Data61 PhD Scholarship Project

September 2022

Contents

10. Taming Uncertainty in AI Driven Exploration of Complex Systems

Description

Granular matter is an example of a complex system in which the fine details of the interactions between particles at small scales profoundly affects the observed macroscopic behaviour of the material. This has numerous real-world consequences, from the dramatic in for example the sudden onset of landslides, to the mundane (but highly impactful) problems of processing powders in pharmaceuticals and ensuring proper mixing of ingredients in food products.

Numerical simulations lie at the forefront of granular matter research. Like all digital simulations of the real world, these are subject to uncertainty: observational noise, calibration error in simulation parameters, limitations of the numerical solver, and gaps in the model representation. Machine learning and statistical tools to quantify and reduce this uncertainty are relatively advanced for mesh-based methods, but still underdeveloped for the particle-based methods used to study granular matter. Even useful representations of uncertainty in simulations that track millions of individual particles have yet to be devised.

This project will explore new ways of combining AI and physics based computational modelling within a consistent uncertainty quantification framework. This will be applied both to accelerate the simulation of granular matter in applications such as comminution modelling and metal 3D printing, and as a means of consistently quantifying and understanding the ability of the simulated models to accurately represent their real-world counterparts.

Skills and capabilities required for the project

A first-class honours (or equivalent) degree in Computer Science / Physical Sciences / Mathematical Sciences, ideally with a working knowledge of ML principles. Geographic information systems; Data analytics; Scientific computing (Python, R, Matlab); Statistics (spatial is a bonus)

Supervisors

Richard Scalzo and Gary Delaney (CSIRO's Data61)

Location

Clayton, Victoria

11. Uncertainty quantification in multi-step decision making

Description

In many real-world applications, decisions need to be taken at multiple levels where a current decision may be based on several previous decisions. These decisions may be made by several Machine Learning (ML) models along the pipeline, each producing an intermediate decision. Uncertainty from errors made by models in earlier steps of decision making, propagate through the process and the final decision will be impacted from errors made in earlier levels. This is especially complex when considering epistemic (systematic) uncertainty (i.e., uncertainty in the actual model used in the specific decision process). It is necessary to quantify this uncertainty alongside the final decision in such multi-step decision making processes. The following are some research questions that needs to be investigated:

1. How can the error and corresponding uncertainty be quantified along the line of decision making?

2. How can ML be embedded in uncertainty frameworks that represent and propagate uncertainty?

3. How do such uncertainty frameworks scale computationally?

4. How can the decision steps be coherently integrated to drive better decisions? Can we make the learning algorithms along the line of decision making complementary to reduce uncertainty?

For a real–world example of multi–step decision making, consider a prawn farm where we want to use a data driven technology to automatically measure prawn size. This will involve using cameras to capture images of prawns as they are pulled out of water. The multi-step decision making will look like:

(Step 1): use of a machine learning model to segment the image and detect prawns,

(Step 2): Use of a machine learning model to convert pixel length of prawns to actual length of prawns, and

(Step 3): Use of a machine learning model to convert the length into prawn weights.

An error made during detection step (e.g., a partially visible prawn detected) will influence the length and weight calculations in future steps. A framework to quantify such uncertainty at multiple steps and combine them can certainly help with better decision making. We are currently conducting a project on prawn measurement under Digiscape funding. We have collected data for the prawn measurement across multiple Digiscape projects and have a use case. The PhD student will have access to real world dataset to validate the developed algorithms as part of PhD project.

Skills and capabilities required for the project

1. Machine learning

2. Statistics

3. Python/MATLAB

4. Desirable: Image processing/Computer vision

Supervisors

Ashfaqur Rahman and Joel Dabrowski (CSIRO's Data61)

Location

Sandy Bay, Tasmania

12. A probabilistic data-driven framework to scenario planning

Description

Strategic foresight is a popular tool that policy makers use to anticipate potential opportunities and challenges that could arise when a policy is implemented. Traditionally, researchers often consider qualitative approaches in strategic foresight due to the complexity of data (i.e., data comes from various types of information such as spatial data, temporal data, and image data). This project aims to provide a probabilistic data-driven framework for strategic foresight and scenario planning.

In particular, this project will focus on

- Developing a framework that integrates data from various types of information (e.g., spatial data, temporal data, and image data); and handles with missing data and mixed frequency data, which are common in strategic foresight.
- Developing a probabilistic data-driven model in Bayesian framework to i) account for uncertainties (data error, model misspecification, sampling error); and ii) incorporate some priori information relating to policies in modelling to provide a better prediction.
- Developing an efficient algorithm to estimate the probabilistic data-driven model.

This project will provide policy makers a probabilistic data-driven tool for scenario planning.

Skills and capabilities required for the project

- A first-class honours (or equivalent) degree from a well-recognised university in the physical, natural or social sciences
- An ability to comprehend complex strategy, foresight and policy problems encountered by government, industry and community organisations and to support decision makers using quantitative techniques
- Advanced practical skills, knowledge and expertise in computer coding, data science, statistical forecasting and machine learning
- Strong programming experience in Python, MatLab or C++ or other relevant language

Supervisors

Stefan Hajkowicz and Kelly Trinh (CSIRO's Data61)

Location

Sydney, Melbourne, Brisbane, Canberra or remote where requirements can be met

13. Using AI and NLP to assist in scientific and technical document creation

Description

Government and industry are increasingly interested in AI-assistance in document preparation. Recent methods in artificial intelligence (AI) and natural language processing (NLP) have shown the tremendous potential of neural network language models to capture language use and even some knowledge/societal norms. The PhD student will work with open domain data sets to explore key NLP problems in the context of this NLP application space, such as: (1) neural language modelling, (2) information extraction for scientific documents, (3) discourse modelling, (4) improving language models with domain knowledge, and (5) data augmentation and data set preparation for NLP in low-resource scenarios. Applicants should have a research background in one or more of the following areas: linguistics, computational linguistics, natural language processing, artificial intelligence, machine learning.

Skills and capabilities required for the project

One or more of the following: Natural Language Processing, Computational Linguistics, Artificial Intelligence, Machine Learning

Supervisors

Stephen Wan (CSIRO's Data61)

Location

Eveleigh, NSW

14. High-dimensional state inference for dynamic systems with uncertainty quantification

Description

The project involves developing novel methods for state inference in complex dynamical systems with later focus on the estimation of joint model and state from observations of the state over time. For example, one may have observations in time from a weather system and be interested in both estimating the relevant climate model describing the system as well as forecasting future states of the system. In very high-dimensional systems (like weather systems), computational complexity versus estimation accuracy and uncertainty quantification become key considerations in algorithm design and analysis. The student will look at novel methods for inference in these settings that consider these factors with an emphasis on rigorous algorithm development and performance characterisations. Expected outcomes will be methods as just detailed with software and published articles to be produced. The student will likely work closely with practitioners in application fields (like climate scientists) as well as with data scientists and mathematical statisticians and machine learners with the aim of producing practical methods that may be tested in real systems.

Skills and capabilities required for the project

The skills required for this project include basic statistics, math, and/or computer science with applications in data science, analytics and/or machine learning. Research experience at the level of a dedicated honours-type project is necessary. Experience in programming in Python and Matlab will be beneficial but more general programming skills are sufficient.

Supervisors

Edwin Bonilla (CSIRO's Data61) and Adrian Bishop (University of Technology Sydney and CSIRO’s Data61)

Location

Eveleigh and University of Technology Sydney, NSW

15. Mechanistic Machine Learning

Description

Machine learning (ML) has revolutionized many aspects of our daily lives with incredible breakthroughs in computer vision, speech analysis and natural language processing. However, modern ML techniques such as deep learning are notorious for being data hungry (requiring large amounts of labelled data to train) and difficult to interpret. In areas involving scientific discovery, such large amounts of data may not be available and, more importantly, the ability to understand and interpret what ML models learn is a must. In this project we will investigate different approaches to combining ML with mechanistic models in order to make ML methods more data efficient and provide a better understanding of the world through their mechanistic counterpart. We will focus on three different approaches involving state-space models, emulator-based approaches and (stochastic) differential equations. Applications where such approaches are not only beneficial but necessary include ecology, biosecurity, economics, and physics but we aim at addressing problems in climate modelling through our collaborations with CSIRO's Ocean and Atmosphere.

It is expected that the student, under the guidance of his university and Data61 supervisors, will develop new frameworks for learning and inference in such hybrid models that are scalable and efficient. For this purpose, we will build upon the supervisory team's strong expertise and track record in probabilistic inference, statistics and machine learning. The outcomes of this project will not only have an impact in machine learning and statistics but also have the potential to revolutionise significant areas of science such as those mentioned above where the combination of physical models with data-driven approaches is key.

The student is expected to develop the research and methods for the above problems, publish and present the corresponding outcomes at top machine learning venues (such as NeurIPS, ICML, ICLR, AISTATS) and contribute to specific applications involving our collaborations with Oceans and Atmosphere. The student will also be given the opportunity to work alongside our collaborators at The University of Warwick (UK) and EURECOM (France) who have been working on similar problems.

Skills and capabilities required for the project

- A first-class honours (or equivalent) degree in computer science, statistics or related quantitative fields
- Excellent knowledge of machine learning techniques (e.g., popular supervised and unsupervised learning methods and fundamental concepts such as generalization, regularization, overfitting)
- Expertise in high-level programming languages such as Python
- Desirable: Knowledge of probabilistic inference techniques and deep learning frameworks such as Pytorch and TensorFlow

Supervisors

Edwin Bonilla (CSIRO's Data61)

Location

Eveleigh, NSW

16. Hybrid physics-based and machine learning modelling for real-time injury risk assessment in workplaces, clinics, and sports

Description

Background: Musculoskeletal injuries and disorders (MSDs) comprise more than 70% of workplace injury claims and affect most people who enjoy sports and physical pastimes. Whilst very common and highly costly to the individual, country and businesses, MSD risk cannot yet be predicted. There are two key barriers to solving this problem: (1) body movements and external forces typically can't be non-invasively measured during activity, (2) these types of data, which have only been measured in the gait lab, haven't been correlated to injury rates (to produce a measure of injury risk). Additionally, motion capture systems are too expensive for most organisations to own and require highly trained technicians to produce suitable outputs.

Team: The Digital Human (DH) team have developed a marker-less motion capture (MMC) system prototype that can non-invasively measure human movement and predict external body loads in any environment. The La Trobe University (LU) team perform studies with healthy and injured participants during which motion capture, health and strength data are measured in the laboratory. They have hundreds of datasets of joint surgery patients and suburban athletes (with and without MSDs) from which models can be developed.

Project Description: The PhD students will develop new hybrid biomechanical physics-based and machine learning (PB-ML) models that will predict injury risk from data collected in the gait lab. Each of the students will apply the modelling process to a different domain with different motivations and outcomes but with the benefit of overlap in developed methods and tools. The models will be incorporated into the DH MMC pipeline, enabling MSD risk quantification in the native environment for each domain.

Domain 1: Manual handling in the workplace. Upper limb and back injuries are common in many workplaces that require repetitive actions such as lifting and carrying. New non-invasive strategies are required to monitor MSD risk. This project will seek to investigate the extent of workplace injuries, the aetiology of these injuries, and develop and validate PB-ML models to quantify musculoskeletal loading and injury risk.

Domain 2: High intensity activity (sports and workplace). High intensity movements are associated with increased injury risk, but they are common in sports and some workplace activities. Running, one of the most popular recreational activities worldwide, has enormous health benefits, but a high risk of lower-limb overuse injuries. At least one in four runners will sustain an injury, commonly knee injury and osteoarthritis and/or leg bone fractures. student will build a PB-ML model from the LU datasets to enable the prediction of (1) external forces and (2) the likelihood of injury from measured movement patterns for high intensity activity.

Domain 3: Joint surgery. LU have hundreds of datasets for movements by patients who have received joint surgeries and data specifying the incidence of injury or recurrent surgery after the initial surgery. For instance, more than 40% of patients who have one total knee replacement will need their other knee replaced and this comes with economic cost and negative impacts on quality of life. This student will develop a PB-ML model that can predict the likelihood of re-injury or need for re-surgery of joint surgery patients from their body movements.

In each domain the student will implement the model in the MMC pipeline and validate the ability of the MMC system for measuring the required movement parameters for predicting MSD risk "in the wild".

Skills and capabilities required for the project

Numerical and or ML predictive modelling, strong coding skills

Desirable: biomechanics/physiology knowledge

Supervisors

Simon Harrison and Raymond Cohen (CSIRO's Data61) in collaboration with Kay Crossley, Jodie McClelland and Kane Middleton (La Trobe University)

Location

Clayton, Victoria

17. Scalable structure learning using Markov chain Monte Carlo (MCMC) and variational inference (VI) algorithms.

Description

Algorithms for inferring the structure of Graphical models (e.g., Bayesian networks) from data have become increasingly popular methods for uncovering the direct and indirect influences among random variables and facilitate causal discovery. The main challenge is that in complex systems, there is a vast number of possible combinations of the network structures which strongly limits the inference scalability.

The aim of this project is to tackle the scalability problem via transforming the original discrete/continuous hybrid parameter space into a completely continuous space where scalable approximate inference via Markov chain Monte Carlo (MCMC) and variational inference (VI) techniques are feasible.

Students participating in this project will obtain a deep insight into the theory and applications of the structure learning problem as well as the state-of-the-art approximate inference tools. These tools will be modified and extended to match the requirements of the structure learning problem.

Skills and capabilities required for the project

A first-class honours (or equivalent) degree from a recognised university in Computer Science / Physical Sciences / Mathematical Sciences, with a working knowledge of machine learning principles.

Supervisors

Hadi Afshar and Edwin Bonilla (CSIRO's Data61) in collaboration with Robert Kohn (UNSW)

Location

Eveleigh, NSW

18. Quantum Causal Inference: Scalable computation with variational Bayes and quantum computing

Description

Causal modelling provides decision-makers with a tool to assess the efficacy of a proposed intervention such as resource allocation. However, in many scenarios, computing the likelihood is computationally intractable. This project will tackle this problem by developing approximate inference methodologies for models with intractable likelihood and leveraging the power of quantum computation.

Variational Bayes (VB) has proven extremely useful in inferring big and complex models with a huge number of unknown parameters. Applying VB over causal structure offers potential speed-ups necessary to explore the space of causal structures. Quantum computing can support an additional exponential speed-up by using randomised linear algebra to compute efficient updates of the model.

Students participating in this project will develop variational models to learn and approximate a causal structure that underlies the data. Furthermore, students will explore the use of quantum-inspired randomised linear algebra to compute efficient gradient update for VB-based optimisation.

Skills and capabilities required for the project

A first-class honours (or equivalent) degree from a recognised university in Computer Science / Physical Sciences / Mathematical Sciences, with a working knowledge of machine learning principles.

Supervisors

Roman Marchant and Gilad Francis (CSIRO's Data61) in collaboration with Minh-Ngoc Tran (The University of Sydney).

Location

Eveleigh NSW or Pullenvale, QLD

19. Bayesian Adaptive Trials for Structure Learning

Description

Structure learning in causal inference is a complex challenge which needs large quantities of observational data to infer causal relationships. Generally, Randomised Control Trials are used to determine causal effects in a frequentist setting, affecting large numbers of the population. However, these trials are both costly and inefficient.

Bayesian Adaptive Trials (BATs) are a novel technique for sequentially acquiring relevant information, deciding simultaneously the characteristics of individuals recruited in the trial and the associated intervention. A more efficient causal discovery process can be achieved by quantifying uncertainty in network space, and then making decisions that will trade off reduction of uncertainty through exploration but also exploit current knowledge to apply specific interventions on subjects that have higher chances of benefiting from them.

This project will further develop BATs for specific use in causal discovery. As information is gathered sequentially, the principles behind Bayesian optimisation will be leveraged to reduce uncertainty in causal effects while simultaneously achieving beneficial interventions to the population.

Skills and capabilities required for the project

A first-class honours (or equivalent) degree from a recognised university in Computer Science / Physical Sciences / Mathematical Sciences, with a working knowledge of machine learning principles.

Supervisors

Roman Marchant and Gilad Francis (CSIRO's Data61) in collaboration with Robert Kohn (UNSW).

Location

Eveleigh, NSW or Pullenvale, QLD

20. Deep learning on the edge

Description

Performing inference on edge devices via deep-learning models for data-driven tasks such as classification, semantic segmentation, and object detection finds numerous applications, e.g., in environmental sensing, digital agriculture, supply chain integrity, and advanced manufacturing. However, running typical deep-learning models on edge devices or embedded systems is challenging due to their large computational and memory requirements.

This PhD project is around developing new deep neural-network architectures with compact models that can be used to perform end-to-end inference on embedded systems/edge devices with limited memory, energy, and computational resources. Several techniques for minimizing the model complexity of deep neural networks, including quantization, sparsification, low-rank approximation, pruning, neural architecture search, and knowledge distillation, will be studied. Subsequently, the accuracy-complexity trade-offs associated with each approach and their combinations will be analysed. The insights gained from the analysis will guide the development of novel architectures for deep neural networks that enable end-to-end inference of high-level information from raw sensor data on resource-constrained embedded systems/edge devices. The new compact models will be built upon the notion of trading increased training complexity for reduced model complexity, hence will possess considerably different characteristics compared with their conventional counterparts. The performance of the new models will be evaluated using image, audio, and motion sensor data pertaining to real-world Internet of things (IoT) applications within, for example, the digital agriculture domain. Theoretical analysis of the properties of the developed algorithms will also be considered. The research outcomes will be published in top-tier computer-science conferences and journals.

Skills and capabilities required for the project

A first-class honours (or equivalent) degree in computer science, electrical engineering, mathematics, or similar areas.

Basic knowledge of linear algebra, statistics, and probabilities.

Good programming skills with Python and C/C++.

Supervisors

Reza Arablouei, Volkan Dedeoglu and Jiajun Liu (CSIRO's Data61) in collaboration with Yifan Liu (University of Adelaide).

Location

Pullenvale, QLD

21. Trustworthy Digital Twins

Description

Digital twins (DTs) are virtual representations of physical objects or systems. They often rely on sensor data to facilitate forecasting, analysis, and decision support. Extracting actionable higher-level information from sensor data warrants the data as well as the underlying models, whether process-driven or data-driven, to be trustworthy, as it has a significant impact on the level of uncertainty in the DT-based forecasts and decisions.

In this project, we will develop trust mechanisms for DTs to evaluate, quantify, and monitor the trust in the DT sensor data and the related models, and consequently the confidence in the DT-induced insights and predictions. The new mechanisms will enable end-to-end trust evaluation, covering data acquisition, communication, storage, processing, modelling, forecasting, and decision making for Industry 4.0/5.0, digital agriculture, supply chains, sustainable credentialing, future digital manufacturing, and environmental applications.

Skills and capabilities required for the project

A first-class honours (or equivalent) degree in computer science, electrical engineering, mathematics, or similar areas

Basic knowledge of linear algebra, statistics, and probabilities

Good programming skills

Supervisors

Volkan Dedeoglu and Reza Arablouei (CSIRO's Data61) in collaboration with Raja Jurdak (QUT)

Location

Pullenvale, QLD

22. Efficient Scene Understanding on Compressed Video Streams

Description

In recent years, the development of deep learning has brought significant success to the scene of understanding tasks on benchmark datasets, but often with a high computational cost. This task will be even more expensive when extended to the video sequence. To apply to real-world, especially edge applications, such as detecting animals in the wild, tracking the products on the conveyor belt, and monitoring the working environment of miners to prevent hazards, it is important to build an efficient video scene understanding system.

Existing works typically require compressed video bitstreams to be decoded into RGB frames before being processed, which requires local storage space and increases computational cost. Meanwhile, when recognizing localizing the objects in the video sequence, most existing works leverage additional networks to estimate objects' motion, while ignoring the motion information that already exists in the compressed video bitstreams. Furthermore, in some real-world application scenarios, online decoding video and then processing can hardly achieve real-time segmentation because decoding videos in the mobile terminal takes much time and computational cost.

Given the bitstream of the compressed video, the main problem is how to extract features from compressed video for scene understanding tasks. Different from RGB frames, information is stored in a compressed manner. Three data modalities i.e., the predictive frames (P-frame), Motion Vector, and Residual are employed to reconstruct the whole video based on a reference frame (I-frames). For example, Motion Vector describes the displacement of the current frame relative to I-frames, while Residuals in the P-frames maintain the RGB differences between I-frame and its reconstructed frame calculated by Motion Vectors in the P-frames after motion compensation. Simply treat each data modality as a single data bitstream to extract features cannot extract powerful features from compressed video for scene understanding tasks.

This project aims to design efficient scene understanding frameworks based on compressed video streams. To make full use of the dependency of compressed information in the video stream, a series of new network structures will be explored and designed. Efficient training methods will also be designed for the efficient scene understanding frameworks to further improve the accuracy and robustness. The scope of the project falls within DSSG's Edge AI research area.

Skills and capabilities required for the project

1) A first-class honours (or equivalent) degree in computer science, data science or electrical engineering

2) Demonstrated expert knowledge and experience in point cloud-based machine learning.

3) Demonstrated statistical analysis, and manuscript and research proposal preparation skills, including a solid track record of refereed research publications in data mining, machine learning and artificial intelligence.

Supervisors

Jiajun Liu, Olivier Salvado, Mark Hedley, Reza Arablouei (CSIRO's Data61) in collaboration with Yifan Liu (University of Adelaide)

Location

Pullenvale, QLD

23. Adversarial Behaviour Learning for Privacy Protection

Description

The large number of emerging short videos in social media apps and online platforms poses growing challenges to current techniques in video understanding and privacy-preserving techniques [1]. For example, user behaviour and preference information in social video apps can be extracted and leveraged by attackers to launch phishing and ransomware activities or targeting specific app users by recommending and posting social videos to influence election campaigns [2][3]. The existing techniques for detecting and defending privacy leakage issues in social app videos still face the following major challenges: lack of an in-depth understanding of behaviour leakage, and lack of effective defensive techniques for behaviour leakage.

Machine learning techniques, such as video analysis and scene/object detection can be leveraged by attackers to automatically mine and extract user behaviour information (e.g., facial identity, location, activity and preference). Existing methods are still insufficient in understanding how and to what extent information is leaked from large-scale videos to unintended users. More specifically, most existing privacy-preserving methods focus on classification problems. They add perturbations or noise in conventional data (e.g., images or texts) to mislead malicious learning models, thus preventing them from extracting sensitive information. However, such an approach is still vulnerable against attacks due to its limitation on differentiable models, incapable of handling emerging video behaviour leakage attacks using non-differentiable models (e.g., coordinate regression model).

In this project we propose to design and implement a foundational adversarial behaviour learning framework to understand, discover and preserve sensitive video information. Firstly, a video understanding technique is developed to capture the underlying user behaviours in short videos via multi-source information (e.g., events, background scenes, subtitles, voices, and actions) through hybrid representation learning, and then such technique is used to simulate attacks. At second stage, we will aim to develop a new adversarial generation approach by imposing human-imperceptible perturbations on videos to protect against attacks on greater variety of model structures (e.g., regression model) with non-differentiable operations. We will investigate advanced transferable adversarial learning techniques to generate effective privacy-preserving videos for unseen yet malicious video analysis, thereby improving the robustness of the underlying defensive models and raising the bar against new emerging attacks.

In this project, we expect the student to develop a prototype system that can be used to produce secure videos before releasing them to online media platforms. The scientific innovation will lead to high quality paper publication in top conference (e.g., ICCV, ECCV, CVPR) and journals (e.g., CVIU, TPAMI etc.). Meanwhile, the IPs developed in this project can also be used by other business units within CSIRO on cybersecurity and open up potential opportunities for commercialisation. The prototype will be used as a front-end to process and impose imperceptible perturbations on user uploaded videos before publishing them on social video apps. Therefore, our privacy-preserving software can be used to protect private sensitive information against malicious attacks and leakage to untrusted third parties.

[1] R. Poddar, G. Ananthanarayanan, S. Setty, S. Volos, and R. A. Popa. Visor: Privacy-preserving video analytics as a cloud service. In 29th

USENIX Security Symposium (USENIX Security 20), 2020.

[2] How video became a dangerous delivery vehicle for malware attacks. https://securityintelligence.com/articles/how-video-became-a-dangerous-delivery-vehicle-for-malware-attacks/, Aug 2019.

[3] Phishing attacks are targeting your social network accounts.

https://www.bleepingcomputer.com/news/security/phishing-attacks-are-targeting-your-so

Skills and capabilities required for the project

- A first-class honours (or equivalent) degree in computer science.

- Machine learning and computer vision background.

- Ideally with experience in video analysis

Supervisors

Xun Li and David Ahmedt (CSIRO's Data61) in collaboration with Yulei Sui and Xin Yu (University of Technology Sydney)

Location

Marsfield, NSW or University of Technology Sydney, NSW

24. Human action discovery with weakly supervised learning from videos

Description

Real-world scenarios are complicated and encapsulate such a large number of variations that it is impossible to capture them all in a particular dataset. In machine learning, weakly supervised approaches are particularly useful in handling these types of discrepancies in the data as it does not require fine-grain labelling. Action level labelling of videos is a more convenient and economical solution as compared to the frame or pixel-level annotations. Higher-level annotations also provide a cost-performance trade-off when detecting only those actions which are of particular concern. In addition to this, manual screening and analysing large, captured video data is a cumbersome task for humans and it usually requires a lot of focus and time to find actions of interest. For example, in Agriculture and Food (A&F), animal behaviour and welfare are very important for the production of good quality products. Also, in the Ocean and Atmosphere (O&A) unit tracking the activities of fishing vessels on the seas and reporting any irregular conduct is vital. In this project, we aim to develop foundations for automated video action indexing using weakly supervised machine learning approaches and to provide efficient and accurate video analytics.

Modern deep learning techniques have already surpassed human performance in several visual recognition tasks, i.e., object classification, segmentation, and video classification. However, some tasks like simultaneous multiple action detection, recognition, and classification are still outstanding problems in the machine learning community. The challenge lies in the efficient Spatio-temporal analysis of weakly labelled video data and determining important actions of a human or discovering activities that have not been previously observed. In this context, we aim to investigate and propose weakly supervised approaches that can provide an edge over fully supervised counterparts. Furthermore, our goal is to provide a baseline dataset and models that can help solve broader challenges related to the digital twin production environment.

In this project, the student will investigate and develop novel weakly supervised machine learning-based methods to detect, classify and recognize actions from large-scale videos. In this context, both vision-based data and audio-visual data will be analysed and learned in a discriminant manner. Aspirational aims include efficient video analytics for action indexing and to benchmark its performance in a resource-constrained environment. More specifically, the student will investigate 1) multiple object tracking in 2D videos, 2) multiple action detection and classification, 3) video action indexing, and 4) activity discovery of unknown actions. Large-scale analysis of video camera streams will enable us to model occluded objects and their actions in the scene more effectively.

Skills and capabilities required for the project

* Strong linear algebra and computer vision foundations
* Basic understanding of machine learning algorithms, deep learning models, and visualization
* Familiarity with 2D object detection, semantic and/or instance-level segmentation
* Strong expertise in Python and C/C++ in a Linux environment is required

* Development experience using TensorFlow and PyTorch is a must
* Distributed computing using Slurm will be preferred

Supervisors

Zeeshan Hayder (CSIRO's Data61) in collaboration with Dr Ajmal Saeed Mian (University of Western Australia), Dr Jing Zhang (ANU) and Ali Zia (CSIRO's Data61 and ANU).

Location

Canberra, ACT (CSIRO and ANU) or University of Western Australia, WA

25. Digitising the Deep Past - Machine learning, Rock art and Indigenous engagements

Description

The sandstone country of Cape York hosts one of the richest bodies of rock art in the world where spectacular galleries document the life-ways of generations of Aboriginal peoples. Globally, rock art sites have common appeal regardless of the cultural background of the viewer. More than just ‘pretty pictures’, these ‘ancient history books’ capture a wealth of knowledge about past lifeways, peoples and landscapes. Despite their priceless status, such sites are increasingly threatened by climate change, exacerbated in many regions by pollution and development pressures. There is thus an urgent need to document these sites before they are lost.

Classification of motifs in images of rock art provides a unique and challenging machine learning problem. Many sites have organically grown over millennia, resulting in layers of rock art motifs that tell a cultural story spanning hundreds or thousands of years. The motifs individually are intricate and nuanced and rely on local expert cultural knowledge to identify. Embedding this knowledge within a machine learning system has incredible archaeological, cultural and scientific implications.

One of the long-standing challenges for image description tasks is the visual semantics gap. The advances in natural language processing and computer vision techniques helped to partially address this problem. Embedding domain knowledge into the model has emerged as a promising method for improving the performance of the model. However, it remains to embed and leverage the experts' knowledge into deep learning models. In this project the student will investigate attention Graph Neural Networks to interpret the visual and contextual relationship for image captioning in a hierarchical manner. The performance of traditional methods suffers from the domain gap between natural scene images and specific rock art images. Thus, the student will explore self-supervised feature representation algorithms such as contrastive predicting coding and variational autoencoders to extract high quality node embeddings. Domain knowledge may be used to improve accuracy as well interpretability of the model. The system will be capable of categorising rock art motifs from this incredible rock art province using a taxonomy developed in collaboration with the expert Indigenous community.

It is expected that, due to the complex and nuanced nature of the rock art data, the system will advance the state-of-the-art in image recognition and the techniques developed could be applied to other complex problems and datasets, potentially from other types of sensing. The findings and outcomes of this research will be published in top computer vision conferences and will be very useful for CSIRO in multiple business units.

Skills and capabilities required for the project

- Strong undergraduate level knowledge of one or both of computer vision or machine learning

- Strong undergraduate level knowledge of one or more of topics including statistics, probability, computer science, differential equations, Fourier analysis or functional analysis

- Competent written and verbal English language abilities

- Confident Python programming and one or more ML platform such as Pytorch or Tensorflow

Supervisors

Mohamad Ali Armin and David Ahmedt (CSIRO's Data61) in collaboration with Gervase Tuxworth, Lynley Wallis and Paulo de Souza (Griffith University)

Location

Pullenvale and Griffith University, QLD

26. Grab with both hands: Dual arm robot manipulation for robot-vegetation physical interaction

Description

This project intends to develop novel state-of-the-art dexterous dual-arm manipulation capabilities for soft materials (like vegetation) in field robots. Significant achievements have been made in performing dexterous and fine manipulation using a variety of end effectors, soft grippers and hand designs in workplace environments. These approaches soon reach their limit within the context of dense vegetation where the arm has to penetrate the branches to survey a fruit bunch or perform biological sample collection.

Humans perform such acts with ease using dual-arm or bi-manual manoeuvres, like sweeping aside vegetation to “clear up” the visual focus area or stabilising the branch to pluck a fruit. Such advanced manipulation moves in unstructured, deformable environments have been an under-studied problem that this project will investigate further.

Learning methods have demonstrated progress in controlling a high degree of freedom agents and manipulating deformable objects. However, these approaches are yet to be utilised robustly for dual-arm control operations for interaction with natural environments.

This PhD project will look at developing these dual-arm manipulation capabilities in robot-vegetation interaction and will showcase the ability to pluck fruits, sense under dense canopy, and collect leaves and other biological samples for plant health monitoring.

The PhD candidate will work within a team of world-class researchers, engineers, and other PhD students to develop robotic capabilities that enable navigation and interaction in challenging natural environments like forests, caves, etc.

Skills and capabilities required for the project

1. Background in Robotics, Statistical Techniques, Machine Learning
2. Strong programming skills with expertise in C++/Python and experience in ROS/ROS2
3. Ability to work in teams

Supervisors

Tirthankar Bandyopadhyay and Brendan Tidd (CSIRO's Data61)

Location

Pullenvale, QLD or Clayton, Victoria

27. Acrobatic Robots: Multi-limbed Agile Locomotion in Unstructured Terrain

Description

Legged locomotion for robots has become very popular in the recent years, especially with the availability of multiple commercially quadruped platforms. While they are impressive, their operation is largely limited to flat or moderately uneven terrain. This is in stark contrast to multi-legged animals in nature who can effortlessly traverse extremely challenging terrain. The focus of this work would be to address the problem of traversing and climbing on discontinuous terrain where multi-limbed multi-contact locomotion is required for effective traversal. Some traditional approaches dealing with multi-contact dynamics weren't viable for real-time control of legged robots due to computational requirements for analytical solvers as well as limitations in available actuators. However, with the recent advances in actuator technology and learning based methods, some previously overlooked approaches may become viable. This project will bring together the best of traditional control dynamics and mechanism design with state of the are machine learning techniques to design and implement a truly acrobatic multi-legged robot for traversing unstructured terrain.

Skills and capabilities required for the project

The student is required to have very strong programming skills in C++/Python, advanced knowledge in control theory, mechanism design and state of the art ML toolchains and libraries.

Supervisors

Navinda Kottege, Tirthankar Bandyopadhyay (CSIRO's Data61) and Ian Manchester (University of Sydney)

Location

Pullenvale, QLD

28. Share the Surprise: Cross-Robot High Level Information Sharing for Effective Navigation

Description

Sharing knowledge about a certain environment is something humans do on a regular basis. It is common for someone to call a friend providing information about the status of the road they have just driven by: “watch for snow” or “careful with potholes”. That information is usually processed by the recipient of the call, and they adapt their driving behaviour accordingly, so there are no “surprises”. Similarly, when a driver sees a “Kangaroo” sign on the road, they know they should adapt their speed and heighten awareness, even if no kangaroo is currently present or “sensed” at all in the surroundings. This PhD project aims to apply this analogy of high-level knowledge sharing to robots. If a robot moves through or explores an environment, it can share knowledge about that environment to another robot that might be in that area later. This has direct implications on all levels of robot behaviour, from perception, to planning and control. Perception and sensing can be heightened towards certain types of events or obstacles; motion planning can be modified (e.g., speed reduction); and control parameters can be made more or less conservative depending on the information provided. This is important as most current robots use a single behaviour throughout a mission: they go from A to B with fixed speed (unless there are locally sensed obstacles) and fixed perception strategy. This is not how humans intuitively operate and this PhD research will address that limitation in the robotics state-of-the-art.

It is important to differentiate between high-level information sharing (as proposed in this project) and sharing of metric information (e.g., map-sharing) among robot agents. Map-sharing is also an extremely important topic, but has different applications and communication requirements, so it could work complementary to the objectives of this PhD proposal.

The student will study how to extract and prioritise meaningful information for past missions and how to efficiently share that with another robot agent. This will involve machine learning for environmental understanding and fundamentals of robotics to translate that information into new robotic behaviour. This will be achieved through experimentation and theoretical modelling. Publications and very practical outcomes applicable to robotics are expected.

Skills and capabilities required for the project

The student should have very strong mathematical skills and an Engineering or Computer Science background with experience in Robotics and machine learning. C++ and Python programming and knowledge of ROS are also necessary.

Supervisors

Paulo Borges, Jason Williams and Tirthankar Bandyopadhyay (CSIRO's Data61)

Location

Pullenvale, QLD

29. Self-supervised Uncertainty Prediction in Lidar Place Recognition

Description

Globally localising an agent within a map is key to autonomous robotics. Also, reliably finding the robot's current location can be used in Simultaneous Localisation and Mapping (SLAM) to close loops in revisited areas to deal with drift. Place Recognition (PR) approaches, using deep learning, show impressive results in finding the robot's location in a map database without requiring any prior knowledge. However, these PR approaches are based on heuristic supervised learning which restricts the use of the approach in different environments.

This PhD project will jointly address the weaknesses of the existing PR methods to produce a self-supervised, adaptive and reliable PR model that can learn during deployment. This will be achieved by pivoting towards PR as a self-supervised learning task that produces probabilistic predictions.

Skills and capabilities required for the project

A first-class honours (or equivalent) degree in a relevant area in the past 5 years (e.g., Robotics, Computer Science, Electrical Engineering, Mechatronics)

Strong competencies in one or more of the following areas: Robotics and Deep Learning.

Demonstrated strong programming skills in C++ or Python in Linux.

Demonstrated Experience in Pytorch and/or Tensorflow.

Supervisors

Milad Ramezani, Dimity Miller and Peyman Moghadam (CSIRO's Data61)

Location

Pullenvale, QLD

30. An apple in hand: Automating agricultural soft robot design through topology optimisation

Description

Soft and flexible robots are ideally suited to agricultural applications like fruit picking, pest management and animal husbandry. Their soft materials allow them to conform to the shape of target objects and robustly grasp them without damage.

By exploring the ample design space offered by an almost limitless combination of possible gripper geometry, material, and actuation, soft robotic end effectors can be tailored to the requirements of specific tasks. To do so, efficient simulators and optimization algorithms are essential, they must be able to find high-performing designs and accurately model the physical interactions.

This project looks to provide these capabilities for the design of soft robotic components that safely and efficiently perform on-farm tasks. Specifically, it will investigate the use of novel topology optimisation algorithms to automatically design high-performing soft robotic end-effectors for agricultural applications.

Existing soft robot topology optimization coarsely captures the physics in their optimisation models, missing critical features such as large-displacements, material nonlinearities, contact interactions, and friction.

The successful candidate is expected to extend the framework to of traditional compliant mechanism design (Sigmund) and level set topology optimization approaches (Wang) to incorporate physical features critical to agricultural soft robotic end effectors. The initial phases of the project will investigate non-linear soft robot topology optimisation and aim to efficiently solve the problem through sequential linearisation. It will then look at incorporating contact using a fictitious material method. The end goal is holistic topology optimisation algorithm for agricultural grasping, capable of transforming task requirements (workspace, object mechanical properties, deformation profile, etc) into a high-performing soft robotic tool. Its expected outcomes include prototype soft grippers and publications in high impact journals

The successful candidate will have access to multimaterial 3D printing and robotic testing facilities to prototype and evaluate designs and will work with PhD students and senior researchers within CSIRO's Robotics and Autonomous Systems Group and Monash Robotics to make an impact on robotic agriculture.

Skills and capabilities required for the project

- A first-class honours (or equivalent) degree in a field relevant to the project (mechanical/mechatronic engineering, computer science, maths, physics or similar)
- Experience with finite element modelling/multiphysics simulation (E.g., COMSOL, Ansys, Abaqus, Moose)
- Programming ability, preferably in one or more of Python, C++ and Matlab
- An overall H1 (80–100%) grade (either undergraduate or master's degree), or be in the top 5% of the applicant's graduating cohort

- Have completed a research project, component, subject, or group of subjects that accounts for at least 25% of one year's work at Honours level, or 25% of one year accumulated over the length of a masters course
- Proficient English

Desirable:

- Knowledge of design exploration/optimisation methods, especially topology optimisation.

Supervisors

Josh Pinskier and David Howard (CSIRO's Data61) in collaboration with Michael Wang and Chao Chen (Monash University)

Location

Pullenvale, QLD

31. SWAM visual SLAM

Description

3D scene understanding and 3D mapping are core topics in the field of computer vision. Specifically, current approaches to creating large scale maps typically rely on a single, or at most a few, high-end, expensive sensors that scan an area in a relatively controlled and restricted way. What this project attempts to explore is to radically reduce the cost of each individual sensor, and the platforms that carries them, and instead significantly increase the number of sensors/platforms used. For example, a swarm of 100 inexpensive drones streaming to a base a standard resolution camera feed will be able to quickly build a large area equivalent to one large high-resolution drone. In addition, the redundancy of the sensing will allow to lose many agents while still achieving the mapping task.

If successful, this has the benefit of being vastly more fault tolerant as the system may not rely on any individual sensor, but rather exploit the collective sensing capability. A system with such properties is useful in a range of environments, for example, exploration in dangerous areas, search, and rescue, and in industrial environments where regular mapping is needed, and system maintenance can be kept at a small, ongoing, non-urgent level rather than less frequently but more critical.

This project will use multiple cameras for 3D scene reconstruction complementing the approaches used by the Robotics Group based on Lidar and single high resolution sensing platform.

Scientific questions to be explored are around 1) efficient sensor fusion from sensor platforms with uncertainly in location and orientation, 2) efficient distributed algorithms, 3) addressing the “Where To Look Next?” coordination problem of the fleet.

It is expected that the student collaborates closely with members of the Robotics group as they are working on systems addressing problems in a related domain and synergies are likely. Furthermore, the student is expected to have between 2 and 4 conference publications at the end of this project. The student will develop strong written and verbal communication through presentations at internal and external conferences, workshops and meetings.

Skills and capabilities required for the project

- Strong undergraduate level knowledge of one or both of computer vision or machine learning
- Competent written and verbal English language abilities
- Confident Python programming and one or more ML platform such as Pytorch or Tensorflow
- Desired knowledge of 3D geometry

Supervisors

Lars Petersson, Olivier Salvado, Kasra Khosoussi (CSIRO's Data61) and Hongdong Li (ANU)

Location

Canberra, ACT or Pullenvale, QLD

32. Collaborative Perception for Heterogeneous Resource-Constrained Robot Teams

Description

There is a rapidly expanding multi-billion dollar market for intelligent robotics and machine perception. These technologies provide invaluable assets for addressing numerous societal challenges and industrial needs in various sectors. The last two decades have witnessed tremendous progress in single-robot machine perception. However, individual robots are fundamentally limited by their onboard mission-critical resources, narrow "field of view", capacity to perceive the world, and limited set of experiences to learn from. Therefore, we inevitably have to rely on teams of collaborating robots working together to solve the increasingly complex tasks of tomorrow.

This PhD project will focus on developing fundamental collaborative machine perception capabilities such as Collaborative Simultaneous Localization and Mapping (CSLAM) for a fleet of heterogeneous robots operating in large-scale unknown GPS-denied environments. As part of this project, new decentralized and distributed optimization methods will be designed for collaborative perception tasks in challenging time-varying communication regimes. Our work will enable robots to efficiently share and learn from self and peer experiences, while seamlessly adapting to and leveraging heterogeneity in resources and sensing modalities. The algorithms developed in this project will be implemented and tested on CSIRO Data61's fleet of ground and aerial robots.

Skills and capabilities required for the project

Must have a Bachelor’s degree with the first-Class Honours or a Master’s degree with Research in a relevant area in the past 5 years (e.g., Computer Science, Electrical Engineering, Mechatronics, Mathematics).

Strong background in Linear Algebra, Optimization, Probability Theory and Statistics.

Background in Computer Vision and/or Robotics.

Supervisors

Kasra Khosoussi, Navinda Kottege and Paulo Borges (CSIRO's Data61)

Location

Pullenvale, QLD

33. Robust and Lossless Compression for On-device Machine Learning Models

Description

On-device machine learning (ML) is rapidly gaining popularity on mobile devices. Mobile developers can use on-device ML to enable ML features on users' mobile devices, such as face recognition, augmented virtual reality, voice assistance, and medical diagnosis. Compared to cloud-based machine learning services, on-device ML is privacy-friendly, of low latency, and can work offline.

On-device ML requires models to be deployed at the local mobile devices, thereby inevitably creating a requirement of model compression and security to prevent new attack surfaces. This project aims to combine the adversarial robustness, model compression and our novel compression encryption into one task.

- Research Task I: Investigate model lossless compression algorithms and optimizations that are friendly to on-device ML, including but not limited to quantization, distillation, transfer learning, pruning, latent variable models, etc.

- Research Task II: Investigate whether and how model compression and our novel compression encryption (e.g., a single compression algorithm additionally providing encryption) can facilitate an active defence technique, i.e., the adversarial training, by relaxing the network capacity requirement.

- Research Task III: Conduct theoretical and experimental analysis for the compression performance of the proposed designs and the security against attacks on on-device ML.

Skills and capabilities required for the project

Strong programming skills (C/C++, Python), familiarity with machine learning (ML) models (PyTorch, Tensorflow), understanding of ML security, and data mining skills, along with handling large datasets.

Supervisors

Shuo Wang and Jason Xue (CSIRO's Data61)

Location

Marsfield, NSW

34. Application of Differential Privacy in Accurate Quantum (State) Computing

Description

Quantum computing can be a game-changer in fields such as, cryptography, chemistry, material science, agriculture, and pharmaceuticals once the technology is more mature. As a result, a growing number of entities race to benchmark, stabilise, and ultimately commercialise this technology. However, the 2022 revenues for QC hardware, software, and QC-as-a-service will likely be less than US$500 million. The main reason is that the promised "quantum supremacy" of QCs is yet far from a true accomplishment. Briefly, the quantum supremacy of a QC is demonstrated by its ability to quickly solve a problem that never has been solved by a classical computer, e.g., breaking TLS encryptions. Quantum supremacy requires a reliable set of measurements but in quantum mechanics, measurement is, famously, an inherently destructive process.

One promising direction is to leverage the state-of-the-art findings in stochastic control theory to achieve a new type of measurement called gentle measurements. Recently, the very interesting joint work of Scott Aaronson (Physicist) and Guy N. Rothblum (Computer Scientist) showed that the problem setting (including the threat model) of Differential privacy (DP), which has been an exceptionally successful concept when it comes to providing provable security guarantees for classical computations, is quite similar to the quantum state estimation problem. They studied this connection and concluded that the DP mechanism could provide gentle measurements. Specifically, in differential privacy (DP), we want to query a database about n users in a way that "leaks at most ε about any individual user," even conditioned on any outcome of the query. Meanwhile, in gentle measurement, we want to measure n quantum states in a way that "damages the states by at most α," even conditioned on any outcome of the measurement. In both cases, we can achieve the goal by techniques like deliberately adding noise to the outcome before returning it.

Therefore, in line with this breakthrough, the aims of this project are:

(1) to apply the state-of-the-art DP to QC algorithms and identify the perturbation mechanism that maximises QC algorithms' accuracy;

(2) to measure the privacy guarantee of such quantum computation as privacy concerns in applying quantum algorithms to individuals' datasets will be an immediate need to address right after QC commercialisation; and

(3) finally, using accurate DP (gentle) quantum measurements to bridge the gap between responsible AI and quantum computing by providing basic foundations for quantum AI algorithm which respects privacy.

Skills and capabilities required for the project

Prospective students must have:

- Graduated from Physics, Electrical Engineering and Math

- Excellent academic track record (CGPA >5.5)
- Strong background in probability theory and quantum computing.
- Strong familiarity with Python and strong coding and data handling skills
- Strong English writing and oral skills to communicate
- Strong research skills demonstrated through honours/master thesis, publication or drafts.

Supervisors

Meisam Mohammady, Dongxi Liu and Muhammad Usman (CSIRO's Data61)

Location

Marsfield, NSW

35. Physics inspired ML for critical infrastructures

Description

This project aims to develop a novel synthesis of physical and computational approaches to ML powered cyber-physical system security. Synthesis of computation, communication and control models with machine learning for practical relevance. Software plays a fundamental role in modern control (e.g., cruise control, self-driving cars), communication (e.g., 6G), and power systems (e.g., DER). However, complex software architectures used in real-world critical infrastructure are usually not captured by abstract system-theoretic models. This project will combine data from (simulation of) critical cyber-physical systems with principles and strategies derived from control and coding theories. Data-oriented machine learning methods will act as a glue between computing and system-level abstractions to fulfill cybersecurity functions such as vulnerability assessment, resilience, attack detection, and attack response. From physical system models to computational ones, system models will highlight which aspects of actual systems are vulnerable to attacks, how resilience can be improved at a system level, and which data features can be used for attack detection. In the reverse direction, computational models will inform which system parameters attackers can realistically affect and what type of system errors they can introduce.

Skills and capabilities required for the project

Strong programming skills (C/C++, Python) and computer system knowledge, familiarity with machine learning (ML) models (PyTorch, Tensorflow), understanding of ML security and codata mining skills.

Supervisors

Mohammed Bahutair and Jason Xue (CSIRO's Data61)

Location

Marsfield, NSW

36. Knowledge-Sharing among AI: Privacy-Preserving Federated Transfer Learning

Description

Recent advancements in federated learning (FL) allow multiple parties to collaboratively train to learn an artificial intelligence (AI) model without sharing raw data. However, FL is challenging in real-world applications because datasets may differ in both the sample and feature spaces. Transfer learning is an effective way to solve the difficulty of data annotation by transferring knowledge from a related source domain to the target domain. For example, a bank might train an AI model of clients' debt repayment behaviour using the knowledge in another AI model trained by a fashion retail company. Supply-chains in the AgriFood and Waste Management sectors are other highly relevant domain applications for Transfer learning. For example, prediction or classification models built for on a specific mandarin-focused chain may be re-used for some other related citrus chains. However, data privacy and confidentiality become a serious concern when cross-organizational transfer learning is conducted. Thus, in this project, we aim to address the challenges in dealing with heterogeneous data from multiple sources to perform model training safely and efficiently without violating data privacy and confidentiality. As the main tasks, the PhD student will focus on:

1. conducting a comprehensive study on the existing federated and transfer learning techniques,

2. developing algorithms/techniques with strong privacy guarantees to deal with heterogeneous feature spaces using transfer learning considering efficiency and scalability,

3. conducting extensive experiments on real-world use-cases related to CSIRO's initiatives in Trusted Agrifood Exports, and Ending Plastic Waste, such as transfer learning between different supply chains for different produces, or between different circular chains of waste/resource recovery for different regions or materials, respectively.

The expected outputs include publications in high-impact conferences/journals on data privacy, such as the International Conference on Machine Learning (ICML), IEEE Symposium on Security and Privacy (S&P), and Privacy Enhancing Technology Symposium (PETS). The project outcomes may also benefit a wide range of science fields, as well as many sectors.

Skills and capabilities required for the project

- A first-class honours (or equivalent) degree in Computer Science or relevant field.
- Programming experience in Python.
- Knowledge in artificial intelligence and machine learning is preferable.
- Participation of publications in top-rated venues such as ACM/IEEE would be a plus.

Supervisors

Ming Ding, Thilina Ranbaduge and Thierry Rakotoarivelo (CSIRO's Data61)

Location

Eveleigh, NSW

37. Privacy Enhanced AI-based Learning Analytics

Description

This project will advance academic knowledge and industry practice on how to design and implement AI-based Learning Analytics, which adheres to the principles of Responsible and Trustworthy AI. Ethical and safe use of data for Learning Analytics (LA) is poorly understood and varies between institutions due to the lack of transparent and widely agreed industry standards and practices. In the context of Privacy, one of the principles of Responsible AI, concerns of ethical use of data, breaching learners' sensitive information and the impact of proposed analytics interventions on learners' well-being are all real threats to the deployment of LA systems. This project will focus on technical approaches to orchestrate effective, ethical, and trustworthy incorporation of AI in Learning Analytics. More specifically, it aims at preserving the utility of employed AI-based LA methods in concrete industry scenarios, while providing assurance on Responsible AI principles, such as learners' Privacy, methods' Fairness and Explainability.

Data61 has been collaborating with industry and academic partners to develop privacy-only enhanced AI-based LA methods for use-cases in the Education sector. This PhD project will build on this existing work, and extend it to the study of LA methods that provide formal guarantees on multiple principles, such as Privacy, Fairness, Explainability, and study their joint interactions on data utility. This project will leverage existing relationship with industry partners to elicit real-world specific scenarios and source existing data about learners', educators, and learning programs, which will be used to validate its novel responsible AI-based LA methods.

This project will include the following milestones:

- a survey of the existing literature and deployed tools on the interactions between data privacy, method fairness and Explainability in the context of AI-based LA.
- the definition of specific LA use-cases from discussions with existing partners, focusing on their responsible and ethical aspects and their required utility in the data.
- the development of different approaches to enable such scenarios with the required constraints in terms of utility, privacy, fairness and explainability. These approaches (e.g., novel/modified Machine Learning algorithms for LA, novel data treatment) will provide strong formal and/or empirical guarantees on these constraints and will be evaluated in real-world contexts.
- a discussion on generalising these approaches beyond the specific project use-case to a broader set of LA applications, and some prototype implementations in proof-of-concept solutions in collaboration with existing partners.

Specific outcomes will include research publication in targeted journals and conferences, internal technical reports, source code for the developed approaches, and prototype demonstrations (e.g., live/recorded demo, posters).

Skills and capabilities required for the project

- A first-class honours (or equivalent) degree in Computer Science or related fields.

- Programming experience in data analytics and machine learning (e.g., R, Python, TensorFlow, PyTorch, etc).
- Some experience in privacy-enhancing, machine learning, or other related technology.

Supervisors

Thierry Rakotoarivelo (CSIRO's Data61)

Location

Eveleigh, NSW

38. Trustworthy and Resilient Distributed Learning

Description

Artificial intelligence (AI) has gained significant attention because of the achievements of machine learning (ML) and deep learning algorithms that rapidly accelerate research and transform practices in multiple fields, including health, agriculture, cybersecurity, and advanced manufacturing. Given the constraints of data sharing, distributed learning has emerged as a strategy for effective collaboration between data owners while enhancing privacy, ensuring governance, and complying with regulatory aspects. Distributed ML techniques, including federated learning and split neural networks (or split learning), enable machine learning without directly accessing raw data, which can often be personal and sensitive, held by clients such as hospitals or end devices such as the Internet of Things (IoT). Nevertheless, privacy remains to be a significant issue for distributed ML since the shared information in distributed learning can still reveal certain knowledge of the underlying data.

Generally, there are significant security and privacy challenges when deploying distributed learning techniques. First, privacy attacks, such as data/feature inversion or model inversion, pose a significant threat to data privacy even without direct access to them. Second, distributed learning techniques are inherently vulnerable to poisoning and backdoor attacks considering the facts that several participants could be malicious and the host organization has no right and means to examine the local data. Third, new privacy threats are emerging in split/vertical federated learning, such as the label interference attack. Fourth, there is still a significant gap on the empirical end-to-end evaluations of different privacy enhance techniques in distributed systems and practical implementation issues.

In this project we aim to investigate and address aforementioned challenges. To the best of our knowledge, no work has so far addressed the distributed learning problem systematically considering different attack scenarios, which we believe is important for future distributed learning applications, such as food supply chains, smart energy networks, digital health, and advanced manufacturing.

The main objectives and outcomes of the project are:

1. Explore and analyse the potential of threats for each information exchange level based on an overview of the current state-of-the-art attack mechanisms, and then discuss the possible defence methods against such threats – 1 survey paper

2. Design and develop novel distributed ML algorithms using privacy-enhancing technologies such as differential privacy against the recently discovered privacy attacks – 2-3 technical papers

3. Compare and evaluate (both theoretically and empirically) the trade-off between privacy and accuracy provided by the algorithms to evaluate their practicality under different application scenarios, such as edge computing – 1-2 technical papers

Skills and capabilities required for the project

A first-class honours (or equivalent) degree in Computer Science or relevant field.

Programming experience in Python.

Knowledge in privacy-enhancing technologies and machine learning is preferable.

Supervisors

Thilina Ranbaduge, Ming Ding and Youyang Qu (CSIRO's Data61)

Location

Eveleigh, NSW

39. Digital Agriculture AI: Privacy-Preserving Federated-ML Analytics for Trusted Supply Chains

Description

In food supply chains – capturing, interpreting, and disseminating data analytics from farmer to food-plate, is a vital digital solution for the maintenance and growth of Australian agriculture. For trusted supply chains – to enable participation, provision, and dissemination of analytics – a key driver is preserving privacy/confidentiality of participants (such as farmers, processors, and distributors) while providing improved utility for all stakeholders. Building on existing pilot work this PhD project will design for, and investigate, such trusted supply chains for a range of primary produce important to the Australian agriculture sector, such as, grains, legumes, honey, and beef. Privacy-preserving analytics will be applied to organic-compound, geochemical and lipidomic profiles of the produce, for instance; as well as privacy-preserved analytics of time-of-transaction, price, and food quantities. Inference of provenance, quality, and product verification, with transaction statistics and forecasts, will be enacted in localised subgraphs of the supply chain, close to, or at the point of data capture. Thus, reliable information can be conveyed to all stakeholders, which doesn't mitigate commercial-confidentiality or personal-privacy. Privacy-preserving federated machine learning (ML) will be applied over distributed food-supply chain ledgers (which could, e.g., be blockchain-based). Thus, in this project the student will complete the following tasks: (i) Design appropriate supervised ML of primary produce suitable to produce-type and chemical profiling; (ii) Adopt supervised ML with provably private transformations of food and transaction characteristics that maintains accuracy across the vertical partitions of the supply chain; (iii) Create suitable graph-theoretic models of the supply chain, for food and transactions, applying techniques, such as, edge differential privacy at sub-graph level, and novel node centrality measures; with federated ML over sub-graphs; (iv) Build to dynamic-representation for supply-chain time-series data, enabling near-real-time implementations; (iv) Working from smaller test-cases, update design as necessary to enable scalability and interoperability amongst supply chain participants (e.g., different commercial entities, farmers in different sectors). This work is linked to CSIRO's Trusted Agrifood Exports (TAE) Mission.

Skills and capabilities required for the project

A first-class honours (or equivalent) degree in Computer Science and/or related Engineering/Science. Skills in applied mathematics required. Programming skills such as Python languages and/or Matlab desirable.

Supervisors

David Smith (CSIRO's Data61) in collaboration with Robert Barlow (CSIRO Agriculture and Food)

Location

Eveleigh, NSW

40. Privacy Preservation in Deep Generative Networks

Description

Artificial intelligence (AI), especially deep neural networks, has transformed our life and society, e.g., digital manufacturing, personalized healthcare, smart agriculture, etc. The domain applications, such as equipment maintenance, disease diagnostics, supply chain optimization, usually require a significant amount of training data samples, which is unfortunately not always available. Hence, it is preferrable to generate more applicable data using deep learning generative networks such as generative adversarial network (GAN), variational autoencoder (VAE). However, deep learning models are prone to be vulnerable to privacy inference attacks (such as membership inference, model inversion, and model extraction attacks). These privacy inference attacks have been well studied in recent years on deep learning based discriminative models (classification models). In contrast, the privacy leakage risks from deep learning based generative models (generation models) are overlooked. For example, by attacking a GAN used in digital manufacturing or smart healthcare/agriculture, it might expose individuals'/companies' sensitive information such as production status, health conditions, goods distribution networks, etc.

In this project, we aim to systematically analyse the privacy leakage of the training data and/or the model itself from two popular generative models, i.e., generative adversarial network (GAN) and variational autoencoder (VAE). Various assumptions will be considered, such as whether it is explainable (black-box or white-box), which granularity of auxiliary information the attacker has access to, etc. Based on the analytical results, we then work to develop a privacy-preserved deep generative model, which mitigates the privacy inference attacks. Besides, there is a potential to implement federated generative models to enlarge the size of training data samples by accommodating deep generative models into federated learning paradigm. Upon completion of this project, we aim to deliver a trustworthy and privacy-preserved GAN/VAE network and may extend to other application domains such as digital manufacturing. This will fundamentally solve the critical issue of lack of data in the mentioned application domains.

The candidate PhD student will be involved in the tasks of i) investigation of a general attack model targeting deep learning generative networks, ii) analysis of the privacy leakage of the training data and/or the model, and iii) study of new privacy-preserving deep learning generative models.

Skills and capabilities required for the project

A first-class honours (or equivalent) degree in computer science, electrical engineering, or related area.

Programming experience in Python, MATLAB, etc.

Knowledge in deep learning and/or data privacy

Supervisors

Youyang Qu and Ming Ding (CSIRO's Data61)

Location

Eveleigh, NSW

41. Responsible Data Lifecycle Management and Analytics

Description

The use of Automated Decision Systems (ADS) is increasing at a rapid pace in both Government and non-Government contexts. For instance, the Centrelink established Debt Program in July 2016 which relies on AI and algorithmic techniques for raising and recovering debts. Automated commercial tools are increasingly used by companies in the hiring and employment process.

The importance of incorporating responsibility into ADS is broadly recognized to guarantee fairness, accountability, interpretability and transparency in the socio-legal-technical systems. In the EU and Australia, the General Data Protection Regulation (GDPR) and the Australian Human Rights Commission are introduced in recently years. In research communities, significant efforts have been dedicated to the responsible development of ADS. Nevertheless, existing efforts often focus on the last mile of data analytics, namely, on ensuring responsibility in model designing and deployment. The key question is how to embody trust and responsibility into the holistic lifecycle of data management and analytics in ADS? In this project, we aim to answer this question and identify critical opportunities of improving data representativeness and controlling bias.

Activity 1: Literature review on the existing research and industry approach: Related literatures in the areas of decision-making automation pipeline, statistic analytics, DB operation and optimization, AI/ML will be reviewed. In addition, we will investigate the challenges from the existing ADS.

Activity 2: Representativeness-enhancing Data Acquisition and Pre-processing: Data used for analytics tasks is often acquired for a different purpose and does not represent the true distributions. Low absolute representation of minority groups in data leads to potential substantial bias in the ADS system. The key challenge here is to make the data-acquisition task aware, setting coverage objectives based on performance requirements of downstream analytics tasks rather than enforcing a simple global threshold as in existing works. Furthermore, technical bias can be introduced at any stage of the ADS lifecycle especially in pre-processing such as missing value imputation, selection, join or ranking. Sophisticated methods will be developed to truly reflect distribution in pre-processing steps by incorporating domain knowledge.

Activity 3: Fairness-aware Data Analytics: In this activity, we aim to enhance responsibility of ADS by considering associational fairness and causal fairness when conducting data analytics tasks. In associational fairness, we will investigate fairness measures based on data alone, such as conditional statistical parity and predictive value parity, which pose consistency constraints on different groups to avoid spurious correlations. Causal fairness captures context knowledge as causal relationship represented as causal DAGs. Correlations are measured by inference of particular entities along causal paths that are considered to be socially or legally unacceptable. The fairness-enhancing strategies above will be investigated and integrated in the ADS for responsible data analytics processing.

Activity 4: Establish a generic framework: Will investigate how to incorporate the above innovated approach with the existing the ADS and tested using real-world cases. Will design a generic framework to support Responsible Data Lifecycle Management and Analytics.

Expected Outcomes: The scientific outcome will develop novel techniques that enable a trusted data analytics pipelines in which representativeness-enhancing interventions are integrated into a fair ADS lifecycle. The methods will be applied to solve the responsible challenges of current ADS and published in peer reviewed high-quality journals and international conferences.

Skills and capabilities required for the project

- Demonstrated capability to carry out research work independently including literature search, concept development, experimental work planning and implementation
- Good knowledge in database systems, machine learning, algorithms, and graph theory
- Good programming skills Python and Java
- Good communication and writing skills and capable to prepare oral presentation and written report, articles
- Good interpersonal skills and a team player, hardworking and willing to strive for excellence

Supervisors

Sherry Xu, Qing Liu (CSIRO's Data61) in collaboration with Wenjie Zhang (UNSW)

Location

Eveleigh, NSW

42. Diversity and Inclusion in designing Crisis Management Apps

Description

Australia is currently facing a huge challenge in addressing natural disasters such as bushfires, floods, and the ongoing pandemic. These events have high social, financial, environmental and human costs. Australian Government and organisations are increasingly relying on investing in technological solutions to better prepare for unprecedented crisis events. It is predicted that by 2025 almost 21.5 million Australians will have a smartphone. The diversity of utilities, simplicity of use, ease of access, personalisation, ubiquity and flexibility of mobile technologies and apps make them a valuable tool in current times. However, these apps must work within a socio-technical ecosystem during a crisis event where users' diversity and their social attitude toward the technology play almost as critical a part as the technical excellence and effectiveness of the app itself.

Developing crisis-driven mobile apps is challenging for various reasons: 1) technical demographics of the target users for devices/platforms, 2) time urgency, 3) uncertainty, 4) quality of the app, and 5) other non-functional requirements including data privacy, security, and user trust. One of the least explored dimensions of target users during the development of these apps is diversity and inclusion. Crisis events impact all citizens and hence the technology has to offer an inclusive design to address the diversity requirements in the resulting products.

The overall aim of this project is to address the gaps in the existing body of research and practice by devising an effective framework that would incorporate predictive analysis of the design features within a socio-technical context, and diversity of the target users to make the apps more inclusive, and easily adaptable depending on the nature of the crisis. This will be achieved via the following activities:

(1) Data collection and analysis – chronological data analysis of existing apps for crisis events in the past to determine human-centred success/failure factors for these apps;

(2) Dataset creation – create datasets that will be used by AI-powered tools to perform predictive modelling for the optimal design of features for future mobile apps that are based on the analysis of historical results from step 1;

(3) Methodology development – design a methodology for eliciting and specifying the diversity and inclusion requirements for crisis-driven mobile apps;

(4) Design and evaluate a framework for crisis-driven mobile apps powered by AI systems

Skills and capabilities required for the project

- a first-class honours (or equivalent) degree in Computer Science or related discipline with outstanding grades from the top universities
- shows evidence of research ability – e.g., an Honours thesis, a minor thesis, publication(s), and a well-developed preliminary research proposal, developed under the guidance of the supervisor

- has a strong interest in one or more of the following areas: software engineering, AI, mobile apps
- has strong programming and analytical skills
- has proficient English skills, particularly in reading, writing and speaking

Supervisors

Didar Zowghi (CSIRO's Data61) in collaboration with Chetan Arora (Deakin University) and Muneera Bano

Location

Eveleigh, NSW or Clayton, Victoria

43. Explainable Comprehensive Software Vulnerability Prediction and Protection through Diversified Software Vulnerability Knowledge Graph

Description

With sheer amount of software vulnerabilities and high dependency on third-party libraries, traditional rule-based or human-forced vulnerability detection approaches are challenged by heterogeneous data resources, complex library dependencies and conflicts, limited vulnerability support, and heavy human workloads that cannot ensure the security of complex cyber systems. The development of machine learning has improved automation and effectiveness of vulnerability detection to a new level. However, the black-box machine learning suffers from high false positive rate with unexplainable and unreliable detection results. Security experts still need to manually find key clues of the vulnerability to validate AI outcomes.

Knowledge graph opens a new door to solve the problems. Its structure can efficiently integrate heterogeneous resources from different databases, supporting automatic knowledge induction, threat modelling, knowledge retrieval and risk analysis. Besides, knowledge graph provides direct visualization to clients, offering explainable information to AI decisions that improve trustworthiness and responsibility of AI outcomes.

We explore following topics:

1. Multi-faceted vulnerability knowledge searching: Current vulnerability database only support rough word-based vulnerability searching and can only show single-faceted knowledge. By entity extraction and data integration, we can construct comprehensive vulnerability knowledge graph containing finer-grained vulnerability key aspects supporting multi-faceted vulnerability knowledge searching that can reveal inner relation among heterogeneous vulnerability data resources and support complex knowledge induction.

2. Software supply chain vulnerability portrait and detection: With attack of Log4shell to software supply chains, the importance of supply chains has been enhanced to national strategy. U.S. president Biden has announced an executive order to highlight necessity of robust supply chain security system against international terrorist. By combining knowledge from Maven, GitHub and MITRE CVE, we can construct vulnerability knowledge containing module reference knowledge, supporting explainable and traceable software supply chain vulnerability portrait and detection.

3. Dependency conflict cognitive software supply chain protection: By combining software supply chain vulnerability knowledge graph with library dependency conflict detection techniques, it is possible to transparently recommend exchangeable up-stream libraries considering both security factors and deployment factors.

4. Zero-day vulnerability collection and integration: Current security advisories suffer from long validation processes, causing delayed vulnerability publication. Platforms including Twitter and security forums of different products usually public zero-day vulnerabilities at the earliest time. However, their data are disorder textual descriptions scattered in different locations that cannot be used by AI for downstream tasks. By NLP techniques, it is possible to extract and clean disorder data into regular knowledge that can be combined with KG and downstream tasks.

5. Malware and APT cognitive threat modelling: Real cyber-attack can be complex with multiple exploiting steps in a very long period with different hacker groups, and finally cause severe damages. By combining vulnerability knowledge graph with MITRE's ATT&CK knowledge graph, we can integrate vulnerability basic knowledge with real exploiting knowledge together to support finer-grained threat modelling and warning. By historical data, it is possible to induct the attack intention, goals and hacker groups combining AI techniques, which provide better protection to cyber systems.

Skills and capabilities required for the project

Has strong interest and motivation in researching and learning new things.

Has strong programming skills in Python.

Has solid understanding of ML/DL, NLP and Knowledge Graph techniques, or shows the potential of understanding these techniques.

Supervisors

Jiamou Sun and Zhenchang Xing (CSIRO's Data61)

Location

Eveleigh, NSW

44. Analysing Security of Closed-source Unmanned Aerial Vehicle Firmware: Vulnerability Detection, Repair, and Simulation

Description

Unmanned Aerial Vehicles (UAVs) or drones, are flying mini robots and have been widely deployed in various scenarios (e.g., surveillance, delivery) because of their features of pilotless, a small physical shape, and fast speed. As more and more critical missions are assigned, the requirements of their reliability and security are increased, towards completing missions successfully. Therefore, flight code of UAVs become one of the most essential phases to determine whether a mission can be completed successfully.

When implementing the flight code, developers not only design the customized flight code, but also invoke the APIs provided in the Dronekit, a development kit containing APIs developed by third parties to support advanced use cases including computer vision, path planning, 3D modelling, etc. Although the Dronekit provides a lot of convenience for developers, it introduces the potential maintenance issue, which can cause security problems. When a vulnerability is reported to a third party, it will fix the issue and publish the updated version immediately. However, for the developers utilizing the APIs provided by the third party, they might not be notified, and the vulnerable versions are still being utilized until the problem is identified. Since 2017, some efforts in detecting security issues of flight control programs are proposed, but none of them help developers fix the problem.

This project aims to address the security issues of UAVs from two perspectives: Dronekit vulnerability repair and closed-source emulation. By analysing the flight code of a drone, our proposed approach identifies the invoked APIs in the flight code and further checks whether the API versions match the newest version of the publisher. If it is not the newest version, the approach then identifies how the third-party repairs the issue and further applies the fix code correspondingly. Since the code logic of different programs are different, the approach customizes the repaired code after learning the code dependencies of each flight code. Different from the other devices, it is hard to verify the experiment result on the real drones, which might cause drone trajectory deviation or even drone crash. Hence an emulator to simulate the flight status is necessary. Unfortunately, existing emulators can only simulate the status after revising the flight code to match the interface provided by the emulators. Such a code modification might change the data processing logic of the original flight code. Therefore, our proposed approach will build an emulator targeting on simulating the compiled flight code that are well-established.

Two promising students are expected to be recruited. One will be involved in developing the emulator targeting to simulate the flight code in compiled version. Another will be involved to analyse the security problems in Dronekit and propose the corresponding repair strategies. The expected outcomes are 1-2 core A* conference/journal papers and 2-3 core A/A* conference/journal papers.

Skills and capabilities required for the project

Binary Analysis, Protocol Analysis, Software Programming

Supervisors

Zhi Zhang, Surya Nepal (CSIRO's Data61) in collaboration with Siqi Ma (UNSW)

Location

Marsfield, NSW

45. Privacy Attacks and Defences in Cross-cyber physical domains

Description

As data exists in different modalities in the real world, viable interactions and combinations among multimodal data feature the creation and discernment of multimodal information in deep learning research. However, pre-trained large multimodal models can often carry more information (because they have much stronger data representation ability) than single modal models and they are usually applied in sensitive scenarios such as medical report generation and disease identification. Thus, multimodal models may lead to severe data privacy problems, for example, the attacker can recover the patient information from pre-trained models. This project studies the privacy leakage of large-scale pre-trained multimodal models through the lens of membership inference attack, the process of determining whether a data record belongs to the training dataset of a model or not.

Skills and capabilities required for the project

Strong programming skills (C/C++, Python) and computer system knowledge, familiarity with machine learning (ML) models (PyTorch, Tensorflow), understanding of ML security and codata mining skills.

Supervisors

Jason Xue (CSIRO's Data61)

Location

Marsfield, NSW

46. Knowledge-driven Data Integration for Causal Analysis

Description

Causality plays an integral role in all forms of decision making, particularly it is important for responsible decision making by machine learning systems. Whatever we consider the potential effects of our decisions, we are thinking about cause. Causal inference goes beyond making associations and observations about what happens in our world. We can also reason about interventions ("what will happen if something is changed?") and counterfactuals ("what would have happened if something were different?"). The importance of causal inference for making informed decisions has long been recognised in health, medicine, social sciences and other domains. The availability of "big data" in today's world further presents opportunities to unleash the power of causal analysis to transform decision-making systems.

However, in many cases, it is not clear what data should be used for analysis, let alone how they are suitable. For example, when analysing a natural or social event where the exact cause is unknown, scientists typically start from some hypothesis, then search for data that may help verify the hypothesis; when other possibilities arise, they repeat the process. Obviously, in cases like this, the data to be used are uncertain and bounded by individuals' knowledge and understanding of events. There are chances that critical data are omitted and thus biased or even erroneous conclusions are reached. Also, data may come from different sources, with different formats (e.g., texts, images and relational tables). How to handle such data with existing causal frameworks is a challenge as they typically assume that data are homogenous and can be represented in a single flat table.

This project aims to address these two challenges and support causal analysis in heterogeneous and dynamic settings through a knowledge-driven causal framework. The student is expected to investigate (1) the role of knowledge in causal inference and the use of knowledge graphs for guiding the data discovery process, (2) multimodal data integration for causal analysis, (3) representation of causal background knowledge and assumptions in heterogeneous settings, and (4) answering complex causal queries in heterogeneous and dynamic settings.

Skills and capabilities required for the project

- A first-class honours (or equivalent) degree in Computer Science or related areas.
- Solid understanding of database and machine learning techniques.
- Good programming skills in one or more languages: Python, Java or C/C++.

Supervisors

Yanfeng Shu and Chen Wang (CSIRO's Data61)

Location

All sites including Sydney, Hobart, Melbourne, Canberra, Adelaide, and Brisbane

47. Uncertainty-guided Lifelong Machine Learning

Description

Uncertainty modelling describes what a machine learning (ML) model does not (and does) know, and measures goodness of prediction. Representing and estimating ML model uncertainty is of crucial importance for general robust ML systems and real-world ML applications in the fields such as manufacturing, biology, and scientific studies. Especially for the deep learning (DL) model with powerful representation abilities and less explainability, uncertainty modelling is more critical for making DL trustworthy and effective. On the other hand, for further pushing forward the applications of ML/DL systems in the real world with dynamic requirements, lifelong learning (LL), i.e., continual learning (CL), attracts tremendous attention to incrementally learn and update the DL models from the online data and task streams. However, DL models suffer from catastrophic forgetting issues, where learning on the new data (in the stream) quickly overwrites the learned parameters (on old data). This project will study novel uncertainty estimation methods for DL in the LL system and tackle the forgetting issue with guidance from uncertainty.

The lifelong learning system will be guided by uncertainty modelling in the DL model parameters and structures, i.e., the model (epistemic) uncertainty with the awareness of aleatoric uncertainty from data noise. The uncertainty can be modelled via data density estimation and Bayesian probabilistic modelling on deep neural network parameters. The DL models can identify novel concepts to learn in the streaming data based on the uncertainties reflected by the data density estimation. As uncertainty naturally describes what a DL model does not (and does) know, the uncertainty helps to identify what is in parameters to remember and what to change, alleviating the forgetting issue. Specifically, the probabilistic modelling and updating of the uncertainty models on both data and parameters will consider the temporal correlations in the learning stream and the relationships among the tasks.

Students participating in this project will develop novel models with efficient estimation methods for describing and predicting the uncertainties for data and DL models, considering the scalability and efficiency in real-world applications. The students will then design the uncertainty-guided lifelong/continual learning system in various application scenarios.

Skills and capabilities required for the project

- A first-class honours (or equivalent) degree from a well-recognised university in Computer Science, with a working knowledge of machine learning principles.

Supervisors

Sally Cripps (CSIRO's Data61) in collaboration with Dong Gong (UNSW) and Lina Yao.

Location

Eveleigh, NSW

48. Towards Edge AI: Efficient Deep Learning for Resource-constrained Edge Devices

Description

Machine learning at the edge (ML@Edge) aims to bring the capability of running ML models locally to edge devices. It is important for many scenarios where raw data is collected from sources far from the cloud. However, Deep neural networks (DNNs) are both computationally and memory intensive, making them difficult to deploy on embedded systems with limited hardware resources. This project will tackle the problems by developing efficient deep learning methods, i.e., pruning, factorization, quantization as well as compact model design.

To step further, we will enable our model to support diverse architectural settings by decoupling training and search. We can quickly get a specialized sub-network by selecting from large, generalized network without additional training.

The implementation of machine learning models in edge AI will decrease the latency rate and improve the network bandwidth. Edge AI helps applications that rely on real-time data processing by assisting with data, learning models, and inference. The success of this project brings significant benefits to IoT edge computing, cyber-physical systems, smart factories, autonomous vehicles, smart healthcare, etc.

Students participating in this project will develop efficient deep learning models for devices with limited resources. Furthermore, the students will be able to employ the proposed framework to deploy state-of-the-art AI models to edge computing devices.

Skills and capabilities required for the project

- A first-class honours (or equivalent) degree from a well-recognised university in Computer Science, with a working knowledge of machine learning principles.

Supervisors

Sally Cripps (CSIRO's Data61) in collaboration with Xiaojun Chang (University of Technology Sydney) and Lina Yao.

Location

Eveleigh NSW

49. Scalable structure learning for graphical models

Description

Graphical models are an important component in explainable Machine Learning. They are widely used to understand the interactions, even causal relationships, among complex systems in many real-world scenarios. They are supporting tools to make robust and interpretable decisions.

Structure learning of graphical model is an essential step to support decision-making. However, the scalability of existing structure learning algorithms has hindered applicability of graphical models in practice. Computational burden grows super-exponentially with respect to the complexity of the system, i.e., the number of variables in the graphical models. Theory has shown that incapability to include enough variables in the graphical models may lead to misleading conclusions about the interactions of the complex system of interest.

In this project, we aim to design scalable and efficient algorithms of structure learning for large-scale datasets. The problem will be tackled by using a "divide-and-conquer" strategy, which is common and effective for scalable algorithm design. The success of this project enables us to conduct structure learning for systems which include a large number of variables. This will help us understand the large complex systems, and more importantly, provide us more accurate dependency/causal structure than using small dataset. The students participating in this project will learn and master state-of-the-art structure learning algorithm for graphical models, and then develop new algorithms for structure learning.

Skills and capabilities required for the project

- A first-class honours (or equivalent) degree from a recognised university in Computer Science / Physical Sciences / Mathematical Sciences, with a working knowledge of machine learning principles.

Supervisors

Andy Zhu (CSIRO's Data61) in collaboration with Robert Kohn (UNSW) and Lina Yao.

Location

Eveleigh NSW

50. Structure Learning via Sampling from a Manifold

Description

In many applications for Bayesian inference, the probability density is distributed on a manifold that is embedded in an ambient space of higher dimension. Such distributions are usually caused by imposing constraints on the parameters. Only very recently, effective inference algorithms have been proposed for this setting and designing more robust inference tools is an open problem.

Another important and challenging problem in Machine Learning and Statistics is Structure learning i.e., estimating the topology of directed acyclic graphs (DAGs) a.k.a. Bayesian networks.

This PhD project links these two problems via converting the structure learning setting into Probabilistic inference on a continuous space. Due to the acyclicity constraint, the target probability density is distributed on a manifold and the inference is carried out by sampling from it.

The students participating in this project should have a reasonably strong mathematical background and throughout their PhD, they will obtain a deep insight into the structure learning problem and causal inference as well as the theory of Markov chain Monte Carlo and sampling from a manifold.

Skills and capabilities required for the project

- A first-class honours (or equivalent) degree from a recognised university in Computer Science / Physical Sciences / Mathematical Sciences, with a working knowledge of machine learning principles.

Supervisors

Hadi Afshar and Yanan Fan (CSIRO's Data61) with Minh Ngoc Tran (University of Sydney)

Location

Eveleigh NSW

51. The R-index: quantifying standards for reproducible research

Description

Reproducibility and replicability (R&R) in scientific research are essential to validate results and confirm new knowledge. They are even more important when these results inform policy, future scientific studies, or health-based decisions. Their importance is global, and over recent years many national scientific bodies* and coalitions of learned institutions have produced expert reports, statements, and findings about what makes and defines reproducible research (in which the outcome of a study can be exactly repeated) and replicable research (in which a different study can validate the outcomes of earlier work).

So, for a given study, how can one quickly understand the reproducibility or replicability of its findings? This question can be answered, but it will involve a painstaking process of assessing the various R&R components, including e.g., availability of data, the process of data cleaning, code for analysis, random number seeds, the availability of the computational framework, the clarity of experimental settings, the level of detail in a proof, and so on, and this assessment will necessarily be presented in a long and detailed description.

This PhD research will build the framework and structure from which an R-index can be constructed across research in Science, Engineering and Medicine. This work will be of global importance and potentially extremely significant impact.

* – e.g., "Reproducibility and Replicability in Science" (2019). National Academies of Sciences, Engineering and Medicine. https://nap.nationalacademies.org/catalog/25303/reproducibility-and-replicability-in-science

Skills and capabilities required for the project

- A first-class honours (or equivalent) degree from a well-recognised university in a STEM discipline with strong quantitative background

Supervisors

Yanan Fan (CSIRO's Data61) with Scott Sisson (UNSW).

Location

Eveleigh NSW

52. Bayesian inference on short texts

Description

Short texts in the form of surveys are now ubiquitous, these are easy and cheap to collect, and often come with meta-data in the form of individual level covariates. Examples include restaurant and hotel reviews, banking customer satisfaction surveys. For companies to make use of these surveys, they must understand the important drivers of satisfaction: these will come both from the meta-information, as well as text responses.

This project will consider ways in which expert knowledge can be included to guide the creation of topics groups in short texts; the incorporation of meta information; and the related issues of Bayesian model choice. For large scale problems, development of new computational methods will be required, particularly in cases for when new information becomes available. This project aims to address these two challenges and support causal analysis in heterogeneous and dynamic settings through a knowledge-driven causal framework. The student is expected to investigate (1) the role of knowledge in causal inference and the use of knowledge graphs for guiding the data discovery process, (2) multimodal data integration for causal analysis, (3) representation of causal background knowledge and assumptions in heterogeneous settings, and (4) answering complex causal queries in heterogeneous and dynamic settings.

Skills and capabilities required for the project

- A first-class honours (or equivalent) degree from a well-recognised university in Statistics or Computer Science with strong mathematical background

Supervisors

He Zhao and Yanan Fan (CSIRO's Data61) in collaboration with Tom Stindl (UNSW)

Location

Eveleigh NSW