1. Deep Learning for Inverse Image Estimation

Supervisor: Stephen Gould, CECS (stephen.gould@anu.edu.au) and Draogmir Neshev, RSPE (dragomir.neshev@anu.edu.au)

Deep learning applied to contemporary problems in artificial intelligence has resulted in significant advances in areas such as image understanding, natural language processing, and autonomous driving. An area where deep learning is particularly relevant is in modelling inverse problems, i.e., problems where we can efficiently compute or simulate a forward model from X to Y but where determining X from Y is difficult.

This project investigates one such a problem relating to the design of a novel optical system. Specifically, the project looks at using deep learning techniques to determine what stimulus is needed to produce a given holographic image under a certain physical model. Successful completion of the project may result in close industry collaboration and involvement.

Requirements/Prerequisites:

- Students should have taken a course on computer vision and/or machine learning. They should have strong mathematical and programming skills (in Python). Familiarity with deep learning techniques is desirable.
- Independent worker willing to tackle a challenging research problem with potential follow through beyond the summer scholar program.

Students will gain:

- Experience working on cutting-edge problems in computer vision using deep learning.
- Skills in conducting and presenting research, both written and oral.
- Be involved with state-of-the-art researchers with the goal of producing impactful science.
- Experience in seeing the path from fundamental research to commercial application.

2. Generic Reinforcement Learning Agents

Supervisor: Marcus Hutter [Marcus.Hutter@anu.edu.au]

Agent applications are ubiquitous in commerce and industry, and the sophistication, complexity, and importance of these applications is increasing rapidly; they include speech recognition systems, vision systems, search engines, planetary explorers, auto-pilots, spam filters, and robots [RN03]. Existing agent technology can be improved by developing systems that can automatically acquire during deployment much of the knowledge that would otherwise be required to be built in by agent designers. This greatly reduces the effort required for agent construction, and results in agents that are more adaptive and operate successfully in a wide variety of environments [LH07].

Goals:
Technically, the project is about a recent general approach to learning that bridges the gap between theory and practice in reinforcement learning (RL). General-purpose, intelligent, learning agents cycle through sequences of observations, actions, and rewards that are complex, uncertain, unknown, and non-Markovian [RN03]. On the other hand, RL is well-developed for small finite state Markov decision processes (MDPs) [SB98]. Extracting the right state representations out of bare observations, that is, reducing the general agent setup to the MDP framework, is an art that involves significant effort by designers. The project is to investigate (by simulations or theoretical) recent models [Hut09] that automate the reduction process and thereby significantly expand the scope of many existing RL algorithms and the agents that employ them.

Requirements/Prerequisites:
- background in Artificial Intelligence and Machine Learning
- good programming skills

- performing (computer) experiments and analysing results
- good math skills; linear algebra at the very minimum
- mastering elementary probability calculus

Students will gain:
- acquaintance with state-of-the art RL algorithms
- improved math skills: linear algebra, statistics, probability, and information theory

Background literature:

- [RN10] S.J. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach* Prentice-Hall, Englewood Cliffs, NJ, 3rd edition, 2010.
- [SB98] R.S. Sutton and A.G. Barto. *Reinforcement Learning: An Introduction* MIT Press, Cambridge, MA, 1998.
- [LH07] S. Legg and M. Hutter. Universal intelligence: A definition of machine intelligence. Minds & Machines, 17(4):391-444, 2007.
- [Hut09] M. Hutter. Feature reinforcement learning: Part I: Unstructured MDPs. Journal of Artificial General Intelligence, 1:3-24, 2009.
- [Hut14] M. Hutter. Extreme state aggregation beyond MDPs. In Proc. 25th Intl. Conf. on Algorithmic Learning Theory (ALT'14), LNAI 8776, pages 185--199.

3. Mathematical Foundations of Artificial Intelligence

Supervisor: Marcus Hutter [Marcus.Hutter@anu.edu.au]

The first decade of this century has seen the nascence of the first mathematical theory of general artificial intelligence. This theory of Universal Artificial Intelligence (UAI) has made significant contributions to many theoretical, philosophical, and practical AI questions. In a series of papers culminating in book [Hut05], an exciting sound and complete mathematical model for a super intelligent agent (AIXI) has been developed and rigorously analysed. The model is actually quite elegant and can be defined in a single line:

$$\text{AIXI} \qquad a_k := \arg\max_{a_k} \sum_{o_k r_k} \ldots \max_{a_m} \sum_{o_m r_m} [r_k + \ldots + r_m] \sum_{q\,:\,U(q,a_1..a_m)=o_1 r_1..o_m r_m} 2^{-\ell(q)}$$

(*k*=now, *a*=action, *r*=reward, *o*=observation, *U*=Universal TM, *q*=program, *m*=lifespan, *l*=length)

Goals:
The fundamentals of UAI are already laid out, but there are literally hundreds of fundamental theoretical/mathematical open questions [Hut05,Hut09] in this approach that have not yet been answered. The complexity ranges from suitable-for-short-projects to full PhD theses and beyond.

Requirements/Prerequisites:
- excellent math skills, ideally in information theory or probability or statistics
- creativity in finding and constructing proofs
- ability to clearly interpret the meaning of mathematical theorems

Students will gain:
- acquaintance with the most promising mathematical approach to general AI
- experience working on interdisciplinary research questions
- experience in proving theorems
- advancements in active math skills

Background literature:

- [Hut05] M. Hutter. Universal Artificial Intelligence: Sequential Decisions based on Algorithmic Probability. Springer, Berlin, 2005.
- [Hut09] M. Hutter. Open problems in universal induction & intelligence. Algorithms, 3(2):879-906, 2009.

4. Universal Artificial Intelligence

Supervisor: Marcus Hutter [Marcus.Hutter@anu.edu.au]

The dream of creating artificial devices that reach or outperform human intelligence is an old one. Most AI research is bottom-up, extending existing ideas and algorithms beyond their limited domain of applicability. The information-theoretic top-down approach (UAI) pursued in [Hut05] justifies, formalizes, investigates, and approximates the core of intelligence: the ability to succeed in a wide range of environments [LH07].

All other properties are emergent. Recently, effective approximations of UAI have been derived and experimentally investigated [VNHUS11]. This practical breakthrough has resulted in some impressive applications. For the first time, without providing any domain knowledge, the same agent is able to self-adapt to a diverse range of interactive environments. For instance, it is able to *learn* from scratch to play TicTacToe, Pacman, Kuhn Poker, and other games by trial and error. These achievements give new hope that the grand goal of Artificial General Intelligence is not elusive.

Goals:

The theoretical [Hut05], philosophical [LH07], and experimental [VNHUS11] foundations of UAI are already laid out, but plenty remains to be done to solve the AI problem in practice. The complexity of the open problems ranges from suitable-or-short-projects to full PhD theses and beyond.

Requirements/Prerequisites:

- background in AI or ML or statistics or information theory.
- excellent programming or writing or math skills

Students will gain:

- acquaintance with the most comprehensive theory of rational intelligence to date.
- experience in writing a literature survey
- experience in advancing the state-of-the art implementation of AIXI and apply it to new (toy/game) problems
- experience in proving non-trivial theorems.

Background literature:

- [Hut05] M. Hutter. *Universal Artificial Intelligence: Sequential Decisions based on Algorithmic Probability*, Springer, Berlin, 2005.
- [LH07] S. Legg and M. Hutter. *Universal intelligence: A definition of machine intelligence.* Minds & Machines, 17(4):391-444, 2007.
- [VNHUS11] J. Veness, K. S. Ng, M. Hutter, W. Uther, and D. Silver. *A Monte Carlo AIXI approximation.* Journal of Artificial Intelligence Research, 40:95-142, 2011.

5. A New Cross-layer Approach to Communication Security

Supervisor: Dr. Xiangyun (Sean) Zhou [xiangyun.zhou@anu.edu.au]

We send and receive private and sensitive information over wireless networks every day. Do we really trust that the communication is perfectly secured with zero information leakage? In this project, we look at how communication is secured using the classical method of cryptography and explore new ways of providing high security using relatively simple cryptographic techniques. We take a new angle to look at the security problem, that is, to consider the impact of intrinsically imperfect and noisy nature of the communication medium, which has been largely ignored in the prior study on cryptography. The study on the performance of simple cryptographic techniques over noisy communication medium will lead to new insights on security design.

Requirements/Prerequisites:

- Strong mathematical skills.
- Experience in programming and use of Matlab.
- Basic knowledge of cryptography and communication theory is desirable.

6. Rethinking the Design of Communication Receiver

Supervisor: Dr. Xiangyun (Sean) Zhou [xiangyun.zhou@anu.edu.au]

This projects addressed a very fundamental and “old” problem, i.e., how a communication receiver processes its received signal for information detection. Communication engineers and researchers have well established two separate receiver designs in the past decades, one based on coherent processing and the other based on non-coherent processing.

Our recent research revealed a surprising result that a huge improvement on data rate can be obtained by jointly using coherent and non-coherent processing at the receiver, which has never been considered before. This makes the “old” problem new again with lots of opportunities in exploiting the new receiver structure. This project will look into some of the interesting open problems for obtaining better understanding on the performance improvement that the new receiver structure can realise.

Requirements/Prerequisites:

- Strong mathematical skills.
- Experience in the use of Matlab.
- Basic knowledge of communication theory

7. Atomically thin Nano-materials and devices

Supervisor: Yuerui Lu [yuerui.lu@anu.edu.au]

Two-dimensional (2D) nano-materials, such as molybdenum disulfide (MoS_2) and graphene, have atomic or molecular thickness, exhibiting promising applications in nano-electro-mechanical systems. Graphene is a one-atom thick carbon sheet, with atoms arranged in a regular hexagonal pattern. Molybdenum disulfide (MoS_2) belongs to transition metal dichalcogenides (TMD) semiconductor family YX_2 (Y=Mo, W; X=S, Se, Te), with a layered structure. These 2D nano-materials can be integrated into nano-electro-mechanical systems, enabling ultra-sensitive mechanical mass sensors, with single molecule or even single atom sensitivities.

Moreover, the mechanical resonators based on these 2D nano-materials would be a perfect platform to investigate quantum mechanics, opto-mechanics, material internal friction force, nonlinear physics, etc.

Requirements/Prerequisites:

- Specific requirements and/or prerequisites to be determined by the Supervisor.

8. Nano Biomedical Devices

Supervisor: Yuerui Lu [yuerui.lu@anu.edu.au]

The ability to detect bio-molecule at ultra-low concentrations (e.g. atto-molar) will enable the possibility of detecting diseases earlier than ever before. A critical challenge for any new bio-sensing technology is to optimize two metrics --- shorter analysis time, and higher concentration sensitivity in clinically relevant small volumes. Moreover, practical considerations are equally important:

- simplicity of use

- mass producible (low cost), and
- ease of integration within the clinical structure.

Compared with other methods, nano-electro-mechanical system (NEMS) based bio-sensors are promising in clinical diagnostics because of their extremely high mass sensitivity, fast response time and the capability of integration on chip. We have demonstrated a low concentration DNA (atto-molar sensitivity) optically interrogated ultrasonic mechanical mass sensor, which has ordered nanowire (NW) array on top of a bilayer membrane. This method represents a mass-based platform technology that can sense molecules at low concentrations, which could be useful for early-stage disease detection. We can develop this sensor further to measure an array of biomarkers (e.g. DNA or proteins), by providing both the needed specificity and sensitivity in physiological disease (e.g. cancer) detection.

Requirements/Prerequisites:
- Specific requirements and/or prerequisites to be determined by the Supervisor.

9. Human Centred Computing / Bioinspired Computing projects

Supervisor: Tom Gedeon [tom@cs.anu.edu.au]

Please see my own student projects page at http://cs.anu.edu.au/people/Tom.Gedeon/projects.html for the current list of my projects. They range from human centred projects using eye gaze, physiological signals, EEG, fNIRS and other sensors, and effectuators such as electrical body and brain stimulation; bioinspired projects using neural networks, deep learning or evolutionary algorithms; or a combination of both human centred and bio-inspired. There is scope for multiple students in these areas.

Requirements/Prerequisites:
- An interest in human experiments particularly relating to human internal states such as emotion, stress, anxiety and so on;
- An interest in bio-inspired computing techniques;
- A willingness to have fun learning how computers can be made to be more useful and responsive to human beings.

Students will gain:
- Practical experience in design and conduct of human experiments
- Experience in advanced data analysis and prediction.

10. Emotional Reaction Dataset

Supervisor: Tom Gedeon [tom@cs.anu.edu.au]

Most of the top 10 cited papers in ICMI below to our annual "Emotion recognition in the wild challenge". We wish to extend that dataset by recording physiological signals from humans responding to the dataset. This project will be to collect data breadthwise and depthwise, to prepare the data for publication, and to participate in writing an initial paper reporting on the new dataset.

Requirements/Prerequisites:
- An interest in human experiments, particularly relating to human internal states such as emotion
- An interest in bio-inspired computing techniques, particularly neural
- A willingness to have fun learning how computers can be made to be more useful and responsive to human beings.

Students will gain:
- Practical experience in design and conduct of human experiments
- Experience in advanced data analysis and prediction.

11. Personalising Handwriting Recognition

Supervisor: Tom Gedeon [tom@cs.anu.edu.au]

Handwriting recognition using deep learning is well understood and remains an active area of research interest. This project is to construct an end-to-end system which recognises real world handwriting, using existing code as far as possible. The goal is to develop personalisation adaptations needed to work well on single individuals. A substantial handwritten corpus is available, along with access to the author.

Requirements/Prerequisites:

- An interest in computer vision or deep learning
- A willingness to have fun learning how computers can be made to be more useful and responsive to human beings.

Students will gain:

- Practical experience in design and personalisation of deep learning models
- Experience in advanced data analysis and prediction.

12. Prototype a Parallel SAT Solving Algorithm for Cyber Security Applications

Supervisor: Charles Gretton [charles.gretton@anu.edu.au]

Develop and evaluate a parallel implementation of a Boolean SAT(isfiaiblity) procedure.
The solution procedure shall combine CDCL (Conflict Directed Clause Learning) and DLS (Dynamic Local Search) heuristics.

Target decision problems shall be from applications in cyber security -- e.g. 'how do I perturb the input to a program to make a particular block of code execute?'.
Existing CDCL and DLS codebases in C++ are available to be leveraged in this exercise.
Strong students can use this project as a platform towards a strong honours thesis, or a PhD project.

Requirements/Prerequisites:

- Passionate, high performing, and strong algorithmic-programming skills.

13. Traffic Congestion Prediction Using Origin-to-Destination Pattern Matching

Supervisors: Dr. Marian-Andrei Rizoiu (ANU) and Dr Adriana- Simona Mihaita (Data61)

Predicting traffic congestion is a very hard topic due to various external and stochastic events. Current traffic simulation model can only be used for traffic congestion prediction if they are alimented using updated and representative Origin-to-Destination (OD) matrices (number of trips from any origin to any destination in the city).

This work aims to study a collection of time-dependent OD matrices for an urban network – the city of Sydney – and to build a pattern analysis framework for:

a) Choosing the most representative OD based on updated information (weather, public events, etc.) and
b) Predicting the OD matric for the next 30 min or 1 hour in the future

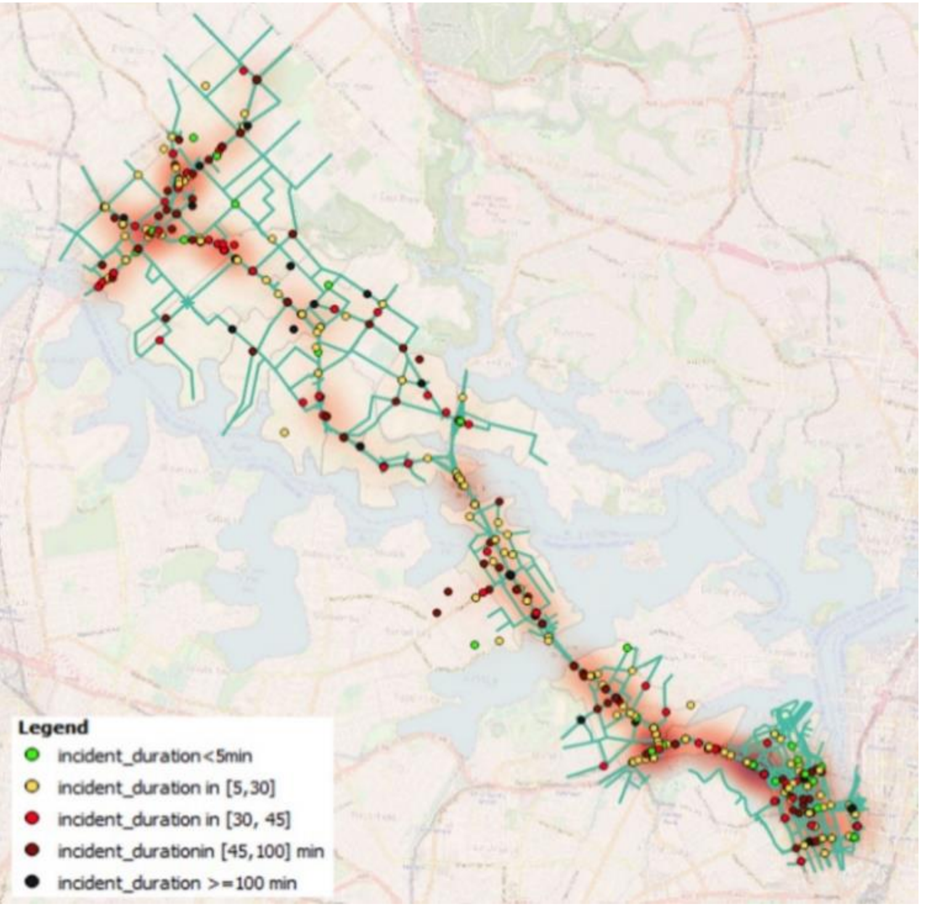

Keywords:
Prediction, pattern analysis, machine learning, OD matrix

Requirements/Prerequisites:

- background Machine Learning and/or Data Science methods;
- good programming skills
- performing (computer) experiments and analysing results
- good math skills and linear algebra (matrix operations)
- **desirable:** Git, R/Python, desire to make sense of real data and solve real issues.

14. Dynamic Knowledge Tracing

Supervisor: Qing Wang [Qing.Wang@anu.edu.au]

Knowledge tracing refers to the problem of modelling knowledge of students over time as they interact with a learning system, which have been studied extensively in computer supported education. However, the knowledge tracing problem is inherently difficult as human learning is grounded in the complexity of both the human brain and human knowledge. Recently, deep neural networks have been successfully used to build several knowledge tracing models, including the Deep Knowledge Tracing (DKT) model developed by the researchers from Stanford University (Piech et al., Deep knowledge tracing, NIPS 2015) and Dynamic Key-Value Memory Network (DKVMN) model (Zhang et al., Dynamic key-value memory networks for knowledge tracing, WWW 2017), which have been shown to be able to predict student performance better than the previous knowledge tracing models.

Students will gain:

- a solid understanding of deep learning models
- skills to implement and apply these techniques in a real-world application setting.

Requirements/Prerequisites:

- Strong skills in programming (Python) are required.
- Solid knowledge background in data mining and machine learning are desired.

Background Literature:

- Piech et al., Deep knowledge tracing. In Advances in Neural Information Processing Systems, 2015
- Xiong et al., Going deeper with deep knowledge tracing. In Educational Data Mining, 2016
- Zhang et al., Dynamic key-value memory networks for knowledge tracing. In Proceedings of the 26th International Conference on World Wide Web, 2017

15. Discovering Inconsistent Data in a Dynamic World

Supervisor: Qing Wang [Qing.Wang@anu.edu.au]

Inconsistent data exists everywhere but not always evident. However, using inconsistent data may lead to poor decision making, expensive mistakes, communication chaos, etc. There is an increasing industry-driven demand for tools that can efficiently identify and reduce inconsistent data. One important aspect in developing such tools is to find efficient and effective ways of capturing the structure and semantics of data which will enable the detection of inconsistent data.

The project aims to analyse and mine patterns of inconsistencies occurring among data, and then develop efficient algorithms to capture and resolve inconsistent data based on such patterns.

Students will gain:

- a solid understanding of data mining and database technologies,
- skills to use these techniques to analyse data and improve data quality.

Requirements/Prerequisites:

- Strong skills in software development (Java or Python) are required.
- Solid knowledge background in data mining and relational database are desired.

16. Australian Signals Directorate (ASD) – Various Projects

Supervisors: To be determined

The Australian Signals Directorate (ASD) is a vital member of Australia's national security community, working across the full spectrum of operations required of contemporary signals intelligence and security agencies: intelligence, cyber security and offensive operations in support of the Australian Government and Australian Defence Forces (ADF).

ASD values strong critical thinking and research skills, with an emphasis on an ability to present and reason analysis in diverse forums. ASD analysts work on a broad range of tasks which include: investigating large and complex data sets, developing new methods of analysing data for intelligence or information security purposes and solving cryptological problems using advanced mathematical concepts.

ASD is pleased to partner with ANU to sponsor up to ten Summer Research Scholarships in the following areas:

- number theory
- cryptography
- statistics
- data science

- secure systems,
- vulnerability research

To apply, submit a short paper describing your work. There will also be an opportunity to deliver a short presentation about your work to ASD and ANU staff. Please note on your application if you have a preference for one of the nominated topics.

Requirements/Prerequisites:

- Australian Citizens only
- Minimum GPA of 5.5/7.0 (preference will be given to applicants with strong academic performance in a cognate background)
- Nominate your area/s of interest, based on the list above.

17. Feature tracking using event cameras

Supervisor: Robert Mahony, CECS (Robert.Mahony@anu.edu.au)

Event cameras provide a new paradigm for robotic vision. The low temporal latency and high dynamic range of event camera data sequences are ideal for high speed vision applications necessary for modern dynamic robotic systems. However, event cameras do not provide traditional image sequences and there is a need to develop new processing paradigms to extract the necessary information for robotic perception algorithms. This project will investigate point feature detectors for event cameras. In addition to single camera features, the project will consider stereo camera features where temporal synchronisation between events from different cameras is exploited as one of the cues for data association.

Requirements/Prerequisites:

- Students should have strong programming skills, both prototyping in MATLAB or python, and coding in C.
- Background in systems and control is desirable.
- Independent worker willing to tackle a challenging research problem.

Students will gain:

- Experience working on cutting-edge problems in computer vision.
- Skills in conducting and presenting research, both written and oral.
- Be involved with state-of-the-art researchers with the goal of producing impactful science.