

The study of cooperation between intelligent agents has a long history. Despite often being framed in human terms, many existing models used for exploring the evolution of cooperation (e.g. the Iterated Prisoners' Dilemma) are based on assumptions found in ecological systems: minimal cognition, and in-the-moment decisions with minimal broader context. Under these assumptions, common outcomes include the prevalence of defection-based strategies, which seemingly justify or explain unsustainable, selfish behaviour, including notable results like Hardin's [2] Tragedy of the Commons.

Cooperation runs deeper than simple ecological models. Empirical work including by Nobel Prize winner Elinor Ostrom has shown [3] that societies can avoid antisocial outcomes, such as the Tragedy of the Commons, by establishing institutional rules that govern their interactions. Essentially, agents use their cognitive abilities to collectively act such as to 'change the rules of the game'. But this raises the question: when will group members establish and sustain such rules, so that they continue to enjoy from socially beneficial outcomes?

Models to explore this question have been formalised using evolutionary agent-based modelling and evolutionary game theory [4], and these techniques provide complementary insights. An important question here is how populations evolve new, stable institutional rules and compliance behaviour, moving from non-cooperative equilibria to cooperative ones, thereby achieving sustainable cooperation. Under the assumption of well-mixed populations, these models have provided insights into exactly when individuals will be incentivised to maintain their institution and enforce its rules, including the critical amount that a group must invest into incentivising agents to monitor rule compliance, and requirements for initial 'bootstrapping' of an institution from zero.

In reality, however, institution formation occurs on networks, and there is a wealth of literature on social learning in networks; see [1]. The interplay of this topic with evolutionary algorithms arises in economic theory, biology, and social science [6]. This body of literature outlines how one's personal strategy evolves based on observed strategies of some limited number of agents who are one's neighbours in an underlying social network, as opposed to evolving based on the entire population (as under a well-mixed assumption). A wide range of evolutionary rules can be used (e.g. *copy-majority*, *copy-satisfied-if-dissatisfied*), producing a variety of possible equilibria and evolutionary dynamics. Limited work considers the evolution of cooperation in a heterogeneous population: in [5], the evolution of the frequency of cooperative strategies is studied for simple two-player games on a network. We remark that the literature focuses predominantly on the 'base game' in a networked setting, rather than exploring how Ostrom-style institutional rules and collective action may evolve in the population.

This project will develop evolutionary and agent-based models of institution formation in networked populations. This extends the work of [4] to networks, and simultaneously extends the consideration of social learning in networks beyond base-game strategies and towards collective action and institutional rule formation. To do this, we will build evolutionary models of institution formation on several fixed networks. Using these models, the researcher will be able to extract rigorous results and statistical measures concerning the evolution of norms and cooperation, that provide new insight into the evolution of institutions, and thus 'deeper' forms of cooperation, on networks. We focus on the particular mechanisms of evolution in this setting, and future work may study the role of the network structure, how initial conditions affect the dynamics and equilibria, and how strengths of social ties can impact the outcomes.

References:

- [1] B. Golub and E. Sadler. Learning in social networks. *Available at SSRN 2919146*, 2017.
- [2] G. Hardin. The tragedy of the commons. *Science*, 162(3859):1243–1248, 1968.
- [3] E. Ostrom. *Governing the commons: The evolution of institutions for collective action*. Cambridge University Press, 1990.
- [4] S. T. Powers, A. Ekárt, and P. R. Lewis. Modelling enduring institutions: The complementarity of evolutionary and agent-based approaches. *Cognitive Systems Research*, 52:67–81, 2018.
- [5] M. Scatà, A. Di Stefano, A. La Corte, P. Liò, E. Catania, E. Guardo, and S. Pagano. Combining evolutionary game theory and network theory to analyze human cooperation patterns. *Chaos, Solitons & Fractals*, 91:17–24, 2016.
- [6] G. Szabó and G. Fath. Evolutionary games on graphs. *Physics reports*, 446(4-6):97–216, 2007.

Description of Research Group

The researcher will be based at the **Trustworthy Artificial Intelligence Lab**, at Ontario Tech University, the research lab led by Canada Research Chair Dr. Peter Lewis. We are an interdisciplinary lab hosted in the Software and Informatics Research Centre, and part of the Faculty of Business and Information Technology. The lab hosts a range of research projects that contribute to an overarching programme that explores how to make the relationship between AI and society work better. A key question is: how can we conceive of and build AI systems that meet this challenge?

One core part of the answer to this question lies in the social intelligence of agents, and in an understanding of how socially-sensitive agents interact with each other and people, and what emerges from this. This project therefore contributes to a fundamental theoretical underpinning for much of this work.

Our work draws on extensive experience in leading AI adoption projects in commercial and non-profit organizations across several sectors, as well as faculty research expertise in artificial intelligence, nature-inspired / evolutionary computing, artificial life, agent-based modelling, trust, network science, and computational self-awareness.

The lab, established thanks to grants from the Canada Foundation for Innovation (CFI) and Ontario Research Fund for Research Infrastructure, provides a state-of-the-art environment for collaborative interdisciplinary research with a computational focus. This includes:

- Access to state-of-the-art high performance computing facilities for simulation, numerical computation, machine learning, etc., through the provision of an on-site high performance CPU/GPU server. Further, the Faculty has committed technician time to support the lab and its research programs, therefore the researcher would have access to dedicated technical support in implementing and deploying experiments.
- A physical test-bed and demonstrator using small but powerful autonomous vehicular robots in a reconfigurable environment. This could be used to demonstrate the outcomes of the foundational work in real robot cooperation tasks, and also presents an opportunity for the researcher to supervise short undergraduate projects during their stay.
- A dedicated collaborative lab environment designed to support training, experiential learning, joint work, and socialising. This includes the provision of Google Jamboards for collaborative virtual/hybrid meetings and workshops, a large whiteboard wall, reconfigurable seminar/meeting space, workbenches for hardware-related projects (including robotics and other physical demonstrators), and office-style work cubicles for researchers.

The lab will be the primary base for the researcher. They will have a dedicated workspace and will be entitled to access all lab facilities and events. Co-location with other funded projects by the supervisors and other colleagues, whose students and postdoctoral researchers are also primarily based in the lab, will ensure cross-pollination of ideas and knowledge, and further contribute to a vibrant lab culture and environment. This is achieved through weekly lab meetings, broader inter-lab seminars, including with the nearby Social Robotics Lab; Security, Artificial Intelligence, and Networks Lab; Games And Media Entertainment Research (GAMER) Lab; and Institute for Digital Life, and regular inter-lab social events. All researchers at the lab, including the proposed researcher in this project, participate (including presenting) in the Computer Science Seminar Series.

The researcher will be jointly supervised by Dr. Peter Lewis and Dr. Jane Breen, who bring expertise in evolutionary modelling of institutions and complex networks respectively.



The Software & Informatics Research Centre. Our lab is based on the third floor.

Curriculum Vitae - Peter Lewis

Dr. Peter Lewis holds a **Canada Research Chair in Trustworthy Artificial Intelligence** at Ontario Tech University. He has a background in computer science, with expertise in several areas at the intersection of **artificial intelligence (AI)** and **socio-technical systems**. He has published extensively on the application of principles from psychology, economics, biology and sociology to computing systems (>75 papers; >1600 citations; h-index = 21).

His research has been covered in the international press, including the *BBC* and *Daily Mail Online* (UK), *Hindustan Times* (India), and *The Global Times* (China), and recognized through awards: Best Paper Award at the *12th International IEEE Conference on Self-Adaptive and Self-Organizing Systems* (SASO) in 2018; and *BT Gordon Radley Best Author of Innovation* (Runner Up) in 2010.

He has led research projects worth over \$7.5m (CAD equivalent), attracting funding from the Canadian Natural Sciences & Engineering Research Council (NSERC) and Canada Foundation for Innovation (CFI), the European Union (FP7 & Horizon 2020), Innovate UK, and commercial partners. His track record of academic and industrial research collaborations (>100 co-authors and >70 company partners) generates both scientific and business value (e.g., 60% reduction in mobile device battery usage by applying multi-objective optimization to mobile-cloud architectures; 44% reduction in vehicle usage through novel scalable bio-inspired fleet optimisation techniques). His work also led to valuable impact in diverse areas such as smart camera networks, interactive music, mobile robotics, and manufacturing. He has extensive experience of delivering projects and leading teams of researchers, including being Director of the £3.3M *Think Beyond Data* initiative, which provides an AI R&D capability to SMEs in the Midlands of England. He also co-founded *Beautiful Canoe*, a software house and social enterprise, whose vision is to develop the technology leaders of the future.

Employment & Education

- **Ontario Tech University, Canada** (2021–present): Canada Research Chair and Associate Professor.
- **Aston University, UK** (2013–2021): Lecturer then Senior Lecturer in Computer Science.
- **University of Birmingham, UK** (2010–2013): Postdoctoral Research Fellow, Centre of Excellence in Computational Intelligence and Applications (CERCIA), School of Computer Science.
- **University of Birmingham, UK** (2006–2010): Ph.D. Computer Science. Supervisor: Professor Xin Yao. Thesis title: Evolutionary market-based resource allocation in decentralised computational systems.

Positions of Responsibility

- Associate Editor, *IEEE Technology and Society Magazine* and *ACM Trans. Autonomous and Adaptive Systems*.
- Steering Committee Member, *IEEE Int. Conference on Autonomic Computing and Self-Organizing Systems*.
- Board Member & Social Impact Chair, *International Society for Artificial Life*.

Research Interests & Key Publications

• Evolutionary Modelling of Socio-Technical Systems:

One important contribution of my work is the characterisation of a range of modelling techniques in terms of their ability to capture and explain social and cognitive behaviour, such as the formation of norms. This led to a set of recommendations for modellers of socio-technical systems.

- CM Barnes, A Ghouri, and PR Lewis. Explaining evolutionary agent-based models via principled simplification. *Artificial Life*, 27(3), 2021
- ST Powers, A Ekárt, and PR Lewis. Modelling enduring institutions: The complementarity of evolutionary and agent-based approaches. *Cognitive Systems Research*, 52:67–81, 2018

• Computational Self-Awareness & Reflection:

I played a significant role in establishing the field of self-aware computing, making several key contributions, including concepts, reflective architectures, and algorithms, for systems that can learn and reason about their own behaviours, interactions, capabilities, and motivations.

- PR Lewis, A Chandra, F Faniyi, K Glette, T Chen, R Bahsoon, J Torresen, and X Yao. Architectural aspects of self-aware and self-expressive computing systems. *IEEE Computer*, 48:62–70, 2015
- PR Lewis, M Platzner, B Rinner, J Torresen, and X Yao, eds. *Self-Aware Computing Systems: An Engineering Approach*. Springer, 2016.

• Trust and Artificial Intelligence:

My analysis of the role of trust in intelligent socio-technical systems has helped to shape understanding of the field and clarify terminology concerning different forms of trust at play between humans and intelligent machines. One specific recent contribution of note is a model for the systematic analysis of trust decisions with machines.

- PR Lewis and S Marsh. What is it like to trust a rock? A functionalist perspective on trust and trustworthiness in artificial intelligence. *Cognitive Systems Research*, 72:33–49, 2022.
- P Andras et al. Trusting intelligent machines. *IEEE Technology and Society Magazine*, 37, 2018

Employment & Education

- **Ontario Tech University** (2019–present): Assistant Professor, Applied Mathematics, Faculty of Science.
- **Iowa State University** (2018–2019): Postdoctoral Scholar, Mathematics.
- **University of Manitoba** (2014–2018): Ph.D. Mathematics.
 - Supervisor: Dr. Steve Kirkland
 - Thesis title: Markov Chains under Combinatorial Constraints: Analysis and Synthesis
 - Winner of the Distinguished Dissertation Prize

Research Interests & Key Publications

- **Markov chain theory and its applications:**

A Markov chain is a type of mathematical model which has been used to describe countless types of dynamical systems, from traffic behaviour in an urban road network, to molecular dynamics and their role in drug design. My research focuses on the combinatorial structure of the transition matrix, and its influence on the behaviour of the system. This extends to random walks on graphs and networks (the fundamental idea behind Google PageRank), and I am particularly interested in random-walk-based connectivity measures such as Kemeny's constant.

- [A structured condition number for Kemeny's constant.](#)
Jane Breen, Steve Kirkland.
SIAM Journal of Matrix Analysis and its Applications 40:4 (2019), 1555-1578.
- [Clustering behaviour in Markov chains with eigenvalues close to one.](#)
Jane Breen, Emanuele Crisostomi, Mahsa Faizrahnemoon, Steve Kirkland, Robert Shorten.
Linear Algebra and its Applications 555 (2018), 163–185.

- **Spectral graph theory and combinatorial matrix theory:**

The field of spectral graph theory explores the question of what structure can be determined of a graph from the eigenvalues of an associated matrix, while combinatorial matrix theory determines what can be said of a matrix based on its associated graph.

- [Computing Kemeny's constant for barbell-type graphs.](#)
Jane Breen, Steve Butler, Nicklas Day, Colt DeArmond, Kate Lorenzen, Haoyang Qian, Jacob Riesen.
Electronic Journal of Linear Algebra 35 (2019), 583–598.
- [Maximum spread of graphs and bipartite graphs](#)
Jane Breen, Alex Riasanovsky, Michael Tait, John Urschel. (2022+)

- **Complex networks:**

A complex network is essentially a graph, but is assumed to have inherent properties from the domain of applications, such as social networks, epidemiological networks, neurological networks, urban road networks. Network science is a hugely interdisciplinary field, and my recent interest in this field is a deliberate move to increase the interdisciplinarity in my research program, and ties together the two primary research interests listed above. My expertise in random walk methods can be applied to better understand certain centrality metrics and their influence on dynamic systems on the network.

- [On node ranking in graphs and disease spread in networks.](#)
E. Dudkina, M. Bin, J. Breen, E. Crisostomi, P. Ferraro, S. Kirkland, J. Maráček, R. Murray-Smith, T. Parisini, L. Stone, S. Yilmaz, R. Shorten. (2022+)

Additional information

Additional funding and Scholarship Extension

We are pleased to offer:

- Match funds in order to extend the internship up to an additional three months at the equivalent rate (thereby making a six-month research visit), if the researcher so desires.
- Additional funding for reasonable travel expenses to and from Canada at the beginning and end of the internship.
- Funding to cover travel costs associated with attending the EvoStar conference in order to present the results of the work.
- A waiver of all bench fees for the researcher, and the researcher will have access to the lab's research infrastructure and facilities as laid out in the research group description as an in-kind benefit.

About Ontario Tech University

Ontario Tech University has grown out of a bold, ambitious vision to take on the grand challenges we face as a society and find solutions to meet tomorrow's needs. Under the vision of 'tech with a conscience', Ontario Tech continually and strategically invests resources in research that creates knowledge, solves problems, results in economic and social innovation, and engages students. Founded in 2002, Ontario Tech University has quickly established itself as a vibrant young university with a strong focus on collaborative, interdisciplinary research.

Ontario Tech's research community continues to be enhanced by new and existing infrastructure and resources as well as the organization of numerous ongoing workshops, colloquia, and seminar series open to the University and the surrounding community. The Trustworthy AI Lab is housed in the Software and Informatics Research Centre (SIRC), opened in 2017. This four-story, 7,600m² space is a modern facility devoted to research, study, and collaboration. The Faculty of Business and Information Technology (FBIT) hosts their own Data Centre in SIRC to support faculty research, and this houses the lab's dedicated high-performance computing server. SIRC was designed to promote interdisciplinary collaborations between computer scientists, information technology (IT) specialists, and engineers by providing a hub for research and training activities related to the University's expertise in health and business analytics, IT security, networking, gaming, and software engineering as well as interactions with community and industry partners in these areas.

About Oshawa and Environs

Ontario Tech is situated in Oshawa, on the Eastern edge of the Greater Toronto Area. This provides easy access to downtown Toronto (including via local train) as well as plenty of opportunity to explore and enjoy the many forests and lakes that Canada is known for. The campus area contains ample modern student and researcher housing, as well as diverse restaurants and shops, and the researcher will be supported to find suitable accommodation for their stay.

