



Applied Artificial Intelligence Institute

Annual Report 2022

Artificial intelligence
built by humans, for humans



A²I²
APPLIED ARTIFICIAL
INTELLIGENCE INSTITUTE



We respectfully acknowledge the Wurundjeri and Wadawurrung people of the Kulin Nation, and the Peek Whurrong people of the Maar Nation, as the Traditional Custodians of the lands on which Deakin University's Burwood, Melbourne City, Geelong, and Warrnambool campuses stand. We pay our respects to their Elders past and present, and to their Ancestors, and acknowledge their continuing custodianship, care and connection to Country over many thousands of years. We recognise that sovereignty over Country was never ceded.



TABLE OF CONTENTS

Our mission	4
Director's message	5
A2I2 at a glance	6
Our unique AI expertise	7
Our researchers	8
Our networks and partnerships	9
Our research	11
Our Alumni	11
A world-first approach to mental health treatment	12
Providing life-saving decision support to hospital trauma teams	15
HDR insights	16
SmartHeart	19
Lighter, safer vehicles for our troops	20
SunSmart Global UV App endorsed by the United Nations	22
Cutting edge dementia technology	25
Grants	26
Projects	28
Awards	29
Partners	29
Staff	30
HDR Students	33
Events	34
Publications	38

OUR MISSION


Deakin University's Applied Artificial Intelligence Institute (A2I2) is unique in Australia, a university AI institute that combines world-class fundamental research with cutting-edge software development.

Bringing together decades of research and industry expertise, we deliver proof of concept solutions for problems too hard for many others to solve.

A2I2 explores new frontiers in AI and develops novel methods of implementing safe and effective AI solutions.

Our unique, multi-disciplinary expertise makes A2I2 the ideal partner to help industry, government and the community take full advantage of the benefits that AI offers.

Directors' message



In 2022, the team at the Applied Artificial Intelligence Institute (A2I2) continued its leading work in transformational industry-focused artificial intelligence.

We are immensely proud of the work of our leading academics, researchers, software engineers, supervisors, and students.

The exceptional quality of outcomes produced by our team at the Institute continues to be recognised both here at Deakin and internationally. Some of the achievements this year include:

Research outcomes

- **5 papers** accepted at Thirty-sixth Conference on Neural Information Processing Systems (**NeurIPS 2022**) in New Orleans, USA.
- **4 papers** accepted at European Conference on Computer Vision (**ECCV**).
- **23 papers** in total accepted at **A/A Rank conferences**.
- **14 papers** published in **Q1 Journals**.

Translation outcomes

- Launch of the **BrainTrack app** with **Dementia Australia**, which is supporting early dementia diagnosis and has been downloaded by over 55,000 users globally.
- **Cancer Council Victoria** partnered with A2I2 to deliver a global version of the SunSmart App. The SunSmart Global UV app is one of only a handful of apps worldwide to be officially endorsed by the United Nations.
- **Defence Science and Technology Group (DSTG)** tactical event recognition and tracking technology was deployed on exercise in November 2022, successfully enabling the identification and recognition of tasks in real time in a search and rescue exercise.
- **Uniting AgeWell Limited** (Victoria and Tasmania) used AI to identify at-risk participants who needed a higher level of care through their care package.
- **Brand Influence Group** rolled out on-device processing technology to identify stock levels of products accurately and efficiently on supermarket shelves using hand-held devices.

For the third year in a row, an A2I2 student was awarded an Alfred Deakin Medal for doctoral theses for 2022. Graduates of the Institute continue to be sought after by leading institutes and companies and many are forging their careers here at the Institute.

By fostering strong collaborations across fields such as health, renewable energy, manufacturing, defence, education, and space, our researchers and project teams are advancing the interdependence of technology and human expertise. This work is hard and challenging, but also immensely rewarding.

2022 has seen the Institute deepen its relationship with existing partners and welcomed new partners, including HomeGuardian and Control Bionics.

The pivotal work being undertaken with the Black Dog Institute and the University of New South Wales continued across 2022. Black Dog undertook the Vibe Up trial, evaluating intervention methods for anxiety and depression in university students across Australia. The A2I2 team developed a novel AI solution that allowed personally targeted interventions to be trialled in shorter timeframes with fewer participants.

Institute Co-Director Professor Kon Mouzakis was awarded an Alfred Deakin Professorship by the Deakin University council. The title of Alfred Deakin Professor is the highest honour Deakin University can bestow upon a member of the academic staff.

In September 2022, ARC Laureate Professor Svetha Venkatesh joined a group of leading Australian artificial intelligence professors in the foundation meeting of the Kingston AI Group. The group met to discuss how to better coordinate our national research and education effort so that Australia benefits from the AI revolution.

2022 was a very solid year of industry and academic achievements and this would not have been possible without the hard work and dedication of the staff at A2I2.

Alfred Deakin Professor Svetha Venkatesh

Australian Laureate Fellow
Co-Director, A2I2

Alfred Deakin Professor Kon Mouzakis

Co-Director, A2I2

80+

AI RESEARCHERS AND
SOFTWARE ENGINEERS

TOP 15

GLOBAL RESEARCHERS
BASED AT A2I2

7

RESEARCHERS RANKED IN
THE TOP 3 IN AUSTRALIA

100+

Q1 JOURNAL
PUBLICATIONS

TOP 1%

AI PHD CANDIDATES
BASED AT A2I2

35+

AWARDS RECEIVED
ACROSS OUR WORK

250+

INDUSTRY
COLLABORATIONS

100+

1ST RANK CONFERENCE
PAPERS

Our unique AI expertise

A2I2 is one of the only AI institutes that combines world-class fundamental research with cutting-edge software development. Bringing together decades of research and industry expertise, we deliver proof of concept solutions to problems too hard for many others to solve. We are experts in the field of small-data AI research and translation. We are really good at these things:



Bayesian Optimization

Optimising expensive-to-evaluate functions - reducing expensive experimentation, and research and development processes, which allows for advanced results delivered faster and with less resources.



Reinforcement learning

Machines learn by playing. Teaching AI systems to learn from experiences and make smarter decisions on their own — which means you can increase productivity by automating repetitive tasks without human supervision.



Deep learning

Machines learn from previous data. Identifying efficient ways for machine learning to occur with minimal human supervision — solving critical real-world problems such as drug design, chemical reactions, and more.



Natural language programming

Think ChatGPT. Bringing human-like communication to AI systems, allowing virtual assistants like Siri and Alexa to respond to your voice commands with ease — creating a more interactive and personalised user experience.



Domain generalization

Machines learn from other machines. Enabling AI models to operate effectively in new and challenging environments - enabling self-driving vehicles to navigate and adapt in real-world conditions.



Abstraction & representation

Machines learning more like humans. Building AI models that can understand vast amounts of complex data and identify trends or anomalies — enabling financial institutions to identify fraudulent transactions within seconds.



Decision making using offline data

A new field of research. Building AI models that can make decisions with historical data and without real-time input — streamlining your processes, allowing faster and more automated decisions without human intervention or new data.



Human machine teaming

Machines and humans learn together. Developing AI systems that work collaboratively with human experts to deliver better performance than either could achieve on their own, such as supporting surgical decision making in real time or reducing waste in manufacturing processes.


Our researchers

True impact results from combining industry leading experience with world-leading research capabilities — which is why we bring together 100+ of the world's brightest minds to solve society's biggest challenges. Our team of AI experts are led by Directors that are committed to nurturing your project to success



ARC Laureate Professor Svetha Venkatesh

Svetha Venkatesh is an ARC Australian Laureate Fellow, Alfred Deakin Professor and a co-Director of Applied Artificial Intelligence Institute (A2I2) at Deakin University. She was elected a Fellow of the International Association of Pattern Recognition in 2004 for contributions to formulation and extraction of semantics in multimedia data, a Fellow of the Australian Academy of Technological Sciences and Engineering in 2006, and a Fellow of the Australian Academy of Science in 2021 for ground-breaking research and contributions that have had clear impact.



Alfred Deakin Professor Kon Mouzakis

Professor Kon Mouzakis is a Professor of Software and Technology Innovation at Deakin University. He is a Decision Support and Software Development expert with over 30 years of experience and currently holds the position of Co-Director of the Applied Artificial Intelligence Institute (A2I2). In 2018, Kon was appointed as the Director, Australian Research Council, Industrial Transformation Research Hub for Digital Enhanced Living, and in 2022, was conferred with the title of Alfred Deakin Professor.

Some of our world-leading experts:

- Professor Sunil Gupta is ranked #2 nationally, and #13 globally in Bayesian optimization
- Professor Rajesh Vasa is ranked #6 globally and #2 nationally in Automated software engineering
- Associate Professor Santu Rana is ranked #3 nationally, and #23 globally in Bayesian optimization
- Associate Professor Shannon Ryan is ranked #3 globally and #1 nationally in Hypervelocity impact
- Dr Hung Le is ranked #2 globally, and #1 nationally in Memory Networks
- Dr Alistair Shilton is ranked #1 nationally and #25 globally in Kernel methods



Our research

We are dedicated to the development of AI that adds positive value to the world and enhances human expertise. We attract some of the world's brightest young minds working in AI. Our institute delivers solutions to challenges in:



Health



Renewable energy



Manufacturing



Defence



Education



Space



Finance



IoT

Our Alumni

Our A2I2 graduates have secured highly sought after jobs in leading organisations, including:

Google

Amazon

Canva

Johns Hopkins Whiting School of Engineering

Deloitte

Zendesk

Trusting Social

Deakin University

Monash University

Trang University, Vietnam

Victoria University, Australia

Central Queensland University

REA Group

Vietnam National University

Keep Flying, Singapore

A world-first approach to mental health treatment

The pivotal work being undertaken in partnership with the **Black Dog Institute** in NSW continued across 2022.

Black Dog Institute launched The Optimise Project in 2020, as a collaborative research project working with Deakin University's Applied AI Institute.

It started with a question: *How can AI-driven modelling improve mental health research outcomes for participants?*

Funded by the Medical Research Future Fund (MRFF) Artificial Intelligence in Health scheme, The Optimise Project is aiming to advance personalised approaches to mental health care to match individuals with the tools and strategies that are most likely to help them.

A growing problem

In Australia, more than 1 in 4 students experience high-stress levels that negatively impact their studies, relationships, and daily lives.

Randomised Control Trials (RCT) have shown that a range of interventions targeting mindfulness, exercise and sleep can be effective. RCTs are the gold-standard for studying the effectiveness of single treatments or interventions across the health sector, but are also expensive, take significant periods of time to run and the results are typically broad-based.

One of the major problems in managing psychological distress is recognising in real time the best combination of treatments or interventions for a particular individual, given their symptoms, severity, and demographics.

Vibe Up: a novel approach to large scale trials

The Vibe Up study forms part of The Optimise Project, using artificial intelligence to analyse real-time data within an adaptive trial method.

Although adaptive trial methods were first discussed over three decades ago, they have only recently been utilised in health settings. Through the Vibe Up study, the research team are pioneering a world-first approach to mental health treatment using adaptive experimental design to optimise treatment for people experiencing psychological distress.

For the Vibe Up study, the research team recruited over 1000 student participants, each enrolled in tertiary courses and with a unique experience of university life.

The participants were provided with a mobile app that offered remote access to psychoeducational interventions.

A series of artificial intelligence-based 'mini-trials' were conducted, aimed at discovering the most effective interventions.

The perceptions of mental health professionals have been compared with the decisions made by AI in the Vibe Up study.

The Mobigene Study

The next trial in The Optimise Project is the Mobigene Study, using the same methodology as the Vibe Up Study with a clinical sample. The Mobigene Study has assessed the feasibility of combining genetic data (already collected through the Australian Genetics of Depression Study) with smartphone digital data and self-reported mental health information.

The bigger picture

New developments in AI offer a unique and scalable way to address the limitations of RCTs in health settings where a treatment type might be uncertain.

The new AI being developed under The Optimise Project aims to offer personalised treatments to patients and allow the optimisation of treatments across a range of disciplines.

Imagine the potential of that.

The results of the Vibe Up/The Optimise Project study are due to be announced in 2023.



“Through Vibe Up, we aim to provide a richer and more streamlined healthcare experience, to help people get the most effective treatment as fast as possible.”

Professor Jill Newby
Head of Clinical Research, Vibe Up Study Lead
Black Dog Institute



Providing life-saving decision support to hospital trauma teams

In the bays of the Alfred Trauma Service in Melbourne, a critical resuscitation decision is made every 72 seconds.

One of the premier trauma centres in the world, The Alfred began working with the National Trauma Research Institute (NTRI) and Deakin's Applied Artificial Intelligence Institute (A2I2) in 2019 to develop an advanced iteration of the Trauma Reception and Resuscitation decision support system (TRR2.0©). Across 2022, Professor Mark Fitzgerald, Director of Trauma Services at The Alfred Hospital, and his team, have continued to work with A2I2 on further enhancing the system.

In 2022

Software engineers at A2I2 developed updates to the system that allow the National Trauma Research Institute to create and support different types of decision support algorithms. Alongside the trauma algorithms, the system now supports decision making algorithms created for cardiac and stroke care, with other targeted algorithms currently in development.

"The TR&R has entrenched The Alfred Hospital's goal of error-free resuscitation in the treatment and management of trauma patients. We have observed a 21% reduction in errors of omission, 30% reduction in blood transfusions, and significant reductions in the time spent in ICU"

Professor Mark Fitzgerald,
Director of Trauma Services
The Alfred Hospital

The Need for Decision Support

Stabilising a patient within a trauma bay places immense pressure on trauma bay staff.

Each year, more than 1,100 major trauma patients – those patients with injuries that can cause death or disability – present to the Alfred Trauma Service. For each patient, any mistake by trauma bay staff could mean death, and in the first 30 minutes of arriving at the Trauma Centre every moment counts. Staff are highly trained and experienced to deal with situations like this, following procedures and process. Even so, with multiple staff, highly stressful environments and the pressure to stabilise patients, sometimes errors do occur. These are errors of omission – where a step or process is left out.

How the system works

The Trauma Reception and Resuscitation System is designed to aid and support highly trained staff to stabilise and treat major trauma patients within the first 30 minutes of arrival at the Trauma Centre. The system monitors patient data, such as vital signs and diagnosis information, and provides advice and treatment options based on hospital procedures and guidelines created in the accompanying Medical Algorithm Builder. These recommendations are displayed on a large screen visible to all staff, and act as a decision support system to reduce errors of omission.

A two-year independent study following the deployment of the TR&R system found the system contributed to:

- **A 21 per cent reduction in errors by omission**
- **30 per cent reduction in blood transfusions**
- **Reduced the time required for treatment in the Intensive Care Unit by 26 hours**

HDR insights:

Third year running for Alfred Deakin Medal for doctoral theses

The outstanding quality of the theses produced by our A2I2 doctoral students continues, with Dr Thanh Nguyen-Tang being awarded an Alfred Deakin Medal for doctoral theses for 2022.

Dr Nguyen-Tang's medal-winning thesis follows three years of acknowledgement of the best doctoral theses from across the university for the Applied Artificial Intelligence Institute (A2I2).

In 2020, the first year A2I2 was eligible for the awards after being founded in November of 2019, Dr Hung Le was awarded the best thesis medal. In 2021, Dr Thao Le was also awarded the prestigious Deakin University thesis medal.

We are incredibly proud of the quality of our students' work, and of the high level of support and guidance provided by our expert supervisors.

A2I2's Alfred Deakin Medal winners

2022:

Dr Tang Thanh Nguyen

Research Fellow,
Johns Hopkins Whiting School of Engineering

Executive Supervisor: Prof Sunil Gupta
Co-Supervisor: A/Prof Santu Rana
Associate Supervisor: A/Prof Truyen Tran
Prof Svetha Venkatesh

2021:

Dr Thao Minh Le

Research Fellow, A2I2

Principal Supervisor: A/Prof Truyen Tran
Associate Supervisor: Dr Vuong Le
Prof Svetha Venkatesh

2020:

Dr Hung Le

Research Lecturer, A2I2

Principal Supervisor: A/Prof Truyen Tran
Associate Supervisor: Prof Svetha Venkatesh

Five Minutes with Dr Thanh Nguyen-Tang

What did you learn during your PhD that has supported your academic career?

Asking fundamental questions is something that I think very important since I pursued my research career. Not all questions are important and worthy of an answer, and fundamental ones will take you far and make the journey much more fun and rewarding. Learning to ask questions that matter is perhaps one of the most valuable lessons I have learned so far. However, it is not always trivial to be able to ask the 'right' fundamental questions - this is something that I am still constantly learning to do.

You had several achievements during your PhD study, what stands out to you as those you are most proud of?

Thank you for your kind words. To be honest, though my PhD study time was enjoyable, I do not consider myself to have an achievement, and I think the best work I will be proud of is yet to come (but I hope it will come soon). But I do have fun when some of my works are perceived quite well in the community. For example, there is a Reddit thread discussing my AAAI 2021 paper by a senior AI engineer at Microsoft, while my TMLR 2022 paper is frequently credited by top researchers as one of the first papers that formally start out the study of the theoretical benefits of deep neural networks for offline reinforcement learning.

What can you tell current students about the benefits of serving the research community during a PhD?

Serving the research community such as reviewing for conferences has many benefits and I highly recommend PhD students to take this opportunity whenever possible. It gives you a high sense of community and connection with other peers in your field. As a reviewer, you will be one of the first few people to be exposed to new ideas that are not even available to the public. This can help you quickly develop a good taste of what's going recently in your field. Reviewing also forces you to read a paper critically - to the point that you will be able to judge its scientific merit, a skill that I think is very important to develop as an independent researcher.

What were the key challenges for you undertaking your PhD research?

The key challenges were to ask the right questions - questions that are interesting and fundamental enough yet doable enough given the limited time. My research was/is to tackle foundational questions of modern machine learning settings - questions such as what can and cannot be learned from observations and how much observations are necessary and sufficient for performing a learning task at hand successfully. My general idea is to break a question into substantial pieces that are small enough to be tackled. I also enjoy drawing new perspectives from different fields (e.g., combinatorics, statistics, and optimization) to shed light on an answer to the question at hand.

[Cheeky question] Was it just that you were an amazing student, or was Sunil really that good as a supervisor?

I think Sunil was a great supervisor. He gave me the freedom to do research without constraint to any specific pre-determined project and he is being also a good friend. In addition, A2I2 has a supportive environment with many good friends of mine who support me not just academically but personally.

Can you tell us a little bit about your life in the USA and your work at Johns Hopkins Whiting School of Engineering?

I am currently living near the main campus of Johns Hopkins in Baltimore, Maryland. It is only 1 driving hour to Washington DC and 4 driving hours away from New York City. Life here has been fantastic so far. I am lucky to have been surrounded by many nice and smart people here. I am working with my fantastic advisor and other collaborators on a set of problems that aim to lay the theoretical foundations of modern machine learning, with key topics including reinforcement learning, transfer learning, and robust machine learning. In particular, we study the theoretical foundations of several important modern learning paradigms when the data-generating distributions might change over time, due to either the nature of the learning tasks or due to an adversary. We also study the theoretical foundations of algorithmic constraints such as robustness, fairness, and privacy. Understanding these constraints will ensure the trustworthiness of machine learning methods and allow them to be adopted in broader contexts and applications.

Five minutes with Dr Thao Minh Le

How do you describe Computer Vision to people who know nothing about computer science?

Computer vision is all about teaching computers to understand our visual world from what they see by analysing digital images and videos captured by cameras.

What did you learn during your PhD that has supported your academic career?

Reflecting on my PhD journey, I see it as a continuous learning experience marked by personal growth in conducting independent research. I guess the most valuable lesson I gained was the art of problem solving. Research is not just about finding answers, but it is all about asking the right questions.

You had several achievements during your PhD study, what stands out to you as those you are most proud of?

What I am most proud of is seeing how my research findings have facilitated subsequent research by other researchers, advancing the entire field of interest. There were only a few researchers worldwide working on my topic, but it has since attracted a significant interest within the research community.

What were the key challenges for you undertaking your PhD research?

I have to say the COVID-19 pandemic posed a significant challenge during my PhD research. A major portion of my PhD studies occurred during the pandemic, and the abrupt transition from working in an office to working from home was really challenging. Managing mental health and maintaining focus during this long period proved to be difficult. Of course, relocating from Tokyo to Geelong presented another challenge, given the stark contrast in the living style between the two places. I am deeply grateful for the support I received from everyone here, especially my incredible supervisors, Truyen, Svetha and Vuong, who made these difficult times more manageable.


[Cheeky question] Was it just that you were an amazing student, or was Truyen really that good as a supervisor?

All of my publications are the result of collaborative efforts between myself and the supervisory team. The success of my thesis would not have been possible without the insights gained during our regular meetings and brainstorming sessions. In my opinion, while student's abilities and the excellence of their supervisor play a role, what truly matters is how efficiently they work together as a team. Effective communication between supervisor and student is the key to the success of any PhD research.

Can you tell us a little bit about your life here in Australia and how your work at A2I2 has shaped your research career?

In addition to the tremendous support from my colleagues and supervisors at work, I am fortunate to have my little family here with me, providing me with strength and courage during challenging times. Australia is a beautiful country with amazing natural landscapes and a diverse culture, and my family and I love being here.

To be entirely honest, my work at A2I2 has been a turning point in my research career as it has significantly boosted my confidence in pursuing a research career post PhD. It has been an honour to collaborate with brilliant minds to tackle the most exciting challenges in the field of AI.



SmartHeart

The SmartHeart pilot project is a first step in investigating a new co-designed way of supporting people at risk of heart failure in the home in Australia.

A2I2's head of translational research, Prof Rajesh Vasa, continued his work as a Principal Investigator on the SmartHeart project.

"The idea is to make a house a smart home for people recovering from heart disease," Prof Vasa said. "So, it could be post-surgery or post-incident, or they may just have a chronic condition. They get a Samsung smart watch, which provides continuous data around heart rate and so on, we put sensors in their medicine cabinet, as well as other sensors around the house, and we give them a smart speaker that gives them support and guidance on how to change their lifestyle a little bit."

The idea behind the project is to investigate whether combining continual data monitoring with a system to flag sudden or increased risk factors can help alert medical teams that a patient with heart disease needs a timely intervention. A weighing scale, for example, sends an alert if there is a sudden change in weight that may indicate a need for an intervention needed in a patient's medication.

"The sensor and smart watch data is run through an algorithm in the background that was developed with input from cardiologists. The smart speaker in the home can talk to the patient and get additional information when required, so an intervention, if needed, is better informed," Prof Vasa said.

"Typically, people will not be aware of all the indicators that something with their condition has changed and needs to be addressed fast. You will probably not be meeting your cardiologist or even your GP regularly, so this system is looking to fill that gap."

SmartHeart is a proof of principal validation. If it is shown to be effective through the pilot program, the next step would be a small-scale trial implementation. The hope is that if the system works, something like it could be implemented across the community at scale.

About SmartHeart

Heart failure is a leading cause of preventable death and disability in Australia with increasing prevalence. Complex medication regimens, monitoring of symptoms, ongoing medical review, exercise, and attention to food and fluid intake are all important components of heart failure management but require education and support to achieve. New approaches are urgently needed to support people with heart failure to better self-manage this condition, and to enable the health system to optimise the delivery of care for the rapidly increasing need and demand on services.

The SmartHeart research team is working with people with heart failure, their carers and healthcare providers to produce an intelligent smart ecosystem (SmartHeart) to support the co-management of people with heart failure. SmartHeart will deploy state-of-the-art sensing technologies to understand people with heart failure's behaviours and symptoms, and deliver individualised support to improve self-management, health outcomes, and reduce hospital admissions and mortality in this priority population.

We are using co-design and participatory research methods to develop the SmartHeart ecosystem with end users and healthcare providers as active participants in the process. The feasibility and acceptability of implementing our Smart-Heart ecosystem alongside existing health services for the management of heart failure will be explored so that Smart-Heart can be rapidly translated into practice.

SmartHeart is a four-year NHMRC grant funded project. The multi-disciplinary project team is headed by Prof Ralph Maddison, Professor of Digital Health and Disease Prevention at Deakin University's Institute for Physical Activity and Nutrition (IPAN) and includes leading researchers and practitioners spanning the fields of digital health, disease prevention and nutrition, artificial intelligence, cardiology, nursing, general practice, computer science, and health economics.

Lighter, safer vehicles for our troops

How can we make the Australian Army's vehicles lighter, stronger, and better able to protect our troops?

In Defence applications, this is termed 'survivability'. Associate Professor Shannon Ryan, Head of Defence and Space Research, has been working on this problem with defence industry protection specialists, Plasan.

The challenge:

Armour on defence vehicles is limited by how much weight the vehicle can carry.

The solution:

Use a novel human-machine teaming approach to inspire the design of lighter, more capable armour designs.

The outcome:

In 2 weeks and less than 100 ballistics tests, we improved the performance of an armour system by 10%, something that would normally take years.

Plasan have designed the armour for the Australian Army's Hawkeye vehicles, upgraded the survivability of the Bushmaster vehicles, and are the designers for the new Redback infantry fighting vehicles, to be built in Geelong under the Commonwealth government's multi-billion-dollar Land 400 Phase 3 Program. They know a lot about armour design. The question was, how to make better designs, faster, cheaper and with less testing?

A2I2 are world-leading experts in Bayesian optimization – a machine-learning approach that can be used to find the best design in a 'sample efficient' manner – in other words, for the fewest number of attempts. However, this approach would not typically take into account the decades of expertise in armour design that a company like Plasan holds.

"There are reasons for that. It's hard to extract expertise from a human in a way that you can ingest it in a machine learning model," A/Prof Ryan said.

"The research team here developed something called BO-Muse [Bayesian optimization Muse] for real world applications such as this. It is basically a framework for how you can team up a human expert or experts with AI to do experimental design."

The approach taken with BO-Muse has its roots in psychology. Humans are unmatched at tweaking existing solutions to squeeze out a few percent better performance. However, high-level experts find it difficult to imagine novel solutions in their field because of a concept known as cognitive entrenchment. The more you know about a domain, the harder it is to think creatively about it. In the BO-Muse framework, the role of the AI is to make out-of-the-box suggestions in a systematic way. These creative ideas inspire the human expert, who uses their deep expertise to tweak them into workable solutions.

"We tested this with Plasan as a proof-of-concept and it worked brilliantly."

Advancing the field of AI: What is BO-Muse?

BO-Muse is a new way for human experts and AI to work together in experimental design.

- Bayesian Optimisation (BO) is a machine learning based optimisation technique that can be used in experimental design to discover the best solution to a complex problem in the fewest number of experiments.
- Conventionally, when we apply Bayesian optimization, whether it's to solve a health problem or to accelerate the design of a new material, the algorithm runs independently without any other inputs.
- However, if we have access to an expert and can incorporate their expertise into the optimisation, then we can perform the optimisation even faster!
- Humans are unbeatable in their ability to understand complex systems, particularly when they have some past experience with the problem (e.g., they know some designs that work well). However, they tend to fixate on solutions that they know and understand. As their degree of expertise increases, their ability to think of creative solutions to problems in their field decreases ("this is called cognitive entrenchment").
- In BO-Muse the role of the AI is to generate out-of-the-box ideas to inspire creativity in experts, hence the name "Muse".



SunSmart Global UV App endorsed by the United Nations

Cancer Council Victoria partnered with Deakin's Applied Artificial Intelligence Institute (A2I2) to deliver a global version of the SunSmart App.

Cancer Council Victoria's Head of Prevention, Adjunct Associate Professor Craig Sinclair, said SunSmart Global UV app allows users to identify the times of the day when high UV puts them at risk of sun damage.

"We want people to use UV and sun protection times to help guide behaviours, and not to rely on the temperature, as you can't see or feel UV radiation, the damage is often being done before it is noticed"

SunSmart Global UV is one of very few Apps officially endorsed by the United Nations (UN).

The app was developed by Cancer Council Victoria, in partnership with A2I2 and funded by the ARC Research Hub for Digital Enhanced Living.

A2I2 Product Manager, Andrew Vouliotis, said this latest version of the SunSmart App was taking the SunSmart messages that Australians are so familiar with: **Slip** on clothing, **Slop** on sunscreen, **Slap** on a hat, **Seek** shade and **Slide** on sunglasses; to the rest of the world.


"It was our job to bring Cancer Council Victoria's vision of a global SunSmart App into reality," Mr Vouliotis said.

While speaking at a briefing at the United Nations in Geneva, WHO spokesperson, Carla Drysdale added that the app "aims to help people around the world know when to use sun protection, in an effort to reduce the global burden of skin cancer and UV-related eye damage".

The app was truly a collaborative effort, and the result is a single App that provides daily sun protection times, and, where the forecast data is available, a five-day weather and UV forecast. SunSmart Global UV is one of only a few apps in the world to be officially endorsed by the United Nations (UN).

The A2I2 project team collaborated with national and international agencies to gather UV and forecast data. Live UV data comes from the Australian Radiation Protection and Nuclear Safety Agency (ARPANSA), with global UV drawn from the European Centre for Medium-Range Weather Forecasts (ECMWF). The Bureau of Meteorology (the Bureau) worked with the ECMWF, helping to guide them on method of forecasting UV data in Australia. While weather data is provided by the Hong Kong Observatory.

To bring the vision of a global UV alert app to life, Cancer Council Victoria partnered internationally with the World Health Organization (WHO), the World Meteorological Organization (WMO), the International Labour Organization, and the UN Environment Programme (UNEP).







Cutting edge dementia technology

Dementia Australia has continued to pursue a mission of developing new technologies to improve the experience of dementia for everyone.

In October 2022, Dementia Australia released the BrainTrack app, the latest addition to a suite of ground-breaking virtual reality, artificial intelligence and app technologies developed in partnership with Deakin's Applied Artificial Intelligence Institute (A2I2).

Dementia Australia CEO, Maree McCabe AM, said the use of technology is changing and improving the support, care practice, knowledge and awareness for all people impacted by dementia.

"The relationship with A2I2 has truly been the fulfilment of some of our dreams about how we can use technology to make a profound and lasting difference to the lives of people living with dementia."

Maree McCabe AM, Dementia Australia CEO

BrainTrack is a free app and has been developed as a tool for users to learn about brain health and track cognition over time through a series of fun, travel-themed games. Users are prompted to log in every month, and if they have concerns they can generate a pdf report of results which can be shared with their GP to use as a conversation starter.

To date, over 55,000 people have downloaded the app around the world.

Since first setting out to create change through the use of immersive technology, Dementia Australia's product offerings have brought dementia to life for thousands of Australians, who develop insight and empathy by experiencing first-hand its symptoms and impacts.

"Professional and family carers tell us how profound it is to enter the world of a person with dementia. And we know from formal evaluations the power of this style of training compared to traditional 'chalk and talk,'" Ms McCabe said.

She described the tech journey at Dementia Australia as an audacious one; harnessing the power of gaming technologies and other high-tech tools to transform dementia understanding and care.

Other multi-award-winning projects developed by Dementia Australia and A2I2 include EDIE (Educational Dementia Immersive Experience); the virtual reality experience, Talk with Ted; an artificial intelligence-powered avatar to help professional carers practice their communication skills, and Ask Annie, the care worker in your pocket, who delivers micro lessons to care workers through the convenience of an app with an anytime, anywhere approach.

Grants

GRANT NAME	SOURCE	GRANT PERIOD	ALLOCATION (AUD)
ARC Research Hub for Digital Enhanced Living	ARC Industrial Transformation Research Hubs 2017	2017-2022	\$2,962,655
Automated Wheelchair	Innovation Connections Grant (Commonwealth + Control Bionics)	2021-2022	\$112,955
Carer Manager Plus Scheduler	Innovation Connections Grant (Commonwealth + SEQUEL)	2021-2022	\$99,724
Learning a Taxonomy	Innovation Connections Grant (Commonwealth + SURROUND)	2021-2022	\$149,711
Statistical Estimator for Resource Calculation Workflow	Innovation Connections Grant (Commonwealth + SantieX Global)	2021-2022	\$68,181
Developing a Comprehensive Care Pathway For those at Risk of Suicide But Not in Care: The Under the Radar Project	MRFF Million Minds	2020 - 2023	\$167,000
Conductor - MRFF Optimise MIA Variation	Black Dog Institute & MRFF	2020 - 2023	\$164,000
Implementing innovative technology promoting self-awareness of brain health and self-determination in obtaining a timely dementia diagnosis	MRFF Dementia Ageing and Aged Care Mission	2020 - 2023	\$193,021
A generic framework for verifying machine learning algorithms	ARC Discovery Projects	2021-2023	\$361,000
Development of Bayesian optimisation tools for accelerated design and discovery of non-magnetic structural damping alloys	DSTG	2021-2022	\$119,505
Assuring an off-the-shelf AI algorithm	DSTG	2021-2022	\$253,034
Machine Learning for Rapid Materials Design and Assessment	DSTG	2021-2024	\$541,847
PRECIS: PRecision Evidence for Childhood obesity prevention Interventions	NHMRC - Ideas Grant	2021-2024	\$310,414
Machine-learning based trajectory modelling	Department of Defence	2022	\$44,828

GRANT NAME	SOURCE	GRANT PERIOD	ALLOCATION (AUD)
Applied Artificial Intelligence Research in Health - Optimising Treatments in Mental Health Using AI (UNSW led)	Medical Research Future Fund	2020-2023	\$3,121,410
ARC Research Hub for Digital Enhanced Living	ARC Industrial Transformation Research Hubs 2017	2017-2022	\$2,962,655
Centres of Research Excellence - Centre of Research Excellence in Suicide Prevention: CRESP II (UNSW led)	NHMRC Project Grant -APP1152952	2018-2023	\$189,345
Developing a Comprehensive Care Pathway for those at Risk of Suicide But Not in Care: The Under the Radar Project (UNSW led)	Medical Research Future Fund – Million Minds Mission	2020-2024	\$335,158
Early Moves: A Prospective Cohort Study to Identify an Early Biomarker for Cognitive Impairment (Curtin University led)	NHMRC	2020-2024	\$230,175
Molecular Modelling and Machine Learning to Optimize the Cellulose Nanocrystal Interface with Water and Polymers	Ford USA	2020-2022	\$180,471
"Pattern Analysis for Accelerating Scientific Innovation"	Australian Laureate Fellowship Grant	2017-2022	\$3,016,065
Machine Learning for Design & Evaluation of Complex Armour Systems	Plasan Sasa Ltd	2021-2024	\$138,624
Coupled self-supervised learning and deep reasoning for improved processing of noisy and dynamic multimodal data from multiple sources - Phase 1	Department of Defence	2022	\$105,365
Movement Detection for Robust and Interpretable Cerebral Palsy Risk Assessment	The Research Foundation Cerebral Palsy Alliance	2022-2025	\$298,678
SatPing - a tracking beacon for spacecraft (scoping study)	SmartSat CRC Ltd	2022-2023	\$30,000
Coupled self-supervised learning and deep reasoning for improved processing of noisy multimodal data - Phase 2	Department of Defence	2022-2025	\$758,275

Projects

PARTNER	PROJECT	YEAR BEGAN	YEARS ACTIVE	TOTAL VALUE
Black Dog Institute	ARC DEL Hub	2017	2017-2022	\$50,000
Black Dog Institute	Applied Intelligence Research in Health - Optimising Treatments in Mental Health Using AI	2020	2020-2023	\$2,941,010
Black Dog Institute	Developing a Comprehensive Care Pathway for Those at Risk of Suicide But Not in Care: The Under the Radar Project	2020	2020-2024	\$335,158
Brand Influence Group	Planogram Compliance	2022	2022	\$99,932
Cancer Council Victoria	SunSmart Global UV App	2021	2021-2023	\$47,160
Compass	Griddle Exam Engine	2022	2022	\$80,000
Control Bionics	Automated Wheelchair	2021	2021-2022	\$112,955
Deakin University	Living Knowledge	2021	2021-2022	\$500,000
Deakin University (IPAN)	Smart Heart	2022	2022	\$59,093
Dementia Australia	Ask Annie Phase 3	2022	2022	\$184,880
Dementia Australia	Brain Passport	2021	2021-2022	\$250,000
Dementia Australia	Brain Passport Phase 2	2022	2022	\$250,000
DST Group	Context Aware Search	2021	2021-2022	\$143,338
DST Group	Context Aware Search - extension	2022	2022	\$232,838
DST Group	Exploring Digital Techniques to Enhance Military Effectiveness	2019	2019-2022	\$552,646
DST Group	Extensible Wargaming Platform	2022	2022	\$216,427
DST Group	Research Support for Information Displays and Human-System Interfaces	2021	2021-2023	\$799,478
DST Group	Tactical Event Recognition and Tracking Phase 2	2022	2022-2023	\$222,799
EPFL	Muse Phase 3	2022	2022	\$50,000
GoAct	ARC DEL Hub	2017	2017-2022	\$25,000
Interrelate	ARC DEL Hub	2017	2017-2022	\$37,500
Lockheed Martin	SPARC Phase 2	2021	2021-2022	\$150,000
NTT	Phase 4 Alerts & Customisation	2022	2022	\$451,492
Sequel	Carer Manager Plus Scheduler	2021	2021-2022	\$99,724
Surround	Learning a Taxonomy	2021	2021-2022	\$149,711
Uniting AgeWell	ARC DEL Hub	2017	2017-2022	\$100,000
Uniting NSW	ARC DEL Hub	2017	2017-2022	\$50,000

Awards

DATE	RECIPIENT	AWARD
May 2022	Professor Kon Mouzakis	Alfred Deakin Professor
August 2022	Dr Leonard Hoon, Dr Rena Logothetis, Dr Scott Barnett	Best Poster at the 18th IEEE International Conference on eScience
October 2022	Dr Arun Kumar Anjanapura Venkatesh	NeurIPS 2022 Scholar Award
October 2022	Dr Thommen Karimpanal George	Top reviewer (Top 8% of reviewers) at NeurIPS 2022

Partners

Alfred Health
 ARC Research Hub for Digital Enhanced Living
 Ambulance Victoria
 Atlas Zoo
 Australia India Trauma Systems Collaboration
 Black Dog Institute
 Brand Influence Group
 CAE
 Cancer Council Victoria
 Cerebral Palsy Alliance
 Compass Education
 Control Bionics
 CSIRO
 Curtin University
 Deakin University - IFM
 Deakin University - IPAN
 Deakin University - Faculty of Health
 Dementia Australia
 Department of Defence
 DST Group
 École Polytechnique Fédérale de Lausanne (EPFL)
 Education Services Australia (ESA)
 Ford Motor Company

Flinders University
 HomeGuardian
 Interrelate
 IOOF
 Karista
 Lockheed Martin
 NSW State Super
 Nippon Telegraph and Telephone Corporation (NTT)
 Plasan Sasa
 Pondr
 Regis
 Samsung
 Sequel
 Smartsat CRC
 Surround
 Sugati
 Thought Green Technologies
 Uniting AgeWell Ltd
 Uniting NSW
 University of New South Wales
 University of Technology Sydney
 Western Sydney University
 Western Health

Staff

Directors

Alfred Deakin Professor Svetha Venkatesh
ARC Laureate Fellow,
Co-Director of A2I2

Alfred Deakin Professor Kon Mouzakis
Co-Director of A2I2

Executive Team

Dr Simon Parker
Head of Operations, A2I2

Professor Rajesh Vasa
Head of Translational Research and Development

Professor Sunil Gupta
Head, AI, Optimisation and Materials Discovery

Associate Professor Santu Rana
Head, AI and Robotics

Associate Professor Truyen Tran
Head, AI, Health and Science

Associate Professor Shannon Ryan
Head, Defence and Space Research

Dr Scott Barnett
Deputy Head, Translational Research
and Commercialisation

Dr Antonio Giardina
Deputy Head, Translational Research
and Commercialisation

Dr Leonard Hoon
Deputy Head, Translational Research
and Commercialisation

Research Team

Professor Ted Goranson
Professorial Research Fellow

Associate Professor Payal Mukherjee
Associate Professor

Dr Vuong Le
Senior Research Lecturer

Dr Stewart Greenhill
Senior Research Fellow

Dr Thin Nguyen
Senior Research Fellow

Dr Rena Logothetis
Research Fellow

Dr Srikanth Thudumu
Research Fellow

Dr Alistair Shilton
Research Fellow

Dr Angie Simmons
Research Fellow

Dr Hung Le
Research Fellow

Dr Thao Minh Le
Research Fellow

Dr Kien Do
Research Fellow

Dr Phuoc Nguyen
Research Fellow

Dr Prishant Patil
Research Fellow

Dr Jessica Rivera Villicana
Research Fellow

Dr Tran The Hung
Research Lecturer

Dr Thommen Karimpanal George
Alfred Deakin Postdoctoral Research Fellow

Dr Dang Nguyen
Alfred Deakin Postdoctoral Research Fellow

Dr Romero Almeida Barata De Morais
Associate Research Fellow

Dr Dung Nguyen
Associate Research Fellow

Dr Julian Berk
Associate Research Fellow

Mr Stephan Jacobs
Associate Research Fellow

Dr Priya Rani
Associate Research Fellow

Dr Haripriya Nguyen
Associate Research Fellow

Mr Mahdi Babaei
Associate Research Fellow

Mr Yashvir Grewal
Associate Research Fellow

Dr Adrian Bingham
Associate Research Fellow

Dr Anupam Chaudhuri
Associate Research Fellow

Dr Arun Kumar Anjanapura Venkatesh
Associate Research Fellow

Dr Zafaryab Rasool
Associate Research Fellow

Mr Simon Vajda
Associate Research Fellow

Dr Praful Hambadre
Associate Research Fellow

Mr Chaitanya Lukka
Research Engineer

Mr Mahad Rashid
Research Engineer

Mr Hung Du
Research Engineer

Ms Shangeetha Sivasothy
Casual Research Fellow

Ms Manisha Senadeera
Research Assistant

Mr Neeraj Mohan
Research Assistant

Mr Johnny Assaf
Research Assistant

Ms Reshma Mathew
Casual Research Assistant

Ms Hanna Navissi
Casual Research Assistant

Engineering Team

Mr Allan Jones
General Manager, Software Engineering

Mr Michael Abbott
Principal Software Engineer

Mr Alasdair Grant
Principal Software Engineer

Mr Damian Bailey
Principal Software Engineer

Mr Stuart Cameron
Principal Software Engineer

Mr Rohan Liston
Principal Software Engineer

Mr Andrew Vouliotis
Product Manager

Mr Joost Funke Kupper
Senior Software Engineer

Mr Quentin Cooke
Senior Software Engineer

Mr David Willie
Senior Software Engineer

Mr Blake Preston
Senior Software Engineer

Engineering Team (continued)

Mr Stephen Rosman
Senior Software Engineer

Mr Stefanus Kurniawan
Software Engineer

Mr Akash Agarwal
Software Engineer

Mr Robert Hyndman
Software Engineer

Ms Tanya Frank
Software Engineer

Mr Taylan Selvi
Software Engineer

Mr Johnahan van Zyl
Software Engineer

Mrs Namrata Mehare
Software Engineer

Mr Richard Alysandratos
Software Engineer

Mr Rhys Hill
Software Engineer

Mr Justin John
Software Engineer

Mr Rodney Pilgrim
Software Engineer

Mr Zachary Brannelly
Software Engineer

Mr Nicholas Manoglou
Software Engineer

Mr Matthew Power
Junior Software Engineer

Mr Michael Elliot-Jones
Junior Software Engineer

Mr Sankhya Singh
Junior Software Engineer

Operations Team

Mr Jeff Bilman
Institute Manager

Mr Peter Dafnas
Senior Partnerships Manager

Mr Kevin Hoon
ARC Hub Manager

Dr Trang Tran
HDR Support Coordinator

Mrs Kim Ngan Pham
Project Manager, Research Grants and Contracts

Ms Natasha Beavis
Executive Officer

Ms Shweta Gupta
Administrative Officer

Mr Anthony Vouliotis
Business Analyst

Ms Monique Ellemor
Research Administrator

Mr Joshua Cole
AI Systems Officer

Mr Aiden Molluso
AI Systems Officer

Mrs Davina Montgomery
Communications Coordinator

Ms Ananya Krishna
Casual Professional

Mr Robert Ruge
Casual Professional

Mr Liam Shuster
Casual Professional

Mr Ammar Nouredin
Casual Intern

HDR Students

HDR Commencements in 2022

Mr Xuan-Tuyen Tran
Mr Maxence Pierre Henri Hussonnois
Mr Aleksandar Pasquini
Mr Shereef Ajsal
Mr Khanh-Toan Nguyen
Mr Tri-Nhan Vo
Mr Pham Van Linh Le
Mr Nam-Giang Ngo
Mr Quang-Duy Tran

HDR Completions

Mr Thanh-Tung Hoang
Mr Arun Kumar Anjanapura Venkatesh
Mr Dung Nguyen
Mr Rodney Pilgrim

Ongoing HDR students (by December 2022)

Mr Quang Duy Tran
Mr Tri Nhan Vo
Mr Nam Giang Ngo
Mr Pham Van Linh Le
Mr Ajsal Shereef Palattuparambil
Mr Maxence Pierre Henri Hussonnois
Mr Gia Hy Nguyen
Mr Xuan Tuyen Tran
Mr Khanh Toan Nguyen
Mr Aleksandar Pasquini
Ms Ragja Palakkadavath
Mr Kishan Reddy Nagiredla
Mr Tiru Naga Raja Venkat Munagala
Mr Trong Dat Phan
Mr Nguyen Thai Bao Duong
Mr Tien Kha Pham
Mr Hoang Anh Pham
Mr Banibrata Ghosh
Ms Manisha Senadeera
Mr Azharuddin Mohammed
Ms Hala Abdelkader Mahmoud Abdelkader
Mr Matthew Van Zyl
Mr Jai Kannan
Mr Tuan Dung Lai
Mr Sheng Fung Wong
Mr Hoang Dang
Mr Duy Hung Tran
Mr Ben Joseph Philip
Mr Buddhika Semage
Ms Shangeetha Sivasothy
Mr Xuan Duc Nguyen
Mr Simon Vajda
Mr Minh Tri Nguyen

Events

DATE FROM	DATE TO	EVENT	WHO ATTENDED
21/02/2022	25/02/2022	ARC ITRP Symposium for Hub Managers (Brisbane)	Kevin Hoon
25/02/2022	26/02/2022	AAAI Conference on Innovative Applications of Artificial Intelligence	Dr Alistair Shilton, Dr Hun Le
28/03/2022	30/03/2022	The 25th International Conference on Artificial Intelligence and Statistics	Dr Hung Tran The
31/03/2022		DST Field Trials in Adelaide	Michael Abbott, Damian Bailey
31/03/2022		Deakin research finance team presentation	Prof Kon Mouzakis
1/04/2022		2022 Women in AI Awards	Prof Kon Mouzakis, Dr Jess Rivera Villicana, Dr Priya Rani, Dr Rena Logothetis, Shangeetha Sivasothy
7/04/2022		Technical presentation at AI4Space workshop	A/ Prof Shannon Ryan
8/04/2022		Defence College presentation Topic - Leveraging AI in a contemporary military context	Prof Rajesh Vasa
11/04/2022	13/04/2022	1st Conference on Causal Learning and Reasoning (CLeaR 2022)	Bau Doang
21/04/2022		ARC DEL Hub Advisory Board meeting	Prof Kon Mouzakis, Prof Svetha Venkatesh, Kevin Hoon, Peter Dafnas
26/04/2022		Geelong Manufacturing Council Technology & Innovation Summit	Prof Rajesh Vasa (presenter) Scott (presenter)
4/05/2022		Keynote at Kinetic Space Safety Workshop	A/ Prof Shannon Ryan
10/05/2022	12/05/2022	Technical presentation at International Symposium on Ballistics	A/ Prof Shannon Ryan
10/05/2022	12/05/2022	International Semantic Web Conference (ISWC 2022)	Dr Angie Simmons
16/05/2022	19/05/2022	26th Pacific-Asia Conference on Knowledge Discovery and Data Mining	Dr Preeti Gopal
17/05/2022		Visit to A2I2 by MineExcellence	Prof Kon Mouzakis, Prof Rajesh Vasa, Dr Srikanth Thudumu

DATE FROM	DATE TO	EVENT	WHO ATTENDED
18/05/2022		2022 Indo-Pacific Maritime Exposition (Sydney)	Prof Rajesh Vasa (presenter)
19/05/2022		Guest speaker at Black Dog Institute (Sydney) Topic - Conductor architecture	Dr Leonard Hoon
23/05/2022		Guest speaker at MARCS Institute, Western Sydney University (Sydney)	Dr Leonard Hoon
6/06/2022		Invited Presenter, Jericho Disruptive Innovations at Royal Australian Air Force	Prof Sunil Gupta
16/06/2022		Ministerial visit to A2I2 - Trailblazer announcement	Gladys Liu MP, Stuart Robert MP, VC Iain Martin, DVCR Julie Owens
16/06/2022		Workshop with Black Dog Institute (Sydney)	Dr Leonard Hoon, Dr Scott Barnett
19/06/2022	24/06/2022	Presentation at 2022 Conference on Computer Vision and Pattern Recognition	Duy Hung Tran
28/06/2022		A2I2 Advisory Board Meeting	Prof Kon Mouzakis, Prof Svetha Venkatesh, Dr Simon Parker Jai Kannan (presenting) Duy Hung Tran (presenting)
28/06/2022		Visit to A2I2 Wairu Ponds by ONR (Office of Naval Research Global) and US Army International Technology Center	Prof Kon Mouzakis, Prof Svetha Venkatesh, Dr Simon Parker, A/Prof Shannon Ryan
12/07/2022		Meeting with Nobel Laureate Stanley Whittington	Prof Kon Mouzakis
18/07/2022		A2I2 3MT Competition	A2I2 HDR students and supervisors
18/07/2022	23/07/2023	Presentation at IEEE World Congress on Computational Intelligence - IJCNN	Dr Minh Tri Nguyen
19/07/2022		Bionics Institute 2022 Graeme Clark Oration Dinner Melbourne Convention & Exhibition Centre	Prof Kon Mouzakis, Dr Rena Logothetis, Dr Priya Rani
19/07/2022	22/07/2022	AISA Australian Cyber Conference 2022	Dr Leonard Hoon
4/08/2022	5/08/2022	2022 IEEE eScience International Conference (Salt Lake City, USA)	Dr Rena Logothetis

DATE FROM	DATE TO	EVENT	WHO ATTENDED
15/08/2022		IFM Board Meeting	Prof Kon Mouzakis
16/08/2022		AmCham Lunch - 2022 A Space Odyssey, The Commonwealth Club (Canberra)	Prof Kon Mouzakis
21/08/2022	25/08/2022	26th International Conference on Pattern Recognition	Laknath Semage
1/09/2022		2022 ADSTAR Conference	Dr Simon Parker, Dr Antonio Giardina, A/Prof Shannon Ryan
12/09/2022	13/09/2022	2022 ARC DELH Symposium, Burwood	Prof Kon Mouzakis, Prof Rajesh Vasa, Andrew Vouliotis, Kevin Hoon, Dr Leonard Hoon, Dr Scott Barnett, Peter Dafnas, Dr Angie Simmons
21/09/2022		Keynote at Hypervelocity Impact Symposium	A/Prof Shannon Ryan
26/09/2022		ATSE/AAS Webinar on Innovation in AI Powered Technologies	Prof Svetha Venkatesh
1/10/2022		2022 IEEE Women in Engineering Leadership Summit Monash University (Clayton)	Dr Priya Rani
6/10/2022		Invited talk at National Economics University	A Prof Truyen Tran
11/10/2022		2022 ESEC-FSE Conference (Singapore)	Dr Scott Barnett
14/10/2022		Keynote talk at Australian Academy of Health and Medical Sciences	Prof Svetha Venkatesh
23/10/2022	27/10/2022	European Conference on Computer Vision (ECCV)	Hoang-Anh Pham
27/10/2022		Invited Speaker at Technology in Psychiatry Summit 2022	Prof Sunil Gupta
4/11/2022		Invited Speaker at Clinical Trials @ Deakin	Prof Sunil Gupta
8/11/2022	9/11/2022	Opening of the Scale Facilitation Office at One World Trade Centre (New York City)	Prof Kon Mouzakis, Prof Rajesh Vasa
14/11/2022	18/11/2022	Visit to C4V Battery Production Hub (Endicott New York)	Prof Kon Mouzakis, Prof Rajesh Vasa
14/11/2022		Invited Speaker, Department of Health, Victoria	Prof Sunil Gupta
21/11/2022		Keynote, The 16th International Conference on Advanced Computing and Analytics	A/Prof Truyen Tran
24/11/2022		Guest Lecture at the Defence College, Canberra	Prof Kon Mouzakis, Prof Rajesh Vasa

DATE FROM	DATE TO	EVENT	WHO ATTENDED
24/11/2022		IFM Annual Conference & Dinner	Prof Kon Mouzakis
28/11/2022	30/11/2022	Advanced Data Mining and Applications 2022	Azhar Mohammed
29/11/2022	3/12/2022	Thirty-sixth Conference on Neural Information Processing Systems (NeurIPS) - 5 papers presented	Dr Hung Le, Dr Kien Do, Dr Hung Tran The, Dr Arun Kumar, Ragja Palakkadavath
30/11/2022		Presentation at IEEE International Conference on Data Mining (ICDM)	Dr Thin Nguyen
3/12/2022		Deakin Uni Cyber Security Board meeting	Prof Kon Mouzakis
4/12/2022	8/12/2022	23rd International Society for Music Information Retrieval Conference (ISMIR 2022)	Dr Stewart Greenhill
5/12/2022	8/12/2022	35th Australasian Joint Conference on Artificial Intelligence	Dr Hung Le
5/12/2022		2022 ICARS (Intelligence Community Academic Research Symposium)	Dr Angie Simmons

Publications

Conference papers

1. Kannan, J., Barnett, S., Simmons, A., Cruz, L., & Agarwal, A. (2022). **MLSmellHound: A Context-Aware Code Analysis Tool.**

Conference name: 2022 IEEE/ACM 44rd International Conference on Software Engineering: New Ideas and Emerging Results (ICSE-NIER)

Conference ranking: A*

Abstract: Meeting the rise of industry demand to incorporate machine learning (ML) components into software systems requires interdisciplinary teams contributing to a shared code base. To maintain consistency, reduce defects and ensure maintainability, developers use code analysis tools to aid them in identifying defects and maintaining standards. With the inclusion of machine learning, tools must account for the cultural differences within the teams which manifests as multiple programming languages, and conflicting definitions and objectives. Existing tools fail to identify these cultural differences and are geared towards software engineering which reduces their adoption in ML projects. In our approach we attempt to resolve this problem by exploring the use of context which includes i) purpose of the source code, ii) technical domain, iii) problem domain, iv) team norms, v) operational environment, and vi) development lifecycle stage to provide contextualised error reporting for code analysis. To demonstrate our approach, we adapt Pylint as an example and apply a set of contextual transformations to the linting results based on the domain of individual project files under analysis. This allows for contextualised and meaningful error reporting for the end user.

2. Nguyen, T. T., Gupta, S., Nguyen, A. T., & Venkatesh, S. (2022). **Offline Neural Contextual Bandits: Pessimism, Optimization and Generalization.**

Conference name: The Tenth International Conference on Learning Representations (ICLR)

Conference ranking: A*

Abstract: Offline policy learning (OPL) leverages existing data collected a priori for policy optimization without any active exploration. Despite the prevalence and recent interest in this problem, its theoretical and algorithmic foundations in function approximation settings remain under-developed. In this paper, we consider this problem on the axes of distributional shift, optimization, and generalization in offline contextual bandits with neural networks. In particular, we propose a provably efficient offline contextual bandit with neural network function approximation that does not require any functional assumption on the reward. We show that our method provably generalizes over unseen contexts under a milder condition for distributional shift than the existing OPL works. Notably, unlike any other OPL method, our method learns from the offline data in an online manner using stochastic gradient descent, allowing us to leverage the benefits of online learning into an offline setting. Moreover, we show that our method is more computationally efficient and has a better dependence on the effective dimension of the neural network than an online counterpart. Finally, we demonstrate the empirical effectiveness of our method in a range of synthetic and real-world OPL problems.

3. Tran-The, H., Gupta, S., Rana, S., & Venkatesh, S. (2022). Regret Bounds for Expected Improvement Algorithms in Gaussian Process Bandit Optimization.

Conference name: The 25th International Conference on Artificial Intelligence and Statistics (AISTATS)

Conference ranking: A

Abstract: The expected improvement (EI) algorithm is one of the most popular strategies for optimization under uncertainty due to its simplicity and efficiency. Despite its popularity, the theoretical aspects of this algorithm have not been properly analyzed. In particular, whether in the noisy setting, the EI strategy with a standard incumbent converges is still an open question of the Gaussian process bandit optimization problem. We aim to answer this question by proposing a variant of EI with a standard incumbent defined via the GP predictive mean. We prove that our algorithm converges, and achieves a cumulative regret bound of $(Y_T - T)$, where Y_T is the maximum information gain between T observations and the Gaussian process model. Based on this variant of EI, we further propose an algorithm called Improved GP-EI that converges faster than previous counterparts. In particular, our proposed variants of EI do not require the knowledge of the RKHS norm and the noise's sub-Gaussianity parameter as in previous works. Empirical validation in our paper demonstrates the effectiveness of our algorithms compared to several baselines.

4. Shilton, A., Rana, S., Gupta, S., & Venkatesh, S. (2022). TRF: Learning Kernels with Tuned Random Features.

Conference name: Thirty-Sixth Conference on Artificial Intelligence (AAAI)

Conference ranking: A*

Abstract: Random Fourier features (RFF) are a popular set of tools for constructing low-dimensional approximations of translation-invariant kernels, allowing kernel methods to be scaled to big data. Apart from their computational advantages, by working in the spectral domain random Fourier features expose the translation invariant kernel as a density function that may, in principle, be manipulated directly to tune the kernel. In this paper we propose selecting the density function from a reproducing kernel Hilbert space to allow us to search the space of all translation-invariant kernels. Our approach, which we call tuned random features (TRF), achieves this by approximating the density function as the RKHS-norm regularised least-squares best fit to an unknown "true" optimal density function, resulting in a RFF formulation where kernel selection is reduced to regularised risk minimisation with a novel regulariser. We derive bounds on the Rademacher complexity for our method showing that our random features approximation method converges to optimal kernel selection in the large N, D limit. Finally, we prove experimental results for a variety of real-world learning problems, demonstrating the performance of our approach compared to comparable methods.

5. Pham, K., Le, H., Ngo, M., Tran, T., Ho, B., & Venkatesh, S. (2022). Generative Pseudo-Inverse Memory.

Conference name: International Conference on Learning Representations (ICLR)

Conference ranking: A*

Abstract: We propose Generative Pseudo-Inverse Memory (GPM), a class of deep generative memory models that are fast to write in and read out. Memory operations are recast as seeking robust solutions of linear systems, which naturally lead to the use of matrix pseudo-inverses. The pseudo-inverses are iteratively approximated, with practical computation complexity of almost $O(1)$. We prove theoretically and verify empirically that our model can retrieve exactly what have been written to the memory under mild conditions. A key capability of GPM is iterative reading, during which the attractor dynamics towards fixed points are enabled, allowing the model to iteratively improve sample quality in denoising and generating. More impressively, GPM can store a large amount of data while maintaining key abilities of accurate retrieving of stored patterns, denoising of corrupted data and generating novel samples. Empirically we demonstrate the efficiency and versatility of GPM on a comprehensive suite of experiments involving binarized MNIST, binarized Omniglot, Fashion/MNIST, CIFAR10 & CIFAR100 and CelebA.

6. Philip, B., Abdelrazek, M., Barnett, S., Bonti, A., & Grundy, J. (2022). Towards Better mHealth Apps: Understanding Current Challenges and User Expectations.

Conference name: 2022 IEEE/ACM 9th International Conference on Mobile Software Engineering and Systems (MobileSoft),

Conference ranking: N/A

Abstract: Mobile health (mHealth) apps have become ubiquitous and offer several different features to provide a better health outcome for end-users. While the availability of thousands of mHealth apps offers a great many options for consumers, they also introduce several challenges if needing to use more than one app. We designed an anonymous survey based on constructs of the Technology Acceptance Model (TAM), the Mobile App Rating Scale (MARS) and the Value Proposition Canvas to collect data on the user experience (UX) around these challenges. We surveyed 70 people over the age of 18 having experience with mHealth apps and found issues such as limited customizability, unwanted and redundant features, and data entry challenges that lead to a degraded UX overall. These challenges are also valid from a developer's point of view where they spend significant efforts in developing these redundant or unneeded features for more than one platform. In this paper, we discuss these user challenges and emerging implications for mHealth app developers.

7. Le, H., Abdolshah, M., George, T. K., Do, K., Nguyen, D., & Venkatesh, S. (2022). Episodic Policy Gradient Training.

Conference name: Thirty-Sixth Conference on Artificial Intelligence (AAAI)

Conference ranking: A*

Abstract: We introduce a novel training procedure for policy gradient methods wherein episodic memory is used to optimize the hyperparameters of reinforcement learning algorithms on-the-fly. Unlike other hyperparameter searches, we formulate hyperparameter scheduling as a standard Markov Decision Process and use episodic memory to store the outcome of used hyperparameters and their training contexts. At any policy update step, the policy learner refers to the stored experiences, and adaptively reconfigures its learning algorithm with the new hyperparameters determined by the memory. This mechanism, dubbed as Episodic Policy Gradient Training (EPGT), enables an episodic learning process, and jointly learns the policy and the learning algorithm's hyperparameters within a single run. Experimental results on both continuous and discrete environments demonstrate the advantage of using the proposed method in boosting the performance of various policy gradient algorithms.

8. Le, H., & Venkatesh, S. (2022). Neurocoder: Learning General-purpose Computation Using Stored Neural Programs.

Conference name: International Conference on Machine Learning (ICLR)

Conference ranking: A*

Abstract: Artificial Neural Networks are uniquely adroit at machine learning by processing data through a network of artificial neurons. The inter-neuronal connection weights represent the learnt Neural Program that instructs the network on how to compute the data. However, without an external memory to store Neural Programs, they are restricted to only one, overwriting learnt programs when trained on new data. This is functionally equivalent to a special-purpose computer. Here we design Neurocoder, an entirely new class of general-purpose conditional computational machines in which the neural network "codes" itself in a data-responsive way by composing relevant programs from a set of shareable, modular programs. This can be considered analogous to building Lego structures from simple Lego bricks. Notably, our bricks change their shape through learning. External memory is used to create, store and retrieve modular programs. Like today's stored-program computers, Neurocoder can now access diverse programs to process different data. Unlike manually crafted computer programs, Neurocoder creates programs through training. Integrating Neurocoder into current neural architectures, we demonstrate new capacity to learn modular programs, handle severe pattern shifts and remember old programs as new ones are learnt, and show substantial performance improvement in solving object recognition, playing video games and continual learning tasks. Such integration with Neurocoder increases the computation capability of any current neural network and endows it with entirely new capacity to reuse simple programs to build complex ones. For the first time a Neural Program is treated as a datum in memory, paving the ways for modular, recursive and procedural neural programming.

9. Nguyen, D., Nguyen, P., Le, H., Do, K., Venkatesh, S., & Tran, T. (2022). Learning theory of mind via dynamic traits attribution.

Conference name: the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS)

Conference ranking: A*

Abstract: Machine learning of Theory of Mind (ToM) is essential to build social agents that co-live with humans and other agents. This capacity, once acquired, will help machines infer the mental states of others from observed contextual action trajectories, enabling future prediction of goals, intention, actions and successor representations. The underlying mechanism for such a prediction remains unclear, however. Inspired by the observation that humans often infer the character traits of others, then use it to explain behaviour, we propose a new neural ToM architecture that learns to generate a latent trait vector of an actor from the past trajectories. This trait vector then multiplicatively modulates the prediction mechanism via a 'fast weights' scheme in the prediction neural network, which reads the current context and predicts the behaviour. We empirically show that the fast weights provide a good inductive bias to model the character traits of agents and hence improves mindreading ability. On the indirect assessment of false-belief understanding, the new ToM model enables more efficient helping behaviours.

10. Nguyen, D., Nguyen, P., Venkatesh, S., & Tran, T. (2022). Learning to Transfer Role Assignment Across Team Sizes.

Conference name: The 21st International Conference on Autonomous Agents and Multiagent Systems

Conference ranking: A*

Abstract: Multi-agent reinforcement learning holds the key for solving complex tasks that demand the coordination of learning agents. However, strong coordination often leads to expensive exploration over the exponentially large state-action space. A powerful approach is to decompose team works into roles, which are ideally assigned to agents with the relevant skills. Training agents to adaptively choose and play emerging roles in a team thus allows the team to scale to complex tasks and quickly adapt to changing environments. These promises, however, have not been fully realised by current role-based multi-agent reinforcement learning methods as they assume either a pre-defined role structure or a fixed team size. We propose a framework to learn role assignment and transfer across team sizes. In particular, we train a role assignment network for small teams by demonstration and transfer the network to larger teams, which continue to learn through interaction with the environment. We demonstrate that re-using the role-based credit assignment structure can foster the learning process of larger reinforcement learning teams to achieve tasks requiring different roles. Our proposal outperforms competing techniques in enriched role-enforcing Prey-Predator games and in new scenarios in the StarCraft II Micro-Management benchmark.

11. Gopal, P., Gupta, S., Rana, S., Le, V., Nguyen, T., & Venkatesh, S. (2022). Real-Time Skill Discovery in Intelligent Virtual Assistants.

Conference name: Advances in Knowledge Discovery and Data Mining: 26th Pacific-Asia Conference (PAKDD).

Conference ranking: A

Abstract: Solution to many real-world problems often involve the use of expert-level knowledge from various specializations. Such inter-disciplinary problems are usually divided into tasks which are then assigned to a set of bots, each specialized in a particular skill. Supervised selection of the right bot each time is cumbersome and not scalable. Hence there is a need for an AI system that identifies the type of task and assigns it to a suitably trained bot. Challenges arise in non-stationary environments when the cost of choosing different bots vary or the bots themselves might evolve in their skills. In this paper, as in Conversational AI, a number of bots are at our disposal, each of which is trained to handle (i.e., answer) a specific type of question in a conversation. We develop a meta-algorithm that learns about capabilities (Skill Discovery) of the available bots in real-time and appropriately selects a relevant bot for the question at hand. We present contextual bandits as a solution in this setting and introduce gradual finetuning of query information to improve Skill Discovery. Using two popular datasets from conversational AI: CoQA and SQuAD, we show promising results of our method on non-stationary environments.

12. Tran, H., Le, V., Venkatesh, S., & Tran, T. (2022). Persistent-Transient Duality in Human Behavior Modeling.

Conference name: The IEEE/CVF Conference on Computer Vision and Pattern Recognition (Precognition Workshop)

Conference ranking: A*

Abstract: We propose to model the persistent-transient duality in human behavior using a parent-child multi-channel neural network, which features a parent persistent channel that manages the global dynamics and children transient channels that are initiated and terminated on-demand to handle detailed interactive actions. The short-lived transient sessions are managed by a proposed Transient Switch. The neural framework is trained to discover the structure of the duality automatically. Our model shows superior performances in human-object interaction motion prediction.

13. Senadeera, M., Karimpanal, T. G., Gupta, S., & Rana, S. (2022). Sympathy-based Reinforcement Learning Agents.

Conference name: The 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS)

Conference ranking: A*

Abstract: As artificial agents become increasingly prevalent in our daily lives, it becomes imperative to equip them with an awareness of societal norms; specifically, the ability to account for and be considerate towards others they may cohabit with. In this work, we explore the ability for an agent trained through reinforcement learning to exhibit sympathetic behaviours towards another (independent) agent in the environment. We propose to achieve such behaviours by first inferring the reward function of the independent agent, through inverse reinforcement learning, and subsequently learning a policy based on a sympathetic reward function - a convex combination of the inferred rewards and the agent's own rewards. The corresponding weighting is determined by a sympathy function which is computed based on the estimated return of the agent's current action relative to that of all possible actions it could have taken. We evaluate our approach on adversarial as well as assistive environment settings, and demonstrate the ability of our sympathetic agent to perform well at its own goal, while simultaneously giving due consideration to another agent in its environment. We also empirically examine and report the sensitivity of our agent's performance to the hyperparameters introduced in our proposed framework.

14. Nguyen, T., Mohammed, A., Duong, B., & Nguyen, T. (2022, July). Causality-aided Recommendation Systems.

Conference name: International Joint Conference on Neural Networks (IJCNN)

Conference ranking: B

Abstract: Recommendation systems often face the challenge of cold-start problem where the two parties are completely new to the systems. Previous works introduce the integration of side information of items and users into the systems so that they can learn a representation of the inputs with better generalization, thus potentially solving the cold-start problem. However, using users' or items' side information brings another problem to deep learning based collaborative filtering models known as spurious correlations. In this case, deep learning models fail to adequately generalize and simply capture the idiosyncrasies of the training datasets. In short, deep learning models find a 'short-cut' created by the hidden bias in the datasets and make the spurious correlations between specific features and the outcome. In this work, we propose the counterfactual constraint on user/item attributes, allowing the systems to avoid hidden bias in the datasets, thus having the robustness and generalization. The goal is to discourage the disruptive distribution change in the prediction when the distribution shift occurs in the user/item attribute spaces. The proposed model compares favorably with conventional baseline on both synthetic and real-world datasets.

15. Duong, B., & Nguyen, T. (2022). Bivariate causal discovery via conditional divergence.

Conference name: Conference on Causal Learning and Reasoning.

Conference ranking: N/A

Abstract: Telling apart cause and effect is a fundamental problem across many science disciplines. However, the randomized controlled trial, which is the golden-standard solution for this, is not always physically feasible or ethical. Therefore, we can only rely on passively observational data in such cases, making the problem highly challenging. Inspired by the observation that the conditional distribution of effect given cause, also known as the causal mechanism, is typically invariant in shape, we aim to capture the mechanism through estimating the stability of the conditional distribution. In particular, based on the inverse of stability—the divergence—we propose Conditional Divergence based Causal Inference (CDCI), a novel algorithm for detecting causal direction in purely observational data. By doing this, we can relax multiple strict assumptions commonly adopted in the causal discovery literature, including functional form and noise model. The proposed approach is generic and applicable to arbitrary measures of distribution divergence. The effectiveness of our method is demonstrated on a variety of both synthetic and real data sets, which compares favorably with existing state-of-the-art methods.

16. Pham, H. A., Le, T. M., Le, V., Phuong, T. M., & Tran, T. (2022, October). Video dialog as conversation about objects living in space-time.

Conference name: European Conference on Computer Vision (ECCV)

Conference ranking: A*

Abstract: It would be a technological feat to be able to create a system that can hold a meaningful conversation with humans about what they watch. A setup toward that goal is presented as a video dialog task, where the system is asked to generate natural utterances in response to a question in an ongoing dialog. The task poses great visual, linguistic, and reasoning challenges that cannot be easily overcome without an appropriate representation scheme over video and dialog that supports high-level reasoning. To tackle these challenges we present a new object-centric framework for video dialog that supports neural reasoning dubbed COST—which stands for Conversation about Objects in Space-Time. Here dynamic space-time visual content in videos is first parsed into object trajectories. Given this video abstraction, COST maintains and tracks object-associated dialog states, which are updated upon receiving new questions. Object interactions are dynamically and conditionally inferred for each question, and these serve as the basis for relational reasoning among them. COST also maintains a history of previous answers, and this allows retrieval of relevant object-centric information to enrich the answer forming process. Language production then proceeds in a step-wise manner, taking the context of the current utterance, the existing dialog, and the current question. We evaluate COST on the AVSD test splits (DSTC7 and DSTC8), demonstrating its competitiveness against state-of-the-arts.

17. Nguyen, D., Gupta, S., Do, K., & Venkatesh, S. (2022, October). Black-box few-shot knowledge distillation.

Conference name: European Conference on Computer Vision (ECCV)

Conference ranking: A*

Abstract: Knowledge distillation (KD) is an efficient approach to transfer the knowledge from a large “teacher” network to a smaller “student” network. Traditional KD methods require lots of labeled training samples and a white-box teacher (parameters are accessible) to train a good student. However, these resources are not always available in real-world applications. The distillation process often happens at an external party side where we do not have access to much data, and the teacher does not disclose its parameters due to security and privacy concerns. To overcome these challenges, we propose a black-box few-shot KD method to train the student with few unlabeled training samples and a black-box teacher. Our main idea is to expand the training set by generating a diverse set of out-of-distribution synthetic images using MixUp and a conditional variational auto-encoder. These synthetic images along with their labels obtained from the teacher are used to train the student. We conduct extensive experiments to show that our method significantly outperforms recent SOTA few/zero-shot KD methods on image classification tasks. The code and models are available at <https://github.com/nphdang/FS-BBT>.

18. Patil, P. W., Gupta, S., Rana, S., & Venkatesh, S. (2022, October). Video Restoration Framework and Its Meta-adaptations to Data-Poor Conditions.

Conference name: European Conference on Computer Vision (ECCV)

Conference ranking: A*

Abstract: Restoration of weather degraded videos is a challenging problem due to diverse weather conditions e.g., rain, haze, snow, etc.. Existing works handle video restoration for each weather using a different custom-designed architecture. This approach has many limitations. First, a custom-designed architecture for each weather condition requires domain-specific knowledge. Second, disparate network architectures across weather conditions prevent easy knowledge transfer to novel weather conditions where we do not have a lot of data to train a model from scratch. For example, while there is a lot of common knowledge to exploit between the models of different weather conditions at day or night time, it is difficult to do such adaptation. To this end, we propose a generic architecture that is effective for any weather condition due to the ability to extract robust feature maps without any domain-specific knowledge. This is achieved by novel components: spatio-temporal feature modulation, multi-level feature aggregation, and recurrent guidance decoder. Next, we propose a meta-learning based adaptation of our deep architecture to the restoration of videos in data-poor conditions (night-time videos). We show comprehensive results on video de-hazing and de-raining datasets in addition to the meta-learning based adaptation results on night-time video restoration tasks. Our results clearly outperform the state-of-the-art weather degraded video restoration methods. The source code is available at https://github.com/pwp1208/Meta_Video_Restoration.

19. Do, K., Harikumar, H., Le, H., Nguyen, D., Tran, T., Rana, S., ... & Venkatesh, S. (2022, October). Towards Effective and Robust Neural Trojan Defenses via Input Filtering.

Conference name: European Conference on Computer Vision (ECCV)

Conference ranking: A*

Abstract: Trojan attacks on deep neural networks are both dangerous and surreptitious. Over the past few years, Trojan attacks have advanced from using only a single input-agnostic trigger and targeting only one class to using multiple, input-specific triggers and targeting multiple classes. However, Trojan defenses have not caught up with this development. Most defense methods still make inadequate assumptions about Trojan triggers and target classes, thus, can be easily circumvented by modern Trojan attacks. To deal with this problem, we propose two novel “filtering” defenses called Variational Input Filtering (VIF) and Adversarial Input Filtering (AIF) which leverage lossy data compression and adversarial learning respectively to effectively purify potential Trojan triggers in the input at run time without making assumptions about the number of triggers/target classes or the input dependence property of triggers. In addition, we introduce a new defense mechanism called “Filtering-then-Contrasting” (FtC) which helps avoid the drop in classification accuracy on clean data caused by “filtering”, and combine it with VIF/AIF to derive new defenses of this kind. Extensive experimental results and ablation studies show that our proposed defenses significantly outperform well-known baseline defenses in mitigating five advanced Trojan attacks including two recent state-of-the-art while being quite robust to small amounts of training data and large-norm triggers.

20. Mohammed, A., Nguyen, D., Duong, B., & Nguyen, T. (2022, November). Efficient Classification with Counterfactual Reasoning and Active Learning.

Conference name: 14th Asian Conference on Intelligent Information and Database Systems (ACIIDS)

Conference ranking: B

Abstract: Data augmentation is one of the most successful techniques to improve the classification accuracy of machine learning models in computer vision. However, applying data augmentation to tabular data is a challenging problem since it is hard to generate synthetic samples with labels. In this paper, we propose an efficient classifier with a novel data augmentation technique for tabular data. Our method called CCRAL combines causal reasoning to learn counterfactual samples for the original training samples and active learning to select useful counterfactual samples based on a region of uncertainty. By doing this, our method can maximize our model’s generalization on the unseen testing data. We validate our method analytically, and compare with the standard baselines. Our experimental results highlight that CCRAL achieves significantly better performance than those of the baselines across several real-world tabular datasets in terms of accuracy and AUC. Data and source code are available at <https://github.com/nphdang/CCRAL>.

21. Simmons, A., Vasa, R., & Giardina, A. (2022). SignalKG: Towards reasoning about the underlying causes of sensor observations.

Conference name: the ISWC 2022 Posters, Demos and Industry Tracks (CEUR Workshop Proceedings)

Conference ranking: A

Abstract: This paper demonstrates our vision for knowledge graphs that assist machines to reason about the cause of signals observed by sensors. We show how the approach allows for constructing smarter surveillance systems that reason about the most likely cause (e.g., an attacker breaking a window) of a signal rather than acting directly on the received signal without consideration for how it was produced.

22. Semage, B. L., Karimpanal, T. G., Rana, S., & Venkatesh, S. (2022, August). Fast Model-based Policy Search for Universal Policy Networks

Conference name: 26th International Conference on Pattern Recognition (ICPR)

Conference ranking: B

Abstract: Adapting an agent's behaviour to new environments has been one of the primary focus areas of physics-based reinforcement learning. Although recent approaches such as universal policy networks partially address this issue by enabling the storage of multiple policies trained in simulation on a wide range of dynamic/latent factors, efficiently identifying the most appropriate policy for a given environment remains a challenge. In this work, we propose a Gaussian Process-based prior learned in simulation, that captures the likely performance of a policy when transferred to a previously unseen environment. We integrate this prior with a Bayesian Optimisation-based policy search process to improve the efficiency of identifying the most appropriate policy from the universal policy network. We empirically evaluate our approach in a range of continuous and discrete control environments and show that it outperforms other competing baselines.

23. Semage, B. L., Karimpanal, T. G., Rana, S., & Venkatesh, S. (2022, August). Uncertainty Aware System Identification with Universal Policies.

Conference name: 26th International Conference on Pattern Recognition (ICPR)

Conference ranking: B

Abstract: Sim2real transfer is primarily concerned with transferring policies trained in simulation to potentially noisy real world environments. A common problem associated with sim2real transfer is estimating the real-world environmental parameters to ground the simulated environment to. Although existing methods such as Domain Randomisation (DR) can produce robust policies by sampling from a distribution of parameters during training, there is no established method for identifying the parameters of the corresponding distribution for a given real-world setting. In this work, we propose Uncertainty-aware policy search (UncAPS), where we use Universal Policy Network (UPN) to store simulation-trained task-specific policies across the full range of environmental parameters and then subsequently employ robust Bayesian optimisation to craft robust policies for the given environment by combining relevant UPN policies in a DR like fashion. Such policy-driven grounding is expected to be more efficient as it estimates only task-relevant sets of parameters. Further, we also account for the estimation uncertainties in the search process to produce policies that are robust against both aleatoric and epistemic uncertainties. We empirically evaluate our approach in a range of noisy, continuous control environments, and show its improved performance compared to competing baselines.

24. AV, A. K., Rana, S., Shilton, A., & Venkatesh, S. (2022). Human-AI Collaborative Bayesian Optimisation

Conference name: Advances in Neural Information Processing Systems

Conference ranking: A*

Abstract: Abstract Human-AI collaboration looks at harnessing the complementary strengths of both humans and AI. We propose a new method for human-AI collaboration in Bayesian optimisation where the optimum is mainly pursued by the Bayesian optimisation algorithm following complex computation, whilst getting occasional help from the accompanying expert having a deeper knowledge of the underlying physical phenomenon. We expect experts to have some understanding of the correlation structures of the experimental system, but not the location of the optimum. The expert provides feedback by either changing the current recommendation or providing her belief on the good and bad regions of the search space based on the current observations. Our proposed method takes such feedback to build a model that aligns with the expert's model and then uses it for optimisation. We provide theoretical underpinning on why such an approach may be more efficient than the one without expert's feedback. The empirical results show the robustness and superiority of our method with promising efficiency gains.

25. Le, T. H., Karimpanal George, T., Abdolshah, M., Nguyen, D., Do, K., Gupta, S., & Venkatesh, S. (2022). Learning to Constrain Policy Optimization with Virtual Trust Region.

Conference name: Advances in Neural Information Processing Systems

Conference ranking: A*

Abstract: We introduce a constrained optimization method for policy gradient reinforcement learning, which uses two trust regions to regulate each policy update. In addition to using the proximity of one single old policy as the first trust region as done by prior works, we propose forming a second trust region by constructing another virtual policy that represents a wide range of past policies. We then enforce the new policy to stay closer to the virtual policy, which is beneficial if the old policy performs poorly. We propose a mechanism to automatically build the virtual policy from a memory buffer of past policies, providing a new capability for dynamically selecting appropriate trust regions during the optimization process. Our proposed method, dubbed Memory-Constrained Policy Optimization (MCPO), is examined in diverse environments, including robotic locomotion control, navigation with sparse rewards and Atari games, consistently demonstrating competitive performance against recent on-policy constrained policy gradient methods.

26. Pham, K., Le, T. H., Ngo, M., & Tran, T. (2022). Functional Indirection Neural Estimator for Better Out-of-distribution Generalization.

Conference name: Advances in Neural Information Processing Systems

Conference ranking: A*

Abstract: The capacity to achieve out-of-distribution (OOD) generalization is a hallmark of human intelligence and yet remains out of reach for machines. This remarkable capability has been attributed to our abilities to make conceptual abstraction and analogy, and to a mechanism known as indirection, which binds two representations and uses one representation to refer to the other. Inspired by these mechanisms, we hypothesize that OOD generalization may be achieved by performing analogy-making and indirection in the functional space instead of the data space as in current methods. To realize this, we design FINE (Functional Indirection Neural Estimator), a neural framework that learns to compose functions that map data input to output on-the-fly. FINE consists of a backbone network and a trainable semantic memory of basis weight matrices. Upon seeing a new input-output data pair, FINE dynamically constructs the backbone weights by mixing the basis weights. The mixing coefficients are indirectly computed through querying a separate corresponding semantic memory using the data pair. We demonstrate empirically that FINE can strongly improve out-of-distribution generalization on IQ tasks that involve geometric transformations. In particular, we train FINE and competing models on IQ tasks using images from the MNIST, Omniglot and CIFAR100 datasets and test on tasks with unseen image classes from one or different datasets and unseen transformation rules. FINE not only achieves the best performance on all tasks but also is able to adapt to small-scale data scenarios.

27. Duong, B., & Nguyen, T. (2022, November). Conditional Independence Testing via Latent Representation Learning

Conference name: IEEE International Conference on Data Mining (ICDM)

Conference ranking: A*

Abstract: Detecting conditional independencies plays a key role in several statistical and machine learning tasks, especially in causal discovery algorithms, yet it remains a highly challenging problem due to dimensionality and complex relationships presented in data. In this study, we introduce LCIT (Latent representation based Conditional Independence Test) - a novel method for conditional independence testing based on representation learning. Our main contribution involves a hypothesis testing framework in which to test for the independence between X and Y given Z, we first learn to infer the latent representations of target variables X and Y that contain no information about the conditioning variable Z. The latent variables are then investigated for any significant remaining dependencies, which can be performed using a conventional correlation test. The empirical evaluations show that LCIT outperforms several state-of-the-art baselines consistently under different evaluation metrics, and is able to adapt really well to both non-linear and high-dimensional settings on a diverse collection of synthetic and real data sets.

28. Du, H., Thudumu, S., Singh, S., Barnett, S., Logothetis, I., Vasa, R., & Mouzakis, K. (2022, October). A Framework for Evaluating MRC Approaches with Unanswerable Questions

Conference name: IEEE 18th International Conference on e-Science (e-Science)

Conference ranking: B

Abstract: Machine reading comprehension (MRC) is a challenging task in natural language processing that demonstrates the language understanding of the machine. An approach to tackle this challenge requires the machine to answer the question about the given context when needed and abstain from answering when there is no answer. Recent works attempted to solve this challenge with various comprehensive neural network architectures for sequences such as SAN, U-Net, EQuANT, and others that were trained on the SQuAD 2.0 dataset containing unanswerable questions. However, the robustness of these approaches has not been evaluated. In this paper, we propose a data augmentation approach that converts answerable questions to unanswerable questions in the SQuAD 2.0 dataset by altering the entities in the question to its antonym from ConceptNet which is a semantic network. The augmented data is, then, fitted into the U-Net question answering model to evaluate the robustness of the model.

29. Ryan, S.; Berk, J.; Weiss, A.; Yatom, N.; Vardy, M.; Rana, S.; Greenhill, S. & Venkatesh, S. Adaptive experimental optimisation for sample-efficient armour design.

Conference name: the 32nd International Symposium on Ballistics

Conference ranking: N/A

Abstract: We present a Bayesian optimisation methodology intended to support a human expert in the design of armour systems for which limited prior knowledge/data exists and within a limiting, pre-defined experimental budget. We apply the methodology to design an armour configuration consisting of multiple plates, with multiple materials, at varying orientations and spacing, for protection against 12.7 mm APM2 and 20 mm FSP threats. The full-factorial design matrix for the defined solution space exceeds 17,500 possible solutions. With an objective to minimise system weight, we identify a solution within 102 ballistic tests (44 design iterations) that provides a weight reduction of 11.4% over expert-designed reference configurations and a mass efficiency of 1.5 relative to a monolithic RHA Class 1. The value of the demonstrated methodology is expected to increase with increasing armour (or threat) complexity.

30. Le, T. M., Le, V., Gupta, S., Venkatesh, S., & Tran, T. (2023). Guiding Visual Question Answering with Attention Priors

Conference name: the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)

Conference ranking: A

Abstract: The current success of modern visual reasoning systems is arguably attributed to cross-modality attention mechanisms. However, in deliberative reasoning such as in VQA, attention is unconstrained at each step, and thus may serve as a statistical pooling mechanism rather than a semantic operation intended to select information relevant to inference. This is because at training time, attention is only guided by a very sparse signal (i.e. the answer label) at the end of the inference chain. This causes the cross-modality attention weights to deviate from the desired visual-language bindings. To rectify this deviation, we propose to guide the attention mechanism using explicit linguistic-visual grounding. This grounding is derived by connecting structured linguistic concepts in the query to their referents among the visual objects. Here we learn the grounding from the pairing of questions and images alone, without the need for answer annotation or external grounding supervision. This grounding guides the attention mechanism inside VQA models through a duality of mechanisms: pre-training attention weight calculation and directly guiding the weights at inference time on a case-by-case basis. The resultant algorithm is capable of probing attention-based reasoning models, injecting relevant associative knowledge, and regulating the core reasoning process. This scalable enhancement improves the performance of VQA models, fortifies their robustness to limited access to supervised data, and increases interpretability.

31. Mohammed, A., Nguyen, D., Duong, B., Nichols, M., & Nguyen, T. (2022, November). Handling Missing Data with Markov Boundary.

Conference name: International Conference on Advanced Data Mining and Applications (ADMA)

Conference ranking: B

Abstract: In machine learning (ML) applications, high-quality data are very important to train a well-performed model that can provide robust predictions and responsible decisions. A common problem in ML applications e.g., healthcare is that the training dataset often consists of samples (or records) with missing values. As a result, the ML model cannot use such samples in its training phase. Handling missing data is thus an important and open research problem. In this paper, we propose a method to predict missing values by considering a causal graphical model framework. Our method exploits the Markov boundary encapsulating all necessary information about the missing variables. By utilizing the information encoded in the Markov boundary, we formulate a predictive function for each feature that has missing values to predict its missing values. Compared to existing methods, our predictive function is trained with only the features involved in the Markov boundary. To demonstrate the effectiveness of our proposed method, we compare its imputation performance with those of state-of-the-art imputation methods via a comprehensive experiment on seven real-world datasets. Our empirical results highlight that our method is significantly better than those of the baselines in terms of the imputation error thanks to its Markov information.

32. Do, K., Le, T. H., Nguyen, D., Nguyen, D., Harikumar, H., Tran, T., ... & Venkatesh, S. (2022). Momentum Adversarial Distillation: Handling Large Distribution Shifts in Data-Free Knowledge Distillation.

Conference name: Advances in Neural Information Processing Systems

Conference ranking: A*

Abstract: Data-free Knowledge Distillation (DFKD) has attracted attention recently thanks to its appealing capability of transferring knowledge from a teacher network to a student network without using training data. The main idea is to use a generator to synthesize data for training the student. As the generator gets updated, the distribution of synthetic data will change. Such distribution shift could be large if the generator and the student are trained adversarially, causing the student to forget the knowledge it acquired at the previous steps. To alleviate this problem, we propose a simple yet effective method called Momentum Adversarial Distillation (MAD) which maintains an exponential moving average (EMA) copy of the generator and uses synthetic samples from both the generator and the EMA generator to train the student. Since the EMA generator can be considered as an ensemble of the generator's old versions and often undergoes a smaller change in updates compared to the generator, training on its synthetic samples can help the student recall the past knowledge and prevent the student from adapting too quickly to the new updates of the generator. Our experiments on six benchmark datasets including big datasets like ImageNet and Places365 demonstrate the superior performance of MAD over competing methods for handling the large distribution shift problem. Our method also compares favorably to existing DFKD methods and even achieves state-of-the-art results in some cases.

33. Gupta, S., Rana, S., Truong, T., Tran-Thanh, L., & Venkatesh, S. (2022). Expected Improvement for Contextual Bandits.

Conference name: Advances in Neural Information Processing Systems

Conference ranking: A*

Abstract: The expected improvement (EI) is a popular technique to handle the tradeoff between exploration and exploitation under uncertainty. This technique has been widely used in Bayesian optimization but it is not applicable for the contextual bandit problem which is a generalization of the standard bandit and Bayesian optimization. In this paper, we initiate and study the EI technique for contextual bandits from both theoretical and practical perspectives. We propose two novel EI based algorithms, one when the reward function is assumed to be linear and the other for more general reward functions. With linear reward functions, we demonstrate that our algorithm achieves \sqrt{n} -near-optimal regret. Notably, our regret improves that of LinTS [3] by a factor d while avoiding to solve a NP-hard problem at each iteration as in LinUCB [1]. For more general reward functions which are modeled by deep neural networks, we prove that our algorithm achieves a $O^*(d^* T)$ regret, where d^* is the effective dimension of a neural tangent kernel (NTK) matrix, and T is the number of iterations. Our experiments on various benchmark datasets show that both proposed algorithms work well and consistently outperform existing approaches, especially in high dimensions.

34. Palakkadavath, R., Nguyen-Tang, T., Gupta, S., & Venkatesh, S. (2022, October). Improving Domain Generalization with Interpolation Robustness.

Conference name: NeurIPS 2022 Workshop on Distribution Shifts: Connecting Methods and Applications.

Conference ranking: N/A

Abstract: We address domain generalization (DG) by viewing the underlying distributional shift as performing interpolation between domains. We devise an algorithm to learn a representation that is robustly invariant under such interpolation and term it as interpolation robustness. We investigate the failure aspect of DG algorithms when availability of training data is scarce. Through extensive experiments, we show that our approach significantly outperforms the recent state-of-the-art algorithm.

35. Gupta, S. K. (2022). Mining of Frequent and High Utility Fuzzy Sequential Patterns.

Conference name: International Conference on Applied Artificial Intelligence and Computing (ICAAC)

Conference ranking: N/A

Abstract: Sequential pattern mining is one of the very popular data mining techniques to find patterns from databases. The variation in technique is based on the threshold used for mining which can be frequency, utility, or both. The use of fuzzy theory helps to provide natural representation by expressing parameters in linguistic terms. This paper proposes a hybrid fuzzy sequential pattern mining algorithm that expresses quantity in linguistic terms. The proposed method generates less number of patterns than the existing algorithms, according to the results of the experiments.

36. Senadeera, M. (2022, May). Empathetic Reinforcement Learning Agents.

Conference name: the 21st international conference on autonomous agents and multiagent systems

Conference ranking: A*

Abstract: With the increased interaction between artificial agents and humans, the need to have agents who can respond to their human counterparts appropriately will be crucial for the deployment of trustworthy systems. A key behaviour to permit this, one which humans and other living beings exhibit naturally, is empathy. In my research I explore the potential for agents to behave in ways that may be considered empathetic. Empathy is a two stage process involving the identification of the feelings or goals of the other, and having that same feeling be evoked in oneself. I began my work towards this objective by initially designing an agent who exhibits sympathy - the ability to identify the goals of another. Empathy is slightly more complex as it involves a process of projecting the state of the other back onto oneself and observing one's own response. In my research I hope to draw inspiration from this and evoke empathy through a process of mapping the other's goals back to oneself. By drawing upon empathetic responses, the hope is that this will lead to a faster and deeper understanding of the other.

37. Dang, L. H., Le, T. M., Le, V., Phuong, T. M., & Tran, T. (2022, November). Time-Evolving Conditional Character-centric Graphs for Movie Understanding.

Conference name: NeurIPS 2022 Temporal Graph Learning Workshop.

Conference ranking: N/A

Abstract: Temporal graph structure learning for long-term human-centric video understanding is promising but remains challenging due to the scarcity of dense graph annotations for long videos. It is the desired capability to learn the dynamic spatio-temporal interactions of human actors and other objects implicitly from visual information itself. Toward this goal, we present a novel Time-Evolving Conditional cHaracter-centric graph (TECH) for long-term human-centric video understanding with application in Movie QA. TECH is inherently a recurrent system of the query-conditioned dynamic graph that evolves over time along the story and follows throughout the course of a movie clip. As aiming toward human-centric video understanding, TECH uses a two-stage feature refinement process to draw attention to human characters and their interactions while treating the interactions with non-human objects as contextual information. Tested on the large-scale TVQA dataset, TECH clearly shows advantages over recent state-of-the-art models.

38. Nguyen, D., Nguyen, P., & Tran, T. (2021, February). EvSys: A relational dynamic system for sparse irregular clinical events.

Conference name: NeurIPS 2022 Temporal Graph Learning Workshop.

Conference ranking: N/A

Abstract: Clinical events such as clinic visits, hospital admissions, ECG readings and lab tests are recorded in modern healthcare systems. While these offer a great wealth of information about the state of health for a patient, modeling is challenging because (a) the events are sparse and irregular; (b) reading types vary greatly between episodes; and (c) the readings do not directly tell about the underlying continuous biological and mental processes that give rise to these readings. To tackle these challenges, we propose EvSys, a deep recurrent system that disentangles the observed measurement processes from latent health processes. With this design, the model is native to arbitrarily sparse and irregular clinical measurement events, and it captures the interacting underlying health processes that are not directly observed. We validate EvSys on two public datasets, namely PhysioNet 2012 and MIMIC-III, and demonstrate that the proposed model performs favorably against state-of-the-art methods.

Journal articles

1. Schonberg, W., & Ryan, S. (2022). Predicting metallic armour performance when impacted by fragment-simulating projectiles–Model adjustments and improvements.

Journal name: International Journal of Impact Engineering

Journal ranking: Q1

Abstract: In a previous study, a set of nine existing penetration models for blunt projectiles were identified and reviewed for their ability to predict the plugging-mode ballistic limit velocity of monolithic titanium alloy, aluminium alloy, and steel plates. Assessed with a database of more than 650 experimental ballistic limit measurements, often at conditions beyond that for which they had been originally derived or validated, it was found that all nine models could predict the ballistic limit to within +/- 50% for more than 50% of the database entries. In this paper, simple modifications to these nine models are proposed that introduce empirical adjustments, reformulations of the target strength dependency, or a combination of both with the pragmatic goal of identifying a model suitable for application across the wide range of monolithic metallic targets impacted by different calibre fragment-simulating projectiles across a full range of ordnance and sub-ordnance velocities in our database. The modifications were able to improve the performance of all models across this range, with the best performing models able to predict the experimental ballistic limit to within +/- 10% for 50% of the database entries and to within +/- 20% for 76% of the database entries.

-
2. Philip, B.J., Abdelrazek, M., Bonti, A., Barnett, S., & Grundy, J. (2022). Data collection mechanisms in health and wellness apps: review and analysis.

Journal name: JMIR mHealth and uHealth

Journal ranking: Q1

Abstract: Background

There has been a steady rise in the availability of health wearables and built-in smartphone sensors that can be used to collect health data reliably and conveniently from end users. Given the feature overlaps and user tendency to use several apps, these are important factors impacting user experience. However, there is limited work on analyzing the data collection aspect of mobile health (mHealth) apps.

Objective

This study aims to analyze what data mHealth apps across different categories usually collect from end users and how these data are collected. This information is important to guide the development of a common data model from current widely adopted apps. This will also inform what built-in sensors and wearables, a comprehensive mHealth platform should support.

Methods

In our empirical investigation of mHealth apps, we identified app categories listed in a curated mHealth app library, which was then used to explore the Google Play Store for health and medical apps that were then filtered using our selection criteria. We downloaded these apps from a mirror site hosting Android apps and analyzed them using a script that we developed around the popular AndroGuard tool. We analyzed the use of Bluetooth peripherals and built-in sensors to understand how a given app collects health data.

Results

We retrieved 3251 apps meeting our criteria, and our analysis showed that 10.74% (349/3251) of these apps requested Bluetooth access. We found that 50.9% (259/509) of the Bluetooth service universally unique identifiers to be known in these apps, with the remainder being vendor specific. The most common health-related Bluetooth Low Energy services using known universally unique identifiers were Heart Rate, Glucose, and Body Composition. App permissions showed the most used device module or sensor to be the camera (669/3251, 20.57%), closely followed by location (598/3251, 18.39%), with the highest occurrence in the staying healthy app category.

Conclusions

We found that not many health apps used built-in sensors or peripherals for collecting health data. The small number of the apps using Bluetooth, with an even smaller number of apps using standard Bluetooth Low Energy services, indicates a wider use of proprietary algorithms and custom services, which restrict the device use. The use of standard profiles could open this ecosystem further and could provide end users more options for apps. The relatively small proportion of apps using built-in sensors along with a high reliance on manual data entry suggests the need for more research into using sensors for data collection in health and fitness apps, which may be more desirable and improve end user experience.

3. Redley, B., Douglas, T., Hoon, L., White, K., & Hutchinson, A. (2022). Nursing guidelines for comprehensive harm prevention strategies for adult patients in acute hospitals: An integrative review and synthesis.

Journal name: International Journal of Nursing Studies

Journal ranking: Q1

Abstract: Background

Harm prevention research has seldom considered the complex demands on nurses negotiating multiple interrelated factors that contribute to preventable harms common in hospitalized patients. Best practice guidelines are available for individual risk factors, but few consider multiple factors that contribute to risk. As a consequence, duplication, contradiction, gaps, and volume of information limit harm prevention guideline use by nurses in daily practice.

Objectives

To systematically synthesise best-evidence recommendations from clinical practice guidelines to support nurses to deliver comprehensive harm prevention during acute hospitalization.

Design

An integrative review process was used to systematically identify, examine, evaluate and synthesise clinical nursing guidelines to prevent harm to hospitalized patients.

Methods

The search strategy developed with an expert librarian used a combination of targeted searching for guidelines published on websites, and forward and backward citation searching. Guidelines included were those most recently published, relevant to the international nursing context, and addressing one or more of eight factors contributing to preventable harms. The AGREE-REX (Appraisal of Guidelines Research and Evaluation–Recommendations Excellence) tool was used for critical appraisal of guidelines regarding appropriateness to target users (i.e., nurses), trustworthiness, and implementable in acute hospitals. EndNote and NVIVO 12 were used to manage the high volume of extracted data and facilitate analysis. Analyses involved using the framework method to code data for relevance to an eight-factor harm prevention framework; steps for inductive thematic analyses were used for synthesis. Iterations of the thematic model were refined by sharing with hospital patient safety experts, who endorsed the final model.

Results

154 guidelines met inclusion criteria, providing 7,429 recommendations. Synthesis involved mapping of recommendations across the eight-factor framework that informed a hierarchy of risk for harm prevention activity. Six themes represented nursing care strategies across the eight-factors that could be integrated into local practice contexts. The themes are framed into a model for nurse comprehensive harm prevention.

Conclusions

The complexity and volume of guidance for comprehensive harm prevention necessitates contemporaneous, integrated, and accessible guidance to support nurses' decision-making in their daily care provision. This research provides an integrated model to assist nurses to identify patients most vulnerable to multiple preventable harms during hospitalization and guide a comprehensive harm prevention strategy to keep them safe in hospital.

4. Patil, P. W., Gupta, S., Rana, S., & Venkatesh, S. (2022). Dual-frame spatio-temporal feature modulation for video enhancement.

Journal name: Pattern Recognition

Journal ranking: Q1

Abstract: Current video enhancement approaches have achieved good performance in specific rainy, hazy, foggy, and snowy weather conditions. However, they currently suffer from two important limitations. First, they can only handle degradation caused by single weather. Second, they use large, complex models with 10–50 millions of parameters needing high computing resources. As video enhancement is a pre-processing step for applications like video surveillance, traffic monitoring, autonomous driving, etc., it is necessary to have a lightweight enhancement module. Therefore, we propose a dual-frame spatio-temporal feature modulation architecture to handle the degradation caused by diverse weather conditions. The proposed architecture combines the concept of spatio-temporal multi-resolution feature modulation with a multi-receptive parallel encoders and domain-based feature filtering modules to learn domain-specific features. Further, the architecture provides temporal consistency with recurrent feature merging, achieved by providing feedback of the previous frame output. The indoor (REVIDE, NYUDepth), synthetically generated outdoor weather degraded video de-hazing, and de-raining with veiling effect databases are used for experimentation. Also, the performance of the proposed method is analyzed for night-time de-hazing and de-raining with veiling effect weather conditions. Experimental results show the superior performance of our framework compared to existing state-of-the-art methods used for video de-hazing (indoor/ outdoor) and de-raining with veiling effect weather conditions. The code is available at <https://github.com/pwp1208/PR2022>

5. Kuttichira, D. P., Gupta, S., Nguyen, D., Rana, S., & Venkatesh, S. (2022). Verification of integrity of deployed deep learning models using Bayesian Optimization

Journal name: Knowledge-Based Systems

Journal ranking: Q1

Abstract: Artificial Intelligence (AI) is hugely driven by machine learning (ML). Popular ML algorithms like deep neural networks (DNN) are used in wide variety of applications. Due to excellent predictive capabilities, these models are used in many real-world applications. To accommodate the huge service demands from these models, they are hosted on cloud platforms. Hosting a model in cloud raises security concerns. Malicious third party can modify the model when in transit from local system to cloud platform or after being uploaded in the cloud. We propose Bayesian Compromise Detection (BCD) algorithm to detect such model compromises by generating sensitive samples. Finding a sensitive sample is essentially an optimization problem that aims to maximize the prediction differences between the original and the compromised model. The optimization problem is challenging as (1) a cloud customer can only have a blackbox access to the compromised model (2) sensitive sample has to be searched in the distribution of the training inputs, which is most likely a high-dimensional space and (3) the optimization problem is non-convex. To tackle these obstacles, we use Variational Autoencoder (VAE) to map high-dimensional data to a non-linear low-dimensional space and then use Bayesian Optimization (BO) to find the globally optimum sensitive sample. Proposed algorithm generates a sensitive sample that can detect model compromise without incurring much cost by multiple queries. Experiments with multiple datasets demonstrate the effectiveness of our proposed method. Our method outperforms the state-of-the-art method in terms of detection rate of the compromised models on cloud platforms.

6. Harikumar, H., Rana, S., Gupta, S., Nguyen, T., Kaimal, R., & Venkatesh, S. (2022). Prescriptive analytics with differential privacy.

Journal name: International Journal of Data Science and Analytics

Journal ranking: Q2

Abstract: Prescriptive analytics is a mechanism that provides the best set of actions to be taken to prevent undesirable outcomes for a given instance. However, this mechanism is prone to privacy breaches if an adversary with subsidiary data is allowed multiple query access to it. So, we propose a differential privacy mechanism in prescriptive analytics to preserve data privacy. Differential privacy can be achieved with the help of sensitivity of the given actions. Roughly speaking, sensitivity is the maximum change in the given set of actions with respect to the change in the given instances. However, a general analytical form for the sensitivity of the prescriptive analytics mechanism is difficult to derive. So, we formulate a nested constrained optimization to solve the problem. We use synthetic data in the experiments to validate the behavior of the differential privacy mechanism with respect to different privacy parameter settings. The experiments with two real-world datasets—Student Academic Performance and Reddit dataset, demonstrate the usefulness of our proposed method in education and social policy design. We also propose a new evaluation measure called the prescription success rate to further investigate the significance of our proposed method.

7. Nguyen, T. M., Nguyen, T., & Tran, T. (2022). Mitigating cold-start problems in drug-target affinity prediction with interaction knowledge transferring.

Journal name: Briefings in Bioinformatics

Journal ranking: Q1

Abstract: Predicting the drug-target interaction is crucial for drug discovery as well as drug repurposing. Machine learning is commonly used in drug-target affinity (DTA) problem. However, the machine learning model faces the cold-start problem where the model performance drops when predicting the interaction of a novel drug or target. Previous works try to solve the cold start problem by learning the drug or target representation using unsupervised learning. While the drug or target representation can be learned in an unsupervised manner, it still lacks the interaction information, which is critical in drug-target interaction. To incorporate the interaction information into the drug and protein interaction, we proposed using transfer learning from chemical-chemical interaction (CCI) and protein-protein interaction (PPI) task to drug-target interaction task. The representation learned by CCI and PPI tasks can be transferred smoothly to the DTA task due to the similar nature of the tasks. The result on the DTA datasets shows that our proposed method has advantages compared to other pre-training methods in the DTA task.

8. Nguyen, D. H., Nghiem, N. V. D., Nguyen, B. S., Le, D. T. T., Sabahi, S., Nguyen, M. T., & Le, H. (2022, July). Make The Most of Prior Data: A Solution for Interactive Text Summarization with Preference Feedback

Journal name: Computational Linguistics and Natural Language Processing-Findings

Journal ranking: Q1

Abstract: For summarization, human preference is critical to tame outputs of the summarizer in favor of human interests, as ground-truth summaries are scarce and ambiguous. Practical settings require dynamic exchanges between human and AI agent wherein feedback is provided in an online manner, a few at a time. In this paper, we introduce a new framework to train summarization models with preference feedback interactively. By properly leveraging offline data and a novel reward model, we improve the performance regarding ROUGE scores and sample-efficiency. Our experiments on three various datasets confirm the benefit of the proposed framework in active, few-shot and online settings of preference learning.

9. Nguyen, D., Nguyen, D. T., Zeng, R., Nguyen, T. T., Tran, S. N., Nguyen, T., ... & Fookes, C. (2021). Deep auto-encoders with sequential learning for multimodal dimensional emotion recognition.

Journal name: IEEE Transactions on Multimedia

Journal ranking: Q1

Abstract: Multimodal dimensional emotion recognition has drawn a great attention from the affective computing community and numerous schemes have been extensively investigated, making a significant progress in this area. However, several questions still remain unanswered for most of existing approaches including: (i) how to simultaneously learn compact yet representative features from multimodal data, (ii) how to effectively capture complementary features from multimodal streams, and (iii) how to perform all the tasks in an end-to-end manner. To address these challenges, in this paper, we propose a novel deep neural network architecture consisting of a two-stream auto-encoder and a long short term memory for effectively integrating visual and audio signal streams for emotion recognition. To validate the robustness of our proposed architecture, we carry out extensive experiments on the multimodal emotion in the wild dataset: RECOLA. Experimental results show that the proposed method achieves state-of-the-art recognition performance.

10. Nguyen, T. M., Quinn, T. P., Nguyen, T., & Tran, T. (2022). Explaining Black Box Drug Target Prediction Through Model Agnostic Counterfactual Samples.

Journal name: IEEE/ACM Transactions on Computational Biology and Bioinformatics

Journal ranking: Q2

Abstract: Many high-performance DTA deep learning models have been proposed, but they are mostly black-box and thus lack human interpretability. Explainable AI (XAI) can make DTA models more trustworthy, and allows to distill biological knowledge from the models. Counterfactual explanation is one popular approach to explaining the behaviour of a deep neural network, which works by systematically answering the question "How would the model output change if the inputs were changed in this way?". We propose a multi-agent reinforcement learning framework, Multi-Agent Counterfactual Drug-target binding Affinity (MACDA), to generate counterfactual explanations for the drug-protein complex. Our proposed framework provides human-interpretable counterfactual instances while optimizing both the input drug and target for counterfactual generation at the same time. We benchmark the proposed MACDA framework using the Davis and PDBBind dataset and find that our framework produces more parsimonious explanations with no loss in explanation validity, as measured by encoding similarity. We then present a case study involving ABL1 and Nilotinib to demonstrate how MACDA can explain the behaviour of a DTA model in the underlying substructure interaction between inputs in its prediction, revealing mechanisms that align with prior domain knowledge.

11. Ryan, S., Berk, J., Rana, S., McDonald, B., & Venkatesh, S. (2022). A bayesian optimisation methodology for the inverse derivation of viscoplasticity model constants in high strain-rate simulations.

Journal name: Defence Technology

Journal ranking: Q1

Abstract: We present an inverse methodology for deriving viscoplasticity constitutive model parameters for use in explicit finite element simulations of dynamic processes using functional experiments, i.e., those which provide value beyond that of constitutive model development. The developed methodology utilises Bayesian optimisation to minimise the error between experimental measurements and numerical simulations performed in LS-DYNA. We demonstrate the optimisation methodology using high hardness armour steels across three types of experiments that induce a wide range of loading conditions: ballistic penetration, rod-on-anvil, and near-field blast deformation. By utilising such a broad range of conditions for the optimisation, the resulting constitutive model parameters are generalised, i.e., applicable across the range of loading conditions encompassed by those experiments (e.g., stress states, plastic strain magnitudes, strain rates, etc.). Model constants identified using this methodology are demonstrated to provide a generalisable model with superior predictive accuracy than those derived from conventional mechanical characterisation experiments or optimised from a single experimental condition.

12. Nguyen, M. T., Nguyen, T., & Tran, T. (2022). Learning to discover medicines.

Journal name: International Journal of Data Science and Analytics

Journal ranking: Q2

Abstract: Discovering new medicines is the hallmark of the human endeavor to live a better and longer life. Yet the pace of discovery has slowed down as we need to venture into more wildly unexplored biomedical space to find one that matches today's high standard. Modern AI-enabled by powerful computing, large biomedical databases, and breakthroughs in deep learning offers a new hope to break this loop as AI is rapidly maturing, ready to make a huge impact in the area. In this paper, we review recent advances in AI methodologies that aim to crack this challenge. We organize the vast and rapidly growing literature on AI for drug discovery into three relatively stable sub-areas: (a) representation learning over molecular sequences and geometric graphs; (b) data-driven reasoning where we predict molecular properties and their binding, optimize existing compounds, generate de novo molecules, and plan the synthesis of target molecules; and (c) knowledge-based reasoning where we discuss the construction and reasoning over biomedical knowledge graphs. We will also identify open challenges and chart possible research directions for the years to come.

13. Tawfik, S. A., Nguyen, P., Tran, T., Walsh, T. R., & Venkatesh, S. (2022). Machine learning-aided exploration of ultrahard materials.

Journal name: The Journal of Physical Chemistry C

Journal ranking: Q1

Abstract: Ultrahard materials are an essential component in a wide range of industrial applications. In this work, we introduce novel machine learning (ML) features for the prediction of the elastic moduli of materials, from which the Vickers hardness can be calculated. By applying the trained ML models on a space of ~110,000 materials, these features successfully predict the elastic moduli for a range of materials. This enables the identification of materials with high Vickers hardness, as validated by comparing the predictions against the density functional theory calculations of the moduli. We further explored the predicted moduli by examining several classes of materials with interesting mechanical properties, including binary and ternary alloys, aluminum and magnesium alloys, metal borides, carbides and nitrides, and metal hydrides. Based on our ML models, we identify a number of ultrahard compounds in the B-C and B-C-N chemical spaces and ultrahard ultralight-weight magnesium alloys Mg₃Zn and Mg₃Cd. We also observe the inverse of the hydrogen embrittlement effect in a number of metal carbides, where the introduction of hydrogen into metal carbides increases their hardness, and find that substitutional doping of Al in transition-metal borides can yield lighter materials without compromising the thermodynamic stability or the hardness of the material.

14. Quinn, T. P., Jacobs, S., Senadeera, M., Le, V., & Coghlan, S. (2022). The three ghosts of medical AI: Can the black-box present deliver?

Journal name: Artificial intelligence in medicine

Journal ranking: Q1

Abstract: Our title alludes to the three Christmas ghosts encountered by Ebenezer Scrooge in A Christmas Carol, who guide Ebenezer through the past, present, and future of Christmas holiday events. Similarly, our article takes readers through a journey of the past, present, and future of medical AI. In doing so, we focus on the crux of modern machine learning: the reliance on powerful but intrinsically opaque models. When applied to the healthcare domain, these models fail to meet the needs for transparency that their clinician and patient end-users require. We review the implications of this failure, and argue that opaque models (1) lack quality assurance, (2) fail to elicit trust, and (3) restrict physician-patient dialogue. We then discuss how upholding transparency in all aspects of model design and model validation can help ensure the reliability and success of medical AI.

15. Deane, K., Yang, Y., Licavoli, J.J., Nguyen, V., Rana, S., Gupta, S., ... & Sanders, P. G. (2022). Utilization of Bayesian Optimization and KWN Modeling for Increased Efficiency of Al-Sc Precipitation Strengthening.

Journal name: Metals, Multidisciplinary Digital Publishing Institute

Journal ranking: Q1

Abstract: The Kampmann and Wagner numerical model was adapted in MATLAB to predict the precipitation and growth of Al3Sc precipitates as a function of starting concentration and heat-treatment steps. This model was then expanded to predict the strengthening in alloys using calculated average precipitate number density, radius, etc. The calibration of this model was achieved with Bayesian optimization, and the model was verified against experimentally gathered hardness data. An analysis of the outputs from this code allowed the development of optimal heat treatments, which were validated experimentally and proven to result in higher final strengths than were previously observed. Bayesian optimization was also used to predict the optimal heat-treatment temperatures in the case of limited heat-treatment times.

16. Kulkarni, A., Patil, P. W., Murala, S., & Gupta, S. (2022). Unified Multi-Weather Visibility Restoration.

Journal name: IEEE Transactions on Multimedia.

Journal ranking: Q1

Abstract: Automated surveillance is widely opted for applications such as traffic monitoring, vehicle identification, etc. But, various weather degradation factors such as rain and snow streaks, along with atmospheric veil severely affect the perceptual quality of an image, eventually affecting the performance of these applications. There exist weather specific (rain, haze, snow, etc.) methods focusing on respective restoration task. As image restoration is a preprocessing step for high level surveillance applications, it is practically inapplicable to have different architectures for different weather restoration. In this paper, we propose a lightweight unified network, having 1.1M parameters (1/40 and 1/6 of the existing state-of-the-art rain with veil removal, and snow with veil removal methods respectively) for removal of rain and snow along with the veiling effect present in the images. In this network, we propose two parallel streams to handle the degradations and restoration: First, degradation removal stream (DRS) focuses mainly on removing randomly repeating degradations i.e., rain and snow streaks, through the proposed adaptive multi-scale feature sharing block (AMFSB) and stage-wise subtractive block (SSB). Second, feature corrector stream (FCS) mainly focuses on refining the partial outputs of the first stream, reducing the veiling effect and acts supplementary to the first stream. Finally, we leverage contrastive regularization for better convergence of the proposed network. Substantial experiments on synthetic as well as real-world images, along with extensive ablation studies, demonstrate that the proposed method performs competitively with the existing state-of-the-art methods for multi-weather image restoration. The code is available at <https://github.com/AshutoshKulkarni4998/UVRNet>.

17. Nguyen-Tang, T., Gupta, S., & Venkatesh, S. (2022). On Sample Complexity of Offline Reinforcement Learning with Deep ReLU Networks in Besov Spaces.

Journal name: Transactions on Machine Learning Research

Journal ranking: Q1

Abstract: Offline reinforcement learning (RL) leverages previously collected data for policy optimization without any further active exploration. Despite the recent interest in this problem, its theoretical results in neural network function approximation settings remain elusive. In this paper, we study the statistical theory of offline RL with deep ReLU network function approximation. Notably, our sample complexity holds under two novel considerations: the Besov dynamic closure and the correlated structure. While the Besov dynamic closure subsumes the dynamic conditions for offline RL in the prior works, the correlated structure renders the prior works of offline RL with general/neural network function approximation improper or inefficient in long (effective) horizon problems. To the best of our knowledge, this is the first theoretical characterization of the sample complexity of offline RL with deep neural network function approximation under the general Besov regularity condition that goes beyond the linearity regime in the traditional Reproducing Hilbert kernel spaces and Neural Tangent Kernels.



Email a2i2@deakin.edu.au

Deakin University
Waurin Ponds Campus
NA Building
75 Pigdons Road
Waurin Ponds Victoria 3216

Deakin University
Melbourne Burwood Campus
BC Building, Level 7
221 Burwood Highway
Burwood Victoria 3125

a2i2.deakin.edu.au