PHD PROJECTS 2021

Project title:
Neural memory for life-long agents
Supervisor Name and E-Mail:
A/Prof. Truyen Tran, truyen.tran@deakin.edu.au Dr Hung Le, lethai@deakin.edu.au
Project Description:
Current machines are not persistently intelligent. Computer scientists are able to create smartagents that defeat human on certain tasks such as games, translation and visual recognition,yet fall short of maintaining that intelligence over time and across domains. That shortcoming is attributed to the lack of efficient memory mechanisms that hasten the learningof the temporal task at hand whilst mitigating catastrophic forgetting over tasks. To equip agents with such a life-long memory, it is critical to study an existing successful memory model—the human brain. In human brain, the hippocampus plays an outstanding role in various memory and reasoning processes and thus, may hold the key to approach life-long intelligence capability. This project focused on building deep learning memory mechanisms for life-long agents. The ultimate goal is retaining adequate performance of reinforcement learning (RL) agents in coping with a stream of incoming temporal tasks (100-10,000 tasks). An example of fulfilling that goal is learning to play all Atari games in a continuous manner with a minimal number of playing experiences. The project marries reinforcement learning with transfer/continual learning. It is known that humans, when armed with the hippocampus, learn to master at this type of tasks. A hippocampal memory system for RL agents is therefore definitely beneficial. In light of hippocampus memory design, two prerequisites are wide-bandwidth relational working memory and long-term modular episodic memory mechanisms. The former is responsible for quickly acquiring decision-making skills from the current observations whilstthe latter manages the memorisation and consolidation of these skills into a form of representations that is persist and transferable to novel tasks. We aim to design a memory system that consists of the two memory mechanisms and a unified online learning algorithmthat facilitates their operations harmoniously.

Project title:
Verification of Machine Learning Algorithms
Supervisor Name and E-Mail:
A/Prof Sunil Gupta, sunil.gupta@deakin.edu.au
Project Description:

Machine learning (ML) algorithms are being increasingly used to automate decision makingaffecting our daily lives. A growing discomfort that is arising in this new era of ML algorithms is whether these algorithms actually behave as intended, that is, conform to theirdesign specifications.

A unique aspect of ML algorithms is that they are often trained using opportunistically available data. For example, consider an ML algorithm that is trained using Electronic Medical Records (EMR) data for healthcare risk prediction. EMR data is routinely collected at hospitals mainly for billing and administrative purposes. It records patients' demographic information along with a history of any diagnoses, medical procedure and pathology. This information is quite effective for developing prognostic models for predicting the risk of a disease in future. However, since the data are mostly collected from only a part of the worldwith a different purpose (e.g. billing), it may suffer from a bias caused by the incompleteness or mismatch in distribution. There may be other countless reasons for an ML algorithm to not behave as intended e.g. improper training, noisy data, mismatch in training and deployment scenario etc. It is, therefore, impractical to verify ML algorithms against all possible misbehaviours. Also, verifying the decision making of an ML algorithm requires itscomparison with a gold standard, which in turn requires querying the gold standard for truelabels. Obtaining true labels are usually expensive. Thus an important question arises: howto verify an ML algorithm effectively in a data-efficient manner?

This PhD will investigate into verifying ML algorithms in different learning paradigms e.g. verifying a supervised learning algorithm such as a classification or a regression algorithm,verifying unsupervised learning algorithms such as verifying a clustering algorithm, and verifying a reinforcement learning algorithm. Our initial work in this area has used Bayesian optimisation [1] as the first solution where an ML algorithm-under-test is verified against a gold standard by querying the gold standard at inputs suggested by Bayesian optimisation [2]. Bayesian optimisation is an efficient approach for this problem because it is one of the best method for finding global optimum of expensive, black-box functions.

The essential background for this position includes a solid foundation of fundamental concepts in machine learning, statistics, mathematical optimisation and linear algebra along with provable programming skill in Python or MATLAB.

References:

1. Shahriari, Bobak, Kevin Swersky, Ziyu Wang, Ryan P. Adams, and Nando DeFreitas. "Taking the human out of the loop: A review of Bayesian optimization." *Proceedings of the IEEE* 104, no. 1 (2015): 148-175.

2. Gopakumar, Shivapratap, Sunil Gupta, Santu Rana, Vu Nguyen, and Svetha Venkatesh. "Algorithmic assurance: An active approach to algorithmic testingusing Bayesian optimisation." In *Advances in Neural Information Processing Systems*,

Project title:

Securing AI systems from Trojan Attacks

Supervisor Name and E-Mail:

A/Prof Santu Rana, santu.rana@deakin.edu.au

Project Description:

Artificial Intelligent (AI) systems, especially the deep learning variants, have achieved exceptional performance in many complex tasks. However, with success, came the threats as well. The black-box nature of the deep learning models makes them particularly vulnerable to malicious attacks. Specifically, it has been recently shown that it is extremelyeasy to infect a model with Trojan viruses such that its behaviour can be controlled by simplyadding a small, and seemingly benign alteration to the incoming data.

For example, a face-based authentication system can be infected such that when it encounters faces with a particular form of tattoo, it grants access even to the strangers, or an autonomouscar can be made to wrongly classify a stop sign to be a high-speed limit sign by pasting a small, inconspicuous looking sticker on the signboard. Since, the response of the infected model remains unchanged for clean samples, the Trojan remains undetected until it is purposefully triggered to cause a harmful incident.

In this project, we will be aiming to build necessary tools and technology to detect such contaminated models. Specifically, we will be looking at the intrinsic properties of the deeplearning models to identify if there exist any patterns that makes the infected model behave differently than the benign ones. Knowledge in deep learning and optimisation will be crucial. The area of AI safety is an emerging field and there is a scope of dynamically adjusting the project as we go.

References

Wang, B., Yao, Y., Shan, S., Li, H., Viswanath, B., Zheng, H. and Zhao, B.Y., 2019, May. Neural cleanse: Identifying and mitigating backdoor attacks in neural networks. In 2019 IEEE Symposium on Security and Privacy (SP) (pp. 707-723). IEEE.

Project title:
Computational models for drug repurposing
Supervisor Name and E-Mail:
Dr Thin Nguyen, thin.nguyen@deakin.edu.au A/Prof Truyen Tran, truyen.tran@deakin.edu.au
Project Description:
It costs about 2.6 billion US dollars to develop a new drug, and can take up to 17 years for FDA approval. Finding new uses for already approved drugs avoids the expensive and lengthy process of drug development. E.g., currently, nearly 70 existing FDA-approved drugs are investigated if they can be repurposed to treat COVID-19. In order to repurpose drugs effectively, it is useful to know which proteins are targeted by which drugs. High- throughput screening experiments are used to examine the affinity of a drug toward its targets; however, these experiments are costly and time-consuming, and an exhaustive searchis infeasible because there are millions of drug-like compounds and hundreds of potential targets. As such, there is a strong motivation to build computational models that can estimatethe interaction strength of new drug–target pairs based on previous drug–target experiments.In this project, we aim to develop computational models for treatment discovery, including those that depict drugs and proteins as molecular graphs and adapt graph neural networks. These models (1) estimate the interaction strength of drug–target pairs to expedite drug repurposing and (2) predict drug-cell response to estimate the efficacy of drugs.

Project title:
Sample-efficient training for real-world reinforcement learning agents
Supervisor Name and E-Mail:
Dr Thommen Karimpanal George, thommen.karimpanalgeorge@deakin.edu.au A/Prof Santu Rana, santu.rana@deakin.edu.au
Project Description:
The primary focus of this project will be to develop and implement novel and scalable algorithms for the sample efficient training of reinforcement learning (RL) agents. The project should be undertaken with the aim of making significant improvements to the state-of-the-art, ultimately resulting in high quality conference publications. The approaches thatmay be explored include, but are not limited to: 1. Safe RL 2. Model-based RL 3. Meta-learning/Transfer learning in RL 4. Representation learning 5. Incorporation of prior knowledge in RL Candidates should be comfortable with python/MATLAB programming, and should preferably be familiar with RL literature to a fair extent. Successful candidates will be provided with the option of validating their work through implementation on a variety of robotics platforms, which will be available for use. As the project does not involve coursework, successful candidates will be expected to commence research immediately, and complete the project within a period of approximately 3 years.

Project title:
Machines that learn to talk about what they see
Supervisor Name and E-Mail:
A/Prof Truyen Tran, truyen.tran@deakin.edu.au Dr Kien Do, k.do@deakin.edu.au
Project Description:
Modern AI has achieved ground breaking successes. However, many problems that are relatively easy for human but have proved to be very challenging for machine. One of suchproblems is building a machine that learns to reason on a dynamic scene and talks about it.As it turns out, the main workhorse of modern AI – deep learning – while effective in learning one-step mapping from an input to an output in a fixed form, is limited when facing new forms and multi-step inference, for example, answering unfamiliar compositional questions. This project focused on learning to reason visually. We aim to design and improve an artificial mental faculty that produces new knowledge from the previously acquired visual knowledge in response to a natural question. This task therefore sits at the intersection of four separate AI subfields: computer vision, natural language processing, machine learningand machine reasoning. A powerful demonstration of such a capability is answering unseennatural questions about an image or a video, especially when in a multi-turn dialog. In this setting, the task of visual question answering boils down to learning to acquire and manipulate visual knowledge distributed through space-time and conditioned on the compositional linguistic cues. Our approach is based on bridging the gap between symbol grounding and symbolic reasoning through neural networks. This project comprises several sub-projects. One sub-project is to design a dual-process of reasoning, which consists of a reactive visual representation module interacting (System 1)with a deliberative reasoning module (System 2). The second sub-project aims to constructa dynamic relational working memory that temporarily links and refines visual concepts related to the question as reasoning proceeds. The third sub-project is to invent a UniversalNeural Turing Machine capable of constructing neural programs on-the-fly in response to aquery.

Project title:

Towards learning human-like intelligent systems that can grow their mind

Supervisor Name and E-Mail:

A/Prof Truyen Tran, truyen.tran@deakin.edu.au

Dr Kien Do, k.do@deakin.edu.au

Project Description:

Artificial Intelligence (AI) systems powered by deep learning have reached or surpassed human on many challenging tasks such as object detection, voice recognition, and languagetranslation. However, these systems demand a huge, expensive amount of labelled data for training and do not generalize well to new environments. Humans, by contrast, learn in a very different way. We only need a few labelled samples to learn a concept. Upon seeing unlabelled data, we can still distil concepts related information to adjust our estimation.

The learning mechanism that leverages unlabelled data to improve a model's performance is known as semi-supervised learning (SSL). Recent advances in SSL have shown that it ispossible to train models with hundreds to thousands times fewer labels to achieve competitive results compared to those using all labels, making SSL an exciting research topic. In fact, SSL has been drawing more and more attention from NLP and computer vision communities. In some extreme cases, it takes only a few labels to learn a completelynew concept. This gives rise to the few-shot learning and meta-learning (a.k.a. learning to learn) problems.

Reducing the dependency of AI systems on human's supervision is a key for scaling up learning, but this does not guarantee generalization in general. To develop their own thinking, AI systems must be aware of what they do not know. They should be able to model uncertainty, detect out-of-distribution samples, and be equipped with intrinsic motivations. Learning intrinsically motivated agents has long been an important research topic in reinforcement learning for dealing with situations in which rewards are sparse andenvironments may change over time (the task still remains the same).

In the process of learning new things, it is critical that intelligent systems are able to decidewhich knowledge they should acquire among numerous pieces of new information to achieve better performances on tasks they are solving. This requires abilities to identify meaningful concepts (both old and new) in the data based on the systems' prior knowledge/inductive bias (e.g., disentanglement), and to do reasoning on these concepts (e.g., relational/compositional/causal/contrastive inference).

In this project, PhD students will be guided towards building a human-like AI system that can learn by itself with little supervision and can actively acquire new knowledge
necessary for its jobs. Such autonomous systems are critical in the future, e.g., self-drivingcars and virtual assistants. Students can choose one or more topics from the aforementioned set of topics to pursue. On the application side, students are encouraged toapply their proposed techniques to computer vision, reinforcement learning and structureddata (e.g., graphs) with the help of experienced researchers and engineers in the lab.

Project title:
Learning to explore and extrapolate
Supervisor Name and E-Mail:
A/Prof Truyen Tran, truyen.tran@deakin.edu.au Dr Phuoc Nguyen, phuoc.nguyen@deakin.edu.au
Project Description:
Most state of the art machine learning algorithms are only powerful at modelling and generating data they have been trained on. When presented with unseen data, they often fail to capture its distribution and can only generate data that are similar to the training examples. It is essential that an intelligent agent should be able to generalise to unseen data. In medical domain for example, an AI doctor needs to diagnose correct diseases and gives proper treatments when faced with the symptoms and medical conditions of a patient.However, it is also critical that when facing a novel combination of symptoms of a potentially new disease the AI should ideally be able to recognise it and give the predictionwith proper justifications of uncertainty. In materials design, we want to learn the representation and generation mechanisms that support not only the training data and its interpolation but also its extrapolation. This capacity helps us in speeding up the discovery processes of new materials, such as alloys, chemical compounds, and drugs, which are usually expensive and laborious. This project focuses on learning to explore and extrapolate. We aim to develop novel learning methods that can learn to represent and reason about the data in a way that supports exploration and extrapolation to unseen data. Therefore, this project will touchmodern unsupervised and semi-supervised learning for out of distribution tasks. This project is divided in to three sub-projects. The first one is to develop a disentangled representation at model level beside the data level. The second is to develop an efficient and expressive model architecture that supports exact inference. The third project is to develop a regularisation framework to assist in the exploration and extrapolation learning.

Project title:

Provable Generalization in Deep Learning with Structural Inductive Biases

Supervisor Name and E-Mail:

A/Prof Sunil Gupta, sunil.gupta@deakin.edu.au

Project Description:

One of the ultimate goals of an AI system is the ability to generalize any learnings derived from training data to unseen situations. This is commonly referred to as *Generalization*. Deep learning (DL) is a subarea in machine learning (ML) that has empirically shown strong generalization ability for many tasks. However, theoretical reasons behind this success remains a mystery – e.g. why over-parameterised deep learning models defy the problems of overfitting? Additionally, data at test time may come from different (but relevant) distributions, and the current DL models are not robust to such changes in distribution. How can we provably improve the generalization ability of current DL models?

Inductive biases (e.g., composability in convolutional neural networks) – a set of cognitive structures that prioritizes the learning toward some desirable properties, are a promising direction to improve DL generalization [1]. This research proposals aim at systematically developing new structural inductive biases that provably guarantee robust generalization in DL. The significance of this project is that it will fundamentally and algorithmically advance the generalization ability of DL systems which will in turn benefit vast downstream DL application tasks that require stronger forms of cognitive abilities (e.g., sequential decision making and reasoning) rather than large-scale pattern recognition as in many current DL systems.

The fundemental challenges in DL generalization are embodied into two following questions: (i) What inductive biases can improve generalization in deep neural networks?, and (ii) In particular, how much does a structure induced by inductive biases improve the generalization bound? For (i), we will draw inspirations from cognitive science and information theory to design new scalable deep neural networks with inductive biases. For (ii), a new generalization theory is required as the capacity-based generalization bounds obtained by classical statistical learning (e.g., via uniform convergence with VC dimension and Rademacher complexity) are vacuous for overparametrized models such as deep neural networks [2]. This PhD will investigate into understanding and improving DL generalization guided by these two questions.

The essential background for this position includes a solid foundation of statistics, linear algebra, machine learning, and DL as well as proven programming skills (e.g., Python).

References:

[1] Anirudth Goyal, Yoshua Bengio. “Inductive Biases for Deep Learning of Higher-Level Cognition”, 2021.

[2] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, Oriol Vinyals. “Understanding deep learning requires rethinking generalization”, 2017.

Project title:

Generative AI for Human-Machine Co-creation

Supervisor Name and E-Mail:

A/Prof Sunil Gupta, sunil.gupta@deakin.edu.au

Dr Vuong Le, vuong.le@deakin.edu.au

Project Description:

In the last decade, the field of artificial intelligence has been greatly advanced by machine learning (ML), in particular, deep learning. In many applications, ML algorithms are able to perform at human-level and being adopted to solve problems concerning our daily lives. With such promise, ML algorithms are now being perceived as important partners to humans in most areas e.g. decision making, search and optimisation, creative applications, scientific discovery and so on. An important problem in this pursuit is *human-machine co-creation,* which is the next frontier of machine learning.

Thus far machine learning has primarily focused on tasks underpinned by pattern recognition in large corpora of data. Such systems have a clear goal, and are trained to achieve a predefined objective. But the creative process is different. It requires explorations in unchartered territory, the end goal is often abstract at the start and gets refined during the creative process. Each creative task is somewhat different, and in the novel spaces, data is scarce. Previous creative tasks may not fully inform the current task. To be a useful partner, the machines must learn from parsimonious human feedback. *Building machine learning systems that can overcome these challenges is widely open.*

The problem of human machine co-creation requires data generation with human interpretable controls, ability to generate novel data, ability to adjust the output based on human feedback and continuous evolution of the machine. These requirements intersect multiple areas of machine learning - disentangled representation learning [1], controlled data generation [4], novelty generation [2] and inverse models [3]. This PhD position will focus on advancing the research frontier in generative models along above-mentioned aspects.

The essential background for this position includes a solid foundation of statistics, linear algebra, mathematical optimisation and basic understanding of machine learning and deep learning along with proven programming skills (e.g. Python).

References:

[1] Bengio, Y., Courville, A., & Vincent, P. (2013). Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, *35*(8), 1798-1828.

[2] Lake, B. M., Salakhutdinov, R., & Tenenbaum, J. B. (2015). Human-level concept learning through probabilistic program induction. *Science*, *350*(6266), 1332-1338.

[3] Nguyen, P., Tran, T., Gupta, S., Rana, S., Dam, H. C., & Venkatesh, S. (2020). HyperVAE: A Minimum Description Length Variational Hyper-Encoding Network. *arXiv preprint arXiv:2005.08482*.

[4] Hu, Z., Yang, Z., Liang, X., Salakhutdinov, R., & Xing, E. P. (2017, July). Toward controlled generation of text. In *International Conference on Machine Learning* (pp. 1587-1596). PMLR.

Project title:

Offline Reinforcement Learning with Provable Statistical Efficiency

Supervisor Name and E-Mail:

A/Prof Sunil Gupta, sunil.gupta@deakin.edu.au

Project Description:

Reinforcement Learning (RL) is an active research area in machine learning and becoming increasingly ubiquitous. An RL agent interacts with an environment to simultaneously learn the state dynamics and a policy to make optimal decisions under uncertainty. In many practical settings, interactions with environment are expensive, unsafe or even prohibited, making it difficult for the RL agent to obtain any online data. However, an offline dataset from historical interactions (e.g., demonstration data from human experts or data from any previous policies) is often available in such environments. This research proposal aims to develop new algorithms to leverage offline data with provable statistical efficiency. The significance of this project is that it will constitute fundamental understanding (of the statistical limits and benefits) of offline RL [1] and provide practical algorithms to efficiently and reliably accelerate real-life decision-making problems such as audit, marketing, dynamic pricing, personalized medicines, recommender systems, matching markets, and new material discoveries.

The fundamental challenges of offline RL are embodied into two following questions: (i) What is a minimal, realistic condition that guarantees sample efficiency in offline RL? and (ii) How can we design a practical algorithm that can efficiently leverage various offline dataset scenarios with provable guarantee? This PhD will significantly advance the current literature of offline RL by investigating into these fundamental questions. Our initial effort toward answering these questions include our recent works [2,3] where in [2], we have established a sample complexity of offline policy evaluation - an instance of offline RL, under deep ReLU network function approximation – a commonly encountered scenarios in real applications, while in [3], we have proposed a novel way to leverage offline data and combine it with online learning to overcome the problem of support deficiency in offline learning.

The essential background for this position includes a solid foundation of statistics, linear algebra, mathematical optimisation and basic understanding of machine learning, reinforcement learning, and deep learning along with proven programming skills (e.g. Python).

References:

[1] Sergey Levine, Aviral Kumar, George Tucker, Justin Fu. “Offline Reinforcement Learning: Tutorial, Review, and Perspectives on Open Problems”.

[2] Thanh Nguyen-Tang, Sunil Gupta, Hung Tran-The, Svetha Venkatesh. “On Finite-Sample Analysis of Offline Reinforcement Learning with Deep ReLU Networks”.

[3] Hung Tran-The, Sunil Gupta, Thanh Nguyen-Tang, Santu Rana, Svetha Venkatesh. “Combing Offline Learning and Online Learning for Contextual Bandits with Deficient Support”.

Project title:

Structured Reasoning on Videos

Supervisor Name and E-Mail:

Dr Vuong Le, vuong.le@deakin.edu.au

A/Prof Truyen Tran, truyen.tran@deakin.edu.au

Project Description:

Videos have been traditionally treated as pixel-based signals and modeled with grid-based models such as CNNs. However, the videos inherently contain under-lying structures both in spatial [3] and temporal [6] dimensions. Such structure consists of active agents [7], passive objects[1] and background areas. It also includes the motions, relations and interactions of these entities. Considering these structures in video modeling will reduce complexity, support symbolic reasoning and provide semantic rich representation that bring interpretability.

The structured representation specifically supports connecting the basic facts together to deduce more complex knowledge beyond perceptive functionalities like detection and recognition. These higher-level knowledge includes relational and causal, collaborative and adversarial, analogical and counterfactual information. Modeling videos to understand these aspects requires advancing from visual recognition to visual reasoning.

Application wise, learning to reason on structured representation of videos allow new capabilities of AI systems including future anticipation [5, 7], anomaly detection [4], counterfactual prediction, and question answering [2]. Research in this topic also expand to unsupervisedly discovering the video structure according to the specific reasoning problem.

This PhD scholarship will explore this research theme and ground the ideas to particular projects that are expected to result in impactful publications. The potential PhD student will work with a team of supervisors with strong expertise in deep learning and computer vision in a dynamic team-work environment.

A successful candidate must have firm mathematical background, strong programming skill with an interest in computer vision using deep learning method. Having previous research experience is a plus point.

References:

[1] Hoang Long Dang, Thao Minh Le, Vuong Le, Svetha Venkatesh, and Truyen Tran. Object-centric relational reasoning for video question answering. ECCVW, 2020.

[2] Thao Minh Le, Vuong Le, Svetha Venkatesh, and Truyen Tran. Hierarchical conditional relation networks for video question answering. CVPR 2020.

[3] Thao Minh Le, Vuong Le, Svetha Venkatesh, and Truyen Tran. Dynamic language binding in relational visual reasoning. IJCAI-20.

[4] Romero Morais, Vuong Le, Truyen Tran, Budhaditya Saha, Moussa Man-sour, and Svetha Venkatesh. Learning regularity in skeleton trajectories for anomaly detection in videos. CVPR 2019.

[5] Romero Morais, Vuong Le, Truyen Tran, and Svetha Venkatesh. Learning to abstract and predict human actions. BMVC 2020.

[6] Romero Morais, Vuong Le, Svetha Venkatesh, and Truyen Tran. Learning asynchronous and sparse human-object interaction in videos. CVPR 2021.

[7] Hung Tran, Vuong Le, and Truyen Tran. Goal-driven long-term trajectory prediction. WACV 2021.

Project title:

Human-behaviour Understanding through Visual Analysis

Supervisor Name and E-Mail:

Dr Vuong Le, vuong.le@deakin.edu.au

A/Prof Truyen Tran, truyen.tran@deakin.edu.au

Project Description:

Observing and understanding human and their behaviour through images and videos have been a long-term important topics in computer vision. Most effort to date concentrate on perception tasks like skeleton detection and action recognition. Recent advances in human behaviour analysis provide a look into the underlying patterns of the behaviour including internal purpose (or intention) [4], regularity [1], planning [2], and interacting with environment [3]. Toward a more advanced capability of AI and useful real-life applications, many more steps are needed toward deeper understanding into the factors of human behavior and applying them in collaborative tasks. Human behaviour understanding aims at putting the observed recognized behaviour into a larger context so that more fundamental knowledge about the subjects and the situation are explored.

The bigger, better picture of the human behaviour that this project aims at have vast applications in science and everyday life. The AI doctors can infer important information about the patient's health and overall wellbeing by observe their motions. Similarly, coliving robotic agents can understand the long-term story, reasons and background of the events therefore predict future more accurately, interact more sensibly and support more helpfully to their human partners. These grounded applications will be the motivation and experimental test beds for the the theoretical findings of the project.

This PhD scholarship will explore this research theme and ground the ideas to particular projects that are expected to result in impactful publications. The PhD student will work with a team of supervisors with strong expertise in deep learining and computer vision in a dynamic team-work environment.

Besides the common requirement, a successful candidate must have firm mathematical background, strong programming skill with an interest in computer vision using deep learning method. Having previous research experience in related fields is a plus point.

References:

[1] Romero Morais, Vuong Le, Truyen Tran, Budhaditya Saha, Moussa Man-sour, and Svetha Venkatesh. Learning regularity in skeleton trajectories for anomaly detection in videos. CVPR 2019.

[2] Romero Morais, Vuong Le, Truyen Tran, and Svetha Venkatesh. Learning to abstract and predict human actions. BMVC 2020.

[3] Romero Morais, Vuong Le, Svetha Venkatesh, and Truyen Tran. Learning asynchronous and sparse human-object interaction in videos. CVPR 2021.

[4] Hung Tran, Vuong Le, and Truyen Tran. Goal-driven long-term trajectory prediction. WACV 2021.

Project title:

Human-centric Computer Vision

Supervisor Name and E-Mail:

Dr Vuong Le, vuong.le@deakin.edu.au

A/Prof Truyen Tran, truyen.tran@deakin.edu.au

Project Description:

AI research are reaching impressive performance in vision tasks, such as perception [1] and reasoning [2]. These methods are usually developed as standalone automated systems and ignoring aspects that are particular to humans that are operating in the same spaces with them. These aspects such as creativity, emotion, and ethics are particular to humans and variate among different individuals. As automated machines are increasingly designed to operate in the world shared with human, they must have the ability to observe and customize based on the interaction and collaboration with their human-partners.

Human-centric computer vision advance this new paradigm to find theories and methods to analyze visual data according to the particular context set by human. Research in this topics includes modeling of visual objects, actions, and scenes with respect to a viewpoint set by human. A human-centric computer vision system is able to allow human to inject their intention, opinion, preference into the AI's perception, reasoning and generation processes. As an example, it finds customized representation of a public space scene that is accurate and matched with a particular request of “is it safe to cross the street?” or “what is the weather like today?”

Further from that, a true machine partner should be able to automatically understand the context that it is sharing with human and by itself analyze the scene up to such context so that it can interact and collaborate the best with human. This involves human behaviour analysis, contextualized representation and visual reasoning. The ultimate goal of these research is the world where machine and human contribute their own advantages to the joint effort in achieving a complex task that neither can do alone. Machine has the edge in speed, consistency and objectiveness, while human has the strength in implicit reasoning, personality and purpose.

This PhD scholarship will explore this research theme and ground the ideas to particular projects that are expected to result in impactful publications. The PhD student will work with a team of supervisors with strong expertise in deep learining and computer vision in a dynamic team-work environment.

Besides the common requirement, a successful candidate must have firm mathematical background, strong programming skill with an interest in computer vision using deep learning method. Having previous research experience in related fields is a plus point.

References:

[1] Romero Morais, Vuong Le, Svetha Venkatesh, and Truyen Tran. Learning asynchronous and sparse human-object interaction in videos. CVPR 2021.

[2] Thao Minh Le, Vuong Le, Svetha Venkatesh, and Truyen Tran. Hierarchical conditional relation networks for video question answering. CVPR 2020.

University
Deakin University
Faculty/Institute
Applied Artificial Intelligence Institute (A^2I^2)
Project Start Date
Monthly intake
Application Deadline:
1/10/2021
Location (City/Campus)
Burwood (Melbourne, VIC) or Waurn Ponds (Geelong, VIC)
Funding Information
A^2I^2 offers full tuition + stipend scholarships available (the 2021 annual stipend rate is $28,600 tax-free), as well as top up scholarships for those who have already received other funding. Each scholarship also includes attractive HDR funding to cover the cost of presenting at the major international conferences in the field and research facilities to support student's candidature. Scholarship holders can claim up to $1,500 for relocation expenses. More details of DUPR can be found in this link: https://www.deakin.edu.au/courses/fees-scholarships/scholarships/find-a-scholarship/rtp-and-duprs
Selection criteria
All applications will go through a rigorous assessment process and shortlisted applicants will be interviewed. We are looking for self-motivated candidates who meet the following requirements: • a bachelor's or master's degree in the area of Computer Science, Electrical Engineering or related discipline. A high GPA with first class honours result is necessary to apply for scholarship candidature; • evidence of research ability - e.g. research experience, an honours thesis, peer-reviewed publication(s), and a preliminary research proposal; • strong background in some of the following subjects: Machine Learning; Linear algebra; Probability and statistics; Optimisation; • programming skills in one of C/C++, Java and Python. • overseas applicants should have a minimum IELTS score of 6.5 (with no band under 6) or a minimum TOEFL iBT score of 70 with minimum writing score of 21. The certificate of English language competency must be provided before commencing the course.
How to apply
Interested applicants should email applications to: Dr Trang Tran, HDR Coordinator at trang.tran@deakin.edu.au. The application should include applicant's Resume (or CV), and other supporting documents (if any): Expression of Interest form; Degree certificates, Academic transcripts; Published papers; Research proposal; and Referral reports. Access Expression of Interest form via this link: https://bit.ly/2PaOq8X