

Systems Projects 2026-27

Contents

Game Theory and Learning in Multi-Agent Systems.....	2
Game theory, learning and mechanism design in cryptoeconomic systems	3
Understanding the Complexity of Negotiations	4
Exploring Interactive Multi-Dimensional Approaches to Patient Communication in Oral Health Education	5
Data Science and Optimization for Sustainable Last-Mile Logistics	6
Goal-based explanations for autonomous systems and robots	7
Software sustainability analysis and improvement.....	8

Game Theory and Learning in Multi-Agent Systems

Supervisor: Stefanos Leonardos

Areas: Artificial Intelligence, Machine learning / Deep learning, Game theory, Multi-agent systems

(Back to [Scholarship Not Allocated](#))

Project Description

This project is designed for students interested in research at the intersection of game theory, learning dynamics, and multi-agent systems, with applications in economics, machine learning, and artificial intelligence. The aim is to study how complex patterns and behaviors emerge when many agents interact and adapt over time. We will explore phase transitions in strategic interactions, investigate the role of chaos and dynamical systems, and develop or analyze novel learning algorithms. The project will combine tools from mathematics, game theory, and AI to understand coordination, competition, and long-term dynamics in real-world systems. Students will have the opportunity to contribute to both theoretical advances and practical applications, helping to shape the future of intelligent learning systems.

References

1. I. Sakos, S. Leonardos, S. A. Stavroulakis, W. Overman, I. Panageas, G. Piliouras. Beating Price of Anarchy and Gradient Descent without Regret in Potential Games, 12th International Conference on Learning Representations (2024).
2. S. Roesch, S. Leonardos & Y. Du. Selfishness Level Induces Cooperation in Sequential Social Dilemmas, 23rd Conference on Autonomous Agents and Multiagent Systems (2024).
3. Leonardos, S., Sakos, J., Courcoubetis, C. and Piliouras, G. (2023). Catastrophe by Design in Population Games: A Mechanism to Destabilize Inefficient Locked-in Technologies. ACM Trans. Econ. Comput. 11, 1–2, Article 1 (June 2023), 36 pages. doi:10.1145/3583782
4. Leonardos, S., and Piliouras, G. (2022). Exploration-exploitation in multi-agent learning: Catastrophe theory meets game theory, Artificial Intelligence, Volume 304, 103653, doi:10.1016/j.artint.2021.103653.
5. Leonardos, S., Piliouras, G., and Spendlove, K. (2021). Exploration-Exploitation in Multi-Agent Competition: Convergence with Bounded Rationality, in Advances in Neural Information Processing Systems, volume 34, pp. 26318--26331, Curran Associates, Inc., https://proceedings.neurips.cc/paper_files/paper/2021/file/dd1970fb03877a235d530476eb727dabPaper.pdf.
6. Leonardos, S., Overman, W., Panageas I., and Piliouras, G. (2022). Global Convergence of Multi-Agent Policy Gradient in Markov Potential Games, in International Conference on Learning Representations (ICLR 2022), <https://openreview.net/forum?id=gfwON7rAm4>.

Game theory, learning and mechanism design in cryptoeconomic systems

Supervisor: Stefanos Leonardos

Areas: Machine learning / Deep learning, Cryptoeconomics, Game theory

(Back to [Scholarship Not Allocated](#))

Project Description

This project is for students interested in exploring the fast-growing field of cryptoeconomics through the lens of game theory. The focus is on modeling and analyzing blockchain-based systems to understand the incentives that drive participants and the mechanisms that shape outcomes. Topics include transaction fee mechanisms (TFMs), miner extractable value (MEV), automated market makers (AMMs), transaction censorship, and many more. Students will investigate how cryptoeconomic mechanisms evolve, how strategic behavior impacts blockchain performance, and how new mechanisms can be designed for greater efficiency and fairness. The project combines tools from game theory, economics, computer science, and dynamical systems, with opportunities to build models, run simulations, and connect theory with real-world blockchain applications.

References

1. Buterin, V, Reijersbergen, D, Leonardos, S, Piliouras, G. Incentives in Ethereum's hybrid Casper protocol. Int J Network Mgmt. 2020; 30:e2098. <https://doi.org/10.1002/nem.2098>
2. Leonardos, S., Reijersbergen, D., Monnot, B., and Piliouras, G., "Optimality Despite Chaos in Fee Markets", arXiv e-prints, 2022. doi:10.48550/arXiv.2212.07175, (2025).
3. Performative Market Making, C Kleitsikas, S Leonardos, C Ventre, arXiv preprint arXiv:2508.04344.
4. MEV Sharing with Dynamic Extraction Rates, P Braga, G Chionas, P Krysta, S Leonardos, G Piliouras, C Ventre, Proceedings of the Workshop on Decentralized Finance and Security, 1-10, (2024).
5. W. Wu, T. Thiery, S. Leonardos, C. Ventre. Strategic Bidding Wars in On-chain Auctions. IEEE ICBC 2024, <https://arxiv.org/abs/2312.14510>.
6. Leonardos, S., Monnot, B., Reijersbergen, D., Skoulakis, E., and Piliouras, G. (2021). Dynamical analysis of the EIP-1559 Ethereum fee market. In Proceedings of the 3rd ACM Conference on Advances in Financial Technologies (AFT '21). Association for Computing Machinery, New York, NY, USA, 114—126. <https://doi.org/10.1145/3479722.3480993>.

Understanding the Complexity of Negotiations

Supervisor: Alfie Abdul-Rahman

Areas: Human-centred computing (human-computer interaction), Natural Language Processing, Systems (software engineering, programming)

(Back to [Scholarship Not Allocated](#))

Project Description

Negotiated texts are the outcome of formal decision-making processes in which groups of actors deliberate, draft, and refine documents over time. Many foundational texts of the modern world—including international treaties such as the Universal Declaration of Human Rights or the Treaty of Versailles, and constitutions such as that drafted by the American states in 1787—were not authored by individuals but produced through complex collective negotiation. Even relatively short documents of this kind often embody thousands of individual proposals, revisions, and decisions, reflecting a dense history of interactions among negotiating parties.

To understand such processes, it is essential to track the involvement of delegations and assess their influence on the development of negotiated texts. Traditional close reading can provide detailed insights, while distance reading and computational methods such as machine learning enable the analysis of patterns across large corpora. However, to fully capture the complexity of negotiations, these approaches need to be integrated with innovative visual analytics methods that make it possible to explore the interplay of events, actors, and evolving drafts at scale.

This project will apply a visual analytics approach to illuminate the complexity of negotiations and the influence of delegations throughout the drafting process. It will combine computational text analysis, natural language processing, and advanced visualization to bridge formal records (such as meeting minutes and draft revisions) with informal archival materials (e.g., private diaries, correspondence, or social media records in contemporary contexts).

Possible Research Questions

- How can new static and interactive visualizations be designed to support data discovery and insight generation in large datasets of negotiation events represented as interacting timelines?
- What novel approaches can reveal the evolution of complex, technical documents over extended periods of negotiation?
- How can indexing and retrieval methods for negotiation datasets be improved, and how can their results be presented in more intuitive and interpretable displays?
- In what ways can natural language processing be used to link information from informal archives (e.g., diaries, letters, or social media feeds) with the formal records of a negotiation, thereby enriching our understanding of the decision-making process?

Expected Contributions

- A visual analytics framework for analysing negotiation processes at multiple scales.
- New visualization techniques for representing evolving texts, timelines, and actor interactions.
- Computational tools for indexing, linking, and exploring negotiation-related datasets.
- Methodological insights into integrating close reading, distance reading, and machine learning with visual exploration.
- Case studies demonstrating how these methods enhance our understanding of historical and contemporary negotiated texts.

Exploring Interactive Multi-Dimensional Approaches to Patient Communication in Oral Health Education

Supervisor: Informatics: Dr Alfie Abdul-Rahman & Dr Lin Gui FoDOCS: Dr Melanie Nasseripour & Dr Ana Angelova

Areas: Artificial Intelligence, Machine learning / Deep learning, Human-centred computing (human-computer interaction), Natural Language Processing, Systems (software engineering, programming)

(Back to [Scholarship Not Allocated](#))

Project Description

This joint project between the Department of Informatics and the Faculty of Dentistry, Oral & Craniofacial Sciences (FoDOCS) seeks to transform how patient communication is taught in oral health education. Effective patient-clinician communication is central to encouraging oral hygiene practices, supporting behaviour change, and demonstrating professionalism. However, conventional training methods are resource-intensive, requiring significant time, staff, and cost to deliver authentic clinical practice scenarios.

To address these challenges, the project proposes the design and evaluation of interactive, multi-dimensional learning environments, with a particular emphasis on immersive Virtual Reality (VR). VR offers an unparalleled ability to create engaging, realistic, and repeatable patient encounters where students can practice communication skills in a safe but lifelike clinical setting. Complementary approaches—including text-to-text and voice-to-voice communication systems—will expand accessibility, providing adaptable tools for different learners and educational contexts.

Generative Language Models (GLMs) will be harnessed to generate customised patient cases and conversational responses, ensuring that scenarios adapt dynamically to student inputs. These responsive simulations allow students to repeatedly practice communication tasks such as explaining oral hygiene, negotiating dietary changes, or discussing treatment options, thereby reinforcing learning and professional skill development.

Research Questions

- How can immersive VR be designed to simulate realistic patient-clinician interactions that enhance students' communication, empathy, and professionalism?
- What are the comparative benefits and limitations of VR, text-based, and voice-based simulation approaches in oral health education?
- In what ways can GLMs be integrated to generate personalised, contextually relevant, and adaptive patient scenarios that respond in real time to learners' actions?
- How do students perceive and engage with VR-based versus non-immersive approaches, and how does this affect learning outcomes?
- What metrics and evaluation frameworks can be developed to measure improvements in communication competence, confidence, and behaviour change skills using immersive and multi-dimensional tools?
- How can these approaches be scaled sustainably across curricula and adapted to different cultural and institutional contexts in oral health education?

Expected Contributions

- Development of a VR-based immersive training system for patient communication in oral health education.
- Comparative evaluation of VR, text, and voice-based simulations, identifying best practices for different learning settings.
- Integration of GLMs for adaptive, customised case study generation.
- A multi-dimensional toolkit for cost-effective, scalable communication training.
- Evidence-based recommendations for embedding immersive and interactive communication training within oral health curricula.

Data Science and Optimization for Sustainable Last-Mile Logistics

Supervisor: Dimitrios Letsios

Areas: Data Science and Optimization (*not* Foundations of Computing, more Machine Learning and Algorithm Engineering), Systems (software engineering, programming)

(Back to [Scholarship Not Allocated](#))

Project Description

Electric Vehicles (EVs) and cargo bikes are considered the main contemporary alternative to fossil fuel vehicles for reducing carbon emissions in urban environments. They are an integral part of transport decarbonization strategies and widely adopted as an environmentally friendly transport mode by logistics services during the last 5 years.

However, their adoption at scale poses emerging operational challenges that do not prevail in conventional distribution services with motor vehicles. On one hand, EVs require battery charging that can be time consuming and subject to accessibility constraints. On the other hand, cargo bikes are subject to rider cycling effort and safety challenges due to inadequate infrastructure and unfriendly geographies.

This project will develop methodologies for supporting stakeholders in addressing the aforementioned challenges and enhancing reliability, scalability, and financial sustainability of last-mile logistics services with sustainable transport modes, notably EVs and cargo bikes. To this end, the project will advance spatial data science and logistics optimization methodologies for assessing and the EV readiness and cargo bikeability of the London road network, determining suitable infrastructure interventions, and effectively running delivery operations. The project will also involve working with input (data and subject matter expertise) by a London council and a delivery company.

An applicant recruited for this project is expected to have good data analysis, operations research, and machine learning skills, ideally evidenced by relevant completed courses, projects, and/or work experience. Python programming skills will be essential for working on this PhD topic. Work is to be devoted to data processing, mathematical modeling, algorithm design, optimization, and performance evaluation of computationally obtained results.

Goal-based explanations for autonomous systems and robots

Supervisor: Gerard Canal

Areas: Artificial Intelligence, Robotics

(Back to [Scholarship Not Allocated](#))

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 who 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 on-going task. In this case, the new plan may be different to the original one, thus changing the behavior that the robot was performing.

This project will explore how to generate goal-based 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.

In addition to the available support by the CDT, the candidates will have the opportunity of contributing to the REXAR (UK) and COHERENT (international) research projects, while collaborating with and being supported by a network of researchers in aligned areas. These projects focus on reasoning for autonomous robots in assistive scenarios, dealing with explanations at different levels of the robotics system, and reasoning about goals and plans.

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).

Software sustainability analysis and improvement

Supervisor: Kevin Lano

Areas: Machine learning / Deep learning, Systems (software engineering, programming), Foundations of computing (algorithms, computational complexity)

(Back to [Scholarship Not Allocated](#))

Project Description

The project would consider techniques for analysing software sustainability (in the sense of energy use and energy efficiency) using either rule-based analysis and refactoring, or by the use of deep learning techniques such as LLMs to identify energy use flaws and potential refactorings.

Energy-efficiency improvement of machine learning systems is particularly important and could be the focus of the research. Equally, energy-efficiency improvement of mobile apps is another possible focus.

There is the potential for industrial collaboration in this area.

References

(Lano et al., 2024a) K. Lano et al., "Software modelling for sustainable software engineering", STAF 2024.

(Lano et al., 2024b) K. Lano et al., "Design Patterns for Software Sustainability", PLoP 2024

(Lano et al, 2025) K. Lano et al, "Sustainable Software Re-engineering", ECMFA 2025.