EDINBURGH CENTRE FOR
ROBOTICS

Innovation Ready

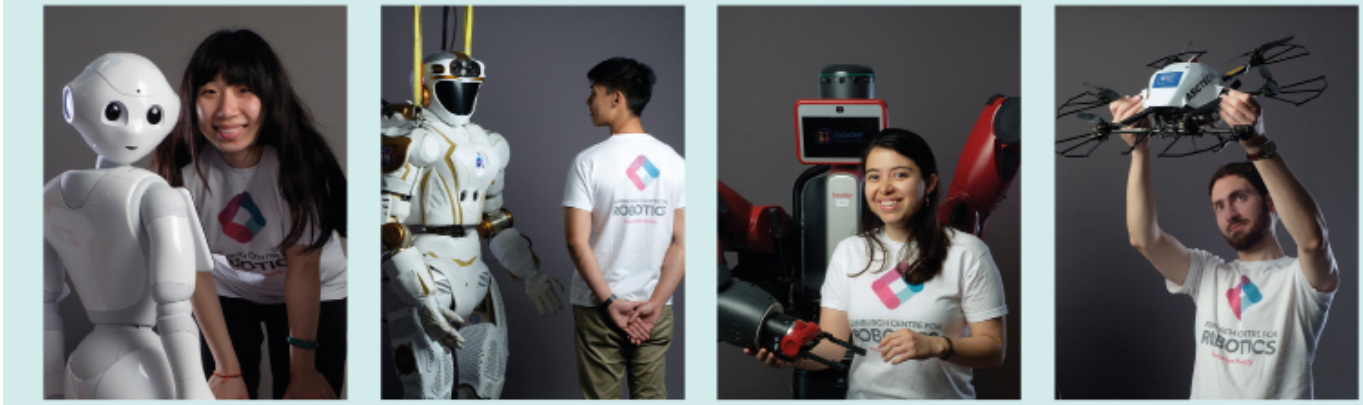

EPSRC Centre for Doctoral Training in Robotics and Autonomous Systems Newsletter

Spring 2021

Student Achievements

Congratulations to Marian Andrecki who passed his viva on the 29th of January, 2021.
His thesis is entitled "Predictive and reactive reinforcement learning from images" and explores the use of deep reinforcement learning techniques and the role of prediction in complex agent environments like games.

Abstract: Recent years have seen significant improvements in model-free reinforcement learning (RL). This is especially apparent in simulated domains: game playing at human level or learning locomotion from scratch. Unfortunately, these advances did not translate into major breakthroughs for robotics. It is often argued that the key reason for that is the inability to provide enough experience for the data-hungry RL methods. Agents acting in the real-world are expensive to run, susceptible to damage and operate in more complex environments.

This research aims to advance current RL by enabling model-based techniques. The approach is to use unsupervised learning to obtain time predictive models for high dimensional sensory percepts. With such models the agent: (a) can predict results of different actions, (b) is able to determine which actions are particularly unsafe and should not be explored, (c) has a suitable representation of an environment for later value-learning. These capabilities promise that the agent will have to explore less in order to complete the learning (because it uses data more efficiently) and when it does explore it can avoid risky experiments.

Marian's examiners were:
External - Dr Mike Preuss from Leiden University
Internal - Dr Ron Petrick

Marian is currently working in London in a fintech startup - Proportunity, where he forecasts prices of properties and helps first time buyers to get on the property ladder.

Congratulations to João Moura who passed his viva on the 8th of February. His thesis is entitled "Controlling and Learning Constrained Motions for Manipulation in Contact".

Abstract Summary: In contrast with many robotic applications where the environment poses an obstacle to the robot, in some scenarios we require the robot to be in contact with it while executing a desired motion. That interaction imposes a constraint to the robot motion and, therefore, the robot needs to be able to explicitly handle that contact interaction when controlling its motion or when learning it from demonstration. My thesis focuses on studying constrained manipulation motions. The three main topics covered are: modelling the dynamics of a rigid-body robot subject to task-based constraints; controlling a train cab front cleaning robot for wiping an uneven surface with unknown geometry; and finally learning underlying motion policies, given a set of demonstrations performed under different constraints.

Joao's examiners were:
External - Dr Matthew Howard from Kings College
Internal - Dr Matt Dunnigan

João is now working on the Orca hub with Prof. Sethu Vijayakumar in the SLMC group which is also affiliated with the Alan Turing Institute (ATI) hub at Edinburgh.
He has been collaborating with colleagues from the SLMC group and setting up a new KUKA IIWA dual arm set-up for collaborative manipulation tasks.

Congratulations to Calum Imrie who passed his viva on the 25th of Feburary 2021.
His thesis is entitled "The Robustness of Interaction Control in Robot Swarms".

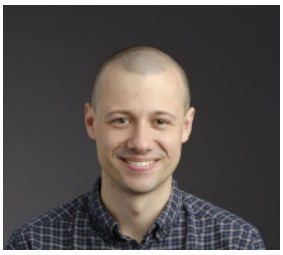

Abstract: Robotics is making significant strides in its capabilities and is being integrated into many real-world tasks. Some of these are carried out in unpredictable and dynamic environments, for example, search and rescue operations in a disaster area. Other problems, such as medical applications that require the robots to be very small, have led to significant progress in nanorobotics, in which the processor and actuating abilities are limited due to physical constraints. To solve these problems, we must look to methods that are robust to changes in the environment as well as solutions that draw maximum effect from the abilities of the robot. This may be achieved by utilising the potential of a group of robots often related to the principle of swarm intelligence. In this context, we study emergent phenomena in a group of agents each having simple low-level rules and strongly rely on local communication, resulting in increased flexibility and versatility. Individually the agents cannot perform well. As a collective, however, they are able to complete complex tasks and display implicit cognitive abilities that are critical in the applications of swarm robotics. This Thesis aims to put forward the argument for the strengths of swarm intelligence, and furthermore how it may be incorporated in applications. This Thesis first looks at a swarm having the aim to maximise energy consumption in a fixed-sized environment. The agents can sense the whole environment and augment sensory information with local communication as input for a neural network which is trained by an evolutionary strategy. Although successful for one source, as the number of sources increases frustration builds within the swarm and its effectiveness declines. It is shown that by limiting information through evolving the physical properties, the swarm can avoid this confusion. We have studied models and techniques that do not require explicit learning, and that can adapt as the situation changes. The Thesis first looks at reaction-diffusion equations being deployed onto a swarm and form Turing patterns, which are stable periodic patterns usually either honey-comb spotted patterns or stripes. We then explore how patterns can be induced via the environment. This is deployed onto a virtual Kilobot swarm to demonstrate how Turing patterns can guide robotic systems, even within the limitations of the Kilobot platform. We then present another system which has the aim to separate the agents in the swarm to allow maximum coverage of the environment. Starting with a one-dimensional dynamical system we show how this system operates given a periodic boundary condition. Using this as a foundation, we include a fixed environment, limited range of communication, and finally how the rules can be employed onto individual agents in a two-dimensional environment for maximal coverage. We compare this to the separation rule in classical swarm systems and show that our system has a smoother distribution allowing it to cover the environment quicker. This can then be taken further to circumvent obstacles or surround objects using the principle of separation.

Calum's examiners were:
External - Takashi Ikegami from the University of Tokyo
Internal - Prof Taku Komura

He is currently in the process of writing his chapters into papers to be submitted, as well as looking into postdoc positions.

Congratulations to Emmanuel Kahembwe who passed his viva on the 26th of March 2021. His thesis is entitled "Efficient Methods and Architectures for Deep Neural Network Sequence Models"

Abstract: The most recent resurgence of neural networks, termed "Deep Learning", has led to a reinvigoration of the artificial intelligence research field and all related sub-fields; from robotics and vision to natural language processing and understanding. In the last decade, this field has seen incredible breakthroughs, primarily driven by improvements to computing capability that have allowed for the training of ever larger neural network architectures.

It has also become clearer that representation learning is a crucial step in pre-training models that generalise well; offering high performance on downstream tasks of interest such as classification, regression, clustering, anomaly detection, etc.

My thesis focuses on democratising this representation learning step by lowering the cost of training and inference; particular for high-dimensional sequential neural network models. These models incorporate some aspect of temporal modelling that if properly treated, allows for more efficient training and better performance at test-time.

Our recent work published to a special issue in the journal – "Neural Networks" investigates training efficiency with respect to video modelling. We analyse the dynamics of the optimisation process and unravel previously unknown phenomena that are encountered when training neural network models on high dimensional data. We also propose some solutions to these issues resulting in the low-cost training of high performance neural networks. As a result we can match the performance of models trained with 512 TPUs (a specialised neural network training chip) using a single commodity GPU.

The other related problem is on how to get the best performance from neural network models once they are fully trained. We tackle this problem within the framework of Manifold transduction, whereby we exploit the temporal dimension of the task to gain extra performance for relatively little cost. For example during motion prediction, we can exploit the fact that once a robot has taken a particular number of steps in a particular direction, the probability that it will encounter things (e.g. states, dynamics, etc) related to other trajectories lowers. Learning this with neural networks is difficult unless such characteristics are sufficiently represented in the training dataset, We developed real-time online algorithms that can down-weight the probability of things occurring that aren't related to the current trajectory/task and up-weight everything else that is related. This enables improved test-time performance and adaptation for very minimal cost; resulting in state-of-the-art performance for many tasks.

Manny's examiners were:
External: Shakir Mohamed from DeepMind
Internal: Dr. Hakan Bilen

Manny now plans to take some time out and spend it with family. He plans to continue pursuing the line of work on manifold transduction as it may offer a way to deeply merge the current work in adaptive robot control and online machine learning.

Papers Accepted at Conferences

Hugo Sardinha, a student in the 2016 cohort had his paper titled "Towards an Adaptive Lévy Walk Using Artificial Endocrine Systems" accepted at the twelfth International Conference on Adaptive and Self-Adaptive Systems and Applications which took place in October.

This paper proposes an adaptation mechanism, that draws inspiration from the regulatory function of hormones, allowing artificial agents to modulate their behaviour between local and global searches in cluster-based foraging scenarios.
http://www.iaria.org/conferences2020/ADAPTIVE20.html

Filippos Christianos**,** who is in the 2nd year of his PhD, had a paper accepted at Neural Information Processing Systems (NeurIPS) 2020 which took place in December. The paper is titled "Shared Experience Actor-Critic for Multi-Agent Reinforcement Learning" and proposes a new algorithm for efficient multi-agent learning in sparse-reward environments.

This paper shows that our algorithm learns faster and converges to higher returns than several baselines and other state-of-the-art algorithms, in a series of tested environments. The paper was co-authored with Lukas Schäfer and Stefano Albrecht as part of the Autonomous Agents Research Group.
The full paper text can be found here
https://nips.cc/

Alessandro Suglia, a student from the 2017 cohort had his long paper presented at the International Conference on Computational Linguistics. His paper is titled "Imagining Grounded Conceptual Representations from Perceptual Information in Situated Guessing Games".
You can read the paper here .
https://coling2020.org/

Alessandro also will be presenting a paper at EACL 2021 titled "An Empirical Study on the Generalization Power of Neural Representations Learned via Visual Guessing Games" which is part of a collaboration between Heriot-Watt University, Carnegie Mellon University and University of California, Los Angeles.

In this paper he studies the power of visual guessing games as a holistic procedure for grounded language learning tasks. Guessing games are a prototypical instance of the "learning by interacting" paradigm. This work investigates how well an artificial agent can benefit from playing guessing games when later asked to perform on novel NLP downstream tasks such as Visual Question Answering (VQA). He proposes two ways to exploit playing guessing games: 1) a supervised learning scenario in which the agent learns to mimic successful guessing games and 2) a novel way for an agent to play by itself, called Self-play via Iterated Experience Learning (SPIEL).

This work formalises a self-play learning framework that allows artificial agents to learn more effective representations that are useful for other multi-modal downstream tasks. The learning procedure is generic and versatile because it does not require any human supervision and can be applied to any vision-language downstream tasks.
EACL 2021

Papers accepted at ICRA 2021 Conference

These SLMC papers co-authored by current and former RAS students have been accepted at the International Conference on Robotics and Automation (ICRA 2021) to be held in Xi'an, China.

Ran Long, **Christian Rauch**, Tianwei Zhang, Vladimir Ivan and **Sethu Vijayakumar,** RigidFusion: Robot Localisation and Mapping in Environments with Large Dynamic Rigid Objects, IEEE Robotics and Automation Letters (RAL), 2021
Presented at: IEEE International Conference on Robotics and Automation (ICRA 2021), Xian, China (2021)
RigidFusion is a state-of-the-art dense SLAM method that is robust to large dynamic occlusion (over 65%) in the scene, without requiring prior shape or appearance of the dynamic objects. It also contributes a pipeline to simultaneously segment, track and reconstruct the static background and one dynamic rigid body from RGB-D sequences. Importantly, we publish the dataset with the camera and object ground truth trajectories for benchmarking future work in the area of SLAM in dynamic environments.

Henrique Ferrolho, Vladimir Ivan, **Wolfgang Xaver Merkt**, Ioannis Havoutis and **Sethu Vijayakumar**, Inverse Dynamics vs. Forward Dynamics in Direct Transcription Formulations for Trajectory Optimization, Proc. IEEE International Conference on Robotics and Automation (ICRA 2021), Xian, China, 2021.
Direct transcription is a powerful technique that uses numerical optimisation to solve motion planning problems. Such numerical formulations use mathematical constraints to enforce motion requirements; in robotics, those constraints are used for, e.g., body placement, contact positions, system dynamics. We discuss two possible approaches to enforce nonlinear whole-body dynamics of robots in direct transcription: forward dynamics vs. inverse dynamics. Results show that using inverse dynamics is faster, requires less iterations, and is more robust to coarse problem discretisations.

Traiko Dinev, Wolfgang Xaver Merkt, Vladimir Ivan, Ioannis Havoutis and **Sethu Vijayakumar,** Sparsity-Inducing Optimal Control via Differential Dynamic Programming , Proc. IEEE International Conference on Robotics and Automation (ICRA 2021), Xian, China, 2021.
Our paper describes how to use sparse controls in dynamic motion planning. We apply sparsity-inducing costs to plan satellite manoeuvres, where thrusters using liquid propellants can only be switched on and off and can not provide variable thrust. We also apply sparsity in controls to a humanoid reaching task, which allows us to select the required number of joints for this lower-dimensional motion. We analyze the properties of a family of soft sparse costs and give insight into how to tune their free parameters.

Carlo Tiseo, Vladimir Ivan, **Wolfgang Xaver Merkt**, Ioannis Havoutis, **Michael Mistry and Sethu Vijayakumar,** A Passive Navigation Planning Algorithm for Collision-free Control of Mobile Robots, Proc. IEEE International Conference on Robotics and Automation (ICRA 2021), Xian, China, 2021.
A passive planning algorithm capable of autonomous obstacle avoidance in a domain with small concavity is presented. This novel method uses a passive controller that enables the navigation of complex dynamic maps without relying on numerical optimisation. Simulation and experimental results show that the technique can generate smooth, stable trajectories in drones and wheeled robots. The small computational cost enables scalability to swarm applications where the agents' movements are synchronised by issuing coordinated targets.

Further Information
ICRA 2021

Student Placements

Alessandro Suglia has started a research internship with Amazon Alexa AI - Natural Understanding team working with Dr. Jesse Thomason and Prof. Gaurav Sukhatme. On the 5th July, he will be starting another research internship with Facebook AI Seattle working with Dr. Shane Moon on grounded video understanding.

Academic Achievements

Ignazio Maria Viola awarded a European Research Council Consolidator Grant

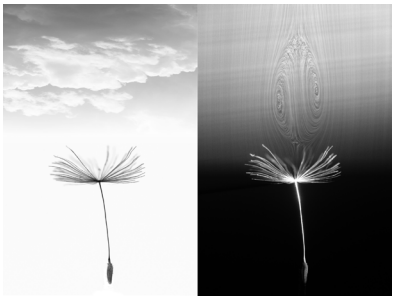

Ignazio Maria Viola, an academic at the School of Engineering, University of Edinburgh has been awarded a £2 Million grant over 5 years. The project, which will start in June 2021 will study the fluid dynamics of dandelion-inspired drones that are transported by the wind.

What is the Dandidrone project?
Imagine we could fly small drones that remain airborne for days: monitoring the environment, searching for pollutants, tracking airborne coronaviruses.
Unfortunately, this is not yet possible, because insect-scale drones can remain airborne for only a few minutes before the power runs out.
But some natural flyers, by contrast, remain airborne for days without using any power. One of the most extraordinary examples is the dandelion seed: it takes off from just 30 cm away from the ground, and travels for hundreds of kilometres powered only by the wind.
In a recent study led by Viola and Naomi Nakayama (School of Biology), of which some findings were published in the leading journal Nature [1], the researchers revealed that the flight of the dandelion is enhanced by a new flow feature that had never been observed before: the separated vortex ring.
In this ERC-funded five-year project, Viola will next investigate how the amazing flight capabilities of the dandelion can be replicated by manmade flyers. The project seeks to understand and establish proof of principle of a completely new fluid mechanics mechanism that might enable small flyers to passively hover in turbulent wind.
The funding will enable the recruitment of three postdoctoral research associates, two PhD students, the construction of a one-of-its-kind wind tunnel to study small flyers in gusts, and a period of study at the California Institute of Technology (Caltech) in Pasadena, USA.

About ERC Consolidator Grants
The ERC Consolidator Grants are awarded to outstanding researchers with seven to twelve years of postdoctoral experience, and a scientific track record showing great promise.
The ERC received 2,453 research proposals in this round, and approximately 12% were funded, across 24 countries in Europe, covering a wide range of academic disciplines.

References
[1] Cummins, C., Seale, M., Macente, A., Certini, D., Mastropaolo, E., Viola, I. M. & Nakayama, N. A separated vortex ring underlies the flight of the dandelion. Nature 562, 414–418 (2018).

Links
Ignazio's research group: https://voilab.eng.ed.ac.uk

New publication from Dr Francesco Giorgio-Serchi

Dr Francesco Giorgio-Serchi, one of our academics based at the University of Edinburgh, has a new publication out on Science Robotics. You can read it here!

Abstract
Elasticity has been linked to the remarkable propulsive efficiency of pulse-jet animals such as the squid and jellyfish, but reports that quantify the underlying dynamics or demonstrate its application in robotic systems are rare. This work identifies the pulse-jet propulsion mode used by these animals as a coupled mass-spring-mass oscillator, enabling the design of a flexible self-propelled robot. We use this system to experimentally demonstrate that resonance greatly benefits pulse-jet swimming speed and efficiency, and the robot's optimal cost of transport is found to match that of the most efficient biological swimmers in nature, such as the jellyfish Aurelia aurita. The robot also exhibits a preferred Strouhal number for efficient swimming, thereby bridging the gap between pulse-jet propulsion and established findings in efficient fish swimming. Extensions of the current robotic framework to larger amplitude oscillations could combine resonance effects with optimal vortex formation to further increase propulsive performance and potentially outperform biological swimmers altogether.

Professor Subramanian Ramamoorthy speaks at the ESSS 12th National Undergraduate Surgery & Trauma Conference

On March 13 Professor Ramamoorthy gave a talk entitled "Robotics in Surgery: What's Under the Hood?"

This talk was aimed at introducing key ideas about robotics and AI, and how they are rapidly changing the landscape of surgical systems. Drawing on work done within the Turing Institute sponsored project on Safe AI for Surgical Assistance, this talk covered topics including dexterous manipulation of soft tissue and models learned from detailed observation of human performance in surgical tasks, machine learning methods for interpreting imaging and other sensory information, and approaches to closing the loop towards tasks such as autonomous excision. This talk was aimed at the audience of medical students specialising in surgery and trauma.

Professor Barbara Webb awarded Fellowship

We would like to congratulate Barbara Webb, an academic at the University of Edinburgh who has been awarded a five year EPSRC Established Career Fellowship to start from September 2021. The topic is "An insect-inspired approach to robotic grasping" and she will examine how ants are able to reliably grasp arbitrary unknown objects in clutter, and how these skills can be transfered to robot applications such as e-commerce and environmental clean-up.

ECR & National Robotarium in the News!

The National Robotarium was featured on BBC newsround to announce that it will open it's doors in Spring of 2022. You can read about it here.

The National Robotiarum's development of robotics for the offshore renewable sector is looked at in this article from Heriot-Watt News.

Prof Thusha Rajendran, an academic at Heriot-Watt University was interviewed at the Scottish Sun about the National Robotarium, you can read the article here.

Dr Alex Li, an academic at the University of Edinburgh was featured in an article on BBC around his work with robotic dogs. You can read the article here.

Prof Sethu Vijayakumar, co-director of RAS I at the University of Edinburgh, was interviewed by Ruth Davidson for the Podcast An Inconvenient Ruth. You can listen here.

Events

Robotics & Care Mash-Up

The National Robotarium and Product Forge are hosting a five day online "mashup" event bringing together health & care professionals, academics and technology providers to prototype new solutions to support the care sector from the 10th - 14th May.

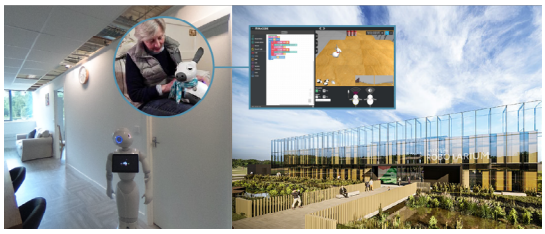

How can you get involved?

Join a team and develop a new product prototype! Participants will be provided with coaching and mentoring by Product Forge and the Open Ambient Assisted Living project, and will be able to access real robotics hardware via a set of "remote hands" at the National Robotarium laboratory at Heriot-Watt. Participants on the CPD track can obtain a digital certificate on completion validated by the CPD Standards Office.

More information is available here.

Newsletters are available on our website and posted on twitter @EdinRobotics. In addition to this newsletter, information about our most recent achievements and activities can also be found in the news section on our website. Our annual reports are also available for review and download here. The centre is very active on social media and we would encourage you to follow us on Twitter @EdinRobotics to keep up-to-date with our activities.

https://www.edinburgh-robotics.org/

Engineering and Physical Sciences Research Council

HERIOT WATT UNIVERSITY

THE UNIVERSITY of EDINBURGH