployed in various agricultural environments and are able to complete tasks such as crop scouting, pest and weed control, or harvesting [2]. Automated mapping of farm and field environments is an essential stage towards their commercialisation [1]. Currently, predefined topological maps are used to navigate up and down crop rows to deliver crop treatments [3]. Farms are constantly changing, crop rows and farm structure will vary over time requiring new topological maps which are typically created manually [4]. Our proposed solution addresses that problem by automating the topological map creation. Although automated waypoint creation from a map is a well-studied problem, it has not been applied in agricultural applications. Current crop row segmentation algorithms often rely on assumptions of straight, parallel, equally spaced crop rows, using e.g. Hough transforms [5]. These perform well on straight rows but fail when the crop rows are curved, which often occurs in fields with trees, pylons or ditches, or when the crop rows change direction in irregularly shaped fields [6].

Karoline Heiwolt and Willow Mandil are co-first authors.

Current solutions also commonly use sensors positioned close to the ground and thus focus on local guidance [6], [7].

The key contributions of this paper are a novel method of crop row detection and topological mapping from aerial images, which takes into account common crop row variations found in fields and evaluation of the method on both real world fields and simulated digital twins.

II. METHODOLOGY

The proposed approach uses aerial images of a farm, such as hi-resolution UAV or satellite images. In this work we use real images of a Lincolnshire farm captured by UAV, as well as a digital twin of a farm environment simulated in Gazebo [8]. The ground plane of this environment is covered with textures, consisting of real images of soil and rows of different types of crop, collected from a camera mounted on a mobile platform deployed at the real farm facilities of the University of Lincoln.

To avoid harm to the crop, mobile robots should travel across the field along the centre line of crop rows only. For our approach to automate deployment in new fields, a topological map is created by converting the captured aerial image into a set of waypoints, connected with traversable edges.

First, we find the locations of crops in the image by colour-based segmentation. We then determine the principal angle of parallel crop rows visible in the segmented binary image. We construct a set of oriented graphs (0 to 180°) resulting from the sum of intensities across interval lines perpendicular to the orientation (Fig. 1). The principal angle α is determined as perpendicular to the graph with the highest mean peak. A dense set of waypoints is then placed on the centre of crop clusters along the intensity profile lines perpendicular to the principal crop row direction a (Fig. 2, left). The waypoints are clustered into individual rows and ordered by their distance along each row, to produce a continuous safe route for travelling along each separate crop row. Additionally, a safe turning point is appended to the start and end of each crop row, parallel to α (Fig. 2, centre).

Next, we remove redundant waypoints from the dense set by omitting waypoints for which the deviation from the previous direction of travel is only within some permitted perpendicular distance, l . The result is a sparser set containing only waypoints in locations where the direction of travel changes by more than l (Fig. 2, right). This down-sampling procedure introduces a sparsity-accuracy trade-off and majorly influences




Fig. 1: Perpendicular lines drawn on the binary image at angle 0° (left) and 90° (right), along with their oriented graphs of intensity sums.




Fig. 2: Left: Placement of waypoints (red) on peak clusters, Centre: dense waypoints sorted into individual rows, Right: a sparse variant of the same topological map.

the performance of the finished topological map. This trade-off is evaluated in Sec. III.

III. EVALUATION

A. Experimental Setup

The quality of the topological map was assessed by its coverage, measured as the proportion of the area in which crops grow (as manually annotated), that has successfully been surveyed by a simulated Thorvald robot [9] after it visited every way point on the topological map once. We report results for three test scenarios consisting of rectangular fields with different row crops (basil, lettuce, and onions), and one set of non-uniform rows with gentle and severe bends simulating situations where there are environmental obstacles present in the field. Additionally, we also applied the method to an aerial image taken by UAV of a real Lincolnshire farm¹ (see Fig. 4 and Fig. 5) growing winter wheat for validation of our approach.

¹courtesy of Jonathan Trotter and SAGA Robotics

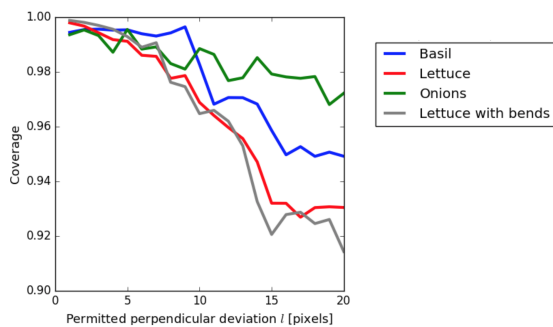


Fig. 3: Coverage in four scenarios dependent on the maximum permitted perpendicular deviation l from previous travel direction.

B. Results

The proposed algorithm deals well with variability in crop placement and curvature within crop rows. For straight crop rows (see Onions in Fig. 3), even very sparse maps achieve near optimal coverage. The approach also translates well to real-world images (Fig. 4). However our algorithm's limitations become apparent when the algorithm is applied to a larger, more irregularly shaped field, in which the general direction of crop rows changes significantly (Fig. 5). The principle crop row angle found across the entire image is only suitable for part of the image. The algorithm fails to pick up on the crop rows on the left side.

IV. CONCLUSIONS AND FUTURE WORK

We presented a novel topological mapping algorithm, which is robust to curvature within single crop rows. We also demonstrated its applicability to real-world examples. However, this algorithm is presented as a baseline for future development. To map large fields (Fig. 5), we propose to repeatedly apply the algorithm in a hierarchical quadtree procedure, repeatedly partitioning the image and dynamically increasing the resolution in uncertain areas, thus evaluating the principal angles accurately for subsections of the field. In future work this system should be extended to create complete semantic maps of entire farm environments, enabling efficient automated fleet deployment for the next generation of agricultural robots.




Fig. 4: Dense (left) and sparse (right) topological map generated from an aerial image of a real farm of wheat crops.




Fig. 5: The algorithm applied to a real world scenario with changes in crop row direction (Composite image).

REFERENCES

- [1] T. Duckett, S. Pearson, S. Blackmore, B. Grieve, W.-H. Chen, G. Cielniak, J. Cleaversmith, J. Dai, S. Davis, C. Fox *et al.*, "Agricultural robotics: The future of robotic agriculture," UK-RAS Network White Papers, Tech. Rep., 2018.
- [2] D. Longo and G. Muscato, "Design and simulation of two robotic systems for automatic artichoke harvesting," *Robotics*, vol. 2, pp. 217–230, 12 2013.
- [3] T. D. Le, V. R. Ponnambalam, J. G. O. Gjevestad, and P. J. From, "A low-cost and efficient autonomous row-following robot for food production in polytunnels," *Journal of Field Robotics*, vol. 37, no. 2, pp. 309–321, 2020.
- [4] A. Bechar and C. Vigneault, "Agricultural robots for field operations. part 2: Operations and systems," *Biosystems Engineering*, vol. 153, pp. 110 – 128, 2017.
- [5] R. Ji and L. Qi, "Crop-row detection algorithm based on random hough transformation," *Mathematical and Computer Modelling*, vol. 54, no. 3–4, pp. 1016–1020, 2011.
- [6] I. García-Santillán, J. M. Guerrero, M. Montalvo, and G. Pajares, "Curved and straight crop row detection by accumulation of green pixels from images in maize fields," *Precision Agriculture*, vol. 19, no. 1, pp. 18–41, 2018.
- [7] A. Ahmadi, L. Nardi, N. Chebrolu, and C. Stachniss, "Visual servoing-based navigation for monitoring row-crop fields," 2019.
- [8] N. Koenig and A. Howard, "Design and use paradigms for gazebo, an open-source multi-robot simulator," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, Sendai, Japan, Sep 2004, pp. 2149–2154.
- [9] L. Grimstad and P. J. From, "The Thorvald II Agricultural Robotic System," *Robotics*, vol. 6, no. 4, 2017.

Trajectory Tracking and Control of Multiple Robot Arms on a Free-Floating Spacecraft for Debris Removal

1st Ashith Shyam
Surrey Space Center
University of Surrey
Guildford, United Kingdom
0000-0001-8990-672X

2nd Arunkumar Rathinam
Surrey Space Center
University of Surrey
Guildford, United Kingdom
a.rathinam@surrey.ac.uk

3rd Zhou Hao
Surrey Space Center
University of Surrey
Guildford, United Kingdom
z.hao@surrey.ac.uk

<https://doi.org/10.31256/Jp9Yd3M>

Abstract—In this work, we propose the idea of using dedicated small spacecrafts carrying multiple robot arms for space debris removal. Future space missions are expected to have computers with high computational capabilities. Hence, a model predictive control architecture is used to generate the control commands. The small spacecraft in consideration carries two identical 3-DoF robot arms attached diagonally opposite to each other. Thus the workspace covered by the robot arm increases. We show trajectory tracking simulation results of the proposed free-floating spacecraft to demonstrate the concept.

Index Terms—Trajectory tracking, Model predictive control, Spacecraft manipulator, Debris removal

I. INTRODUCTION

NASA estimates that there are around half a million pieces of space junk currently floating around in the Lower Earth Orbit (LEO) [7]. These debris not only possess a huge threat to the existing functional satellites but also for any future missions. Quantitatively, on an average calculated during 2004–12, 72 objects were placed into LEO per year and this rose to 125 post 2012. It is estimated that the space environment can be stabilised when on the order of 5–10 objects are removed from LEO per year.

Space environment poses some unique challenges such as latency in communication, extreme safety requirements and reduced processing powers due to which the level of autonomy in decision making is far less than its counterparts on earth. To enable fully autonomous decision making, future space missions are expected to operate under autonomy level E4 according to the European Co-operation for Space Standardization (ECSS) [5]. This would enable future robots to have more computational resources for carrying out complex optimizations which until now are very limited and also to perform various activities (like space exploration, servicing and repair.) without any or minimal human intervention.

The main contribution of this work is the idea of using spacecrafts dedicated for debris removal equipped with multiple robot arms. Also, perform trajectory tracking and control

This work has been carried under the consortium FAIR-SPACE and obtained funding from UKRI, UK Space Agency and Industrial Strategy.

of two 3-DoF arms attached to the spacecraft using a model predictive control algorithm. These spacecrafts are assumed free floating (with no active attitude disturbance control system) and hence any motion of the robot arm would induce a movement on the spacecraft. We present the simulations of trajectory tracking and control of two robotic arms attached to the spacecraft using a model predictive control algorithm. Robotic arms on spacecraft are preferred over other debris removal methods because it can be easily extended to other orbital applications like on-orbit servicing and assembly, and autonomous rendezvous and docking.

II. FORMULATION

A. Kinematic and Dynamics of free-floating spacecrafts

Prior works [6, 9, 10] on the kinematic and dynamic equations of free-floating spacecrafts are available in the literature. We provide the main equations here for completeness.

$$\begin{aligned} v_{eff} &= (J_m - J_s I_s^{-1} I_m) \dot{\phi}_m \\ &= J^* \dot{\phi}_m \end{aligned} \quad (1)$$

where v_{eff} , J_m , J_s , I_s , I_m and $\dot{\phi}_m$ are respectively the end-effector velocity, manipulator Jacobian, spacecraft Jacobian, 3×3 satellite inertia matrix, $3 \times n$ DoF manipulator inertia matrix and the manipulator joint velocities. Here nDoF corresponds to the number of actuated joints. This formulation assumes that the momentum is conserved and is zero at the beginning. The dynamic formulation can be expressed as

$$M(\phi) \ddot{\phi} + C(\phi, \dot{\phi}) = \tau \quad (2)$$

where τ is the control torque to be applied at the manipulator joints, $M(\phi)$ and $C(\phi, \dot{\phi})$ are respectively the mass and Coriolis and centripetal matrices.

B. Feedback Linearization

Substituting the joint torques, τ with a fictitious input, $\alpha\tau' + \beta$ and substituting $\alpha = M(\phi)$ and $\beta = C(\phi, \dot{\phi})$ [2] without loss of generality, we could linearize the eq. (2) and can write each individual joint variables in state space form as

$$\dot{x} = \begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} + \begin{bmatrix} 0 \\ 1 \end{bmatrix} \tau' \quad (3)$$

$$\dot{x} = A'x + B'\tau'$$

where x_1 and x_2 corresponds to the angular displacement and velocity of the actuated joints.

C. Model Predictive Control

For our MPC [1] formulation, we use a prediction horizon of 5 seconds. The formulation can be shown as in eq. (4)

$$J = x^T P x + \int_{t_0}^{T-1} x^T Q x + u^T R u dt \quad (4)$$

subject to the constraints

$$\dot{x} = Ax + B\tau$$

$$x(0) = x_0$$

where $x = \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$. Here $Q (\geq 0)$ and $R (> 0)$ are respectively the time varying state and control cost matrices (usually diagonal matrices), P is the stabilizing matrix obtained by the solution of algebraic Riccati equation at every time instant, t_0, T are the initial and final time respectively, A and B are obtained by diagonally stacking up A' and B' equal to the number of joints and u is the control input. Interested readers may refer to [4] for a detailed explanation. Since we are dealing with a free-floating spacecraft, it is important to note that only the robotic arms are actuated. The motion of the spacecraft is due to the reaction forces induced on it from the arms. Thus the P matrix only acts as a stabilizer to the arm motion

III. SIMULATIONS

We show simulations done in Mayavi [8] of two identical robot arms which are attached diametrically opposite to each other on the spacecraft as shown in Fig 1. The Denavit-Hartenberg [3] values and the dynamic parameters of the robot arm used in the simulations are as shown in table I and table II.

TABLE I
DH PARAMETERS OF THE ROBOT ARM

Joint	α (rad)	a (m)	d (m)	θ (rad)
1	$-\pi/2$	0.0	0.5	θ_1
2	$\pi/2$	0.0	0.0	θ_2
3	0	1.0	0.0	θ_3
flange	0	1.0	0.0	θ_3

TABLE II
DYNAMIC PARAMETERS

	Satellite Link 0	Link 1	Link 2	Link 3
Mass (kg)	200.0	20.0	50.0	50.0
l (m)	2.10	0.25	2.5	2.5
I_x	1400.0	0.1	0.25	0.25
I_y	1400.0	0.10	26.0	26.0
I_z	2040.0	0.10	26.0	26.0

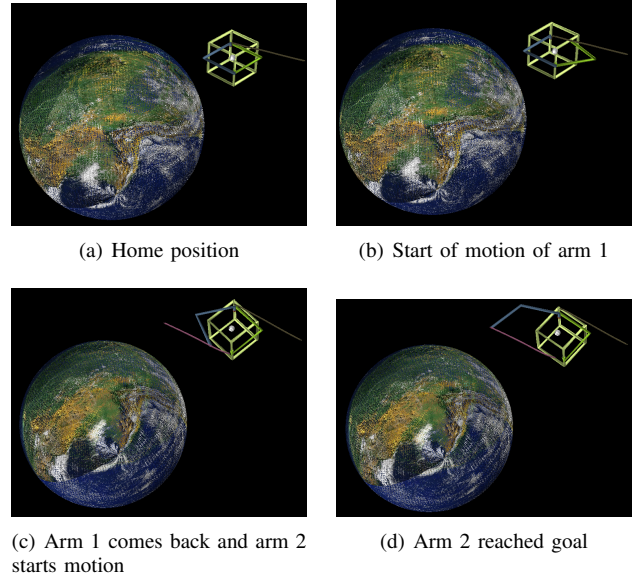


Fig. 1. Two arm trajectory tracking to reach debris

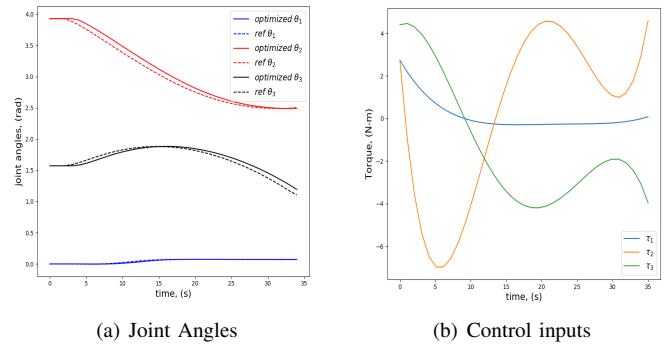


Fig. 2. Tracking performance of the controller for arm 1

IV. RESULTS

Fig 2 gives the tracking performance of the model predictive controller for arm 1. The average error in tracking the reference trajectory for joint 1, joint 2 and joint 3 are respectively 0.03%, 0.18% and 0.2% respectively. The simulation video is available at <https://youtu.be/xEQBiDpLLg>

V. CONCLUSIONS AND FUTURE WORK

From simulations, we have demonstrated the possibility of having multiple robot arms on the same spacecraft. Further, it is shown that a model predictive tracking controller could be developed to track a trajectory. The slight error in tracking could be attributed to the hand tuning of control gain matrices, Q and R .

For future work, it would be interesting to see how more than two arms could be used and simultaneous actuations of the arms with minimal attitudinal changes to the spacecraft.

REFERENCES

- [1] Eduardo F Camacho and Carlos Bordons Alba. *Model predictive control*. Springer Science & Business Media, 2013.
- [2] Ashitava Ghosal. *Robotics: fundamental concepts and analysis*. Oxford university press, 2006.
- [3] Richard S Hartenberg and Jacques Denavit. A kinematic notation for lower pair mechanisms based on matrices. *Journal of applied mechanics*, 77(2):215–221, 1955.
- [4] Donald E Kirk. *Optimal control theory: an introduction*. Courier Corporation, 2004.
- [5] Fabrício Kucinskis and Maurício Ferreira. Taking the ECSS autonomy concepts one step further. In *SpaceOps 2010 Conference Delivering on the Dream Hosted by NASA Marshall Space Flight Center and Organized by AIAA*, page 2364, 2010.
- [6] Kostas Nanos and Evangelos G Papadopoulos. On the dynamics and control of free-floating space manipulator systems in the presence of angular momentum. *Frontiers in Robotics and AI*, 4:26, 2017.
- [7] NASA. Space debris and human spacecraft. https://www.nasa.gov/mission_pages/station/news/orbital_debris.html, 2013.
- [8] Prabhu Ramachandran and Gaël Varoquaux. Mayavi: 3d visualization of scientific data. *Computing in Science & Engineering*, 13(2):40–51, 2011.
- [9] Yoji Umetani, Kazuya Yoshida, et al. Resolved motion rate control of space manipulators with generalized jacobian matrix. *IEEE Transactions on robotics and automation*, 5(3):303–314, 1989.
- [10] Markus Wilde, Stephen Kwok Choon, Alessio Grompone, and Marcello Romano. Equations of motion of free-floating spacecraft-manipulator systems: An engineer’s tutorial. *Frontiers in Robotics and AI*, 5: 41, 2018.

Improving Quadrupedal Locomotion on Granular Material Using Genetic Algorithm

Jakub Hulas
School of Mechanical Engineering
University of Leeds
 Leeds, United Kingdom
 jakubhulas@gmail.com

Chengxu Zhou
School of Mechanical Engineering
University of Leeds
 Leeds, United Kingdom
 C.X.Zhou@leeds.ac.uk

<https://doi.org/10.31256/Qs1Ub7E>

Abstract—This paper describes process of generating a stable gait for a quadruped robot on granular material such as sand. To achieve this goal a simulated environment was created, to model the kinematics and dynamics of the robot. Genetic algorithm was used as a gait generation technique. The simulation model of Laikago was built in MATLAB and Simscape. The contact model between the foot and ground, was based on existing results from sand penetration testing. The genetic algorithms optimised the joint space trajectory, to evaluate of the gait performance on granular terrain. The proposed method required to first create a simulation constrained to planar motion to simplify the problem. The 2D gait was then used as the initial guess for the optimisation in 3D simulation and allowed for creation of an acceptable gait. The final gait has demonstrated the effectiveness of the proposed approach for quadruped robot walking on granular terrain.

Index Terms—quadruped robot, genetic algorithm, legged robot, locomotion

I. INTRODUCTION

The legged robot technology has experienced a large development in the last few decades mostly caused by the development of smaller high-torque motors. Legged robots have lower efficiency than the wheeled ones, when moving on flat ground. Their advantage is, however the flexibility when moving across complex environments [1]. Adaptivity of the robot to the environment is crucial for its practical utility. The hardware cannot be tested on all the possible grounds, that is why there is a need for a quick generation method which yields an appropriate gait. One of possible methods is central pattern generation (CPG), Which uses the data from animal walking patters and generates appropriate gaits using deep learning [2]. This approach requires gathering and processing a large normalized data set. The approach chosen in this paper was reinforcement learning with GA. It uses a simulated environment to generate gaits and evaluate and them using a cost function [3].

Laikago is a new affordable quadruped robot, which is lacking the ability of locomotion through complex environments. It can walk up the hills or on the grass but the ability of walking on granular material such as sand and mud has not been demonstrated, even though it is a feature included in the more expensive quadruped robots such as ANYmal.¹

¹Unitree.cc, [online], Available at: <http://www.unitree.cc/>, [Accessed 2 Apr. 2019]

II. SIMULATION

A. Quadruped Robot Model

The Laikago model is imported into Simscape. It represents relationships between the individual links and joints. For the initial gait generation, the chassis is connected to the world using a planar joint. This allows the robot to move upwards and forwards without the risk of falling to the side. This reduces the degrees of freedom (DoF) from 6 to 3, reducing the complexity of the problem. In the succeeding simulations the planar joint is replaced with a unconstrained joint block to freeing all 6 DoF. This procedure allowed convergence in the first generation. The results of that could be later used as an initial guess for the gait generation of the unconstrained robot, significantly improving the resulting gait.

B. Sand Contact Modeling

A non-linear contact model is chosen for the simulation. As shown in the eq. 2 the normal force is a function of penetration depth and velocity. It is dependent on the stiffness and damping constants which are based on empirical results [4].

$$F_n = kz^n + c\dot{z} \quad (1)$$

An additional feature of the model, is a transition depth, at which the slope of the force-penetration curve is changing its gradient. The transition depth is dependent on the sand properties such as the volume fraction. It occurs because of the way the normal and shear forces are distributed in sand.

To verify the model, before applying it to the whole robot, a simulation consisting of one leg fixed at the base, is created. The leg moves vertically downwards with a constant end-effector velocity and penetrates the sand. When the ground contact is detected and the appropriate reaction forces are applied to the foot. The results of the simulation are shown on the force penetration graph (fig. 2). The loading case clearly shows a non-linear nature of the contact. For this simulation the assumed volume fraction of the sand is 0.6. Based on this value the transition depth is 0.3 m and the stress/depth value is equal to $10^6 Nm^{-3}$ For one foot, the stiffness constant k , is equal to $3927N/m$ [4]. The damping coefficient is set to $250Ns/m$. The exponential constant is set to 1.25 [5].




Fig. 1. Sand model force-penetration curve.

III. GAIT GENERATION

GA is an optimisation method based on natural selection. The algorithm creates a generation of solutions, which consists of the best solutions from the previous generation, a crossover between them, and random mutations. This allows to check a wider range of possibilities compared to other optimisation methods, such as linear regression. It reduces the problem of misjudging local minima for a global minimum.

The simulated robot is actuated by 6 joint space positions during one repeating gait period. The Laikago robot has 3 motors in each leg, which means there are 72 variables to optimise. To simplify the problem, the front and rear pair of legs are actuated using the same trajectory with half a period delay between the right-left pairs. This allowed to reduce the optimisation to 18 variables, significantly increasing the convergence.

A. Cost Function

Cost function (CF) is used to evaluate the performance of the generated gait. The value of the cost function is fed back to the GA. The algorithm uses the value as a variable to optimise. The main role of the CF is to constrain and push the robot gait towards a solution which a human would judge as successful. That's why the optimisation had to be initially repeated to find proper penalties and adjust their weights accordingly. The positive penalties represent what is the goal of the robot gait, in this case the distance and the time walked without falling. The negative ones are all the constraints put on the robot. They include: excessive motion of the robot body in the undesired directions, robot walking off a straight path, aggressive joint moves and a dynamic walk (when less than 2 feet are on the ground).

$$cost = -\Pi(+ve)/\Pi(-ve)(2) \quad (2)$$

Eq. 2 indicates the CF is always negative, to ensure finding the minimum.




Fig. 2. Snapshots of the final gait movement on sand.

B. Final Gait

The gait was first generated in 2D because the additional constraints applied on the body simplified the overall problem. This allows the optimisation to converge faster to stable results. The 2D trajectory was used as an initial guess for the gait generation in 3D simulation. That yields much higher convergence rate than a pure 3D generation. The velocity of the final gait was $0.4m/s$ which is 30% of the max speed. During walking, the robot is slightly drifting to the side from the straight path. The final gait can be further improved with adjusting the weights of the CF.

IV. CONCLUSION

The presented GA is shown to generate a stable walking gait for the Laikago robot on granular material. Gait simulation can be used as a first step for generation of new gaits. It could be applied to different robots, or environmental scenarios. In the next steps a reinforced learning algorithm could optimise the gait in an experimental setup, using the simulation gait as a starting point. The contact model should be validated by performing experiments using the hardware.

The procedure of choosing the right reward system for the gait evaluation is time consuming, because it is based on trial and error approach. It might be hard to judge if the given penalty system is right or the success is caused by random mutations. It is difficult to avoid reward hacking of the model and predict the exact outcome. This could be solved with reward modelling. This approach uses a feedback from a human, who chooses a preferred solution. The model then adjust the cost function to satisfy the human choices [6].

REFERENCES

- [1] C. Zhou, X. Wang, Z. Li, D. Caldwell and N. Tsagarakis, "Exploiting the redundancy for humanoid robots to dynamically step over a large obstacle," *IEEE/RSJ IROS*, pp. 1500–1604, 2015.
- [2] D. Bucher, G. Haspel, J. Golowasch and F. Nadim, "Central Pattern Generators," In *eLS*, John Wiley & Sons, Ltd (Ed.), 2015.
- [3] J. A. Carr, "An Introduction to Genetic Algorithms," 2014.
- [4] A. Peter, "Establishment of Bekker's Model for Predicting the Pressure-Sinkage Behaviour of a Loamy Sand Soil," 2014.
- [5] J. Aguilar and D. Goldman, "Robophysical study of jumping dynamics on granular media," *Nature Physics*, 12(3), 2015, pp. 278–283.
- [6] P. Christiano, J. Leike, T. Brown, M. Martic, S. Legg and D. Amodio, "Deep reinforcement learning from human preferences," 2017.

Analysis of two-wheeled robot morphology for a slope environment

Robert Woolley, Jon Timmis, Andy M. Tyrrell
University of York, UK
rw1445@york.ac.uk

<https://doi.org/10.31256/Sf6Zi7L>

Abstract—This paper considers a mathematical and simulated analysis of the Cylindabot, a robot design with an adaptive morphology with minimal actuation. The paper compares simulation against a mathematical model to both understand the dynamics of such a robot design and also provide a mathematical foundation for future designs.

I. INTRODUCTION

Two-wheeled robots are a common form of robot design due to their minimal number of actuators. A more specialised type is discussed in this paper, having large wheels that encompass the main body of the robot. The consequences of this are that the robot will be able to survive falls, handle bumps that are large compared its size and never be the wrong way up due to symmetry. Several similar robots have been previously [1], [2], [3] and [4]. In [4] the robot was able to create a mathematical model of step clearance.

In this work a set of mathematical inequalities are derived and values are then substituted into these equations to give theoretical results. Finally they are compared with results from V-rep simulations [5]. These mathematical models help confirm results from simulation, give a theoretical derivation and can be quickly adapted for different robots.

II. INTERNALLY WEIGHTED WHEELS

For continuous climbing of a slope, the centre of mass has to be further forward than the point of contact (Fig 1). Once the robot is rolling up the slope, the motors will keep the mass at this point so it will continue to climb.




Fig. 1. Wheel on slope to show balance point

$$\frac{r}{R} > \sin\theta \quad (1)$$

In eqn. (1) r = radius of center of mass and R = radius of wheel. If this condition holds, the robot will be able to climb that slope.

III. TAILED TWO-WHEELED ROBOT

The addition of a tail allows the robot to apply more torque through its wheels but creates an extra drag behind the robot. This extra point of contact makes the robot statically balanced and the position the centre of mass less important. The mass of the robot M_t is split between M (main body) and m (tail).

For the robot moving on the flat, assuming the robot needs only to maintain a steady speed, then the friction on the wheel simply needs to be bigger than the drag of the tail.

$$\mu_1 M > \mu_2 m \quad (2)$$




Fig. 2. Diagram of forces applied to tailed robot. D = drive force of wheels, G = ground reaction, F = friction

The moments around the two points of contact on the slope are used to calculate the dynamic requirement. Taking moments around G_1 (Fig 2) and then resolving forces the follow inequality can be derived. Here γ to represent $\tan^{-1}(0.5)$.

$$0 < \mu_1(M + m)g \cos\theta - (\mu_1 + \mu_2)G_2 - (M + m)g \sin\theta \quad (3)$$

$$\text{where: } G_2 = \frac{Mg \sin\theta + \sqrt{5}mg \cos(\theta - \gamma)}{2} \quad (4)$$

The final inequality (eqn. 3) assumes that it is desirable for the wheel not to slip. The drive force is set to zero to find the maximum angle that the wheel could grip while maintaining speed.

IV. MATHEMATICAL PREDICTIONS

The inequalities above were derived from sound theory but to obtain appropriate results, realistic parameters need to be substituted into them (listed in Table I). The most important being radius of the centre of mass (r) and frictional coefficient (μ). They are given two possible values each to inform possible designs.

TABLE I
PARAMETER TO SUBSTITUTE INTO INEQUALITIES

Term	Value
R	150mm
r	$\frac{R}{2}$ or $\frac{R}{4}$
$\mu_1\mu_2$	0.5 or 1
m	$0.1M_t$

Firstly, consider the friction required to stop the robot sliding down the slope ignoring drive forces. Coulomb friction acts as a second requirement that must be satisfied in eqn. 5.

$$\tan \theta < \mu \quad (5)$$

$$\begin{aligned} \theta &< 26.57^\circ (\mu = 0.5) \\ \theta &< 40^\circ (\mu = 1) \end{aligned}$$

Without a tail the inequality (eqn. 1) can be used with the two centres of mass:

$$\begin{aligned} 0.5 > \sin \theta & \quad 0.25 > \sin \theta \\ \theta < 30^\circ & \quad \theta < 14.48^\circ \end{aligned} \quad (6)$$

The first of these would not be possible with $\mu = 0.5$ but is a theoretical limit for this design if the frictional coefficient was increased.

With a tail the inequality (eqn. 3) can be used as $D \rightarrow 0$:

$$\begin{aligned} \tan \theta &< \frac{4}{15} \\ \theta &< 14.9^\circ \end{aligned} \quad (7)$$

Recalculated with $\mu = 1$, this gives an angle of $\theta = 21.8$ with a tail.

V. SIMULATION RESULTS

In this section these mathematical models are compared with a similar set up created in V-rep simulations. In simulation, the slope varies from 0 to 45 degrees in increments of one with 10 runs at each angle. This means that there are 460 runs for each result given in Table II. The results are estimated by adding together the number of successful runs and using that to calculate the middle of the region where the robot fails. Two offset masses are used to set the centre of mass to different locations in Fig 3.




Fig. 3. Simulation setup. Left: the angle of the slope was varied and the robot with tail deployed Right: the robot with offset mass and no tail.

A. Comparison of Results

The results from Table II match well to the mathematical predictions. The mathematical predictions are an upper limit and therefore having results being just below the predictions is to be expected.

TABLE II
SLOPES - SIMULATED VERSES MATHEMATICAL PREDICTIONS

Centre of Mass	μ	Target Speed (rad/s)			Mathematical Prediction (deg)
		1	2	3	
		Max Slope Angle (deg)			
Quarter	1	10	13	14.1	14.8
	0.5	10	13	14	14.8
Half	1	20	28	29	30
	0.5	20	26	26	26.6
Tail	1	20.9	21.9	21.6	21.8
	0.5	13.5	13.8	13.7	14.9

The speed of the robot was input into the V-rep simulation as a value of intrinsic target velocity. A proportional controller was used when the tail was not present to balance the robot. It was tuned to a target speed of 3 rad/s hence the lower speeds were not as close to the prediction.

The results for the tail were less affected by the target speed of the robot, this could be because the tail allows the movement to be more stable. With the results matching it is promising that they may translate to the real world, and be useful in simulation-based path-planning models.

VI. CONCLUSION

In this work a mathematical analysis of an internally weighted two wheeled robot, with and without a tail, was undertaken. The results from these calculations were compared to results from V-Rep simulations. Three factors were varied so that a comparison could be made. Firstly how the robot design is balanced, where the centre of mass is or whether it has a tail. The second was the frictional coefficient of wheels and finally the intrinsic target velocity. A similar approach was used for steps, however the results did not line up with predictions. The results from slope mathematical predictions are encouragingly accurate and could applied to help plan routes in unstructured terrain.

REFERENCES

- [1] D. Hougen, S. Benjaafar, J. Bonney, J. Budenske, M. Dvorak, M. Gini, H. French, D. Krantz, P. Li, F. Malver, B. Nelson, N. Papanikolopoulos, P. Rybski, S. Stoeter, R. Voyles, and K. Yesin, "A miniature robotic system for reconnaissance and surveillance," in *ICRA*, 2000.
- [2] C. Zheng and K. Lee, "WheelLeR: Wheel-Leg Reconfigurable Mechanism with Passive Gears for Mobile Robot Applications," in *ICRA*, 2019.
- [3] P. Abad-Manterola, J. A. Edlund, J. W. Burdick, A. Wu, T. Oliver, I. A. Nesnas, and J. Cecava, "Axel," *IEEE Robotics and Automation Magazine*, vol. 16, no. 4, pp. 44–52, 12 2009.
- [4] K. Nagatani, M. Kuze, and K. Yoshida, "Development of a Transformable Mobile Robot with a Variable Wheel Diameter," *Journal of Robotics and Mechatronics*, vol. 19, no. 3, pp. 252–257, 6 2007.
- [5] Coppel Robotics, "V-Rep." [Online]. Available: <https://www.coppeliarobotics.com/>

The Goods and Bads in Dyadic Co-Manipulation: Identifying Conflict-Driven Interaction Behaviours in Human-Human Collaboration

Illimar Issak
School of Computer Science
University of Lincoln
Lincoln, UK
illimarissak@gmail.com

Ayse Kucukyilmaz
School of Computer Science
University of Nottingham
Nottingham, UK
Ayse.Kucukyilmaz@nottingham.ac.uk

<https://doi.org/10.31256/Fv3Gn1L>

Abstract—One of the challenges in collaborative human-robot object transfer is the robot’s ability to infer about the interaction state and adapt to it in real time. During joint object transfer humans communicate about the interaction states through multiple modalities and adapt to one another’s actions such that the interaction is successful. Knowledge of the current interaction state (i.e. harmonious, conflicting or passive interaction) can help us adjust our behaviour to carry out the task successfully. This study investigates the effectiveness of physical Human-Human Interaction (pHHI) forces for predicting interaction states during ongoing object co-manipulation. We use a sliding-window method for extracting features and perform online classification to infer the interaction states. Our dataset consists of haptic data from 40 subjects who are partnered to form 20 dyads. The dyads performed collaborative object transfer tasks in a haptics-enabled virtual environment to move an object to predefined goal configurations in different harmonious and conflicting scenarios. We evaluate our approach using multi-class Support Vector Machine classifier (SVMc) and Gaussian Process classifier (GPc) and achieve 80% accuracy for classifying general interaction types.

Index Terms—Classification, Feature Extraction, Haptics, Physical Human-Human Interaction, Physical Human-Robot Interaction, Learning and Adaptive Systems

I. INTRODUCTION

Physical human-human interaction (pHHI) is complex; it involves good interpersonal coordination and mutual role adaption Melendez-Calderon et al. [1]. These help humans to determine how and when their partner’s goals and the overall interaction states change, allowing them to enhance their movements Takagi et al. [2]. Learning how and when the interaction states change in pHHI has important implications for physical human-robot interaction (pHRI). A robot which can accurately infer the current interaction state can use that information to adjust its behaviour to better complement the human partner during pHRI. In [3], we presented a feature extraction method to perform online classification for distinguishing between interaction states during pHHI as an effort to understand how two human partners’ interactive states change over physical collaboration. This paper summarizes our classification results using the data collected during a dyadic object transfer using Madan et al.’s behavior taxonomy [4].

II. BACKGROUND

The data was collected using a virtual environment where human dyads interact through the haptic channel [4]. 40 volunteers, who got randomly matched to form dyads, participated in the study. The dyads collaborated in order to move an object in between target configuration. Two scenes with various scenarios that involve rotation and translation movements were created to provoke a range of different interaction patterns inducing conflicts and harmony.

Madan et al.’s taxonomy assumes that there are three main types of interaction in any collaborative task between humans: T1. Harmonious interaction, T2. Conflicting interaction, and T3. Neutral interaction. Using this assumption we observed the frequently emerging patterns from the interaction and classed them into 6 task dependent interaction pattern classes as shown in Figure 1. The interaction pattern classes fall under interaction type classes as follows: C1: Harmonious translation, C2: Harmonious rotation with translation, C3: Harmonious braking, C4: Persistent conflict, C5: Jerky conflict, and C6: Passive agreement

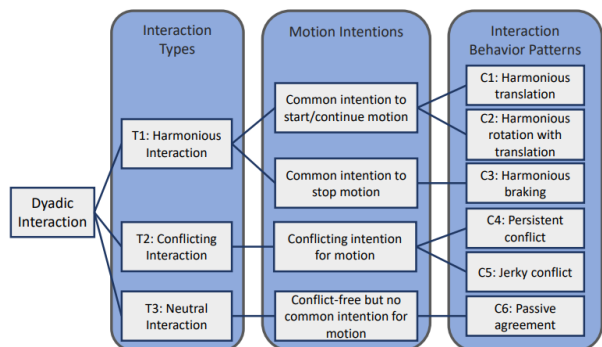


Fig. 1. Interaction taxonomy proposed by Madan et al. [4].

III. METHODOLOGY

The dataset consists of variable length annotated interaction segments. In order to perform online classification we set a

short window to extract a small sequence to be used for feature extraction. We extract 2 seconds worth of features every 1.5 seconds. The window parameters are empirically set such that enough data is mined for accurate classification while also taking into account human reaction times.

A. Online feature extraction

Before we begin with the extraction process, we prepare the raw timeseries data for processing. We assign every data point a label that matches the annotation defining the class of the corresponding interaction segment. As we iterate through the interactions, our window encounters segments that belong to more than one class. This can result in ambiguity of the segment's class and training on such segments can reduce the performance of our model. We deal with such ambiguity by checking the prominence of each class, and drop the window if the difference between the most prominent and second most prominent label is less than 20%. From windows that are not dropped, we extract features using the feature definitions in [4]. For each window, we compute the mean, standard deviation, median, and interquartile range for each of the variables. The feature set contains 48 features, which are normalized before being used for training and testing.

IV. RESULTS

We investigate online classification performance of Support Vector Machine (SVMc) and Gaussian Process (Gpc) classifiers in two layers of Madan et al.'s hierarchy, namely on both task-dependent and task-independent behaviours. The performance of our model is evaluated using confusion matrices and by reporting the correct classification rates.

A. Experiment 1: Online classification of interaction patterns

In the first experiment, we investigate how our approach performs in distinguishing task-dependent interaction patterns. Figure 2 shows the confusion matrices. Our results indicate that SVMc reaches a 78.04% accuracy, whereas Gpc achieves an accuracy of 80.79% on the online feature set, with 2.75% improvement on the performance of SVMc.




Fig. 2. Confusion matrices for SVMc and Gpc for the online classification of interaction patterns

B. Experiment 2: Online classification of interaction types

In the second experiment we look at our model's performance for distinguishing task independent behaviors, namely harmonious, conflicting and neutral interactions. The SVMc and Gpc achieved 83.31% and 83.40% accuracy respectively. The confusion matrices are shown in Figure 3




Fig. 3. Confusion matrices for SVMc and Gpc for the online classification of interaction types

The following table summarises the previously discussed online classification performances and compares them to offline classification performance.

A COMPARISON OF SVMc AND Gpc CLASSIFICATION PERFORMANCES

	Offline classification	Online classification	
	Int. Patterns	Int. Patterns	Int. Types
SVMc	86.1%	78.0%	83.3%
Gpc	87.2%	80.8%	83.4%

V. FINDINGS

Our experiments indicate that haptic data can be used for accurate classification of human interaction types and patterns in real-time. We also demonstrate our windowing method as a viable online feature extraction method for timeseries classification to identify interaction states during ongoing physical collaboration. The results indicate that both Gpc and SVMc perform well at online classification of interaction states with our feature extraction technique. Gpc achieves a slightly better accuracy but at the cost of much longer training time.

VI. FUTURE WORK

This study acts as a first step to build a proactive robotic partner, which can assist a human, while being aware of the interaction state that the partners are in. Our study also demonstrates that haptic data is extremely useful for physical interaction inference. In future work we intend to design and experiment with more sophisticated haptic features, to see how much useful information can be carried through the haptic channel. We also aim to combine haptics with other modalities, such as vision and muscle activity in order to build a more comprehensive model for interaction and individual user states and intentions. This model could then be used in HRI to define proactive robot behaviours and/or role arbitration as described in [5].

REFERENCES

- [1] A. Melendez-Calderon, V. Komisar, and E. Burdet, "Interpersonal strategies for disturbance attenuation during a rhythmic joint motor action," *Physiology & Behavior*, vol. 147, pp. 348–358, 2015.
- [2] A. Takagi, G. Ganesh, T. Yoshioka, M. Kawato, and E. Burdet, "Physically interacting individuals estimate the partner's goal to enhance their movements," *Nature Human Behaviour*, vol. 1, no. 3, 2017.
- [3] A. Kucukyilmaz and I. Issak, "Online Identification of Interaction Behaviors From Haptic Data During Collaborative Object Transfer," *IEEE Robotics and Automation Letters*, 5(1), pp. 96-102, 2019.
- [4] C. E. Madan, A. Kucukyilmaz, T. M. Sezgin, and C. Basdogan, "Recognition of Haptic Interaction Patterns in Dyadic Joint Object Manipulation," *IEEE Transactions on Haptics*, vol. 8, no. 1, pp. 54–66, 2015.
- [5] A. Kucukyilmaz, S.O. Oguz, T.M. Sezgin, and C. Basdogan. Improving human-computer cooperation through haptic role exchange and negotiation. *Immersive Multimodal Interactive Presence*, pp. 229-254, 2012.

Does Expression of Grounded Affect in a Hexapod Robot Elicit More Prosocial Responses?

Luke Hickton
EECAIA Lab, ASRG,
Dept. of Computer Science
Univ. Hertfordshire, UK
l.hickton2@herts.ac.uk

Matthew Lewis
EECAIA Lab, ASRG,
Dept. of Computer Science
Univ. Hertfordshire, UK
m.lewis4@herts.ac.uk

Kheng Lee Koay
ASRG,
Dept. of Computer Science
Univ. Hertfordshire, UK
k.l.koay@herts.ac.uk

Lola Cañamero
EECAIA Lab, ASRG,
Dept. of Computer Science
Univ. Hertfordshire, UK
l.canamero@herts.ac.uk

<https://doi.org/10.31256/Hz3Ww4T>

Abstract—We consider how non-humanoid robots can communicate their affective state via bodily forms of communication, and the extent to which this can influence human response. We propose a simple model of grounded affect and kinesic expression and outline two experiments (N=9 and N=180) in which participants were asked to watch expressive and non-expressive hexapod robots perform different ‘scenes’. Our preliminary findings suggest the expressive robot stimulated greater desire for interaction, and was more likely to be attributed with emotion. It also elicited more desire for prosocial behaviour.

Index Terms—Human Robot Interaction (HRI), Situated Robots, Expression, Kinesics, Embodied Affect

I. INTRODUCTION

Research in the field of social psychology relating to the expression and interpretation of affect has typically focussed on facial expressions [1]. Most Human Robot Interaction (HRI) research tends to reflect this trend, with the generation and interpretation of facial expressions gaining more attention than studies of bodily forms of expression [2]–[5]. Furthermore, much of this work pertains to humanoid morphologies.

This paper describes how animal-like forms of bodily expression, coupled with a grounded model of affect, could enable situated robots of varying shapes and sizes to effectively communicate their needs in a socially evocative manner [6]. Our approach differs from works such as [7]–[9] in that we have adopted a robot-centred approach [10] by seeking to model the underlying substrate of emotion and ground expression in actions that provide adaptive benefits to the robot.

The topics of emotion, expression and context of interpretation are referred to throughout this paper, and therefore they will be introduced briefly below.

A. Emotion

Emotion can be described broadly in terms of physiological arousal, expressive behaviours and conscious experience [11]. There are two predominant perspectives in terms of the classification of emotion: discrete and dimensional. Discrete theories, which include [12]–[15], propose that there are a finite number of distinguishable basic emotions whilst dimensional models, such as [16]–[18], seek to represent the key aspects of emotion using a series of continuous axes. The dimensional model utilised in this work will focus on arousal only, leaving

valence to be inferred by the observer from the environmental context.

B. Expression

Expression can be considered as the communication of emotion via facial and bodily movement. Some argue that such signals are principally aimed at influencing the behaviour of others within a social group [19], whilst others suggest they are accurate indicators of underlying emotional state [12]. Darwin was amongst the first to suggest that expressive communication may have arisen from mechanisms that provide adaptive benefits [20]. This work is consistent with the Darwinian perspective in that we propose kinesic responses that are primarily intended to provide adaptive benefits to the robot, rather than attempting to convey the outward aspects of discrete human emotions.

C. Context

The interpretation of kinesic signals does not occur in a vacuum, and the broader environmental context will determine how such information is processed by an observer. Heider and Simmel first noted the importance of situational context, noting that this element was seldom considered in studies of kinesics [21]. Whilst certain characteristics, such as origin of movement [21] and changes of speed or direction [22], have been shown to create a perception of animacy, the nature of the robot’s interactions with its environment will also determine whether it is attributed with motivations, beliefs or desires: a mode of interpretation Dennett describes as the Intentional Stance [23].

II. RESEARCH QUESTIONS

In consideration of the points outlined above, the following hypotheses were defined in order to examine the processes humans use to make judgements about robots, make sense of their behaviour, and determine how to respond to them:

- The actions of a robot are more likely to be interpreted as those of an intentional agent if it is able to express arousal.
- A robot that is able to express arousal will evoke greater empathy and emotional response from human observers.
- A robot that is able to express arousal will ultimately provoke greater desire for prosocial interaction.




Fig. 1. Diagram illustrating the model of affect employed.

III. ARCHITECTURE

Fig 1 illustrates our model of affect, which is loosely based on a mammalian stress response. The model features two hormones, E and C, which broadly correspond to epinephrine and cortisol in mammals. The first provides a rapid, but brief, response to relevant external stimuli whilst the second is a longer-term response to repeated stressful episodes or deficits in internal variables. These hormones directly modulate five kinesic properties: stance radius; stance height; step length; step height and movement speed [24]. Each of these properties affect the movement of the robot, providing both an adaptive benefit and an associated cost. For example, faster movement speed enables rapid response to potential threats, but also consumes more energy and places strain on the robot's actuators.

IV. EXPERIMENTS AND PRELIMINARY RESULTS

Two related experiments were conducted to test the hypotheses outlined above. The first was a qualitative study that provided detailed insights and identified areas of particular interest. A second study enabled us to build on these insights and capture data from a much larger group of participants. The dependent variables in both experiments were the participant's general perception of the robot, their emotional response towards it, their understanding of its behaviour and motivations; and, ultimately, their willingness to actively assist it. The independent variable was the robot's expressive capability. Therefore a between-group design was adopted for both experiments to facilitate control of this attribute, with participants being divided evenly into two groups, A and B. The additional forms of expressive responses were enabled for group B, but not for the control group A.

In the first experiment, a total of nine participants were asked to observe a hexapod's behaviour as it performed in six discrete episodes, each lasting between two and four minutes. These episodes were designed to tell a story by creating a situation for the robot that an observer could interpret and respond to: an approach that has often been adopted using human actors [25]. The stories were also intended to be comprehensible from the situational context alone, enabling them to be usable for both groups. After each episode, a brief semi-structured interview was conducted, during which participants were asked to describe what happened during the scene, any particularly key moments, how they felt about the scenario and the robot's behaviour and whether they would have liked to have interacted with either the robot or its

TABLE I
TABLE SUMMARISING THE PRELIMINARY RESULTS OF OUR SECOND
(N=180) EXPERIMENT

Theme	Group					
	Yr4 A	Yr4 B	Yr5 A	Yr5 B	Yr6 A	Yr6 B
Attribution of Emotion	1.7%	8.6%	9.5%	12.9%	9.8%	15.5%
Empathy Toward Robot	0%	2.6%	0.9%	4.3%	14.3%	16.4%
Intentional Stance Adopted	11.2%	14.7%	17.2%	30.2%	37.5%	45.7%
Interaction Envisaged	56.0%	67.2%	69.8%	74.1%	78.6%	84.5%
Prosocial Disposition	26.7%	29.3%	22.4%	36.2%	27.7%	52.6%

environment. These interviews were intended to ascertain the mode of interpretation they had adopted when watching the robot, their feelings towards it and whether they would have liked to intervene in order to assist or hinder it.

The results of this study indicated that the expressive robot was attributed with emotion roughly three times more frequently than the non-expressive one, and that expression also appeared to significantly influence desire for interaction. However, we found no evidence of a link between expression of arousal and desire for prosocial behaviour on the part of the observer.

Our second experiment was conducted at a local primary school. A total of 180 children took part, selected from year groups 4-6 (age range 8-11). The event was run over six days, with a single class of approximately 30 children taking part each day. As with the previous experiment, the participants were asked to watch a hexapod robot perform a number of 'scenes'. Group B observed the robot with its arousal model enabled, whereas the control group A viewed it with the model dormant. Videos were used to ensure repeatability. Between each scene, the children were asked to complete a brief questionnaire consisting of eight questions. The first five were multiple choice questions that were designed to establish the participant's broad disposition toward the robot, whilst the remaining three requested short written responses describing how the video made them feel, what they would have liked to do if they were in the video and why.

Our preliminary findings, summarised by Table I, suggest that group B participants were much more likely to adopt an intentional stance when describing the behaviour of the robot, and more likely to experience emotional empathy towards the robot. Consequently, they were far more likely to suggest prosocial forms of behaviour intended to help the robot when asked how they would have liked to have interacted with it. A comprehensive analysis of the results is currently underway.

V. CONCLUSIONS AND FUTURE WORK

To date, our work has focussed on kinesics in the context of open-loop interaction: participants describe how they would like to interact with the robot, but there is no continuous feedback cycle. Future work will close the loop by engaging participants in a shared task that requires continuous interaction with the robot. This task will be designed to create tension between the robot's need to maintain homeostasis and the participant's desire to achieve other objectives. This sets the stage for us to determine how bodily forms of expression can influence the willingness of humans to accommodate the robot's needs, even when they may conflict with their own.

REFERENCES

- [1] B. de Gelder, "Why bodies? Twelve reasons for including bodily expressions in affective neuroscience," *Philosophical Transactions of the Royal Society B: Biological Sciences*, vol. 364, no. 1535, pp. 3475–3484, 12 2009. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pmc/articles/PMC2781896/>
- [2] A. Beck, L. Cañamero, A. Hölle, L. Damiano, P. Cosi, F. Tesser, and G. Sommavilla, "Interpretation of emotional body language displayed by a humanoid robot: A case study with children," *International Journal of Social Robotics*, vol. 5, 08 2013.
- [3] A. Beck, B. Stevens, K. A. Bard, and L. Cañamero, "Emotional body language displayed by artificial agents," *ACM Trans. Interact. Intell. Syst.*, vol. 2, no. 1, Mar. 2012. [Online]. Available: <https://doi.org/10.1145/2133366.2133368>
- [4] M. Lewis and L. Cañamero, "Are discrete emotions useful in human-robot interaction? Feedback from motion capture analysis," *Proceedings - 2013 Humaine Association Conference on Affective Computing and Intelligent Interaction, ACII 2013*, pp. 97–102, 2013.
- [5] K. L. Koay, G. Lakatos, D. S. Syrdal, M. Gácsi, B. Bereczky, K. Dautenhahn, A. Miklósi, and M. L. Walters, "Hey! There is someone at your door. a hearing robot using visual communication signals of hearing dogs to communicate intent," in *2013 IEEE Symposium on Artificial Life (ALife)*, April 2013, pp. 90–97.
- [6] C. Brazeal, *Designing sociable robots*. Cambridge, MA: MIT Press, 2002.
- [7] F. Kaiser, K. Glatte, and M. Lauckner, "How to make nonhumanoid mobile robots more likable: Employing kinesic courtesy cues to promote appreciation," *Applied Ergonomics*, vol. 78, pp. 70–75, 07 2019.
- [8] G. Lakatos, M. Gácsi, V. Konok, I. Brúder, B. Bereczky, P. Korondi, and A. Miklósi, "Emotion attribution to a non-humanoid robot in different social situations," *PloS one*, vol. 9, p. e114207, 12 2014.
- [9] J. Novikova and L. Watts, "A design model of emotional body expressions in non-humanoid robots," in *Proceedings of the Second International Conference on Human-Agent Interaction*, ser. HAI '14. New York, NY, USA: Association for Computing Machinery, 2014, pp. 353–360.
- [10] K. Dautenhahn, "Socially intelligent robots: dimensions of human-robot interaction," *Philosophical Transactions of the Royal Society B: Biological Sciences*, vol. 362, no. 1480, pp. 679–704, 04 2007.
- [11] K. Scherer, "What are emotions? And how can they be measured?" *Social Science Information*, vol. 44, no. 4, pp. 695–729, 2005.
- [12] P. Ekman, *Emotion in the human face*, 2nd ed. Cambridge: Cambridge University Press, 1982.
- [13] J. P. Scott, "The function of emotions in behavioural systems: A systems theory analysis," in *Emotion: Theory, Research, and Experience*, H. Plutchik, R. Kellerman, Ed. New York: Academic, 1980, vol. 1, pp. 35–56.
- [14] C. Izard, *Human Emotions*. New York: Springer US, 1977.
- [15] S. Tomkins, *Affect, Imagery, Consciousness*. Springer, 1962, vol. 1.
- [16] J. Mehrabian, A. Russell, *An Approach to Environment Psychology*. Cambridge, MA: The MIT Press, 1974.
- [17] D. Watson and A. Tellegen, "Toward a consensual structure of mood," *Psychological bulletin*, vol. 98, pp. 219–35, 10 1985.
- [18] H. Lövhelm, "A new three-dimensional model for emotions and monoamine neurotransmitters," *Medical Hypotheses*, vol. 78, no. 2, pp. 341 – 348, 2012.
- [19] A. Fridlund, *Human facial expression: An evolutionary view*. San Diego: Academic Press, 1994.
- [20] C. Darwin, *The Expression of the Emotions in Man and Animals*. London: John Murray, 1872.
- [21] F. Heider and M. Simmel, "An experimental study of apparent behaviour," *The American Journal of Psychology*, vol. 57, no. 2, pp. 243–259, 1944.
- [22] P. Tremoulet and J. Feldman, "Perception of animacy from the motion of a single object," *Perception*, vol. 29, no. 8, pp. 943–951, 2000.
- [23] D. Dennett, *The Intentional Stance*. Cambridge, Mass: MIT Press, 1989.
- [24] L. Hickton, M. Lewis, and L. Cañamero, "A flexible component-based robot control architecture for hormonal modulation of behaviour and affect," in *Towards Autonomous Robotic Systems*, Y. Gao, S. Fallah, Y. Jin, and C. Lekakou, Eds. Cham: Springer International Publishing, 2017, pp. 464–474.
- [25] H. Wallbot, "Bodily expression of emotion," *European Journal of Social Psychology*, vol. 28, pp. 879–896, 1998.

Optimising Soft Fin Ray Robotic Fingers using Finite Element Analysis to Reduce Object Slippage

Joshua Emerson
School of Engineering
University of Lincoln
Lincoln, United Kingdom
josh.d.emerson@gmail.com

Dr Khaled Elgeneidy
School of Engineering
University of Lincoln
Lincoln, United Kingdom
KElgeneidy@lincoln.ac.uk

<https://doi.org/10.31256/Dy2Bn4P>

Abstract—This research aims to optimise soft fin ray robotic fingers by finding relationships between design parameters and the performance of the fingers using simulation and eventually experimental validation. The design parameters that were chosen to optimise were the rib angle increments of the fingers. This was done by simulating an object being displaced in the x-direction moving into the finger to simulate gripping, then being displaced in the y-direction out of the gripper to simulate slipping. To measure the gripping performance of the fingers, the deformation of the fingertip, as well as force reaction of the finger were recorded. This paper includes initial results from the conducted simulations.

Keywords—soft robotics, optimisation, simulation, FEA

I. INTRODUCTION

This project involves the optimisation of a soft Fin Ray robotic fingers using Finite Element Analysis and experimental validation. The Fin Ray Effect® is based on the deflection of fish fin ribs that allows them to conform to different shapes [1]. Designing robotic fingers based on the Fin Ray effect enables passive adaptation to the shape of a grasped object which increases the contact area for a more stable grasp [2]. Creating Fin Ray fingers from soft materials enables the picking of delicate objects that are could otherwise be damaged by rigid fingers [3]. This has many applications especially in the agriculture/food industry which involves moving or harvesting delicate produce with varying delicacy and shapes. The fingers need to be able to adapt to the shape of the target object while applying the minimal amount of force that is necessary so as not to damage the object. They also need to be able to exert larger forces for denser targets or for harvesting tasks. This is where the concept of layer jamming, which is the friction between deforming ribs, is used to create variable stiffness in the soft material fingers [3]. For this reason, the aim is to optimise for both shape adaptation and force generation while still allowing minimal contact forces for delicate objects by changing key design parameters of the fingers.

II. LITERATURE REVIEW

A. Soft Robotic Grippers

Soft robotic grippers have many advantages over rigid-body robotic grippers, one such advantage is their ability to handle delicate objects without damaging them. Soft grippers are usually constructed of flexible materials that allow them to conform to many different shapes in different circumstances. With more demand to automate agricultural and industrial processes there is a need for manipulators that

are adaptable and can safely handle delicate objects, soft robotic grippers are capable of filling these roles [4].

B. Hyperelastic Materials and Mathematical Models

Modelling of the materials used for soft robotic grippers is very important as it allows researchers to quickly and efficiently test changes to the geometry or parameters of the grippers without having to produce a physical copy and test it experimentally. This research uses Ninjaflex which is “a thermoplastic polyurethane” [5]. It is a material used in FDM (Fused Deposition Modelling) 3D printers, and is a good candidate to create soft grippers due to its material properties that offer flexibility and strength.

Previous work showed that only two pairs of parameters are needed to render the measured experimental data in a proper way [5] using the Ogden’s material model. The values that were found were then applied to a simulation software and the results that were produced were consistent with the experimental data [5]. This was the model used for the simulations in this research.

C. Fin Ray Robotic Fingers

The Fin Ray effect works well for passive grippers, these only require power to open and close and not to sustain grip on an object. Previous research has shown that these fingers are well suited to picking up “convex and/or light” objects [6], however, struggled with thin, heavy objects. It has also been shown that a gripper using fin ray fingers was much more effective at holding objects when it was orientated perpendicular to the ground rather than parallel to it, being able to hold double the weight. Changing the geometry of the fingers, such as the angle and position of the ribs, effects the force generation and shape adaptation of the finger. Previous research has shown that when these ribs are orientated in specific ways friction between the ribs can occur causing an increase to the stiffness of the finger and in turn increases the force generation of the finger [7]. This is called the layer jamming effect and presents opportunities for fin ray fingers to be able to have variable force generation depending on the amount of layer jamming occurring in the finger [7]. Further work was done looking at using a fixed increment when increasing rib angle as well [8].

The contribution of this research is to explore the effect of rib angle increment variation on force generation and tip displacement of the finger. A particular interest here is to evaluate how the angle increments affect the slippage of a grasped object out of the fingers.

III. FEA SIMULATION

The initial stage of this research required many simulations to be run to attempt to find the most important design features that have the largest impact on the targeted performance characteristics, those being shape adaptation and force generation. The incremental changes in rib angle were one of these such design features found to influence the performance of the finger. A twelve ribbed finger had been used for previous research with a fixed angle increment. It was not possible to optimise for the twelve rib angles simultaneously this finger due a limitation of ANSYS software in relation to the number of input parameters. Therefore, the geometry was reduced to only eight ribs to simplify the optimisation process. The original increment of 3 degrees between each rib gap was used as the initial angle and a boundary of three degrees was applied for each rib angle to attempt to optimise the finger. The simulation involves a circular object being displaced in the x-axis into the centre of the finger (stage 1), then being displaced in the y-axis up out of the finger (stage 2). This was intended to simulate the finger grasping an object then having the object slipping out from the finger.

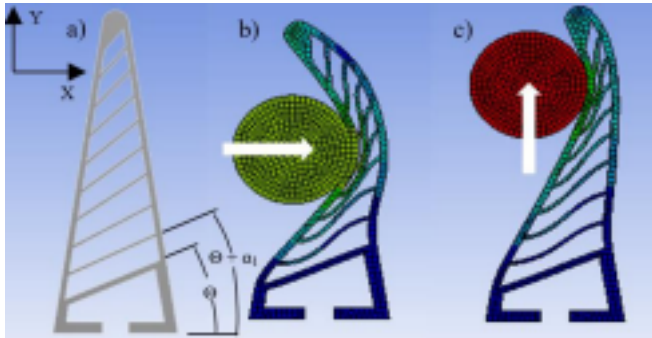


Figure 1: Shows (a) before deformation with labelled design parameters tabulated below (b) Stage 1 in the x-direction and (c) Stage 2 in the y-direction.

Table 1: This shows the optimised increment between each of the ribs

	Alpha 1	Alpha 2	Alpha 3	Alpha 4	Alpha 5	Alpha 6	Alpha 7
Angle Increment (degrees)	3.33	5.39	3.02	2.99	2.53	3.43	2.95

The new angle increments found by the optimisation can be found in Table 1. They didn't follow a regular progression, but rather tended to remain around $3^\circ \pm 0.5^\circ$. There was an exception as the increment between ribs 2 and 3 was over 1.5 times larger than any of the other increments. This may be due to the contact of the object with the finger, but further research is needed to determine the true cause of this.

IV. RESULTS AND DISCUSSION

The maximum force reactions in the x and y-axis were measured during both stages of the simulations so that there could be a comparison made between the original design and the optimised design to highlight improvements. Table 2 shows the results of the two simulations during the first stage and the percentage difference between the values found. Table 3 shows the same during the second stage.

Table 2: Simulation results for object displacement in the x-axis

	Deformation of Tip Y-axis (mm)	Deformation of Tip X-axis (mm)	Maximum Force Reaction X-axis (N)	Maximum Force Reaction Y-axis (N)
Optimised	-7.9272	-3.1923	2.8885	-0.1714
Percentage Difference	4%	22%	11%	-2%

Table 3: Simulation results for object displacement in the y-axis

	Deformation of Tip Y-axis (mm)	Deformation of Tip X-axis (mm)	Maximum Force Reaction X-axis (N)	Maximum Force Reaction Y-axis (N)
Optimised	-0.0063581	-0.0009634	2.4732	0.35513
Percentage Difference	2%	52%	13%	4%

These results show that changing the rib angles can have a large effect on the performance of the fin ray finger. Every aspect of the finger performance was improved by the optimisation except for the force reaction in the y-direction during the first stage of the simulation. An increase of 22% in the directional deformation of the tip of the finger in the x-axis during the first stage shows a large improvement in the fingers ability to grasp an object more securely and reduce the likelihood of the object slipping. There were also improvements in the force reaction on the object in the x-axis during both stages of the simulation showing a potentially firmer gripping finger. Finally, there was an increase of 4% in the force reaction in the y-axis during stage 2 of the simulation. This suggests the finger would be less susceptible to having the object slip out of its grasp.

V. CONCLUSIONS AND FUTURE WORK

The optimised geometry of the finger was shown to improve shape adaptation as well as produce small improvements in the force applied in the y-axis during the simulation of object slippage. Possible future work that could be done on this subject is researching the effects of using variable spacing between the ribs to continue to maximise the force generation and shape adaptation of the fingers.

VI. ACKNOWLEDGEMENTS

I would like to express my appreciation for my research supervisor Dr Khaled Elgeniedy for his guidance and continued support as well Dr Jonathan Griffiths for his advice on simulation techniques. I would also like to thank the School of Engineering at the University of Lincoln for their funding and support.

VII. REFERENCES

- [1] W. Crooks, G. Vukasin, M. O'Sullivan, W. Messner and C. Rogers, "Fin Ray Effect Inspired Soft Robotic Gripper: From the RoboSoft Grand Challenge Toward Optimization," *Front. Robot. AI*, 2016.
- [2] C. I. Basson, G. Bright and A. J. Walker, "TESTING FLEXIBLE GRIPPERS FOR GEOMETRIC AND SURFACE GRASPING CONFORMITY IN RECONFIGURABLE ASSEMBLY SYSTEMS," *THE SOUTH AFRICAN JOURNAL OF INDUSTRIAL ENGINEERING*, vol. 29, no. 1, pp. 128-142, 2018.
- [3] J. Shintake, V. Cacucciolo, D. Floreano and H. Shea, "Soft Robotic Grippers," *Advanced Materials*, vol. 30, no. 29, 2018.
- [4] J. Hughes, U. Culha, F. Giardina, F. Guenther, A. Rodendo and F. Iida, "Soft Manipulators and Grippers: A Review," *Frontier in Robotics and AI*, vol. 3, 2016.
- [5] T. Reppel and K. Weinberg, "Experimental Determination of Elastic and Rupture Properties," *TECHNISCHE MECHANIK*, vol. 38, no. 1, pp. 104-112, 2018.
- [6] W. Crooks, S. Rozen-Levy, B. Trimmer, C. Rogers and W. Messner, "Passive gripper inspired by *Manduca sexta* and the Fin Ray® Effect," *International Journal of Advanced Robotic Systems*, 2017.
- [7] K. Elgeneidy, P. Lightbody, S. Pearson and G. Neumann, "Characterising 3D-printed Soft Fin Ray Robotic Fingers with Layer," in *2nd IEEE International Conference on Soft Robotics (RoboSoft)*, Seoul, Korea, 2019.
- [8] K. Elgeneidy, A. Fansa, I. Hussain and K. Goher, "Structural Optimization of Adaptive Soft Fin Ray Fingers with," in *IEEE International Conference on Soft Robotics (RoboSoft) (In press)*, 2020.

A Cable-based Manipulator for Chemistry Labs

Lupo Manes, Sebastiano Fichera, David Marquez-Gamez, Andrew I. Cooper, Paolo Paoletti

<https://doi.org/10.31256/Cd1Vr80>

Abstract—This paper presents the design of an end-effector for handling of supplies commonly found in chemistry labs. The system uses a cable loop capable of providing an effective grasp of any prismatic or cylindrical object, making it ideal for handling vials and other containers commonly used in laboratories. When compared to the more common parallel jaw gripper design, the proposed cable based end-effector is able to handle a larger variety of objects without interfering with the surrounding objects even in a crowded environment (minimal footprint). The payload capability of the gripper has been tested on a load test apparatus with different materials, demonstrating its effectiveness.

Keywords—Robotic Gripper, Cable-based Manipulator, Chemistry Automation

I. INTRODUCTION

The current trend in industry and research is pursuing effective human-robot interaction and cooperation, where the two would participate to the same workflow both efficiently and safely. This would make for easier set up and inspection of automated plants, since the safety features and procedures would be embedded in the robotic system. Currently, robot manufacturers provide cooperative robots in the form of low-power robotic arms (<35kg max payload), sometimes mounted on mobile bases. These systems come equipped with collision detection and avoidance and compliance features to avoid harming humans or damaging themselves or the surrounding environment [4].

Chemistry research, and more specifically material discovery, relies on exploring a large number of chemical combinations. Artificial intelligence trained for the task can skim most combinations through simulation and narrow down the possible solutions to a few hundred composites. This software combined with a robotic system makes a robotic scientist that can run the reactions in a lab and has provided some great results [1]. Current systems require a highly controlled environment and can only perform specific tasks[2]. As a result, a significant amount of resources goes into the design and possible modifications of a physical system. The creation of a robot that can work in any chemistry lab with the available supplies is extremely appealing to reduce the set-up costs of chemistry automation and improve its accessibility.

In order to be effective in an environment with variable features, like position of supplies or layout of the lab benches,

All authors are part of the Leverhulme Research Centre for Functional Materials Discovery, Material Innovation Factory, University of Liverpool, Liverpool L69 7ZD. LM, PP and SF are part of the School of Engineering, University of Liverpool, Liverpool L69 3GH. LManes@liverpool.ac.uk, aicooper@liverpool.ac.uk, Seba84@liverpool.ac.uk, David.Marquez-Gamez@liverpool.ac.uk, P.Paoletti@liverpool.ac.uk

the robotic system needs to be extremely robust. Specifically, when it comes to manipulation, a balance needs to be struck between computation (the software) and embodiment (the hardware). Computation has traditionally been the main focus for manipulation task because of the quicker turnaround when compared to mechanical design. However, thanks to the advance in rapid manufacturing technology, the hardware can nowadays be updated at similar pace to the software [5]. This has led to further exploration of potential mechanical designs in an effort to find the right tool for the job instead of finding a workaround for sub-optimal equipment. This is particularly true for the development of end-effectors, where many of the newer designs have become better adapted for their work environment [3]. The end-effector proposed in this paper aims to outperform current designs in their flexibility, footprint and ease of grasping. The use of a cable loop mechanism allows for all of these requirements while also keeping costs low.

II. PROPOSED DESIGN

The proposed design, shown in Fig. 1, consists of a main body containing the drive components and a vertical beam (finger) with a cable loop at its end. A 0.5 mm polyamide cable is used to envelop objects and pull them against the finger. To keep the object aligned to the end-effector, the contact surface on the finger is concave. One end of the cable is fixed at the end of the finger while the other can be pushed or pulled by a two wheel arrangement, shown in Fig. 2, with wheel 1 driven by a geared DC motor and wheel 2 used to keep the wire in contact with wheel 1. Both wheels use a 3D printed flexible material (Ninjaflex™) external cover with a V groove to better grip and feed the cable. The cable is fed into a channel that drives it to the desired position.




Figure 1: Proposed end-effector (left) and a bottom view of the grasping mechanism (right).

All the mechanical components for the design have been 3D printed using PolyLactic Acid (PLA) for rigid components

and Ninjaflex™ for the flexible ones. This made it possible to create channels for the cable to go through inside the structural components.




Figure 2: Cable drive mechanism top view and section. Drive wheel (1), idler (2) and cable in blue in the section view.

An embedded capacitive force sensor (SingleTact™10N) mounted behind the contact surface collects information about the grasping force, and a rotary encoder mounted to the motor shaft is used to estimate the radius of the cable loop. Two PI closed-loop control systems, implemented on an Arduino™Mega board, are used to control position and torque of the motor connected to the cable loop.

III. TESTING

The system has been tested by grasping cylindrical objects of different diameters attached to a load cell to check the maximum payload weight, see Fig. 3.




Figure 3: Assembled test rig.

The tests were conducted with both glass and plastic cylindrical objects of different diameters and materials clamped to the load cell, performing 5 runs for each test sample to insure reliability. For each test, the cable loop was tightened around the cylinder until a grasping force of 5N was registered by the sensor in the end-effector. The drive wheel was then locked in place and the end-effector was slowly lifted using the linear actuator of the load test apparatus shown on the right of Fig. 3.

The results of these tests are summarised in Fig. 4. The end-effector could lift more weight with plastic specimens because of their higher surface roughness and therefore extra friction. In spite of the lower friction of glass, the end-effector was still able to reliably lift 400g of payload. Tests with both materials show no significant direct correlation between diameter of the

object and weight lifted, which suggests that the design could prove to be highly versatile. The maximum load exceeds the requirements for use in a chemistry lab where most vials will weigh less than 100g.




Figure 4: Maximum weight lifted as the diameter of the samples increases. Error Bars represent one standard deviation.

IV. CONCLUSIONS

The proposed end-effector design provides a small footprint and can handle variable payload sizes, thanks to an unconventional grasping mechanism. The cable loop design compromises on the type of shape the system can grasp but has potentially unprecedented flexibility for payload size. Thanks to its small dimension the mechanism could be easily fitted to existing 6 DoF Robotic arms with good reach and dexterity.

To improve its reliability, the prototype will require better materials and electromechanical components and more thorough testing. Moreover, software integration with a robotic manipulator can be developed to allow for testing in a realistic environment.

REFERENCES

- [1] Ross D. King et al. "The robot scientist adam". In: *Computer* 42.8 (2009), pp. 46–54. DOI: 10.1109/MC.2009.270.
- [2] J. Pan. "Engineering Chemistry Innovation". In: *ACS Medicinal Chemistry Letters* 10.5 (2019), pp. 703–707. DOI: 10.1021/acsmchemlett.9b00096.
- [3] K. Tai et al. "State of the art robotic grippers and applications". In: *Robotics* 5.2 (2016), pp. 1–20. DOI: 10.3390/robotics5020011.
- [4] A. Vysocky and P. Novak. "Human - Robot Collaboration in Industry". In: *MM Science Journal* 2016.OCTOBER (2016), pp. 1066–1072. DOI: 10.17973/MMSJ.2016_10_201640.
- [5] S. Wade-McCue et al. "Design of a Multi-Modal End-Effector and Grasping System: How Integrated Design helped win the Amazon Robotics Challenge". In: (2017). arXiv: 1710.01439. URL: <http://arxiv.org/abs/1710.01439>.

An Open Source Seeding Agri-Robot

Harry Rogers

School of Computer Science

University of Lincoln, United Kingdom

15623886@students.lincoln.ac.uk

Charles Fox

School of Computer Science

University of Lincoln, United Kingdom

chfox@lincoln.ac.uk

<https://doi.org/10.31256/Or6Mf2T>

Abstract—To enable wider experimental uptake and build interest in autonomous precision agriculture, we present a low-cost, open source hardware and software platform for automated precision seed planting, for use by researchers and practical agri-hackers. The robot platform is about to plant seeds with higher within-row precision than conventional systems, and using high end Global Navigation Satellite System (GNSS) is able to store the geospatial location of each seed to accuracies sufficient to enable revisiting of plants during growth. It can also operate with lower cost GNSS using a swappable interface for other applications. We provide software up to the level of pose-to-pose driving using Dubins paths. The robot is based on a low-cost closed-source commercial off the shelf (COTS) vehicle, low-cost COTS seeder, and Raspberry Pi, with our open hardware design showing how to link them to each other and to the software. The hardware and software is being open-sourced at <https://github.com/Harry-Rogers/PiCar> as part of this publication.

Index Terms—Agriculture, Raspberry Pi, GNSS, Mechatronics, Python, Raspbian, Open Source

I. INTRODUCTION

Automated precision agriculture is becoming important around the world due to increased population and demand for food, declining fuel reserves, and climate change [1]. Currently there are many solutions to automate agriculture that involve harvesting crop, spraying pesticide, spraying weeds and planting seeds. There are many solutions like the IBEX robot that sprays weeds [2], the Rubion robot that picks strawberries¹, and Fendt's project MARS robot² that can monitor and accurately document precise planting of corn.

However, these are expensive, large-scale systems, which do not easily allow smaller scale innovators such as researchers, farmers in developing countries, and the maker community, to experiment and extend their designs. Innovation in both farming practice and engineering design often emerges bottom up from these communities. Therefore, this paper makes such a system available to the community.

Like most open source hardware (OSH), the system is based on off-the-shelf components which may not be OSH themselves. It combines low costs COTS vehicle³, seed planter, and embedded computer and provides physical build instructions

¹Rubion robot (Octinion; <http://octinion.com/products/agricultural-robotics/rubion>)

²Fendt's project MARS robot (Fendt; <https://www.fendt.com/int/fendt-mars>)

³COTS vehicle (Sunfounder Ltd; <https://www.sunfounder.com/smart-video-car-kit-v2-0.html>)

and software to run. It can plan and control an efficient point-to-point route including speeding up on long straights and slowing to control turns. Using a GNSS module it stores the Geographic Information System (GIS) coordinates of each seed planted so the system can return to it. This system is cheap enough to be reproduced at scale to create swarms that can complete many processes in a very small amount of time the total cost of the system is shown in table 1. Seed planting is chosen as a first task because even current agri-business systems such as large tractor driven drills are not able to plant seeds at this precision, with spacing between seeds in a wheat crop typically varying by several centimeters. The system is legal to operate autonomously under the guidelines given by [3].

II. SYSTEM DESIGN

A. Open source Hardware

The main components consist of a Raspberry Pi 3 model B+, Robot HATS³ a hat board that is on top of the Raspberry Pi, a PCA9865 [4] board for controlling servos, a TB6612 [5] board to control the DC motors and a camera for the system to see. Alongside this, a COTS seed dispenser [6] and GNSS module [7] have been added to the kit to complete the functionality for the project. Robot HATS sends power to the other boards as well as signals to turn the servos to steer or move the camera. PCA9865 sends signals to the servos to steer the car as well as pan or tilt the camera, it also sends signals to the TB6612 board to either turn the wheels or not. TB6612 sends signals to the DC motors depending on what is needed from the other boards. The seed dispenser is modified to be shorter therefore making it lighter it is attached with glue. Fig. 1 shows the built system. Table 1 shows the per-item cost, the system overall costs £206.96.




Fig. 1. Built system

TABLE I
PER-ITEM COST

Item	Category	Cost(GBP)
COTS Robot	Material	99.99
COTS Seed dispenser	Material	14.18
Raspberry Pi	Material	32.75
Power Adapter	Material	8.29
SD card	Material	5.29
Batteries	Material	16.99
Battery Charger	Material	11.99
GNSS Module	Material	12.48

B. Open Source Software

When turning the system on, the Raspberry Pi starts a script using the Django (<https://www.djangoproject.com/>) web framework so that an application on the user's device, such as a smart phone, can connect and operate the system. There are two main areas of software: movement and data.

1) *Movement*: Pose to pose path planning and control is provided, which can be used to plant seeds at regular intervals along one or more paths. Planning operates by using a Dubins path algorithm. There are three steps to completing this path: a starting turn, going straight, and ending turn. The Dubins [8] path is calculated by creating two arcs around the start and end point that are connected by the shortest straight line between them. The modification over the standard Dubins method is that during the straight, the controller speeds up to its maximum speed to reduce the time taken to be completed. During curves the controller slows down to its minimum speed to reduce skidding.

2) *Data*: The data collected from the GNSS module contains GIS coordinates of where seeds have been planted. The seeds' location is found by calculating the circumference of the dispenser and its location relative to the robot center. The open hole that dispenses the seed is at the top each time the system starts a new task. This is so that when half a rotation is completed the seed is dispensed. National Marine Electronics Association (NMEA) data collected from the GNSS module is processed and currently stored in CSV format though we plan to switch to a GIS database such as PostGIS.

III. TESTING AND RESULTS

A. Design of experiment

An experiment was carried out to test whether or not the system is effective at planting seeds and recording their accurate GIS coordinates. In the experiment, the system must drive to a pose that is 6.8m away using a Dubins path. Along this path it must plant 20 seeds, at regular spacial intervals, and record their GIS coordinates. It is assumed that the system will drive over a cut trench and the seeds will then be covered after it completed its task. Whilst on the path it must also speed up when driving straight and slow down when turning, without affecting the seed spacing.

B. Results

All 20 seeds were planted. The system travelled 11.2m in total to get to the pose that was 6.8m away because of the

Dubins Curve requirements. The planting rate was 1.1 seeds per second. Two seeds were planted at both of the first and second points due to traveling slowly. Seeds also once leaving the seed dispenser were moving at a rate that could cause them to roll or bounce off of the ground. Fig. 2 shows where each seed was planted on the Dubins path. Planted seed ground truth locations were then measured manually with a ruler and found to have a standard deviation of 139mm from their intended and logged locations. Minimum distance between seeds was on average 399mm apart from each other. The top speed of the system was 0.49m/s and the minimum speed of the system was 0.26m/s. The average speed the system was travelling at was 0.39m/s. The coordinates on average were 8.23m away from the actual location, before differential GNSS correction. The system was also able to do the path in reverse completing the path with the same speeds.




Fig. 2. Mapped seeds





IV. CONCLUSION

The observed seed planting accuracy of the system was 139mm. This was due to seeds bouncing off the ground as the trundle was moving very quickly on the straight distance therefore causing seeds to move from their original planting point. One way this system could make farming more efficient is if a swarm was created and the trundle size was reduced. This would mean that multiple rows could be completed at once. The system plants much less seeds per second than a large tractor system, but as it is a system it does not require to be paid by the hour as labourers would it is also open source so everyone can use it. However, some seeds at the earlier locations were not planted individually but were dispensed in groups in error, due to the low speed of the robot when setting off. Minimum distance between seeds was too large (399mm) for crops such as wheat, but could be reduced by using a smaller trundle or trundle gearing. The differential GNSS accuracy is such that driving the robot to revisit specific plants appears possible for seed spacings up to 139mm. The system can follow the exact path taken in reverse to do this. The results suggest that the system would be a useful test bed for research and hobbyist platforms and if scaled up to swarms could make seeding and other operations more efficient. As an open source hardware and software project, we welcome future contributions from the community.

REFERENCES

- [1] T. Duckett, S. Pearson, S. Blackmore and B. Grieve (2018) *Agricultural robotics: the future of robotic agriculture*. arXiv:1806.06762
- [2] A. Binch and C. Fox (2017) *Controlled comparison of machine vision algorithms for Rumex and Urtica detection in grassland*. Computers and Electronics in Agriculture 140:123-138.
- [3] S. Basu, A. Omotubora, M. Beeson and C. Fox (2020) *Legal framework for small autonomous agricultural robots*. AI & SOCIETY volume 35, pages 113–134.
- [4] Sunfounder Ltd; <https://bit.ly/3401mRX>
- [5] Sunfounder Ltd; <https://bit.ly/3dGy7Ih>
- [6] WOLF-Garten GmbH & Co.KG; <https://bit.ly/2WXyGaU>
- [7] u-blox; <https://www.u-blox.com/en/product/ubx-g7020-series>
- [8] Dubins, L.E. (July 1957). "On Curves of Minimal Length with a Constraint on Average Curvature, and with Prescribed Initial and Terminal Positions and Tangents". American Journal of Mathematics. 79 (3): 497–516.

Investigating PID Control for Station Keeping ROVs

Kyle L. Walker*[‡] , Adam A. Stokes* , Aristides Kiprakis[†] , Francesco Giorgio-Serchi* 

*Institute for Integrated Micro and Nano Systems

[†]Institute for Energy Systems

School of Engineering

University of Edinburgh

Edinburgh, UK

[‡]Correspondence: *K.L.Walker-3@sms.ed.ac.uk*

<https://doi.org/10.31256/Ky3Xg3B>

Abstract—For controlling Unmanned Underwater Vehicles (UUVs) in deep water, Proportional-Integral-Derivative (PID) control has previously been proposed. Disturbances due to waves are minimal at high depths, so PID provides an acceptable level of control for performing tasks such as station-keeping. In shallow water, disturbances from waves are considerably larger and thus station-keeping performance naturally degrades. By means of simulation, this letter details the performance of PID control when station keeping in a typical shallow-wave operating environment, such as that encountered during inspection of marine renewable energy devices. Using real wave data, a maximum positional error of 0.635m in the x-direction and 0.537m in the z-direction at a depth of 15 m is seen whilst subjected to a wave train with a significant wave height of 5.404m. Furthermore, estimates of likely displacements of a Remotely Operated Vehicle (ROV) are given for a variety of significant wave heights while operating at various depths. Our analysis provides a range of operational conditions within which hydrodynamic disturbances don't preclude employment of UUVs and identify the conditions where PID-controlled station keeping becomes impractical and unsafe.

Index Terms—station keeping, PID, ROV, underwater robotics, shallow water, thruster dynamics.

I. INTRODUCTION

Offshore industries are becoming increasingly interested in operating with a higher degree of autonomy than currently available. The offshore energy and marine renewable energy sectors, in particular, need to perform systematic maintenance operations and accurate sensor deployment in order to improve structure survivability and reduce overall running costs of the plant [1], [2]. Unmanned Underwater Vehicles (UUVs) have previously been deployed for inspection and maintenance, but are often not equipped for undertaking more complex missions especially close to the sea surface. The task of operating in perturbed sea states remains largely unsolved due to the difficulty of enabling safe station keeping (holding a stationary position) of the vehicle when operating in proximity with submerged structures. This prevents any systematic maintenance operation of offshore structures as well as accurate surveillance of any submerged environments [3]–[5].

This letter aims to investigate the use of one of the most classical feedback control methods, Proportional Integral-Derivative (PID) control, for station keeping in these shallow

water environments. The dynamic response of a vehicle subject to varying wave disturbances is tested via numerical simulations. The performance of the controller is evaluated when the vehicle is subjected to a typical, realistic wave field by monitoring the positional error in the surge, x , and heave, z , directions. Upon evaluation, results confirm that utilising PID control for station keeping yields excessive positional error outside of a narrow band of wave disturbances and operational depth. Therefore, an alternative control method is required to improve the reliability and accuracy of UUV station keeping to assist in the continued drive for fully automated operation and management of offshore structures.

II. MODELLING

The simulated scenario entails a vehicle located at varying depths D within a water column of depth $D_w = 50m$ and attempting to station keep whilst subjected to an oncoming wave train of varying significant height H_s . The wave field simulated in this work was taken from [6] and is representative of a typical wave field seen at the National Northwest Marine Renewable Energy Center (NNMREC) North Energy Test Site (NETS) and correlates with data collected by a buoy deployed in the area; the composition is formed utilising the fundamentals of Airy Wave Theory [7], [8]. The particle velocities and accelerations at the vehicle location are then deduced using widely detailed wave theory [9] and input as a disturbance to the vehicle dynamics, discussed below.

The vehicle modelled in this work is the SeaBotix vLVB300 Remotely Operated Vehicle (ROV) [10], a typical type of underwater vehicle which is controlled by a pilot; in this work the vehicle is controlled autonomously. Only the surge and heave motions are considered as this letter aims to provide evidence for the requirement of more advanced control methodologies. Following the methodology outlined in [11], the vehicle is modelled as a rigid body with 2DOF whilst assuming the vehicle is neutrally buoyant and neglecting the influence of the Coriolis effect. The dynamic equation can then be simplified and expanded to give:

$$m_d \dot{v}_a + m_a \dot{v}_r - \frac{1}{2} \rho_f A_i C_D v_r |v_r| = T \quad (1)$$

where m_d and m_a are the dry and added mass terms respectively, $v_a = (v_{ax}, v_{ay})$ and $v_r = (v_{rx}, v_{ry})$ are the absolute and relative velocity of the vehicle with respect to the water,

This work was supported by the Engineering and Physical Sciences Research Council [grant number EP/R513209/1], the Research Council United Kingdom (RCUK) and the ORCA Hub.

TABLE I

THRUSTER MODEL PARAMETERS FOR THE SIMULATIONS UNDERTAKEN IN SECTION III

Parameter	Nomenclature	Value
Thrust Co-efficient	K_T	0.464
Propeller Diameter	D	0.1 m
Time Step	Δt	0.2 s
Motor Time Constant	T_m	0.1 s
Drag Coefficient, x	$C_{D,x}$	0.84
Drag Coefficient, z	$C_{D,z}$	1.06
Added Mass, x	$m_{a,x}$	8.1 kg
Added Mass, z	$m_{a,z}$	36.7 kg




Fig. 1. A 100s segment of the wave field generated using the parameters detailed in [6] and the displacement of the vehicle caused by this wave field when the vehicle is attempting to remain stationary at a depth of 15m.

ρ_f represents the density of the fluid (in this case sea water), A_i represents the area of the incident side to the flow, C_D represents the drag coefficient and T represents the thrust produced by the propellers.

To accurately describe the behaviour of the vehicle and the controller performance, a thruster model was utilised based on [11], [12] which considers the Bilinear Thruster Model in conjunction with a reduction term, approximating the propeller angular velocity as a first order system. This reduction term accounts for the effects of the fluid flow through the propeller; this model is also utilised in [13], [14] and reads:

$$T = K_T \rho_f n |n| D^4 - \frac{1}{3} v_f \rho_f D^3 |n| \quad (2)$$

where K_t , n , v_f and D respectively represent a thrust constant, the propeller angular velocity, the fluid speed into the propeller disk and the propeller disk diameter.

III. RESULTS

Simulations were performed over a 240s temporal segment and the magnitude of the positional error was recorded. An example of the time history of the vehicle response while performing station-keeping at $D = 15m$ (i.e. for $D/D_w = 0.3$) and subject to a wave train with a significant wave height of $H_s = 5.404m$ (i.e. for $H_s/L = 7.72$) is presented in Fig. 1. From this simulation, the maximum positional error witnessed was $0.635m$ in the x -direction and $0.537m$ in the z -direction.

A series of test cases were performed using the same wave train but for $0.075 < D/D_w < 0.9$ and $0.25 < H_s/L < 3$, where $L = 0.7m$ represents the reference dimension of the vehicle. Maximum positional error in surge and heave are reported in Fig. 2 and 3 and it is observed that for $H_s \approx 3L$ then the positional error approaches $0.4m$ at low depth. Fig. 2-3 highlight the existence of a region of D/D_w and H_s/L values where the profile of error displacement shows a markedly non-linear trend with a steep gradient.




Fig. 2. Maximum error in the surge, x , direction when subject to the wave field in Fig. 1 over a range of H_s/L and D/D_w .




Fig. 3. Maximum error in the heave, z , direction when subject to the wave field in Fig. 1 over a range of H_s/L and D/D_w .

IV. DISCUSSIONS & CONCLUSIONS

If the positional error shown in Fig. 1-3 is considered, the results show that in a typical shallow water environment the level of control offered by PID is unacceptable. All cases operating near the free surface correlate to an increase in positional error as the effects from the waves are increased significantly. Similarly, if the depth is held constant and the wave height is increased, the positional error increases for the same reason. For missions which involve precise manipulation or inspection of fine-scale structural elements, the positional error estimated through these simulations is too large and therefore using PID control is unsuitable.

For this reason, an alternative more advanced control method is required which can offer higher performance; non-linear model-based PID has the potential to substantially improve performance without increasing the required computation power too drastically [15]. Furthermore, the use of Model Predictive Control (MPC) is envisioned; the preliminary work of [6] will be extended to include additional DOF and subsequently a controller will be developed and tested at the FloWave facility at the University of Edinburgh [16], [17].

An alternative solution is to develop a suitable manipulator to constrain the motion of the vehicle when subjected to wave and current disturbances. Unlike the MPC approach, this would rely on the capability of hardware to simply withstand the forces exerted on the vehicle, but would require a structure to grip onto. Hence, the MPC approach is preferred.

REFERENCES

- [1] *The ROV Manual: A User Guid for Observation Class Remotely Operated Vehicles*. Butterworth-Heinemann, 2007.
- [2] *The technology and applications of autonomous underwater vehicles*. UK: Taylor & Francis, 2003, vol. 2.
- [3] J. Elvander and G. Hawkes, "ROVs and AUVs in support of marine renewable technologies," in *OCEANS 2012 MTS/IEEE, Hampton Roads, VA, October 14-19, 2012*.
- [4] S. Sivčev, E. Omerđić, G. Dooly, J. Coleman, and D. Toal, "Towards inspection of marine energy devices using rovs: Floating wind turbine motion replication," in *Advances in Intelligent Systems and Computing*, 2018, pp. 196–211.
- [5] R. Capocci, E. Omerdic, G. Dooly, and D. Toal, "Development and testing of remotely operated vehicle for inspection of offshore renewable devices," in *Technological Innovation for Resilient Systems*, L. M. Camarinha-Matos, K. O. Adu-Kankam, and M. Julashokri, Eds. Cham: Springer International Publishing, 2018, pp. 282–289.
- [6] D. C. Fernandez and G. A. Hollinger, "Model Predictive Control for Underwater Robots in Ocean Waves," *IEEE Robotics and Automation Letters*, vol. 2, no. 1, pp. 88–95, 2017.
- [7] R. G. Dean and R. A. Dalrymple, "Water wave mechanics for engineers and scientists." 1984.
- [8] O. M. Phillips, *The Dynamics of the Upper Ocean*, 2nd ed., 1966.
- [9] D. Reeve, A. Chadwick, and C. Fleming, "Coastal Engineering: Processes, Theory and Design Practice," 2004.
- [10] Teledyne-Seabotix, "vLBV300 ROV Datasheet," in *vLBV & vLBC vectored Little Benthic Vehicles*, 2015.
- [11] T. I. Fossen, *Guidance and control of ocean vehicles*. Chichester: Wiley, 1994.
- [12] A. J. Healey, S. M. Rock, S. Cody, D. Miles, and J. P. Brown, "Toward an improved understanding of thruster dynamics for underwater vehicles," *IEEE Journal of Oceanic Engineering*, vol. 20, no. 4, pp. 354–361, 1995.
- [13] L. Medagoda and S. B. Williams, "Model predictive control of an autonomous underwater vehicle in an in situ estimated water current profile," *OCEANS 2012 MTS/IEEE, Yeosu, South Korea, May 21-24, 2012*.
- [14] K. L. Walker, A. A. Stokes, A. Kiprakis, and F. Giorgio-Serchi, "Impact of thruster dynamics on the feasibility of rovs station keeping in waves," in *OCEANS 2020 MTS/IEEE, Singapore, August 11-14, 2020*.
- [15] C.-J. Wu, "6-dof modelling and control of a remotely operated vehicle," Master's thesis, Flinders University, 2018.
- [16] D. Ingram, R. Wallace, A. Robinson, and I. Bryden, "The design and commissioning of the first, circular, combined current and wave test basin," in *OCEANS'14 MTS/IEEE, Taipei, Taiwan, April 7-10, 2014*, pp. 1–7.
- [17] S. Draycott, B. Sellar, T. Davey, D. R. Noble, V. Venugopal, and D. M. Ingram, "Capture and simulation of the ocean environment for offshore renewable energy," *Renewable and Sustainable Energy Reviews*, vol. 104, pp. 15–29, 2019.

Control of hydraulically-actuated manipulators with dead-band and time-delay uncertainties*

Olivia Albrecht
Engineering Department
Lancaster University
Lancaster, UK
o.albrecht@lancaster.ac.uk

Manuel Bandala
Engineering Department
Lancaster University
Lancaster, UK
m.bandala@lancaster.ac.uk

Stephen D. Monk
Engineering Department
Lancaster University
Lancaster, UK
s.monk@lancaster.ac.uk

C. James Taylor
Engineering Department
Lancaster University
Lancaster, UK
c.taylor@lancaster.ac.uk

<https://doi.org/10.31256/Wd9Qw9E>

Abstract—The research behind this article is motivated by robotic operations in radiologically contaminated environments, notably for nuclear decommissioning. However, the experiments reported within are based on a recently reconfigured, hydraulically-actuated, dual manipulator robot that is being used for R&D into both tele-operation and autonomy in a non-active laboratory setting. One element of this research concerns the development of novel control systems to address time-delay and deadband uncertainties. The article briefly discusses some preliminary results and plans in this regard. Recent improvements to the hardware demonstrator are also described.

Index Terms—nuclear decommissioning, hydraulic actuators, deadband, time-delay, uncertainty

I. INTRODUCTION

A significant number of nuclear facilities around the world have reached the end of their useful life and hence are in the process of decommissioning. Since it is environmentally unfriendly and dangerous for plant workers, many decommissioning tasks are accomplished with robots, for which direct tele-operation is standard [1]. With constrained spaces and highly-contaminated facilities, fully autonomous solutions are unlikely to be considered safe or cost-effective in the near future. Nonetheless, with the advent of more efficient and robust embedded electronics and sensors, there is significant interest in semi-autonomous capabilities [2]–[4].

The present article concerns a previously developed dual-manipulator robotic platform [5]. The system has recently been reconfigured, hence the new hardware framework and control software are described (section II). The broad aim is to develop semi-automatic control systems that reduce operator workload, speeding up task execution and reducing operator training time, whilst minimizing the introduction of additional sensors and other components. Due to limited sensor data availability in nuclear environments, a system for grasping generic objects could be unreliable. As a result, the developed approach is based on the concept of multiple subsystems for common tasks under one user interface: one for pipe cutting, one for pick and place operations, and so on. This aims to reduce the complexity of the problem, potentially leading to improved performance and reliability. Furthermore, cognitive workload

is reduced by tailoring the information shown to the operator. The research focuses on pipe cutting as an illustration of the generic approach, since this is a common repetitive task [6].

Motivated by preliminary testing that highlights limitations in the performance, one aspect of the research programme concerns the development of improved ‘low-level’ control systems for hydraulic manipulators, such that they can more effectively achieve the ‘higher-level’ task orientated objectives. In this regard, it is notable that uncertainties and nonlinearities, including actuator deadbands and time-delays, are not always fully addressed in the literature [7]. In fact, the two major challenges in high performance positioning and tracking stabilisation of robot manipulators, are the friction between moving parts and the deadband of the actuators.

The present work utilises a state-dependent parameter (SDP) framework to characterise the manipulators. The parameters of SDP models are functionally dependent on measured variables, such as joint angles and velocities, normally defined in discrete-time terms [5]. However, in contrast to other recent research into SDP systems for the same machine [6], the present work utilises a new *continuous-time* SDP model that is not dependent on the sampling interval, and uses this to investigate uncertainties (including time-variations) in the system time-delay and deadband (section III).

II. RECONFIGURED HARDWARE

The laboratory demonstrator used in this article consists of two HYDROLEK HLK-7W manipulators, each a 6-degrees-of-freedom articulated arm, with a seventh actuator for the gripper. Whilst the original set-up is described by reference [6], a ball valve, pressure gauge and new pressure pump were added in 2019. Fig. 1 shows the location of these new elements. The hydraulic system was upgraded with a Bosch Rexroth Pressure & Tank Circuit Hydraulic Power Unit, providing 5.5 L/min at 220 bar and has a 15 L oil tank.

The manipulators are now controlled via a NI Compact DAQ 9132 system. The cDAQ 9132 is a 1.33 GHz Dual-core atom computer with 4 slots for I/O modules. The system runs both Windows 7 Embedded Edition and Labview 2018 for programming and interfacing. The cDAQ 9132 utilises three I/O modules: one NI 9205 i.e. a 32-channel analogue-to-digital

The authors are grateful for the support of the Engineering and Physical Sciences Research Council (EPSRC), grant number EP/R02572X/1, the National Centre for Nuclear Robotics.

(ADC) converter and two NI 9264 i.e. 16-channel digital-to-analogue converters (DAC). The two NI 9264 modules are used to actuate the P02AD1 valves in the two manipulators. The position angle sensors are rotary linear potentiometers.

A dedicated box was recently installed next to the robot in order to hold the controller and associated equipment. A monitor, mouse and keyboard are externally connected such that an operator can control or program the robot from outside the safety cell (to some degree, representing the situation on a nuclear site where the robot will be remote from the operator). Since the present article focuses on the low-level joint control problem, inverse kinematics and the human-machine interface are not described here: see [6] for a recent reference.

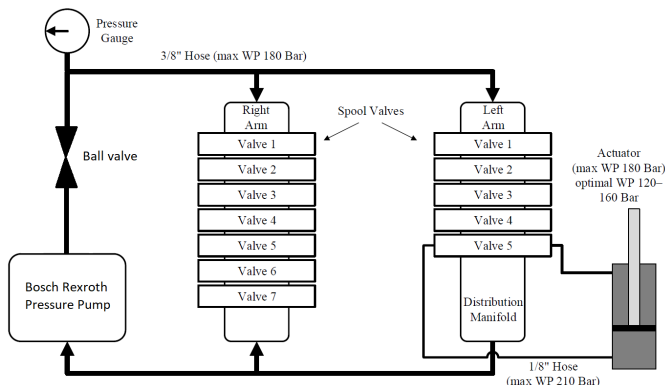


Fig. 1. Schematic diagram of the reconfigured hydraulic system.

III. METHODOLOGY AND PRELIMINARY RESULTS

The new continuous time SDP model for hydraulic manipulators is identified in three stages, as follows:

Step 1. Open-loop step experiments, such as those shown in Fig. 2, suggest that a first order linear differential equation,

$$\dot{\theta}(t) = -a_1\theta(t) + b_1u(t - \tau) \quad (1)$$

provides an approximate representation of individual joints, with a_1 , b_1 and the time-delay τ estimated using the SRIVC algorithm in the CAPTAIN toolbox [8]. Here, $u(t)$ and $\theta(t)$ represent the control input and joint angle respectively, where the former is a scaled signal in the range ± 10 .

Step 2. Further analysis of experimental data using SDP methods, suggests that $a_1 \approx 0$ is time invariant, whilst b_1 is a state dependent parameter. Hence,

$$\dot{\theta}(t) = q\{u(t - \tau)\} \quad (2)$$

where $q\{u(t - \tau)\}$ represents a static nonlinear function of the input. For brevity, further details are omitted from this article, but see [6] for an example of this static nonlinearity, albeit expressed in discrete-time terms. However, these equations and the prior work cited above, all assume time-invariant τ , whereas Fig. 2 illustrates how the actual recorded time-delays can vary from experiment to experiment (in the case of Fig. 2) or during normal operation (more generally), hence introducing a substantial challenge for control design.




Fig. 2. Open-loop experiments for an illustrative manipulator joint using step inputs for a range of magnitudes, with the initially estimated time-delay shown as a solid vertical line (i.e. sample 30). These graphs show the actual recorded time-delays range from 18 samples to 29 samples because of variations in the deadband. The upper and lower subplots show the manipulator being raised and lowered respectively, each trace representing a different experiment.

IV. CONCLUSIONS

This article has described the updated configuration of a robotic platform used for R&D. The SDP approach to system identification for hydraulic manipulators is briefly reviewed, here with the equations adapted into a new continuous-time form. This new formulation is designed to facilitate research into models and control systems that address time-delay and related deadband nonlinearities. In the latter regard, the authors are developing and presently evaluating both conventional and nonsingular terminal sliding mode control systems, in addition to various forms of SDP based control.

REFERENCES

- [1] I. Tsitsimpelis, C. J. Taylor, B. Lennox, and M. J. Joyce, "A review of ground-based robotic systems for the characterization of nuclear environments," *Prog. Nuclear Energy*, vol. 111, pp. 109–124, 2019.
- [2] M. Talha, E. Ghalamzan, C. Takahashi, J. Kuo, W. Ingamells, and R. Stolkin, "Robotic decommissioning of legacy nuclear plant," in *IEEE International Symposium on Safety, Security, and Rescue Robotics*, 2006.
- [3] N. Marturi, A. Rastegarpanah, C. Takahashi, M. Adjigble, R. Stolkin, S. Zurek, M. Kopicki, M. Talha, J. A. Kuo, and Y. Bekiroglu, "Towards advanced robotic manipulation for nuclear decommissioning," in *IEEE Robotics & Automation for Humanitarian App.*, 2016.
- [4] A. Rastegarpanah, N. Marturi, and R. Stolkin, "Autonomous vision-guided bi-manual grasping and manipulation," in *IEEE Workshop on Advanced Robotics and its Social Impacts*, 2017.
- [5] C. J. Taylor and D. Robertson, "State-dependent control of a hydraulically-actuated nuclear decommissioning robot," *Control Engineering Practice*, vol. 21, no. 12, pp. 1716–1725, 2013.
- [6] M. Bandala, C. West, S. Monk, A. Montazeri, and C. J. Taylor, "Vision-based assisted tele-operation of a dual-arm hydraulically actuated robot for pipe cutting and grasping in nuclear environments," *Robotics*, 2019.
- [7] H. Deng, J. Luo, X. Duan, and G. Zhong, "Adaptive inverse control for gripper rotating system in heavy-duty manipulators with unknown deadzones," *IEEE Transactions on Industrial Electronics*, 2017.
- [8] C. J. Taylor, P. C. Young, W. Tych, and E. D. Wilson, "New developments in the CAPTAIN Toolbox for Matlab with case study examples," in *18th IFAC Symposium on System Identification (SYSID)*, Stockholm, Sweden, July 2018.

A Recurrent Encoder-Decoder Network Architecture for Task Recognition and Motion Prediction in Human-Robot Collaboration based on Skeletal Data

1st Dianhao Zhang

Centre For Intelligent Autonomous
Manufacturing Systems
Queen's University Belfast
Belfast, UK
dzhang07@qub.ac.uk

2nd Ngo Anh Vien

The Institute of Electronics, Communications
and Information Technology
Queen's University Belfast
Belfast, UK
v.ngo@qub.ac.uk

3rd Seán McLoone

Centre For Intelligent Autonomous
Manufacturing Systems
Queen's University Belfast
Belfast, UK
s.mcloone@qub.ac.uk

<https://doi.org/10.31256/Qa6Qg1Q>

Abstract—To achieve more accurate early prediction of human motion and enable robots to respond in a safe manner, a human-robot collaboration (HRC) architecture is proposed to predict future human activities and their trajectories using skeletal data collected from a Kinect. The architecture is the combination of two models. The first model is designed to predict motion trajectories and is based on a recurrent Encoder-Decoder long-short-term-memory (LSTM) network that takes a historical trajectory as input and predicts a future trajectory as its output. The second model predicts the next activity to be performed using a combined LSTM and conditional random field network (LSTM-CRF). Preliminary results are presented showing the efficacy of the approach with the LSTM-CRF able to achieve high-quality human activity classification, and the encoder-decoder LSTM able to accurately predict the coordinates of future human motion trajectories.

Index Terms—human-robot collaboration, skeletal data, sequence-to-sequence, LSTM

I. INTRODUCTION

Efficiency and safety are among the most important considerations in manufacturing automation. With collaboration between humans, aided by their ability to anticipate each other's actions, high-level tasks with multiple sub-activities can easily be handled to satisfy these criteria. In contrast, in human-robot collaboration (HRC) efficiency and productivity are frequently negatively impacted by safety related conservative robot operation, as well the computational overhead of dynamically computing actions sequences. It is therefore desirable to develop robots that can anticipate human actions in order to facilitate safe and effective collaboration. To achieve this target, an integrated HRC architecture is proposed consisting of real-time human dynamic motion tracking, human motion recognition and human trajectory prediction modules. The overall architecture of the proposed system is summarised in Fig 1. The input to the system is a sequence of subtasks expressed as trajectories of human skeletal joint coordinates, with each subtask corresponding to one action. These are processed by a long-short-term-memory Encoder-Decoder neural network (LSTM-ED) to give a prediction

of the human joint coordinates for the next subtask. This information on human pose is then processed by a LSTM Conditional Random Fields (CRF) model to generate the label of the future subtask. The predicted human motion trajectory and subtask label are then passed to the robot motion planning and control algorithms. These algorithms are currently under development. The contribution of this paper is to introduce the architecture for human motion and activity prediction for HRC, and to present some preliminary results demonstrating its efficacy.

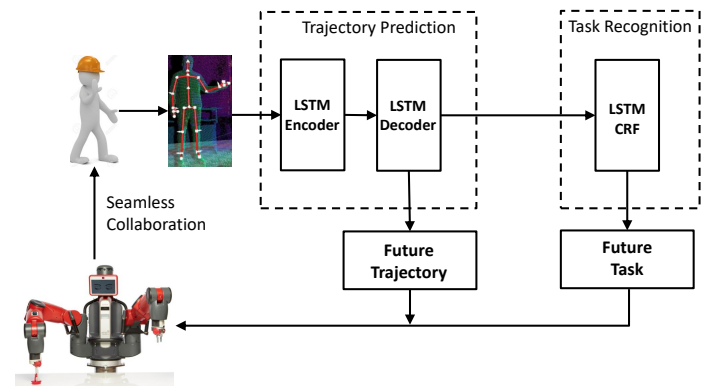


Fig. 1: Human robot collaboration architecture.

II. NETWORK TYPOLOGIES

A. LSTM-ED Network for motion prediction

Long Short-Term Memory (LSTM) is a recurrent neural network (RNN) architecture that was designed by Hochreiter and Schmidhuber [2] to address the vanishing and exploding gradient problems of conventional RNNs, and are ideally suited to processing sequential (time-series) data. In our work we employ a LSTM-ED architecture to implement a motion prediction model. The architecture, which has previously been successfully applied to video prediction and intelligent translation applications [4], [6], is a many-to-many LSTM implementation consisting of a multi-layer encoder and a multi-layer

decoder (Fig 2). The encoder computes a representation s for each input sequence that represents a past motion sequence. Based on that input representation, the decoder generates an output sequence that represents a sequence of future motion. The input is a sequence of coordinates with 50 values ($25 \text{ joints} \times 2\text{D coordinates}$) from a set of 90 frames (3 seconds) of human skeleton reference points, as captured by a Kinect. The attention vector is the sum of hidden states of the encoder, weighted by attention scores. The input to the decoder is the concatenation of the previous hidden state and the attention vector. The first prediction is used as the input to the next LSTM cell.

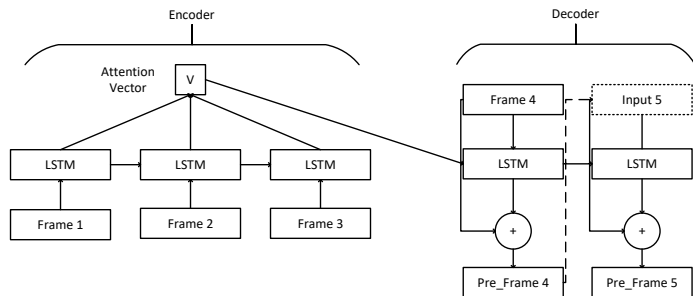


Fig. 2: LSTM sequence-to-sequence model architecture

B. LSTM-CRF for activity classification

Conditional random fields (CRF) are discriminative models for sequence labeling, which have been shown to be a powerful model for sequence tagging problems [3]. The motivation for introducing this model in combination with an LSTM, as first suggested by [1], is that it enables both previous and future context to be considered, and is therefore more suited to representing the diversity and time-varying nature of human activity. An LSTM-CRF network is obtained by employing the hidden states of the an LSTM as the inputs to a CRF layer, where the role of the CRF is to learn a mapping from the hidden state values to the subtask labels [5].

III. RESULTS

We evaluate the performance of the proposed LSTM-ED and LSTM-CRF models for a basic activity involving the use of a screwdriver. The activity is partitioned into three sub-activities; sitting-down, bending over to pick up the screwdriver, and using the screwdriver to tighten a screw (simulated action). As the frame rate of the Kinect is 30 frames per second, one sample of the full screwdriver task consists of 180 frames (2 seconds per sub-activity). A dataset consisting of 120 repetitions of this task was recorded and used as training and test data for the models. Twenty percent of the data was retained as test data. Models were implemented in Python and trained using PyTorch.

First, we evaluate the LSTM-CRF by reporting its activity classification accuracy. Second, we evaluate the LSTM-ED by comparing the difference between the predicted future coordinates and the ground-truth in terms of the mean square error. In addition, we plot the predicted human pose for 90

future frames based on the previous 90 frames to check if the behaviour corresponding to the sub-activity can be predicted accurately. The results are shown in Fig 3.

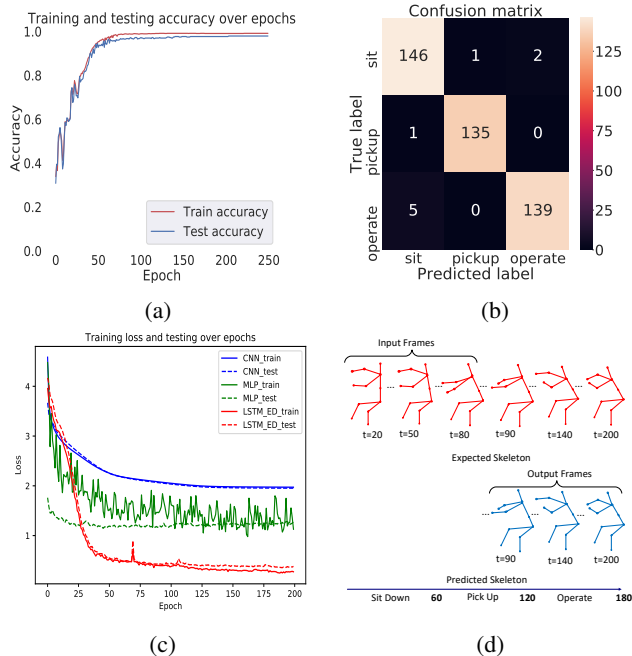


Fig. 3: Experimental results: (a) the LSTM-CRF model learning curves; (b) The test data confusion matrix for sub-activity recognition in the screwdriver usage task; (c) MSE comparison between MLP, CNN and LSTM-ED models for motion prediction; (d) Selected expected and predicted skeletal data frames for one instance of the screwdriver task using the LSTM-ED model.

For the screwdriver usage activity the LSTM-CRF model is able to achieve an accuracy of 99.24% and 98.76% on the training and test datasets, respectively, and for motion prediction the LSTM-ED substantially outperforms MLP and CNN based alternatives, which were also trained and tested for this activity.

IV. CONCLUSION

In this paper, we propose an architecture for providing robots with the capacity to anticipate human actions as an enabler for more effective human-robot collaboration. The architecture involves the use of skeleton-based human tracking data and LSTM-CRF and LSTM-ED based recurrent neural network models to perform human activity classification and human motion prediction, respectively. Preliminary results demonstrate the potential of the approach for predicting human activity subtasks ahead of time, offering the possibility of designing robot path planning and motion control algorithms that are more responsive and attuned to human interaction.

In future work, bespoke path planning and motion control algorithms will be developed and integrated with the proposed architecture and the overall system evaluated on more complex human-robot collaboration scenarios.

REFERENCES

- [1] Tuan Do and James Pustejovsky. Fine-grained event learning of human-object interaction with LSTM-CRF. *arXiv preprint arXiv:1710.00262*, 2017.
- [2] Sepp Hochreiter and Jürgen Schmidhuber. Long Short-Term Memory. *Neural Computation*, 9(8):1735–1780, November 1997.
- [3] Zhiheng Huang, Wei Xu, and Kai Yu. Bidirectional LSTM-CRF Models for Sequence Tagging. *arXiv:1508.01991 [cs]*, August 2015.
- [4] Minh-Thang Luong, Hieu Pham, and Christopher D. Manning. Effective Approaches to Attention-based Neural Machine Translation. *arXiv:1508.04025 [cs]*, September 2015.
- [5] Ariadna Quattoni, Michael Collins, and Trevor Darrell. Conditional Random Fields for Object Recognition. In L. K. Saul, Y. Weiss, and L. Bottou, editors, *Advances in Neural Information Processing Systems 17*, pages 1097–1104. MIT Press, 2005.
- [6] Nitish Srivastava, Elman Mansimov, and Ruslan Salakhudinov. Unsupervised learning of video representations using LSTMs. In *International conference on machine learning*, pages 843–852, 2015.

Topological Robot Localization in a Pipe Network

Rob Worley, Sean Anderson

Department of Automatic Control and Systems Engineering, University of Sheffield, Sheffield, UK
rfworley1@sheffield.ac.uk

<https://doi.org/10.31256/Zw1Wq5M>

Abstract—Topological localization is advantageous for robots with limited sensing ability in pipe networks, where localization is made difficult if a robot incorrectly executes an action and arrives at an unknown junction. Novel incorporation of measurement of distance travelled is used in a Hidden Markov Model based localization method, which is shown to improve accuracy.

Index Terms—Robot Localization, Topological Localization, Pipe Inspection Robots.

I. INTRODUCTION

Water pipe infrastructure is in regular need of maintenance, the cost of which may be reduced by precisely locating faults using robots for autonomous, persistent monitoring of a network. A principal challenge for this robotic system is to localize itself and faults in the network. This work is on *topological* localization for a single robot in a network of pipes. While metric information would be required for precise localization of a fault, topological localization to a single discrete pipe or junction would be sufficient for navigation and for localizing a fault to a part of the network. Metric localization is poorly suited to pipes, as parametric methods like Kalman filters poorly describe the multimodal probability distribution of robot position, and non-parametric methods like particle filters require high computational power from the robot with limited power and size.

Early work in robot localization was done in a topological map [1], as was early work on localization in a pipe network [2]. Recent work on topological localization incorporates some geometric information [3], and recent work on localization in pipes also uses both metric and topological information [4].

This work investigates challenges to localization by the possibility of the robot incorrectly executing an action, and presents the incorporation of measurement of distance into the localization method. The rest of the paper will describe the model of uncertainty in robot motion, and describe the novel addition to the typical localization method. Simulation has been used to evaluate the method, and to investigate the effect of uncertainty parameters on the localization accuracy.

II. PROBLEM DEFINITION

The robot moves in a network of pipes shown in Fig. 1. The network has a range of topologies at smaller scales. It is assumed that the topology and approximate geometry of the network are known. At a junction, the robot chooses a direction at random relative to its own unknown orientation.

This work is supported by an EPSRC Doctoral Training Partnership Scholarship. S. Anderson acknowledges the support of EPSRC grant EP/S016813/1 (Pipebots)

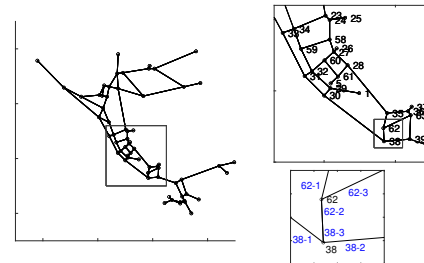


Fig. 1. The example network of pipes used in simulation, consisting of 63 nodes, each connected to 1, 3, or 4 other nodes. Inset is a smaller part of the network with labelled junction indices. The further inset shows the labelled pipes in blue, where each pipe with two labels, one for each node.

This action could be chosen to best inspect the network, however this would not affect localization so is neglected.

The robot state is defined by three components. The first component is the robot's discrete position, which is the junction index. The second component is the robot's discrete direction which is the index of the pipe which it has arrived from, allowing use of information about the robot's choice of action. The third component is the robot's previous position, allowing information about the length of the journey between junctions to be used, as described later. The latter two components are distinct when there are multiple paths between two positions. The robot state is only updated at junctions or at ends of pipes,

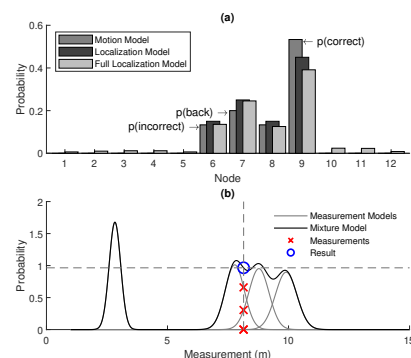


Fig. 2. (a) An example of the discrete probability distribution for robot motion, in this case from node 7 in a 12 node network. Shown is the motion model used to simulate the robot motion, the estimate of this distribution used for localization, and the full localization model considering the probability of missing a node. (b) An example continuous probability distribution over possible measurements of distance between a pair of nodes.




Fig. 3. An example of the use of the localization method in the map in Fig. 1. Each column shows the belief vector over a subset of the discrete states at that time step, where each value represents the probability of being in the corresponding state. The darkness of the colour corresponds to the value of belief in the state. The largest belief value is highlighted in blue if it is correct and red otherwise, where the correct value is shown in orange. (a) The result found without using the measured distance between states, (b) The improved result found when using the distance.

and the robot’s position and orientation are not considered in transitions between these states.

There are four sources of uncertainty in the robot motion: *Incorrect action* execution, *return* to the previous junction, not detecting a junction and *missing* it without updating the state, and normally distributed *noise* in the time taken to travel between junctions. The three discrete components of this model are illustrated as a discrete probability distribution in Fig. 2(a). As the state transition model is difficult to compute exactly, for a given network a Monte Carlo method is used to approximate the transition probability between each state.

The robot makes two observations: the number of exits from a junction, and the distance moved since its last state update. For a given state transition there are a number of routes and corresponding distances. The probability distribution over possible noisy distance measurements is given by a sum of Gaussian distributions, shown in Fig. 2(b). Odometry used to observe this distance could be done using wheel encoders, vision, or simply using the control input and time taken.

III. METHODS

The discrete probability distribution, or belief, over the possible robot states is desired. The belief is a vector summing to one where each value represents the probability of being in the corresponding state. The forward algorithm is used to compute the belief as a Hidden Markov Model (HMM). Using the given state definition, this is equivalent to a second order HMM. The localization model parameters are set to be somewhat incorrect estimates of the values in the motion model described previously, so that the robot does not have exact knowledge of the true motion model. The typical form in Equation 1 computes the updated belief \mathbf{b}' over states s' , based on the belief \mathbf{b} over states s , the observation o , action a , transition and observation models T and O , and a new term for incorporating measured distance m , M .

$$\mathbf{b}'(s') = M(m|s')O(o|s')T(s'|s, a)\mathbf{b}(s) \quad (1)$$

TABLE I

EFFECT OF EACH PARAMETER ON THE LOCALIZATION ERROR METRICS.

Error Metric	incorrect action		return probability		miss probability		noise magnitude	
	ρ^a	g^b	ρ	g	ρ	g	ρ	g
Total Error	-0.41	0.00	-0.20	-0.04	0.98	0.83	0.99	1.96
Mislocalization	-0.46	0.00	-0.61	-0.05	0.98	0.37	0.99	0.66
Relocalization Time	-0.01	0.00	0.69	1.50	0.95	7.47	0.96	13.4

^aThe correlation coefficient ρ between the metric and the parameter.

^bThe linear fit gradient g is the magnitude of the effect of the parameter.

IV. RESULTS

An example of the localization performance is shown in Fig 3. This illustrates the improvement found when incorporating measured distance between junctions. The robot is simulated moving 1000 times between junctions in the network shown in Fig. 1 using the robot definitions given previously, and the state is estimated after each move. The *total* error is measured as the proportion of steps at which the estimation is incorrect. The effect of the four parameters is investigated by performing the simulation for different values of each, giving 256 sets of measurements in total. Over the 16 parameter sets representing lower uncertainty, the median total error without use of measured distance is 0.60 (with an interquartile range of 0.17). This is reduced to 0.18 (with an interquartile range of 0.11) with the use of measured distance.

The *total* error can be decomposed into two parts: the proportion of steps where an initial *mislocalization* occurs, and the mean number of steps before successfully *relocalizing*. Table I shows two measures of the effect of each parameter on the result for these metrics: the correlation coefficient and the gradient of a linear fit. The probability of missing a junction and measurement noise have a strong effect on all metrics, and the probability of incorrectly returning to the previous node has an effect on the relocalization. The probability of correctly executing an action does not affect the accuracy. These results give a measure of the hardware requirements for localization.

V. CONCLUSION

Simulations show that a Hidden Markov Model based method is able to effectively localize a robot in a discrete network where there is a possibility of the robot missing nodes in the network, using noisy measurements of distance travelled between nodes. Variation in the measurement noise and the probability of missing a node is shown to have a large effect on the localization effectiveness.

REFERENCES

- [1] D. Kortenkamp and T. Weymouth, “Topological mapping for mobile robots using a combination of sonar and vision sensing,” *Proceedings of the National Conference on Artificial Intelligence*, vol. 2, pp. 979–984, 1994.
- [2] J. Hertzberg and F. Kirchner, “Landmark-based autonomous navigation in sewerage pipes,” *Proceedings of the 1st Euromicro Workshop on Advanced Mobile Robots, EUROBOT 1996*, pp. 68–73, 1996.
- [3] C. Gomez, A. C. Hernandez, L. Moreno, and R. Barber, “Qualitative Geometrical Uncertainty in a Topological Robot Localization System,” *Proceedings - 2018 International Conference on Control, Artificial Intelligence, Robotics and Optimization, ICCAIRO 2018*, pp. 183–188, 2018.
- [4] D. Alejo, F. Caballero, and L. Merino, “A robust localization system for inspection robots in sewer networks,” *Sensors (Switzerland)*, vol. 19, no. 22, pp. 1–28, 2019.

The role of ‘urban living labs’ in ‘real-world testing’ robotics and autonomous systems.

Rachel Macrorie
Urban Institute
University of Sheffield
Sheffield, U.K.
r.m.macrorie@sheffield.ac.uk

Mateja Kovacic
Nissan Institute of Japanese Studies
University of Oxford
Oxford, U.K.
mateja.kovacic@nissan.ox.ac.uk

Andy Lockhart
Urban Institute
University of Sheffield
Sheffield, U.K.
a.m.lockhart@sheffield.ac.uk

Simon Marvin
Urban Institute
University of Sheffield
Sheffield, U.K.
s.marvin@sheffield.ac.uk

Aidan While
Urban Institute
University of Sheffield
Sheffield, U.K.
a.h.while@sheffield.ac.uk

<https://doi.org/10.31256/Om6Ry1P>

Abstract—As the world rapidly becomes more urban, alongside opportunities to enhance social life, cities face pressing challenges in terms of; congestion, air pollution, food production, provision of care services, infrastructural upgrades, and increased demands for energy, water and mineral resources. Concurrently, technological advances are beginning to extend robotics and autonomous systems technologies (RAS) outside of controlled laboratory environments. In response to these trends, new collaborations are emerging in the form of RAS urban living labs (RAS-ULL) – sites devised to design, test and learn from social and technical RAS innovation in real time. In the U.K., and internationally, there is growing policy advocacy for governments to support development of these experimental spaces to refine and de-risk RAS applications, but this goal is challenging for urban governance bodies and regulators protecting the public realm. This paper analyses RAS-ULLs from three emblematic national-urban contexts leading in urban robotic experimentation – the U.K., U.S.A. and Japan – adopting a focus on delivery robots, service robots and maintenance robots. Developing a comparative analytical framework, we demonstrate the significant potential to learn from international RAS-ULL exemplars to enable extended field experiments in U.K. cities and build strategic capacity around a proactive policy landscape for responsible urban robotic experimentation.

Keywords— *Urban living labs, experiments, field robotics, regulation, human-robot interactions, social acceptability.*

I. INTRODUCTION

In the context of contemporary urbanism, there is mounting interest amongst researchers, technologists and policymakers in re-shaping city infrastructures, services and aspects of social life through advances in robotics and autonomous systems (RAS) [1–7]. The possibilities for RAS restructuring of the city reflect significant developments in materials engineering, communications, artificial intelligence and machine learning, entwined with data analytics and socio-technical platforms that use robotics to augment and re-bundle service infrastructures [8]. This potential is reflected in national economic strategies, for example Japan’s Robot Revolution Initiative, which seeks ‘to make Japan the world’s most advanced robot showcase and achieve a society in which robots are utilised more than anywhere in the world’ [9], and in proposals for a new generation of utopian city projects – such as the proposed mega-city of Neom in Saudi Arabia [10], or Toyota’s plans for a smaller-scale Woven City in Japan [11]. Alongside these flagship visions and initiatives, there is growing pressure for existing cities to open up public spaces

for new RAS experiments and applications, as emerging technologies reach the point of potential real-world application. Set against existing policy and regulatory frameworks, experimental RAS Urban Living Labs (RAS-ULL) are being advanced as an explicit form of intervention to trial, de-risk and improve new technologies, build public support, and appraise the possibilities, realities and implications of this new phase of urban restructuring.

To date, RAS systems have been primarily developed in tightly regulated contexts such as research laboratories [4]. As recognised by the U.K. Government Office for Science White Paper ‘Technology and Innovation Futures’, there is a need to support RAS, and particularly robotic, applications outside of these controlled environments and ‘establishing further ‘test beds’ to experiment with emerging technologies in carefully supervised real-world systems, like cities’ [13]. Creating living laboratories ‘provide[s] a sharp focus to aim developments from basic RAS scientific research into first prototype demonstrators’ [12] because urban test beds are ‘more open and complex, less predictable, and... where [human-machinic interactions] are less controlled’ [3]. However, creating these spaces and conditions raises critical challenges for urban decision-makers [3, 12]. RAS infrastructures are expensive, and there are technical, trust, safety and ethical challenges in bringing emerging technologies into complex and dynamic urban spaces alongside humans. Decisions will need to be made about the granting of licences to R&D organisations and selective changes in regulation. New collaborations between universities, private companies and public agencies will be needed to undertake and learn from experiments. The wider public will need to be actively involved in understanding the purposes and potential of testing in their local communities, as well as how experimental findings will be used [3].

Urban living labs (ULL) are sites devised to design, test and learn about the wider issues involved in the application of new technologies in real world conditions [14]. While the notion of ULL is broad, at its core is the idea that urban sites can provide a learning arena within which the co-creation of innovation can be pursued between research organisations, public institutions, private sector and community actors [15]. ULL have 3 key characteristics:

1. *Experimentation*: The testing of new technologies, solutions and policies in real world conditions, often in highly visible ways.
2. *Participation and user involvement*: Co-designing, collaboration and engagement with many stakeholders is often central to the experimental approach.
3. *Evaluation of actions and impact*: Systemic processes of evaluation underpins the ability of ULL to facilitate formalised learning and upscaling of applications.

For leading practitioners and advocates, ULL are seen not only as a means through which to gain experience, demonstrate and test technologies, but also as a step towards developing responses that have the potential to be scaled up across different domains, in order to support more systemic change. Notwithstanding the current Japanese service-industry context, real-world experiments in autonomous vehicles (AVs) and unmanned aerial vehicles (UAVs) underway in a number of U.S. states, and trials of delivery ‘bots’ in the U.K. city of Milton Keynes, there remains relatively little analysis of what types of ULL might be needed for different aspects of RAS applications. Most ULL have to date been focussed on innovations in digital technologies, smart cities and urban infrastructure. Consequently, there is an urgent need to consider how such robotic-ULLs might address key deficits and challenges in the UK to accelerate the application of RAS, and review how international experience might provide lessons for effective urban experimentation and application domestically [3].

II. METHODOLOGY

To develop the evidence base that can support RAS-ULLs as a means for ‘responsible urban innovation’ [4], a systematic and comparative analysis of their rationales and interests, operations, outcomes and challenges is required. To contribute to this agenda, this article adopts an internationally comparative analytical approach to examine; (i) the design and enabling conditions for RAS-ULLs, (ii) the processes through which RAS-ULLs are implemented, and (iii) the effectiveness with which, and learnings from how, these interventions reshape and augment city infrastructures and services, societal practices, and urban governance. We develop a robust framework for analysing these three dimensions comparatively across three global contexts that are leading in robotic urban experimentation – Japan, U.S.A. and U.K. The research focuses on ULL that are experimenting with (i) service robots, (ii) delivery robots, and (iii) maintenance robots. We have selected these RAS innovations on the basis that they; are already being trialled in semi-public and public realms, seek to reshape and augment the delivery of urban services and operation of city infrastructures, challenge existing urban policy frameworks and regulation, and will involve interaction with citizens when deployed in cities. The case studies were researched through a combination of documentary review, and approximately 50 one-hour long semi-structured interviews conducted with policymakers, robotics firms, and professionals working with and/or employing robots in these locations. These data were subsequently coded and thematically analysed.

III. ANTICIPATED OUTCOMES & IMPLICATIONS

Whilst this research is ongoing, we anticipate that this project will highlight how robots are being materialised in specific ‘early mover’ cities, and that their processes and effects will be uneven. We seek to appraise the extent to which these RAS innovations become ‘embedded’ into, shaped by, and themselves shape infrastructures and urban services, social relations and policy arrangements [16] in order to inform future research priorities. To date, these initiatives are limited in material extent, with a focus on visions more than application, and discrete trials rather than holistic urban robotic restructuring [7]. This research will highlight; (i) how different social, technical and political contexts create conditions for, limit, and lead to contestations around urban robotic experimentation, (ii) the necessary coevolution of spatial planning, urban regulation, urban design and human-robotic interaction in the future ‘infrastructuralisation’ of robotically augmented cities, and (iii) the need to link national innovation priorities for future cities and industrial strategy to pressing urban issues, to responsibly create a social context for RAS applications in contemporary cities.

REFERENCES

- [1] V. Del Casino Jr., “Social geographies II: robots,” *Progress in Human Geography*, vol 40(6), pp.846-855, 2016.
- [2] R. Macrorie, S. Marvin and A. While, “Robotics and automation in the city: a research agenda”, *Urban Geography*, DOI:10.1080/02723638.2019.1698868, 2019.
- [3] S. Marvin, A. While, M. Kovacic, A. Lockhart, R. Macrorie, “Urban robotics and automation: Critical challenges, international experiments and transferable lessons for the U.K.”, UK-RAS Network White Paper, 2018.
- [4] M. Nagenborg, “Urban robotics and responsible urban innovation”, *Ethics and Information Technology*, 2018, DOI:10.1007/s10676-018-9446-8.
- [5] L. Royakkers and R. van Est *Just Ordinary Robots: Automation from love to war*, London: CRC Press, 2015.
- [6] I. Tiddi, E. Bastianelli, E. Daga., M. d’Aquino, E. Motta “Robot-city interaction: mapping the research landscape – a survey of the interactions between robots and modern cities,” *International Journal of Social Robotics*, DOI:10.1007/s12369-019-00534-x
- [7] A. While, S. Marvin and M. Kovacic, “Urban robotic experimentation: San Francisco, Tokyo and Dubai” (in press).
- [8] M.R. Frank, L. Sun, M. Cebrian, H. Youn, I. Rahwan, “Small cities face greater impact from automation,” *Journal of the Royal Society Interface* 15: 20170946. DOI:10.1098/rsif.2017.0946, 2018.
- [9] Robot Revolution Realization Council, “The Prime Minister in Action”, retrieved from http://japan-kantei.go.jp/97_abe/actions/201501/23article3.html, January 2015.
- [10] O. Hassan, “Artificial intelligence: Neom and Saudi Arabia’s economic diversification from oil and gas,” *The Political Quarterly*. DOI:10.1111/1467-923X.12794, 2020.
- [11] J. McCurry, “Toyota unveils plan to build ‘city of the future’ near Mount Fuji”, *The Guardian*, 7 January 2020.
- [12] Lloyd’s Register Foundation, “Foresight review of robotics and autonomous systems: Serving a safer world”, Report Series No. 2016.1, October 2016.
- [13] Government Office for Science (2017) p.10 “Technology & Innovation Futures 2017”, London: Government Office for Science, p.10.
- [14] S. Marvin, H. Bulkeley, L. Mai, K. McCormick, Y. Voytenko Palgan, (Eds) *Urban living labs: Experimenting with city futures*, New York: Routledge, 2018.
- [15] C. Liedtke, M.J. Welfens, H. Rohn, and J. Nordmann, “Living lab: user-driven innovation for sustainability,” *International Journal of Sustainability in Higher Education* 13(2), pp.102-118, 2012.
- [16] S.L. Star, “The ethnography of infrastructure” *American Behavioral Scientist*, 43(3), pp.377-391, 1999.

RICA: Robocentric Indoor Crowd Analysis Dataset

Viktor Schmuck and Oya Celiktutan

Centre for Robotics Research, Department of Engineering

King's College London, London, United Kingdom

{viktor.schmuck; oya.celiktutan}@kcl.ac.uk

<https://doi.org/10.31256/Io1Sq2R>

Abstract—In this paper, we introduce an egocentric dataset recorded from a robot's point of view (robocentric), which has been created to serve as a platform for indoor crowd analysis. The dataset features over 100,000 RGB, depth, and wide-angle camera images as well as LIDAR readings, recorded during a social gathering where the robot captured group interactions between participants using its on-board sensors. We evaluated three different human detection algorithms on our dataset to demonstrate the challenges of indoor crowd analysis from a robot's perspective.

Index Terms—Indoor crowd analysis; Multisensory egocentric dataset; Group recognition

I. INTRODUCTION

Crowd analysis can enable robots to navigate in indoor spaces, approach groups or individuals, and through human-robot interaction assist them in their tasks or in achieving their goals. The research conducted during the past decade on crowd analysis and group detection shows promising results as it utilises the concept of F-formations [1] in order to determine interaction spaces. Most approaches have relied on head and/or body posture detection to build models [2], based on top-down or a bird-eye viewpoint images.

As highlighted by Taylor and Riek [3], these techniques do not keep a robotic context in mind, as they often do not consider the unpredictability of human spaces. Moreover, they do not deal with the different types of noise introduced by the robot's sensors and movement [4], nor do they approach the problem from a robot's point-of-view, which makes them less accurate when applied to an egocentric view.

As shown in Fig. 1, to address the aforementioned gaps, we collected a novel Robocentric Indoor Crowd Analysis (RICA) Dataset¹ using Toyota's Human Support Robot (HSR) [5] as a robotic platform. In particular, we recorded a crowded, semi-public indoor event using robot's on-board cameras as well as LIDAR sensor. In comparison to the existing datasets such as the JackRabbit Dataset [6], the RICA dataset was acquired with less high-end sensors, and we annotated it to enable human detection and group recognition. In this paper, we discuss the challenges of crowd analysis from a robot's perspective and compared three benchmark human detection methods on our dataset.

We thank Toyota Motor Europe for providing the Toyota HSR robot as a development platform.

¹The dataset will be made available at <https://sairlab.github.io/rical/>.



Fig. 1. RGB-D (a1, b1) and Wide-angle camera (a2, b2) samples from two different timestamps of the RICA dataset.

TABLE I

SUMMARY OF THE COLLECTED DATA USING ROBOT'S ON-BOARD SENSORS COMPARED TO THE RELEVANT RECORDINGS OF JRDB.

Sensor Type	Num. of Samples		Average Framerate	
	RICA	JRDB	RICA	JRDB
RGB camera	43,060	57,713	10.542	15.116
Depth camera	39,909	57,714	9.771	15.116
Wide-angle camera	17,877	58,313	4.377	15.273
Joint position	63,569	38,476	15.563	10.078
IMU	127,324	74,234	31.172	19.443
LIDAR	50,926	56,844	12.468	14.888

II. ROBOCENTRIC INDOOR CROWD ANALYSIS DATASET

The proposed dataset was recorded during a reception-style semi-public event in an indoor environment with Toyota's Human Support Robot (HSR) [5]. The robot recorded the event with an "ASUS Xtion PRO LIVE" – RGB-D – camera, a wide angle camera (Nippon Chemi-Con NCM13-J-02), and a "Laser measuring range sensor (UST-20LX)" – LIDAR – sensor. The dataset contains over an hour-long recording of 50 people conversing at a departmental party. Attendees were provided written informed consent, and the data collection protocol was approved by the Ethical Committee of King's College London, United Kingdom. Moreover, for privacy-preserving reasons, the face of the attendees was blurred and only distance and image data was collected.

To obtain a diverse dataset, the robot was driven around at different speeds, following varying paths. Using the height and head adjustment of the robot, its cameras were raised to different elevations, and its head was set to record at a variety of tilt and roll angles. Examples of the camera's captured images can be seen in Fig. 1, where the image data was captured at a resolution of 640×480 . We also recorded IMU measurements of the robot and the joint positions of its head while moving, which can be used to find correspondence between image modalities and LIDAR readings (963 samples from -2.098 to 2.098 radians per sample). The number of samples and average rate per modality are given in Table I.



Fig. 2. An annotated image recorded with the RGB camera, showing a person (ID 21 – blue bounding box on the right hand side) not belonging to any group, and two individuals (IDs 19-20 – red bounding boxes in the middle) belonging to group ID 57 (green bounding box in the middle), where the group formation of group ID 57 is annotated as *face-to-face*.

We labelled the dataset by using a modified version of the Actanno annotation tool [7]. All RGB images of the dataset have been labelled at a group-level, plus identifying the group formations (i.e. L-arrangement, face-to-face, side-by-side, semi-circular, and rectangular) [8]. In addition, person-level labelling, marking people and assigning them to the identified groups, has been done for a total of 8,148 RGB images. These person-level annotations show that in each frame there are 1 to 8 people with an average of 3.92 individuals per frame. The annotations of the remaining modalities can be automatically derived from the labelled bounding boxes based on the timestamps and the joint positions. A sample annotated image can be seen in Fig. 2.

A. Challenges

The RICA dataset was collected without providing participants a script, therefore capturing the natural behaviour of attendees when a mobile support robot was navigating the event floor. The robot was driven at different speeds, on randomised paths, while its cameras were raised to different elevations and its head was held in different angles as it observed the interaction groups. Our manual inspection of the data shows that this resulted in high variation in the camera-to-subject distance (0.1-25m), and participants were often occluded by static objects or each other. It was not ensured that all participants of a single group were in the field of view of the robot and the observation length of each group was varied. The height variation introduces colour changes in the observations of the RGB-D and Wide-angle camera’s images. Due to these factors, the recognition of individuals, groups, and group types from the robot’s viewpoint is a challenging task, and datasets dedicated to robocentric settings are crucial to advancing the state-of-the-art.

III. EVALUATION

We define a series of tests to evaluate the performance of state of the art human detection algorithms on the collected dataset. In particular, we test three methods on the RICA dataset, without fine-tuning: (1) Histogram of Oriented Gradients (HOG) [9] combined with non-maxima suppression

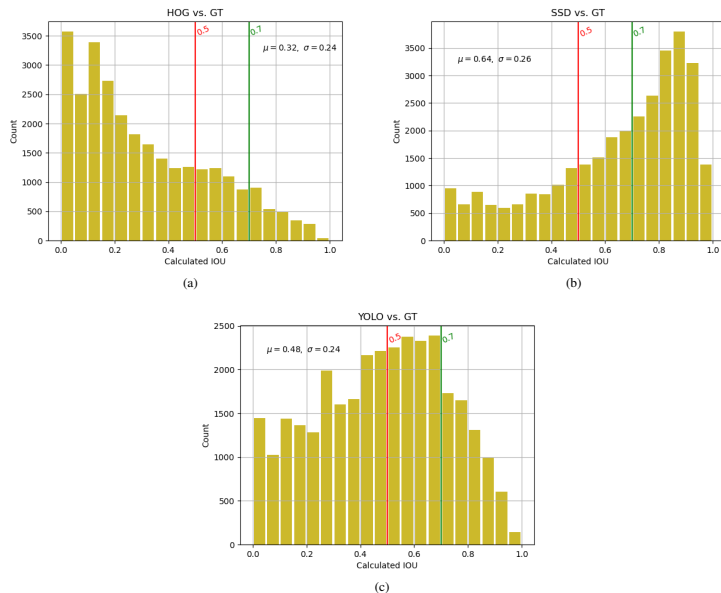


Fig. 3. Histograms of IOU values for between GT and (a) HOG; (b) SSD; and (c) YOLO. The red vertical lines show the minimum IOU and overlap scores to consider a bounding box as a True Positive detection. Green vertical lines indicate the IOU and overlap scores above which the detection is considered as successful.

(NMS); (2) MobileNet-SSD (SSD) [10] – trained on MS-COCO [11], and then fine-tuned on VOC0712 [12] – with centroid tracking, and (3) YOLO [13] – trained on MS-COCO [11]. After retrieving the bounding boxes with each of the human detection methods from the person-level annotated RGB images of the RICA dataset, we computed their intersection over union (IOU) values against ground truth (GT).

Even though there is minimum a single person in each frame, the HOG+NMS detector failed to detect any humans in over 11% of the images, whereas the MN-SSD and YOLO exhibited a similar, better performance – not detecting any humans in 0.7% of all frames. The results of the IOU comparisons are given in Fig. 3. The best mean IOU score ($\mu = 0.64$) was obtained with the SSD detector.

IV. CONCLUSION AND FUTURE WORK

In this paper, we introduced a novel robocentric dataset for indoor crowd analysis, called RICA. Our preliminary analysis shows that the state-of-the-art human detectors fall short and sometimes are unable to detect any humans in the scene due to a list of challenges as summarised in Section II-A. As future work, we will investigate how we can improve human detection and tracking, e.g. by employing occlusion handling techniques in tracking [14]. Moreover, we aim to design an unsupervised approach to group detection in indoor crowded scenes – by adding modalities other than RGB image inputs, and utilising F-formations – based on the RICA dataset.

REFERENCES

- [1] A. Kendon, *Conducting interaction: Patterns of behaviour in focused encounters*. Cambridge University Press, 1990.

- [2] C. Raman and H. Hung, "Towards automatic estimation of conversation floors within F-formations," *arXiv:1907.10384 [cs]*, Jul. 2019.
- [3] A. Taylor and L. D. Riek, "Robot Perception of Human Groups in the Real World: State of the Art," in *2016 AAAI Fall Symposium Series*, Sep. 2016.
- [4] A. Tapus, A. Bandera, R. Vazquez-Martin, and L. V. Calderita, "Perceiving the person and their interactions with the others for social robotics – A review," *Pattern Recognition Letters*, Cooperative and Social Robots: Understanding Human Activities and Intentions, vol. 118, pp. 3–13, Feb. 2019.
- [5] T. Yamamoto, K. Terada, A. Ochiai, F. Saito, Y. Asahara, and K. Murase, "Development of Human Support Robot as the research platform of a domestic mobile manipulator," *ROBOMECH Journal*, vol. 6, p. 4, Apr. 2019.
- [6] R. Martín-Martín, H. Rezatofighi, A. Shenoi, M. Patel, J. Gwak, N. Dass, A. Federman, P. Goebel, and S. Savarese, *JRDB: A Dataset and Benchmark for Visual Perception for Navigation in Human Environments*. 2019.
- [7] C. Wolf, E. Lombardi, J. Mille, O. Celiktutan, M. Jiu, E. Dogan, G. Eren, M. Baccouche, E. Dellandréa, C.-E. Bichot, C. Garcia, and B. Sankur, "Evaluation of video activity localizations integrating quality and quantity measurements," *Computer Vision and Image Understanding*, vol. 127, pp. 14–30, 2014.
- [8] P. Marshall, Y. Rogers, and N. Pantidi, "Using F-formations to analyse spatial patterns of interaction in physical environments," in *Proceedings of the ACM Conference on Computer Supported Cooperative Work, CSCW*, 2011, pp. 445–454.
- [9] N. Dalal and B. Triggs, "Histograms of oriented gradients for human detection," in *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, Jun. 2005, 886–893 vol. 1.
- [10] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and A. C. Berg, "Ssd: Single shot multibox detector," in *Computer Vision – ECCV 2016*, B. Leibe, J. Matas, N. Sebe, and M. Welling, Eds., Cham: Springer International Publishing, 2016, pp. 21–37.
- [11] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft COCO: Common Objects in Context," en, in *Computer Vision – ECCV 2014*, D. Fleet, T. Pajdla, B. Schiele, and T. Tuytelaars, Eds., ser. Lecture Notes in Computer Science, Cham: Springer International Publishing, 2014, pp. 740–755.
- [12] R. Mottaghi, X. Chen, X. Liu, N.-G. Cho, S.-W. Lee, S. Fidler, R. Urtasun, and A. Yuille, "The role of context for object detection and semantic segmentation in the wild," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014.
- [13] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You Only Look Once: Unified, Real-Time Object Detection," *arXiv:1506.02640 [cs]*, May 2016, arXiv: 1506.02640.
- [14] F. Camara, N. Bellotto, S. Cosar, D. Nathanael, M. Althoff, J. Wu, J. Ruenz, A. Dietrich, and C. W. Fox, "Pedestrian Models for Autonomous Driving Part I: Low level models, from sensing to tracking," *arXiv:2002.11669 [cs]*, Feb. 2020, arXiv: 2002.11669. [Online]. Available: <http://arxiv.org/abs/2002.11669> (visited on 04/08/2020).

Robots Producing Their Own Hierarchies with DOSA; The Dependency-Oriented Structure Architect

Benjamin Hawker

*The Department Of Computer Science
The University of Sheffield
Sheffield, United Kingdom
ben.hawker@sheffield.ac.uk*

Roger K Moore

*The Department Of Computer Science
The University of Sheffield
Sheffield, United Kingdom
r.k.moore@sheffield.ac.uk*

<https://doi.org/10.31256/Qt1Os7P>

Abstract—Hierarchical Control Structures are employed in a variety of robots. However, the choice of hierarchy to use affects the ability to control the environment. If a robot could help define the required control hierarchy, this would help avoid the engineer picking an invalid hierarchy and secure effective control. This paper proposes a methodology named DOSA (Dependency Oriented Structuring Architect) which allows a robot to define their own control hierarchy from experience. DOSA brings consistency to hierarchical control, using a clear understanding of hierarchies to encapsulate the heuristic approach an engineer uses and automate a process that requires specialised expertise.

Index Terms—Hierarchical, Control, Automation, Learning

I. INTRODUCTION

Hierarchical Structures are used across robotic solutions with variable success [1]–[3]. Hierarchies when built well have competence that outdoes non-hierarchical approaches [4]. It is a challenge to find a definition of a hierarchy that gives the required insight for an engineer to be able to produce a valid control hierarchy. Prescott identifies the difference between a layered system and a hierarchical system [5], but no clearer definition exists. Solutions to robotic problems that employ hierarchical controllers show that a hierarchy is formed from dependencies, where the control of one signal requires the control of another [6]. Dependencies therefore are key to understanding a hierarchical structure, which are something an engineer and perhaps an autonomous agent can deduce. The question then becomes whether these dependencies can be identified. This paper proposes the use of dependencies between input signals as a means of understanding the constraints on possible hierarchies. Then, these constraints can be used to identify a suitable plan of experimentation in the environment that would allow a suitable hierarchy to be deduced. This framework for deriving hierarchies can be used to identify valid hierarchies, improving a control engineer’s derivation of the hierarchy and in turn improving control.

II. DOSA; THE DEPENDENCY-ORIENTED STRUCTURING ARCHITECT

A. How Dependencies Could Identify the Hierarchy

To briefly consider closed loop control, a controller should output to all actuators that it needs control over. It should not be outputting to actuators that do not affect the input signal being controlled. This would be at best inconsequential, or more likely counter-productive. This principle can reduce the possible hierarchies before attempting any combinations. Identifying which actuators affect which inputs is possible either with simple observation or automatically with Input-Output Analysis [7]. If a controller for a specific input is to hierarchically depend on another controller, it should be to control actuators that it is affected by and no other actuators. After all, a controller should not be directly or indirectly outputting to actuators that are inconsequential to the control task of that controller.

B. Stage One: Identifying Possible Hierarchical Constraints

For each input, the set of actuators that must be controlled can be identified. A controller for that input should only hierarchically delegate to a controller whose set of actuators is a subset of its own actuator set. This way, it will not end up outputting to an actuator that it does not require for control. If a robotic arm with motors and sensors at each joint are considered, there would be some clear findings.

- The position of the hand is affected by the wrist joint, the elbow joint and the shoulder joint.
- The position of the wrist is affected by the elbow and shoulder joints, but not the wrist joint
- The position of the elbow is only affected by the shoulder joint.
- Nothing changes the position of the shoulder, so the position of the shoulder need not be considered.

By examining these subsets, invalid hierarchies can be identified, as detailed in Fig. 1. The possible combinations of hierarchically arranged controllers is reduced by understanding the inputs they affect.

	Shoulder Joint	Elbow Joint	Wrist Joint
Elbow Position	Yes	No	No
Wrist Position	Yes	Yes	No
Hand Position	Yes	Yes	Yes

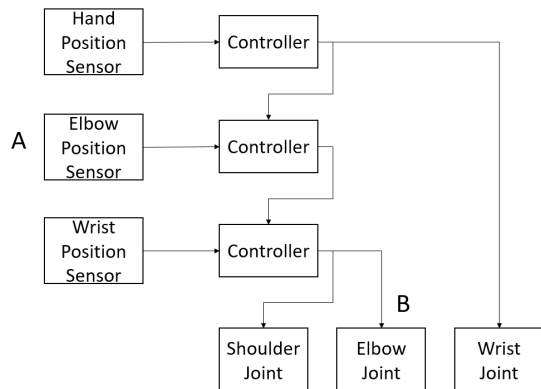


Fig. 1. A table showing which inputs are affected by which actuators above and an example controller for those inputs and actuators below. This controller is invalid when considering the constraints from the table above. The elbow position control signal (A) outputs to the elbow joint actuator (B) indirectly. The table indicates the elbow position is not dependent on the elbow joint, making the hierarchy invalid.

C. Stage Two: Identify Resolution Order

The agent can now developmentally experiment to work out which hierarchy is best, but a suitable order is required to allow this developmental progression. This becomes a simple logic problem, since no controller should be processed before any possible subordinate controller is processed. Therefore, controllers need to be added to the queue such that no possible subordinate is added after it.

- Make a list of every controller and the set of actuators it must output to.
- For each controller's actuator set, identify how many other controller's actuator set is a subset.
- Every controller that has no subsets among the other controller's actuator sets is safe to process, as it has no possible subordinates in the remaining controllers. Add all those controllers to the resolution queue, and remove them from the list.
- If the list is empty, the task is complete. Otherwise, return to the second step.

D. Section Three: Identifying a Suitable Hierarchy Through Progressive Experimentation

Each controller can now be progressively added, attempting the possible different arrangements that would meet the constraints of that controller. Fig. 2 demonstrates the addition of a second controller and the valid arrangements including this controller. Given an optimisation algorithm to tune the controller's parameters, the best arrangement can be selected that minimises error over all controllers. This could be either hierarchically placed on top of one or more existing controllers or entirely aside from them. The process is complete once

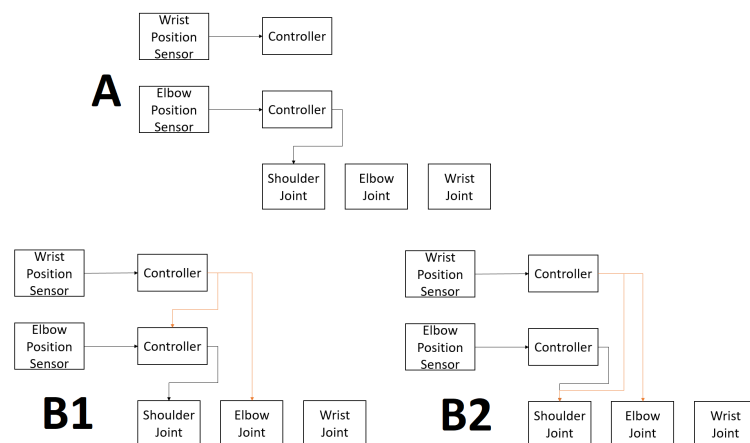


Fig. 2. A controller that is in the process of being progressively derived (A) and two possible options for where the current controller could connect to (B1 and B2). Given the constraints in Fig. 1, reaching the shoulder joint and elbow joint actuator can be achieved in two ways. This is either by controlling an existing controller (B1) or connecting directly to the actuators (B2) and can be decided by which produces the least error.

all controllers have been added in the resolution order. This process can streamline the complex process of trying all hierarchical configurations and make obtaining hierarchies practical and led by unifying methodology.

III. LIMITATIONS, DISCUSSION AND FUTURE WORK

This paper details a methodology named DOSA that can derive hierarchies from experience in the environment. Also, this paper clarifies the nature of a hierarchy and how dependencies are critical to the definition and deployment of a hierarchy. This paper only considers simple scenarios, so more complex scenarios need to be considered to allow a wide application for DOSA. Such examples would be inputs that have complex relationships with actuators (such as multiple actuators being required to change an input value). Input-Output analysis covers understanding many complex input-output relationship [8], [9] but testing that these work in control specific environments is required. Furthermore, this paper doesn't cover the required action if inputs have an identical actuator set. A variety of solutions have been derived by hand, including a hierarchy for the arm used in Living Control Systems 3. The next steps are to automate DOSA and apply it on a simple robotic control system that would typically be solved by a cascade controller. Input-Output analysis shall be used for stage one and evolutionary algorithms to optimise in step three. The long-term objective of this work is to have agents be able to fully derive and optimise their own control hierarchies independent of a human engineer. This would save required time, expertise and guarantee a consistent and effective control hierarchy.

REFERENCES

- [1] A. Sahrin and P. M. H, "Cascade Control With PID-PSO Method on The Stabilizer Unit," in *The 2nd International Conference on Applied Electromagnetic Technology (AEMT) 2018*, 2018.

- [2] R. Brooks, "A robust layered control system for a mobile robot," *IEEE Journal on Robotics and Automation*, vol. 2, no. 1, pp. 14–23, 1986. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=1087032>
- [3] B. Digney and M. Gupta, "A distributed adaptive control system for a quadruped mobile robot," in *IEEE International Conference on Neural Networks*. IEEE, 1993, pp. 144–149. [Online]. Available: <http://ieeexplore.ieee.org/articleDetails.jsp?arnumber=298516>
- [4] Y. Zhong and Y. Luo, "Comparative study of single-loop control and cascade control of third-order object," in *Procedia Engineering*, vol. 15. Elsevier, jan 2011, pp. 783–787.
- [5] T. J. Prescott, P. Redgrave, and K. Gurney, "Layered Control Architectures in Robots and Vertebrates," *Adaptive Behavior*, vol. 7, no. 1, pp. 99–127, jan 1999. [Online]. Available: <http://adb.sagepub.com/content/7/1/99.short>
- [6] R. Kennaway and L. R. Date, "Control of a multi-legged robot based on hierarchical PCT <http://www2.cmp.uea.ac.uk/~jrk/PCT/jp9913p.pdf>," 1999.
- [7] P. J. Schroeder and B. Korel, "Black-box test reduction using input-output analysis," in *Proceedings of the ACM SIGSOFT 2000 International Symposium on Software Testing and Analysis*, vol. 25, no. 5, sep 2000, pp. 173–177. [Online]. Available: <http://portal.acm.org/citation.cfm?doid=347636.349042>
- [8] H. Ying, "Deriving analytical input-output relationship for fuzzy controllers using arbitrary input fuzzy sets and Zadeh fuzzy AND operator," *IEEE Transactions on Fuzzy Systems*, vol. 14, no. 5, pp. 654–662, oct 2006.
- [9] P. Roy, R. Ray, C. Wang, and W. F. Wong, "ASAC: Automatic sensitivity analysis for approximate computing," in *ACM SIGPLAN Notices*, vol. 49, no. 5. New York, New York, USA: Association for Computing Machinery, may 2014, pp. 95–104. [Online]. Available: <http://dl.acm.org/citation.cfm?doid=2597809.2597812>

Development of Tools and Methods for Autonomous Fixed-wing UAV Research

Mahmud Safat Khan
Materials and Engineering Research
Institute
Sheffield Hallam University
Sheffield, U.K.
mahmud.s.khan@student.shu.ac.uk

Lyuba Alboul
Materials and Engineering Research
Institute
Sheffield Hallam University
Sheffield, U.K.
l.alboul@shu.ac.uk

Jonathan Potts
Materials and Engineering Research
Institute
Sheffield Hallam University
Sheffield, U.K.
acesjrp@my.shu.ac.uk

<https://doi.org/10.31256/Mk8Ck9Q>

II. TEST BENCH SETUP

A. Target Airframe: FMS SkyTrainer 182



Fig 1: Target RC Airframe

A 1410mm wingspan RC airframe [2] based on the Cessna 182 has been acquired to be used as the target airframe.

B. X-Plane Modelling and Simulation:




Fig 2: Plane Maker & Airfoil Maker X-Plane modelling

X-Plane 10 is an FAA (Federal Aviation Administration) certified flight simulator which comes with tools such as Plane Maker and Airfoil Maker [3]. Plane Maker enabled the construction of a flight model of the target airframe to be used by the simulator. Airfoil Maker enables generation and/or editing of airfoil data to enhance the realism of the simulation. X-Plane allows communication over the User Datagram Protocol (UDP), which makes it possible to control the aircraft externally (e.g. with Simulink) and capture the live flight data (simulated).

C. XFLR5 Wing Analysis:

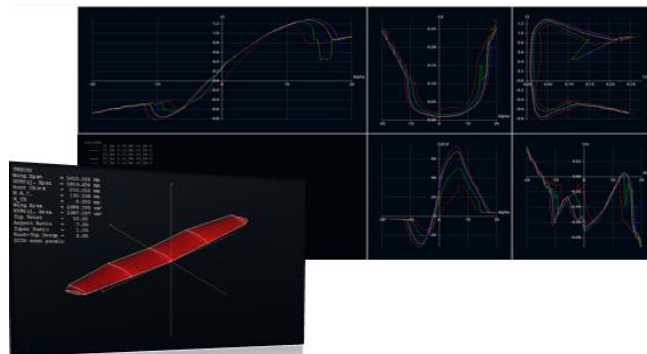


Fig 3: Wing Aerodynamics analysis on XFLR5

Abstract—Development of tools and methods to aid in the research of fixed-wing Unmanned Aerial Vehicles (UAVs) is presented. A Remote-Controlled (RC) model aircraft is used as the target airframe. The flight simulator X-Plane 10 is used for the grey-box modelling and MATLAB/Simulink is used for the white-box modelling. XFLR5 is utilised to obtain crucial aerodynamic data. The process of constructing the flight model on X-Plane 10 via a set of software tools and the development of a 12-state, six-degrees-of-freedom MATLAB/Simulink based simulation are demonstrated. The method presented demonstrates useful applicability.

Keywords—UAV, HITL, GNC, Simulation, Modelling

I. INTRODUCTION

Autonomous systems, like fixed-wing UAVs, employ two major operational modules: high-level Guidance and Navigation algorithms and low-level control laws. Whereas the former module, Guidance and Navigation, is concerned with tasks such as localisation, path-planning, state estimation, etc., the latter module focuses on transcribing these higher level outputs to low-level outputs (for example: deflecting a control surface to alter the flight trajectory). These modules together can be thought of as the Flight Control System (FCS). The field of Guidance Navigation & Control (GNC) is where various disciplines within Robotics and Aerospace Systems converge. The increasing global demand for autonomous flying vehicles with ever increasing applications has highlighted certain shortcomings pertaining to fixed-wing UAVs. While fixed-wing UAVs offer greater operational range, fuel efficiency and payload carrying capabilities, they have yet to benefit from appropriate level of attention from researchers.

There are various challenges a researcher may face with in this complex interdisciplinary area of research. The inability of fixed-wing UAVs to hover or fly at low speeds (relative to rotary wing design) translates to a very unforgiving (high risk of crashing) situation for flight tests. To avoid collisions and crashes of fixed-wing UAVs, the researchers then have to conduct every flight test outdoors, which is not necessary for rotary wing designs. Even with that problem resolved, researchers are left with the challenge of devising autonomous precision landing and take-off from designated airstrips. In the absence of that, whether the UAV carries scientific-research payloads or is being deployed to deliver emergency medical payloads, it remains a high risk endeavour. This calls for greater emphasis to be placed on the need for higher fidelity mathematical modelling and simulation along with means of testing both the software and hardware prior to real flight test [1]. In this paper, the development of a test bench which utilises common software packages, open source codes and hardware, simulation software etc. to aid researchers in the field are presented.

The aerofoil data in X-Plane is typically applied to high Reynolds' number flows and therefore does not accurately model the behaviour of a small aircraft. XFLR5 (based on XFOIL) is a freely available and open source tool that produces reliable aerofoil/wing data for aircrafts in low Reynolds' number regimes (50k to 250k) [4]. Given the specific geometry of the target aerofoil, it generates reliable curves/graphs representing its aerodynamic characteristics (polars). This reliability refers to the consistent polars and plots that match the experimental (wind tunnel) data for the same aerofoil [4]. Therefore, the resulting polars were utilised to generate aerofoil data to be used as inputs to the X-Plane model. The parameters obtained also aid in the construction of the simulation discussed in the following section (section D).

D. MATLAB/Simulink Models and Simulation:




Fig 4: Simulink Simulation of the target airframe

A 12-state, 6 degrees-of-freedom MATLAB/Simulink simulation was developed. The Forces & Moments acting on the UAV and associated Flight Dynamics parameters were modelled with appropriate first order ODE's. S-function blocks [5] were utilised to make custom-blocks for the Forces & Moments and Dynamics. This allows for complete control over the underlying mathematical equations. This model includes the capability to communicate over UDP, a visual output, environmental modelling (e.g. Dryden Gust Model [6]), and interfacing with a flight controller. Beard & McLain (2012) presents much of the general outline of this simulation in the form of skeletal MATLAB and Simulink files. However, the researcher must write appropriate codes based on suitable mathematical models to make these skeletal files functional.

E. HITL Simulation:




Fig 5: Hardware-In-The-Loop (HITL) simulation test

The open-source flight controller Pixhawk and flight control software PX4 has been used alongside QGroundControl (GQC) [7] to establish communication between X-Plane and the controller. By modifying the code of the FCS, it is possible to test the Pulse width modulation (PWM) servo outputs during the Hardware-In-The-Loop (HITL) simulation. It is also possible to develop an FCS in Simulink and deploy that on the Pixhawk through the Pixhawk PSP package [8] for Simulink's Embedded Coder.

III. RESULTS, DISCUSSION & CONCLUSION

The goal was to present a test-bench that addresses many of the problems faced by researchers whilst working with fixed-wing UAVs. The mathematical models were implemented in the form of a Simulink simulation. X-Plane 10's built in modelling tools were utilised to create an X-Plane simulation of the target airframe. One problem with such modelling is the absence of quality aerodynamic data on the aerofoil being used. It was demonstrated that via the use of tools like XFLR5, it is possible to obtain appropriate aerodynamic data and improve the simulation fidelity without the reliance on resource intensive wind tunnel testing or complex Computational Fluid Dynamics modelling.

Finally, a practical application was demonstrated in the form of a HITL simulation. Such simulations allow researchers to test the behaviour and performance of the flight code in a safe context where modelling errors, hardware faults and various other bugs in the overall system can be identified and dealt with without risking damage to the actual UAV and its systems in real flight. Once actual flight tests are conducted and appropriate flight data are recorded, the researcher benefits from important insights from juxtaposition of the flight data against the two simulation's predictions. This allows for fine-tuning of the models and higher simulation fidelity. The overall setup presented here can also be exploited to conduct studies on a range of topics such as Guidance Algorithms, State Estimation, Parameter Identification and Sensor Fusion.

REFERENCES

- [1] J.R. Raol, J. Singh, Flight Mechanics Modelling and Analysis, , CRC Press, Boca Rayton Florida, 2009, chapter 6-9
- [2] FMS Models, FMS SkyTrainer 182, <http://www.fmsmodel.com/fms-1400mm-sky-trainer-182-5ch-at-red-pnp>, retrieved on 28/03/2018
- [3] FAA Certified X-Plane, <https://www.x-plane.com/pro/certified/> retrieved on 28/03/2019
- [4] Drela M. XFOIL: an analysis and design system for low Reynolds number airfoils., University of Notre Dame, 1989
- [5] Simulink S-Function, Mathworks User Guide, <https://uk.mathworks.com/help/simulink/sfg/what-is-an-s-function.html>, retrieved on 28/03/2019
- [6] Dryden Wind Turbulence Model, Mathworks <https://uk.mathworks.com/help/aerobkls/drydenwindturbulencemodel/continuous.html> retrieved on 28/03/2019
- [7] The Pixhawk Project, PX4, QGC and Pixhawk, <http://pixhawk.org/>, retrieved on 28/03/2019
- [8] Pixhawk PSP package, Pixhawk Support From Embedded Coder <https://uk.mathworks.com/hardware-support/px4-autopilots.html> retrieved on 03/03/2020
- [9] BEARD, R., & McLAIN, T., "Small Unmanned Aircraft: Theory and Practice". Princeton, Oxford: Princeton University 2012

A review of manufacturing systems for introducing collaborative robots

RuiDong Ma
Department of ACSE
The University of Sheffield
 Sheffield, UK
 rma17@sheffield.ac.uk

Jingyu Chen
Department of ACSE
The University of Sheffield
 Sheffield, UK
 jchen118@sheffield.ac.uk

John Oyekan
Department of ACSE
The University of Sheffield
 Sheffield, UK
 j.oyekan@sheffield.ac.uk

<https://doi.org/10.31256/Zb5Dy3B>

Abstract—Industry 4.0 highlights a new industrial revolution for the manufacturing system. This work aims to provide a review of different types of manufacturing systems and present motivations of introducing collaborative robots into manufacturing. We start with a discussion about the existing research of human-robot collaboration as well as its perception and control strategies. Then, we give a review of the current applications of swarm robots in manufacturing. Finally, we propose some insights for future directions of human-robot society.

Index Terms—manufacturing system, human-robot collaboration, swarm robotics, cognitive model

I. INTRODUCTION OF MANUFACTURING SYSTEM

The manufacturing system, which is defined as a collection of labour resources and integrated equipment, is utilized to process and assemble the raw production materials [1]. In this section, five types of manufacturing systems are discussed and compared. We also present the features and potential robot usage of them as shown in Table I.

Flow shop is a product-oriented system while scheduling the sequence of order is difficult. Sadik and Urban [2] introduced a case study which optimizes the scheduling problem with Human Robot Collaboration (HRC). Cellular manufacturing which groups similar parts into families and assigns the associated machines located in each cell into groups [3] implements the small scale to produce part of a production with one worker in [4]. Its people-oriented character emphasizes the human operator's versatility and flexibility. However, to improve production efficiency, robot assistance should be added into the system as another step. Flexible manufacturing system (FMS) is defined as a production method which is adaptable for production type and size. Krüger et al. [5] proposed Intelligent Assist Systems for more flexible assembly tasks. Reconfigurable manufacturing system (RMS), combining the flexibility of FMS with the high throughput of a dedicated manufacturing system, is designed for adapting to the rapid changes of the market within the same part family. The project shop aims for large scale products which require multiple components in the layout like aeroplane manufacturing. Bauda et al. [6] proposed 'Air-Cobot robot' for vision inspecting of production quality.

TABLE I
 MANUFACTURING SYSTEM

Manufacturing system	Features	Potential cobot usage
Cellular manufacturing	High product variation Highly skilled labour	Task-based HRC to improve efficiency
Flexible manufacturing	High product variation Highly skilled labour	Intelligent assist system for the variate product
Flow shop	Low product variation Low skilled labour	Solving scheduling problem and manual labour shortage
Reconfigurable manufacturing	Customized flexibility Adaptability	Reconfigurable machine tools
Project shop	Large products Low variation	Air-Cobot for vision inspection

II. COLLABORATIVE ROBOT IN MANUFACTURING SYSTEMS

A. Industrial tasks for human-robot collaboration

The main advantage of human-robot collaboration in the manufacturing system is that robots can assist human operators with sophisticated tasks. In this manner, machines do not replace humans, but they supplement their ability by getting rid of heavy work for workers. Unlike the traditional industrial robots, collaborative robots (cobots) in the manufacturing system can offer more safety and dexterity. Such Robots, for instance, rethink or universal robots, can combine the precision and speed of machines with the proportions and flexibility of human hands. Another feature is that the robot can learn from demonstration due to its simplification in programming for specific tasks.

To allow the robot to better understand the human, several perceptions are utilized to collect external data to the internal representation system. Robot vision combines the camera and software toolkit to enable the robot to obtain visual data from the world and execute responding physical actions [7],[8],[9]. The impedance control is used to measure the force between the manipulator and human, and hence infer the relationship between the force and position [10],[11]. Audition, as sounds or voice, is another common modality which can be used to guide intelligent system or communication [12].

In manufacturing tasks nowadays, the cobot becomes more competitive when compared with the human operator and traditional industrial robots. Many manufacturers are eager to adopt HRC technology to enhance the effectiveness and flexibility of their production. Table II demonstrates some industrial scenarios working with cobot.

As seen in the table, the main tasks where cobots are involved with the industry are manual assembly tasks. The human operator is able to operate variant productions while the work-ability can be restricted by ergonomic factors and hence influence the accuracy and production volume [13]. Traditional industrial robots can handle high repetitive and payload tasks, for instance, the ABB IRB 7600 can handle up to 500kg materials [14]. However, in complex manual assembly tasks, it is too expensive to achieve and dangerous to human operators [15]. The cobot can combine the repeatability from industrial robots and flexibility from workers. Meanwhile, as the safety control strategy of cobot is designed for operating among humans, it can also save work space [16].

TABLE II
SOME STATE-OF-THE-ART IN USING COBOT FOR INDUSTRIAL TASKS

Industrial scenarios	Tasks	Advantages
BMW[17]	Equipping insulation insider door	Replace human worker
Audi[16]	UR3 cobot for adhesive on car roof	Save space
Volkswagen[13]	KuKA cobot for screwing on drive train	Easier to reach locations
ARM[18]	Prepreg for composite layup	Reduce human operator's workload

III. SWARM ROBOTS IN MANUFACTURING SYSTEM

The industrial robots have been successfully deployed in manufacturing for the last decades. They can be represented by the static robot arms which are programmed to execute the heavily manual, complex and hazardous tasks. However, the setup of the layout and the controller for these inflexible machines often cost much time and money when the design of the product changes. Thus, mobile robots like unmanned ground vehicle (UGV) and unmanned aerial vehicle (UAV) with good maneuverability can be appropriately utilized to make a difference. Moreover, as the manufacturing environment is dynamic and uncertain, we cannot expect one single robot to fulfil all the tasks. Therefore, to enhance the efficiency and robustness of the system, the concept of swarm robotics which is inspired by the collective behaviours of social insects can be introduced. The swarm engineering [19] aiming at overcoming the current limitations of swarm robots is also addressed to make the robot team collaboratively solve the real-world challenges in manufacturing.




Fig. 1. Cooperative transportation

Although the swarm robots already make some achievements on the surveillance, mapping and navigation, the applications in the industry mainly focus on the manipulation [20], transportation [21], and assembly. A team of drones transportation scenario is shown in Fig.1. In [22], the authors develop a new tool which is capable of co-localizing holes and fasteners for robust insertion and fastening. In the experiment, a heterogeneous team of four robots with different skills are assigned to align and fasten a panel to a corresponding box. The transportation of materials is achieved in [23] using neural network synthesised by an evolutionary algorithm and [24] using the leader and follower scheme. Another application where swarms have been used is logistics as well as the sorting task in the warehouse. As self-organisation is a well-known behaviour of swarm intelligence, collective behaviour is explored for the autonomous goods classification using ground robots in the real world. In [25], controlled by the neural network, the swarm of agents called ants are evolved to perform the patch sorting and annular sorting for the objects with different shapes in the environment. Currently, aerial robots haven't been massively deployed in the factories due to the problem of stability and battery life except for the project shop manufacturing system. For example, in the shipbuilding or aircraft assembly industry, as the position of the layout is fixed, the material components can be transported into the product by the aerial swarm [26].

IV. CONCLUSION AND FUTURE PLAN

From the previous review of the collaborative robot, it is observed that the cobot can combine the repeatability and flexibility from industrial robot and human operator to enhance the production efficiency. However, in the existing cobot application scenario, most of them are task-based and operated in constrained static environments. Moreover, the material and information flow of the manufacturing system is often a problem [1]. For instance, in a project shop, the work may be interrupted if the material supply is late. The potential solution can be the combination of swarm robots and static robots. For instance, if the static cobot can sense the current workflow and any absence of components based on its cognitive architecture, it will inform the swarm robots of the transportation task. In this manner, the manufacturing system will become more automatic. This paper investigates and analyzes the current and potential applications of collaborative robots in manufacturing systems. To deal with the uncertainty of the market, increasing the automation level of the system by a massive deployment of the robots in factories is no longer the priority. It might be necessary to pay more attention to human-robot and robot-robot interactions so that robots can be easily reconfigured to collaborate better with the human, which makes the manufacturing system move closer towards the standard of Industry 4.0.

ACKNOWLEDGMENT

The authors would like to thank the University of Sheffield for the use of equipment and lab space to conduct this research

REFERENCES

- [1] T. Vamos, "Automation production systems and computer integrated manufacturing. Mikell P. Groover," *Automatica*, vol. 24, no. 4, p. 587, 1988.
- [2] A. R. Sadik and B. Urban, "Flow shop scheduling problem and solution in cooperative robotics—case-study: One cobot in cooperation with one worker," *Future Internet*, vol. 9, no. 3, 2017.
- [3] Y. Yin and K. Yasuda, "Similarity coefficient methods applied to the cell formation problem: A taxonomy and review," *International Journal of Production Economics*, vol. 101, no. 2, pp. 329–352, 2006.
- [4] J. T. C. Tan, F. Duan, Y. Zhang, K. Watanabe, R. Kato, and T. Arai, "Human-robot collaboration in cellular manufacturing: Design and development," *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2009*, pp. 29–34, 2009.
- [5] J. Krüger, R. Bernhardt, and D. Surdilovic, "Intelligent assist systems for flexible assembly," *CIRP Annals - Manufacturing Technology*, vol. 55, no. 1, pp. 29–32, 2006.
- [6] M.-a. Bauda, A. Grenwelge, S. Larnier, M.-a. Bauda, A. Grenwelge, S. Larnier, M.-a. Bauda, A. Grenwelge, and S. Larnier, "3D scanner positioning for aircraft surface inspection To cite this version : 3D scanner positioning for aircraft surface inspection," 2019.
- [7] B. Teke, M. Lanz, J. K. Kämäräinen, and A. Hietanen, "Real-time and Robust Collaborative Robot Motion Control with Microsoft Kinect ® v2," *2018 14th IEEE/ASME International Conference on Mechatronic and Embedded Systems and Applications, MESA 2018*, pp. 1–6, 2018.
- [8] F. Fang, Q. Xu, Y. Cheng, L. Li, Y. Sun, and J.-H. Lim, "Self-Teaching Strategy for Learning to Recognize Novel Objects in Collaborative Robots," pp. 18–23, 2019.
- [9] K. T. Song, Y. H. Chang, and J. H. Chen, "3D vision for object grasp and obstacle avoidance of a collaborative robot," *IEEE/ASME International Conference on Advanced Intelligent Mechatronics, AIM*, vol. 2019-July, pp. 254–258, 2019.
- [10] L. Rozo, S. Calinon, D. G. Caldwell, P. Jiménez, and C. Torras, "Learning Physical Collaborative Robot Behaviors From Human Demonstrations," *IEEE Transactions on Robotics*, vol. 32, no. 3, pp. 513–527, 2016.
- [11] E. C. Townsend, E. A. Mielke, D. Wingate, and M. D. Killpack, "Estimating Human Intent for Physical Human-Robot Co-Manipulation," 2017. [Online]. Available: <http://arxiv.org/abs/1705.10851>
- [12] P. Gustavsson, A. Syberfeldt, R. Brewster, and L. Wang, "Human-robot Collaboration Demonstrator Combining Speech Recognition and Haptic Control," *Procedia CIRP*, vol. 63, pp. 396–401, 2017. [Online]. Available: <http://dx.doi.org/10.1016/j.procir.2017.03.126>
- [13] KUKA, "Many wrenches make light work: KUKA flexFELLOW will provide assistance during drive train pre-assembly," 2016. [Online]. Available: <https://www.kuka.com/en-de/press/news/2016/10/vw-commits-to-human-robot-collaboration-in-wolfsburg>
- [14] ABBRobotics, "IRB7600 Opening a new world of possibilities," 2019. [Online]. Available: <https://new.abb.com/products/robotics/industrial-robots/irb-7600>
- [15] M. Hägele, W. Schaaf, and E. Helms, "Robot assistants at manual workplaces - Effective co-operation and safety aspects," *Proceedings of the 33rd International Symposium on Robotics (ISR)*, p. 6, 2002.
- [16] Audi, "New human-robot cooperation in Audi production processes," 2015. [Online]. Available: <https://www.audiusa.com/newsroom/news/press-releases/2015/02/new-human-robot-cooperation-in-audis-production-processes>
- [17] N. Giles and M. Hatzel, "Innovative human-robot cooperation in BMW Group Production," 2013. [Online]. Available: <https://www.press.bmwgroup.com/global/article/detail/T0209722EN/innovative-human-robot-cooperation-in-bmw-group-production?language=en>
- [18] Advanced Robotics for manufacturing, "ROBOTIC ASSISTANTS FOR COMPOSITE LAYUP," 2020. [Online]. Available: <http://arminstitute.org/projects/robotic-assistants-for-composite-layup/>
- [19] R. T. Anderson, L. Neri, R. T. Anderson, and L. Neri, "Swarm Engineering," *Reliability-Centered Maintenance: Management and Engineering Methods*, no. May, pp. 97–206, 2000.
- [20] D. Mellinger, M. Shomin, N. Michael, and V. Kumar, "Cooperative grasping and transport using multiple quadrotors," *Springer Tracts in Advanced Robotics*, vol. 83 STAR, pp. 545–558, 2012.
- [21] V. Spurny, M. Petrlik, V. Vonasek, and M. Saska, "Cooperative Transport of Large Objects by a Pair of Unmanned Aerial Systems using Sampling-based Motion Planning," *IEEE International Conference on Emerging Technologies and Factory Automation, ETFA*, vol. 2019-Septe, pp. 955–962, 2019.
- [22] M. Dogar, R. A. Knepper, A. Spielberg, C. Choi, H. I. Christensen, and D. Rus, "Multi-scale assembly with robot teams," *International Journal of Robotics Research*, vol. 34, no. 13, pp. 1645–1659, 2015.
- [23] R. Groß and M. Dorigo, "Towards group transport by swarms of robots," *International Journal of Bio-Inspired Computation*, vol. 1, no. 1-2, pp. 1–13, 2009.
- [24] T. Sugar, "Control and coordination of multiple mobile robots in manipulation and material handling tasks," *Experimental Robotics VI*, 2000. [Online]. Available: <http://www.springerlink.com/index/P8U250420L347P12.pdf>
- [25] V. Hartmann, "Evolving Agent Swarms for Clustering and Sorting," pp. 1–8, 2005. [Online]. Available: papers://82ac23f7-2eaf-4339-a5e1-4600c19d7f01/Paper/p2524
- [26] Q. Lindsey, D. Mellinger, and V. Kumar, "Construction of cubic structures with quadrotor teams," *Robotics: Science and Systems*, vol. 7, pp. 177–184, 2012.

Smart Suitcase Implementation Using Fuzzy Logic and Deep Learning

KIM TIEN LY

School of Computer Science

University of Nottingham

UNITED KINGDOM

<https://doi.org/10.31256/Fn7Jw8E>

Abstract—This paper proposes a real-time control application using deep learning and fuzzy controller to create a smart suitcase. This device is a mobile robot that automatically follows its owner with the ability to avoid obstacles on its way. The deep learning model used in this project is based on the efficient technique and state-of-the-art object detection algorithm – Mobilenet-SSD convolutional neural network. The fuzzy logic controller with logical if-then rules forms an effective automatic control strategy. This paper presents a proof of concept illustrating the integration of fuzzy controller with real-time object detection and tracking using deep learning.

Keywords—autonomous robot, mobile robot, deep learning, fuzzy controller, embedded system

I. INTRODUCTION

Machine learning for computer vision, especially convolutional neural networks have become popular for its diverse applications including face or object detection [1] [2], or even signal analysis [3]. Since AlexNet was released and won the ImageNet Challenge in 2012 [4], the general trend has been to make deeper and more complicated networks in order to achieve higher accuracy. However, these advances to improve accuracy are redetrimated in terms of size and speed. In 2017, Google has announced a class of efficient models called MobileNets for mobile and embedded vision applications. MobileNets are based on a streamlined architecture that uses Depthwise Separable Convolutions to build lightweight deep parameters that efficiently choose a tradeoff between latency and accuracy [4].

Fuzzy logic controller [5] allows flexible and optimal operation defining the rules to change vague explanations to obvious definitions. Fuzzy sets can be built from human experience and perception, which helps to deal with real-life situations. The vague boundaries enhance the ability to handle uncertainty in mobile robots, which makes fuzzy logic be widely implemented from daily products [6] to process industry [7] or decision making system [8].

Although deep learning techniques have long been applied in computer vision, robot vision has a specific challenge with real-time operation due to limited onboard resources. Similarly, fuzzy logic controller has been used in various robot applications, but only with basic automotive sensors for robots to compute immediate actions. This paper points out a clear approach to combine deep learning and fuzzy controller in an autonomous robot. The idea is to design a smart suitcase that can follow its user automatically in real-time, which gives the owner freedom to do another task when walking on an airport platform. This model is practically built and tested with the aim to contribute to develop devices for human utilities.

II. METHODS

A. Overview

A carry-on suitcase model is designed with four wheels, two of them are controlled by motors and the others are orientable. A Raspberry Pi 3 B+ is responsible for owner detection, which is a logo, for tracking. However, due to its limitation in real-time image processing, an Intel Movidius Neural Compute Stick is integrated with Raspberry Pi and an ARM STM32F407 microcontroller is used to control the robot to follow its owner and simultaneously avoid obstacles, which are all other objects apart from the specified owner. SPI protocol is implemented for communication between them.

B. Raspberry Pi 3 B+

1) Training deep learning neural network model for object detection

Dataset is created using taken pictures of the specified logo with the help of data augmentation, which involves creating transformed versions of collected images using techniques such as flip, blur adding, rotation, image's characteristic adjustment, etc. Positions of the objects are labelled and saved as XML files in default PASCAL VOC format. Two classes are defined in order to classify the proposed object and background respectively. Cross-validation technique is applied in creating Lightning Memory-Mapped Database (LMDB). The training phase takes place on Google Colaboratory to make use of the computational power of GPU Tesla K40 supported by Google Corporation.

2) Tracking result transfer

SPI protocol is implemented for data transfer from Raspberry Pi to ARM, where Raspberry Pi is the master-side of the connection.

C. ARM

1) *Fuzzy for owner following*: takes the inputs from owner detecting result to compute two outputs for the difference in the two motors' PWM.

Input:

- Position error (ePosition): the difference between the center of the camera and the center of the detected frame. The purpose of this variable is to ensure that the owner is right in front of the suitcase.
- Distance error (eDistance): the difference between the height of the detected frame and the expected height. The purpose of this variable is to ensure a specific distance between the suitcase and its owner. Since there can be a lot of interference from

people passing by, using an ultrasonic sensor for distance measurement was crossed out.

Output: PWM difference value for each motor.

Rules: Fuzzy rules for owner following are defined in TABLE I.

TABLE I. Object following fuzzy rules for two motors. Noted abbreviations: negative big (NB), negative medium (NM), negative small (NS), zero (ZE), positive small (PS), positive medium (PM), positive big (PB), low (LO), medium (ME), high (HI), negative (NE), positive (PO).

Right motor Left motor		ePosition				
		NB	NS	ZE	PS	PB
eDistance	NE	NS	ZE	PS	PM	PB
		PB	PM	PS	ZE	NS
	ZE	NM	NS	ZE	PS	PM
		PM	PS	ZE	NS	NM
	PO	NB	NM	NS	ZE	PS
		PS	ZE	NS	NM	NB

2) Obstacle Avoidance

a) Sensors and algorithm

The implemented algorithm is the bubble rebound algorithm [9]. Accordingly, robots require a ring of ultrasonic sensors covering a certain range returning results of detected objects on its way.




Fig. 1. Four ultrasonic sensors cover an angle of 120° with defined sensitivity bubble.

The robot is expected to move towards the area having the lowest density of obstacles [9].

b) Fuzzy for obstacle avoidance: takes the calculated angle from the bubble rebound algorithm together with current PWM of the two motors to compute two outputs for the difference in the two motors' PWM.

Input:

- Angle: the direction where the suitcase needs to follow to avoid obstacles.
- Current PWM: the current PWM on each motor.

Output: PWM difference value for each motor.

Rules: Fuzzy rules for obstacle avoidance are defined in TABLE II.

TABLE II. Obstacle avoidance fuzzy rules for two motors. Noted abbreviations: negative big (NB), negative medium (NM), negative small (NS), zero (ZE), positive small (PS), positive medium (PM), positive big (PB), low (LO), medium (ME), high (HI), negative (NE), positive (PO).

Right motor Left motor		Angle				
		NB	NS	ZE	PS	PB
Current_PWM	LO	NS	ZE	PS	PM	PB
		PB	PM	PS	ZE	NS
	ME	NM	NS	ZE	PS	PM
		PM	PS	ZE	NS	NM
	HI	NB	NM	NS	ZE	PS
		PS	ZE	NS	NM	NB


III. DISCUSSION

A. Result


The autonomous suitcase is tested with both indoor and outdoor environment. It has proved to be able to move smoothly in real-time owner tracking and obstacle avoidance. The recorded video is available at [this link](#).

The tracking feature is examined in many cases including different distances, different rotations, different background or even when a part of the logo is hidden. The training phase gains 92% accuracy in logo detection. Failure cases are found in low-light environment. The detection speed is around 15-18fps, which allows the Raspberry Pi to transfer data to ARM with the sample rate 100ms.

Results of fuzzy controller for owner following are shown in Fig.2. In Fig.2(a), the suitcase starts to move when it is far from the tag, it moves straight forward as the tag is not too far to the left or right of the suitcase. Fig.2(b) indicates the situation when the owner is far to the left, the right motor moves faster to turn left.



(a)



(b)

Fig. 2. Results of fuzzy controller when the owner is a little far and around the middle (a), when the owner is far and to the left (b).

The suitcase model uses only low-cost devices.

B. Conclusion

The approach in this paper has provided a practical application by combining machine learning and fuzzy controller theory. Things become easier for engineers to make real robots in industry with basic principles and low-cost electrical components.

IV. REFERENCES

- [1] C. Garcia and M. Delakis, "Convolutional face finder: a neural architecture for fast and robust face detection," in *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2004.
- [2] A. Frome, G. Cheung, A. Abdulkader, M. Zennaro, B. Wu, A. Bissacco, H. Adam, H. Neven and L. Vincent, "Large-scale privacy protection in Google Street View," in *IEEE 12th International Conference on Computer Vision*, Kyoto, Japan, 2009.
- [3] U. R. Acharya, S. L. Oh, Y. Hagiwara, J. H. Tan and H. Adeli, "Deep convolutional neural network for the automated detection and diagnosis of seizure using EEG signals," in *Computers in Biology and Medicine*, 2018, pp. 270-278.
- [4] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto and H. Adam, "MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications," April 2017.
- [5] C. C. Lee, "Fuzzy logic in control systems: fuzzy logic controller. I," *IEEE Transactions on Systems, Man, and Cybernetics*, pp. 404-418, 1990.
- [6] M. Akram, S. Habib and I. Javed, "Intuitionistic Fuzzy Logic Control for Washing Machines," *Indian Journal of Science and Technology*, vol. 7(5), no. 654-661, 2014.
- [7] E. K. Juuso, "Fuzzy Control in Process Industry: The Linguistic Equation Approach," *Fuzzy Algorithms for Control*, pp. 243-300.
- [8] J. Zeng, M. An and N. J. Smith, "Application of a fuzzy based decision making methodology to construction project risk assessment," *International Journal of Project Management*, vol. 25, no. 6, pp. 589-600, 2007.
- [9] I. Susnea, V. Minzu and G. Vasiliu, "Simple, Real-Time Obstacle Avoidance Algorithm for Mobile Robots," Vols. ISBN: 978-960-474-144, January 2009.

Game Theory For Self-Driving Cars

Fanta Camara^{1,2} and Charles W. Fox^{1,2,3}

<https://doi.org/10.31256/Sk9Zg2D>

Abstract—Pedestrian behaviour understanding is of utmost importance for autonomous vehicles (AVs). Pedestrian behaviour is complex and harder to model and predict than other road users such as drivers and cyclists. In this paper, we present an overview of our ongoing work on modelling AV-human interactions using game theory for autonomous vehicles control.

I. INTRODUCTION


Autonomous Vehicles (AVs) also called “self-driving cars” are appearing on the roads. The technology is claimed by many automotive companies and their arrival on the market was announced for 2020 [1]. But their future interaction with other road users raise some concerns. Autonomous vehicles currently lack of the ability of human drivers to read the personality of other road users, predict their future behaviour and then interact with them. A comprehensive review of pedestrian modelling techniques for AVs was recently proposed, ranging from low level sensing, detection and tracking models introduced in [3] to high level interaction and game theoretic models of pedestrian behaviour presented in [4].

A European project called CityMobil2¹ used transport data science [12] to reveal a drawback of highly safe and perfect autonomous vehicles. This project launched a trial with an autonomous minibus in two European cities, in La Rochelle (France) and in Trikala (Greece). After a few days of driving, people became used to the minibus and they learnt its driving behaviour, the AV’s behaviour was easily predictable, as it would avoid any obstacle by a stop. Thus, pedestrians started stepping intentionally in front of the minibus [14]. In most of these cases, the minibus was slowed down or stopped for fun.

This inability of current AVs to accurately predict pedestrian crossing intent is known as “the big problem with self-driving cars” [2]. Pedestrians do not exhibit the same behaviour with human drivers, hence the European project interACT², to which this work is part of, is investigating current human drivers-road users interactions. From these observations, we are trying to understand how these interactions occur in order to develop new behavioural models for road users e.g. [13] [15] [16] and new eHMI (external Human-Machine Interface) solutions that could facilitate the communication for autonomous vehicles in mixed traffic environments, i.e. with human-driven cars, cyclists and pedestrians.



(a) Game of Chicken



(b) Sequential Chicken Model

Fig. 1: Game of Chicken: two agents try to cross over an intersection as quickly as possible while avoiding a collision. The first agent to pass wins the game (reward), the second loses (small penalty) and they are both bigger losers if there is a collision (large penalty).

II. GAME THEORY MODEL

As a solution to the minibus problem, we started using a game theory model called the game of chicken, as shown in Fig. 1a. Game theory is a well-known framework used for modelling decision-making between rational agents. We proposed a mathematical model for the game of chicken [13], a discrete sequential game theory model called the Sequential Chicken Game, for negotiations between an autonomous vehicle and a pedestrian at an unsignalized intersection, as shown in Fig. 1b. This model shows that not only the first agent to yield is more likely to lose the game but also if the AV only uses its position to signal its intent, there must exist a small probability for a collision to occur. This collision probability can be used as a threat for the pedestrian, preventing them from stepping intentionally in front of the AV.

III. EMPIRICAL EXPERIMENTS

A. Board Game Experiment

A first empirical study [11] expanded the sequential chicken model using empirical data to measure behaviour of humans in a controlled plus-maze experiment with participants playing the game of chicken as a board game. This study provided an empirical understanding of the human factors required by future autonomous vehicles. In the first three games, i.e. natural games, players were simply told to cross over the intersection as quickly as possible. After playing the natural games, each group played a further three games in which specific chocolate rewards were specified in advance, i.e. chocolate games. With these two game types, we found that more collisions occurred during the chocolate game than in the natural game. The results showed that participants had a preference for saving time U_{time} rather than avoiding a

¹ Institute for Transport Studies (ITS), University of Leeds, UK

² Lincoln Centre for Autonomous Systems (LCAS), School of Computer Science, University of Lincoln, UK

³ IbeX Automation Ltd, UK

¹https://www.youtube.com/watch?v=ls6xsj_fCWU

²<https://www.interact-roadautomation.eu/>

collision U_{crash} . Such parameters (U_{time} , U_{crash}) of the model could be inferred via a Gaussian Process regression.

B. Physical Experiment

We later developed a novel empirical method [5] based on tracking real humans in a semi-structured environment, in order to model and predict their behaviour with game theory. We made use of dynamic programming to compute the optimal game theoretic solution form, then found the behavioural parameters via empirical observation and a Gaussian Process regression analysis. This model formed a step towards game-theoretic controllers for autonomous vehicles in similar real-world situations such as negotiations over priority at un-signalised road-crossings. This second study showed that participants were globally playing rationally, 11% of them deviated from their optimal behaviour. It also confirmed participants preference for time saving rather than collision avoidance, this unusual result was due to the high safety conditions of the experiment.

IV. SEQUENCE ANALYSIS OF PEDESTRIAN-VEHICLE INTERACTIONS

A. Sequence Patterns Recognition

We collected a large scale data from real-world human road crossings at the intersection near the University of Leeds, UK. Pedestrian-vehicle interactions were decomposed into sequences of independent discrete events [9]. We looked for common patterns of behaviour that can predict the winner of an interaction, which can thus be integrated into game-theoretic AV controllers to inform real-time interactions. We used probabilistic methods – logistic regression and decision tree regression – and motif analysis to analyse sub-sequences of actions used by both pedestrian and human drivers while crossing. We found predictive features that could inform the AV about the eventual winner of an interaction.

B. Filtration Analysis

We then used the same dataset of pedestrian-vehicle interaction sequences to study the temporal orderings (filtration) in which features (including signals from the pedestrian) can be revealed to an autonomous vehicle and their informativeness over time during pedestrian-vehicle interactions [8]. This framework suggests how optimal stopping controllers may then use such data to enable an AV to decide when to act (by speeding up, slowing down, or otherwise signalling intent to the pedestrian) or alternatively, to continue at its current speed in order to gather additional information from new features, including signals from that pedestrian, before acting itself. In this study, we found that the AV should wait and observe about 7 to 10 features before acting/making its decision.

V. PEDESTRIAN INTENTION ESTIMATION

To optimally interact with pedestrians, autonomous vehicles must be able to predict their crossing intent. Thus, we developed a model inspired by the Sequential Chicken model. It appeared that a heuristic method, simply based on

tracking data, was found to be very efficient in estimating crossing intent for most of the interactions [10]. However, this heuristic model would fail in more complex and maybe critical interaction scenarios.

VI. VR EXPERIMENTS

As virtual reality (VR) offers the opportunity to experiment on human behaviour in simulated real world environments that can be dangerous or difficult to study, we used it to develop three simple experiments about pedestrian-AV interaction at non signalized crossings. VR allows us to better understand pedestrian crossing behaviour in more realistic conditions than in our previous artificial laboratory experiments and also to improve the AV game theoretic behaviour model.

In a first experiment [6], we asked participants to cross the road as they would do in every day life. We recorded their trajectories in order to learn their behaviour preferences, i.e. time delay vs collision avoidance. The virtual AV's decision-making was based on the Sequential Chicken model [13], which is a discrete model, thus the car had a slow and a fast speed. Our analysis of the data showed that participants were more cautious in crossing and often yielded for the AV.

In a second and third experiments [7], we wanted to learn from the participants which combination of space and time parameters (from the game theory model) would make the car behave more “naturally” and also to discover if there is any behavioural change in crossing in different environments and with different car models. Experiment 2's environment was a wide tarmac road with a narrower pathway and the AV was a normal sized-car whereas in Experiment 3 the environment looked more like a park/garden and the car looked like a small podcar. Participants were presented each time with an AV that had different parameters, they were asked whether they found the interaction with the virtual AV “natural” or “un-natural”, they had to rate it on a scale from 1 (un-natural) to 5 (natural). Two methods were used to change the parameters of the car:

- Brute Force: we used predefined orderly sets of parameters one after the other
- Gradient Descent: we started with a hypothetical optimal parameter and then changed the parameters following the preferences expressed by each participant.

The results show that pedestrians prefer an AV that makes its decisions quickly and that pedestrians behave similarly in different environments.

VII. CONCLUSIONS AND FUTURE WORK

This is a work in progress on self-driving car technology. We present game theory as a tool to model future human interactions with autonomous vehicles. Semi-structured empirical and VR experiments with human participants and interaction sequence analysis provide a better understanding of human behaviour by inferring their behaviour parameters using Gaussian Process regression. Future work will look into developing the game theory model on a real self-driving car and to test its validity by performing some experiments with human participants.

REFERENCES

- [1] B. Bontrager. The race to fully autonomous cars, 2018. Medium.com.
- [2] R. Brooks. The big problem with self-driving cars is people and we'll go out of our way to make the problem worse, 2017. IEEE Spectrum.
- [3] F. Camara, N. Bellotto, S. Cosar, D. Nathanael, M. Althoff, J. Wu, J. Ruenz, A. Dietrich, and C. W. Fox. Pedestrian models for autonomous driving part i: low level models, from sensing to tracking. *Under submission to IEEE Transactions on Intelligent Transportation Systems*, 2020.
- [4] F. Camara, N. Bellotto, S. Cosar, F. Weber, D. Nathanael, M. Althoff, J. Wu, J. Ruenz, A. Dietrich, A. Schieben, G. Markkula, F. Tango, N. Merat, and C. W. Fox. Pedestrian models for autonomous driving part ii: high level models of human behaviour. *Under submission to IEEE Transactions on Intelligent Transportation Systems*, 2020.
- [5] F. Camara, S. Cosar, N. Bellotto, N. Merat, and C. W. Fox. Towards pedestrian-av interaction: method for elucidating pedestrian preferences. In *IEEE/RSJ Intelligent Robots and Systems (IROS) Workshops*, 2018.
- [6] F. Camara, P. Dickinson, N. Merat, and C. W. Fox. Towards game theoretic av controllers: measuring pedestrian behaviour in virtual reality. In *IEEE/RSJ International Conference on Intelligent Robots and Systems Workshops*, 2019.
- [7] F. Camara, P. Dickinson, N. Merat, and C. W. Fox. Examining pedestrian behaviour in virtual reality. In *Transport Research Arena (TRA) 2020*, 2020.
- [8] F. Camara, O. Giles, R. Madigan, M. Rothmüller, P. H. Rasmussen, S. A. Vendelbo-Larsen, G. Markkula, Y. M. Lee, L. Garach, N. Merat, and C. W. Fox. Filtration analysis of pedestrian-vehicle interactions for autonomous vehicles control. In *Proceedings of the 15th International Conference on Intelligent Autonomous Systems (IAS-15) workshops*, 2018.
- [9] F. Camara, O. Giles, R. Madigan, M. Rothmüller, P. H. Rasmussen, S. A. Vendelbo-Larsen, G. Markkula, Y. M. Lee, L. Garach, N. Merat, and C. W. Fox. Predicting pedestrian road-crossing assertiveness for autonomous vehicle control. In *Proceedings of IEEE 21st International Conference on Intelligent Transportation Systems*, 2018.
- [10] F. Camara, N. Merat, and C. W. Fox. A heuristic model for pedestrian intention estimation. In *Proceedings of IEEE 22nd International Conference on Intelligent Transportation Systems*, 2019.
- [11] F. Camara, R. Romano, G. Markkula, R. Madigan, N. Merat, and C. W. Fox. Empirical game theory of pedestrian interaction for autonomous vehicles. In *Measuring Behavior 2018: 11th International Conference on Methods and Techniques in Behavioral Research*. Manchester Metropolitan University, March 2018.
- [12] C. Fox. *Data Science for Transport: A Self-Study Guide with Computer Exercises*. 2018.
- [13] C. W. Fox, F. Camara, G. Markkula, R. Romano, R. Madigan, and N. Merat. When should the chicken cross the road?: Game theory for autonomous vehicle - human interactions. In *Proceedings of VEHITS 2018: 4th International Conference on Vehicle Technology and Intelligent Transport Systems*, January 2018.
- [14] R. Madigan, S. Nordhoff, C. Fox, R. E. Amini, T. Louw, M. Wilbrink, A. Schieben, and N. Merat. Understanding interactions between automated road transport systems and other road users: A video analysis. *Transportation Research Part F: Traffic Psychology and Behaviour*, 66:196 – 213, 2019.
- [15] G. Markkula, R. Romano, R. Madigan, C. W. Fox, O. T. Giles, and N. Merat. Models of human decision-making as tools for estimating and optimizing impacts of vehicle automation. *Transportation Research Record*, 2018.
- [16] J. Wu, J. Ruenz, and M. Althoff. Probabilistic map-based pedestrian motion prediction taking traffic participants into consideration. In *IEEE Intelligent Vehicles Symposium (IV) June 26-30, 2018, Changshu, Suzhou, China*, 2018.

Deep learning for robotic strawberry harvesting

1st Xiaodong Li
School of Computer Science
University of Lincoln
Lincoln, United Kingdom
doli@lincoln.ac.uk

2nd Charles Fox
School of Computer Science
University of Lincoln
Lincoln, United Kingdom
scoutts@lincoln.ac.uk

3rd Shaun Coutts
Lincoln Institute of Agri-Food Technology
University of Lincoln
Lincoln, United Kingdom
scoutts@lincoln.ac.uk

<https://doi.org/10.31256/Bj3Kl5B>

Abstract—We develop a novel machine learning based robotic strawberry harvesting system for fruit counting, sizing/weighting, and yield prediction.

Index Terms—machine vision, cascade detector, Harr/LBP feature, yolo.

I. INTRODUCTION

Strawberries are a high-value crop all around world, but during harvest season, due to the weather fluctuations, strawberry yields can vary greatly on a daily basis. It remains a challenge for farmers to efficiently manage labour and transport which rely heavily on accurate prediction of near future production.

Traditional yields estimation is time consuming and labour intensive. With the maturation of low cost camera sensing and corresponding vision processing technology, machine vision has become a potential alternative to traditional way. It has high adaptivity to variant image quality.

The Viola-Jones cascade detector used here is well supported in the OpenCV library with both Haar-like feature [1] and Local Binary Patterns(LBP) feature [7] [6]. Notably it was the first real-time (CPU-based) face detector and recognizer [1].

The project was initially low cost CPU-based platform which was later upgraded with an Nvidia GPU card (GTX TITAN X). The required expanding and enhancement on existing functionality, e.g, ripeness prediction, brought our attention to Deep learning.

Although there are already success [3] with deep learning, this quick development integrated with existing system is notable. By passing the results of the Viola-Jones detector to a YOLO (You Only Look Once) [2] [4] based deep learning system, the classification is achieved in various color category and improvement on detection accuracy, etc.

II. SYSTEM ARCHITECTURE AND METHODOLOGY

A. System design

The system is shown in Fig.1, where "Video input" of 'back2back' opposite facing 2-camera provides consecutive image frame (fps: 30); "Image Tracking" is to track individual images simply by a template matching strategy to resolve overlapping, and it reduces the computational cost by operation conducted only on parsed none overlapped image, which improved the running speed for detectors as detection doesn't need apply on each frame, this makes real time detection

possible and overcomes double counting; "Fruits Detection" applies a trained cascade detector for strawberry detection. The cascade detector model has a single class for the full range of strawberry including flowers; To further split this class into different categories, the "Classification" used a trained YOLO model to classify color difference, i.e. green, white and red, including a leaf model to further remove false detection; "Count/Size/weight" can provide different metrics for different purpose, e.g the number of strawberries in different color, with a pre-defined fruit shape model (e.g sphere, cylinder and cubic, etc) the metrics such as diameter, circumference, or volume can be obtained, further combined with strawberry's density, the overall harvest weights can be estimated, etc. currently, a simple method based on the average distance between the camera and the strawberry rack is applied, which combines the calibrated camera parameters to calculate the dimension or volume according to the strawberry shapes; according to weather condition and farmer's experience, the "Yields Prediction" provides guide on the amount of ripe strawberries for the harvest in different time, etc.




Fig. 1: System Structure

B. Viola-Jone Cascade Detector

Viola-Jones [1] detector is a Harr-like cascade detector, which can reduce the computational cost on the image intensities. Using a sliding window over the image, Harr-like features are calculated and compared by the difference from a learned threshold. While single features are weak learners, a classifier cascade is used to obtain stronger classifiers by combining them. Viola-Jones takes a variant of adaptive boosting (AdaBoost) for feature selection and training a detector using object and background images. A single classifier consists of a weighted sum of many weak classifiers, where each weak classifier is a threshold on a single Haar-like rectangular feature. The weight associated with a given sample is adjusted based on whether or not the weak classifier correctly classifies the sample. Haar-like features boost real-time detection for human faces, but still infeasible for larger image.

1) *Local Binary Patterns features*: LBP features [7] [6] utilized here provide a suitable alternative. It is often a powerful (speed) feature for texture classification. A LBP vector can be simply calculated in an image cell (e.g 16x16 pixels of sub-window), where the pixel is compared to each of its 8 neighbors along a circle. The pixel's value is the concatenation of a binary "0" and "1", which is assigned by comparing (less or greater) with each neighbor. This 8-digit binary number is usually converted to decimal for convenience. The histogram over the cell for the frequency of each "number" occurring is computed and regarded as a 256-dimensional feature vector.

C. Deep Learning

YOLO [2] [4] is one of the most effective accurate real time object detection algorithms. It applies a single neural network to the full image, and then divides the image into regions and predicts bounding boxes and probabilities for each region. These bounding boxes are weighted by their predicted probabilities, which output the recognized object after non-max suppression applied.

Because of the small size of strawberry image (24x24 pixels) in Viola-Jones, it can not provide enough information for precise classification with Haar/LBP feature. We then utilized a self-trained YOLO model for color-based object classification to predict the ripeness of strawberries. Without manual relabeling the existing datasets, an effective and innovative integration of state-of-the-art advanced YOLO with existing system is achieved successfully.

D. Data Sets and Annotation for YOLO

The initial labeled training image datasets (45K) were collected across the world e.g Fig.2 and cropped ones like Fig.3. Inspired by YOLO training mechanism, without standard YOLO labeling, the innovative auto-'one4one' self annotation method on existing data is shown below:

$$\text{center} - x = x/w(0.5); \text{center} - y = y/h(0.5); \text{width} = w/W(1.0); \text{height} = h/H(1.0);$$

x : x -coordinate of center of the bounding box;

y : y -coordinate of center of the bounding box;

w : width of the bounding box;

h : height of the bounding box;

W : width of the whole image;

H : height of the whole image;



(a) Indoor

(b) Outdoor

Fig. 2: Typical Strawberry Image




Fig. 3: Strawberry Training Images


III. RESULTS AND SUMMARY

System was developed with C/C++, and the snapshots in Fig.4 depict system configuration&monitoring(a), detection image(b). In (a), the metrics(unit), which is configurable as request, for sizing and weighting are diameter(mm) and weights(gram) respectively.


The classification examples in Fig.5 are red strawberry (a), green one (b) and Leaves (c).

The system was tested on a 80m strawberry growing rack, through the absolute difference of the numbers between manual and machine counts divided by manual ones, we have overall (all categories together) accuracy of about 90% on counting/detection, 95% on classification for red strawberries, and 80% for green and white color respectively with the model trained on 300 samples for each category, which can be further improved by more carefully selected samples added for training.

The sizing has a threshold above 15mm in diameter, together with other functionality, we believe the weight prediction can be improved with more accurate assumed depth measurements, shape model (cylinder used here) and density setting, etc.



(a) System GUI



(b) Detection Outputs

Fig. 4: System Features



(a) Red Strawb

(b) Green Strawb

(c) Leaves

Fig. 5: System Classification

REFERENCES

- [1] Paul Viola , Michael Jones, “Rapid Object Detection using a Boosted Cascade of Simple Features”, IJCV 2001.
- [2] Joseph Redmon, Ali Farhadi, “YOLOv3: An Incremental Improvement”, Computer Vision and Pattern Recognition, April, 2018.
- [3] M. Mahmud, M. S. Kaiser, A. Hussain and S. Vassanell, “Applications of Deep Learning and Reinforcement Learning to Biological Data”, IEEE Transactions on Neural Networks and Learning Systems, vol. 29, no. 6, pp. 2063-2079, June 2018.
- [4] Joseph Redmon, Ali Farhadi, “YOLO9000: Better, Faster, Stronger”, Computer Vision and Pattern Recognition, December, 2016.
- [5] N. Cristianini , J. Shawe-Taylor, “An introduction to support Vector Machines: and other kernel-based learning methods,” , Cambridge University Press, 2000.
- [6] L. Wang , DC. He, “Texture Classification Using Texture Spectrum,”IEEE Transactions on Pattern Recognition,vol. 23, pp.905–910, 1990.
- [7] DC. He , L. Wang, “Texture Unit, Texture Spectrum, And Texture Analysis,” IEEE Transactions on Geoscience and Remote Sensing,vol. 28, pp.509–512, 1990.

3D Printed Variable Infill Soft Fingers for the SIMPA Prosthetic Arm

Daniel De Barrie
School of Engineering
University of Lincoln
Lincoln, United Kingdom
ddebarrie@lincoln.ac.uk

Khaled Goher
School of Engineering
University of Lincoln
Lincoln, United Kingdom
kgoher@lincoln.ac.uk

Khaled Elgeneidy
School of Engineering
University of Lincoln
Lincoln, United Kingdom
kelgeneidy@lincoln.ac.uk

<https://doi.org/10.31256/Wj4Jc8Q>

Abstract— The SIMPA (Soft-grasp Infant Myoelectric Prosthetic Arm) prototype prosthetic arm for toddlers features soft ‘fingers’ made from silicone-rubbers. The manual multi-stage moulding process to produce those soft-fingers presents several practical challenges that limit the design freedom and degree of customisation. This work demonstrates the potential for 3D printing highly flexible and soft fingers (Shore hardness 60A) that can achieve the desired grasping performance of the original gripper fingers by introducing variable infill densities. This initial result presented here highlights the potential of this approach to further miniaturise and improve upon the initial soft gripper design

Keywords—Soft-Robotics, Medical Devices, 3D-Printing

I. INTRODUCTION

Soft-grippers are increasingly being explored within the field of robotics [1] with a particular focus on delicate grasping tasks [2], such as the handling food products like strawberries [3]. These soft-grippers have many advantages over the traditional rigid systems, primarily the ability to passively adapt to geometric features, allowing a wide range of objects to be securely grasped. The concept of incorporating soft-grippers into the design of a prosthetic arm has only recently been explored. The device (Fig. 1), “SIMPA: Soft-grasp Infant Myoelectric Prosthetic Arm” [4], presented an alternative to the rigid approach taken by other comparable devices [5]. This approach proved successful, resulting in a device capable of grasping a wide range of objects securely, whilst also incorporating safety features such as removing sharp edges and pinch points.



Fig. 1. SIMPA Prototype

The prosthetic was primarily 3D-printed, with only the grippers being moulded. This proved to be a major manufacturing challenge, due to the intricate nature of the small-scale design, and the requirement for an internal structure through which cables are threaded. The result was many failed attempts and a forced redesign of the device in order to consider the difficulty experienced during gripper prototyping. Recently, access to novel soft-filament (COEX-

Flex TPU 60A¹), has opened up the possibility of additively manufacturing the grippers directly. The material is one of the most flexible fused deposition moulding (FDM) filament currently available and has yet to be explored as a means of manufacturing soft-grippers. The feasibility of producing wire driven soft-grippers for the SIMPA prototype directly via (FDM) 3D-printing will be explored here, with the specific focus on the effect of infill parameters on the adaptability of the soft-fingers.

II. DESIGN, MANUFACTURE AND INITIAL VERIFICATION

The grippers featured in the prototype device are simple 3-segmented (2-segments in the case of the ‘thumb’) wire-driven “finger” like structures (Fig. 2) [6]. The “finger” utilises the elastic energy stored within a soft material to extend the gripper once the cables driving the flexion are relaxed. The gripper is constructed out of two silicone rubbers, Smooth-Sil® 960 (Smooth-On, Inc.) and Dragon Skin® 30 (Smooth-On, Inc). The material properties are such that the Smooth-Sil® 960 acts as the spline of the gripper, providing the elastic storage required for the extension. Conversely, the Dragon Skin® 30 provides a skin like feel and subtleness.




Fig. 2. Original Silicon-rubber Soft-gripper Design

In the case of additively manufacturing the grippers, the performance of this composite structure must be considered. By altering the infill density throughout the design, the effective rigidity of the material can be altered; with lower infill percentages increasing flexibility. By using this approach, it is hoped that the gripper can be additively manufactured, whilst maintaining comparable elasticity and grasping performance.

The CAD model of the grippers was imported into Cura (Ultimaker BV) for pre-print processing. This model is the inverse of the mould used for the original grippers. For the initial test a line infill of 25% was used for the more rigid spline of the device (Fig. 3). The top surface of the gripper was conversely given 0% infill (Fig. 4) to approximate the feel of the soft contact pads. The grey box seen in Figures 3 and 4, represents a custom infill area, allowing multiple infill settings to be used on the same model.

¹ Data-sheet can be found at <https://flexionextruder.com/shop/x60-ultra-flexible-filament-black/>

TABLE I. OVERVIEW OF PRINT PARAMETERS USED

Property	Value
Base Infill	25% (Line)
Pad Infill	0%
Number of Top/Bottom Layers	2
Number of Walls	3
Top/Bottom Layer Pattern	Concentric
Layer Height	0.2mm
Nozzle Size/Line Width	0.5mm




Fig. 3. Reinforcement Infill Structure, Lines 25%




Fig. 4. Contact Pad Infill Structure, 0%

The gripper was manufactured on a Lulzbot Taz 6, modified with a Flexion Extruder (Diabase Engineering). The full detail of the grippers is captured in the print (Fig. 5), and takes approximately 2 hours to complete; this compares to an average of 2 days when manufactured using silicon, this is due to the multi-stage nature of the process and the required cure times of the materials used.




Fig. 5. Printed Gripper Sample with an in-built Passage for Tendons.

The additively manufactured grippers from an initial manual inspection (pulling on a cable threaded through the gripper), appear to be slightly stiffer than the silicon version. As this is the first attempt, there is no optimisation in terms of the print properties or CAD model yet. It is assumed that with further refinement it should be possible to improve on this performance, in terms of reducing the force required for actuation.

For the supporting structure (replicating the Smooth Sil®), it would be simple to set-up a rig that compares the force required for a set amount of deformation, characterising the relationship between infill and apparent stiffness for grippers produced with the COEX-Flex material. The procedure should first determine the force required for full contraction of the original silicone fingers, to act as a benchmark value. The same experimental procedure should then be performed on the printed grippers, with various modifications to print parameters such as infill percentage and pattern, number of walls, line thickness, layer height, etc. A relationship between these several parameters and the performance of the grippers can then be determined.

To the touch, the grasping surface is considerably more rigid than the Dragon Skin® 30 surface of the silicone grippers. A Shore Hardness test or comparable procedure would need to be used to determine the actual surface hardness, though this has not yet been performed. The internal tubes are predominately reason for the increased apparent stiffness, as they effectively add additional infill on the top surface. The size of these could be reduced as a simple method of improving the design. Alternatively, a redesign utilising the flexibility in design that 3D-printing offers can potentially negate these issues.

III. CONCLUSIONS AND FUTURE WORK

This paper presented preliminary results on 3D printing highly flexible and soft fingers for the previously developed SIMPA prototype to address the manufacturing challenges with the manual multi-stage moulding of the original silicone-rubber fingers. The initial gripper design, which was unable to be produced by moulding, featured the same geometry but was only 10mm wide, rather than 20mm. The small scale prevented the inner tubes from being inserted correctly and ultimately led to a redesign of the hand utilising 3 grippers, instead of 5. The smaller gripper has successfully been printed, opening up the possibility for this design concept to be re-examined, producing a 5 fingered variant of the SIMPA. Further optimisation of the design will be investigated to benefit from the design freedom offered by 3D printing. The validation would then consist of performing the same experimental procedures used to verify the SIMPA prototype, using the 3D-printed grippers. The inclusion of 3D-printing remains consistent to the ethos of the SIMPA prototype, which aims to be predominantly additively manufactured, moving away from the traditional approach to producing such devices, which still relies heavily on the skill of the craftsman.

The proposed relationship study between print parameters and actuation force would serve a more general purpose to characterise the material's use in FDM printing. This could be used in the design of general purpose soft-grippers and in other applications where variable stiffness may be required.

REFERENCES

- [1] D. Rus and M. T. Tolley, "Design, fabrication and control of soft robots," *Nature*, vol. 521, no. 7553, pp. 467–475, 2015.
- [2] K. Elgeneidy, P. Lightbody, S. Pearson, and G. Neumann, "Characterising 3D-printed Soft Fin Ray Robotic Fingers with Layer Jamming Capability for Delicate Grasping," in *2019 2nd IEEE International Conference on Soft Robotics (RoboSoft)*, 2019, pp. 143–148.
- [3] M. Schaffner, J. A. Faber, L. Pianegonda, P. A. Rühls, F. Coulter, and A. R. Studart, "3D printing of robotic soft actuators with programmable bioinspired architectures," *Nat. Commun.*, vol. 9, no. 1, 2018.
- [4] D. De Barrie, R. Margetts, and K. Goher, "SIMPA : Soft-Grasp Infant Myoelectric Prosthetic Arm," in *2020 International Conference on Robotics and Automation (ICRA)*, 2020.
- [5] P. Slade, A. Akhtar, M. Nguyen, and T. Bretl, "Tact: Design and performance of an open-source, affordable, myoelectric prosthetic hand," *Proc. - IEEE Int. Conf. Robot. Autom.*, vol. 2015–June, no. June, pp. 6451–6456, 2015.
- [6] M. Manti, T. Hassan, G. Passetti, N. D'Elia, C. Laschi, and M. Cianchetti, "A Bioinspired Soft Robotic Gripper for Adaptable and Effective Grasping," *Soft Robot.*, vol. 2, no. 3, pp. 107–116, 2015.

Task Delegation and Architecture for Autonomous Excavators

Jake Rankin
Wolfson School
Loughborough University
LE11 3TU
j.rankin@lboro.ac.uk

Laura Justham
Wolfson School
Loughborough University
LE11 3TU
l.justham@lboro.ac.uk

Yee Mey Goh
Wolfson School
Loughborough University
LE11 3TU
y.goh@lboro.ac.uk

James Morley
JCB Research
JCB Excavators Ltd
ST14 5JP
james.morley@jcb.com

<https://doi.org/10.31256/Ew3Zn4Z>

Abstract — The construction industry is required to deliver safe, productive machines. One method being considered by heavy equipment manufacturers is autonomy. Implementing autonomy to heavy machines is unique, due to the highly skilled nature of a machine’s operation meaning that different levels of autonomy may be more suitable for different tasks. Therefore, effective collaboration strategies between human operators and machines are needed. This paper proposes a machine architecture that considers the task delegation between the operator and machine.

Keywords — *Autonomy, System design, Task delegation*

I. INTRODUCTION

Many autonomous achievements can be seen within the automotive industry, with features such as cruise control and emergency braking becoming more common place, guided by SAE J3014 [1]. This has been essential for standardising a roadmap towards full autonomy for the whole sector for safety reasons. A difficulty that is faced by the automotive industry is the clarity of task delegation between SAE J3014 levels 2 and 3, with the end-user being the most notable example. There are several examples, such as Tesla’s Autopilot [2], where users assume a higher capability than stated.

The shrinking skilled workforce in the construction industry and task complexity means that task delegation between machines and operators is a challenge that needs to be understood when implementing autonomy. Heavy plant manufacturers face this challenge without a standard framework to follow. Autonomy can be achieved using established strategies like A* [3] and Rapidly-Exploring Random Tree but these struggle with task complexities [3].

Having a standardised architecture will help heavy plant manufacturers implement solutions across several types of machines whilst providing clarity of the machine capability to the construction industry. This paper addresses the issue of task delegation through the development of a novel autonomous excavator architecture and identifies opportunities to use technology such as Building Information Modelling (BIM) and Reinforcement Learning (RL) for a more integrated implementation of heavy plant machinery in the construction site.

II. EXISTING LITERATURE

One of the first autonomous excavator architectures was in LUCIE [4]. However, this was solely focused on trenches and doesn’t seem to have a layered architecture to enable fast and slow reactions. LUCIE also required hard-coding actions, such as when to curl the bucket, which resulted in less flexibility and, potentially, more processing during run-time.

One company that provides autonomy solutions, ASI, proposes a three-stage autonomy system which can be applied to different machines [5]. This doesn’t discuss how task

delegation was decided for automation tasks nor does it discuss the scenario of driver assist. It also seems more focused on bulk-digging for mining.

Mastalli et al. investigated a control system that used learning and simulation [6] but didn’t investigate driver roles. Stentz et al. proposed an autonomous loading system for truck loading [7] but this architecture is complex and task specific. An aspect that these authors focused on was the visual aspect which has advanced significantly. Mastalli et al also discussed the advantages of using RL for working in highly constrained environments without the need for hard-coding.

To date, little work has been done on task delegation for excavators, despite its importance for developing autonomy. Kim investigated task-planning for excavation, which would divide the excavator’s project into tasks to be completed [8]. This could have been an opportunity to mention how tasks could be delegated to operators as autonomy developed.

There are several examples in literature of human-machine task delegation in robotics that can be used to address the shortcomings of the automotive sector when applied to heavy plant. Sheridan [9] defined four main application areas of Human-Robot Interaction (HRI), supervisory control and automated vehicles were most applicable to this project. Task delegation is often applied to manufacturing with heavy or repetitive movements allocated to a robot and infrequent dexterous tasks given to a human. Another framework of delegation included Task Specification Trees [10], which used auctions to delegate tasks as actioned by human requests. To delegate, a Delegation Protocol, an extension of a Contract Net Protocol, was used. This relied on an internal manager or subcontractor that auctioned tasks to several agents. These could be based on distance and skillset, depending on the requirements.

One final aspect to consider is the use of turn-taking. With the operator and machine working on the same task, sometimes it is useful for a robot to take the driver’s turn when fatigue is noticeable. Turn-taking is also important as it retains driver skill and attention, which makes it a valuable consideration for level 2 and 3 autonomy. Turn-taking was applied to a social machine [11] where good turn-taking was described as having little overlap of the human and minimal time spent between turns.

A crucial consideration is for the machine to know when to take control and how to hand it back, which is a notable problem for semi-autonomous cars [12]. This is critical in cars as the tasks the machine performs are usually at speed, which means that unexpected changes can lead to a clumsy handover. One review paper [13] discussed task allocation in the mining industry, highlighting the importance of human operators due to their judgement. Although this paper was aimed towards mining, meaning trenching isn’t considered,

there are several transferable concepts that can be applied to construction.

III. DRIVER ROLES

An expert operator performed 15 trenching tasks in different soil conditions at different operating speeds in Mevea [14] to understand driver behaviour and to identify autonomy opportunities. The driver over-corrected when aligning the bucket to the trench, as seen in Figure 1. This was identified as an initial task for a machine to consider when operating alongside an expert operator.




Fig. 1. Slew and boom during trenching, with hesitations circled

IV. PROPOSED DESIGN

A. Architecture

To accompany the task delegation, a three-layered architecture made up of planning, control and hardware layers has been developed, see Figure 2. The control layer has also been designed with human and machine operation in mind with a task delegation module (Delegator). The operator and controller work together on a task, based on the user needs, exchanging information as instructions or feedback.

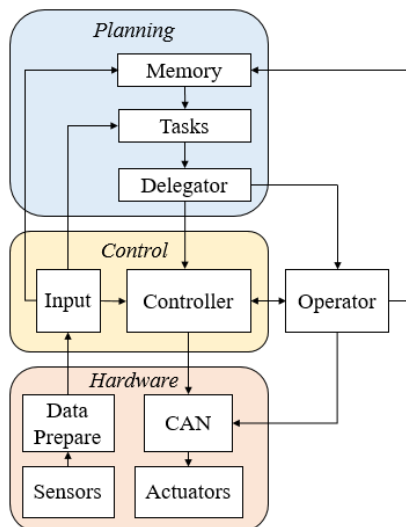


Fig. 2. Autonomous Excavator Architecture

The planning layer is where long-term plans and delegation occur. By using turn-taking, control alternates between the operator and the controller. Memory contains the excavation plans and can be the link-up to BIM, reducing site setting-out and providing managers with work information. These are divided into tasks that are then delegated to the machine and operator.

In the control layer, designers have the option of implementing their own control system and determining their own inputs. This split allows the introduction of machine-learning algorithms, such as RL, to be easily implemented. The operator can interrupt the controller and the controller informs the operator on its progress as well as identifying if it

is uncertain for a decision. Training can be done with an excavator digital twin, using software like Mevea [14] before testing on a real-world machine.

The hardware layer is where control commands enter the CAN bus and control the requested actuators. Sensors receive data which then undergo any necessary preparation prior to being inputted into the controller. This modular approach can then be applied to other machines, reducing costs and providing a standard architecture. This is where safety features must also be included to ensure a reactive response.

B. Task Delegation

Performance is not the only factor to consider with autonomy; autonomous features could lead to operator mind wandering, leading to accidents [15]. The transition between tasks is also important [12] as there is an adjustment period.

A turn-taking system addresses operator mind wandering, with both the machine and operator given tasks, based on the job and autonomy level. Additionally, the machine can ask to takeover a task, repeating the task five times and returning to position, based on the performance drop of the operator. By taking turns, the machine can help to guide a novice and help alleviate the workload of an expert operator without diminishing expert skills. The actuations that an excavator can perform depend on autonomy level, with slew and boom identified as the initial tasks. The proposed structure for analysing performance is based on actor-critic methods to review both the driver and RL's performance. The turn-taking module is shown in Figure 3 along with example operator and learning algorithm behaviour.




Fig. 3. Turn-taking policy (left) and learner predicting with operator (right)

V. CONCLUSIONS

A standardised architecture for autonomous excavation that considers task-delegation has been proposed within this paper, which is simple to implement and flexible for different applications. One of the most promising control strategies is RL, as it has been applied to several robotics tasks [16]. Therefore, it is the focal point for future work because of its operator-like behaviour. It also has the potential of being flexible enough to be applied to new machines outside of excavators and the construction industry.

The next stage in this research is to implement the architecture into real-world machines to confirm transferability. Although future work is focused on investigating the feasibility of RL, it is also important to consider how it can be implemented. Task delegation will also include the further investigation of human factors such as mind wandering to ensure safety and well-being.

VI. ACKNOWLEDGEMENTS

The research work presented is sponsored by the EPSRC-CDT Embedded Intelligence, grant reference EP/L014998/1.

VII. REFERENCES

- [1] SAE International, "Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles (J3016 Ground Vehicle Standard) - SAE Mobilus," 2018.
- [2] S. Levin, "Tesla fatal crash: 'autopilot' mode sped up car before driver killed, report finds | Technology | The Guardian," *The Guardian*, 2018.
- [3] G. Ye and R. Alterovitz, "Demonstration-guided motion planning," in *Springer Tracts in Advanced Robotics*, 2017, vol. 100, pp. 291–307.
- [4] D. A. Bradley and D. W. Seward, "The Development, Control and Operation of an Autonomous Robotic Excavator," *J. Intell. Robot. Syst.*, vol. 21, pp. 73–97, 1998.
- [5] ASI, "Robotic Excavator | ASI," 2019. [Online]. Available: <https://www.asirobots.com/mining/excavator/>. [Accessed: 28-May-2019].
- [6] C. Mastalli and G. Fernández-López, "A proposed architecture for autonomous operations in backhoe machines," in *Advances in Intelligent Systems and Computing*, 2015, vol. 302, pp. 1549–1559.
- [7] A. Stentz, J. Bares, S. Singh, and P. Rowe, "Robotic excavator for autonomous truck loading," *Auton. Robots*, vol. 7, no. 2, pp. 175–186, 1999.
- [8] J. Kim, D. eun Lee, and J. Seo, "Task planning strategy and path similarity analysis for an autonomous excavator," *Autom. Constr.*, vol. 112, p. 103108, Apr. 2020.
- [9] T. B. Sheridan, "Human-Robot Interaction: Status and Challenges.," *Hum. Factors*, vol. 58, no. 4, pp. 525–32, Jun. 2016.
- [10] P. Doherty and F. Heintz, "A Delegation-Based Cooperative Robotic Framework," 2011.
- [11] C. Chao and A. L. Thomaz, "Turn Taking for Human-Robot Interaction," 2010.
- [12] H. E. B. Russell, L. K. Harbott, I. Nisky, S. Pan, A. M. Okamura, and J. C. Gerdes, "Motor learning affects Car-To-Driver handover in automated vehicles," *Sci. Robot.*, vol. 1, no. 1, Dec. 2016.
- [13] W. P. Rogers *et al.*, "Automation in the Mining Industry: Review of Technology, Systems, Human Factors, and Political Risk," *Mining, Metallurgy and Exploration*, vol. 36, no. 4. Springer, pp. 607–631, 15-Aug-2019.
- [14] Mevea, "Mevea Simulation Solutions," 2019. [Online]. Available: <https://mevea.com/>.
- [15] M. Cunningham and M. Regan, "Driver inattention, distraction and autonomous vehicles," in *4th International Driver Distraction and Inattention Conference 2015*, 2015.
- [16] T. Haarnoja, V. Pong, K. Hartikainen, and S. Levine, "Soft Actor Critic—Deep Reinforcement Learning with Real-World Robots – The Berkeley Artificial Intelligence Research Blog," 2018. [Online]. Available: <https://bair.berkeley.edu/blog/2018/12/14/sac/>. [Accessed: 25-Jun-2019].

Understanding human responses to errors in a collaborative human-robot selective harvesting task

Zhuoling Huang¹, Adrian Salazar Gomez¹, Genki Miyauchi^{1,2} Simon Parsons^{1,3}, Elizabeth I. Sklar^{1,3}
 Richie Bird¹, Amar Singh Kalsi¹, ²University of Sheffield, UK ³University of Lincoln, UK
 Chipp Jansen¹, Zeyang Liu¹ g.miyauchi@sheffield.ac.uk sparsons,esklar@lincoln.ac.uk
¹King's College London, UK
 zhuoling.huang@kcl.ac.uk

<https://doi.org/10.31256/Xp9Yb4H>

Abstract—A human-robot approach for the farm of the future motivates robotics researchers to consider ways in which automated devices and intelligent systems can work alongside farmers to address a wide range of highly specialised but often repetitive tasks. The work presented here investigates a collaborative task in which a human and robot share decision making about the readiness of strawberries for harvesting. Preliminary experimental results with two different robot behaviours and two different user groups are compared.

Index Terms—Human-robot collaboration, Agriculture, Computer vision, Machine learning

I. INTRODUCTION

Family farms provide 70% of the food consumed worldwide [1]. However, because young labour is moving out of the agricultural sector, the existing labour force is ageing and the labour pool is not being replenished, the sector faces many difficult challenges. One strategy for addressing some of these challenges is to develop intelligent human-robot solutions in which agriculture involves *human farm workers* collaborating with *robots* to perform a wide variety of tasks. The work presented here introduces a strawberry harvesting system, but the techniques could also apply to other high-value crops. A collaborative human-robot solution to high-value crop harvesting could entail tasks such as identifying which fruits are ready to pick, selecting an appropriate position for a robot manipulator, gently removing the fruit from its stalk and packing the fruit for shipping.

Along this selective harvesting pipeline, detecting the target is, without a doubt, a very important step. Recent research has applied *machine learning (ML)* to detect fruits among which variations of *YOLO* [2]–[4] and *R-CNN* [2], [5]–[8] have been proven to work well. However, none of the existing detection methods can guarantee perfect *precision* (percentage of selected answers that are correct) and *recall* (percentage of all correct answers that are selected) for the detection task—there are always some number of false positive (type I, incorrectly selected) and false negative (type II, missed selection) errors. Here we test the hypothesis that different human users will respond differently to these two different types of errors within the context of selective harvesting.

Supported by China Scholarship Council (CSC).

II. EXPERIMENTAL SYSTEM DESIGN

Basic structure of the setup. Our experimental system for collaborative strawberry detection includes several components: an emulated farm, a mobile robot, a vision server and a graphical user interface (GUI), as shown in Fig. 1. In the *emulated strawberry farm*, 43 high-resolution images of strawberry plants taken on four different real strawberry farms are presented. Of the images, 35 contain 51 mature strawberries in total, 3 contain green strawberries only, and 5 contain only farm background (i.e., leaves, grass, soil). Full-colour prints of these images are attached to the office corridor walls near our robotics lab at a height appropriate for the robot. The *robot* is a Turtlebot2 holding an Asus Xtion RGB camera, controlled using the MRTeAm framework [9] and Robot Operating System (ROS) [10]. The framework enables the robot to move following a command from the user, and to take images of the environment around it, which it passes to the vision server with twin RTX2080Ti GPUs. The *vision server* then uses these data to identify ripe strawberries.

Different methods can be applied to the object identification process. Here, we employed two different deep learning detection methods (details below): *Faster R-CNN*, a typical two-stage detection network, and *YOLOv3*, a widely used one-stage detection network. The *user interface* receives strawberry identification and ripeness estimates, as well as raw images. The user confirms or corrects the robot's estimate and sends their decision back to the vision server.

Here we are interested in comparing the responses of different user groups to different types of errors produced by different ML methods used to identify ripe fruit. We hope the results can help us to distinguish more suitable users to provide feedback for future reinforcement learning tasks and to choose preferred classifiers for these users.

Neural networks used for identification. Many different object detection methods can be applied to this system. Since one of our aims is to test the overarching hypothesis that people will respond differently to two common types of classification errors (*false positive* versus *false negative*), we chose two widely used methods for our pilot study.

YOLOv3 is a one-stage detection method, which does the localization and classification at the same time in a single




Fig. 1. Components and structure of the setup.

end-to-end network. One stage methods are usually faster and better able to support real-time operation than two-stage detection methods. YOLOv3 predicts bounding boxes and a corresponding “objectness” score for each (best predicted overlap with ground truth). We started with a YOLOv3 model with pre-trained weights from Darknet-53 that had been trained on the COCO data set [11]. We then trained it using 145 images drawn from a set of images taken by the authors at 4 strawberry farms in the UK and China. These images include 1497 ripe strawberries and 2047 unripe strawberries.

Faster R-CNN is a two-stage object detection approach that uses convolutional neural networks (CNNs) with two key components: a Regional Proposal Network (RPN) and a detection network. These share convolutional layers, making the model faster and more efficient [12], but can lead to mismatched goals in feature learning [13]. The Faster R-CNN used in our experiments follows the network structure proposed in [12], while the detection network following the RPN is based on a Fast R-CNN structure [14]. The network was trained on 540 images from our dataset, containing 3738 ripe strawberries, and 3735 unripe strawberries.

III. EXPERIMENT DESIGN AND RESULTS

Human subjects were instructed to complete two *missions*, each with a different robot: one which employs Faster R-CNN and the other YOLOv3. For this experiment, the major difference is that Faster R-CNN produces more false positives while YOLOv3 produces more false negatives when applied to our experiment setup. To avoid human bias against the named algorithms, when talking about the two robots with the human subjects, we named them Robot *Felisa* for the Faster R-CNN behaviour and Robot *Yasmin* for the YOLOv3 behaviour. The order of missions assigned was randomised. Thirty (30) human subjects participated, primarily postgraduate students, 15 with a background in deep learning (the experienced group), 15 with other backgrounds including other engineering communities, law and linguistics (the non-experienced group). They were asked to complete a pre-survey, then ran two missions (each with a different robot behaviour), completed a survey after each mission and a final survey after both missions. The survey questions were grouped according to four features: perceived *success*, *collaboration*, *trust* and *speed*. The two post-mission surveys have the same questions, but specify the name of each robot rather than just generically “robot”. The order of survey questions was randomised. Answers indicate any predisposed user bias.

Comparison of classifiers. According to *objective data* collected by the system, we found that for the time-related measures and the total number of strawberries detected, there were no statistically significant differences between the two classifiers (using Student’s *t*-test). However, there were statistically significant differences in the number of true positive (TP), false positive (FP) and false negative (FN) results. This trend, with Felisa providing more FPs and Yasmin providing more FNs, follows our experiment design and was also noticeable to the users. According to *subjective data* collected from surveys, users trust the two robot behaviours equally, but they prefer Felisa for completing this task faster as well as both in complex environments and environments with clear pictures, since users believe that collaboration is quicker with Felisa.

User background comparison. The objective data shows that there is no significant difference in performance between the experienced group and the non-experienced group when collaborating with either robot behaviour. This means that our system is suitable for users with different backgrounds. However, the subjective data shows different attitudes towards the classifiers from different groups of users. As shown in Fig. 2, the groups of users have different expectations for collaboration and speed, but the two missions reduced the difference between them. Overall, experienced users trust Yasmin more while non-experienced users trust Felisa more. Considering preferences for different tasks, the same difference holds for clear pictures. However, both groups prefer Felisa for complex environments and when speed is prioritised.

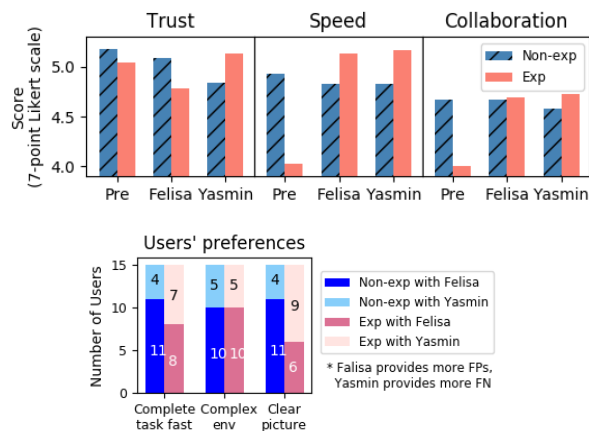


Fig. 2. Subjective data from different user groups. Trust, speed and collaboration are reported on a 7-point Likert scale.

IV. SUMMARY

Our results show that our human-robot collaborative strawberry harvesting system could be successfully used for comparing people’s preferences between the two error types when working with robots. Experienced users trust the robot that provides more false negative results and non-experienced users are the opposite. However, the majority of both groups prefer to work with the robot that produces more false positive results when completing tasks in complex environments.

REFERENCES

- [1] "Proc of the Conf on Family Farming: A dialogue towards more sustainable and resilient farming in Europe and the world," European Commission, Nov 2013, ec.europa.eu/agriculture/events/family-farming-conference-2013_en.htm.
- [2] H. Kang and C. Chen, "Fast implementation of real-time fruit detection in apple orchards using deep learning," *Computers and Electronics in Agriculture*, vol. 168, p. 105108, 2020.
- [3] Y.-P. Liu, C.-H. Yang, H. Ling, S. Mabu, and T. Kuremoto, "A visual system of citrus picking robot using convolutional neural networks," in *2018 5th International Conference on Systems and Informatics (ICSAI)*. IEEE, 2018, pp. 344–349.
- [4] C. Liang, J. Xiong, Z. Zheng, Z. Zhong, Z. Li, S. Chen, and Z. Yang, "A visual detection method for nighttime litchi fruits and fruiting stems," *Computers and Electronics in Agriculture*, vol. 169, p. 105192, 2020.
- [5] Y. Yu, K. Zhang, L. Yang, and D. Zhang, "Fruit detection for strawberry harvesting robot in non-structural environment based on Mask-RCNN," *Computers and Electronics in Agriculture*, vol. 163, p. 104846, 2019.
- [6] M. Halstead, C. McCool, S. Denman, T. Perez, and C. Fookes, "Fruit quantity and ripeness estimation using a robotic vision system," *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 2995–3002, 2018.
- [7] X. Liu, D. Zhao, W. Jia, W. Ji, C. Ruan, and Y. Sun, "Cucumber fruits detection in greenhouses based on instance segmentation," *IEEE Access*, vol. 7, pp. 139 635–139 642, 2019.
- [8] S. Bargoti and J. Underwood, "Deep fruit detection in orchards," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 3626–3633.
- [9] E. Schneider, E. I. Sklar, S. Parsons, and A. T. Özgelen, "Auction-based task allocation for multi-robot teams in dynamic environments," in *Towards Autonomous Robotic Systems: 16th Annual Conference (TAROS)*. Springer, 2015.
- [10] M. Quigley, K. Conley, B. P. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, and A. Y. Ng, "ROS: an open-source robot operating system," in *ICRA OSS Workshop*, 2009.
- [11] J. Redmon and A. Farhadi, "YOLOv3: An incremental improvement," *arXiv preprint arXiv:1804.02767*, 2018.
- [12] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards real-time object detection with region proposal networks," in *Advances in neural information processing systems*, 2015.
- [13] B. Cheng, Y. Wei, H. Shi, R. Feris, J. Xiong, and T. Huang, "Revisiting RCNN: On awakening the classification power of Faster RCNN," in *Proc of European Conference on Computer Vision (ECCV)*, 2018.
- [14] R. Girshick, "Fast R-CNN," in *Proc of IEEE intl conf on computer vision*, 2015.

A Point Cloud Semantic Segmentation Framework for Embedded Systems in Agricultural Robots

Gary Storey

Department of Computer Science
Loughborough University
Loughborough, UK
G.Storey@lboro.ac.uk

Lei Jiang

Department of Computer Science
Loughborough University
Loughborough, UK
L.Jiang2@lboro.ac.uk

Qinggong Meng

Department of Computer Science
Loughborough University
Loughborough, UK
Q.Meng@lboro.ac.uk

<https://doi.org/10.31256/Bk3Wa5K>

Abstract—In this paper the task of point cloud semantic segmentation from RGB-D data in outdoor agricultural environments is addressed specifically for the challenges of detecting trees, obstacles and the safe ground to traverse. A multi-step framework is proposed to enable real-time processing on an embedded system for use in wheeled agricultural robots. The initial step uses a MobileNetV2 based deep learning model to perform semantic segmentation on an RGB image. A point cloud is created from the segmentation and depth images which is then down-sampled, finally RANSAC plane segmentation refines the final segmentation output. Initial findings show the method performs well at labelling the desired targets while also running at up to 9fps on the embedded system.

I. INTRODUCTION

Within computer vision and robotics, point cloud semantic segmentation (PCSS) has in recent years become an activate area of research [1]. The aim of PCSS is to provide semantic labelling to each point in the cloud, this information can then be used to identify specific objects/areas within an environment for example pedestrians, trees and roads. This can provide a robot with contextual understanding of the environment and therefore can plan and act accordingly.

A variety of different approaches have been applied for the task of PCSS from a range of different data sources including RGB-D and Lidar. Methods such as PointNet [2] and PointSeg [3], perform segmentation in the point cloud based upon 3D shape matching and edge features, without the consideration of RGB data. These methods are reliant on available ground truth data for model training, while available data sets like Semantic3D.net [4] are limited to urban and indoor environments. In comparison the RGB domain contains much richer data sets like the ADE2K [5] which also contains agricultural scenes.

In this paper a novel multi-step framework for PCSS from RGB-D data is presented. The application domain for the framework is wheeled field robots, specifically those operating in outdoor agricultural environments. The targeted hardware is an embedded system the Nvidia Jetson TX2, with data from an RGB-D stereo depth camera the Intel Realsense D435. The PCSS framework proposed aims to perform segmentation of trees, obstacles and the ground. This task provides three

distinct challenges firstly unlike their urban counterparts which benefit from distinct roads with relatively flat surfaces, an agricultural wheeled robot may be required to drive on multiple ground types, i.e gravel, mud, grass some of which have deeply uneven surfaces. Therefore not only is it important to identify the ground but also to detect ground areas which may unsuitable to traverse, to allow safe path planning. Secondly there needs to be robustness to unknown or incorrectly labelled obstacles as segmentation models can be prone to error on unseen object types. Finally computational efficiency is required to run on the embedded hardware platform. To tackle these challenges the proposed framework initially employs a deep learning model to provide semantic segmentation on a RGB image to exploit the rich data sets in this domain, this is then refined within the point cloud initially applying intelligent down-sampling followed by RANSAC plane segmentation [6].

II. PROPOSED METHOD

The proposed multi-step framework for PCSS initially captures aligned RGB-D frame data from the Realsense D435 camera before performing the three main steps as detailed in this section. An overview of this framework is shown in Figure 1. The Open3d [8] Python library was used for creating and working with point cloud data.




Fig. 1. Framework Steps Overview. (1) Captured RGB and Depth Data, (2) RGB Semantic Segmentation, (3) Point Cloud Creation, (4) Refined Semantic Segmentation.

A. RGB Semantic Segmentation

Semantic segmentation is the process of predicting a class label for each pixel in the input image. The proposed framework applies a lightweight model to enable fast performance on the Nvidia Jetson TX2. A MobileNet2 [7] encoder with dilated convolutions [9] in tandem with a single layer decoder is used and implemented in PyTorch. The model was trained using the ADE2K training data set [5]. Given an input RGB




Fig. 2. Point Cloud Semantic Segmentation Results. (Left) RGB Image, (Centre) RGB Segmentation Output, (Right) Final output, blue denotes safe ground, red for obstacles and green for trees.

image of $n \times m \times 3$, the output of the RGB segmentation model S is $n \times m \times 3$ where each predicted class correlates to a specific RGB value.

B. Point Cloud Creation

Given the RGB segmentation output of S , the corresponding depth image D , a depth scale λ and the intrinsic camera parameters $C = \{c_x, c_y, f_x, f_y\}$, a point cloud $PC = \{pc_1, pc_2, \dots, pc_n\}$ with n points is then generated as:

$$x_i = (u_j - c_x) \times z / f_x \quad (1)$$

$$y_i = (v_j - c_y) \times z / f_y \quad (2)$$

$$z_i = d / \lambda \quad (3)$$

where u_j and v_j are the j th location in D . While c_x , c_y , f_x and f_y are the camera intrinsic parameters. Each point pc_i is a vector given by $\{x_i, y_i, z_i\}$. An associated RGB colour matrix $C = \{c_1, c_2, \dots, c_n\}$ also exists where $c_i = \{r_i, g_i, b_i\}$.

C. Point Cloud Down-Sampling and Segmentation

Processing point cloud data can be computationally expensive depending upon the number of points, to reduce this point removal via down-sampling is used. To retain sufficient detail the following down-sampling method is proposed. Firstly points corresponding to classes of high confidence are made exempt from calculation. Confidence is established via validation of the RGB segmentation model with ADE2K validation data set [5]. Specific points can be established by the RGB value of the class e.g. trees in C and removed from PC . Uniform nearest neighbour down-sampling is then applied. The PCSS is refined using RANSAC plane segmentation [6].

Final RGB class values are assigned to the relevant points in C where blue denotes flat terrain and red obstacles.

III. EXPERIMENTS

To produce an initial evaluation two experiments were performed. One to assess PCSS accuracy and a second to evaluate computational efficiency on the target hardware. The experimental parameters were as follows, RGB and depth image size was 640 by 480. Nearest neighbour down-sampling was set to 10. RANSAC applied a distance threshold of 0.1, initial points of 100 and 50 iterations.

A. Segmentation Accuracy

Five rural images were evaluated, varying in ground and objects types. Figure 2 shows the images and outputs. The RGB segmentation proves adept at classifying trees, sky, ground and some obstacles classes such as people. Where deficiencies in RGB segmentation exist, like a missed bench object in one image the point cloud based method successfully detects the obstacle. The PCSS method can also distinguish flat from uneven ground surfaces, though depth data sparsity has an impact on some predictions further into the point cloud.

B. Computational Efficiency

Point clouds were tested at depths of 1, 5 and 10 metres. The five images were each processed 10 times and final mean values are given in Table I. When depth is 1 meter the performance is 9fps at 10 meters this is reduced to 5-6fps. Prior to down-sampling the average points per cloud were 69788 and 220722 for 1 and 10 meters respectively.

TABLE I
MEAN COMPUTATIONAL EFFICIENCY BREAKDOWN IN SECONDS

Step	1 Meter	5 Meter	10 Meter
RGB Segmentation	0.064s	0.064s	0.064s
Point Cloud Creation	0.017s	0.034s	0.038s
Point Cloud Down-Sampling	0.018s	0.049s	0.046s
RANSAC	0.008s	0.018s	0.021s
Total	0.11s	0.16s	0.17s

IV. CONCLUSION

In this paper a framework for PCSS is proposed, the initial findings are promising providing a foundation for future research. Using both RGB and point cloud semantic segmentation methods leverages the current strengths of each domain and aid robustness, for example in detecting missed obstacles from the RGB domain. The ability to distinguish flat from uneven terrain is extremely relevant for field robots especially as this framework runs on an embedded system and RGB-D camera costing less than £600. There are limitations due to the qualitative nature of the results, generating ground truth data regarding ground contours is an area for future research so more substantial quantitative results can be generated, while opening up the possibility of applying direct point cloud segmentation techniques like PointSeg in fusion with RGB segmentation.

REFERENCES

- [1] J. Zhang, X. Zhao, Z. Chen and Z. Lu, "A Review of Deep Learning-Based Semantic Segmentation for Point Cloud," in *IEEE Access*, vol. 7, pp. 179118-179133, 2019.
- [2] R. Q. Charles, H. Su, M. Kaichun and L. J. Guibas, "PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation," 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Honolulu, HI, 2017, pp. 77-85.
- [3] C. Li, X. Wei, H. Yu, J. Guo, X. Tang and Y. Zhang, "An Enhanced SqueezeNet Based Network for Real-Time Road-Object Segmentation," 2019 IEEE Symposium Series on Computational Intelligence (SSCI), Xiamen, China, 2019, pp. 1214-1218.
- [4] H. Timo, N. Savinov, L. Ladicky, J. D. Wegner, K. Schindler and M. Pollefeys. "Semantic3D.net: A new Large-scale Point Cloud Classification Benchmark." *ArXiv abs/1704.03847*, 2017.
- [5] B. Zhou, H. Zhao, X. Puig, S. Fidler, A. Barriuso and A. Torralba, "Scene Parsing through ADE20K Dataset," 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Honolulu, HI, 2017, pp. 5122-5130.
- [6] M. A. Fischler, R. C. Bolles, "Random Sample Consensus: A Paradigm for Model Fitting with Applications to Image Analysis and Automated Cartography," *Comm. of the ACM*, Vol 24, 1981, pp 381-395.
- [7] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov and L. Chen, "MobileNetV2: Inverted Residuals and Linear Bottlenecks," 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition, Salt Lake City, UT, 2018, pp. 4510-4520.
- [8] Q-Y. Zhou, J. Park and V. Koltun, "Open3D: A Modern Library for 3D Data Processing," *arXiv:1801.09847*, 2018.
- [9] F. Yu and V Koltun, "Multi-scale context aggregation by dilated convolutions," *arXiv preprint arXiv:1511.07122*, 2015.

A Robotic Environment for Cognitive Assessment

Scott MacLeod, Mauro Dragone

*Department of Engineering and Physical Sciences
Heriot-Watt University
Edinburgh, United Kingdom
{sam19, m.dragone}@hw.ac.uk*

Mario Parra Rodriguez

*Department of Psychology
Strathclyde University
Glasgow United Kingdom
mario.parra-rodriguez@strath.ac.uk*

<https://doi.org/10.31256/Ss5Kx6T>

Abstract—Dementia is a broad category of brain diseases that is defined as a decline in memory and other cognitive abilities, severe enough to impede carrying out activities of daily living (ADLs). There is currently no cure, and accurate diagnosis and effective intervention is hampered by a lack of widely available, reliable, and effective forms of assessment. This paper presents an overview of ongoing research to develop a framework for continuous assessment of users' cognitive abilities based on observed performance of ADLs using non-intrusive sensing, robotic and AI technology. The ultimate goal is to pave the way for robotic and other assistive systems constantly in tune with the needs of their users, as they will be designed using the psychological insights of cognitive trajectories of normal and pathological ageing.

Index Terms—Cognitive Assessment, Dementia, Cognitive Assistance, Ambient Assisted Living, Robotic

I. INTRODUCTION

The goal of this work is to create an interactive robotic environment able to carry out continuous and fine-grained profiling of the users' cognitive status by observing and assisting them during specific Activities of Daily Living (ADLs).

Spotting cognitive deterioration early and finding correlations between cognitive and functional ability over time are key enablers for better informing diagnosis and personalised, continuous, and pro-active interventions [2], ultimately leading to more successful and sustainable treatment [1]. However, cognitive assessment today is usually performed by trained clinicians using standardised pen and paper clinical tests [3]. A specialist typically sees a patient every 6 months or annually [1], leaving large periods of time between assessments.

A number of technologies have already been proposed to assist people with cognitive impairments and reduce caregiver burden. For instance, both the COACH and the CogWatch [7] systems employ sensors and artificial intelligence (AI) techniques to observe and guide older adults suffering from dementia or from the effects of a stroke through an ADL using multi-modal feedback (such as audio/video prompts). User acceptance and thus effectiveness of these technologies relies on their ability to fit the different situations, varying preferences and changing needs of their users. However, continuous customisation remains a significant challenge in the field, which hinders widespread adoption [4].

Our approach is informed by psychologically valid assessment practices. We envisage that the system described here could be hosted in a clinic, initially, as a walk-in testing

facility, where patients' cognitive abilities can be assessed for their ability to remember and correctly execute pre-defined ADLs in a specific context (a test kitchen). However, the same system should be flexible enough to be eventually fitted in individual homes, to promote the completion of personally meaningful activities in naturalistic contexts – ideally to inform caregiving strategies, and enable the continuous and fine-grained personalisation of other assistive technology.

II. RELATED WORK

Ecological Momentary Assessment (EMA) is one example of a technology-enabled system designed to reduce dependency from clinicians time. With EMA, patients use phones or tablets to update the clinicians on a regular basis by completing a questionnaire remotely. Technologies like EMA mitigate biases in the clinical setting, and aid in tracking the changes of cognitive decline over time [1]. However, they require active and diligent patient participation. Crucially, like their pen and paper counterparts, their results may not correlate to ADL performance.

This issue of adherence to periodic assessments is a significant barrier to the wider adoption of this type of cognitive assessment. The MARIO project [5] has shown an increase in the adherence by using a physical robot to prompt questions to the users. Demonstrating promising results in terms of increased acceptance and potential benefits of robots employed as part of dementia care practice [5].

Virtual Reality (VR) systems have also been proposed, as they can be used to place the users in easy-to-control virtual scenarios. A particularly interesting example is the VR Kitchen assessment [6], upon which our approach is based. The assessment is itself informed by the Kitchen Task (KT) [8], in which people are assisted by an expert while performing pre-defined tasks in a test kitchen. KT involves evaluating the cognitive processes that affect task performance and recording the level of cognitive support necessary for successful completion of meaningful tasks. VR settings in general make it very easy to create repeatable conditions and record users' behaviours. However, they have been found to poorly reflect real world performance for the user group, especially because the sensory perceptions and motor skills exercised in VR are significantly different from those at play in a real environment [2].

III. SYSTEM ARCHITECTURE:

Figure 1 illustrates the modular architecture of our system, and highlights the research questions we are investigating for its realisation, namely:




Fig. 1. Proposed architecture for automatic cognitive assessment

C - Context Awareness. How to track primitive actions (“pick spoon”, “move mug to worktop”...) and complex activities (“make tea”, “make sandwich”...) by interpreting data captured from heterogeneous sensors in actual environments? Different settings may warrant different approaches. For instance, it would be possible for a walk-in-testing facility to sense human-object interactions by attaching wireless sensor devices to all relevant objects, and also install fixed cameras or employ a robot assistant using optical and depth cameras to track human actions. The same approach would be too costly, unpractical and/or undesirable for a home setup, where solutions should employ non-intrusive and simpler sensors, and also adopt machine learning techniques, such as Hidden-Markov Models (HMMs) to infer events from incomplete or incorrect observations.

A - Analysis. How to decide what (if any) assistance needs to be provided, and consequently infer the cognitive health of the user? The component will employ a user model - essentially a model of users’ cognitive capabilities - and planning techniques to infer possible user’s errors by analysing the differences between expected and observed steps, before deciding whether to assist the user. This may involve providing clues or issuing prompts, but also requesting some information from the user, especially in cases when the system does not have enough confidence on the exact problem this is experiencing, but also to leverage information that can be extracted from the verbalisation aspect. The component will then need to update the user model, by weighting the type of errors, but also the type of assistance that was necessary for the user to complete each task.

I - Interaction. How to interact more effectively with the users? Rather than pre-recorded audio/video feedback, we are developing a conversational agent to support multiple, bi-directional conversation flows, and plan to investigate the effectiveness of both a voice assistant and a social robotic embodiment. The latter will cover different roles: An observer and assistant, providing clues, prompts and reminders, but also to motivate and help the user to carry out and complete the test. We will investigate suitable affective agent architectures, to leverage expressive moods and emotions as an integral part of social interaction, and Reinforcement Learning (RL) tech-

niques, such as Partially Observable Markov Decision Process (POMDP), to tailor interaction style to each individual.

IV. METHODOLOGY AND EVALUATION STRATEGY

Testbed - The system described in this paper is being developed at the Robotic Assisted Living Testbed (RALT) hosted at Heriot-Watt University, Edinburgh Centre for Robotics. The testbed is a 60m² fully sensorised smart home hosting a number of assistive technologies and domestic robots.

Co-Design - Workshops and focus-groups will be used to bring together stakeholders (patients, carers, healthcare professional) with the research team, to inform the design of successive prototypes. A series of user studies will be carried out to evaluate their technical effectiveness, feasibility, reliability, acceptability and usability.

Prototype - An initial prototype has been developed using the smart kitchen and the humanoid robot Pepper (from Soft-bank Robotics) in the testbed. The OpenHAB smart home middleware is used to collect information from sensors installed in the smart kitchen. These include occupancy and magnetic switch sensors to detect users’ presence and opening/closing of drawers, and an energy monitor, to detect the use of kitchen appliances (coffee machine, kettle, toaster). The main application (a finite state machine where state transitions are triggered by binary sensor events) is executed on top of the Robotic Operating System (ROS). The robot assumes the role of the occupational therapist, instructing and assisting the user in carrying out the assessment

Next steps - Work is ongoing to record a dataset to enable further development of the context awareness and analysis components. The dataset will include sensor data from both the smart home and the robots’ 3D and optical cameras, together with video footage of volunteers carrying out multiple assessment sessions. A first user study is also being planned, to gather information on users’ experience and interaction requirements, including a body of training examples to develop the conversational agent.

REFERENCES

- [1] M Schmitter-Edgecombe et al. “Technologies for health assessment, promotion, and assistance: Focus on gerontechnology”. Positive Neuropsychology Springer New York, 2013, pp. 143–160.
- [2] Carolyn Parsey et al. Applications of Technology in Neuropsychological Assessment. The Clinical neuropsychologist 27 (2013).
- [3] Ismail Zahinoor et al. Brief cognitive screening instruments: an update. International Journal of Geriatric Psychiatry 25.2(2010), pp. 111–120
- [4] Meiland F et al. Technologies to Support Community-Dwelling Persons With Dementia: A Position Paper on Issues Regarding Development, Usability, Effectiveness and Cost-Effectiveness, Deployment and Ethics. JMIR Rehabilitation Assist Technology 4(1)e1 (2017).
- [5] Grazia D’Onofrio et al. MARIO Project: Experimentation in the Hospital Setting. In: Ambient Assisted Living. Springer International Publishing, 2019, pp. 289–303.
- [6] Heriot-Watt Neurophysiology lab, VRAIS kitchen, accessed 05/08/2018, <https://www.neurophysiologylab.hw.ac.uk/vrais-kitchen.html>
- [7] E. M. D. Jean-Baptiste and A. Mihailidis, “Analysis and comparison of two task models in a partially observable Markov decision process based assistive system,” 2017 IEEE 4th ISCOMI, pp. 183-187.
- [8] Carolyn Baum, Dorothy F. Edwards; Cognitive Performance in Senile Dementia of the Alzheimer’s Type: The Kitchen Task Assessment. Am J Occup Ther 1993;47(5):431–436

Robotic ignition systems for oil fields

Mohammed Lami

Materials and Engineering Research
Institute
Sheffield Hallam University
Sheffield, U.K.
b8018683@my.shu.ac.uk

Lyuba Alboul

Materials and Engineering Research
Institute
Sheffield Hallam University
Sheffield, U.K.
l.alboul@shu.ac.uk

Issa Ahmed Abed

Engineering Technical College –
Basrah
Southern Technical University, Basrah,
Iraq
Issaahmedabed@stu.edu.iq

<https://doi.org/10.31256/Ln9Sf6F>

Abstract—In the oil extraction industry, igniting the flare stacks is an essential operation. Oil sites have two kinds of flares, ground flares and flares that installed on towers. The ignition systems generate electrical sparks to burn the gases blowing out of the flares. Due to the permanent high operating temperature and the need for special thermal isolation, classical igniters have low reliability and high cost. In this work, two novel ignition systems have been implemented, the first is the robotic ignition system for ground flares, it utilises a mobile robot which moves toward the flare, avoiding the obstacles in its way and stops after detecting the gas, then it starts igniting the flare before heading to a safe point with no gas and low temperature. The second solution is the automated ignition system to light up the flares on the towers, which is a car that moves on a rail vertically, and begins igniting once it arrives at the tip of the tower, then it comes back to its starting point. As the igniters in both suggested systems are movable, so the system will be exposed to the heat generated by the flame within a very short time, this new feature increases the reliability of the igniter and reduces the complexity and the cost of the system.

Keywords—igniters, microcontrollers, mobile robots, stack flares

I. INTRODUCTION

The work focusses on the oil industry, and particularly on the oil fields where natural gases are occasionally produced during the oil extraction process [1]. These gases are transported to special torches called flare stacks to be burned. Different methods are used to achieve this task, but the most popular method uses the electrical spark to ignite the flares [2]. These sparks are generated by special electrical devices, and then they are delivered to an ignition probe through an electrical cable that is located close to the flare [3]. The parts of the ignition system that are fixed near to the flare are thermally protected using thermal insulators which adds more complexity to the system. However, these systems usually have a short lifetime because of high operating temperature, hence, this reduces the system reliability [4]. The main aim of this research is to implement a new ignition system with a movable igniter, so it can be positioned close to the gas cloud around the flare when it is required to perform the ignition task, and then returns to a safe point away from the high temperature. Consequently, the system complexity will be reduced as no more thermal insulators will be needed, besides, the reliability of the system will be raised because the exposure time of the igniter to the flame will be very short, virtually in terms of seconds. For ground flares, the robotic ignition system has been introduced, which is a robotic car equipped with motorized wheels and different kinds of sensors. With the aid of a special microcontroller and a suitable algorithm, it can move toward the flare stack avoiding the obstacles; the robot uses a gas sensor to measure the density in real-time of the natural gas to determine whether the goal is reached. Once the measured value of the density equals or more than a threshold value stored in the program, the robot decides to stop and start igniting the flare for a while. A fire sensor is utilised detect the flame generated after performing

ignition operation, this helps to verify that the ignition action is completed. Consequently, the robot changes its direction to move away from the flare and searches for a safe point, and of course, it also keeps avoiding the obstacles until the final position is reached. For the flares on towers, an automated ignition system has been suggested, it employs a car which carries the igniter device to be moved on a rail vertically toward the flare and ignite it. The automated car also has a fire sensor to verify the ignition action just like in the robotic ignition system. Both ignition systems mentioned above are designed to work with a secondary robot synchronously, the secondary robot is responsible for opening a gas valve that controls the flowing of the gas from the source (oil well) to the flare stack. The robotic or the automated ignition system firstly orders the secondary robot to open the valve, to ensure the gas is blowing out of the flare before starting the ignition mission.

II. MATLAB/SIMULATION

The artificial potential field (APF) approach has been employed in both systems [5]; in this case, the robot will be guided in the field to the goal (the flare). The mathematical representations for the attractive \vec{F}_{att} and repulsive \vec{F}_{rep} forces are shown in equations (1) and (2) which are retrieved from [5] and [6] respectively.

$$\vec{F}_{att} = K_g \cdot (q - q_{goal}) \quad (1)$$

$$\vec{F}_{rep} = \frac{1}{K_{rep}} e^{-\frac{dist(rob,obs)}{K_{rep}}} \quad (2)$$

The mathematical formulas of robot displacement due to both forces used in MATLAB simulation are given in equations (3) and (4).




Fig 1: MATLAB simulation of the robotic ignition System

$$dx_{rep} = (robot - Obs_n) exp \frac{-norm(robot - Obs_n)}{K_{obs_n}} \quad (3)$$




Fig 2: MATLAB simulation of the automated ignition System

Where n is the number of obstacles, $Kobs_n$ is a scaling factor, robot and Obs_n are the positions for the robot and the obstacle.

$$dx_{att} = K_{goal} \frac{(Goal-robot)}{norm(Goal-pos)} \quad (4)$$

Where K_{goal} is a scaling factor, $robot$ and $Goal$ are the positions for the robot and the goal.

Both systems were successfully tested and simulated. The simulation results of the robotic and automated ignition systems are shown in Fig. 1 and Fig. 2.

III. IMPLEMENTATION

The robotic ignition system (the robotic car) has been prototyped as shown in Fig 5, using the embedded system shown in Fig. 3. The two microcontrollers control robot's motors, obstacle avoidance and ignition action based on the real-time data obtained from the equipped sensors. The automated ignition system (the automated car on the rail) has been prototyped as shown in Fig. 5, using the embedded system shown in Fig. 4. The PLC control car's motors and ignition action based on the real-time data obtained from the equipped sensors.

IV. RESULTS, DISCUSSION AND CONCLUSION

To sum up, the output of this work is the implementation of the two ignition systems. The first one is the robotic ignition system which is designed to serve the ground flares in oil fields. Secondly, the automated ignition system has been introduced as a solution for the flare stacks on the towers in oil fields. In the classical ignition systems, the direct exposure of some components to the flame for long periods reduces the operating lifetime of the ignition device, the heat usually damages these components and then they need to be replaced. Moreover, in some cases, there is no opportunity for service between major shutdown and this leads to a reduction in system reliability [4]. Due to the high operating temperature, the igniter and other parts of the system are protected using special insulators which bear temperature up to 1100 C° [7], and this increases the complexity and the cost of the system dramatically. The unique property of the suggested ignition systems in this work is that the end terminal (the part of the system that responsible for delivering the spark to the flare) is not fixed at the position of the flare, but it is movable, it is utilised only when it is needed to ignite the gas cloud, so it will be exposed to the fire within a short time in terms of few seconds during ignition operation. Hence, this feature could

reduce system complexity and cost as the igniter no longer need a high thermal isolation. Since the exposure time of the igniter to the hazardous temperature becomes short, the igniters will be more reliable with long lifetime. However, although the exposure time becomes short, the external parts of the robot and the automated car should be chosen to bear the high temperature within ignition operation. In addition, the time needed to accomplish the ignition operation by the two suggested systems will be longer as compared with current ignition systems, it depends on the speed of the robot, number of obstacles in its way and how much the ground is sloped. Furthermore, the algorithm used in implementing the robotic ignition system assumes that the robot always succeed in igniting the flare at the first time, but practically it may fail if the amount of the gas is low even if its density is high, a future work, an extended algorithm is needed to execute several attempts that may be required to complete the task.

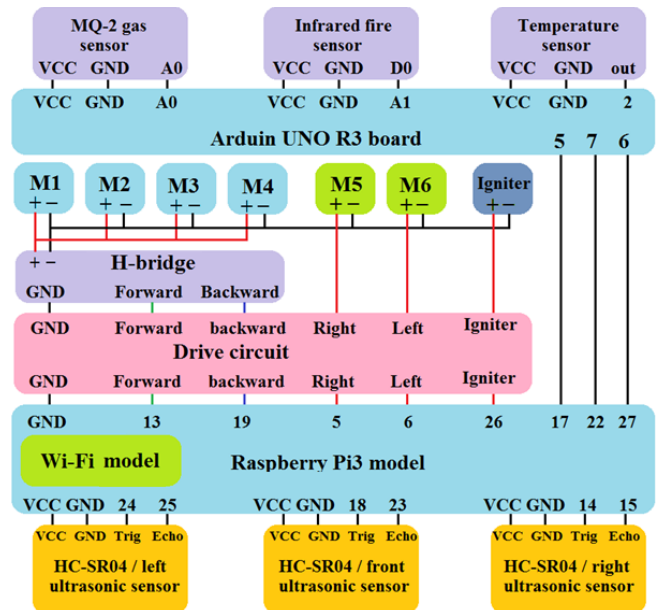


Fig 3: The detailed wiring diagram for the robotic ignition system




Fig 4: The detailed wiring diagram for the robotic automated system




Fig 5: The prototyped robot and the automated car

References

- [1] Eman, "GAS FLARING IN INDUSTRY: AN OVERVIEW," Petroleum & Coal ISSN 1337-7027, December 2015.
- [2] Profire Energy, "FLARE STACK IGNITER," 2015. [Online]. Available: https://www.profireenergy.com/product_documents/FSL_Manual_and_Install_Guide_0.1_print.pdf.
- [3] Honeywell, "Two-rod Pilot Burner/Igniter-Sensor," 2017. [Online]. Available: <https://customer.resideo.com/resources/Techlit/TechLitDocuments/69-0000s/69-0768.pdf>.
- [4] Argo, "Electronic Spark Ignition," 2019. [Online]. Available: <http://www.argoflares.com/research/introduction/pilot-burners-ignition-systems/electronic-spark-ignition/>.
- [5] R. Siegwart, I. Nourbakhsh and D. Scaramuzza, Introduction to autonomous mobile robots, vol. 2nd ed., 2011.
- [6] D. S. Neculescu, E. Pruner and e. al., "Self-organizing robot formations using velocity potential fields commands for material transfer," 2014.
- [7] Kraus, Kurt, Richmond, Dusty, Jennings, Jay and Bietto, Stefano. USA Patent 8986000, 2015.

Towards the development of a deposition technology for an automated rail repair system

D. De Becker
Wolfson School

Loughborough University
Loughborough, Leicestershire
D.De-Becker@lboro.ac.uk

J. Dobrzanski
Wolfson School

Loughborough University
Loughborough, Leicestershire
J.Dobrzanski@lboro.ac.uk

Y. M. Goh
Wolfson School

Loughborough University
Loughborough, Leicestershire
Y.Goh@lboro.ac.uk

L. Justham
Wolfson School

Loughborough University
Loughborough, Leicestershire
L.Justham@lboro.ac.uk

<https://doi.org/10.31256/Vz2Jt4I>

Abstract—The work presented in this paper explored the use of a laser line scanner to generate robotic deposition paths for the repair portion of an automated rail repair system. Currently surface defects cost the UK around £4 million per annum [1], with little traceability being available throughout the repair process. This paper proposes a robotic repair system primarily focussed on the development of the deposition system. The deposition system utilised two different deposition strategies, the first extracted the weld prep from the point cloud to generate the deposition paths for the robot and the second measured the height of the previously deposited material and adjusted the generated path. This paper focusses on the use of the two algorithms and the testing completed on a representative geometry, utilising a caulking gun as a reusable material replacement for the additive welding system.

Index Terms—Path generation, Rail repair, Hybrid manufacturing, Robotic welding

I. INTRODUCTION

The UK rail infrastructure requires constant repair and maintenance. A key contributor to the cost incurred are maintenance costs in the form of rail defects. The work presented in this paper concentrated on the repair of surface defects found on the rail. One such defect is named a squat defect and the UK estimates the cost of squat repair on plain rail to be around £4 million per annum [1]. An example of a squat defect, taken at Great Central Railway line in Quorn station, Leicestershire, UK, can be seen in Fig 1.



Fig. 1. Squat defect taken at Quorn, Leicestershire, UK railway station

Considerable work has been done on automating the identification and registration of both surface and subsurface defects, however, little work has been done on automating the repair




Fig. 2. Schematic of the proposed solution for the rail repair system

procedure once the defect has been found [2] [3] [4] [5]. In the UK this initial defect identification and detection is done by the New Measurement Train (NMT) run by Network Rail, coupled with its Plain Line Pattern Recognition (PLPR) system [6]. The current state of the art in the field of defect repair is the ARR (Automated Rail Repair) system developed during the Shift2Rail project [6]. This system uses a predetermined path and submerged arc welding to deposit material onto an already existing weld prep. This system, although requiring a much lower preheat, cannot achieve a full automated repair process in its current state. Therefore, the research presented in this paper has suggested a complete rail repair system using an industrial robotic arm, as shown in Fig 2.

The 2019 report produced by Network rail for Welding process and technology developments states that weld failure in Wales is approximately 15%, with the key driver being the lack of automation and traceability within the welding process [9]. Therefore, the purpose of this paper is to show the development, testing and validation of an adaptive deposition system. The system can be used to add material onto an already existing weld prep. This weld prep would be generated by the subtractive subsystem within a fully automated rail repair system [10].

II. ADAPTIVE WELDING ALGORITHMS GENERATION

Fig 3 shows the weld prep, used for this work which was taken from the British Standard for Restoration of Rails by Electric Arc Welding [11]. The weld prep dimensions used for the trials were; $L = 100\text{mm}$, $S = 25\text{mm}$ and $D = 3\text{mm}$. The overall width was 70mm, which was equal to the width of the rail head used.




Fig. 3. Schematic of the weld prep used to generate the welding path taken from BS 15594:2009 [11]

During the adaptive welding process, there were two algorithms working together in order to create a robust additive welding process. The initial weld path generation algorithm was driven by the scan data generated by the laser line scanner as shown in Fig 4a. This algorithm identified the weld prep surface within the scan data as shown in Fig 4b, applying a translation technique for deposition optimisation. The technique used to extract this weld prep surface utilised the eigenvalue of the covariance matrix to identify local changes in the data. The covariance matrix is a statistical tool where the largest eigenvector indicates the direction of the data and the second largest eigenvector indicates the second largest data trend. The magnitude of these eigenvectors are defined as the eigenvalue. In this case, the covariance matrix was calculated over a 3mm neighbourhood and a ratio of the two largest eigenvalues of the covariance matrix was calculated. This ratio stayed consistent either side of the weld prep and a large change in this ratio was observed when the neighbourhood partially included data from the weld prep. Once the weld prep was extracted, it was possible to use the points adjacent to the weld prep to create a linear interpolation between the two sides. This yielded a surface which was used to recalculate the Z values for the points within the machined area. This ensured that the repair not only filled the area which had been extracted, but also accounted for any wear which would have occurred over time. The outcome of this linear interpolation and remapping of the points can be seen in Fig 4c, which shows how the final surface should appear once both the additive system, described in this paper, and final finishing operation was completed. Finally, the build plane was extracted and the fill strategy was implemented, an example of the first layer can be seen in Fig 4d. This fill algorithm was based on the user inputs, including the estimated bead width and height along with the step oversize. The fill algorithm worked by raising the build plane based on the estimated bead height and then intersecting the build plane with the point cloud. This generated a layering effect where the subsequent layer deposition height was adjusted by the second algorithm, to match the actual bead height deposited on the measured layer. Such a scan is seen in Fig 5. This allowed, for any

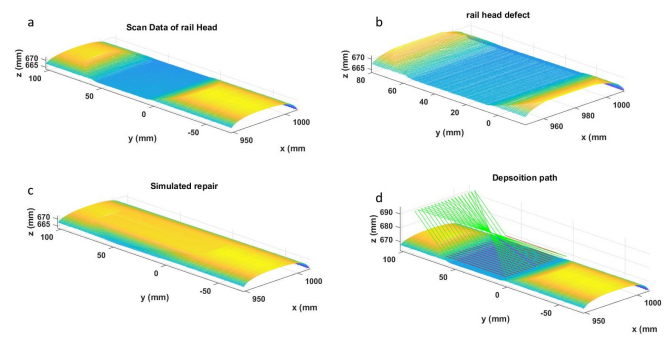


Fig. 4. Scans showing path generation sequence. (a)Original scan data. (b)The area of interest. (c)Repaired rail. (d)Deposition path for the first layer

difference between the bead height input by the user and the real-world bead height to be corrected for. This algorithm works on the highest 20% of Z values within the deposition area and fits a plane to those points. This plane was then used to adjust the next layer's Z values in order to adjust the weld path.

These adaptive welding algorithms undertook preliminary testing using a caulking gun, with caulk used as an additive test replacement to the metal of the welding process. This allows multiple attempts on the same piece without the need for part re-manufacturing. Both algorithms can be viewed functioning, with their outcomes in Fig 5, which shows the penultimate layer deposited. This Figure also shows the height adjustment algorithm working.




Fig. 5. Scan data showing the path being updated to avoid collisions

III. CONCLUSION

The initial testing proved that it was possible to extract the weld prep from a rail head using the covariance matrix. Moving forward this path generation will be tested on a representative flat plate in order to fine tune the welding parameters including the amperage, voltage and the pre-heat temperatures. Once these parameters have been tuned, the path generation will be tested on the representative geometry using the welding equipment. The final step will be to re-validate the testing carried out on the flat plate on a piece of real rail. This will, in turn, provide valuable information and validation of the deposition subsystem of this repair system.

ACKNOWLEDGMENT

The author would like to thank the Intelligent Automation Centre within the Wolfson School of MEME at Loughborough University and the Innovate UK funded CONFIGURE project for the use of the equipment and the continued support in generating this work.

REFERENCES

- [1] Network Rail, "Challenge statement Rail head squats," tech. rep., Network Rail, London, 2015.
- [2] C. Tastimur, O. Yaman, M. Karakose, and E. A. Ö, "A Real Time Interface for Vision Inspection of Rail Components and Surface in Railways," in *IDAP 2017 - International Artificial Intelligence and Data Processing Symposium*, 2017.
- [3] M. Molodova, Z. Li, A. Núñez, and R. Dollevoet, "Automatic Detection of Squats in Railway Infrastructure," *IEEE TRANSACTIONS ON INTELLIGENT TRANSPORTATION SYSTEMS*, vol. 15, no. 5, pp. 1980–1990, 2014.
- [4] X. Yin, X. Wei, and L. Jia, "Detection of Railway Track Squats by Using Bogie Acceleration Measurement," in *Proceedings of the 34th Chinese Control Conference*, pp. 6297–6302, 2015.
- [5] J. Ye, E. Stewart, and C. Roberts, "Use of a 3D model to improve the performance of laser-based railway track inspection," *Proceedings of the Institution of Mechanical Engineers, Part F: Journal of Rail and Rapid Transit*, vol. 233, no. 3, pp. 337–355, 2019.
- [6] E. Kallander, "In2Rail Deliverable D3.1 Investigation of Repair Methods and Welding Techniques," tech. rep., Horizon 2020, 2017.
- [7] F. Javier De La Calle Herrero, D. F. García, and R. Usamentiaga, "Inspection System for Rail Surfaces Using Differential Images," *IEEE Transactions on Industry Applications*, vol. 54, no. 5, pp. 4948–4957, 2018.
- [8] Z. Xiong, Q. Li, Q. Mao, and Q. Zou, "A 3D laser profiling system for rail surface defect detection," *Sensors (Switzerland)*, vol. 17, no. 8, pp. 1–19, 2017.
- [9] Network Rail, "Welding Process and Technology Development," tech. rep., Network Rail, 2019.
- [10] D. De Becker, J. Dobranzski, J. Laura, and I. Y. M. Goh, "A laser scanner based approach for identifying rail surface squat defects (Under review)," *Proceedings of the Institution of Mechanical Engineers, Part F: Journal of Rail and Rapid Transit*, p. 107754631982824.
- [11] BS EN 15594, "Railway applications - Track - Restoration of rails by electric arc welding," tech. rep., 2009.

An Automatic Design Tool for Fluid Elastomer Actuators

Olivia Bridgewater-Smith, Gabriele Maurizi, Sebastiano Fichera,
David Marquez-Gamez, Andrew I. Cooper, Paolo Paoletti

<https://doi.org/10.31256/Xu9Da6Q>

Abstract—Soft robotic actuators are a very promising technology to enable use of robotic manipulators in scenarios that are inaccessible to traditional robots. However, their design and fabrication is a laborious process, especially for users with little knowledge of CAD software and 3D printing. The skills and time necessary for making the moulds used to create such actuators, leaves the process open for human errors and design variations, making accurate and repeatable experimental testing difficult to achieve. To reach a better understanding of this new technology, extensive and detailed experimental work should be undertaken, but this is currently hindered by the time-consuming design process. The design software presented in this paper provides the soft robotic community with a user-friendly design tool for generating 3D printed moulds that the users can customise to their needs. The tool aims to simplify the design process for soft robotics and to also make this technology accessible to users without extensive engineering background.

I. INTRODUCTION

Soft robotics is an emerging field within the robotics community. Their bio-inspired nature allows them to use them for a wide variety of scenarios. Novel soft robots design have been developed based on animal and human biomechanics; for example, an octopus project has shown that robots using soft materials can perform tasks that would not be achievable with ‘hard’ materials [1]. Soft robots can be designed to perform more than one task, so that unstructured or bespoke tasks in chemistry and biology labs can benefit from their adaptability and their ability to bridge the gap between humans and conventional ‘hard’ robotics [2].




Fig. 1. Soft Robot Actuating

OBS, SF, DMG, AC are part of the Leverhulme Research Centre for Functional Materials Discovery, Material Innovation Factory, University of Liverpool, Liverpool L69 7ZD, UK. OBS, PP and SF are part of the School of Engineering, University of Liverpool, Liverpool L69 3GH, UK. GM is part of the Dipartimento di Ingegneria dell’Informazione, University of Florence, Firenze I-50139, Italy. sgobridg@liverpool.ac.uk gabriele.maurizi@stud.unifi.it seba84@liverpool.ac.uk David.Marquez-Gamez@liverpool.ac.uk aicooper@liverpool.ac.uk P.Paoletti@liverpool.ac.uk

II. CORE COMPONENTS AND FABRICATION

The most common design for soft actuators is Fluid Elastomer Actuators (FEA). Such actuators are composed of a series of inflatable chambers mounted on top of a flexible substrate. As each chamber inflates, air exerts a pressure on the chamber’s walls. Given that the inner walls of the chambers are thinner than the other walls, they bulge outwards and make contact with each other, thus providing a force that causes the actuator to bend. A typical example of FEA undergoing actuation is shown in Fig. 1.

The typical fabrication process for producing a FEA consists of three steps: mould design, mould manufacturing, and moulding process [2]. When designing the mould, the actuator geometry needs to be determined according to the desired properties and mechanical response of the actuator itself. Typically, the mould is designed on 3D CAD software and it is sectioned into components. The main body consisting of a union of parts A and B (see Fig. 2) and a tray is used to enclose the bottom of the actuator. The main body and tray are manufactured separately and then joined together using a thin layer of silicone.




Fig. 2. Example of 3D CAD of a mould produced the automatic design tool: Part A (left) and Part B (right).

There are four essential components that need to be designed to create such moulds, as shown in Fig. 2: chambers (P4), air channel (P5), supports (P3/6) and the base [3]. The air channel runs parallel through the actuator and it is responsible for allowing air to pass through into each chamber. The inner and outer supports are placed to aid the joining process of the main body and the tray, ensuring the air channel and chambers are not impeded by the uncured silicone solidifying. The base encapsulates the chambers and provides a flexible surface to induce bending.

III. AUTOMATIC DESIGN TOOL

A software tool was developed to automate the design and fabrication of the mould process. The software aim was to

provide well-designed moulds based on a limited number of user-defined inputs and to prevent the occurrence of human errors and streamline the design process. The proposed tool uses the 3D model builder OpenSCAD and a bespoke Python-based graphical user interface to automate the design.

The main GUI of the design tool consists of input boxes allowing the user to set: chamber width, length, depth, spacing, number of chambers and tray height (see Fig. 3). The code uses the number of chambers and chamber dimension inputs to calculate the length and width of the whole actuator, and it design layers P1/2 using these values. The chambers are then placed along the length of the actuator using arithmetic progression, alongside a central air channel (P5) connecting all of the chambers. The air channel and inner supports height and width are dependent on a proportional ratio applied to the chamber dimensions, empirically tuned based on testing on a wide range of actuators. The complete Part A mould dimensions, calculated by the code, are then used as inputs to generate Part B and the tray. The tool then outputs three STL and SCAD files in the chosen location which the user can directly use to 3D print their moulds.




Fig. 3. GUI to select the actuators geometry.

IV. EXPERIMENTAL RESULTS

To test the capability of the code, various designs were produced with varying chamber widths, lengths, height and number. Such moulds were then 3D printed using an Ultimaker S5 printer with PLA filament, and then used to fabricate soft actuators. Finally, the soft actuators were tested to show that the automated design tool can create functional actuators (see Fig. 1) [4]. The ability to design moulds and actuators at a fraction of the time required by manual design allowed the creation of a large amount of actuators to perform a parametric study on how geometry affects the actuator's mechanical response [5]. For the purpose of this paper, the software was used to prove how changing the width of the actuator chamber affects the tip force. This is complementary to the linear relationship between height and tip force reported in [3] [6]. Taken together, these results can form the basis for optimising FEA design based on desired mechanical response.

The initial testing was conducted by clamping the actuator and measuring its tip force at a range of pressures. Experiments were conducted on three actuators of the same length and chamber number. The first two actuators were made from silicone M6401 (Bentley Materials, UK), the second one had an increased chamber width of 5mm. The third actuator was made from Dragon Skin (Bentley Materials, UK) and matched

the dimensions of the first. As shown in Fig. 4, increasing the actuator width almost doubles the resulting tip force. This suggests that using materials with lower Young's moduli allows larger variations of tip forces for limited changes in input pressures.




Fig. 4. Force vs. Pressure for three actuators

V. CONCLUSION

In this paper an automatic design tool for soft robotic actuators is presented. The availability of this tool will allow researchers to streamline the design of such actuators and to perform parametric studies on the influence of geometry on the actuator response. As an example, preliminary results regarding the relationship between chamber width and tip force are presented. The modular and open-source nature of the presented tool will enable encompassing a wider range of soft actuators in the future. The tool is available at <http://www.liv.ac.uk/paoletti/public/softdesigntool.zip>

To further validate the tool, analytical results will be linked with the design tool to create an input for the desired tip force range of the actuator. The tool will generate the optimal internal geometry to achieve the required force so that actuator design can be optimised based on the applications they will be used for. A variety of moulds will be produced to experimentally verify the tip force output.

REFERENCES

- [1] C. Laschi, M. Cianchetti, B. Mazzolai, L. Margheri, M. Follador, and P. Dario, "Soft robot arm inspired by the octopus," *Advanced Robotics*, vol. 26, no. 7, p. 709–727, 2012.
- [2] D. Rus and M. T. Tolley, "Design, fabrication and control of soft robots," *Nature*, vol. 521, no. 7553, p. 467–475, 2015.
- [3] P. Polygerinos, Z. Wang, J. T. B. Overvelde, K. C. Galloway, R. J. Wood, K. Bertoldi, and C. J. Walsh, "Modeling of soft fiber-reinforced bending actuators," *IEEE Transactions on Robotics*, vol. 31, no. 3, p. 778–789, 2015.
- [4] H. Zhao, Y. Li, A. Elsamadisi, and R. Shepherd, "Scalable manufacturing of high force wearable soft actuators," *Extreme Mechanics Letters*, vol. 3, p. 89–104, 2015.
- [5] L. Ding, N. Dai, X. Mu, S. Xie, X. Fan, D. Li, and X. Cheng, "Design of soft multi-material pneumatic actuators based on principal strain field," *Materials & Design*, vol. 182, p. 108000, 2019.
- [6] Z. Wang, P. Polygerinos, J. T. B. Overvelde, K. C. Galloway, K. Bertoldi, and C. J. Walsh, "Interaction forces of soft fiber reinforced bending actuators," *IEEE/ASME Transactions on Mechatronics*, vol. 22, no. 2, p. 717–727, 2017.

Adaptive Manipulator Control using Active Inference with Precision Learning

Mohamed Baioumy, Matias Mattamala, Paul Duckworth, Bruno Lacerda, and Nick Hawes

Oxford Robotics Institute

{mohamed, matias, pduckworth, bruno, nickh}@robots.ox.ac.uk

<https://doi.org/10.31256/Gz8Eu7C>

Abstract—Active inference provides a framework for decision-making where the optimization is achieved by minimizing free-energy. Previous work has used this framework for control and state-estimation of a robotic manipulator. This required manual definition of precision matrices which serve as controller gains. This paper provides an implementation for control and state-estimation where the precision matrices are tuned during execution-time (precision learning). Learning the precision matrices means automatically adjusting the controller’s gains which decreases oscillations and overshoot.

Supplementary Material

Code and further material is available at: https://github.com/MoBaioumy/active_inference_panda_paper.

Video: <https://youtu.be/li1IglLt0Xk>

I. INTRODUCTION

Modelling all time-varying dynamics for a robotic manipulator a priori is infeasible. Therefore, intelligent robotic manipulators require adaptive behaviour, e.g. to reject disturbances or to handle objects of unknown masses [1].

Recent approaches in robotics have taken inspiration from active inference [2], a theory of the brain prominent in neuroscience. Active inference provides a framework for understanding decision-making of biological agents. Under the active inference framework, optimal behavior arises from minimising variational free-energy: a measure of the fit between an internal model and (past) sensory observations [2]. Additionally, agents take actions that allow reaching preferred future observations specified a priori.

In [3], an active inference controller (AIC) for joint space control of robotic manipulators is presented, which outperforms the state-of-the-art Model Reference Adaptive Control (MRAC) [4]. This approach performs both state-estimation and control, and avoids scalability issues by requiring a fixed number of parameters, such as precision matrices (inverse covariance matrix). These precision matrices act as controller gains and thus affect performance properties such as oscillations, overshoot and rise-time.

In this paper, we present an approach for precision learning, i.e. learning the precision matrix. We show that by having an adaptation step to learn the precision matrices in execution time, which is derived from the same active inference principle, we are able to automatically adjust the controller’s gains, henceforth decreasing oscillations and overshoot.

II. METHOD

Active Inference considers an agent in a dynamic environment that receives an observation \mathbf{o} about a state \mathbf{s} . The agent then infers the posterior $p(\mathbf{s}|\mathbf{o})$ given a model of the agent’s world. Instead of exactly calculating $p(\mathbf{s}|\mathbf{o})$, which could be computationally expensive, the agents approximate $p(\mathbf{s}|\mathbf{o})$ with a ‘variational distribution’ $Q(\mathbf{s})$ which we can define to have a standard form (Gaussian for instance). The goal is then to minimize the difference between the two distributions. This can be computed using the KL-divergence [5]:

$$KL(Q(\mathbf{s})||p(\mathbf{s}|\mathbf{o})) = \int Q(\mathbf{s}) \ln \frac{Q(\mathbf{s})}{p(\mathbf{s}, \mathbf{o})} d\mathbf{s} + \ln p(\mathbf{o}) \quad (1)$$

$$= F + \ln p(\mathbf{o}).$$

The quantity F is referred to as the (variational) free-energy -or Evidence lower bound- and minimizing F minimizes the KL-divergence. If we choose $Q(\mathbf{s})$ to be a Gaussian distribution with mean $\boldsymbol{\mu}$, and utilize the Laplace approximation [6], the free-energy expression simplifies to:

$$F \approx -\ln p(\boldsymbol{\mu}, \mathbf{o}). \quad (2)$$

Now the expression for variational free-energy is solely dependent on one parameter, $\boldsymbol{\mu}$, which is referred to as the ‘belief state’. The objective is to find $\boldsymbol{\mu}$ which minimizes F ; this results in the agent finding the best estimate of its state.

A. Problem statement

In this work we consider low-level joint control for a robotic manipulator. The state of the robot is given by its joint position, velocities, $\tilde{\boldsymbol{\mu}} = [\boldsymbol{\mu}, \boldsymbol{\mu}']$. Note that in the Active Inference framework the state given by position and higher derivatives is also known as *generalised motions* [7]. We also consider observations given by joint encoders, $\tilde{\mathbf{o}} = [\mathbf{o}, \mathbf{o}']$, which are the main source of information about the state.

As in [3], we aim to obtain a control law for the state $\tilde{\boldsymbol{\mu}}$ from the Active Inference framework as well as the robot joint state $\tilde{\boldsymbol{\mu}}$ given the observations $\tilde{\mathbf{o}}$. However, we also derive learning rules for the parameters of the distributions (precision matrices), which allow us to adapt the controller online.

B. Observation model and state transition model

Following [3], the joint probability from Equation (2) can be written as:

$$p(\tilde{\mathbf{o}}, \tilde{\boldsymbol{\mu}}) = p(\tilde{\mathbf{o}}|\tilde{\boldsymbol{\mu}})p(\tilde{\boldsymbol{\mu}}) = \underbrace{p(\tilde{\mathbf{o}}|\boldsymbol{\mu})}_{\text{Observation model}} \underbrace{p(\mathbf{o}'|\boldsymbol{\mu}')}_{\text{Transition model}} p(\boldsymbol{\mu}'|\boldsymbol{\mu})p(\boldsymbol{\mu}''|\boldsymbol{\mu}'), \quad (3)$$




Fig. 1: Results comparing the active inference controller with and without precision learning.

where $p(\mathbf{o}|\boldsymbol{\mu})$ is the probability of receiving an observation \mathbf{o} while in (belief) state $\boldsymbol{\mu}$, and $p(\boldsymbol{\mu}'|\boldsymbol{\mu})$ is the state transition model (also referred to as dynamic model or generative model). The state transition model predicts the state evolution given the current state. These distributions are assumed Gaussian according to:

$$\begin{aligned} p(\mathbf{o}|\boldsymbol{\mu}) &= \mathcal{N}(\boldsymbol{\mu}, \Sigma_o), & p(\boldsymbol{\mu}'|\boldsymbol{\mu}) &= \mathcal{N}(\boldsymbol{\mu}', \Sigma_{o'}), \\ p(\boldsymbol{\mu}'|\boldsymbol{\mu}) &= \mathcal{N}(f(\boldsymbol{\mu}), \Sigma_\mu), & p(\boldsymbol{\mu}''|\boldsymbol{\mu}') &= \mathcal{N}(f'(\boldsymbol{\mu}'), \Sigma_{\mu'}), \end{aligned} \quad (4)$$

where the functions $f(\boldsymbol{\mu})$ and $f'(\boldsymbol{\mu}')$ represent the evolution of the belief state over time. This encodes the agent's preference over future states (in this case the preferred future state is the target state, $\boldsymbol{\mu}_d$). We assume: $f(\boldsymbol{\mu}) = (\boldsymbol{\mu}_d - \boldsymbol{\mu})\tau^{-1}$ and $f'(\boldsymbol{\mu}') = \tau^{-1}\boldsymbol{\mu}'$, where $\boldsymbol{\mu}_d$ is the desired state and τ is a temporal parameter.

Using the previous equations, we can expand F to:

$$\begin{aligned} F &= \frac{1}{2}(\boldsymbol{\varepsilon}_o^\top \Sigma_o^{-1} \boldsymbol{\varepsilon}_o + \boldsymbol{\varepsilon}_{o'}^\top \Sigma_{o'}^{-1} \boldsymbol{\varepsilon}_{o'}) \\ &+ \boldsymbol{\varepsilon}_\mu^\top \Sigma_\mu^{-1} \boldsymbol{\varepsilon}_\mu + \boldsymbol{\varepsilon}_{\mu'}^\top \Sigma_{\mu'}^{-1} \boldsymbol{\varepsilon}_{\mu'} \\ &+ \ln |\Sigma_o| + \ln |\Sigma_{o'}| + \ln |\Sigma_\mu| + \ln |\Sigma_{\mu'}| + C, \end{aligned} \quad (5)$$

where $\boldsymbol{\varepsilon}_\mu = \boldsymbol{\mu}' - (\boldsymbol{\mu}_d - \boldsymbol{\mu})\tau^{-1}$, $\boldsymbol{\varepsilon}_{\mu'} = \boldsymbol{\mu}'' + \tau^{-1}\boldsymbol{\mu}'$, $\boldsymbol{\varepsilon}_o = \mathbf{o} - \boldsymbol{\mu}$ and $\boldsymbol{\varepsilon}_{o'} = \mathbf{o}' - \boldsymbol{\mu}'$ and C refers to constant terms.

C. Estimation and control

To achieve state estimation and control, we perform gradient descent on F using the following update rules:

$$\dot{\tilde{\boldsymbol{\mu}}} = D\tilde{\boldsymbol{\mu}} - \kappa_\mu \frac{\partial F}{\partial \tilde{\boldsymbol{\mu}}}, \quad (6)$$

$$\dot{\tilde{\mathbf{a}}} = -\kappa_a \frac{\partial F}{\partial \tilde{\mathbf{a}}} = -\kappa_a \frac{\partial F}{\partial \tilde{\mathbf{o}}} \frac{\partial \tilde{\mathbf{o}}}{\partial \tilde{\mathbf{a}}}, \quad (7)$$

where κ_a and κ_μ are tuning parameters depending on the desired behaviour, and D is the temporal derivative operator.

D. Precision learning

Precision learning refers to learning the inverse covariance matrices, also referred to as the precision matrices. This is done using one-step gradient descent [8] as:

$$\dot{\Sigma}_o^{-1} = -\kappa_\sigma \frac{\partial F}{\partial \Sigma_o^{-1}}, \quad \dot{\Sigma}_{o'}^{-1} = -\kappa_\sigma \frac{\partial F}{\partial \Sigma_{o'}^{-1}}. \quad (8)$$

Now, using Equations 6, 7 and 8, the manipulator can perform state-estimation, control and precision learning. Since in Equation 7, we use the chain rule with respect to the observations, Σ_o^{-1} and $\Sigma_{o'}^{-1}$ are the matrices that show up in the control law.

III. RESULTS

We evaluate the presented approach and use the active inference controller (AIC) from [3] as a benchmark since the authors have shown their work outperforms MRAC.

Consider a manipulator moving between three different configurations (see supplementary video or section 5 of [3]). If the AIC is tuned properly ($\Sigma_o^{-1} = 1.5I$, $\Sigma_{o'}^{-1} = 0.5I$, $\Sigma_\mu^{-1} = 0.1I$ and $\Sigma_{\mu'}^{-1} = 0.5I$), this results in satisfactory behaviour. In this case, I refers to the 7×7 identity matrix. However, if we perturb all diagonal elements of Σ_μ^{-1} by random values between 0 and 0.2 during initialization, the controller suffers from oscillations. Our approach overcomes this issue by performing precision learning. We update Σ_o^{-1} and $\Sigma_{o'}^{-1}$ for the first two seconds during operation which allows the controller to 'self-tune'. Results are shown in Figure 1. This shows the performance of the AIC and our approach for a single run. Using our approach, all joints show almost no sign of oscillations.

IV. CONCLUSION

In this work we perform state estimation, control and precision learning under the active inference framework. The approach is tested (on the Panda Emika Franka) against previous work and shows the ability to damp oscillations when the precision matrices are not properly initialized.

REFERENCES

- [1] Bin Wei. Adaptive control design and stability analysis of robotic manipulators. In *Actuators*, volume 7, page 89. Multidisciplinary Digital Publishing Institute, 2018.
- [2] Karl Friston, Thomas FitzGerald, Francesco Rigoli, Philipp Schwartenbeck, and Giovanni Pezzulo. Active inference: a process theory. *Neural computation*, 29(1):1–49, 2017.
- [3] Corrado Pezzato, Riccardo Ferrari, and Carlos Hernandez. A novel adaptive controller for robot manipulators based on active inference. *arXiv preprint arXiv:1909.12768*, 2019.
- [4] Dan Zhang and Bin Wei. A review on model reference adaptive control of robotic manipulators. *Annual Reviews in Control*, 43:188–198, 2017.
- [5] Charles W Fox and Stephen J Roberts. A tutorial on variational bayesian inference. *Artificial intelligence review*, 38(2):85–95, 2012.
- [6] Karl Friston, Jérémie Mattout, Nelson Trujillo-Barreto, John Ashburner, and Will Penny. Variational free energy and the laplace approximation. *Neuroimage*, 34(1):220–234, 2007.
- [7] Christopher L Buckley, Chang Sub Kim, Simon McGregor, and Anil K Seth. The free energy principle for action and perception: A mathematical review. *Journal of Mathematical Psychology*, 81:55–79, 2017.
- [8] Rafal Bogacz. A tutorial on the free-energy framework for modelling perception and learning. *Journal of mathematical psychology*, 76:198–211, 2017.

Towards Human-Chatbot Interaction: A Virtual Assistant for the Ramp-up Process

Melanie Zimmer, Ali Al-Yacoub, Pedro Ferreira, Niels Lohse
Intelligent Automation Centre
Loughborough University
Loughborough, UK
{m.zimmer2, a.al-yacoub, p.ferreira, n.lohse}@lboro.ac.uk

<https://doi.org/10.31256/Qx5Dt5V>

Abstract—Nowadays, we are surrounded by virtual assistants in everyday life. But one domain that is assumed to massively benefit from virtual assistants, is manufacturing. In particular, where activities are reliant on human expertise and knowledge, a virtual assistant could help support the human. The vision of this work is inspired by the need for bringing an assembly system more rapidly to an operational state. To achieve this vision, a decision-support framework that aims to better integrate the human operator into the ramp-up activity is proposed. As part of this framework, natural language processing tools are applied to allow the development of a virtual assistant for the ramp-up process. This paper provides an overview of the current work in progress, which is part of a PhD research undertaken at the Intelligent Automation Centre at Loughborough University. It outlines the initial efforts and future steps that have been completed and are planned.

Keywords—Ramp-up Process, Natural Language Processing, Natural Language Generation, Chatbot, Decision-support, Industry 4.0.

I. BACKGROUND

Being able to adapt to increasing demand and customisation of products rapidly is crucial for the competitive strengths of manufacturing companies. Thus, manufacturers are required to assemble a new or tweak an existing production system and getting it to full production in a very short time. Despite these two processes bearing similarities, the scope of this research work is on the former. This process of bringing a system from a low level to full volume operation takes place during the so-called production ramp-up [1]. Ramping up a system requires the human operator to perform process and equipment adjustments based on his/her knowledge and expertise as part of an iterative process (Fig. 1). Simply put, once the system is in place, it will be tested with certain settings to verify if the required Key Performance Indicators (KPIs), such as functionality, product quality, cycle time, etc., are fulfilled [2]. In the unlikely situation, where this is the case, the ramp-up process is finished. The more common situation is that a change to the physical setup or the process needs to be made. This cycle repeats until the KPIs are ultimately met. As can easily be figured, ramp-up can be a very time-consuming activity [1], as due to the uniqueness of each case it makes it heavily reliant on the expertise and intervention of a human [3] that can vary extremely among different people.

During recent years, the understanding that technology and human require a better interaction to achieve successful production ramp-ups has urged manufacturers to rethink their

The research leading to these results has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 680735, project openMOS (Open Dynamic Manufacturing Operating System for Smart Plug-and-Produce Automation Components). Funding from the Engineering and Physical Science Research Council Centre for Doctoral Training in Embedded Intelligence (grant no. EP/L014998/1) is also acknowledged.

strategies on mainly relying on the human [4], [5]. As more and more data become available in the manufacturing domain nowadays through an increased level of end-to-end digitisation and automation, also referred to as Industry 4.0, improved monitoring of real-time ramp-up processes can be supported [4]. An important research subject that has been identified through literature is the improvement of knowledge capture, reuse and communication during the ramp-up to minimise disturbances due to loss of human knowledge. Surprisingly, few researchers have thus far proposed a suitable learning approach and assisting tools (software) for ramp-up [6]. When asked about the function of technology for future ramp-up processes, studies indicate [7] that it will play an important role during ramp-up management, whereas the human will be tasked with problem-solving.

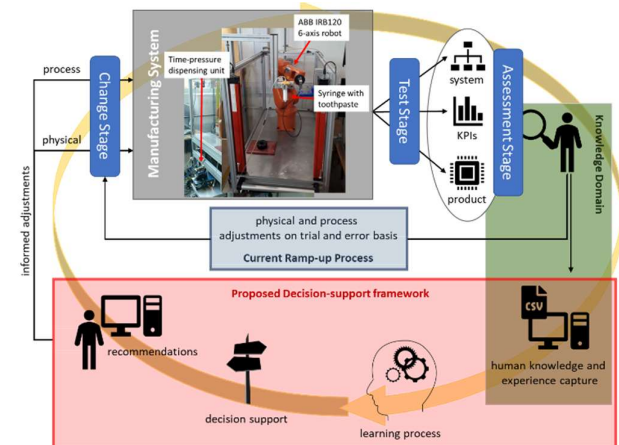


Fig. 1. Overview of the decision-support framework for plug-and-produce assembly systems in contrast to current ramp-up process practice.

The idea presented here is part of an ongoing PhD research work, which mainly aims to reduce the ramp-up effort and ultimately shorten the ramp-up time for plug-and-produce assembly systems. The main objective will be to create a decision-support framework, which will guide a human operator in making adjustments to the equipment and processes of the ramped-up system. One aspect of the proposed decision-support framework includes a module that makes use of Natural Language Processing (NLP) techniques to extract knowledge from captured data and transform it into a meaningful semantic representation. Thus, the hypothesis that is addressed by this work is the following: “Providing shop-floor operators with a virtual assistant during the ramp-up process will reduce the number of trials required to ramp up a system.” As such, relevant information will be made available in a way that the operator’s decision-making for quick system adaptations is supported. This paper presents the idea of a virtual assistant for the operator to provide recommendations in the form of a chatbot.

II. METHODOLOGY

To create a useful decision-support mechanism, formal means of extraction, analysis and interpretation of meaningful data related to human problem-solving are required. For this work, human expertise is captured in natural language during the ramp-up via a developed Graphical User Interface (GUI) as this is a very intuitive way for humans to converse. These data are enriched with data about the system state in terms of process and equipment settings, current and target KPIs. This ultimately allows forming a knowledge base, which contains historical data and information from previous ramp-up scenarios. For this research work, a prototype virtual assistant represented by a chatbot has been developed to provide support to the human in a natural means of communication. A chatbot can be described as a type of conversational or dialogue agent that can engage in a natural language conversation with a human, usually through a keyboard. The general information flow between the chatbot and the human can be described as seen in Fig. 2.




Fig. 2. Simplified Request-Response flow between human and chatbot.

For this research work in progress, a straightforward GUI using Python's Tkinter has been developed to allow the user to engage with the chatbot (Fig. 3). Where the human uses the keyboard to type and send a message, the chatbot uses other Python libraries such as NLTK, Keras and TensorFlow to build the conversation. The underlying model that is used here is a Keras's *Sequential* deep learning model [8]. For example, the operator can enquire about certain settings by asking questions related to the individual equipment, such as the nozzle diameter in the given case study explained later. The chatbot reacts to certain identified keywords based on which it chooses its answer.

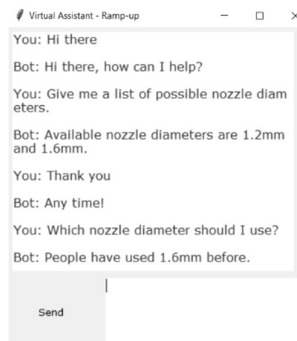


Fig. 3. Current chatbot interface for the virtual ramp-up assistant.

As can be seen, the conversation is currently still very limited and needs enhancements to allow for more detailed questions and answers. One difficulty that was, however, encountered is that many machine learning approaches rely on the availability of sufficient data samples in order to be able to provide meaningful insights about reoccurring patterns. This poses a challenge for the use of these methods during ramp-up as little data are available at that time. Additionally, it has been found that existing chatbot implementations are mostly

trained on movie or other reviews. The manufacturing domain as a corpus is not yet sufficiently addressed. As such, data have previously been collected from a manual dispensing experiment to better understand human decision-making for a dispensing-like task which will serve as an input to the proposed chatbot. The data allowed to extract change actions that achieved the necessary KPIs, but also highlighted certain issues that were encountered. By creating a knowledge base with this information and providing it to the user through the chatbot, it is assumed that the trial and error approach conventionally taken in the ramp-up process can be minimised. More information about the data and the experiment itself can be found in [9].

III. CASE STUDY

In order to be able to test the usefulness of the proposed approach in the near future, an industrial use case of a dispensing process has been developed (cf. *Manufacturing System* in Fig. 1). The objective of this experiment is to tune parameters on the setup to obtain products of good quality. There are three product variations and each participant will have to repeatedly create at least one and until you reach the required quality. Good quality is defined by straight and continuous lines, with no excessive dispensing material and close similarity to the given target pictures. The setup's key component is a single 6-axis industrial robot (ABB IRB120), which is connected to an IRC5 controller. Toothpaste, to simulate a dispensing process, is dispensed in a controlled manner through a nozzle that has been mounted to the surrounding frame as the robot will manipulate the metal workpiece. An automated time-pressure dispensing unit (Fisnar JB1113N) is used. If any of the process or equipment parameters are, however, not set fittingly, the desired product quality and performance will not be achieved, and the adjustment step needs to be repeated. More information about the setup can be found in [10]. To validate the underlying hypothesis of this research given in the introduction of this paper, two sets of participants will be asked to perform a ramp-up scenario on the aforementioned glueing workstation. The first group will solely rely on their knowledge and equipment manuals, whereas the second group will, in addition, have access to the developed chatbot. By comparing the usage and difference in time required to ramp-up the setup to full production state, the usefulness of the proposed virtual assistant for the ramp-up process can be evaluated.

IV. CONCLUSIONS AND FUTURE WORK

This paper is part of an ongoing PhD work looking into reducing the time required to get a system to full production. This stage, also known as ramp-up, is still very human-centric and error-prone. As part of the overall PhD research aim to provide a decision-support framework, one aspect is the development of a virtual assistant for the operator undertaking the ramp-up process. This virtual assistant will be provided in the form of a chatbot, which allows the operator to enquire about certain settings or issues that have previously occurred in similar cases. The chatbot can give recommendations about change actions that can be applied to the system or other useful information. This chatbot is currently under development and its usefulness and, thus, the verification of the introduced hypothesis will be tested on the case study introduced previously. This will be achieved by dividing the participants into two groups, for which one will have access to the chatbot, whereas the other will rely on their sole knowledge and expertise.

REFERENCES

- [1] C. Terwiesch and R. E. Bohn, "Learning and process improvement during production ramp-up," *Int. J. Prod. Econ.*, vol. 70, pp. 1–19, 2001, Accessed: Oct. 12, 2017.
- [2] S. C. Doltsinis, S. Ratchev, and N. Lohse, "A framework for performance measurement during production ramp-up of assembly stations," *Eur. J. Oper. Res.*, vol. 229, no. 1, pp. 85–94, 2013, doi: 10.1016/j.ejor.2013.02.051.
- [3] S. Doltsinis, P. Ferreira, and N. Lohse, "Reinforcement learning for production ramp-up: A Q-batch learning approach," *Proc. - 2012 11th Int. Conf. Mach. Learn. Appl. ICMLA 2012*, vol. 1, pp. 610–615, 2012, doi: 10.1109/ICMLA.2012.113.
- [4] U. Dombrowski, J. Wullbrandt, and P. Krenkel, "Industrie 4.0 in production ramp-up management," in *Procedia Manufacturing*, 2018, vol. 17, pp. 1015–1022, doi: 10.1016/j.promfg.2018.10.085.
- [5] D. Scrimieri, R. F. Oates, and S. M. Ratchev, "Learning and reuse of engineering ramp-up strategies for modular assembly systems," *J. Intell. Manuf.*, vol. 26, no. 6, pp. 1063–1076, 2015, doi: 10.1007/s10845-013-0839-6.
- [6] K. Konrad, M. Hoffmeister, M. Zapp, A. Verl, and J. Busse, "Enabling fast ramp-up of assembly lines through context-mapping of implicit operator knowledge and machine-derived data," in *IFIP Advances in Information and Communication Technology*, 2012, vol. 371 AICT, pp. 163–174, doi: 10.1007/978-3-642-28163-1_20.
- [7] R. Schmitt et al., "On the future of ramp-up management," *CIRP J. Manuf. Sci. Technol.*, vol. 23, pp. 217–225, 2018, doi: 10.1016/j.cirpj.2018.03.001.
- [8] "Guide to the Sequential model - Keras Documentation." <https://keras.io/getting-started/sequential-model-guide/> (accessed Apr. 06, 2020).
- [9] M. Zimmer, A. Al-Yacoub, P. Ferreira, and N. Lohse, "Understanding human decision-making during production ramp-up using natural language processing," *IEEE Int. Conf. Ind. Informatics*, vol. 2019-July, pp. 337–342, 2019, doi: 10.1109/INDIN41052.2019.8972033.
- [10] M. Zimmer, P. Ferreira, P. Danny, A. Al-Yacoub, N. Lohse, and V. Gentile, "Towards a decision-support framework for reducing ramp-up effort in plug-and-produce systems," *Proc. - 2019 IEEE Int. Conf. Ind. Cyber Phys. Syst. ICPS 2019*, pp. 478–483, 2019, doi: 10.1109/ICPHYS.2019.8780369.

Power-aware Fusion of Visual and Wheel Odometry for Mobile Platforms

Mateusz Malinowski
Department of Aerospace Engineering
University of Bristol
Bristol, UK
mateusz.malinowski@bristol.ac.uk

Arthur Richards
Department of Aerospace Engineering
University of Bristol
Bristol, UK
arthur.richards@bristol.ac.uk

Mark Woods
Head of Autonomy and Robotics Group
SCISYS UK Ltd.
Bristol, UK
mark.woods@scisys.co.uk

<https://doi.org/10.31256/Bq8Qz4S>

Abstract—A new approach is presented to estimate the motion of a mobile platform using a fusion of visual odometry (VO) and wheel odometry (WO). An Extended Kalman Filter is used with an augmented state vector including the wheel slip. The new approach allows for flexible use of VO, which is central to slip estimation but consumes significant electrical power. The approach can be used in the future to adaptively optimise the number of VO measurements.

Keywords—mobile platform, Visual Odometry, wheel slip, Extended Kalman Filter

I. INTRODUCTION

Autonomous navigation on other solar system bodies, like Mars, face many interesting challenges. One of the critical problems is precise localisation. Without easy access to on-demand global position measurements, relative localisation techniques, such as Visual Odometry (VO), are of the essence. Furthermore, VO has the additional benefit of allowing to measure wheels slip which in case of Mars can contribute to over 10% of error in the position estimate [1]. Precise and beneficial as it may be, VO requires considerable computation effort and relies on good illumination, resulting in significant electrical power demand. For a Mars rover, just like for any spacecraft, the power budget is carefully balanced, and it cannot spend any more energy than estimated. For future Mars missions, such as European's Sample Fetch Rover, the rover is expected to drive for much longer distances than any other remote missions to date [2]. This objective poses another technical challenge of how to minimise energy usage while maintaining accurate localisation and maximising distance travelled. In this paper, we use a novel Extended Kalman Filter formulation to explore the trade between VO use and navigation accuracy.

II. RELATED WORK

Many different algorithms and strategies may be employed to optimise energy usage and increase localisation accuracy. In their work, NASA presents how VO measurements taken by Mars Exploration Rover (MER) and Mars Science Laboratory (MSL) rover are used to estimate slip and support trajectory control [3], [4]. Different power optimisation strategy was presented in [5] where the optimisation was achieved by introducing an intelligent controller to command wheels based on the soil type. This approach requires parameters tuning but may lead to a more efficient drive. Another solution is to recognise terrain in front of the rover to adapt to the scenery. So far this has been presented as either prediction of slip for better trajectory planning, where rover can avoid high-slip areas to preserve energy, [6] or the overall scenery classification to switch between different sensors which may either optimise energy usage or increase localisation precision [7].




Fig. 1 Block diagram of the proposed model

Often VO is used to estimate wheels slip. In [8] authors present how slip can be estimated using Inertial Measurement Unit (IMU); however, in their approach, calibrated wheels angular velocity is still fused with VO.

III. APPROACH

The proposed model represents only 1D motion and is presented on Fig. 1 with p being the position, v – velocity, θ – wheels accumulated angle, Ω – wheels angular velocity, λ – slip, and R – wheels' radius. Both α and β are treated as process noises that drive the model, where α is interpreted as wheels angular acceleration and β describes the terrain. The system has been discretised with a step size of 0.1 s with state vector defined in (1).

$$x(t) = [p(t) \ \theta(t) \ \Omega(t) \ \lambda(t) \ m(t)]^T \quad (1)$$

To model VO, the state vector is augmented with m , defined as the position at which the last VO image was recorded:

$$m(t^+) = \begin{cases} p(t) & \text{if VO measurement at time } t \\ m(t) & \text{otherwise} \end{cases} \quad (2)$$

Then the two odometry measurements (wheel odometry (WO) which measures wheels accumulated angle and VO which measured delta position) are expressed as:

$$y_{WO}(t) = \theta(t) + \varepsilon_{WO}(t) \quad (3)$$

$$y_{VO}(t) = p(t) - m(t) + \varepsilon_{VO}(t) \quad (4)$$

where ε denotes the respective measurement noise and noting that the VO measurement is not taken at every update.

Wheel slip is modelled using a first-order filter driven by an unknown process noise. This approach means the slip uncertainty grows with time, but not unboundedly as a random walk, and has a tuneable 'forgetting' factor to capture the variability of terrain. The gain parameter of the filter was set to one as it is a scaling factor for the β noise, which models all-terrain properties related to the slip. To simplify this evaluation, we assumed no skidding or breaking; therefore, the slip is defined as presented in (5).

$$\lambda(t) = 1 - v(t) / [R * \Omega(t)] \quad (5)$$




Fig. 2 Mobile platform used for experiments.

IV. RESULTS AND DISCUSSION

The model has been validated using data from real-world experiments where Mars-like rover with Real-Time Kinematic (RTK) GNSS for ground truth (as presented on Fig. 2) was driving on two different types of terrain: tarmac and grass. All graphs from Fig. 3 present sets of trajectories that lasted for 500 s. The rover speed was 0.2 m/s, and VO measurements were done at approximately 10 Hz frequency. Graphs on the left-hand side present how changing VO update rate impacts estimated position errors and their three standard deviation (3σ) boundaries. All trajectories, except for one which happened on a hilly grass terrain, are always within estimated precision. Performing VO measurements more often (e.g. Fig. 3 VO @ 0.2 s) leads to narrower three standard deviation boundaries due to high frequency of precise position measurements. With decreased VO update rate (e.g. Fig. 3 VO @ 3 s), the model maintains its accuracy for most of the trajectories; however, errors tend to be greater since the filter leans more towards WO estimates. Graphs on the right-hand side present how changing σ_β affects position estimation. Low values (e.g. Fig. 3 $\sigma_\beta=0.01$) indicate relatively constant slip, and therefore the filter prefers to trust WO measurements more. In this case, VO is mainly used to measure the slip value indirectly. On the other hand, high σ_β value indicates terrain where slip may change dynamically, therefore the VO measurements are preferred (e.g. Fig. 3 $\sigma_\beta=2$). Also, because of more dynamic changes, more noise into position measurements is introduced, which is visible as saw-tooth on three standard deviation boundaries.

V. CONCLUSIONS

This paper presents a new model for fusing wheel and visual odometry on a mobile robot. We successfully validated the proposed model against the real-world data captured using representative Mars-like rover. The results suggest that the model correctly predicts navigation uncertainty levels as the frequency of VO measurements change. Further research will use this model to pursue an adaptive controller capable of optimising the number of VO measurements to trade accuracy against electrical power usage.

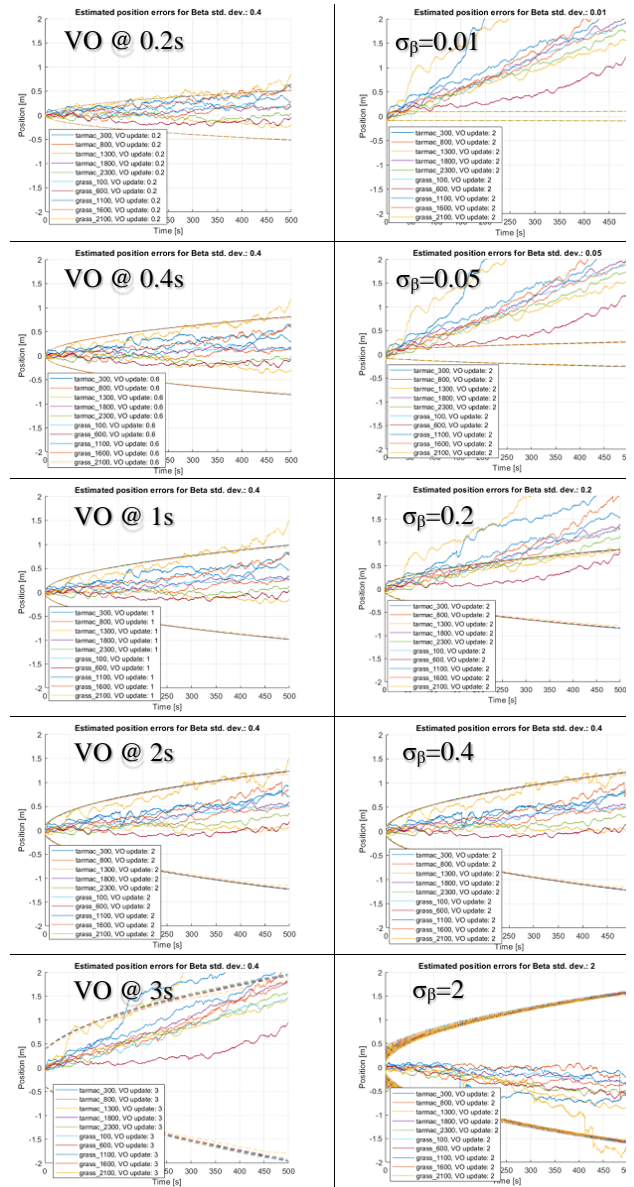


Fig. 3 Position estimation errors for varying VO update time (left-hand side) and β standard deviation (right-hand side).

ACKNOWLEDGEMENT

M. Malinowski thanks SCISYS UK Ltd. and the University of Bristol for the use of equipment and services to conduct this research.

REFERENCES

- [1] M. Maimone, Y. Cheng, and L. Matthies, "Two years of Visual Odometry on the Mars Exploration Rovers," *Journal of Field Robotics*, vol. 24, no. 3, pp. 169–186, 2007, doi: 10.1002/rob.20184.
- [2] L. Duvet, F. Beyer, J. Delfa, and E. Zekri, "ESA Sample Fetch Rover: Heritage and Way Forward," presented at the Second International Mars Sample Return, Apr. 2018, vol. 2071, p. 6122, [Online]. Available: <https://ui.adsabs.harvard.edu/abs/2018LPICo2071.6122D>.
- [3] D. M. Helmick, Yang Cheng, D. S. Clouse, L. H. Matthies, and S. I. Roumeliotis, "Path following using visual odometry for a Mars rover in high-slip environments," in *2004 IEEE Aerospace Conference Proceedings (IEEE Cat. No.04TH8720)*, Mar. 2004, vol. 2, pp. 772–789 Vol.2, doi: 10.1109/AERO.2004.1367679.
- [4] O. Toupet *et al.*, "Traction control design and integration onboard the Mars science laboratory curiosity rover," in *2018 IEEE Aerospace Conference*, Mar. 2018, pp. 1–20, doi: 10.1109/AERO.2018.8396761.

- [5] Jayoung Kim and Jihong Lee, "Intelligent slip-optimization control with traction-energy trade-off for wheeled robots on rough terrain," in *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Sep. 2014, pp. 1938–1943, doi: 10.1109/IROS.2014.6942819.
- [6] A. Angelova, L. Matthies, D. Helmick, G. Sibley, and P. Perona, "Learning to predict slip for ground robots," in *Proceedings 2006 IEEE International Conference on Robotics and Automation, 2006. ICRA 2006.*, May 2006, pp. 3324–3331, doi: 10.1109/ROBOT.2006.1642209.
- [7] M. Woods, A. Shaw, I. Wallace, and M. Malinowski, "The Chameleon field trial: toward efficient, terrain sensitive navigation," in *ASTRA 2015*, May 2015.
- [8] K. Bussmann, L. Meyer, F. Steidle, and A. Wedler, "Slip Modeling and Estimation for a Planetary Exploration Rover: Experimental Results from Mt. Etna," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Oct. 2018, pp. 2449–2456, doi: 10.1109/IROS.2018.8594294.

Debiasing of position estimations of UWB-based TDoA indoor positioning system

Paolo Grasso

*Autonomous Vehicles & Artificial Intelligence Laboratory
Research Institute for Future Transport and Cities
Coventry University
Coventry, UK
grassop@coventry.ac.uk*

Mauro S. Innocente

*Autonomous Vehicles & Artificial Intelligence Laboratory
Research Institute for Future Transport and Cities
Coventry University
Coventry, UK
Mauro.S.Innocente@coventry.ac.uk*

<https://doi.org/10.31256/Ua2Vp3X>

Abstract—When localising an object in a confined environment using an indoor positioning system (IPS) based on ultra-wideband (UWB) technology and an asynchronous time difference of arrival (TDoA) algorithm, systematic errors do occur. Theoretical estimation of these errors can be very hard to make. This study introduces a novel filtering algorithm for reducing the bias of position estimations, therefore increasing their accuracy. The problem is tackled for a two-dimensional IPS by formulating a *debiasing filter* using statistics of real data. Generalisation to the three-dimensional case should be straightforward.

Index Terms—indoor positioning system (IPS), accuracy, bias, precision, time difference of arrival (TDoA), ultra-wideband (UWB)

I. INTRODUCTION

In robotics and autonomy, positioning systems constitute a remarkably important technology. A variety of systems exist, which are aimed at different applications and make use of different algorithms (e.g. Time Difference of Arrival (TDoA) or Two Way Ranging) and different forms of energy (e.g. electromagnetic or sound waves). For example, Global Navigation Satellite Systems (GNSSs) are appropriate for efficient outdoor long-range positioning, while Indoor Positioning Systems (IPS) based on ultra-wideband (UWB) technology present crucial advantages in indoor spaces, namely high accuracy and the ability to penetrate obstacles [3]. It is important to emphasise that UWB positioning constitutes one of the most accurate and promising technologies for IPSs, arguably the best choice at present [3], [4]. A major drawback is given by its susceptibility to interferences, which may be caused by metallic materials and/or by systems working on similar frequencies.

This work builds on a previous study on the signal and geometrical properties of a general IPS [11]. In that study, the unbiased Cramér–Rao Lower Bound (CRLB) analysis for the reference IPS estimated that the precision inside the convex hull defined by the anchors was within ± 3 cm whereas the bias could not be assessed theoretically. Furthermore, an analysis based on [5]–[9] was performed in order to define the localisation boundary – i.e. the bifurcation envelope. Using the available IPS, this work aims to assess if the theoretical prediction of precision in [11] is correct and to estimate the order of magnitude of the bias. These data are subsequently

used for the development of a debiasing filter aimed at increasing the accuracy of the position estimations. The state of the art of currently used IPS is the Extended Kalman Filter (EKF) developed by Mueller et al. [1], [2]. However, there are systematic errors that the EKF does not deal with, as observed in the experiments and acknowledged in [1]. The proposed debiasing filter is an attempt to find a solution to this problem.

II. DESIGN OF EXPERIMENT ON IPS

The purpose of the IPS under study is to localise a moving object based on a spatial distribution of transceivers (anchors) using an asynchronous TDoA algorithm. The designed experiment aims to provide the precision and accuracy maps of the position measurements obtained with an IPS. These maps will be used by the subsequently introduced debiasing filter.

The IPS in question, which is depicted in Fig. 1, consists of a drone to be localised and four transceiver anchors positioned at the vertices and facing the centre of a square domain. All the antennae are at a height of 20 cm from the floor. The drone is on a moving stand equipped with a laser pointer, which is aligned with the onboard UWB antenna in order to achieve reference positioning of high precision (± 1 mm) and accuracy. The equally spaced markers stuck to be floor are the sampling positions, as shown in Fig. 2.




Fig. 1. Diagram of the setup of the studied 4×4 m² IPS for 2D localisation: (a) are the adjustable stands of the (A0–A3) transmitting anchors antennae, (b) the measurement points regularly distributed every 50 cm in both directions, and (c) the mobile stand for the object to be localised.




Fig. 2. Precision mapping ($\pm^x\sigma$) of the x component of the position. The sampling points are the magenta squares.




Fig. 3. Accuracy mapping (xb) of the x component of the position. The sampling points are the magenta squares.




Fig. 4. Debiasing surface ($^x\beta$) for the x component obtained by cubic spline interpolation of the estimated scattered debiasing values ($^x\beta_{ij}$).

In order to build the maps, a large number of measurements ($N = 700$) are taken at a sampling frequency of 100 Hz while keeping the drone still for at least 30 seconds on each marker (X_{ij}). Then, the raw stream of data is post-processed getting rid of the transients corresponding to the movement between markers. The drone is kept aligned with the x axis and parallel to the floor, as the effect of its direction is not being investigated. Finally, the bias (b), standard deviation (σ) and mean squared error (MSE) are computed. For instance, their values in the x direction (superscript x) are as follows:

$$\begin{aligned} ^xb_{ij} &= N^{-1} \sum_{k=1}^N x^{(k)} - X_{ij} \\ ^x\sigma_{ij} &= \left(^x\text{MSE}_{ij} - ^xb_{ij}^2 \right)^{0.5} \\ ^x\text{MSE}_{ij} &= N^{-1} \sum_{k=1}^N \left(x^{(k)} - X_{ij} \right)^2 \end{aligned} \quad (1)$$

where $x^{(k)}$ is the k^{th} position measurement. The resulting maps can be found in Fig. 2 and Fig. 3.

III. FORMULATION OF DEBIASING FILTER

The objective of the debiasing filter is to correct the raw EKF estimation of the position ($\mathbf{x} = [x, y]^T$) so that the filtered position ($\hat{\mathbf{x}}$) is approximately the same as the actual real position ($\mathbf{X} = [X, Y]^T$). The debiasing filter on the sample points described in Section II should provide the real position.

From the definition of bias and variance in (1), the measurement on a reference point can be factorised into (2): the real position, the bias, and a random fluctuation (\mathcal{R}) which is a function of the standard deviation. Note that both b and σ are functions of the real position. For the proposed formulation, \mathcal{R} will be neglected even if this affects the quality of the filter.

$$\begin{aligned} x_{ij} &= X_{ij} + ^xb_{ij} + \mathcal{R}(^x\sigma_{ij}) \\ ^xb_{ij} &= ^xb(X_{ij}, Y_{ij}) \\ ^x\sigma_{ij} &= ^x\sigma(X_{ij}, Y_{ij}) \end{aligned} \quad (2)$$

Since the input of the filter is going to be the estimated position, but the bias map in Fig. 3 is a function of the real position, a change of domain is needed. As implicitly suggested by (2), the change of domain is expressed as in (3).

In this case, the debiasing value will be simply the raw bias with opposite sign, i.e. $^x\beta_{ij} = -^xb_{ij}$.

$$\bar{\mathbf{X}}_{ij} = \mathbf{X}_{ij} + [^xb_{ij}, ^yb_{ij}]^T \quad (3)$$

Hence, the obtained debiasing values are no longer functions of the real position (which is unknown) but are function of the measured position. As shown in (4) for the x direction, the sampled debiasing discrete distributions β_{ij} in both directions can be interpolated over a deformed grid $[\bar{X}_{ij}, \bar{Y}_{ij}]$, obtaining debias surfaces that are functions of the measured positions.

$$[\bar{X}_{ij} \quad \bar{Y}_{ij} \quad ^x\beta_{ij}] \xrightarrow{\text{interp.}} ^x\beta(\mathbf{x}) \quad (4)$$

For instance, in Fig. 4, $^x\beta(\mathbf{x})$ is described by a simple cubic spline. Since the debiasing filter must be applied in real time, a more efficient interpolation method will be considered in the future. Finally, the debiasing values can be directly added to the measured position obtaining the filtered estimation $\hat{\mathbf{x}}$:

$$\hat{x} = x + ^x\beta(\mathbf{x}) \quad \hat{y} = y + ^y\beta(\mathbf{x}) \quad (5)$$

IV. CONCLUSION AND FUTURE WORK

An experimental evaluation of the distribution of variance and bias of positioning estimations was performed. The precision values are bounded by $\pm 3\text{cm}$ as predicted by a previous theoretical study of IPSs [11]. However, the accuracy level was unsatisfactory. In order to increase it, a data-driven debiasing filter was formulated. The debiasing surface obtained is a function of only the biases measured on a discrete number of points. At present, this work is being extended towards a more comprehensive filtering procedure aimed at increasing both precision and accuracy. The debiasing filter being developed will also consider the variance. In addition, interpolating neural networks will enable three dimensions. The experimental platform will be further improved and employed for tests and validation of these filters. Furthermore, the developed IPS with high precision and accuracy will also be used for indoor experiments to test multi-agent self-coordination and collision-avoidance algorithms – e.g. in [12].

REFERENCES

- [1] M. W. Mueller, M. Hamer, and R. D'Andrea, "Fusing ultra-wideband range measurements with accelerometers and rate gyroscopes for quadcopter state estimation," in *International Conference on Robotics and Automation (ICRA)*. IEEE, May 2015.
- [2] M. W. Mueller, M. Hehn, and R. D'Andrea, "Covariance correction step for kalman filtering with an attitude," *Journal of Guidance, Control, and Dynamics*, vol. 40, no. 9, pp. 2301–2306, sep 2017.
- [3] A. Alarifi, A. Al-Salman, M. Alsaleh, A. Alnafessah, S. Al-Hadhrami, M. Al-Ammar, and H. Al-Khalifa, "Ultra wideband indoor positioning technologies: Analysis and recent advances," *Sensors*, vol. 16, no. 5, p. 707, may 2016.
- [4] M. Ghavami, L. Michael, and R. Kohno, *Ultra Wideband Signals and Systems in Communication Engineering*. John Wiley & Sons, 2007.
- [5] M. Compagnoni, R. Notari, F. Antonacci, and A. Sarti, "A comprehensive analysis of the geometry of TDOA maps in localization problems," *Inverse Problems*, vol. 30, no. 3, p. 035004, feb 2014.
- [6] S. O. Dulman, A. Baggio, P. J. Havinga, and K. G. Langendoen, "A geometrical perspective on localization," in *Proceedings of the first ACM international workshop on Mobile entity localization and tracking in GPS*. ACM Press, 2008.
- [7] M. Compagnoni and R. Notari, "Tdoa-based localization in two dimensions: the bifurcation curve," *Fundamenta Informaticae*, vol. 135, no. 1-2, pp. 199–210, 2014.
- [8] R. Kaune, J. Horst, and W. Koch, "Accuracy analysis for tdoa localization in sensor networks," in *14th International Conference on Information Fusion*, 2011.
- [9] R. Kaune, "Accuracy studies for tdoa and toa localization," Aug. 2012.
- [10] M. Compagnoni, A. Pini, A. Canclini, P. Bestagini, F. Antonacci, S. Tubaro, and A. Sarti, "A geometrical–statistical approach to outlier removal for TDOA measurements," *IEEE Transactions on Signal Processing*, vol. 65, no. 15, pp. 3960–3975, aug 2017.
- [11] P. Grasso and M. S. Innocente, "Theoretical study of signal and geometrical properties of two-dimensional uwb-based indoor positioning systems using tdoa," in *6th International Conference on Mechatronics and Robotics Engineering (ICMRE)*. IEEE, Feb. 2020.
- [12] M. S. Innocente and P. Grasso, "Self-organising swarms of firefighting drones: Harnessing the power of collective intelligence in decentralised multi-robot systems," *Journal of Computational Science*, vol. 34, pp. 80–101, may 2019.

The Augmented Agronomist Pipeline and Time Series Forecasting

George Onoufriou

*MLearn, LCAS. University of Lincoln
School of Computer Science
Lincoln, UK
0000-0002-9316-3196*

Marc Hanheide

*LCAS. University of Lincoln
School of Computer Science
Lincoln, UK
0000-0001-7728-1849*

Georgios Leontidis

*University of Aberdeen
School of Natural and Computing Sciences
Aberdeen, UK
0000-0001-6671-5568*

<https://doi.org/10.31256/Qm1Fu7L>

Abstract—We propose a new pipeline to facilitate deep learning at scale for agriculture and food robotics, and exemplify it using strawberry tabletop. We use this multimodal, autonomously self-collected, distributed dataset for predicting strawberry tabletop yield, aiming at informing both agronomists and creating a robotic attention system. We call this system the augmented agronomist, which is designed for agronomy forecasting, and support, maximizing the human time and awareness to areas most critical. This project seeks to be relatively protective of both its neural networks, and its data, to prevent things such as adversarial attacks, or sensitive method leaks from damaging the future growers livelihoods. Toward this end this project shall take advantage of, and further our existing distributed-deep-learning framework Nemesyst. The augmented agronomist will take advantage of our existing strawberry tabletop in our Riseholme campus, and will use the generalized robotics platform Thorvald for the autonomous data collection.

Index Terms—deep learning, database, agriculture, nemesyst, thorvald, strawberries

I. INTRODUCTION

Machine/Deep learning is becoming a bigger and more important part of our daily lives through the rise of an ever-increasing quantity of available data. 3rd-party services use machine learning in combination with user data for tasks ranging from, natural language processing [5], image recognition, diagnosis [3], detection, classification [6], generation, imputation, broadly prediction; medical diagnosis [2], self-driving cars [8], facial recognition [7], etc. However one area with which deep learning has remained relatively stagnant is in agriculture, where data is scarce, forcing the use of remote sensing datasets or the like, as well as the existing research using classical techniques without many of the recent advances. [1, 4, 13] The primary reason why agriculture has remained relatively constant for this long is likely the lack of, and consistency of data, but also the lack of willingness, and trust of the growers/ agriculturalists to release their potentially sensitive techniques latently in any data they provide. Thus if there is little to no data there can be little advancement with deep learning techniques, meaning prospective research will require self collected data to find any meaningful relations between the features and targets with which to predict accurately and far enough ahead to facilitate timely and effective actions.

This work is supported by a BBSRC CTP-FCR studentship and InnovateUK under grant agreements #105151 and #104587

We contribute our work-in-progress methods and results towards creation of a larger and more accurate plant yield prediction framework which both automates but crucially involves experts in a time-effective and prescribed manner. Our work is facilitated by the RASberry research programme¹, which is a collaboration effort between UoL, Saga Robotics, and BerryGardens, funding autonomous strawberry data collection, under our direct control. This involves the generic expandable Thorvald platform, which is an autonomous robot ready for use in many terrains. Thorvald is an ideal candidate platform to use for our own experiments thanks to its autonomy, and available resources. The only drawback of using strawberries is that they are only grown from late June to early October.

II. PLANT YIELD PREDICTION LITERATURE

As it stands there are many existing methods that have been used to attempt to predict crop yield, using data such as remote sensing [16, 4], satellite image, climate conditions, geolocation data, etc. [9] However, there is high variation in the type, quality, and quantity in the datasets used, with very little from a standard dataset with which to use. [14, 16, 15] The vast majority of papers use remote models relying primarily on: temperature, humidity, precipitation, and soil moisture. Some others attempt image based approaches but lack of data is a serious problem for them [13]. This means as far as yield prediction is concerned it is necessary to create a consistent, and granular dataset [14]. All these papers use many different techniques, with a wide variety of data types such that they only marginally narrow the focus for our data collection efforts to things such as climate conditions, [10] meaning we will have to collect a large variety of data and thereafter assess the correlation to achieve the best results.

III. TECHNIQUES

As depicted in Fig. 1, we use Nemesyst [11, 12] to manage MongoDB instances across all desired Thorvalds. We aggregate this data to the distributed database layer, where all the data is made available to offline back-end deep learning sites. These sites are responsible for training, and model evaluation of neural network models (NNs), along with packaging them

¹<https://rasberryproject.com>

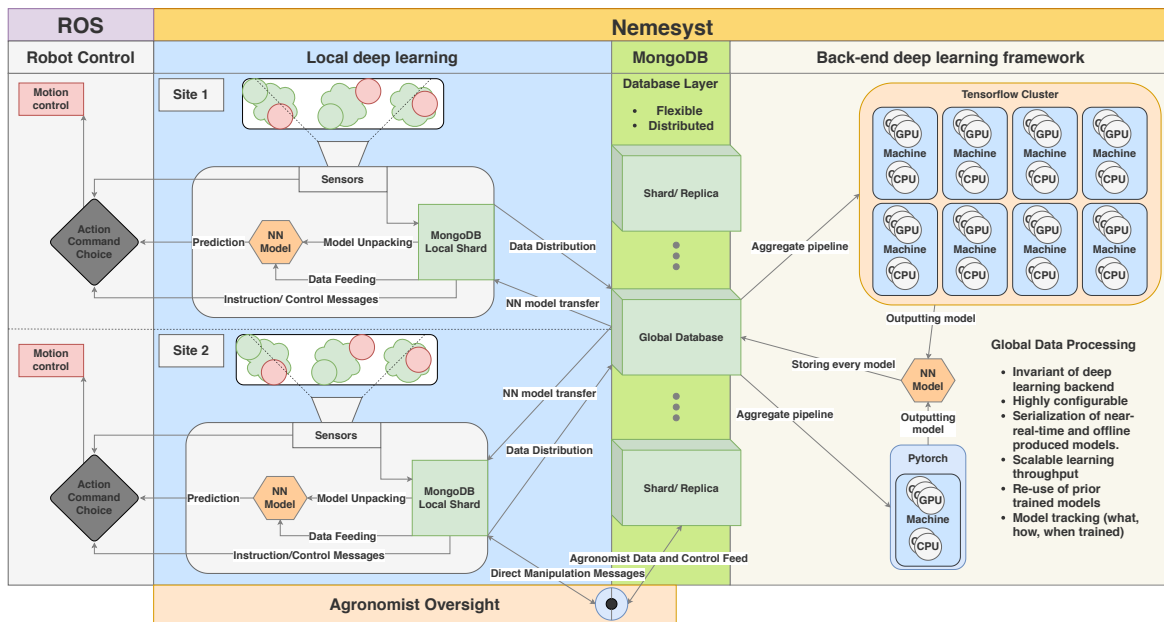


Fig. 1. RASberry data distribution and aggregation pipeline.[11]

back into the database for unpacking and use at a local level. These NNs can then be selected based on their performance and suitability to the application, such as the most performant yield prediction of strawberries versus other berries. The selected NNs used locally can then be used in future to inform decision making processes of the robot, such as attention mechanisms. Attention mechanisms can be used with our databases as a message passing interface to alert and request the attention of specialist agronomists to identify uncertain cases and help with learning along with any immediate control needs.

IV. RESULTS

In the 10-12 fruitful weeks, before the first frosts in October, of June-bearing strawberries, we managed to collect in excess of 40GiB of compressed data (+100GiB if uncompressed), consisting of regular plant imaging, and environment sensing thanks to our Thorvald robots. This data represents a diverse set of scenarios containing:

- 3 still images (RGB, depth, IR, position) every 20cm at various plant angles, and the respective environmental data (temperature, humidity, etc) locally to the camera.
- Continuous environmental data every 15 min. throughout the year, including the lead up to the growing season.
- Stationary camera footage, and its local environment data every 30 minutes for select plants to provide in depth growth and performance data.
- Video based plant image capture during tabletop traversal.
- Yield/ picking data for the number of strawberry punnets collected over time.

However due to constraints in human time, we could only collect yield values twice a week, meaning this last dataset is still relatively small, and would in future need to use

TABLE I
TIME SERIES FORECASTING OF YIELD BY NUMBER OF PUNNETS.

Technique	Mean Absolute Error (Test set)
Vanilla Recurrent Neural Networks	0.210
Long Short-Term Memory	0.381
Gated Recurrent Units	0.155

autonomous pickers to provide more, and consistent labels to our NNs, to train more advanced NNs. Table I thus shows some very early experimental results that demonstrate the ability of various recurrent networks to learn with this limited labeling. Due to the size of the data and how early on in the process we are our results (I) are split plainly 80% training, 20% testing, with around 10-13 epochs for saturation taking less than a few minutes to train using only environmental and yield data at this early stage.

V. CONCLUSION

It has been shown that using distributed Nemesyst database pipelines for data aggregation and modelling as well as distribution in more complex scenarios such as autonomous agricultural data collection, how this can augment the ability of growers to collect data and predict outcomes such as crop yields. A need has been identified for more autonomous data collection to collect more data along with more consistency to feed to NNs to learn more complex representations. Lastly our pipeline can also be used as message passing interfaces for agronomists to monitor, be alerted of any uncertain/ unusual cases, label difficult examples, and potentially control the robots to support their efforts. Our next step is to provide certainty metrics to assess how certain the deep learning models are of the result in such that this can be used for more effective decision making.

REFERENCES

- [1] R Alvarez. "Predicting average regional yield and production of wheat in the Argentine Pampas by an artificial neural network approach". In: *European Journal of Agronomy* 30.2 (2009), pp. 70–77.
- [2] Rachel Anderson et al. "Evaluating deep learning techniques for dynamic contrast-enhanced MRI in the diagnosis of breast cancer". In: *Medical Imaging 2019: Computer-Aided Diagnosis*. Vol. 10950. International Society for Optics and Photonics. 2019, p. 1095006.
- [3] M Biswas et al. "State-of-the-art review on deep learning in medical imaging." In: *Frontiers in bioscience (Landmark edition)* 24 (2019), pp. 392–426.
- [4] Anna Chlingaryan, Salah Sukkarieh, and Brett Whelan. "Machine learning approaches for crop yield prediction and nitrogen status estimation in precision agriculture: A review". In: *Computers and electronics in agriculture* 151 (2018), pp. 61–69.
- [5] Hai Ha Do et al. "Deep learning for aspect-based sentiment analysis: a comparative review". In: *Expert Systems with Applications* 118 (2019), pp. 272–299.
- [6] Hassan Ismail Fawaz et al. "Deep learning for time series classification: a review". In: *Data Mining and Knowledge Discovery* 33.4 (2019), pp. 917–963.
- [7] David Güera and Edward J Delp. "Deepfake video detection using recurrent neural networks". In: *2018 15th IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS)*. IEEE. 2018, pp. 1–6.
- [8] Brody Huval et al. "An empirical evaluation of deep learning on highway driving". In: *arXiv preprint arXiv:1504.01716* (2015).
- [9] Konstantinos G Liakos et al. "Machine learning in agriculture: A review". In: *Sensors* 18.8 (2018), p. 2674.
- [10] Gniewko Niedbała. "Application of Artificial Neural Networks for Multi-Criteria Yield Prediction of Winter Rapeseed". In: *Sustainability* 11.2 (2019), p. 533.
- [11] George Onoufriou. *Nemesyst; Generalised and highly customisable, hybrid-parallelism, database based, deep learning framework*. 2019. URL: <https://github.com/DreamingRaven/nemesyst> (visited on 11/10/2019).
- [12] George Onoufriou et al. "Nemesyst: A Hybrid Parallelism Deep Learning-Based Framework Applied for Internet of Things Enabled Food Retailing Refrigeration Systems". In: *arXiv preprint arXiv:1906.01600* (2019).
- [13] Anup K Prasad et al. "Crop yield estimation model for Iowa using remote sensing and surface parameters". In: *International Journal of Applied Earth Observation and Geoinformation* 8.1 (2006), pp. 26–33.
- [14] Maryam Rahnmooonfar and Clay Sheppard. "Real-time yield estimation based on deep learning". In: *Autonomous Air and Ground Sensing Systems for Agricultural Optimization and Phenotyping II*. Vol. 10218. International Society for Optics and Photonics. 2017, p. 1021809.
- [15] Anna X Wang et al. "Deep transfer learning for crop yield prediction with remote sensing data". In: *Proceedings of the 1st ACM SIGCAS Conference on Computing and Sustainable Societies*. ACM. 2018, p. 50.
- [16] Jiakuan You et al. "Deep gaussian process for crop yield prediction based on remote sensing data". In: *Thirty-First AAAI Conference on Artificial Intelligence*. 2017.

BeetleBot: A Multi-Purpose AI-Driven Mobile Robot for Realistic Environments

Anh Nguyen¹, Erman Tjiputra², Quang D. Tran²

<https://doi.org/10.31256/Gn4Vw9O>

Abstract—We present BeetleBot, a new mobile robot that has been developed to operate in realistic environments. Different from previous state-of-the-art mobile robots, BeetleBot is designed based on the most recent advancements of actuator, mechanical design, and artificial intelligence. These advantages allow the robot to have powerful features including ultramobility, accurate localization, intensive AI ecosystem, and autonomous navigation in normal as well as complex and rough terrain environments. Based on its novel design, BeetleBot can be used in a variety of tasks including but not limited to autonomous delivery, exploration and surveillance, or human-robot interaction. Besides the development of the robot, we also build large-scale open-source simulation models that are fully integrated with ROS and can be used for rapid testing and deployment of any robots. Finally, we show the capability of BeetleBot and the performance of each individual AI module in various tasks and real-world scenarios.

I. INTRODUCTION

Mobile robotics is a long-lasting research and development field in both academia and industry [1]. In general, a mobile robot is designed to autonomously move through its environment without the need of a human operator or on a fixed predetermined path. The mobile robots are usually equipped with an array of sophisticated sensors that enable them to understand, interpret, and act inside the environment, which help them to perform its tasks in the most efficient manner and possible path such as navigating around fixed or moving obstacles (e.g., building, people, and debris). Based on the natural advantages of its design, mobile robots are widely used to work independently or collaboratively with humans in various real-world applications such as logistics (e.g., autonomous delivery), inspection and maintenance (e.g., inspection robot), security and defense (e.g., surveillance robot), agriculture (e.g., fruit picking robot), or urban transportation. Especially, with the recent Coronavirus pandemic, the need of a human-friendly mobile platform that can assist patients in the hospital or be used in healthcare services in general is increasing significantly.

Although there are currently many mobile platforms available, they all share two main limitations: *i)* Most of the mobile robots are designed for a particular problem, hence it is not a trivial task to adapt and use the robot in a new task. *ii)* The integration between hardware, software, and AI for mobile platforms is still a challenge during the deployment. Motivated by these limitations, this paper aims to develop a new mobile platform for remote or autonomous



Fig. 1. An overview of BeetleBot, a multi-purpose AI-driven mobile robot and its external sensor system.

operation in real-world scenarios. Our new robot - BeetleBot - has flexible maneuverability and strong performance to traverse through doorways, over obstacles or rough terrains that may be encountered in indoor, outdoor, or rough terrain environments. Coming with state-of-the-art sensor system and novel AI-driven software, BeetleBot can perform a variety of useful tasks autonomously. Fig. 1 illustrates an overview of BeetleBot and its visual sensor system.

II. DESIGN METHODOLOGY

To build a robot that can perform various tasks in realistic environments, our design methodology focuses on three main aspects: *i)* Robust mechanical design, *ii)* AI-driven software framework, and *iii)* Testing and deployment. We discuss in detail each aspect below.

Design: Maneuverability is the key aspect in mechanic design of a mobile robot. Unlike most of the current mobile robots which use 4-wheels design, we adapt the rocker-bogie mechanism [2] to allow the robot to operate in different terrains. The rocker-bogie design has 6-wheels with no

¹Department of Computing, Imperial College London, UK
a.nguyen@imperial.ac.uk

²AIOZ Pte Ltd, Singapore quang.tran@aioz.io

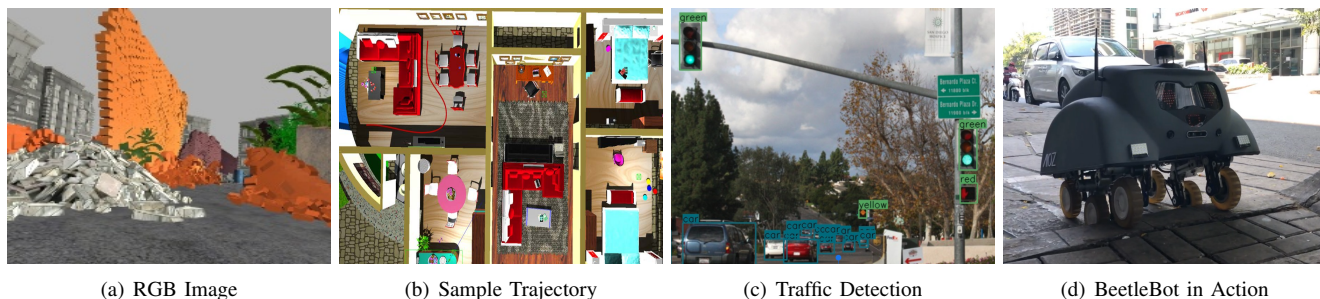


Fig. 2. Illustrations of our simulation environments, detection results and BeetleBot in action. From left to right: (a) A RGB image from our city environment. (b) The path generated by the autonomous navigation network of the robot (denoted as the red line). (c) The traffic recognition result. (d) Our BeetleBot is in action when crossing a slope.

TABLE I
SPECIFICATIONS OF BEETLEBOT

Specifications	
Dimension	L750 × W500 × H600 mm
Environment	Indoor, outdoor, complex terrains
Payload	Max 20kg
Mobility	Max speed: 800 mm/s, 90 degree/s
Wheel	6 wheels, 4 steering drive, 150mm diameter
Battery	10 hours running time, auto charge
Interface	Mobile/Web app, joystick, voice, autonomous
Communication	Wireless 802.11 a/b/g, 4G, digital audio in/out
Computing	8-Core ARM v8.2 64-Bit CPU, NVIDIA Jetson

springs or stub axles for each wheel, allowing the robot to climb over obstacles (e.g., rocks, debris) easily while keeping all 6-wheels on the ground. In order to go over an obstacle, the front wheels are forced against the obstacle by the back two wheels. The rotation of the front wheel then lifts the front of the vehicle up and over the obstacle. The middle wheel is then pressed against the obstacle by the rear wheels and pulled against the obstacle by the front until it is lifted up and over. Finally, the rear wheel is pulled over the obstacle by the front two wheels. During each wheel's traversal of the obstacle, forward progress of the vehicle is slowed or completely halted. Above the rocker-bogie, we currently design a cargo for keeping packages. Most of the visual sensors (e.g, RGB-D, Lidar) are placed in front of the robot. Note that, since we design the robot into separated modules (i.e., wheel mechanism, cargo, and sensor), we can easily customize the design based on the need of different tasks. Table I summarizes the key specifications of our BeetleBot.

AI Ecosystem: Apart from the novel hardware design, our BeetleBot is equipped with state-of-the-art AI methods. Our goal is to close the vision-control loop in robotics and enhance the autonomy of the system. The AI ecosystem of BeetleBot is split into three main applications: Autonomous navigation in normal and complex environments with deep neural network; Visual recognition and localization (e.g., face recognition, pedestrian, traffic light detection, sensor fusion); Human-robot interaction (e.g., emotion estimation and reaction, automatic question and answering). Note that the AI ecosystem of BeetleBot has state-of-the-art accuracy

but is very lightweight and achieves real-time performance. Overall, we can deploy the whole software system on an NVIDIA Jetson Xavier embedded board. Due to the page limitation, we refer the readers to Fig. 2 and the video on our website for the full demonstration:

<https://sites.google.com/site/beetlebotrobot/>

Testing Environments: Since robotics is a multidisciplinary field that requires heavy engineering, software testing before deployment plays a critical role in our framework. In particular, we perform extensive experiments to evaluate the effectiveness of the robot in both simulation and real environments. We employ Gazebo [3] simulation and build large-scale environments to test the robot. In particular, we create 539 3D object models to build three kind of simulation environments: house (indoor), city (outdoor), and natural cave (rough terrain). These objects are used to build 30 environments in total (i.e., 10 instances for each environment). In average, the simulated house environments are built with approximately 130 objects in an area of $400m^2$. The city has 275 objects and spread in $3,000m^2$ while the natural cave environments are built with 60 objects in approximately $4,000m^2$ area. These simulation environments can be used to collect data for different learning-based methods (e.g., autonomous navigation, object detection) as well as deploying the trained model. Note that all of the simulation environment, robot model, and sensors are fully compatible with ROS and Gazebo, hence ease the prototype and deployment process. To encourage further research, these large-scale environments and our collected datasets will be released upon acceptance.

To conclude, we present BeetleBot, a new multi-purpose AI-driven mobile robot. We hope that our experience and open-source resources could benefit the robotic community in the future.

REFERENCES

- [1] R. Siegwart, I. R. Nourbakhsh, and D. Scaramuzza, *Introduction to autonomous mobile robots*. MIT press, 2011.
- [2] D. P. Miller and T.-L. Lee, "High-speed traversal of rough terrain using a rocker-bogie mobility system," in *Space and Robotics*, 2002.
- [3] N. Koenig and A. Howard, "Design and use paradigms for gazebo, an open-source multi-robot simulator," in *IROS*, 2004.

Determining shape of strawberry crops with spherical harmonics

Justin Le Louedec
School of Computer Science
University of Lincoln, UK
jlelouedec@lincoln.ac.uk

Grzegorz Cielniak
School of Computer Science
University of Lincoln, UK
gcielniak@lincoln.ac.uk

<https://doi.org/10.31256/Mc8H11A>

Abstract—Shape descriptor and shape reconstruction are two challenges found in computer vision and graphics as well as in perception for robotics, especially for some fields such as agri-robotics (robotics for agriculture). Being able to offer a reliable description of shape that can also translate directly into a high fidelity model of the shape, would be of high importance for a lot of applications such as phenotyping or agronomy. In this paper we report on our work on using spherical harmonics to offer efficient representation of strawberry shapes and we validate them by reconstructing the fruits. The reconstruction achieve extremely close results to the original shape (less than 1% deviation) and the representation reduce the complexity and improve compactness by a large factor (minimum 100).

Index Terms—3D, spherical harmonics, point clouds, phenotyping

I. INTRODUCTION AND RELATED WORK

An accurate description of crops shape is an important challenge in horticulture. Automating their creation and allowing a complete 3D reconstruction of the objects from it, would improve greatly phenotyping or other agronomy tasks. In this work, we use the mathematical concept of spherical harmonics to create a representation of strawberries' shape and study their fidelity and accuracy by reconstructing the fruits with them. Furthermore, this new representation offers a more compact and efficient representation of the shape than using directly points.

The use of spherical harmonics as a 3D shape representation was first proposed in [1]. The practical use of that representation as a rotation-invariant feature descriptor was later introduced in [2]. Some of the drawbacks related to poor results with discontinuous surfaces were addressed by the introduction of weighted spherical harmonics [3]. Applications of Spherical Harmonics as shape descriptors include [4] who applied them to summarise shapes of sand particles obtained from X-ray micro-tomography. In [5] the authors propose to use Spherical Harmonics to represent the 3D shape of agricultural materials such as grains and show this representation to be a more compact and efficient representation than previous methods and meshes/point clouds. Our work takes inspiration from [5], reducing the process complexity to match the symmetrical characteristics of strawberries.

II. METHODOLOGY

To explain our approach we detail in this section how spherical harmonics transformations are computed and implemented.

Spherical harmonics represent a solution to Laplace's equation, defining a complete set of angular functions in spherical coordinates. We can use it to define the surface f of an object as described in Equation (1). In this equation the surface is formulated as the combination of all spherical harmonics $Y_l^m(\theta, \phi)$ with their respective coefficients c_l^m . In [4] they are computed over three components: $c_l^m = (c_{xl}^m, c_{yl}^m, c_{zl}^m)^T$, and each point of the shape are associated to its spherical coordinates with smoothing methods over each of these components.

$$f(\theta, \phi) = \sum_{l=0}^{\infty} \sum_{m=-l}^{m=l} c_l^m Y_l^m(\theta, \phi) \quad (1)$$

We detail in Equation. 2 the spherical harmonics for degree l and order m , with P being the Legendre functions associated with it. Spherical harmonics are computed on θ (azimuthal coordinates) between $[0, 2 * \pi]$ and ϕ (polar coordinates) between $[0, \pi]$.

$$Y_l^m(\theta, \phi) = \sqrt{\frac{(2l+1)(l-m)!}{4\pi(l+m)!}} e^{im\theta} P_l^m(\cos(\phi)) \quad (2)$$

To go from the spherical mapping of an object to its spherical harmonics representation we need the forward transformation defined in Equation (3). The implementation of these integrals varies from paper to paper [6], [7]. We use the representation of these harmonics as detailed in Equation (4). All the harmonics for a given position on the sphere can be represented as a Matrix \hat{Y} for all degree and orders up to m and l . Using the definition of the inverse spherical harmonics in Equation (1), we can define the dot product of the coefficient by the harmonics matrix. Finally, these coefficients can be expressed as the dot product of our mapping on the sphere by the inverse of the harmonics matrix. Also efficient in terms of memory and simplicity, the forward spherical harmonics transform is not computationally efficient but can be accelerated through lookup tables and down-sampling of our point clouds.

$$c_l^m = \int_0^\pi \int_0^{2\pi} d\phi d\theta f(\theta, \phi) Y_l^m(\theta, \phi) \quad (3)$$



Fig. 1. Example of strawberry point clouds (left example for each pair) and their reconstructions(right example for each pair), for two particular shapes.

$$\begin{aligned}
 \hat{Y} &= Y_l^m(\theta, \phi) \\
 f(\theta, \phi) &= \hat{C}\hat{Y} \\
 \hat{C} &= \hat{Y}^{-1}f(\theta, \phi)
 \end{aligned} \tag{4}$$

As spherical harmonics are intrinsically symmetric over ϕ , one can, similarly to [5] we take separated portions of the shape and generate their spherical harmonics representation independently for more precise control over the shape and its details. We align the strawberries over their principal axis along the horizontal axis and recreate the two halves separately. To recreate each strawberry we change their points to spherical coordinates and using the forward harmonics transform, compute the coefficients up to an arbitrary degree. We then use them to recreate the strawberry shape from the unit sphere. Our sample set is composed of 15 3D point clouds of strawberries from [8], representing all the different shapes as described in the industry phenotyping reports.

III. RESULTS

In this section, we present some of our results and analyse them toward their possible applications.

In Table I we show the quantitative results obtained using spherical harmonics for the reconstruction of our set of strawberries. We compare the average volume and the average surface area, with their deviation. With such a high degree of harmonics, the deviation is very small for most of our samples and his coming from fined details rather than important shape features.

In Figure 1 we present some qualitative results of two different strawberries point clouds reconstructed using spherical harmonics. We use 30 degrees of harmonics, for a very precise and detailed reconstruction (seed bumps and calyx shape with a lot of details). At this level, only a few details are smoothed out and only very sharp angles and edges are not recreated.

In Figure 2 we present an example of the coefficients for one of the strawberries. We showcase the coefficients responsible for most of the shape of the strawberry omitting the one influencing more fine details. As per our way of processing both halves of the strawberry independently, we have them separated by the centre of the chart. We also omit

TABLE I

THE VOLUME AND SURFACE AREA ESTIMATION RESULTS FROM THE RECONSTRUCTION PROCESS OF 15 3D MODELS OF STRAWBERRIES (NOTE THE OBJECTS WERE SCALED DURING CAPTURE PROCESS).

	Original	Reconstructed	Deviation
Volume ($\mu \pm \sigma$) cm^3	45 ± 46	45 ± 46	$\sim 1\%$
Surface area ($\mu \pm \sigma$) cm^2	280 ± 160	277 ± 160	$\sim 1\%$




Fig. 2. Spherical harmonics coefficients responsible for main shape information

the first coefficient, associated with the first harmonic as it is responsible for the round shape of the strawberry, present in every strawberry and non-discriminative other than for scale. The first twenty harmonics are more important for the shape and specificity of the global strawberry, while the following one corresponding to more fine details. This representation offers a compact (maximum a thousand parameters for high precision reconstruction) way to represent the strawberry shape with relatively close fidelity, compared to the hundreds of thousands points composing the object.

IV. CONCLUSION AND DISCUSSION

We have presented an efficient way of describing and reconstructing strawberry shape information. It offers a compact way of representing the shape of such objects and simplifies future work done on processing this shape information. Future work would include reconstruction from partial view, automatic phenotyping from the compact representation, generation of new fruit models etc.

REFERENCES

- [1] Ch. Brechbühler, G. Gerig, and O. Kübler. Parametrization of closed surfaces for 3-d shape description. In *Computer Vision and Image Understanding*, 1995.
- [2] Michael Kazhdan, Thomas Funkhouser, and Szymon Rusinkiewicz. Rotation invariant spherical harmonic representation of 3d shape descriptors. In *Proceedings of the 2003 Eurographics/ACM SIGGRAPH Symposium on Geometry Processing*, 2003.
- [3] Moo Chung, Kim Dalton, Li Shen, Alan Evans, and Richard Davidson. Weighted fourier series representation and its application to quantifying the amount of gray matter. In *IEEE Transactions on Medical Imaging*, 2007.
- [4] Budi Zhao, Deheng Wei, and Jianfeng Wang. Application of spherical harmonics analysis on lbs particles and lbs fragments. In *European Physical Journal Web of Conferences*, 2017.
- [5] Urtė Radvilaitė, Álvaro Ramírez-Gómez, and Rimantas Kačianauskas. Determining the shape of agricultural materials using spherical harmonics. In *Computers and Electronics in Agriculture*, 2016.
- [6] Banday A. J. Wandelt B. D. Hansen F. K. Reinecke M. Bartelmann M. Górski K. M., Hivon E. Healpix: A framework for high-resolution discretization and fast analysis of data distributed on the sphere. In *The Astrophysical Journal*, 2005.
- [7] Kevin M. Huffenberger and Benjamin D. Wandelt. Fast and exact spin-s spherical harmonic transforms. In *The American Astronomical Society*, 2010.
- [8] Bo Li Joe Q. He, Richard J. Harrison. A novel 3d imaging system for strawberry phenotyping. In *Plant Methods*, 2017.

Towards Gamification of the Ramp-up Process for Industry 4.0

Bryony Foley, Pedro Ferreira, Melanie Zimmer
Wolfson School of Mechanical, Electrical and Manufacturing Engineering
Loughborough University
 Loughborough, UK

B.E.Foley-16@student.lboro.ac.uk, p.ferreira@lboro.ac.uk, m.zimmer2@lboro.ac.uk

<https://doi.org/10.31256/Jp9Yd3M>

Abstract—Keeping personnel motivated and overall company morale high is a crucial aspect of employee satisfaction. In manufacturing, various tasks can be found that can put a strain on personnel, one of them being the so-called ramp-up process. Ramp-up requires human intervention to bring a system to full production, but the process of achieving a successful ramp-up can be very cumbersome and lengthy. As this is usually a trial-and-error approach, spirits can get very low over time if the situation appears to only be slowly improving. For this purpose, gamification is introduced into the ramp-up process for Industry 4.0 to improve the ramp-up experience for operators. This paper offers an overview of the idea, current and future work of a web application for ramp-up that includes the major principles of gamification. The presented work is part of an ongoing student project work undertaken with the Intelligent Automation Centre at Loughborough University.

Keywords— *Gamification, Ramp-up, Industry 4.0, Game Design, Web App.*

I. INTRODUCTION

Over the last few years, the focus of the manufacturing industry has been on a new paradigm, often referred to as Industry 4.0 [1]. In essence, the focus of Industry 4.0 is achieving end-to-end digitisation allowing for data creation, sharing and analysis to achieve a more sustainable production [2]. One area in the production life cycle that is known to be very time-consuming, is the so-called ramp-up phase [3], [4]. During that phase a production system is taken from low- to high-volume production, which requires tweaking of system and process settings until the required functionality, product quality and performance are reached [5]. This can be a tedious undertaking as no two systems are the same and no manual for ramping up a system exists [6].

In this context, this ongoing work proposes to use the concept of gamification to assist in the ramp-up process. The idea of increasing the performance during ramp-up using a gamification approach is not new [7]. However, work and knowledge in this area are still very limited and to the best of the authors' knowledge, no web application for ramp-up exists to date. For this research work, gamification is understood as "a tool to increase productivity and morale within a company by using elements of game playing, which can also be used for encouragement". The underlying tenet is that through increasing motivation through the gamification introduction a smoother ramp-up process can be ensured as reinforcing training can be given to the operator. For this purpose, this work will develop a web application to assist the user during the ramp-up process that contains elements from gamification.

A. Project Aims and Objectives

As part of this, research into the design and the development of a useful gamification web app for the ramp-up process, the following aims and objectives will be addressed:

- Identify key aspects of the project for Literature Review including defining Gamification, Ramp-up Process and Industry 4.0.
- As part of the literature review find relevant case studies.
- Determine how Gamification for Ramp-up Process can be linked and how Industry 4.0 links with Gamification and Ramp-up as part of the literature review.
- Design one proposal for experimentation, use case study information found to develop proposal.
- Validate the chosen proposal within the laboratory and collect data from an experiment.
- Gather feedback on ideas through a questionnaire to ensure that game designed matches gaming habits in current society.

II. GAMIFICATION APP DESIGN

To start the project, extensive research was conducted to ensure that all topics were understood. From this, a Product Design Specification (PDS) was created that assists in the design process. The initial design process included designing the social side of the game including the user profile page, the forum page and game mascots. The initial game designs included a quiz-based game, a timer game and 'Build a Toolbox' game. However, the latter seems to have the most potential to be developed. These aspects are briefly described in more detail below.

A. Profile Page

Every user will be able to set up a profile page. Some of the information included will be achievements, skills the user has and any information the user would like to include. Other information they could add may include their job title, education, a fun fact, and their company details. The user will also be required to create an avatar that would act as their virtual representation. The user will be able to choose the sex, skin colour, hair colour, hairstyle, eye colour, clothing etc.

B. Forum Page

The forum page is a way for users to interact with one another. Some of the aspects of the forum page include a news feed for the company's latest news, a way to see other people's posts and a way to react and comment on them.

C. Game Mascots

The idea of a mascot is being considered to capture the user interest in the game. The mascot will fit into the game by giving the user the instructions of how to use the game along with any training the company deem appropriate. The initial ideas shown in **Fig. 1** are "Robbie the Robot", "Ruby Ramp Up" and "Engineer Dave". However, any mascot idea would

REFERENCES

- [1] T. Stock and G. Seliger, "Opportunities of Sustainable Manufacturing in Industry 4.0," in *Procedia CIRP*, 2016, vol. 40, pp. 536–541, doi: 10.1016/j.procir.2016.01.129.
- [2] E. Paravizo, O. C. Chaim, D. Braatz, B. Muschard, and H. Rozenfeld, "Exploring gamification to support manufacturing education on industry 4.0 as an enabler for innovation and sustainability," *Procedia Manuf.*, vol. 21, pp. 438–445, 2018, doi: 10.1016/j.promfg.2018.02.142.
- [3] H. Winkler, M. Heins, and P. Nyhuis, "A controlling system based on cause–effect relationships for the ramp-up of production systems," *Prod. Eng. Res. Dev.*, vol. 1, no. 1, pp. 103–111, 2007, doi: 10.1007/s11740-007-0011-2.
- [4] C. Terwiesch and R. E. Bohn, "Learning and process improvement during production ramp-up," 2001. Accessed: Mar. 06, 2019. [Online]. Available: https://ac.els-cdn.com/S0925527300000451/1-s2.0-S0925527300000451-main.pdf?_tid=39ded4ed-165f-4083-b6cb-5b87372e2eb9&acdnat=1551874315_2063ac645a27f628d3ac2e2e2d73ecab.
- [5] S. C. Doltsinis, S. Ratchev, and N. Lohse, "A framework for performance measurement during production ramp-up of assembly stations," *Eur. J. Oper. Res.*, vol. 229, no. 1, pp. 85–94, 2013, doi: 10.1016/j.ejor.2013.02.051.
- [6] C. H. Glock and E. H. Grosse, "Decision support models for production ramp-up: A systematic literature review," *Int. J. Prod. Res.*, vol. 53, no. 21, pp. 6637–6651, 2015, doi: 10.1080/00207543.2015.1064185.
- [7] A. Kampker, C. Deutskens, K. Deutschmann, A. Maue, and A. Haunreiter, "Increasing ramp-up performance by implementing the gamification approach," *Procedia CIRP*, vol. 20, no. C, pp. 74–80, 2014, doi: 10.1016/j.procir.2014.05.034.
- [8] M. Zimmer, P. Ferreira, P. Danny, A. Al-Yacoub, N. Lohse, and V. Gentile, "Towards a decision-support framework for reducing ramp-up effort in plug-and-produce systems," *Proc. - 2019 IEEE Int. Conf. Ind. Cyber Phys. Syst. ICPS 2019*, pp. 478–483, 2019, doi: 10.1109/ICPHYS.2019.8780369.

Modelling variable communication signal strength for experiments with multi-robot teams

Tsvetan Zhivkov
 Dept of Engineering
 King's College London, UK
 tsvetan.zhivkov@kcl.ac.uk

Elizabeth I Sklar
 Lincoln Institute for Agri-food Technology
 University of Lincoln, UK
 esklar@lincoln.ac.uk

<https://doi.org/10.31256/Ld2Re8B>

Abstract—Reliable communication is a critical factor for ensuring robust performance of multi-robot teams. A selection of results are presented here comparing the impact of poor network quality on team performance under several conditions. Two different processes for emulating degraded network signal strength are compared in a physical environment: modelled signal degradation (*MSD*), approximated according to increasing distance from a connected network node (i.e. robot), versus effective signal degradation (*ESD*). The results of both signal strength processes exhibit similar trends, demonstrating that *ESD* in a physical environment can be modelled relatively well using *MSD*.

I. INTRODUCTION

Reliable communication is one of the key requirements for successful operation of a multi-robot team—without it, team coordination is difficult. Many works focus on network optimisation and communication-aware motion planning for multi-robot systems [1]–[6]. The aforementioned works use either optimisation-based control or consensus algorithms for centralised or distributed multi-robot communication and resilience against either uncooperative robots or maximising communication performance. The majority of research conducted with physical multi-robot teams takes place in controlled indoor laboratory settings supported by robust network infrastructure. However, many recognise the growing need for reliable wireless communication in multi-robot systems in a wide range of environments and applications.

In earlier work [7], we introduced a behaviour-based, network-aware control algorithm which aims to prevent loss of communication by keeping robots within range of networked nodes (i.e. neighbouring robots), even if the network is severely crippled, e.g. due to reduced signal strength or high percentage of dropped network packets. In order to test this behaviour, we previously modelled two of the most common network problems [7]–[9]: *simulated packet-loss (SPL)*, which drops a pre-defined percentage of messages (0%,25%,50%,75%), and a method of estimating network *signal strength*, which periodically examines the distance between robots and based on a pre-determined threshold warns the robots if they are likely to disconnect. The contribution presented here briefly summarises an extension of our experimental framework, adding *modelled signal degradation (MSD)*¹ and *effective signal degradation (ESD)* to allow more

¹In previous works *MSD* was denoted *SSD* (simulated signal degradation)

comprehensive testing of behaviours that respond to network failure.

II. APPROACH

In our previous work [8], we performed preliminary analysis of how multi-robot team performance was impacted by *simulated packet-loss (SPL)*. We then extended this work by introducing our novel *Leader-Follower (LF)* behaviour designed to respond to network weaknesses, described in [7] and evaluated using our Multi-Robot Communication testbed (*MRCComm*) [9]. Our contributions can be separated into two parts: (1) modelling various aspects of network quality (*MRCComm*); and (2) controlling robot behaviour in response to changes in the network quality (*LF*).

The network type is the communication medium that is used to transfer messages between robots. In the experiments demonstrated here we analyse *Ad Hoc (AH)* network type, which is an uncommon network to use in robotics as it has no infrastructure. A robot is used to initialise an *AH* network and the rest of the team connect directly to the that robot. Communication is peer-to-peer, and increasing the proximity of neighbouring robots (i.e. causing signal degradation) negatively impacts communication quality. Moreover, certain assumptions are made about the network type to make our problem more tractable. Firstly, it is assumed that signal-to-noise ratio (SNR) experiences uniform loss and that interference from other devices is negligible. Secondly, the final assumption (limitation) for *AH* is that after 9.0 meters robots can no longer communicate, no matter the experiment configuration. From the tests conducted in our physical environment, a separation greater than 9.0 meters between robots would consistently return a signal strength lower than -60 dBm, which is considered a poor signal.

Here we define two network thresholds that are related to signal strength in order to model different types of communication failure: *MSD* and *ESD*.

MSD is modelled in the same environment that the multi-robot team experiments are conducted in (i.e. indoor office building). Furthermore, it is modelled using two separately trained Support Vector Regression (SVR) models with Radial Basis Function kernels (RBF), for both direct and obstructed line-of-sight signal strength. The data for the models was obtained by dividing and performing two levels of granularity tests (i.e. at a rate of 0.1 m and 1 m steps) on signal strength

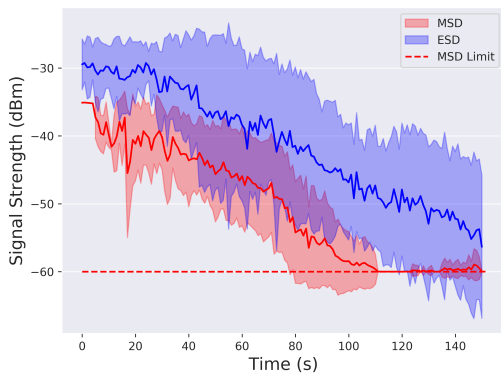


Fig. 1. Signal strength of *ESD* and *MSD* samples, measured over time

with increasing distance². The fine granularity test was to get more accurate data and the coarse granularity test was to allow for acceptable estimation of signal strength with increasing distance.

ESD is a new process implemented in *MRCComm* that queries all network devices connected to the *AH* network at a frequency of 2 Hz, obtaining the signal strength from each device. Executing this process in parallel with our package provided three main benefits. Firstly, if a robot experiences hardware malfunction and loses some/all sensor functionality, assuming the network device is still functional, it would continue pulsing a signal. Secondly, if a robot experiences a software crash caused by an internal or external issue of *MRCComm*, it would continue pulsing a signal. Finally, this functionality of *MRCComm* would allow almost any type of device to run this process and send out a signal, therefore increasing its utility.

Figure 1 shows the mean *MSD* and *ESD* signal strength over a range of time t , measured in seconds from when experiments start at $t_0 = 0$ up until $t = 150$ s (i.e. the first 150 seconds of each experiment). To improve comparability between *MSD* and *ESD* the results were obtained from experiments using the same configuration. It was expected that *MSD* and *ESD* would yield near identical signal strength results. However, as observed in Figure 1 the distributions are different, but the trends are similar and the standard deviations overlap in many cases. The difference between *MSD* and *ESD* is to be expected, since *MSD*'s accuracy is limited to the SVR models used. The SVR models are based on a small sample size of the physical environment and to improve the predicted signal strength either the sample size needs to be expanded to include a data point at every possible location of the physical environment or a more accurate model needs to be introduced. However, the most simple approach can be to apply an offset to the current SVR models to artificially boost the predicted signal strength and thereby improve *MSD*.

MRCComm provides a response to the network parameters in the form of one of two robot behaviours, which are

²Each data point reading was repeated 30 times to get an average result.




Fig. 2. Comparison of the average number of failed tasks for *MSD* vs *ESD*, using *AH* and $\{SPL0, SPL25, SPL50, SPL75\}$, for both *LF* (on the left of each sub-plot) and *NB* (on the right of each sub-plot) behaviours. The *LF* behaviour recovers from all perturbations so there are no failed tasks. *NB* reveals the effects of *MSD* versus *ESD*.

the baseline *Non-responsive Behaviour (NB)* and the novel *Leader-Follower (LF)*. The *NB* behaviour does not react to any network parameters and simply enables robots to attempt to complete tasks assigned to them, even if the network drops out. In contrast, *LF* responds to different network perturbations and thresholds. It uses a high-level grouping technique to alert and force robots to move together in the event that signal strength is too weak (degraded).

III. EXPERIMENTAL RESULTS

A series of experiments were run using the *MRCComm* framework [7]–[9] and an *AH* network, which is initially established and then maintained by the robot team using the Robot Operating System (*ROS*) framework and the multi-master package (*FKIE* [10]). The experiments were run in an office building and each experiment was performed 5 times. Experiments were run with 3 Turtlebot2 robots performing 7 independent (i.e., not constrained by any other task), single-robot observation tasks, starting in a clustered formation. A network perturbation and threshold are applied to each experiment: an SPL (with packet loss varying from 0% to 75%), followed by either *MSD* or *ESD*.

Figure 2 shows a sample of experiment configurations and the number of tasks that failed to be communicated per robot in the team. The results for this performance metric showed that *no* robot using *NB* successfully communicated *all* their task status messages. However, Figure 2 shows that robots using the *LF* behaviour managed to always communicate *all* their task status messages.

IV. SUMMARY

The results presented in this paper show that communication of shared messages was successfully carried out on a physical multi-robot team using the novel *LF* behaviour. Furthermore, we observe that the *MSD* parameter has a similar signal strength trend to the *ESD* parameter. Our objective of having a behaviour capable of reacting and mitigating common network issues has been achieved. In future work we will focus on optimising and expanding the capability of the behaviour to deal with more challenging tasks and environments.

REFERENCES

- [1] J. Le Ny, A. Ribeiro, and G. J. Pappas, "Adaptive communication-constrained deployment of unmanned vehicle systems," *IEEE Journal on Selected Areas in Communications*, vol. 30, no. 5, pp. 923–934, June 2012.
- [2] M. Lindhé and K. H. Johansson, "Adaptive exploitation of multipath fading for mobile sensors," in *2010 IEEE International Conference on Robotics and Automation*, May 2010, pp. 1934–1939.
- [3] A. Ghaffarkhah and Y. Mostofi, "Communication-aware motion planning in mobile networks," *IEEE Transactions on Automatic Control*, vol. 56, no. 10, pp. 2478–2485, Oct 2011.
- [4] L. Sabattini, C. Secchi, N. Chopra, and A. Gasparri, "Distributed control of multirobot systems with global connectivity maintenance," *IEEE Transactions on Robotics*, vol. 29, no. 5, pp. 1326–1332, Oct 2013.
- [5] K. Saulnier, D. Saldaña, A. Prorok, G. J. Pappas, and V. Kumar, "Resilient flocking for mobile robot teams," *IEEE Robotics and Automation Letters*, vol. 2, no. 2, pp. 1039–1046, April 2017.
- [6] A. Mannucci, L. Pallottino, and F. Pecora, "Provably safe multi-robot coordination with unreliable communication," *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 3232–3239, Oct 2019.
- [7] T. Zhivkov, E. Schneider, and E. I. Sklar, "Establishing continuous communication through dynamic team behaviour switching," in *UK-RAS19 Conference: Embedded Intelligence: Enabling & Supporting RAS Technologies*, 2019.
- [8] —, "Measuring the effects of communication quality on multi-robot team performance," in *Towards Autonomous Robotic Systems: 18th Annual Conference (TAROS)*, 2017.
- [9] —, "MRComm: Multi-robot communication testbed," in *Towards Autonomous Robotic Systems: 20th Annual Conference (TAROS)*, 2019.
- [10] S. H. Juan and F. H. Cotarelo, "Technical report: Multi-master ros systems," Institut de Robòtica i Informàtica Industrial (IRI), Barcelona, Spain, Tech. Rep. IRI-TR-15-1, January 2015. [Online]. Available: <http://www.iri.upc.edu/files/scidoc/1607-Multi-master-ROS-systems.pdf>

Towards bio-inspired fruit detection for agriculture

Zeke Hobbs
Sheffield Robotics
The University of Sheffield
Sheffield, UK
z.hobbs@sheffield.ac.uk

Tom Duckett
LCAS
The University of Lincoln
Lincoln, UK
tduckett@lincoln.ac.uk

Simon Pearson
LIAT
The University of Lincoln
Lincoln, UK
spears@lincoln.ac.uk

Michael Mangan
Sheffield Robotics
The University of Sheffield
Sheffield, UK
m.mangan@sheffield.ac.uk

<https://doi.org/10.31256/Le4Ce5S>

Abstract—Automation presents a potential solution to agricultural challenges such as worker shortages, invasive pest species and decreasing profit margins. Many technical challenges remain including visual detection of soft fruits. State-of-the-art fruit detectors increasingly rely on deep learning models and standard imaging devices which achieve excellent performance but require significant effort to train and deploy, limiting their uptake. The fruit-fly species *Drosophila suzukii* successfully pinpoints a host of soft fruits visually presenting an excellent model system which can inspire a new class of fruit detector using sparse computational and training resource. Here we present an outline of the features of fruit fly vision that appear to underlie their fruit finding abilities and present a specification for a novel robot imaging system to verify hypotheses in real agricultural settings.

Index Terms—Agriculture, Fruit Detection, Computer Vision, Bioinspired, Fruit Fly, Multi-Spectral Imaging, Novel Sensing

I. INTRODUCTION

For essential agricultural tasks such as yield prediction, assessment of fruit health and ripeness and for harvesting to be automated, visual fruit detection systems must be: sufficiently robust to function in industrial settings with known challenges of lighting variance and occlusions; computationally efficient to be deployed on small, cheap robot platforms; perform in real-time, and ideally function for a variety of fruits (and their varieties).

A. Engineered approach

Deep neural networks (DNN) represent the state-of-the-art methodology in fruit detection, with models such as MangoY-OLO [1], DeepFruits [2], and most recently L*a*b*Fruits [3] all achieving excellent detection scores when tested on realistic datasets while striving to reduce computational cost (see [1], [3] for discussion). Performance improvements can be traced to innovations in both imaging technologies (RGB [1], [3], 3D depth [4] and RGB + infrared [2] cameras) and network architectures moving from multi stage detectors with coarse feature maps [5] to single-stage detectors [6] and multi-scale feature maps [7], [8]. Most relevant to this work, L*a*b*Fruits [3], demonstrated the utility of looking to nature for inspiration by using a colour opponent process inspired by human visual perception to increase performance.

This work is funded by the PhD Scholarship: AHDB Horticulture Project CP 170 Bioinspired vision systems for automated harvesting.

Yet a concern for DNN models is the need for (re)training for each fruit type or variety not found in the original training dataset. This process often requires manual collection and annotation of images, and expert involvement in retraining models, presenting a potential barrier for fruit growers (see Fig. 1 Upper Panel for a typical DNN training pipeline).

B. Nature's approach

Our research has identified the fruit fly *Drosophila suzukii* (*DS*) as an ideal model for inspiring a new class of low computation, zero-retraining fruit detector. *DS* are an invasive pest species to Europe that can visually locate a variety of soft fruit types of a variety of colours (e.g. strawberries, raspberries, blueberries) despite possessing low-resolution eyes and a highly constrained nervous system [9].

Although *DS* uses a combination of olfactory and vision cues to find fruit, trapping studies demonstrate that bright colours alone actively attracting fruit flies [10]. Underpinning their colour detection abilities are eyes that detect light in spectra outside the range detected by imaging systems used by deep learning algorithms to date. Specifically, *DS* have two photo-receptors sensitive to ultraviolet (UV) light (at 335nm and 355nm), as well as green (530nm) and blue (460nm) light. UV light has already been demonstrated as a powerful cue for segmenting foreground objects from the solar background [11]. Moreover, a recent study has suggested that a similar colour opponent mechanism in this non-visible spectra provides fruit flies with their impressive fruit finding abilities [12]. The short life-span of *DS* would favour a hard-wired visual processing pipeline allowing fruits of various kind to be identified without a costly learning phase with associated benefits for artificial systems (see Fig. 1 Lower Panel for the proposed bioinspired pipeline).

II. DROSOPHILA EYE CAMERA SPECIFICATIONS

To verify whether non-visible light offers benefits for generic fruit detection we propose to construct a novel imaging system inspired by *DS* (See Table I for technical specification and comparison with state-of-the-art), and to collect data in real horticultural settings allowing bench-marking against state of the art models.

III. OUTLOOK

DS offer an excellent inspiration to developing highly resilient, but computationally cheap fruit detection systems. The

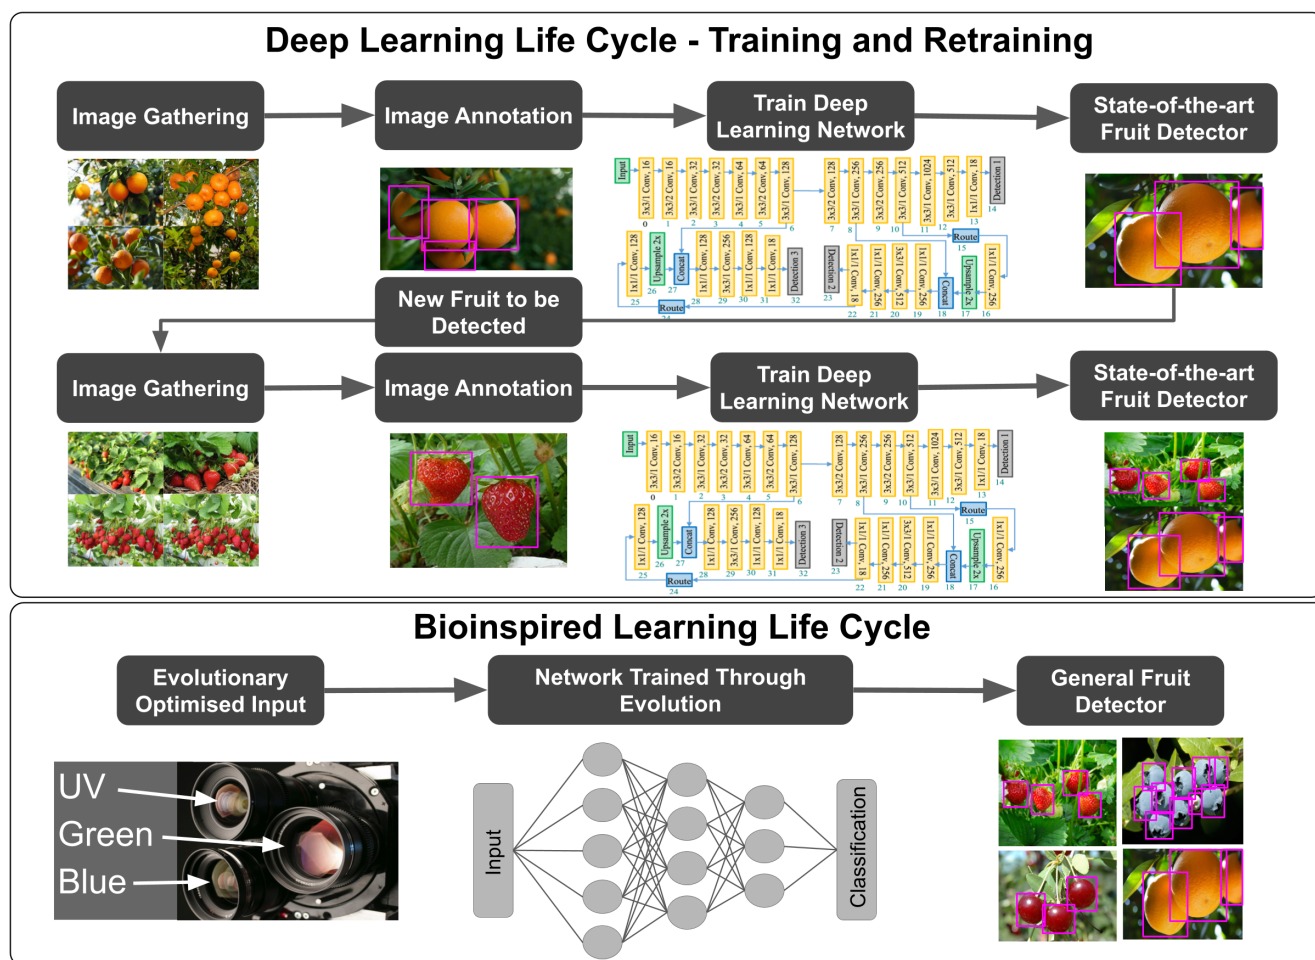


Fig. 1. Deep learning life cycle: Training and retraining - Steps of deep learning, 1) Gather images, 2) Annotate images, 3) Train neural network (Network architecture from [1], 4) High accuracy fruit detector. For additional fruits to be detected not found in the original dataset, repeat steps one, two and three. Bioinspired learning life cycle - 1) Evolutionary optimised input sensitives to UV (335nm and 355nm), green (530nm) and blue (460nm). 2) The network is trained through evolution to find fruits using a small brain requiring no training. 3) General fruit detector able to detect all fruits.

TABLE I
COMPARISON BETWEEN NATURE AND ENGINEERED APPROACHES.

Points of comparison	Detectors			
	DS	[1]	[3]	[2]
Input pixel count	700	4.1MP	0.9MP	2MP
Input type	UVGB	RGB	RGB	RGB + IR
FPS	100 ^[13]	14	26	5
Retraining for new fruit	No	Yes	Yes	Yes

custom camera detailed above and currently under development will play a crucial role in understanding how the humble fruit fly achieves such impressive feats which we will directly apply to solve real agricultural problems. Field data collection scheduled for Spring/Summer 2020.

REFERENCES

- [1] A. Koirala, K. B. Walsh, Z. Wang, and C. McCarthy. Deep learning for real-time fruit detection and orchard fruit load estimation: benchmarking of 'MangoYOLO'. *Precision Agriculture*, pages 1–29, 2 2019.
- [2] Inkyu Sa, Zongyuan Ge, Feras Dayoub, Ben Upcroft, Tristan Perez, and Chris McCool. Deepfruits: A fruit detection system using deep neural networks. *Sensors (Switzerland)*, 16(8), 2016.
- [3] Raymond Kirk, Grzegorz Cielniak, and Michael Mangan. L*a*b*Fruits: A Rapid and Robust Outdoor Fruit Detection System Combining Bio-Inspired Features with One-Stage Deep Learning Networks. *Sensors*, 20(1):275, 1 2020.
- [4] Guichao Lin, Yunchao Tang, Xiangjun Zou, Juntao Xiong, and Yamei Fang. Color-, depth-, and shape-based 3D fruit detection. *Precision Agriculture*, 21(1), 2 2020.
- [5] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks, 2015.
- [6] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You Only Look Once: Unified, Real-Time Object Detection. 6 2015.
- [7] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature Pyramid Networks for Object Detection. *Proceedings - 30th IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017*, 2017-January:936–944, 12 2016.
- [8] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, volume 9351, pages 234–241. Springer Verlag, 5 2015.
- [9] Alessandro Cini, Claudio Ioriatti, and Gianfranco Anfora. A review of the invasion of *Drosophila suzukii* in Europe and a draft

- research agenda for integrated pest management. *Bulletin of Insectology*, 65(1):149–160, 2012.
- [10] Kevin B. Rice, Brent D. Short, Sharon K. Jones, and Tracy C. Leskey. Behavioral responses of *Drosophila suzukii* (Diptera: Drosophilidae) to visual stimuli under laboratory, semifield, and field conditions. *Environmental Entomology*, 45(6):1480–1488, 2016.
- [11] Thomas Stone, Michael Mangan, Paul Ardin, and Barbara Webb. Sky segmentation with ultraviolet images can be used for navigation. Technical report.
- [12] Catherine M. Little, A. Rebecca Rizzato, Lise Charbonneau, Thomas Chapman, and N. Kirk Hillier. Color preference of the spotted wing *Drosophila*, *Drosophila suzukii*. *Scientific Reports*, 9(1):1–12, 12 2019.
- [13] R. C. MIALL. The flicker fusion frequencies of six laboratory insects, and the response of the compound eye to mains fluorescent ‘ripple’. *Physiological Entomology*, 3(2):99–106, 6 1978.

Towards co-creative drawing with a robot

Chipp Jansen¹

¹*Dept of Engineering
King's College London, UK
chipp.jansen@kcl.ac.uk*

Elizabeth I Sklar^{1,2}

²*Lincoln Institute for Agri-food Technology
University of Lincoln, UK
esklar@lincoln.ac.uk*

<https://doi.org/10.31256/LI3Do2J>

Abstract—This paper describes research into the development of a co-creative human-robot drawing system. Based on a pilot user study to survey the drawing practices of artists, various interaction factors have been identified that define example roles that a robot might take as a co-creative drawing partner. A research prototype system which observes an artist drawing with physical media—on paper—through the use of a drawing tablet and multiple cameras. The robotic system observes and captures data in real-time, as the artist draws. The longterm goal is to generate a data-backed model of the artist's drawing process, which in future will respond through projected visual interactions upon the drawn surface. The design and technical details of the observational system are described in this short paper.

Index Terms—human-robot collaboration, co-creative drawing, user study, artist drawing process

I. INTRODUCTION

Inspired by advances in creative AI and human-robot collaborative drawing [1], we envision a system in which physical media and non-invasive observation of an artist contributes to a co-creative mixed digital-physical workflow. Currently, drawing has a large digital tools economy, and the primary workflow for artists to create 2-D content is a digital workflow. Our longterm aim is to develop an intelligent, autonomous system that has the ability to combine active and passive sensing with sophisticated data analysis and active response, designed to help artists move forward in their creative process. Our research involves examination from two perspectives: (1) what is technically feasible through the development and evaluation of a research prototype of a real-time co-creative drawing system; and (2) what artists want with respect to a co-creative drawing partner. Here, we describe the technical design of our prototype system, the functionality of which was guided by input gathered during a pilot study conducted with drawing practitioners [2].

II. BACKGROUND

Traditionally, human-robotic collaboration in the visual arts consisted of artists programming robots to draw imperatively such as AARON [3], or portraiture style through observation from an artist robot, such as PAUL [4]. Collaborative human-robotic drawing is structured around how a robot collaborates with a human. In the DOUG system [1], the robot mimics what the human is drawing and the human responds to what

the robot is drawing in a simultaneous form of collaborative sketching [5]. Research into socially assistive robotics for art therapy has the robot responding to what a human is painting through contributing painting which is a visual metaphor according to a sensed emotional model [6]. In a similar approach, but outside of robotics in the creative computing area, research into co-creative sketching systems involves using visual metaphors to avoid design fixation by presenting imagery that would provide a conceptual shift in what the artist is drawing [7]. The *sketch-based* interaction research provides models of real-time drawing support, such as idealised geometric models [8], processed gradients of drawn images [9], graph-based representations of drawn stroke [10] and neural network representations [11].

To inform development of our human-robot creativity system, we conducted a mixed-methods study of drawing practitioners (e.g. professional illustrators, fine artists and art students) in Autumn 2018 [2]. Our aim was to discover possible roles that technology could play in observing, modelling and possibly assisting an artist with their drawing. A total of 21 participants representing a mix of professional illustrators, part-time drawing enthusiasts and illustration students were interviewed individually. Each participant completed a paper survey about their drawing habits, technology usage and attitude, recorded three drawing exercises and participated in an interview discussing their drawing habits and thoughts about AI/robotics. Three key themes were identified: drawing with physical mediums is a traditional and preliminary way of creation for visual artists; co-creative AI is preferable to didactic AI; and artists share a general discomfort with the idea of automating creative work. Discussion of these themes has been explored elsewhere [2], as has discussion of factors for defining a co-creative drawing robot and the role(s) that robot might take in a co-creative human-robot drawing process [12].

III. RESEARCH PROTOTYPE

Figure 1 shows an early version of our prototype system (*left*) and the corresponding schematic design of its components (*right*). Each component is controlled by a dedicated Raspberry PI with communication coordinated via ROS¹. The sensing components are (as labelled in Figure 1): three Raspberry PI cameras (**1a/C_{TOP}**, **1b/C_{LEFT}** and **1c/C_{RIGHT}**), a depth camera (**2/D_{FRONT}**) and a WACOM Bamboo Slate²

Research is supported through an EPSRC DTP Studentship "Collaborative Drawing Systems", Grant Reference EP/N509498/1

¹<http://ros.org>

²<https://www.wacom.com/en-cl/products/smartpads/bamboo-slate>




Fig. 1. Design of prototype system: (left) the hardware setup, (right) schematic of the design

digital “sketchpad” (3/T), which uses a pressure sensitive pen that tracks movement and produces marks on physical paper. The cameras observe the drawing area at multiple angles and records the textural aspects of the drawing, while the digital sketchpad records a vector representation of the pen’s movements. Our design includes a projector (P_{OVERHEAD}) which overlays the robot’s interaction upon the drawing surface; this component will be implemented in our next generation prototype. Through the use of projection, the artist has the sole physical agency to manipulate the drawing in progress.

Recently, we completed a data-gathering drawing study ($n = 13$) involving full-time drawing practitioners (professionals and students) to test our prototype system and collect an initial set for modelling. Participants did two drawing exercises, each lasting about 10 minutes: draw from *observation* of a still-life, and draw freely from *imagination*.




Fig. 2. Example inventory of datasets gathered for a drawing exercise.




Fig. 3. Drawing in progress.




Fig. 4. Example drawing with time of drawing visualised through color.

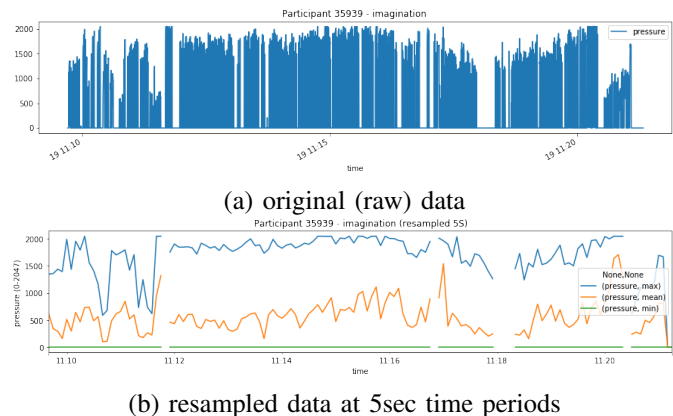


Fig. 5. Pressure of drawing pen as recorded by the WACOM digital “sketchpad” for the same drawing exercise.

Here we present some very preliminary analysis of the data gathered. Figures 2-5 illustrate the data gathered and preliminary modelling for one sample participant. Figure 2 illustrates the amount of data gathered from one drawing exercise completed by our sample participant. Figure 3 shows their drawing in progress. Figure 4 shows the resulting drawing and how it evolved over time. Figure 5a illustrates the level of pressure exerted, over time, by one participant for the same drawing. This data is rather dense, so we have resampled the data in an attempt to build a model of user pressure behaviour (Figure 5b).

IV. SUMMARY

We have described the design of our robotic co-creative drawing system, informed by an initial user study (Autumn 2018) and recently tested (January 2020) to gather drawing data from full-time artists. With this dataset, we are building models of the drawing process that will incorporate dynamics of the artist’s body, their tool use and textural changes to the artwork’s surface. With these models, we plan to experiment by having our robotic system inhabit various roles as a co-creative partner (e.g. suggestive, predictive, referential) and evaluate these roles through a larger user study with drawing practitioners in order to measure impacts of the robot on the artist’s drawing process.

REFERENCES

- [1] S. Chung, "Drawing Operations (DOUG)," <https://sougwen.com/project/drawing-operations>, 2015.
- [2] C. Jansen and E. Sklar, "Co-creative Physical Drawing Systems," in *ICRA-X Robots Art Program at IEEE International Conference on Robotics and Automation (ICRA)*, Montreal QC, Canada, May 2019, p. 2.
- [3] P. McCorduck, *Aaron's Code: Meta-Art, Artificial Intelligence, and the Work of Harold Cohen*. W.H. Freeman, 1991.
- [4] P. Tresset and F. Fol Leymarie, "Portrait drawing by Paul the robot," *Computers & Graphics*, vol. 37, no. 5, pp. 348–363, Aug. 2013.
- [5] E. Sandry, "Creative Collaborations with Machines," *Philosophy & Technology*, vol. 30, no. 3, pp. 305–319, Sep. 2017.
- [6] M. Cooney and P. Berck, "Designing a Robot Which Paints With a Human: Visual Metaphors to Convey Contingency and Artistry," in *ICRA-X Robots Art Program at IEEE International Conference on Robotics and Automation (ICRA)*, Montreal QC, Canada, May 2019, p. 2.
- [7] P. Karimi, M. L. Maher, N. Davis, and K. Grace, "Deep Learning in a Computational Model for Conceptual Shifts in a Co-Creative Design System," *arXiv:1906.10188 [cs, stat]*, Jun. 2019.
- [8] J. Arvo and K. Novins, "Fluid sketches: Continuous recognition and morphing of simple hand-drawn shapes," in *Proceedings of the 13th Annual ACM Symposium on User Interface Software and Technology - UIST '00*. San Diego, California, United States: ACM Press, 2000, pp. 73–80.
- [9] Y. J. Lee, C. L. Zitnick, and M. F. Cohen, "ShadowDraw: Real-time user guidance for freehand drawing," in *ACM SIGGRAPH 2011 Papers on - SIGGRAPH '11*. Vancouver, British Columbia, Canada: ACM Press, 2011, p. 1.
- [10] J. Xing, L.-Y. Wei, T. Shiratori, and K. Yatani, "Autocomplete Hand-drawn Animations," *ACM Trans. Graph.*, vol. 34, no. 6, pp. 169:1–169:11, Oct. 2015.
- [11] D. Ha and D. Eck, "A Neural Representation of Sketch Drawings," *arXiv:1704.03477 [cs, stat]*, May 2017.
- [12] C. Jansen and E. I. Sklar, "Towards a HRI System for Co-creative Drawing," in *Workshop on Creative Content in HRI*, Cambridge, UK, Mar. 2020, p. 4.

Robotic Untangling of Herbs with Parallel Grippers

1st Prabhakar Ray
Department of Engineering
King's College London
 London, United Kingdom
 ray.prabhakar@kcl.ac.uk

2nd Matthew J. Howard
Department of Engineering
King's College London
 London, United Kingdom
 matthew.j.howard@kcl.ac.uk

<https://doi.org/10.31256/Wd8Aj7K>

Abstract—Robotic packaging generally involves picking target objects from a pile consisting of many other similar or random objects. For a pile composed of herbs, the weight picked up can be controlled by varying the opening aperture width of a parallel gripper. However, the individual strands of herbs get entangled with each other, causing more to be picked up than desired. Here, it is shown that using a *spread-and-pick* approach the degree of entanglement in a herb pile can be reduced. Compared to the traditional approach of picking from an entanglement-free point in the pile, the proposed approach reduces the variance in picked weight for homogeneous piles of plastic and real herbs by 36.35 and 23.64 percent, respectively. These results demonstrate that using the proposed *spread-and-pick* approach, the stochasticity of a herb environment can be counterbalanced effectively.

Index Terms—Robotics in Agriculture and Forestry, Agricultural Automation, Computer Vision for Automation

I. INTRODUCTION

Industries manufacturing machinery, transportation equipment and various everyday retail products on a large scale have benefited immensely from intelligent and collaborative assembly-line robots. However, to date, the application of such technologies to the processing of fresh horticultural produce remains mostly dependent on manual labour. The suppliers of fresh herbs, for instance, grow stock under glass or in fields and then must transport them to packaging stations and pack them as per the weight requirements of retailers. The manual packaging process involved is not only costly in terms of labour but also suffers from human errors and low production efficiency.

A more scalable approach could be automation through adaptive robotic systems, however, deploying such a system presents several challenges. Fresh horticultural produce can be highly variable in terms of its handling properties, even within a single plant variety, making it difficult to design robotic controllers for their manipulation. Herbs, in particular, tend to present as a highly stochastic, tangled mass (see Fig. 1(a)), making it difficult for a robotic system to extract a uniform quantity suitable for presentation to the consumer. Smaller, fixed-mass portions must be extracted and fed via conveyor belt for packaging but the tangling (see Fig. 1(b)) makes the mass lifted in a simple pick operation difficult to predict.

Traditionally, researchers have studied bin-picking¹ in the context of two main challenges: (i) *gripper-object collision*



Fig. 1. (a) Handling fresh herbs and (b) entanglement.

and (ii) *object entanglement*. The use of simple geometric primitives such as planes, spheres, cylinders and cones for object recognition in the bin was proposed in [1]. Changes in surface types and depth discontinuities were then used to segment the cluttered scene. A novel vision-based algorithm was proposed in [2], which suggested resolving gripper-object collision by identifying and picking the topmost object in a pile. A deep learning approach for picking individual objects from a cluttered bin was proposed in [3]. These methods prove effective for avoiding gripper-object collision. However, they do not address the issue of object entanglement directly.

In the past, the entanglement between the objects in a pile has been directly addressed through interaction with the pile [4], [5]. Known CAD models of the objects were used in [6] for planning singulation of individual objects from a heterogeneous pile. Although these methods address the issue of entanglement directly, their objective is to *extract a single, rigid object*, rather than a *uniform quantity of deformable, granular media*, such as herbs.

To this end, this paper proposes a *spread-and-pick* method, which reduces entanglement, and in turn, makes the pick operation more predictable in terms of the picking quantity. It has several advantages, including that (i) it does not require any large scale data collection and (ii) it does not depend on any geometrical information about the objects picked.

II. METHOD

A. Collision-free Gripper Pose: Graspability Index

The graspability index (GI) [7] is a vision-based measure for evaluating the candidate grasping poses, which has proved useful in industrial pick and place settings. It uses a single

This work is supported in part by Vitacress Salads UK Ltd.

¹Pile and bin are used interchangeably in this paper.




Fig. 2. Overview of the experimental setup. Red, Green and Blue arrows represent x, y and z axes respectively.

depth map of the scene to estimate the optimal gripper position and orientation for picking an individual object. Along with the parallel gripper, GI is also applicable to general grippers such as multi-finger and vacuum grippers. For an insertion depth (r_z), GI estimates a point $\mathbf{r} = (r_x, r_y, r_\theta)^\top$ in the bin such that the parallel plates of the gripper can be inserted without causing any collision between them and the objects inside. r_θ denotes the orientation of the gripper around its z-axis. A range of r_θ is evaluated using GI, and for the optimal r_θ^* , the optimal collision-free picking point (r_x^*, r_y^*) is estimated.

B. Tangle Reduction

A key problem with the use of GI is that it does not take account of tangling causing undesired objects to be picked alongside the target one. For better control and accuracy, this paper proposes a means to adjust the picking action to reduce tangling and thereby achieve more consistent picking. For this, a *spread-and-pick* approach is proposed, whereby the pick operation is augmented with an extra detangling step using an extension to the GI. Humans frequently use their fingers to separate things while picking, especially when they have to work with one hand. The proposed *spread-and-pick* approach draws inspiration from this behaviour. Specifically, if the target object is between the two fingers of a parallel gripper, instead of moving inwards (closing) and picking the object, the fingers are moved outward to displace other objects entangled or close to the boundary of the target object. The proposed approach extends the GI by identifying regions of high entanglement in the scene and then defining a spreading movement to disentangle them: First, for a fixed insertion depth r_z and gripper orientation r_θ , the optimal collision-free picking point (r_x^*, r_y^*) is estimated using GI. Next, the peak entanglement point $\mathbf{r}' = (r'_x, r'_y)^\top$ is estimated using the collision map as obtained from the vision module. Finally, r_θ is updated as

$$\theta = \arccos\left(\frac{\mathbf{r}^{*\top} \mathbf{r}'}{\|\mathbf{r}^*\| \|\mathbf{r}'\|}\right). \quad (1)$$

TABLE I

Picked weight (mean \pm s.d.) over 20 trials for plastic herbs.

Gripper Width(mm)	Method	Picked Weight(g)
20	Graspability Index	9.844 \pm 11.078
	Spread and Pick	8.247 \pm 7.635
30	Graspability Index	10.482 \pm 9.172
	Spread and Pick	7.501 \pm 5.839
40	Graspability Index	13.267 \pm 11.953
	Spread and Pick	14.333 \pm 8.480

TABLE II

Picked weight (mean \pm s.d.) over 10 trials for real herbs.

Gripper Width(mm)	Method	Picked Weight(g)
20	Graspability Index	3.712 \pm 3.028
	Spread and Pick	15.646 \pm 2.471
30	Graspability Index	8.622 \pm 3.480
	Spread and Pick	15.788 \pm 2.658
40	Graspability Index	19.192 \pm 7.273
	Spread and Pick	18.361 \pm 6.934

C. Procedure

Using the set up as shown in Fig. 2, a series of robotic picking operations are conducted for plastic and real herbs following a simple picking methodology as well as following the proposed *spread-and-pick* methodology. Each picking operation consists of the robot reaching into a pile of tangled media (herbs) of fixed mass, closing its gripper, and lifting what is grasped free of the surface of the bagging area. For simplicity and lower cycle-time, only 3-degrees of freedom of the robot are used for picking movements, and the highest point in the pile is chosen as the target picking location. To ensure a similar physical arrangement of the tangled media between trials, any media picked is returned to the bagging area and the entire quantity is transferred to a container of fixed dimension (18 cm x 13.5 cm) before being placed back on to the bagging area for the next pick. In the case of real herbs, for each method and each opening aperture width w of the gripper, a fresh batch of herbs is used to avoid variations due to changes in their physical state (*e.g.*, due to herbs drying out, or becoming damaged over successive picks). For *spread-and-pick*, the grasping manoeuvre consists of setting the gripper orientation r_θ according to (1) and the gripper aperture w to the chosen width prior to lowering it into the pile. Once lowered into the pile, instead of closing, the gripper plates are moved outwards and opened, leading to dispersion of the herbs. The mass of media picked is recorded for each trial.

III. RESULTS

The results for picking real herbs are shown in TABLE II. As can be seen, there is a clear reduction in the variance of the picked mass for the *spread-and-pick*. The maximum percentage decrease in the variance for plastic and real herbs is 36.35% and 23.64%, respectively, for the intermediate $w = 30$ mm. This suggests that the proposed approach is effective for improving the predictability of picking in this challenging automation task.

REFERENCES

- [1] G. R. Taylor and L. Kleeman, "Robust range data segmentation using geometric primitives for robotic applications," in *Signal and Image Processing (SIP)*, 2003.
- [2] Y. Xu, Y. Mao, X. Tong, H. Tan, W. B. Griffin, B. Kannan, and L. A. DeRose, "Robotic Handling of Surgical Instruments in a Cluttered Tray," *IEEE Transactions on Automation Science and Engineering*, vol. 12, no. 2, pp. 775–780, 2015.
- [3] M. Schwarz and S. Behnke, "PointNet Deep Learning for RGB-D Object Perception in Cluttered Bin Picking," *IEEE International Conference on Robotics and Automation (ICRA)*, no. May, pp. 2–4, 2017. [Online]. Available: http://www.ais.uni-bonn.de/papers/ICRA_2017_WPAW_Schwarz.pdf
- [4] *Resolving Occlusions Through Simple Extraction Motions in Robotic Bin-Picking*, ser. International Manufacturing Science and Engineering Conference, vol. Volume 2: Materials; Biomanufacturing; Properties, Applications and Systems; Sustainable Manufacturing, 06 2016, v002T04A002. [Online]. Available: <https://doi.org/10.1115/MSEC2016-8661>
- [5] C. Papazov, S. Haddadin, S. Parusel, K. Krieger, and D. Burschka, "Rigid 3D geometry matching for grasping of known objects in cluttered scenes," *International Journal of Robotics Research*, vol. 31, no. 4, pp. 538–553, 2012.
- [6] N. Kaipa, S. Shriyam, A. B. Kumbala, and R. K. Gupta, "Automated plan generation for robotic singulation from mixed bins," in *International Conference on Intelligent Robots and Systems (IROS) Workshop on Task Planning for Intelligent Robots in Service and Manufacturing*, 2015.
- [7] Y. Domae, H. Okuda, Y. Taguchi, K. Sumi, and T. Hirai, "Fast graspability evaluation on single depth maps for bin picking with general grippers," *Proceedings - IEEE International Conference on Robotics and Automation (ICRA)*, pp. 1997–2004, 2014.

Towards Insect Inspired Visual Sensors for Robots

Blayze Millward
Sheffield Robotics
The University Of Sheffield
Sheffield, UK
b.f.millward@sheffield.ac.uk

Steve Maddock
Department of Computer Science
The University Of Sheffield
Sheffield, UK
s.maddock@sheffield.ac.uk

Michael Mangan
Sheffield Robotics
The University Of Sheffield
Sheffield, UK
m.mangan@sheffield.ac.uk

<https://doi.org/10.31256/Do2Ik3H>


Abstract—Flying insects display a repertoire of complex behaviours that are facilitated by their non-standard visual system that if understood would offer solutions for weight- and power-constrained robotic platforms such as micro unmanned aerial vehicles (MUAVs). Crucial to this goal is revealing the specific features of insect eyes that engineered solutions would benefit from possessing, however progress in exploration of the design space has been limited by challenges in accurately replicating insect vision. Here we propose that emerging ray-tracing technologies are ideally placed to realise the high-fidelity replication of the insect visual perspective in a rapid, modular and adaptive framework allowing development of technical specifications for a new class of bio-inspired sensor. A proof-of-principle insect eye renderer is shown and insights into research directions it affords discussed.

Index Terms—Novel sensing, Artificial Intelligence and Robotics, Bioinspired, Vision, Rendering

I. INTRODUCTION

For engineers seeking to develop long-range autonomous micro unmanned aerial vehicles (MUAVs) weight and power constraints present a critical design parameter that currently limits their application [1]. Insects provide an existence proof that seemingly simple sensory systems are sufficient for solving complex tasks from navigating natural habitats in 3D; through detection and tracking of food, prey and conspecifics; to rapid flight control to avoid damaging impacts from static and moving objects [2]. We thus propose that novel solutions will arise through revealing the secrets of insect visual systems allowing for abstraction into a new class of low-power, low-weight bioinspired robot sensor.

Insects see the world through a fundamentally different mechanism than humans and most camera systems [2]. Their compound eyes are constructed from hundreds to thousands of self-contained "mini-cameras" known as ommatidia: each comprising a lens, light-guide and light sensitive cells which are physically interlocked over a convex surface per eye (see Figure 1(b)). In addition, the surface structure (e.g. field of view), layout (e.g. density of ommatidia), and ommatidial function (e.g. sensitivity to specific properties of light: wavelength, polarisation) of compound eyes vary across eye regions and between caste, sex and species. Given the vastness of the feature space in which compound eye designs reside and the computational complexity involved in searching that space for possible solutions we define three criteria that any insect eye simulator must meet:



(a) The projection pyramid of a panoramic image, projecting the 3D environment onto a single uniform point or sphere. (b) The sampling rays following the normal of an irregular projection surface, similar to that found on an insect eye.

Fig. 1: Regular and irregular surface projection diagrams. Insect eye photo credit to Matthew Barber, used with Permission.

- 1) Perform beyond real-time.
- 2) Allow arrangement of ommatidia on arbitrary 3D surfaces.
- 3) Allow configuration of individual ommatidial properties.

II. STATE OF THE ART

Insect eye perspective renderers have tended to focus on recreating the panoramic field of view and low spatial resolution properties of compound eyes [3], [4]. Cube-mapping techniques, whereby a camera is rotated to sample images across six viewpoints which are then stitched into a single panoramic image (Figure 1b(a)) combined with downsampling and post-processing approximate these two properties in a computationally efficient way (e.g. [5]). These systems have been used to investigate the impact of these properties on navigational performance in simulated [6] and real-world studies that generate similar perspectives by augmenting standard cameras using convex mirrors [7], [8] or fish-eye lens imaging systems [9]. Yet as such models form images through the projection of a 3D scene onto a uniform viewing surface, they inherently violate *criteria 3*. Attempts to address this issue by modelling compound eyes using multiple small field of view cameras within hardware-accelerated modern game engines have proven unsuccessful due to performance constraints that violate *criteria 1*.

III. INSECT VISION USING RAY-BASED METHODS

Raycasting and *raytracing* techniques are image rendering approaches that produce realistic imagery through the

physically-based simulation of light rays, as opposed to the less realistic projection transform approaches commonly used in real-time graphics processes. By tracing the path of individual rays, ray-based rendering allows for the accurate simulation of optical effects, driving their realism. Technical development has been primarily driven by the film industry (e.g. see advances in the *Toy Story* movie franchise, with the most recent instalment including advanced camera lensing effects).

Due to demand in the video games industry for increasingly realistic graphics generated in real-time, dedicated hardware has been introduced to massively parallelise computationally expensive ray casting and tracing algorithms. These changes in graphics processing ability lend themselves very well to the simulation of compound eyes, as ray-based methods allow for the accurate sampling of light as if refracted through an ommatidium's optical system: something that traditional rendering pipelines struggle to achieve [10].

Moreover, as ray-based methods are inherently designed to handle the simulation of many rays at many locations within an environment, the source position of these rays are immaterial: any projection surface can be used to spawn rays with minimal additional overhead. That is, it should be feasible to render the perspective from any number of ommatidia on any surface.

Finally, modern raytracing hardware is capable of rendering tens of millions of rays per frame, owing to their real-time use on high-definition displays. In comparison, a drone bee's eye consists of only about 10,000 ommatidia [11], the view of each of which could be simulated with 81 rays [9] and the total number of rays would still be less than that used in a 1920-by-1080 pixel (standard HD screen size) rendering.

An initial study conducted by Polster et al. [12] demonstrates some of the benefits to be gained from ray-based simulation of the compound eye, providing support to the approach. Their work does not, however, run in real-time and also lacks tools to explore differing insect eye surface shapes.

IV. PROOF OF CONCEPT

Figure 2 shows sample images using our prototype raytracing-based insect eye renderer. Images were generated in a large 3D environment from an insect's visual perspective at 60 frames per second utilising raytracing hardware in modern consumer-grade NVidia graphics cards (NVidia GeForce RTX 2080Ti) successfully fulfilling *Criteria 1*. Figure 2 (upper) shows a simulated eye with equally spaced ommatidia arranged on a sphere whereas Figure 2 (lower) has ommatidia clustered around the horizon as observed in some insects, demonstrating the ability of the system to fulfil *Criteria 3*.

With the technical challenges of building a ray-tracing based insect eye renderer largely complete (open-source software release expected soon), we will now look to investigate the specific requirements for natural and artificial visual systems in shared tasks such as navigation. Insights gained from more thoroughly exploring the insect visual perspective—and its design relative to visual feature extraction—will help guide the development of visual systems in robotics by considering

not only visual post-processing steps, but also the intrinsic structure and design of the imaging sensor itself.




Fig. 2: Insect perspectives generated using a prototype raytracing-based renderer. Upper: Ommatidial distribution increased on horizon. Lower: Equidistant distribution.

REFERENCES

- [1] V. Kumar and N. Michael, "Opportunities and challenges with autonomous micro aerial vehicles," *The International Journal of Robotics Research*, vol. 31, pp. 1279–1291, 9 2012.
- [2] M. F. Land and D.-E. Nilsson, *Animal eyes*. Oxford University Press, 2002.
- [3] T. R. Neumann, "Modeling Insect Compound Eyes: Space-Variant Spherical Vision," in *Biologically Motivated Computer Vision*, pp. 360–367, Springer, Berlin, Heidelberg, 11 2002.
- [4] A. Wystrach, M. Mangan, A. Philippides, and P. Graham, "Snapshots in ants? New interpretations of paradigmatic experiments," *Journal of Experimental Biology*, vol. 216, pp. 1766–1770, 5 2013.
- [5] B. Baddeley, P. Graham, P. Husbands, and A. Philippides, "A Model of Ant Route Navigation Driven by Scene Familiarity," *PLoS Computational Biology*, vol. 8, p. e1002336, 1 2012.
- [6] A. Wystrach, A. Dewar, A. Philippides, and P. Graham, "How do field of view and resolution affect the information content of panoramic scenes for visual navigation? A computational investigation," *J Comp Physiol A*, vol. 202, pp. 87–95, 2016.
- [7] J. Wessnitzer, M. Mangan, and B. Webb, "Place memory in crickets," *Proceedings of the Royal Society B: Biological Sciences*, vol. 275, pp. 915–921, 4 2008.
- [8] T. Stone, M. Mangan, P. Airdin, and B. Webb, "Sky segmentation with ultraviolet images can be used for navigation," tech. rep., 2014.
- [9] W. Stürzl, N. Boeddeker, L. Dittmar, and M. Egelhaaf, "Mimicking honeybee eyes with a 280° field of view catadioptric imaging system," *Bioinspiration & Biomimetics*, vol. 5, p. 036002, 9 2010.
- [10] T. Whitted, Turner, Whitted, and Turner, "An improved illumination model for shaded display," in *Proceedings of the 6th annual conference on Computer graphics and interactive techniques - SIGGRAPH '79*, vol. 13, (New York, New York, USA), p. 14, ACM Press, 1979.
- [11] J. P. v. Praagh, W. Ribi, C. Wehrhahn, and D. Wittmann, "Drone bees fixate the queen with the dorsal frontal part of their compound eyes," *Journal of Comparative Physiology, A*, vol. 136, pp. 263–266, 1980.
- [12] J. Polster, J. Petrasch, R. Menzel, and T. Landgraf, "Reconstructing the visual perception of honey bees in complex 3-D worlds," tech. rep., 2018.

Vision Navigation System to Manoeuvre Unmanned Aerial Vehicle (UAV)

Patrick Tarek Saleh
Centre for Automation and Robotics
Research
Sheffield Hallam University
Sheffield, United Kingdom
Saleh.Patrick@student.shu.ac.uk

Jacques Penders
Centre for Automation and Robotics
Research
Sheffield Hallam University
Sheffield, United Kingdom
j.penders@shu.ac.uk

Lyuba Alboul
Centre for Automation and Robotics
Research
Sheffield Hallam University
Sheffield, United Kingdom
l.alboul@shu.ac.uk

<https://doi.org/10.31256/Jb6Qv4S>

Abstract— This paper describes the development and the implementation of an omnidirectional multiple stereo cameras vision system. The vision system is compact and light enough to operate on board a commercially available off the shelf miniature UAV (Unmanned Aerial Vehicle)-Quadrotor. The vision system contains several stereo cameras ruggedly fixed on-board the UAV-Quadrotor, it is oriented in such a way that it has a 360 degrees omnidirectional visual coverage. The paper demonstrates that by combining several stereo cameras, it can provide and combine depth information and optical flow data in real time to an on-board image processing computer. One can estimate the position and the orientation such as roll, pitch and yaw of the UAV in 3D (three dimensions) space accurately without the aid of GPS (Global Positioning System), IMU (Initial Measurement Unit) or any other external navigation or orientation aid. This method can be described as “Simultaneous Localization and Exploration Oriented Visual Navigation”. It is a development of a vision system capable of navigating autonomously an UAV-Quadrotor through random free spaces within an unknown complex environment and without any mapping, it simply performs by detecting and avoiding obstacles calculating the required “through flight path”.

Keywords— UAV-Quadrotor, vision system, GPU, stereo camera, visual navigation, Homography planer surface.

I. INTRODUCTION

Airborne mobile robots localisation is the determination of the robot position and orientation in 3D space. Much research has been conducted in order to estimate accurately the robot position and orientation utilising different types of perception methods [1]. GPS and Laser detection and ranging sensor (LIDAR) [2] are among the most popular available positioning and mapping sensors.

Due to the latest advances in cameras sensors technology, cameras have become lighter and more power efficient, they operate at high image resolution and ultra-fast processing power. Hence, this results in increased research in the field of computer vision [1], specifically in visual perception. Many researchers utilise mono and stereo cameras systems for mobile robot localisation and navigation with acceptable result [3] [4]. The proposed vision system has 360 degrees of stereo vision coverage compared with [4][5][6]. It enables an aerial robot to navigate autonomously in a complex non pre-mapped 3D environment. The onboard vision system utilises a robust and light weight low power perception system containing five stereo camera modules, further simply referred as stereo cameras. Each stereo camera consists of one (RGB) Red, Green, Blue camera, one infrared (IR) camera and one IR laser illuminator depicted in Fig.1. Considering the optical flow measurements this approach is

multispectral diverse from the other approaches utilising just one RGB stereo camera dedicated only for depth measurements [4] [5] [6]. To process the visual information obtained from the five stereo cameras, an on-board Nvidia Jetson TX2-256 core GPU (Graphic Processor Unit) Pascal processor board is used. The power requirement of the entire vision system when measured on average does not exceed 50 watts which makes it very suitable for autonomous airborne applications.



Fig. 1. Stereo camera module just 28 grams in weight.

The airborne robot platform used in this research is a miniature custom build Unmanned Aerial Vehicle (UAV) - Quadrotor shown in Fig. 2. The Quadrotor is inherently less unstable than a conventional helicopter. It is a multi-rotor under-actuated helicopter. It has four propellers, each diagonal facing propellers pair rotating in opposite direction cancelling out each other generated rotational torque, all four propellers powered by electric DC brushless motors (BLDC), which are powered from on-board DC voltage.

The UAV-Quadrotor is capable of performing a vertical flight and transition to and from a horizontal flight. Its flight control method depends mainly on controlling the variation of the distribution of the propellers lifting force and propellers generated torques about its center of gravity. Controlling these forces enables the onboard control system to stabilize it during hovering and performing complex flight. The UAV-Quadrotor is capable of lifting a payload weight of up to 1100 gram.



Fig. 2. The research custom build UAV Quadrotor

II. EXPERIMENTAL METHOD

The research method investigates the development of a multiple stereo vision system utilising off-the-shelf hardware and software. The developed system contains five stereo cameras, image processing and navigation computer, off-the-shelf flight controller, wireless communication devices, DC power source and the UAV-Quadrotor frame. The vision system has 360 degrees of stereo vision coverage shown in Fig. 3. It processes in sequence the images obtained from the upward looking stereo camera, forward looking stereo camera, downward stereo looking camera, left-backward looking stereo camera, and right-backward stereo looking camera. The entire vision system operates at up to speed 20 fps at image pixel resolution of (1280 x 720) and given pixel areas of (921,600 pixels) per direction calculating position and orientation of the UAV Quadrotor with the reference to any detected environment features or landmarks. The vision system will compare its position and orientation measurements accuracy during flying, in real time with the on-board high grade IMU and GPS measurements specifically within an indoor low light environment. The obtained experimental data will be sent in real time to a ground station for processing and analysis. Furthermore, the vision system will provide in real time the horizontal and vertical flying control commands to the on-board flight control system, also occasionally for safety reason in conjunction with the UAV-Quadrotor's remote control.

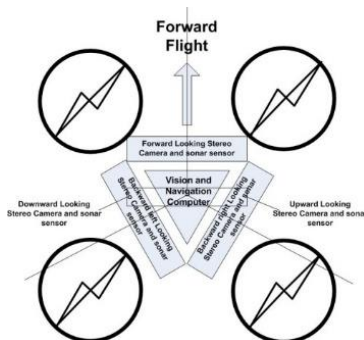


Fig. 3. The UAV-Quadrotor top view showing the stereo cameras orientation.

III. THE VISION AND NAVIGATION ALGORITHM

Considering any 3D environment whether it is known or unknown, empty or equipped, in its simplest description it contains a combination of flat surfaces, edges, corners, and bends or curves. These combinations of environment features present the environment in terms of walls, floor, ground, ceiling, edges, corners, and other possible shapes.

Hence, the vision algorithm will process the images depth and optical flow obtaining five data sets from five stereo cameras. Each data set contains information about the detected environment features such as corners, edges and planer (plane surfaces) estimating their position (X, Y, Z) and the orientation such as roll, pitch and yaw (ϕ , θ , ψ). Visual detection performed mainly by matching (features detector and descriptor) and tracking these features between the sequences of image frames obtained from the five stereo cameras video stream at 20Hz. Homographic Filter is first

level filtering to reduce illumination glare and shadows followed by Bilateral Filter to remove noise and sharpen edges in images. Second level filtering commenced with RANSAC (Random Sampling Consensus), it detects data outliers and excluded them from the measurements, this enables an accurate estimation of the homography matrix and the position of each stereo camera in 3D space. The estimation is based on the assumption that the UAV-Quadrotor is reasonably stable, and it flies or hovers over or under a plan surface such as a floor or ceiling respectively. Finally, the third level filtering performs non-linear estimations utilising Extended Kalman Filter (EKF) to estimate each stereo camera's coordinates (position and orientation) in 3D space. The estimation is calculated with reference to obstacles such as landmarks or an environment features detected during the flight or hovering. The 6-DOF (Degree of Freedom) and altitude of the UAV-Quadrotor are calculated by transforming the translation and rotation corrected components of each of the five stereo cameras local frame coordinates to the UAV-Quadrotor frame coordinates expressed in the North-East-Down (NED) system. Autonomous navigation will be performed by the ability of the vision system to detect, avoid and calculate the flight path required for the UAV-Quadrotor to fly through hindrance free space.

IV. RESULT AND DISCUSSION

As previously discussed, the utilisation of visual perception in localizing the UAV-Quadrotor in an unknown environment appears very promising. As shown the Omnidirection stereo vision system has multidirectional field of view. It is capable to acquire distance, velocity, and rotation components of the environment visually detected obstacles in real time. Such a system is suitable for exploration navigation. It is defined as a method of detecting, avoiding, and navigating available spaces within the environment. Currently the hardware and software integration is completed within the UAV-Quadrotor including flying testing with full payload. Considering urban and industrial areas, the image-processing algorithm will be tested for robustness and the ability of estimating the translation and rotation components and navigating the UAV-Quadrotor during flying and hovering without the aid of the IMU, GPS and any external navigation aid.

V. FUTURE WORK

The multidirectional stereo vision system method for estimating position and orientation will enable the UAV-Quadrotor to fly and avoid obstacles following the required flight path through any random available space within unknown environments without any requirement for mapping. It allows the exploration of the environment at a very fast pace. It is also possible to utilise such a vision system to perform complex flight formation of multiple UAV-Quadrotors. This is achievable by measuring visually the attitude and speed of each UAV in the visual vicinity and coordinate collectively the required flight manoeuvre to fly along a specific flight path to explore a large and complex environment. The proposed system may provide new opportunities of visual oriented navigation and exploration

such exploring un-mapped area, performing ad-hoc fast exploration mission of a very large urban site or navigating planetary landscape where there is uneven landscape with many valleys and hills.

REFERENCES

- [1] Francisco Bonin-Font, Alberto Ortiz and Gabriel Oliver. Visual Navigation for Mobile Robots: a Survey. DPI 2005-09001-C03-02 and FEDER funding, University of the Balearic Islands. November 2008.
- [2] Ryan W. Wolcott and Ryan M. Eustice. Visual Localization within LIDAR Maps for Automated Urban Driving. [Intelligent Robots and Systems \(IROS 2014\), 2014.](#)
- [3] Jos e Gaspar, Member, IEEE, Niall Winters, and Jos e Santos-Victor, Member, IEEE. Vision-based Navigation and Environmental Representations with an Omni-directional Camera. *IEEE Transactions on robotics and automation*, Vol.16, No:6, December 2000.
- [4] Huang, Albert S., Abraham Bachrach, Peter Henry, Michael Krainin, Daniel Maturana, Dieter Fox, and Nicholas Roy. "Visual odometry and mapping for autonomous flight using an RGB-D camera." In *Robotics Research*, pp. 235-252. Springer, Cham, 2017.
- [5] Dryanovski, Ivan, Roberto G. Valenti, and Jizhong Xiao. "Fast visual odometry and mapping from RGB-D data." In 2013 IEEE international conference on robotics and automation, pp. 2305-2310. IEEE, 2013.
- [6]) Handa, Ankur, Thomas Whelan, John McDonald, and Andrew J. Davison. "A benchmark for RGB-D visual odometry, 3D reconstruction and SLAM." In 2014 IEEE international conference on Robotics and automation (ICRA), pp. 1524-1531. IEEE, 2014.

Experimental Analysis of Soft Vacuum Cups for Automated Mushroom Picking

Hasan Husain
School of Engineering
University of Lincoln
Lincoln, United Kingdom
16657424@students.lincoln.ac.uk

Khaled Elgeneidy
School of Engineering
University of Lincoln
Lincoln, United Kingdom
kelgeneidy@lincoln.ac.uk

<https://doi.org/10.31256/BI8Lb9I>

Abstract - In this paper, the use of soft suction cups for automated mushroom picking is studied. The aim is to identify the vacuum level that starts to cause bruising for mushroom cups, in addition to the maximum torque that can be generated at this value to facilitate harvesting via twisting mushrooms. An experimental setup was developed that controls vacuum level, controls rotation of the vacuum gripper, and records resulting torque during picking.

Keywords - *agri-robotics, soft grippers, grasping.*

I. INTRODUCTION

The agriculture sector nowadays is interested in the use of robotics to automate labour-intensive harvesting tasks. Mushroom harvesting is an example which requires lots of human pickers that are becoming difficult to recruit. Picking involves skilful twisting of the mushrooms to separate them from the compost. Soft robotics offers various grasping technologies that are suited for delicate targets such as elastomer actuators, granular jamming, Gecko adhesion, Electro adhesion, and others. Suction based grippers can be a simple solution for mushroom picking, but vacuum can also bruise picked mushroom surface. This project studies the impact of vacuum on mushroom bruising and the maximum torque that can be generated for picking mushroom without bruising using a soft vacuum cup.

II. LITERATURE REVIEW

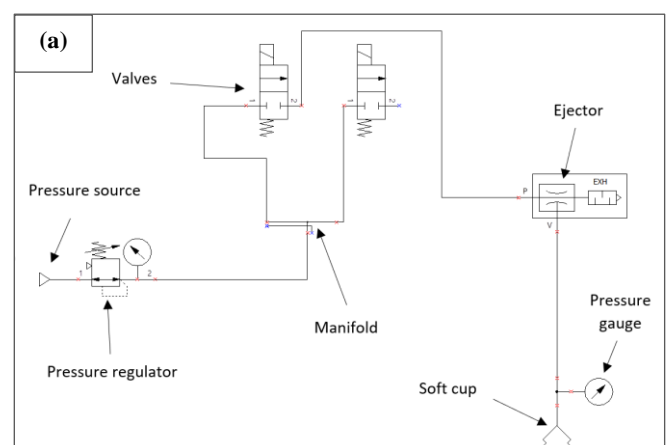
Soft grippers provide excellent shape adaptation to a wide range of objects compared to conventional rigid grippers. Soft grippers can be categories into three main categories. First by using actuation, which can bend and grasp the objects gently similar to human fingers [1][2]. This approach can easily handle convex and non-convex shapes, but it is difficult for picking a flat and deformable object. Second gripper technique is by controlled stiffness [3]. There are four methods for controlling stiffness; The shape memory polymers [4], low-melting-point alloys [5], granular jamming [6], and electrorheological (ER) and magnetorheological (MR) fluids [8]. The controlled stiffness gripper is not ideal for lifting flat or deformable objects, but is usually combined with actuation to handle convex and non-convex objects. The third soft gripper category is controlled adhesion, which also needs to use the actuation method to grip the objects. There are several examples of this method such as: electro-adhesion [9], Gecko adhesion, or simply using suction cups. Adhesion heavily depends on the surface properties of an object, but can be suitable for handling convex, flat and deformable objects, but not ideal for non-convex objects [3][10]. A soft gripper can also combine two technologies to improve performance. The choice of soft gripper type relies on the properties of the object such as shape, weight, and delicacy,

as well as, picking requirements such as speed, force, the power consumption, and biocompatibility.

In this work, soft suction cups are investigated as a simple approach for picking mushroom since they enable picking mushroom cups from the top, which is important since mushrooms grow in dense clusters so there is no space for fingers to reach in between. This makes other actuation based technologies difficult to reach the right position to grasp and twist a mushroom. This paper investigates if the suction cup concept would enable adapting to varying mushroom cup sizes without damaging or bruising. In addition, quantifying the maximum torque that the soft vacuum gripper can generate to harvest by twisting.

III. EXPERIMENTAL SETUP

The pneumatic circuit and the block diagram in figure 1 demonstrate the working principle of the setup controlling the grasping operation of bellows suction cups (SMC ZP2-B15JS). The circuit is supplied with 4 bar pressure input which is adjusted via a pressure regulator based on a potentiometer reading connected to an Arduino board. The regulated pressure flows through an ejector which generates a negative pressure for the suction cup that is proportional to the set positive pressure. The vacuum cup activation and release are controlled by the valve via an Arduino signal from a push button. The Arduino program sets the duration of suction and the analogue signal to the regulator. The pressure gauge provides a visual display of the negative pressure during the operation.



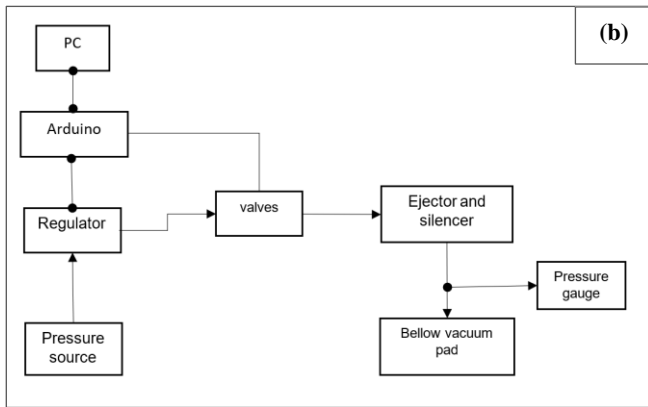


Figure 1: (a) Pneumatic Circuit (b) Block diagram of the system

Figure 2 shows the testing setup with the suction cup connected to the pneumatic circuit. The cup is carried by a stepper motor to enable rotation, which is mounted on a controlled motorized stage that moves the gripper vertically to approach and lift the mushroom. The setup allows mimicking a picking routine that involves lowering the cup until pressing against the mushroom, then activating the rotation motor to twist the suction cup. Sample mushrooms are fixed to a 3D printed piece mounted on a sensitive force-torque sensor (Schunk Mini40) to record the maximum reaction torque that occurs during grasping.




Figure 2: The test setup of the vacuum cup gripper

IV. RESULTS

Figure 3 shows the relationship between the resulting negative pressure at the vacuum cup and the supplied voltage to the pressure regulator for small and large mushrooms sizes. The relationship was mostly linear and was not significantly affected by the difference in mushroom size. Hence, the relationship can be used to estimate the voltage value required to achieve a particular negative pressure value.




Figure 3: the relationship between the input voltage and the resulting negative pressure at the cup

Moreover, the next stage of the experiment tested thirty mushrooms of different sizes at increasing vacuum levels to find out at what negative pressure the surface bruising starts to occur. The tested mushrooms were monitored over a 24-hour period since bruising may occur later. The results showed that mushrooms suffered no damage up until -0.02 bar at the end of the monitored duration. Figure 4 shows two mushrooms as an example to highlight when bruising occurs. A red circle was drawn around the area where the suction cup touched the mushroom to highlight bruising due to the cup.




Figure 4: (a) no damage at a negative pressure of -0.02 bar (b) visible damage at negative pressure -0.08

Finally, a test was performed at the identified vacuum level of -0.02 bar on three mushrooms to evaluate the maximum torque that can be generated during twisting at this value. The results showed that the maximum resulting torque when no bruising to the cup occurs was on average 0.0038 Nm.

Table 1: Torque test values

Pressure in (bar)	Negative pressure (bar)	Voltage (V)	Average Torque (Nm)
1.19	-0.02	0.65	0.0038

V. CONCLUSIONS AND FUTURE WORK

The outcomes of the preliminary work presented here showed that the negative pressure that can lift loose mushrooms using soft suction cups without bruising was -0.02 bar. At this value, the torque generated during twisting fixed mushrooms was on average 0.0038 Nm, which will not always be enough to break the mushroom stalk for automated harvesting tasks. Further investigation into parameters other than size, such as maturity, that could impact the results is still needed. Future work will investigate methods to improve the design of the suction cups to increase the maximum torque without damaging the mushroom for better performance. This could involve combining multiple smaller suction cups or creating custom contact surfaces to better distribute the pressure. Nevertheless, the initial work so far identified an initial benchmark for future development.

VI. ACKNOWLEDGEMENT

This work has been funded by the School of Engineering at the University of Lincoln.

VII. REFERENCES

- [1] B. S. Homberg, R. K. Katschmann, M. R. Dogar, and D. Rus, 'Haptic identification of objects using a modular soft robotic gripper', *IEEE Int. Conf. Intell.*

- Robot. Syst.*, vol. 2015-Decem, pp. 1698–1705, 2015.
- [2] P. Polygerinos *et al.*, ‘Soft Robotics: Review of Fluid-Driven Intrinsically Soft Devices; Manufacturing, Sensing, Control, and Applications in Human-Robot Interaction’, *Advanced Engineering Materials*, vol. 19, no. 12. Wiley-VCH Verlag, 01-Dec-2017.
- [3] J. Shintake, V. Cacucciolo, D. Floreano, and H. Shea, ‘Soft Robotic Grippers’, *Advanced Materials*, vol. 30, no. 29. p. 1707035, Jul-2018.
- [4] J. Shintake, S. Rosset, B. Schubert, D. Floreano, and H. Shea, ‘Versatile Soft Grippers with Intrinsic Electrodehesion Based on Multifunctional Polymer Actuators’, *Adv. Mater.*, vol. 28, no. 2, pp. 231–238, 2016.
- [5] J. Shintake, B. Schubert, S. Rosset, H. Shea, and D. Floreano, ‘Variable stiffness actuator for soft robotics using dielectric elastomer and low-melting-point alloy’, *IEEE Int. Conf. Intell. Robot. Syst.*, vol. 2015-Decem, pp. 1097–1102, 2015.
- [6] E. Brown *et al.*, ‘Universal robotic gripper based on the jamming of granular material’, *Proc. Natl. Acad. Sci. U. S. A.*, vol. 107, no. 44, pp. 18809–18814, 2010.
- [8] T. Nishida, Y. Okatani, and K. Tadakuma, ‘Development of Universal Robot Gripper Using MR α Fluid’, *Int. J. Humanoid Robot.*, vol. 13, no. 4, pp. 1–13, 2016.
- [9] J. Guo, K. Elgeneidy, C. Xiang, N. Lohse, L. Justham, and J. Rossiter, ‘Soft pneumatic grippers embedded with stretchable electroadhesion’, *Smart Mater. Struct.*, vol. 27, no. 5, 2018.
- [10] G. Fantoni *et al.*, ‘Grasping devices and methods in automated production processes’, *CIRP Ann. - Manuf. Technol.*, vol. 63, no. 2, pp. 679–701, 2014.

Mapping the impact of faults in a multi-robot team

Dingdian Zhang¹

¹*Dept of Engineering
King's College London, UK
dingdian.zhang@kcl.ac.uk*

Elizabeth I Sklar^{1,2}

²*Lincoln Institute for Agri-food Technology
University of Lincoln, UK
esklar@lincoln.ac.uk*

<https://doi.org/10.31256/Gp6Fe7B>

Abstract—Hardware failure can have significant impact on the performance of a multi-robot team. The work presented here provides a preliminary assessment of this impact based on experiments conducted with physical robots. Two types of hardware failure are considered: motor failure and laser sensor failure. These can be *permanent* (where the robot does not recover during the mission) or *temporary* (where the robot does recover). Results show that permanent motor failure and laser sensor faults decrease the success rate of task completion to varying extents and laser sensor faults also cause significant errors in measuring distances travelled by the robot experiencing these faults. While these results are not unexpected, the contribution of this work is in laying out a structured baseline that will be used in the near future for comparison of various recovery strategies.

Index Terms—multi-robot teams; fault tolerance

I. INTRODUCTION

A multi-robot team is expected to coordinate behaviours amongst team members in order to complete a set of tasks in such a way that the team outperforms what a single-robot system can achieve. Distribution of tasks amongst multi-robot team members is a well-studied problem. Much work has focused on producing algorithms to optimise *task allocation*, within restricted settings—without evaluating the effectiveness of the allocation when tasks are actually executed and without considering the impact if individual robots experience hardware failure during a mission. Our prior work has pushed beyond the *allocation* stage and concentrated on a systematic and empirical study of the impact of market-based task allocation mechanisms when applied to a diverse landscape of missions and environments [1]–[5]. The emphasis in that work on experimental results, particularly results obtained from physical robot teams, has led to questions of *fault tolerance* in a multi-robot team and motivated our present line of research.

The study presented here investigates the impact on team performance when one of the robots on the team experiences hardware failure, such as loss of ability to move (“motor failure”) or loss of sensing (“sensor failure”). We emulate failure in a controlled way, so that we can measure the impact of (planned) hardware faults on team performance metrics. Results of experiments conducted with physical robots reveal the changes in outcome when faults are “permanent” (robot does not recover during a mission) as compared to “temporary” (robot recovers during the mission).

II. BACKGROUND

Fault tolerance refers to the ability of a system to continue to operate, albeit perhaps at a lower performance level, when

the system suffers from undesired situations (e.g. internal faults and/or interference from the outside environment) [6]. A brief survey of experiments conducted with multi-robot teams examining the effects of faulty team members reveals that a reliable and fault-tolerant multi-robot system comes at a cost. There are three main questions which need to be addressed [7]: (1) how to detect whether robots have faults or not; (2) how to identify or diagnose robot failures; and (3) how to recover from the failures.

Winfield & Nembrini [8] explore fault tolerance in a robot swarm by using their *Failure Mode and Effect Analysis* (FMEA) approach [9]. They test a containment task and analyse the effects hazards have on a robot swarm. They show that partial failure is more harmful than complete failure, e.g. a robot with broken motors will have an anchoring effect, whereas complete failure, which causes the robot to become a static obstacle, has less serious effects on the whole system.

Bjerknes & Winfield [10] study a swarm taxis task in which a team of robots form a coherent group and then move toward a beacon. They consider three failure modes: complete failure of individual robots; failure of a robot’s infrared (IR) sensors; and failure of a robot’s motors only. Results show that IR sensor failure slows down the progress of the swarm, but the whole swarm still reaches the beacon; whereas motor failures cause the swarm to hover around the failed robots, eventually escaping the influence of the failure robots.

Timmis et al. [11] propose a self-healing approach. Using the case study of [10], they consider motor failure due to lack of power, where the remaining power is enough to support communication. Faulty robots emit a faulty signal to their neighbours and then healthy robots share their energy by recharging the drained batteries of faulty robots. As evaluated in simulation, the self-healing system works well even when half the robots in the swarm experience low-energy faults.

Kamel et al. [12] present fault-tolerant cooperative control strategies to deal with actuator faults in a multi-robot formation task, and later work [13] considers partial actuator faults. When complete failure occurs, the formation outcome changes; whereas partial failure results in slower performance but eventual achievement of the desired formation.

Our long term research goal focuses on generating robust strategies for recovering from partial and total motor and sensor failures. Here we present empirical baseline results obtained from a number of experiments conducted on physical robots where faults were emulated.




Fig. 1. Experimental setup

III. APPROACH

In order to investigate the impact of hardware failure in a multi-robot team, we induce different types of failures in the "faulty" robot at a randomly generated time point. We define two sets of parameters to classify each failure mode: *Permanent Failure* (which cannot be recovered by the robot itself and will persist for the remainder of the experiment) versus *Temporary Failure* (which only exists for a short period of time and the robot will recover autonomously without any help); and *Complete Failure* (where the device is entirely broken) versus *Partial Failure* (where the device works but with performance degradation). This work investigates two types of hardware failure (motor and laser sensor) in three failure modes: complete (permanent) motor failure (*CMF*); temporary motor failure (*TMF*); or permanent laser sensor failure (*PLF*). No failure (*NF*) is included for comparison.

To evaluate the impact of each failure type, we consider metrics that measure the performance of both individual robots and the team as a whole. Here we focus on the **distance travelled** by the robots executing allocated tasks; and **success rate** of task completion by each robot.

IV. EXPERIMENTS AND RESULTS

We conducted 320 trials using the MRTEAm framework [4], [14] and Turtlebot 3 Burger platform performing exploration tasks in an office-like environment (Figure 1). Our "faulty" robot employs a modified controller which will induce failure after a randomly generated time, chosen from a Gaussian distribution, and last for a random amount of time¹. Experimental missions have 8 tasks allocated to 3 robots using the *Sequential Single Item (SSI)* auction mechanism [15]. Missions are classified according to scenario and starting condition. Scenarios are defined by several parameters: *single-robot (SR)* versus *multi-robot (MR)* tasks; *independent (IT)* versus *constrained (CT)* tasks; and *static (SA)* versus *dynamic (DA)* task arrival (before or during mission execution, respectively).

A sample set of results is shown in Figures 2 and 3. Figure 2 illustrates trajectories travelled by sample robots impacted by complete motor failure (*CMF*) (Fig. 2a) and partial laser sensor failure (*PLF*) (Fig. 2a) in one mission configuration (scenario 1, distributed start, MR-CT-SA configuration). Figure 3 illustrates performance metrics for the same mission. Figure 3a

¹The parameters used: time of inducing failure: $\mu = 20, \sigma = 5$, and duration of temporary failure: $\mu = 10, \sigma = 2$. Future work will compare other distributions, such as Poisson.







Fig. 2. Trajectories travelled by sample robots impacted by (a) complete motor failure (*CMF*) and (b) partial laser failure (*PLF*) in one mission configuration (scenario 1, distributed start, MR-CT-SA configuration); each mission involves 3 robots performing 8 exploration tasks ($x_1 \dots x_8$).



(a) total distance travelled



(b) distance per robot



(c) success rate per robot

Fig. 3. Statistics from one mission configuration (MR-CT-SA, distributed start, scenario 1)

shows the total distance travelled by the team as a whole, indicating clearly that when there is a partial laser sensor failure, the faulty robot has trouble localising and travels much further than it needs to, contributing to an outside distance for that experimental condition. This is highlighted by the red path shown in Figure 2b and by the tallest (dark) bar in Figure 3b. Figure 3c shows the success rate, which is measured as the percentage of assigned tasks that were completed. When failure is temporary and there are no failures, then all assigned tasks are completed. Complete motor failure has a more detrimental effect than partial laser failure.

V. SUMMARY

We have shown preliminary results of baseline experiments designed to illustrate the impact of various types of hardware failures on the performance metrics achieved by a multi-robot team. A range of scenarios and mission configurations have been produced, though space constraints here allow showing only a sample for one such configuration. A more comprehensive report with complete results is under preparation. Our next steps involve development of response strategies when failures occur, followed by experimental evaluation of our strategies in comparison with others mentioned in Section II.

REFERENCES

- [1] E. Schneider, O. Balas, A. T. Özgelen, E. I. Sklar, and S. Parsons, "An Empirical Evaluation of Auction-based Task Allocation in Multi-Robot Teams (Extended Abstract)," in *Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, Paris, France, May 2014.
- [2] —, "Evaluating auction-based task allocation in multi-robot teams," in *Workshop on Autonomous Robots and Multirobot Systems (ARMS) at Autonomous Agents and MultiAgent Systems (AAMAS)*, Paris, France, May 2014.
- [3] E. Schneider, E. I. Sklar, and S. Parsons, "Evaluating multi-robot teamwork in parameterised environments," in *Towards Autonomous Robotic Systems: 17th Annual Conference (TAROS)*. Springer, 2016.
- [4] E. Schneider, E. I. Sklar, S. Parsons, and A. T. Özgelen, "Auction-based task allocation for multi-robot teams in dynamic environments," in *Towards Autonomous Robotic Systems: 16th Annual Conference (TAROS)*. Springer, 2015.
- [5] D. Zhang, E. Schneider, and E. I. Sklar, "A cross-landscape evaluation of multi-robot team performance in static task-allocation domains," in *Towards Autonomous Robotic Systems: 20th Annual Conference (TAROS)*. Springer, 2019, pp. 261–272.
- [6] L. E. Parker, "Reliability and fault tolerance in collective robot systems," *Handbook on Collective Robotics: Fundamentals and Challenges*, Pan Stanford Publishing, Singapore, 2012.
- [7] D. Crestani, K. Godary-Dejean, and L. Lapierre, "Enhancing fault tolerance of autonomous mobile robots," *Robotics and Autonomous Systems*, vol. 68, pp. 140 – 155, 2015.
- [8] A. F. T. Winfield and J. Nembrini, "Safety in numbers: Fault tolerance in robot swarms," *International Journal on Modelling Identification and Control*, vol. 1, no. 1, pp. 30–37, 2006.
- [9] W. Kenneth, "The fmea pocket handbook," 2004.
- [10] J. D. Bjercknes and A. F. T. Winfield, *On Fault Tolerance and Scalability of Swarm Robotic Systems*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 431–444.
- [11] J. Timmis, A. Ismail, J. Bjercknes, and A. Winfield, "An immune-inspired swarm aggregation algorithm for self-healing swarm robotic systems," *Biosystems*, vol. 146, pp. 60 – 76, 2016, information Processing in Cells and Tissues.
- [12] M. A. Kamel, Y. Zhang, and X. Yu, "Fault-tolerant cooperative control of multiple wheeled mobile robots under actuator faults," *IFAC-PapersOnLine*, vol. 48, no. 21, pp. 1152 – 1157, 2015, 9th IFAC Symposium on Fault Detection, Supervision and Safety for Technical Processes SAFEPROCESS 2015.
- [13] M. A. Kamel, X. Yu, and Y. Zhang, "Fault-tolerant cooperative control design of multiple wheeled mobile robots," *IEEE Transactions on Control Systems Technology*, vol. 26, no. 2, pp. 756–764, March 2018.
- [14] E. Schneider, "Mechanism selection for multi-robot task allocation," Ph.D. dissertation, University of Liverpool, 2018.
- [15] S. Koenig, C. Tovey, M. Lagoudakis, V. Markakis, D. Kempe, P. Keskincak, A. Kleywegt, A. Meyerson, and S. Jain, "The power of sequential single-item auctions for agent coordination," in *21st Natl Conf on Artificial Intelligence (AAAI)*, vol. 2, 2006.

Image Pre-processing vs. Transfer Learning for Visual Route Navigation

William H. B. Smith¹, Yvan Petillot² and Robert B. Fisher³

<https://doi.org/10.31256/Nh4Vy4L>

Abstract—This paper investigates image pre-processing and triplet learning for place recognition in route navigation. The first contribution combines image pre-processing and ImageNet pre-trained neural networks for generating improved image descriptors. The second contribution is a fast, compact ‘FullDrop’ layer that can be appended to an ImageNet pre-trained network and taught to generate invariant image descriptors with triplet learning. The proposals decrease inference time by 8x and parameters by 30x while keeping comparable performance to NetVLAD, the state of the art for this task

I. INTRODUCTION

Neural network features provide state of the art recognition accuracy [1], but are affected by weather, lighting and man-made changes for visual place recognition [2]. This paper compares two techniques to create descriptors that are robust to these visual changes. Firstly, images are pre-processed before their descriptors are extracted by a pre-trained neural network. Secondly a ‘FullDrop’ layer is appended to the pre-trained network and trained with a custom triplet learning scheme to model a chosen route in a variety of conditions 3x faster with 30x fewer parameters than the state of the art for this task: NetVLAD [3]. Not all route appearances can be included in training data for deep learning solutions to this problem. The fast re-training possible with this paper’s approach instead seeks to generalise across just one specified route in multiple conditions. The performance of the embedded descriptors generated by both approaches is compared with those generated by the pre-trained NetVLAD. The Oxford RobotCar Dataset [4] is used for evaluation. This paper’s two contributions are:

- 1) Image pre-processing and neural network combinations for improved feature descriptors.
- 2) Triplet learning scheme and ‘FullDrop’ layer to generate image descriptors 8x faster with 30x fewer parameters than NetVLAD and comparable performance.

II. BACKGROUND

Initially, CNN’s pre-trained on the ImageNet object recognition dataset were used to generate image descriptors for place recognition [5]. Novel architectures [6] were then trained end-to-end from scratch for place recognition. Triplet learning achieved positive results [7], but struggled to compete with off the shelf CNN descriptors.

This work was supported by the EPSRC Centre for Doctoral Training in Robotics and Autonomous Systems.

¹William Smith, The University of Edinburgh and Heriot-Watt University, Edinburgh, EH14 4AS, UK whbsmith@gmail.com

²Yvan Petillot, Heriot-Watt University, Edinburgh, EH14 4AS, UK y.r.petillot@hw.ac.uk

³Robert B. Fisher, School of Informatics, University of Edinburgh, EH8 9LE, UK rbf@inf.ed.ac.uk




Fig. 1. Pre-processing methods (PREP 1-5) on day and night images, left to right: adaptive histogram normalisation; histogram normalisation; histogram normalisation and subtraction of the image’s mean/std. deviation; subtraction of the image’s mean/std. deviation; pixel-wise RGB normalisation.

NetVLAD [3] is the current state of the art in place image retrieval. VGG16 is frozen for training down to, but not including, the Conv5 layer then trained using triplet learning with a final VLAD layer. Performance is evaluated in a city environment at minimum intervals of 12 metres and in other environments at discrete locations. New pre-trained ImageNet networks, such as MobileNet have been released but can’t be used in NetVLAD without lengthy re-training.

III. METHOD

A. Image Pre-processing

Images from a single route in two different conditions were pre-processed using the five techniques (PREP 1-5) in Figure 1 and embedded image descriptors were extracted using VGG16 pre-trained on ImageNet data.

B. Triplet Network

Front-facing, pre-recorded and geotagged example traversal videos of a single route are selected for training data. A custom ‘FullDrop’ layer is appended to VGG16 and MobileNet pre-trained on ImageNet data with the final classification layer removed. The FullDrop layer is trained to model the route with triplet learning [8]. The choice of triplets for learning has a significant effect on performance. The FullDrop layer triplet mining differs from NetVLAD by using positive images from less than or equal to 10 adjacent frames from the anchor and negative images more than 10 adjacent frames away.

Six video traversals of the same route: Early Evening, Morning, Midday, Rain, Afternoon and Overcast were sampled and synchronised for training. Each traversal was 884 frames at intervals of approximate 2.25m, making a total of 5304 training images. Approximately 51,000 triplets were mined for 10 training epochs in batches of 16 with a margin of 1. NetVLAD’s training differed from the FullDrop training

by unfreezing the Conv5 layer of the base network at training time and only used the hardest available triplets. Triplet learning relies on the raw difference between images so the FullDrop layer was not combined with image pre-processing.




Fig. 2. Architecture of the system, including the ‘FullDrop’ layer. W is the width of the layer, which in this case was 512. The dropout parameter was 0.4 and L2 normalization was 0.1. This layer can be appended to any ImageNet pre-trained network and used for this task.

IV. RESULTS

The NetVLAD model ‘VGG-16 + NetVLAD + whitening’ published online by [3] was used as a comparison. NetVLAD was pre-trained on the urban Pittsburgh 30k dataset and shown to be effective on a variety of different urban and suburban locations.

Night and day traversals of a 2km route synchronised and sampled to 878 images from the Oxford RobotCar Dataset taken months apart were used for the evaluation in Table IV-1. Frames from each route were pre-processed and passed through the pre-trained VGG16 or just through the FullDrop or NetVLAD model to produce embedded image descriptors which were compared using the Euclidean distance; the 20 closest matches were generated and the one closest to the ground truth was identified as the localisation prediction in an effort to balance the potential accuracy of a probability-based system with pure place recognition.

1) *Image Pre-processing*: Four of the techniques reduced place recognition error. Specifically PREP 5 reduces mean error by 40.9% and median error by 63.1%.

TABLE I

THE LOCALISATION ERROR (METRES) OF NIGHT AND DAY ROUTE TRAVERSALS, WITH INFERENCE TIME AND TRAINING PARAMETER COUNT, AS DESCRIBED IN SECTIONS III-A AND IV

	μ (m)	Median (m)	σ (m)	Inf.T. (ms)	Params.
MobileNet + FullDrop	13.3	3.32	25.8	7	8×10^5
NetVLAD	15.5	2.39	29.1	80	2.4×10^7
VGG16 + FullDrop	49.4	9.57	73.8	10	5×10^5
PREP 5 + VGG16	50.9	29.0	54.5	11	-
PREP 1 + VGG16	70.1	54.4	66.4	10	-
PREP 2 + VGG16	72.6	65.7	63.6	10	-
PREP 4 + VGG16	73.0	56.4	65.2	11	-
VGG16	83.6	78.5	62.1	10	-
PREP 3 + VGG16	96.8	58.8	94.9	10	-

2) *FullDrop vs. NetVLAD*: The taught VGG16 + FullDrop model reduces the median error from PREP 5 by a further 67.0%. MobileNet + FullDrop reduces mean error by 14.2% compared to NetVLAD. The FullDrop model is capable of generalisation to unseen conditions and produces median localisation predictions that are within 7.2m of NetVLAD’s, which was extensively pre-trained for this task.

Figure 3 illustrates the comparison between the two evaluation routes’ descriptors. The brief, but route specific training for the FullDrop descriptors allow them to represent the similarities between similar sections of route more accurately in comparison to NetVLAD which shows a weaker relationship between adjacent frames. The results suggest FullDrop descriptors are better for utilising the relationship between consecutive frame descriptors for navigation. FullDrop’s more accurate model of the route may introduce confusion between nearby, similar frames.

The MobileNet + FullDrop and NetVLAD models consisted of approximately 8×10^5 and 2.4×10^7 parameters respectively. An Intel i7 CPU and Nvidia GTX 1070 trained the FullDrop layer in 510 seconds once the features had been extracted using the base network. Training a VLAD layer took 3x longer, re-training the full NetVLAD would take far longer.

3) *Inference Time*: NetVLAD (VGG16 + VLAD layer), VGG16 + FullDrop and MobileNet + FullDrop take 80ms, 10ms and 7ms to generate embedded descriptors on an Nvidia GTX 1070.




Fig. 3. Partial difference matrices of night and day route traversals from Table IV-1. Showing the relative embedded descriptor relationships. The darker the blue the lower the match error.

V. CONCLUSION

Deep learning for discrete place recognition provides inspiration for real-time continuous route navigation. Reduced complexity decreases inference and training time while maintaining generalisation across a single route and resilience to the problem of perceptual aliasing. Image pre-processing improves performance and maintains decreased inference time with no need for re-training.

REFERENCES

- [1] A. Sharif, R. H. Azizpour, J. Sullivan, and S. Carlsson, "CNN Features off-the-shelf: an Astounding Baseline for Recognition," in *IEEE Int. Conf. on Computer Vision and Pattern Recognition*, 2014.
- [2] S. Lowry, N. Sunderhauf, P. Newman, J. J. Leonard, D. Cox, P. Corke, and M. J. Milford, "Visual Place Recognition: A Survey," *IEEE Transactions on Robotics*, vol. 32, no. 1, pp. 1–19, 2016.
- [3] R. Arandjelovic, P. Gronat, A. Torii, T. Pajdla, and J. Sivic, "NetVLAD: CNN Architecture for Weakly Supervised Place Recognition," in *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 40, no. 6, 2018, pp. 1437–1451.
- [4] W. Maddern, G. Pascoe, C. Linegar, and P. Newman, "1 year, 1000 km: The Oxford RobotCar dataset," *Int. J. of Robotics Research*, vol. 36, no. 1, pp. 3–15, 2017.
- [5] T. Naseer, W. Burgard, and C. Stachniss, "Robust Visual SLAM Across Seasons," in *IEEE Trans. on Robotics*, vol. 34(2):289-302, 2018.
- [6] Z. Chen, A. Jacobson, N. Sunderhauf, B. Upcroft, L. Liu, C. Shen, I. Reid, and M. Milford, "Deep learning features at scale for visual place recognition," in *Proceedings - IEEE International Conference on Robotics and Automation*, 2017.
- [7] R. Gomez-Ojeda, M. Lopez-Antequera, N. Petkov, and J. Gonzalez-Jimenez, "Training a Convolutional Neural Network for Appearance-Invariant Place Recognition," in *ArXiv pre-print*, 2015.
- [8] D. K. Schroff, Florian and J. Philbin, "Facenet: A unified embedding for face recognition and clustering," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015.

Unsupervised Anomaly Detection for Safe Robot Operations

Pratik Somaiya, Marc Hanheide and Grzegorz Cielniak

Lincoln Centre for Autonomous Systems

University of Lincoln, UK

{psomaiya,mhanheide,gcielniak}@lincoln.ac.uk

<https://doi.org/10.31256/Wg7Ap8J>

Abstract—Faults in robot operations are risky, particularly when robots are operating in the same environment as humans. Early detection of such faults is necessary to prevent further escalation and endangering human life. However, due to sensor noise and unforeseen faults in robots, creating a model for fault prediction is difficult. Existing supervised data-driven approaches rely on large amounts of labelled data for detecting anomalies, which is impractical in real applications. In this paper, we present an unsupervised machine learning approach for this purpose, which requires only data corresponding to the normal operation of the robot. We demonstrate how to fuse multi-modal information from robot motion sensors and evaluate the proposed framework in multiple scenarios collected from a real mobile robot.

Index Terms—anomaly detection, one-class SVM, safety

I. INTRODUCTION

Recently robots have started replacing humans in areas where the jobs are mostly dull, repetitive or dangerous for humans. There are many examples in areas such as agriculture, tourism, logistics and transport where the robots have either fully replaced or they are accompanying humans. One of the most important concerns of operating robots in human-environment is safety. Usually, robots use sensor data to detect the presence of any kind of object or the human but the noise in the sensors or its malfunctioning can cause a disaster. Detection of these kinds of faults at the earliest is very important before causing serious damage. In the real world, it is not feasible to foresee all kinds of possible faults and therefore these can not be modelled easily. Hence, a data-driven approach to detect these kinds of anomalies is required. We propose to use an unsupervised technique which requires only data corresponding to the normal operation of the robot - namely one-class support vector machine (OCSVM). The technique was used in [1] for collision detection and collision point localisation in a humanoid which can help the remote operator to stop the robot in case of an emergency. In [2] the authors used an isolation forest-based anomaly detection method to detect the anomalous behaviour in Unmanned Aerial Vehicles (UAVs). The contribution of our paper is to use OCSVM for anomaly detection in the operation of a mobile robot based on motion data coming from the robot's motion sensors. Furthermore, we also evaluate the proposed framework in multiple real-world scenarios and present the results based on data collected from a real mobile robot.




Fig. 1. An overview of the proposed data-driven anomaly detection system.

II. METHODOLOGY

A. Overview

The general system overview is presented in Fig. 1. The sensor data coming from the robot include odometry, relative motor power and speed for each wheel, linear acceleration from the IMU and issued command velocity. The frequency of received sensor messages varies and therefore in the first instance, these are synchronised. We up-sample all sensor message to 100 Hz. The synchronised messages form a feature vector of 11 values for each time instance which is an input to the one-class SVM classifier.

B. One-class SVM

OCSVM was first proposed by Schölkopf et al. [3] as an extension to the SVM. The method does not require the labelled data from two classes and can be trained using the data from one class only. OCSVM uses a kernel function $k(\mathbf{x}_i, \mathbf{x}_j)$ to map the features \mathbf{x} into a high-dimensional space $\phi(\cdot)$ where it finds a hyper-plane $\mathbf{w} \cdot \phi(\mathbf{x}) - \rho = 0$ separating most of the data from the origin. This is achieved by solving the following quadratic program that maximises the distance between the hyper-plane and the origin:

$$\begin{aligned} \min_{\mathbf{w} \in F, \xi \in \mathbb{R}^\ell, \rho \in \mathbb{R}} \quad & \frac{1}{2} \|\mathbf{w}\|^2 + \frac{1}{\nu \ell} \sum_i \xi_i - \rho \\ \text{subject to} \quad & (\mathbf{w} \cdot \phi(\mathbf{x}_i)) \geq \rho - \xi_i, \quad \xi_i \geq 0 \end{aligned} \quad (1)$$

where $\nu \in (0, 1]$ is an upper bound on the fraction of outliers and a lower bound on the number of training examples used as support vectors, ξ_i are slack variables and ρ is bias. By using Lagrange techniques, the decision function can be given by, $f(\mathbf{x}) = \text{sgn}(\sum_i \alpha_i k(\mathbf{x}_i, \mathbf{x}) - \rho)$, where $i = 1, \dots, \ell$ and α_i are Lagrange multipliers. An incoming datum \mathbf{x}_n is determined as the anomalous if $f(\mathbf{x}_n) < 0$.

III. EXPERIMENTS

A. Experimental Setup

To test the feasibility of the proposed system, we have designed a set of experiments with a real outdoor mobile robot Thorvald [4]. The robot is equipped with wheel encoders, motor controller and IMU. The data collection was performed by driving the robot manually using a joystick and issuing forward or backward command velocities. For validation and evaluation, we have also devised a set of anomalous situations by pushing the robot and recording the ground-truth using a joystick. For the normal operation, we ran the robot at fixed velocities of 0.15, 0.30 and 0.50 m/s (see Table I, dataset 1-3). Dataset 4 was created by merging sets 1-3 to represent a mix of normal behaviours. In the same way, datasets 5-10 represent examples with anomalous cases: sets 5-6, 7-8 and 9-10 correspond to speeds of 0.15, 0.30 and 0.50 m/s respectively.

TABLE I
DATA COLLECTED

number	Samples	Target Class	Anomalies
1	51091	51091	0 (0.00%)
2	19045	19045	0 (0.00%)
3	11037	11037	0 (0.00%)
4	81173	81173	0 (0.00%)
5	18156	17860	296 (1.63%)
6	11940	11681	259 (2.17%)
7	10092	9792	300 (2.97%)
8	7816	7632	184 (2.35%)
9	5508	5455	53 (0.96%)
10	5096	4978	118 (2.32%)

B. Data pre-processing

In the feature vector, some of the features have a higher order of variance than others and they might not allow the classifier to learn from other features as expected. To avoid that, we used standard scaler which scales each feature of the training data such that the mean of each feature is zero and the variance is unit. Later, these mean and variance are used to transform the test data.

C. Hyperparameter Selection

We used radial basis function (RBF) kernel for OCSVM thus the number of hyperparameters becomes two: 1) RBF kernel coefficient gamma (γ) and 2) nu (ν). The performance of OCSVM highly depends on these hyperparameters. The value of ν was set to 0.0001 based on experiments while γ was calculated with the equation, $\gamma = 1/(n * \widehat{\text{Var}}(X))$, where n is the number of features, $\widehat{\text{Var}}$ is the variance and X is the training data.

D. Evaluation Metrics

As the datasets are highly imbalanced, Cohen's Kappa coefficient (κ) and specificity (S) were used to evaluate the performance of the classifier. Cohen's Kappa gives the chance agreement between the observational accuracy and the




Fig. 2. Anomaly detection: true detection (left) and false detection (right)

TABLE II
OCSVM PERFORMANCE EVALUATION

Train	Velocity	Test					
		0.15		0.30		0.50	
		S	κ	S	κ	S	κ
0.15	0.32	0.38	1.00	0.01	1.00	0.01	
0.30	0.01	0.02	0.31	0.39	1.00	0.01	
0.50	0.01	0.02	0.04	0.08	0.33	0.33	
0.15+0.30+0.50	0.13	0.22	0.07	0.12	0.38	0.3	

expected accuracy [5] while specificity measures a classifier's ability to identify the anomalous data.

E. Results

We combined the datasets with the same velocity from Table I and reported the results for the entire dataset in Table II. In the latter, we can see that OCSVM performed better when the training and testing data have the same magnitude of velocity. Further, when the velocity in the training was higher than the testing dataset, OCSVM missed anomalies and when it was lower, OCSVM classified most of the samples as anomalies. It is because as the velocity increases the variance of the features increase and as mentioned in III-C, γ was chosen based on the variance of the training data. For example, when the velocity in the training was lower than the testing dataset, specificity (S) was 1.00 but the value of κ was 0.01, which suggests that the performance of the classifier was poor. Finally, when the model was trained with dataset 4, it was able to predict some anomalies in the case of 0.15 and 0.50 m/s. Fig. 2 shows an example case of the anomaly detection system with the output, ground truth and some of the features.

IV. CONCLUSION AND FUTURE WORK

In this paper, we proposed OCSVM based anomaly detection method which uses multi-modal data fusion to detect the anomalous operation of a robot in human accompanied environments. We evaluated our approach in multiple real-world scenarios.

In future, we are planning to understand more about the temporal aspect of the data which might help the classifier to learn better and potentially improve the classification performance. In addition to that, we are focusing on feature selection approaches for one-class classification.

REFERENCES

- [1] Narukawa, Kaname, Takahide Yoshiike, Kenta Tanaka, and Mitsuhide Kuroda. "Real-time collision detection based on one class SVM for safe movement of humanoid robot." In 2017 IEEE-RAS 17th International Conference on Humanoid Robotics (Humanoids), pp. 791-796. IEEE, 2017.
- [2] Khan, Samir, Chun Fui Liew, Takehisa Yairi, and Richard McWilliam. "Unsupervised anomaly detection in unmanned aerial vehicles." *Applied Soft Computing* 83 (2019): 105650.
- [3] Schölkopf, Bernhard, John C. Platt, John Shawe-Taylor, Alex J. Smola, and Robert C. Williamson. "Estimating the support of a high-dimensional distribution." *Neural computation* 13, no. 7 (2001): 1443-1471.
- [4] Grimstad, Lars, and Pål Johan From. "The Thorvald II agricultural robotic system." *Robotics* 6, no. 4 (2017): 24.
- [5] Cohen, Jacob. "A coefficient of agreement for nominal scales." *Educational and psychological measurement* 20, no. 1 (1960): 37-46.

Intelligent Service Robots to enter the Hospitality Industry: Job Termination or Gospel to Waiters

Jiaji Yang ¹[0000 – 0002 – 1011 – 9676] and Esyin Chew ²[0000 – 0003 – 2644 – 9888]

^{1,2}Cardiff School of Technologies, Cardiff Metropolitan University

EUREKA Robotics Lab

Cardiff, UK

<https://doi.org/10.31256/Yo8Tk2T>

Abstract—This research builds on the existing capabilities of the EUREKA Robotics Lab at Cardiff Metropolitan University and aims to disrupt and enhance humanoid robotic programs by studying the public’s acceptance of the use of AI robots in the hospitality industry in Wales. As the first Humanoid Robotic Waitress pilot in Wales, the research develops a novel robot serving program and allows intelligent robots to serve in a café in Cardiff. In addition, the study first proposed an cutting-edge interactive model for hospitality robots. Data were collected through questionnaires, interview and observation to explore the acceptance of service robots in Wales.

Index Terms—Service robot, Robot acceptance, Robot and human interaction

I. BACKGROUND AND PURPOSE

Since its inception, robots have been widely used in industrial manufacturing and automation control. With the continuous development of technology, robots have gradually begun to shift from large-scale automation in manufacturing to small-scale of autonomous and social services. There are increasing trend for humanoid robots to assist humans to complete simple service tasks in the hospitality industry in China and Japan. However, as the hospitality industry itself is an industry with services as its core product, the acceptance of robots by stakeholders is particularly important. Although robots as an emerging product have been used in the service industry to a certain extent to attract customers, not all of the known cases have shown promising results. The world’s first robot hotel certified by the Guinness Book of Records, Henn-na has been operating in Japan since 2015. Hotel receptionists, messengers, cleaners and luggage carriers are almost robots, however, Japan’s Henn-na Hotel has laid off half of its 243 robots after they created more problems than they could solve, as first reported by The Wall Street Journal [1]. Why is this happening? Is there a deficiency in the “intelligentisation” of the hospitality industry? Are there concerns about these service robots? With the application of robotics blowouts in the Fourth Industrialisation era, exploring public acceptance of service robots is essential: are smart robots entering the hospitality industry lead to job termination or the gospel to waiters and waitress? The research builds on the existing capabilities of the EUREKA Robotics Laboratory at Cardiff

All robotics equipment is sponsored by EUREKA Robotics Lab at Cardiff Metropolitan University

Metropolitan University, with the aim of researching and designing humanoid robot programmes to disrupt and enhance the hospitality sector in Wales. By investigating the public acceptance of AI robots working in the hospitality industry.

II. METHODOLOGY AND RESEARCH APPROACH

Researchers use investigative research method to pilot the conceptual Human-Robot Interaction (HRI) model by experimenting a humanoid service robot, namely Robot EUREKA Gen-1, in a restaurant environment in Cardiff. Researchers collect both qualitative and quantitative data by surveying customers while they were being served by Robot EUREKA. Sitecore [2] is a customer experience management company that provides web content management and multi-channel marketing automation software, with the proven experience and reliable architecture for assessing the maturity of the customer experience. A service robotics programme for managing the HRI is designed and developed (see Figure 1,2,3). The core flow of the service robot model is to employ the information collected by the robot through various sensors, gesture actuators and advanced face recognition capability to respond and provide interaction between the robots and customers. The experiment placed the robot in an obvious position in the restaurant. In-house customers and public passing-by could easily be served by the robot and are attracted to interact with the robot. The core flow of the service robot model is to apply

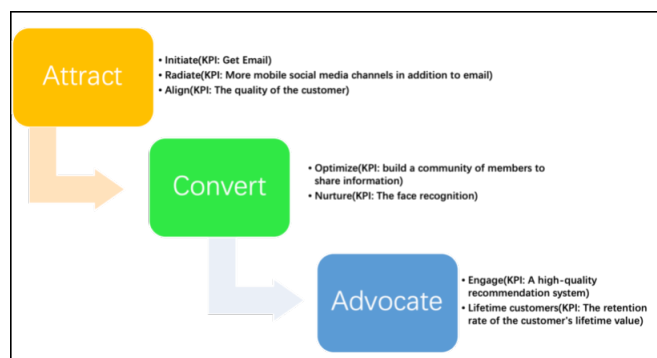


Fig. 1. Intelligent Robot for Customer eXperience Maturity (iRCXM) Model [3]

the information collected by the robot through various sensors, gesture actuators and advanced face recognition capability to

the interaction between the robot and the customers after a certain amount of processing and feedback. The experiment placed the robot in an obvious position in the restaurant. In-house customers and public passing-by could easily be served by the robot and are attracted to the store to interact with the robot.



Fig. 2. Robot and human are interacting)



Fig. 3. Robot interaction system interface

The research used questionnaire and informal interview, triangulated with researchers' observation and automated data collected by Robot Eureka to collect initial stakeholders' experiences. The research was conducted in three-fold:

- Customersrobot interactivities: demo of the humanoid service robot, customers viewed menu on the robot and took order, followed by other interactions such as customers' face registration & FAQs with Robot EUREKA; Gesture control of the robot with hands, Robot-following registered customer and duck- run game; and entertainment such as customer chose the favourite videos or songs for Robot EURKEA to play;

- Both online (iPad) and paper-based questionnaire with ethical consent;
- Automated captured data during the HRI triangulated with informal interview with customers and researchers' observation.

III. INITIAL RESULTS AND PILOT FINDINGS

A total of 23 customers' survey results were obtained from the study. The overall results confirmed that most customers showed a positive attitude towards the acceptance of the service robot. 52% of customers are not worried that robots will replace their jobs in the future, but 43% of customers indicated that they are slightly worried about this situation. All customers are able to accept robots as their colleagues and co-work together with well entertained. Among the customers, 52% are happy to work with robots, and 26% can accept but feel uncomfortable. It is worth noting that 48% of customers can accept occasional errors in robots, and only 9% of customers can accept frequent errors in robots. This means that robots need to improve the accuracy of task execution to gain greater customer acceptance. In addition, 36% of customers worry that their information will be stolen from the robot by others, indicating that data security technology and privacy protection should be considered in the application scenario of intelligent service robots. All customers were happy to provide the consent for Robot EUREKA to capture their face and names for customer registration and their faces captured by the robot showing positive and happy emotion. Researchers observed that all HRI sessions satisfied the "Attract" stage and most HRI session satisfied the "Convert" stage for the iRCXM model (see Figure 1). Only one customer, who refused to interact with Robot EUREKA, and contended the below quote: "I am afraid of your robot and I don't like to see this. He will take over my job!". After a further conversation, that customer is found as a part-time waitress and is resistant against the idea of a "Robotic Waitress". The pilot study demonstrates that users show a positive acceptance of service-type intelligent humanoid robots, but simultaneously they have concerns about information security and the efficiency and reliability of robots completing tasks.

IV. RESEARCH LIMITATIONS AND FUTURE WORK

The initial research population in the survey is limited and at small-scale. Due to the influence of environmental factor of noises, testing in a real environment will reduce the efficiency of the task completion during HRI, which has a certain impact on the results of customer acceptance surveys. Future work will focus to a larger-scale of empirical experiment with the machine learning algorithm to be integrate with iRCXM model for enhancing the "Convert" and "Advocate" stage. The changes of perception from "robot taking over my job" to "service robot enhance my work" and is "the extension of man" [4] need to be catalysed. Policy and ethical recommendations for potential job lost issues and mindset changes for upskilling are necessary.

REFERENCES

- [1] 1.The Verge. (2019). 'Japan's robot hotel lays off half the robots after they created more work for humans'. [online] Available at: <https://www.theverge.com/2019/1/15/18184198/japans-robot-hotel-lay-off-work-for-humans> [Accessed 14 Mar 2020].
- [2] itecore (2017). 'customer experience maturity model'. [online] Available at: http://mediacontent.sitecore.net/webinars/CX_Maturity_Model_NA/CX_Maturity_Model.pdf [Accessed 26 Sep 2019]
- [3] Yang J.J. & Chew E. (2019) MSc Dissertation in Service Robotics and Approved PhD Proposal. Cardiff School of Technologies, Cardiff Metropolitan University.
- [4] McLuhan, M. (2001). *Understanding Media:(Routledge Classics)*.