PhD projects in the Department of Informatics, AY 25-26 — Robotics

The PhD projects listed below will be considered for 2025/26 studentships available in the Department of Informatics to start on 1 October 2025 or later during the 2025/26 academic year.

Please note that this list is not exhaustive and potential applicants can alternatively identify and contact appropriate supervisors outlining their background and research interests or proposing their own project ideas.

Each project is designated for a single student, meaning it can only be assigned to one successful applicant. Some projects come with allocated studentships, while others are eligible for "unallocated" studentships. Applicants who apply for projects with allocated studentships and are selected will be offered a full studentship. In the project list, these are marked as "studentship allocated." Applicants chosen for other projects will compete for the unallocated studentships.

We welcome applications from students who have secured, or are applying for, or plan to apply for other funding (within other studentships internal to the university or external schemes) and from self-funded students. See also this list of funding opportunities available at King's for post-graduate research in Computer Science.

PhD projects

- Allowing autonomous robots to continually learn, generalize, and improve from their experiences (studentship allocated)
- Sustainable Delivery Logistics with Drones and Cargo Bikes
- Goal-based explanations for autonomous systems and robots
- Multi-agent Cooperation with RL and LLMs
- Adaptation and effective communication in collaborative physically Assistive Tasks
- Explaining robotic planning decision points along execution
- Enhancing Safety in Robotics by Tackling Blind-Spots and Bias in AI Models
- Planning and Reinforcement Learning for versatile autonomous robots

Allowing autonomous robots to continually learn, generalize, and improve from their experiences

Supervisor: Dr. Khen Elimelech

Areas: Robotics, Artificial Intelligence (symbolic AI, logic, etc.), Machine learning / Deep learning

Project Description

To perform autonomously tasks such as object rearrangement, assembly, manipulation, and navigation, robots must be able to plan their actions over long horizons. Such planning is usually computationally challenging to perform in real time, especially considering complex robots and task specifications, or large and uncertain planning domains, with many irrelevant objects and distractions. One intuitive approach to support autonomous robots in this challenge is by allowing them to learn to continually improve their planning capabilities over time, based on their experience. This general approach should enable us to build long-lived, multi-purpose robots with human-like versatility and common sense, rather than highly specialized machines.

Unfortunately, despite recent advancement in Machine Learning and "Learning from Demonstrations," existing learning approaches are not suitable for this objective, as these require numerous annotated demonstrations, rendering them unsuitable for online, autonomous learning.

To this end, our recent work introduced a novel algorithmic framework for automatic learning of "planning strategies" by abstracting successful planning experiences. This framework allows a robot to automatically and continually make generalizable conclusions from individual experiences, which can later be adapted for and reused in new contexts, to accelerate the solution of new planning problems—just like humans do, but without human intervention!

Initial results demonstrated the potential of this approach to significantly impact the field of AI-enabled robotics. To achieve that, this project seeks to extend this initial effort in various directions, including: application and adaptation to new platforms, planning domains, and task-types; application to multi-robot and human-robot collaborative systems; integration with (statistical) Machine Learning and Computer Vision techniques, Control, knowledge graphs and other components in the autonomy stack; improving utility and computational tractability through algorithmic development; and improving trustworthiness through formal analysis.

The work on this project is diverse and contains theoretical, computational. and experiential aspects. Students are expected to conduct research, publish papers, develop and release open-source code, and work with physical robots. You will have access to state-of-the-art hardware and resources, and excellent mentorship. Potentially, successful students will have access to collaboration and internship opportunities with industry leaders, such as NASA Robotics, Amazon Robotics and Bosche.

While prior research experience in robotics is recommended, it is not mandatory. Excellent candidates with background in robotics, AI, computer science, algorithms, applied mathematics, engineering, or other relevant background are welcome to apply.

References

[1] Accelerating Long-Horizon Planning with Affordance-Directed Dynamic Grounding of Abstract Strategies, by Khen Elimelech, Zachary Kingston, Wil Thomason, Moshe Y. Vardi, and Lydia E. Kavraki, in IEEE International Conference on Robotics and Automation (ICRA), May 2024.
[2] Extracting generalizable skills from a single plan execution using abstraction-critical state detection, by Khen Elimelech, Lydia E. Kavraki, and Moshe Y. Vardi, in IEEE International Conference on Robotics and Automation (ICRA), May 2023.
[3] Principles of Robot Motion: Theory, Algorithms, and Implementations, by Howie Choset, Kevin M. Lynch, Seth Hutchinson, George A. Kantor, Wolfram Burgard, Lydia E. Kavraki and Sebastian Thrun, The MIT Press, 2005.

Sustainable Delivery Logistics with Drones and Cargo Bikes

Supervisor: Dimitrios Letsios

Areas: Robotics, Foundations of computing (algorithms, computational complexity), Artificial Intelligence (symbolic AI, logic, etc.)

Project Description

Effectively solving inventory optimisation, scheduling, and vehicle routing problems is essential to meet customer demands and manage fleet, human worker, and storage costs for companies in the logistics sector, e.g. courier and delivery services. Drones and cargo bikes offer an alternative to cost-effective and sustainable last-mile deliveries in urban environments. The goal of this project is to develop data science and optimisation methodologies (e.g. metaheuristics and integer programming) for generating solutions to challenging logistics optimisation problems with drones and cargo bikes that attain optimal trade-offs between solution quality and computational effort, while enhancing their performance with simulations and machine/deep learning. Proposed approaches will be evaluated using recent data and benchmark case studies from business partners and in the literature. The project will include a thorough study on effective hyper-parameter tuning (e.g. initialisation components, solution space structure, search operators) of optimisation methodologies for computing the best possible solution within given time frames. Further, it will investigate innovative ways of enhancing performance with dynamic parameter adaptations and machine/deep learning, based on information (e.g. rewards, data properties) collected during the solution process. The project will also assess and strengthen the robustness of proposed approaches via simulations.

Goal-based explanations for autonomous systems and robots

Supervisor: Gerard Canal

Areas: Artificial Intelligence (symbolic AI, logic, etc.), Robotics

Project Description

Autonomous systems such as robots may become another appliance found in our homes and workplaces. In order to have such systems helping humans to perform their tasks, they must be as autonomous as possible, to prevent becoming a nuisance instead of an aid. Autonomy will require the systems or robots to set up their own agenda (in line with the tasks they are meant to do), defining the next goals to achieve and discarding those that can't be completed. However, this may create misunderstandings with the users around the system, who may expect something different from the robot. Therefore, it is important that these autonomous systems are able to explain why they achieved one task and not another, or why some new (unexpected) task was achieved that was not scheduled. Other sources of misunderstandings may come from action failures and replanning, where the robot finds a new plan to complete an ongoing task. In this case, the new plan may be different to the original one, thus changing the behaviour that the robot was performing. This project will explore how to generate goalbased explanations for robots in assistive/home-based scenarios, extracted from goal-reasoning techniques. It will also look at plan repair to enforce cohesion after a replanning to ideally increase the trust and understanding of the users about the system. Those explanations should also contemplate unforeseen circumstances, therefore explaining things based on "excuses" that the robot may give to the user. Finally, we will investigate how to obtain and provide those explanations at execution time, so explaining on the go. The methods developed shall be integrated into a robotic system, in an assistive/service robot scenario.

References

[1] Canal, G., Borgo, R., Coles, A., Drake, A., Huynh, T. D., Keller, P., Krivic, S., Luff, P., Mahesar, Q-A., Moreau, L., Parsons, S., Patel, M., & Sklar, E. (2020). Building Trust in Human-Machine Partnerships. Computer Law & Security Review, 39.
[2] Hawes, N., Burbridge, C., Jovan, F., Kunze, L., Lacerda, B., Mudrova, L., ... & Hanheide, M. (2017). The strands project: Long-term autonomy in everyday environments. IEEE Robotics & Automation Magazine, 24(3), 146-156.
[3] Aha, D. W. (2018). Goal reasoning: Foundations, emerging applications, and prospects. AI Magazine, 39(2), 3-24.
[4] Bercher, Pascal, et al. "Plan, repair, execute, explain—how planning helps to assemble your home theater." Proceedings of the International Conference on Automated Planning and Scheduling. Vol. 24. No. 1. 2014.
[5] Chakraborti, Tathagata, Sarath Sreedharan, and Subbarao Kambhampati. "The emerging landscape of explainable AI planning and decision making. IJCAI 2020.
[6] Gobelbecker, M., Keller, T., Eyerich, P., Brenner, M., & Nebel, B. (2010, April). Coming up with good excuses: What to do when no plan can be found. In Proceedings of the International Conference on Automated Planning and Scheduling (Vol. 20, No. 1).

Multi-agent Cooperation with RL and LLMs

Supervisor: Yali Du

Areas: Artificial Intelligence (symbolic AI, logic, etc.), Machine learning / Deep learning, Natural Language Processing, Robotics

Project Description

Multi-agent systems (MAS) have become increasingly relevant in fields such as robotics, finance, and autonomous systems. However, achieving effective cooperation among multiple agents remains challenging, especially in dynamic and uncertain environments. RL has been a powerful method for training agents, but traditional approaches often struggle with scalability and communication bottlenecks. Meanwhile, LLMs have demonstrated remarkable capabilities in language understanding and generation, which can be leveraged to facilitate communication and strategy development among agents. This study aims to explore how reinforcement learning (RL) can be combined with large language models (LLMs) to improve multi-agent cooperation in complex environments. The goal is to enhance communication, decision-making, and coordination between agents, enabling them to solve tasks that require a high level of collaboration and safety. This project explores the questions of 1) How can LLMs be integrated into multi-agent systems to enhance cooperation and communication among agents trained using RL? 2) What are the optimal communication protocols that maximize the synergy between LLMs and RL in multi-agent scenarios? 3) How can this combination be scaled to large numbers of agents while maintaining efficiency and performance? Dr Du's early attempts explored how to leverage LLMs for communication, and incorporated human instructions to ensure safe and cooperative control, with examples including the game of Werewolf, football, and safe robot control. This research will contribute to the field of multi-agent systems by developing new techniques for improved cooperation using cutting-edge LLMs. The findings could be applicable in various industries, including autonomous vehicles, robotics, and distributed AI systems, where multi-agent cooperation is critical for success.

References

[1] Safe Multi-agent Reinforcement Learning with Natural Language Constraints. Ziyan Wang, Meng Fang, Tristan Tomilin, Fei Fang, Yali Du. Arxiv 2024.
[2] Understanding, Rehearsing, and Introspecting: Learn a Policy from Textual Tutorial Books in Football Games. Xiong-Hui Chen, Ziyan Wang, Yali Du, Meng Fang, Shengyi Jiang, Yang Yu, Jun Wang. NeurIPS 2024 Oral.
[3] Learning to Discuss Strategically: A Case Study on One Night Ultimate Werewolf. Xuanfa Jin, Ziyan Wang, Yali Du, Meng Fang, Haifeng Zhang, Jun Wang. NeurIPS 2024.

Adaptation and effective communication in collaborative physically Assistive Tasks

Supervisor: Gerard Canal

Areas: Artificial Intelligence (symbolic AI, logic, etc.), Robotics

Project Description

Physical robotic Assistance can often be modelled as a collaborative task in which the goal of both the user and the robot is to complete an assistive task together. However, assistive settings have a lot of particularities that differentiate them from traditional Human-Robot Collaboration tasks. For it to be effective, the assistance should be seamless, natural, and without a required effort on the user's side. This means that these robots must be able to communicate with the user in a very natural and intuitive way, but also in an adaptive manner. In this project, we will investigate the development of techniques for the online adaptation of the robot to the human, as well as anticipation of user needs, and seamless communication in the context of assistive tasks such as robotic feeding and dressing.

References

[1] Canal, G., Alenya, G., & Torras, C. (2019). Adapting robot task planning to user preferences: an assistive shoe dressing example. Autonomous Robots, 43(6), 1343-1356.

[2] Canal, G., Alenya, G., & Torras, C. (2016). Personalization framework for adaptive robotic feeding assistance. In Social Robotics: 8th International Conference, ICSR 2016, Kansas City, MO, USA, November 1-3, 2016 Proceedings 8 (pp. 22-31). Springer International Publishing.

[3] Bhattacharjee, T., Lee, G., Song, H., & Srinivasa, S. S. (2019). Towards robotic feeding: Role of haptics in fork-based food manipulation. IEEE Robotics and Automation Letters, 4(2), 1485-1492.

[4] Bhattacharjee, T., Gordon, E. K., Scalise, R., Cabrera, M. E., Caspi, A., Cakmak, M., & Srinivasa, S. S. (2020, March). Is more autonomy always better? exploring preferences of users with mobility impairments in robot-assisted feeding. In Proceedings of the 2020 ACM/IEEE international conference on human-robot interaction (pp. 181-190).

[5] Ondras, J., Anwar, A., Wu, T., Bu, F., Jung, M., Ortiz, J. J., & Bhattacharjee, T. (2022, August). Human-robot commensality: Bite timing prediction for robot-assisted feeding in groups. In 6th Annual Conference on Robot Learning.

Explaining robotic planning decision points along execution

Supervisor: Gerard Canal

Areas: Artificial Intelligence (symbolic AI, logic, etc.), Robotics

Project Description

Explanation of robotic behaviours has been proved to be very important to improve the understanding of the users of such robots, which improves their trust in the robotic system. However, explanations in robotics are tricky as they need to be given at the correct moment and based on what happened in the execution. In robotic-based planning, an interesting explanation is that of decision points, where the robot could have taken a different action with a different outcome. This project focuses on the explanation of such decision points at execution time, integrating information on current and past events that may help explain the decision to a user. For this, we will look into explainability in the space of plans where, knowing the committed plan and what has happened in the execution, we compare it with the other alternatives that the robot had at a certain decision point. This will evolve towards generating explanations along the execution of plans, as well as determining when some decisions may not be obvious to the user, thus warranting explanations.

References

[1] Canal, G., Torras, C., & Alenya, G. (2023). Generating predicate suggestions based on the space of plans: an example of planning with preferences. User modeling and user-adapted interaction, 33(2), 333-357.
[2] Eifler, R., Cashmore, M., Hoffmann, J., Magazzeni, D., & Steinmetz, M. (2020, April). A new approach to plan-space explanation: Analyzing plan-property dependencies in oversubscription planning. In Proceedings of the AAAI Conference on Artificial Intelligence (Vol. 34, No. 06, pp. 9818-9826).
[3] Wachowiak, L., Celiktutan, O., Coles, A., & Canal, G. (2023, June). A Survey of Evaluation Methods and Metrics for Explanations in Human—Robot Interaction (HRI). In ICRA2023 Workshop on Explainable Robotics.
[4] Wachowiak, L., Tisnikar, P., Canal, G., Coles, A., Leonetti, M., & Celiktutan, O. (2022, August). Analysing eye gaze patterns during confusion and errors in human—agent collaborations. In 2022 31st IEEE International Conference on Robot and Human Interactive Communication (RO-MAN) (pp. 224-229). IEEE.

Enhancing Safety in Robotics by Tackling Blind-Spots and Bias in AI Models

Supervisor: Gerard Canal

Areas: Artificial Intelligence (symbolic AI, logic, etc.), Robotics

Project Description

The current revolution of artificial intelligence (AI) is becoming more prominent and its potential is still to be unleashed. In the context of robotics, AI can provide support to multiple scenarios, among them, industry, education and healthcare. It is important to know how these systems can work on these contexts but it is imperative that they can treat people respectfully and equally. There are significant efforts in this direction that focus on the context of fairness and explainability. Several AI models normally employed in robotics, such as computer vision models, have been tested to discover that they still contain blind-spots in their detection capabilities, several of them affecting specifically protected groups, such as children or citizens with disabilities. Even if the models are becoming more explainable these days, the consequences of these blind-spots in their explanations and especially the actions of the robots in the real world still requires deeper studies. This is particularly important due to the safety issues that this may impose, which is specially critical in assistive scenarios where a robot helps a user from a vulnerable group perform activities of daily living. This thesis aims to address these issues by: 1) Identifying use cases where the sensitiveness of fairness issues might have a strong repercussion in the behaviour of the robots, with a special emphasis on when this results in unsafe situations for the user recipient of the assistance. This will consist of collecting different examples for the literature that the student can have access and implementing them with the robots that we have available in the department such as the PAL Robotics' TIAGo or models of smart cars. It will also potentially employ digital twins to create a simulation environment for more complex robots. 2) Create strategies to identify blind-spots. Based on the previous work of adversarial machine learning where blind-spots are normally identified as misclassifications or mis-actions that a robot will execute, this part of the thesis will work on identifying and designing adversarial scenarios that will make the system misbehave. The scenario design will consider potential sensory alterations that the robot will face, especially connected with environment conditions. With this information, the thesis will aim to explain the scenario and the specific conditions that led to the misclassification. This will support redesigning the learning process and will serve for standardising benchmark testing conditions. 3) Based on the previous adversarial scenarios and the specific transformations that led the system to make erroneous decisions, this last part will provide explanations about the system limitations, with an aim to enhance the safety of the system. It will focus on: 1) generalising from the adversarial scenarios to create explanations and 2) inverse the pipeline and create adversarial conditions from specific explanations. These adversarial conditions will be focused on fairness. Besides this last part will put a strong effort on evaluating explanatory systems for robotics under adversarial conditions.

References

Canal, G., Torras, C., & Alenya, G. (2021). Are preferences useful for better assistance? a physically assistive robotics user study. ACM Transactions on Human-Robot Interaction (THRI), 10(4), 1-19.
Vila Abad, M., Canal, G., & Alenya, G. (2018). Towards safety in physically assistive robots: eating assistance. In: Proceedings of the 2018 IROS Workshop on Robots for Assisted Living
Wachowiak, L., Celiktutan, O., Coles, A., & Canal, G. (2023, June). A Survey of Evaluation Methods and Metrics for Explanations in Human—Robot Interaction (HRI). In ICRA2023 Workshop on Explainable Robotics.
Menendez, H. D., Bhattacharya, S., Clark, D., & Barr, E. T. (2019). The arms race: Adversarial search defeats entropy used to detect malware. Expert Systems with Applications, 118, 246-260.
Calleja, A., Martin, A., Menendez, H. D., Tapiador, J., & Clark, D. (2018). Picking on the family: Disrupting android malware triage by forcing misclassification. Expert Systems with Applications, 95, 113-126.
Menendez, H. D. (2022, October). Measuring Machine Learning Robustness in front of Static and Dynamic Adversaries. In 2022 IEEE 34th International Conference on Tools with Artificial Intelligence (ICTAI) (pp. 174-181). IEEE.

Planning and Reinforcement Learning for versatile autonomous robots

Supervisor: Matteo Leonetti

Areas: Artificial Intelligence (symbolic AI, logic, etc.), Robotics

Project Description

Model-based reinforcement learning has been lagging behind initial and exciting model-free results in deep reinforcement learning. In this project we will consider the problem of an autonomous robot required to carry out different tasks in its environment, frequently switching between goals. The research will focus on model learning and effective use of models to drive exploration, hierarchical models, and multi-task heuristics. Examples of previous work in this direction are provided in the reference section

References

1. Probabilistic Inference of Human Capabilities from Passive Observations. Peter Tisnikar, Gerard Canal, Matteo Leonetti. Proc. of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2024.
2. Online human capability estimation through reinforcement learning and interaction. Chengke Sun, Anthony G Cohn, Matteo Leonetti. Proc. of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2023.
3. Reducing the Planning Horizon Through Reinforcement Learning. Logan Dunbar, Benjamin Rosman, Anthony G Cohn, Matteo Leonetti. Proc. of Joint European Conference on Machine Learning and Knowledge Discovery in Databases (ECML-PKDD), 2022.
4. Meta-Reinforcement Learning for Heuristic Planning. Ricardo Luna Gutierrez, Matteo Leonetti. Proc. of International Conference on Automated Planning and Scheduling (ICAPS), 2021.
5. A synthesis of automated planning and reinforcement learning for efficient, robust decision-making. Matteo Leonetti, Luca Iocchi, Peter Stone. Artificial Intelligence, 2016.