AI for Missions PhD Project Proposals

May 2022

Contents

AI4M-PHD-01_3D Hyperspectral analysis of meat to identify contamination

Description

Meat and poultry products are a major part of the agriculture and food industry with the gross value of red meat estimated by the Australian Bureau of Statistics to be $5.2 billion and $770 million for poultry in the last quarter of 2021. However, these products are prone to contamination with pathogens that may go undetected, leading to an inherent lowering of the meat expiration date. Meat contamination leads to billions of dollars of annual loss and significantly contributes to the 600 million cases of foodborne illness and 420 thousand fatalities worldwide. There are several pathogens associated with food-borne illness from contaminated meat including *Salmonella spp.*, *E. coli*, *Campylobacter jejuni* and *Yersinia*. Hyperspectral imaging (HSI) isa powerful technology for food quality assessment and safety since it is a quick, precise, and non-destructive tool for detecting contamination. Furthermore, if accompanied by machine learning strategies HSI can provide an efficient meat surface monitoring solution to avoid human mistakes, and boost public confidence in the meat processing system while maintaining fit-for-purpose production speed. Our proposed PhD project tackles two scientific challenges related to the use of HIS for meat contamination assessment.

Firstly, we want to focus on an explainable vision task that investigates the spatial-spectral properties and 3D structure of meat and its possible contaminants. We will use primal beef and lamb cuts as a proxy for whole carcass measurements as well as broiler chicken carcasses. Typically, convolution (or convolution networks) can be used to get spatial features from hyperspectral images and a variety of approaches ranging from distance measures (like spectral angular measure) to non-negative matrix factorization can be used to exploit spectral responses. Although extracting spectral information and combining it with spatial features is an existing area of research, current limitations revolve around building computer vision systems flexible enough to be applied in real-world scenarios. Here novel machine learning models will be investigated that can extract spectral details and effectively combine them with structural information to explore unique properties present in contamination areas within a meat matrix. We will also explore the contamination spread in a multi-view environment by employing 3D hyperspectral information with weakly supervised machine learning.

Secondly, we will focus on the issues arising from the variable and unbalanced nature of real-world data, which typically follows a long-tailed distribution where few classes are frequently observed while many are only rarely encountered. This creates an inherent challenge when training computer vision models. Additionally, contamination spots can be very small and must be recognised over the highly variable background of the meat matrices when scanned. To tackle this issue, we will investigate transfer learning of spectra in hyperspectral imaging.

Finally, as there is no current recognised data set, an optimal data collection strategy needs to be developed. This strategy will look into issues such as calibration of the camera in a relevant setting, development of classes and hierarchies for storing data, and data labelling.

Relevant publications

- Ali Zia, Jie Liang, Jun Zhou, and Yongsheng Gao. "3D reconstruction from hyperspectral images", In Proceedings of the IEEE Winter Conference on Applications of Computer Vision (WACV), pages 318-325, Waikoloa Beach, Hawaii, 2015.

  Context: This paper explores how different spectra can contribute to the structural properties of an object. This approach helps to understand and explore spectral-spatial properties present in the 3D structure of the meat.

- Ali Zia, Jun Zhou, and Yongsheng Gao. "Exploring chromatic aberration and defocus blur for relative depth estimation from monocular hyperspectral Image", IEEE Transactions on Image Processing (IEEE TIP),30:4357-4370, 2021.

  Context: This paper is useful because it explores depth from a monocular hyperspectral image which in turn can be useful in constructing a 3D model. The paper also explores the manifold learning approach, which could be helpful in contamination detection.

- Ernest Bonah, Xingyi Huang, Joshua Harrington Aheto, Ren Yi, Shanshan Yu, Hongyang Tu. "Comparison of variable selection algorithms on vis-NIR hyperspectral imaging spectra for quantitative monitoring and visualization of bacterial foodborne pathogens in

  Context: This research provides useful cues to build optimal hyperspectral imaging set-ups and experiments for contaminated meat. It also provides some basic optimisation strategies which could be helpful in designing a machine learning framework for our work.

Skillset required

- Knowledge of hyperspectral and multispectral imaging
- Strong linear algebra and computer vision understanding
- Basic understanding of 3D computer vision
- Basic machine learning, deep learning and model visualisation skills
- Good programming skills in Python, Matlab etc
- Understanding of how to work in a Linux environment
- Familiar with machine learning libraries like PyTorch, tensor flow etc.

Contact Details

Dr Vivien (Viv) Rolland, CSIRO, vivien.rolland@csiro.au

Associate Professor Jun Zhou, Griffith University, jun.zhou@griffith.edu.au

AI4M-PHD-02 AI-Driven Digital Twins for Additive Manufacturing

Description

Additive Manufacturing has emerged in recent years as a critical tool, both for the prototyping of ideas during the R&D process and for production of high value items in areas including Aerospace, Biomedical devices and Advanced Manufacturing. The process however remains very time consuming and expensive, and manufacturing defects in completed parts may require them to be discarded and re-printed.

The recent emergence of Digital Twin methodologies, where a virtual representation of a physical asset / process is used to better understand and dynamically optimize and control a process in real-time, has potential for transformational impact in improving additive manufacturing processes. This project will aim to create an Intelligent Digital Twin [1] combining physics-based modelling, real-time sensor measurements and state-of-the-art Bayesian Machine Learning techniques. At the core of such an intelligent Digital Twin [1], are Machine Learning algorithms that allow the system to learn from inputted sets of data to predict outcomes and dynamically adjust when presented with new scenarios or inputs. While the training of these models can be computationally expensive and the acquisition of sufficient amounts of data at times poses a significant challenge, well trained models are generally able to execute and make predictions with high speed and low computational cost. Using an integrated approach that combines data from both experiments and physics based computational models, we will develop an automated system incorporating a Bayesian machine-learning model that will via active sampling determine the best simulation/real world data to obtain and automatically generate and analyse new simulation cases to optimally improve predictive capability. The development of these optimal active sampling techniques and methods for incorporation of both simulated and real world data will be key focuses of the project. In this way we will efficiently develop surrogate models that can be rapidly applied to problems of interest and then used as part of an intelligent Digital Twin control system to dynamically learn and optimize complex processes in real-time.

The main area of application we will consider will be Additive Manufacturing processes (3D printing), with the Intelligent Digital Twin methodologies and Hybrid Bayesian approaches developed also being applicable to broader design, control and optimisation problems.

Relevant publications

1. This recently published paper proposes a hierarchy of Digital Twins for Metal Additive Manufacturing. The hierarchy ranges from simple implicit models used for visualisation and modelling simple variations to the system, up to the type of Intelligent Digital Twin we will seek to create in this project that will be capable of real-time intelligent process control. A. Phua, C.H.J. Davies and G.W. Delaney, "A Digital Twin Hierarchy for Metal Additive Manufacturing", Computers in Industry, 140:103667,2022. https://doi.org/10.1016/j.compind.2022.103667

2. This paper demonstrates the use of sequential Monte Carlo to perform Bayesian optimization on physics-based models, including parametric models of gravity signals and hydraulic simulations of contamination in an aquifer. The models demonstrated here are relatively simple and inference is performed offline, where this PhD will focus on real-time optimisation and possibly on higher-dimensional models. Rafael Oliveira, Richard Scalzo, Robert Kohn, Sally Cripps, Kyle Hardman, John Close, Nasrin Taghavi, Charles Lemckert. "Bayesian optimization with informative parametric models via sequential Monte Carlo.", Data-Centric Engineering 3 (March 2022) https://doi.org/10.1017/dce.2022.5

Skillset required

Degree in Computer Science / Physical Sciences / Mathematical Sciences, ideally with a working knowledge of ML principles.

Contact Details

Gary Delaney (gary.delaney@data61.csiro.au)

Chris Davies, Monash University (Chris.Davies@monash.edu)

AI4M-PHD-03_AI-accelerated breeding of high-protein mungbean

Description

There is a pressing need to produce more protein more sustainably to meet the dual challenges of human population growth and environmental degradation. Plant proteins, especially from nitrogen-fixing legumes are seen as a great solution to these existential problems. In the next 10 years, global protein demand is expected to increase by 20%. Within Australia, there is an emerging pulse protein industry that may grow as much as 6000% by 2030, requiring over 100,000 tons of legume protein (CSIRO Futures, 2022, Protein Roadmap). To meet these challenges, we need to accelerate genetic gains in pulse breeding programs by integrating AI. Here, we will target mungbean, a substantial source of protein grown in QLD and NSW. Availability of a mungbean reference genome enables genomics-based breeding methods such as use of molecular markers for genomic selection, as well as proteomics approaches that can guide breeding to improve seed protein compositional traits. By expanding and leveraging genomic and proteomic resources, there is an opportunity to build comprehensive protein profiles for diverse mungbean genotypes to accelerate mungbean improvement. Currently, there is no proteomics database available for mungbean. In protein crop breeding, a major challenge is how to quickly stack many desirable alleles for disease resistance and stress adaptive traits into new, high yielding varieties with good protein quality. For example, at least 10 genes have been identified that are associated with disease resistance, and 20 genes associated with physiological and component traits contributing to yield. Given that the majority of these were discovered in different and diverse germplasm, there are approximately 2^30=1.07x10^9 ways the lines could be crossed. This practical problem is exacerbated as environmental, as well as other factors contributing to stress, have profound impact on the plant proteome making it particularly challenging to develop high-protein mungbean suitable for production in different growing regions.

The highly combinatorial problem of finding the best crossing combination is well suited to AI approaches. This project will develop a pan-genome and pan-proteome of the founder lines of a nested association mapping population (n = 31; including wild germplasm, exotic and elite breeding lines), and use digital twin and AI-guided approaches to determine optimum strategies that will efficiently stack alleles in the shortest time. This will be accomplished as follows: (1) Develop a pangenome of the founder lines using whole genome sequencing; (2) Develop a pan-proteome of the founder lines using mass spectrometry and associated bioinformatics; (3) Discover the chromosome segments associated with protein and other traits of interest and conduct bioinformatic analysis to explore associated biological pathways; (4) Use our digital twin and AI algorithms (evolutionary computing, CNN, Random forest) to determine the optimal strategy for crossing.

Outcomes of this project include pan-genome and pan-proteome resources as well as an AI-guided strategy to quickly develop elite high-protein mungbean lines.

Relevant publications

- Michelle L. Colgrave, Sonja Dominik, Aarti B. Tobin, Regine Stockmann, Cedric Simon, Crispin A. Howitt, Damien P. Belobrajdic, Cate Paull, and Thomas Vanhercke. Perspectives on Future Protein Production. Journal of Agricultural and Food Chemistry 2021 69 (50), 15076-15083 DOI: 10.1021/acs.jafc.1c05989
- Voss-Fels KP, Stahl A, Wittkop B, Lichthardt C, Nagler S, Rose T, Chen TW, Zetzsche H, Seddig S, Majid Baig M, Ballvora A, Frisch M, Ross E, Hayes BJ, Hayden MJ, Ordon F, Leon J, Kage H, Friedt W, Stützel H, Snowdon RJ. Breeding improves wheat productivity under contrasting agrochemical input levels. Nat Plants. 2019 Jul;5(7):706-714. doi: 10.1038/s41477-019-0445-5.
- Hayes BJ, Shepherd RK, Newman S (1998) Proc. Advanc. Animal. Breed Genet. 12:108.

Skillset required

Expertise in two or more areas: plant genomics, breeding and quantitative genetics, plant physiology, molecular biology and biochemistry, bioinformatics, and computational biology

Contact Details

Professor Michelle Colgrave, CSIRO, michelle.colgrave@csiro.au

Professor Ben Hayes, The University of Queensland, b.hayes@uq.edu.au

Dr Merrill Ryan, Department of Agriculture and Fisheries, Merrill.Ryan@daf.qld.gov.au

AI4M-PHD-04_AI/ML-based risk prediction for zoonotic spillover

Description

Infectious disease emergence from spillover events is of global concern, as highlighted by the COVID-19 global pandemic. More than 70% of emerging infectious diseases are transmitted from animals to humans, and their incidence has tripled over the past 50years. The increase of outbreaks has been driven by various factors, including human encroachment into wildlife habitats, increasing urbanisation and climate change. Predicting future spillover risks are increasingly crucial to human health and wellbeing, and economic stability. Following the One Health approach, it can be said that the health of humans, animals, and shared environments is closely connected. This project focuses on developing a computationally efficient AI/MLalgorithms to leverage large spatiotemporal environmental, human, and animal data sets. Additionally, we are creating our own data sets through the development and deployment of a novel next gen surveillance system, which will also feed into this project. The main objectives of this project are to identify •key reservoir hosts that increase the chance of zoonotic spillovers, •the pathogens that are most likely to emerge, •geographic origins of a potential future pandemic, and •use process-based approaches to supplement and help correct biases of the data-based AI/ML approaches. We expect these models and results to contribute to the early warning and decision support systems on our paths to impact.

Relevant publications

The costs and benefits of primary prevention of zoonotic pandemics (science.org)

Global hotspots and correlates of emerging zoonotic diseases

The Ecology of Nipah Virus in Bangladesh: A Nexus of Land-Use Change and Opportunistic Feeding Behavior in Bats

Skillset required

- Spatiotemporal data analysis
- AI/ML model development
- AI/ML model implementation (python)
- Hybrid AI/ML and process-based modelling development and implementation (python)

  i.e. Degree in Data Sciences, Mathematical Sciences, Spatial Sciences, Computer Sciences, or similar.

Contact Details

Dr Maryam Golchin (maryam.golchin@csiro.au)

A/Prof Roslyn Hickson James Cook University (roslyn.hickson@jcu.edu.au)

AI4M-PHD-05_Anomaly detection using self-supervised learning

Description

In many applications, the detection of rare events is of high value, and the automation of such detections using modern machine learning is a desirable goal and open challenge. This includes diverse applications such as the detection of rare events in diseases and conditions within healthcare and medical imaging settings, the detection of suspect behaviours from video networks or sensors, or the detection of defects in manufacturing approaches. In most of the cases, there is plenty of “normal” data, and a small number of examples capturing the event(s) to be detected; either because those events are highly variable, rare, expensive, or difficult to collect. For example, to detect a suspicious activity in a factory for digital twin applications, it would be almost impossible to record all the possible ways a fire could start and of course it would not be feasible to start an actual fire to provide adequate training data. However, a large amount of video and sensing recordings would exist during normal operations. Similarly, to detect a suspicious activity within a factory, it would be very difficult to record all the possible ways someone might act suspiciously, however, a huge amount of data can be collected of normal activity during regular operations.

The main hypothesis of this project is that a model could be built using only “normal” data which are abundant such that when an abnormal event or anomaly occurs, the model would flag an out of the normal distribution event. This screening could then be used to trigger a review by a human expert or to run a subsequent model to characterise the abnormality.

During this project, the student will investigate self-supervised methods to model distributions from large “normal” datasets, and will learn self-supervised deep representations across a variety of datasets. Importantly, these methods do not require manual labelling, but rather try to match the observed data using small dimensions of latent variables. This class of methods include autoencoders, variational autoencoders, and style conditioned generative adverserial networks. In addition to self-supervised methods, this project will also examine the development of complementary multi-task learning strategies and knowledge distillation approaches which can also support the development of more advanced and generalisable anomaly detection approaches. An additional benefit would be the intrinsic property of those models to provide explanation and transparency, allowing increase trust and easier human in the loop settings.

Relevant publications

- Overview of the various approach to anomaly detection: Anomaly Detection with Machine Learning: An Introduction – BMC Software | Blogs

- Anomaly detection of object using self-supervised learning for tasks detection: CVPR 2021 Open Access Repository (thecvf.com)

- Anomaly detection of brain infarcts in MRI: [2004.03271] Autoencoders for Unsupervised Anomaly Segmentation in Brain MR Images: A Comparative Study (arxiv.org)

Skillset required

The students will require excellent programming skills in Python and knowledge of Pytorch. Image processing/computer vision will be a plus. Background in numerical optimisation, linear algebra, and variational statistics will also be required.

Contact Details

Prof Olivier Salvado, CSIRO (olivier.salvado@csiro.au)

Prof Clinton Fookes, Queensland University of Technology (c.fookes@qut.edu.au)

AI4M-PHD-06_Artificial intelligence(AI)-aided identification of novel bioactive peptides for human health and crop protection

Description

There is an increasing consumer demand to produce evidence-backed healthy food ingredients and natural functional food products. To achieve this goal, exploring raw materials low in antinutritive content, fortified in micronutrients and enriched in proteins and peptides associated with health benefits is an essential step. Bioactive peptides (BPs) are vital components of defence mechanisms in living organisms. They exhibit a broad range of biological activities beneficial to human health and plant protection including antifungal, antibacterial, antiviral, insecticidal, anti-inflammatory or anticancer agents. The potential to harness these bioactive properties has been significantly increased by the exponential growth of available high-quality, well annotated genome sequence data of crop and horticultural species. The diversity of bioactive peptide sequences coupled with the time and costs associated with experimental work in identification and testing of bioactive candidates prevents comprehensive experimental screening of the entire peptide sequence space. Much work has been recently done in leveraging computational tools to identify new bioactive candidates for applications in human health and crop protection, primarily relying on sequence alignments or exploring evolutionary relationships based on sequence similarity. Machine-learning (ML) models are cost-effective and time-saving strategies that can be used to predict biological activities from primary sequences by identifying the essential physicochemical characteristics of bioactive functionality. Next generation sequencing data provide an unexplored resource to uncover peptides with potential bioactive properties in food crops. In this project, we will use a combined approach of computational biology, statistics and mass spectrometry-based proteomics to identify novel peptides with potential bioactive function from agriculturally important crop species. To achieve these goals, first, public bioactive peptide databases will be collected and used for the development of ML model. Models will be trained and validated using experimentally characterized bioactive peptide databases by relying on statistical learning to understand empirical relationships between physicochemical properties and biological activity. The developed model will then be applied in high throughput in silico screening using genomic data resources of food crops to identify novel candidates with the desired biological activity. Expression of predicted bioactive peptides will be then confirmed by mining established RNA seq data sets and at protein level by analysing cereal grains and pulses using discovery proteomics and liquid chromatography mass spectrometry. These foundational tools and knowledge can help in mining the tremendous genomic potential within plants for bioactive peptide diversity, including their unpredictable modifications, and deepen our understanding of their underpinning mechanisms of action, to support their future use in human health and crop protection applications.

Relevant publications

- Discovery and design of peptides with specific function using machine learning: https://www.nature.com/articles/s41598-020-73644-6
- Bioactive peptide discovery and characterisation within a functional food ingredient: https://www.sciencedirect.com/science/article/pii/S2665927121000216
- Discovery and design of peptides with specific bioactive function using machine learning: Machine learning-enabled discovery and design of membrane-active peptides - ScienceDirect

Skillset required

- Data mining
- coding/data science
- bioinformatics
- laboratory skills related to protein analysis

Contact Details

James A. Broadbent (james.broadbent@csiro.au)

Angéla Juhász, Edith Cowan University (a.juhasz@ecu.adu.au)

AI4M-PHD-07 Beyond the Black Box: Developing physically informed ML architectures for accelerating particle based models

Description

Recently much research interest has emerged in the use of Artificial Neural Networks (ANNs) to accelerate simulations of physical systems. Most initial studies in this area have considered the ANN as a black box for learning the underlying physics of the target system, and have not directly incorporated the known physical laws and natural symmetries of the system into the architecture of the network. This leads to networks that are large and slow to train, and that ultimately yield suboptimal performance. However, exploiting prior knowledge of the physics associated with the target function or process, in addition to tailored input-output training set pairs, can significantly improve ANN training and prediction performance. We call this hybrid training.

This project will explore the potential for significant ANN performance gains, including training speed, network size, training data volume and quality, prediction accuracy, and execution speed, through innovations in hybrid training methods. Possible directions include: adding physical constraints as regularizing terms in the loss function; novel strategies to generate more diverse and relevant training data from the target function; and optimising the order in which training points are generated or processed to speed up training.

We will specifically target particle-based models and applications areas including additive manufacturing (3D printing) and ore processing for critical energy metals. Extensions to this work include mathematical analyses of the ANN structures commonly used in practice, which will better inform the hybrid training process, and pursuing novel architectures that are well-suited to fast computations but which are easier to design.

Relevant publications

- This paper provides a comprehensive study of application of Deep Learning to solving multiple physics problems in a range of different application areas. This predominantly focuses around mesh based methods, where this PhD will look to move beyond this to mesh-free particle based methods. Raissi, M., P. Perdikaris, and G. E. Karniadakis. “Physics-Informed Neural Networks: A Deep Learning Framework for Solving Forward and Inverse Problems Involving Nonlinear Partial Differential Equations.” Journal of Computational Physics378 (February 1, 2019): 686–707. https://doi.org/10.1016/j.jcp.2018.10.045.
- This paper is an example of a recent application of machine learning for creating a surrogate model for the Discrete Element Method (a particle based method for modelling granular flows), demonstrating a significant speed with a minimal loss of accuracy on a set of simple test cases. This PhD project will look to a more sophisticated machine learning approach, where the network architecture employed directly takes account of the relevant physical laws and symmetries within the system. Lu, Liqiang, Xi Gao, Jean-François Dietiker, Mehrdad Shahnam, and William A. Rogers. “Machine Learning Accelerated Discrete Element Modeling of Granular Flows.” Chemical Engineering Science245 (December 2021): 116832. https://doi.org/10.1016/j.ces.2021.116832.

Skillset required

Degree in Computer Science / Physical Sciences/ Mathematical Sciences, ideally with a working knowledge of ML principles.

Contact Details

Richard Scalzo (richard.scalzo@data61.csiro.au)

Jonathan Manton, University of Melbourne (j.manton@ieee.org)

AI4M-PHD-08_Data Richness and the impact on geospatial data acquisition

Description

Many scientific programs rely upon publicly accessible geospatial datasets to augment predictions and support greater understanding of nature. Which data to use typically involves complex decision pathways with multiple outcomes possible from a single starting point. The choice of area and what additional data to acquire are critical decisions along this pathway. Risks also need to be considered, such as those associated with imperfect and clustered data, or the costs and timeframe of collecting additional data. Publicly accessible data is widely acknowledged as beneficial with recent estimates being between $20-$23 return for each dollar spent (ACIL Allen, 2015; ECS 2014). However, a quantitative assessment of the impact of data availability (density and distance), quality (relevance and precision) and attributes (clarity and characteristics) is not often performed to support acquisition decisions.

"Data Richness"

Quantitative assessment of data availability, quality and attributes combines to produce a 'data-richness' representation that will address an important gap in decision workflows. These measures can be combined across several data sets to give an overall estimate of 'data richness' (Aitken et al 2018, Ore Geol.Rev). Further, 'coupling', or interdependency, of datasets attempts to measure how much one the acquisition of one data set supports performance improvement in others. A national data richness map that does not yet exist and would provide a key representation of our current state of knowledge based on available datasets. As such, government departments, research organisations and industry use geospatial products without knowing the risk, and metrics that fail to capture the utility of public data for a particular need. This project will produce a workflow to provide national data richness assessments. These assessments will be domain specific, as the relevance of different datasets to scientific question will vary from geoscience, ecology, hydrology, agriculture, marine and atmospheric sciences. A case study examining the evolution of data richness over time will be conducted, that will understand the temporal inter-relationships between mineral discoveries and data collection.

Does data acquisition encourage mineral discovery, or mineral discovery encourage acquisition? Understanding this complex relationship will help to better understand the economic, social and environmental impact of government investment in large-scale publicly available geospatial datasets. This PhD will be supported with a top-up scholarship from the Research Office Science Leader Program.

Relevant publications

- Aitken et al 2018 Ore Geol.Rev. Published work on data richness as applied to mineral exploration in central Western Australia. Authored by CSIRO and University Supervisors
- Swink and Speier 1999 Decision Sciences. Seminal work examining cognitive load and visualisation of multiple data sets in geographic information systems.

Required skillsets (in no particular order)

- Geographic information systems; Data analytics; Scientific computing (Python, R, Matlab); Statistics (spatial is a bonus)

Contact Details

Dr. Mark Lindsay – CSIRO Mineral Resources,
A/Prof Alan Aitken - The University of Western Australia

AI4M-PHD-09_Fine grained spatio-temporal analysis of body/muscular motion for vital signs estimation

Description

Background and problem:

Advances in telehealth technologies with consumer-grade camera phones or webcams may provide valuable assistance to clinicians to not only communicate with patients but also to monitor their vital signs, such as heart rate and respiratory rate, remotely. In a wound care context, vital signs provide valuable indicators to clinicians about a patient's overall wellbeing and can provide key insights to clinicians about wound recovery and signs of infection. In a normal telehealth session, these quantities are often not readily available to the clinician at the time of consultation. Current digital wound care in Australian residential aged care facilities is often a time-consuming off-line assessment, leading to a lack of timely and coordinated care. This process frequently requires peripheral devices that increase expense and need for integration into video calls. To overcome these issues, we propose to estimate the patient's vital signs directly from a consumer-grade camera using spatio-temporal analysis of video streams and reconstructed 3D models of the patient's face and body.

While there has been a rapid development of vision-based systems providing accurate and granular quantification of patients' vital signs, these methods still require a significant advancement to be able to cope with natural conditions such as facial expressions, illumination changes and spontaneous movements. Existing 2D approaches are only able to provide reliable estimates in controlled settings with little motion (head rotations) and steady illumination (i.e. limited to handling translational and in-plane motion).

Proposed approach:

In this project, the student will investigate and develop novel AI-based methods to deal with the shortcomings of measuring vital signs such as heart rate and respiratory rate from video frames captured in a real-world environment before and during a telehealth visit. Aspirational aims include addressing the feasibility of measuring blood pressure and oxygen saturation using body texture and motion.

The student will investigate novel 3D neural field modelling to understand facial motion/feature, as well as upper body, pose 3D shape/stance, to extract vital signs and compare the performance with traditional 2D analysis. 3D modelling enables us to estimate correspondences over the entire video with pixel-level accuracy, even in the presence of out-of-plane or large motions.

External collaboration:

In addition to research skills, the student will develop transferable skills such as project management and communication as a result of the university-industry interaction. This project is done in collaboration with Coviu (https://www.coviu.com/en-au/), a leading telehealth company, under the Medical Research Future Fund program. This offers a unique opportunity for the student to work on cutting edge computer vision and deformable object modelling, to transform aged care facilities through telehealth.

Relevant publications

- Rapczynski, M., Werner, P., & Al-Hamadi, A. (2019). Effects of video encoding on camera-based heart rate estimation. *IEEE Transactions on Biomedical Engineering*, *66*(12), 3360-3370.
- Alafeef, M. (2017). Smartphone-based photoplethysmographic imaging for heart rate monitoring. *Journal of medical engineering & technology*, *41*(5), 387-395.

- Zhang, C., Gebhart, I., Kühmstedt, P., Rosenberger, M., & Notni, G. (2021, June). Real-time multimodal 3D imaging system for remote estimation of vital signs. In *Multimodal Sensing and Artificial Intelligence: Technologies and Applications II* (Vol. 11785, p. 117850F).

Skillset required

- 3D computer vision: 3D reconstruction, bundle adjustment, camera calibration, matrix algebra
- Machine learning: basic machine learning algorithms and deep learning models. Previous knowledge in AI-based 3D reconstruction is a differential.
- 3D visualisation: 3D mesh/point cloud processing and registration, 3D rendering on the web and smart devices
- Programming: Python and C++, in linux environment, and one or more ML platforms such as Pytorch or Tensorflow
- Hardware: cameras and sensors

Contact Details

Mohammad Ali Armin (ali.armin@data61.csiro.au)

Associate Professor Hanna Suominen , ANU (hanna.suominen@anu.edu.au)

AI4M-PHD-10_Global transitions and the rise of new industries: uncovering nascent technology trends to understand future material demand and impacts on stocks and flows

Description

Long term forecasting of material demand pertaining to global-scale changes like the energy transition, or nascent industries like hydrogen, robotics, or space are currently extremely simplistic. For example,2020 reports from CSIRO and The World Bank make metal projections to 2050 based on the current slate of battery chemistries. Already we know these projections to be out of date with recent announcements from BYD on the near-term commercialisation of Sodium-ion batteries possibly supplanting lithium as a core EV battery chemistry. The guiding hypothesis of this work is that AI approaches including ML and more advanced approaches to parse both structured and unstructured data to provide the basis for better prognostication of technology trends and the implications to material flows at a global scale. Sources include Academic literature, trade publications, news articles, social media. All these potentially hold insights for speculating as to the timing of possible commercialisation of new technologies that have ripple effects that alter the global supply and demand picture for materials. The proposed approach is to aggregate various data sources covering R&D development trajectories, market dynamics (industry reports), grey literature including company-level events and announcements, consumer sentiment and behaviours, regulatory shifts, etc. Processing of datasets with a focus on natural language processing such as Latent Dirichlet Allocation used for the discovery of topics and themes from text, and develop algorithms particularly designed to ingest, pre-process, and analyse the variety of data sources previously mentioned. Analysis will be both descriptive and predictive. Unsupervised machine learning methods will be employed for the descriptive component, as well as anomaly detection methods used to identify and describe outlier events and shifts in sentiment. Regression and classification models will be employed, for instance, to forecast product and raw material prices (for the case of regression models), and predict the adoption of or investment in particular technologies from key corporate players in renewable energy technology markets (for the case of classification models).

Relevant publications

- This document shows how simple changes in assumptions about battery technology changes affect out-year metal demand and that a global accounting framework is needed to track these fluctuations. West J., Ford J.A. and Meyers J. (2021). Known unknowns: the devil in the details of energy metal demand. Using an integrated physical framework to explore opportunities and risks for metals in the energy transition. CSIRO, Australia. https://www.csiro.au/-/media/Newsroom/Media-release-images/Metals-demand-report/21-00309_MR_REPORT_CriticalEnergyMinerals_FA.pdf
- This report presents the summary of a joint project between Data61Cyber-Physical Systems Program and UQ that developed a data-driven methodology to identify emerging market opportunities for the program. The text-based machine learning approach employed in the study and corresponding datasets will influence the current proposal. Brea E, Hine D, Green T, Oyola G and Elfes A (2021).Mapping emerging opportunities for the Cyber-Physical Systems Program. CSIRO and The University of Queensland, Australia. https://drive.google.com/file/d/1uLht-G0oMy2obwJRcx9vzzdSrALrzBlq/view?usp=sharing
- This working paper is currently under review by an A* journal. It presents three research studies using unsupervised machine learning techniques and big datasets to examine technology, science and innovation activity. The proposed project will employ these techniques and datasets. Brea E., Hine D. and Kastelle T. (2022). Unsupervised data mining in innovation research: Lessons from three exploratory studies. The University of Queensland,Australia.

https://docs.google.com/document/d/1KdKnZhjYlWmE1sTVh3S5tqWILT0dMCqg/edit?usp=sharing&ouid=108028776307875897054&rtpof=true&sd=true

Skillset required

Familiarity with basic programming languages and statistical tools, and working knowledge of AI and ML packages.

Contact Details

Tim Baynes, CSIRO (tim.baynes@csiro.ua)

Prof. Martie-Louise Verreynne, University of Queensland (m.verreynne@uq.edu.au)

AI4M-PHD-11_Identification of environmental risk factors for antimicrobial resistance using spatio-temporal statistical models

Description

Antimicrobial resistance (AMR) poses a global threat (O'Neill, Welcome Trust 2016). Understanding the drivers and risk factors of AMR prevalence and spread is critical in the fight to maintain society's ability to combat infection using existing antibiotics. The least-well understood component of the one-health perspective to AMR is the role that the environmental component plays in eventual clinical presentations of AMR (Ott et al., 2021). This is due in large part to the complexity of the environmental processes at play, inconsistent observation methods and relatively sparse datasets. This project proposes a statistical approach that will be applied to a specific freshwater case, but in a way that can be readily generalised to other marine or freshwater applications or case studies. In particular, the project will identify generalisable (surrogate) AMR markers and use them as response variable(s) in a spatio-temporal model that can identify the effect of environmental covariates using observations with point (such as water quality) or areal (such as land use) characteristics. The complexity of the environmental process involved and the consequent non-linearities suggest that classical statistical frameworks will be unlikely to succeed. To overcome this, we propose to use some of the most recently developed deep learning algorithms, known for their ability to capture patterns in environmental setups (see for example, Amato et al, 2020 and Kuhnleinet al, 2014). The model will also integrate specific catchment features and demonstrate an approach to tackle the classical change-of-support problem often encountered in ecological and environmental modelling, wherein observations at point locations are related to environmental characteristics defined over areal locations. The three key steps in the project are:1. In collaboration with our UniSA collaborators, identify an appropriate AMR surrogate that is appropriate for freshwater ecosystems but generalisable to other (e.g.marine) contexts. 2.Propose (and code) a spatio-temporal model able to capture the spatio-temporal variations in an observed time series of the AMR surrogate and infer the impact of relevant risk factors. 3.Explore the limitations of the proposed model in terms of data requirements, model choice and the effects of uncertainty on decisions regarding the most efficient risk mitigation strategies.

Relevant publications

- Ickowicz, A., Ford, J. & Hayes, K. A Mixture Model Approach for Compositional Data: Inferring Land-Use Influence on Point-Referenced Water Quality Measurements.JABES.24,719–739 (2019).

Skillset required

- Strong theoretical background in data science / AI / statistics
- Strong interest environmental sciences, acknowledged by previous activities or experience
- Strong computer skills including in programming environment (R, Python)

Contact Details

Keith Hayes, CSIRO (keith.hayes@csiro.au)

Erica Donner, University of South Australia (erica.donner@unisa.edu.au)

AI4M-PHD-12_Implicit Neural Representation for 3D scene modelling

Description

3D scene reconstruction is a classical problem in computer vision with numerous applications from robotic mapping and navigation to entertainment, human-computer interaction, and medical imaging. This project will address the issue of providing fast and affordable 3D Neural scene representations from 3D laser scanning (Lidar) and 360° HD images.

During this project, the student will investigate how to leverage recent learning-based 3D neural representations (such as signed distance functions, NeRFs,...) for the reconstruction of a dynamic scene. A strong emphasis will be placed on the algorithm's robustness to reconstruct a dynamic environment that comprises transient objects whose position can change as images are captured over time (cars, pedestrians,…).

The appointed student will be co-supervised and supported by the CSIRO's robotic group and have access to in-house datasets and have the possibility to acquire new data to test the developed algorithms in a real-world context. The student will be supervised by a multi-disciplinary research team combining experts in Computer Vision, Robotics and applied Maths; The student will also receive access to CSIRO super-computing resources.

The research outcome is expected to be published in top machine learning and computer vision venues (NeurIPS, ICML, AAAI, CVPR, ECCV, ICCV,..) allowing the student to deliver a high-quality PhD dissertation recognised internationally.

Relevant publications

- [ConvONET] A flexible implicit representation for accurate large-scale 3D reconstruction https://pengsongyou.github.io/conv_onet
- [Urban radiance field] 3D reconstructions and novel view synthesis from data captured by mobile scanning platforms https://urban-radiance-fields.github.io/
- [Block-Nerf] Large Scale Scene Neural View Synthesis https://waymo.com/research/block-nerf/
- [Neural Scene Flow Fields] https://www.cs.cornell.edu/~zl548/NSFF/
- [SIRENs] Implicit Neural Representations with Periodic Activation Functions https://arxiv.org/abs/2006.09661

Skillset required

- Degree in Computer Science / Electrical Engineering / Mathematical Sciences
- Strong mathematical background, ideally knowledgeable in computer vision and multiple view geometry: Structure-from-Motion (SfM), Multi-View Stereo (MVS), Bundle adjustment.
- Understanding of working in a Linux environment.
- Good programming skills and knowledge of deep learning libraries (TensorFlow/PyTorch).

Contact Details

Dr Leo Lebrat, CSIRO (Leo.Lebrat@csiro.au)

Prof Clinton Fookes, Queensland University of Technology (c.fookes@qut.edu.au)

AI4M-PHD-13_Linking molecules to sensory traits for next-gen plant-based milk

Description

Plant-based milk is now mainstream. In 2020, sales of the category grew twice as fast as cow's milk in the US. Oat milk sales more than tripled, to $274 million, and almond milk topped $1.6 billion. Still, many consumers still haven't made the switch from traditional milk, in part because of what is deemed inferior taste and texture. In this project, the PhD candidate supported by a team of multi-disciplinary scientists spanning food chemistry, sensory science, analytical chemistry, -omics and health & nutrition, will deconstruct the classical dairy drinking experience. Considerable progress has been made towards understanding the sensory and flavour chemistry/ biochemistry of traditional milk and milk products, but this accumulated knowledge base of decades of research is tied closely with providing key input on challenges involved in production of high-quality milk and milk products at scale consistently. There is also good understanding on the variation in traditional milk flavour and gross chemical composition caused by breed of cow, feed etc. Defining fundamental sensory/flavour and mouthfeel experience of drinking traditional milk based on molecular make-up, presence of variety of macro-molecular structures (milk fat globules/emulsion, casein protein/calcium phosphate complex known as casein micelles) and physico-chemical properties will require a multi-disciplinary approach. This generic knowledge base can then be applied to simulate similar sensory experience in drinking of non-traditional or plant-based milk. One commercial example of this has already appeared in the US market, under the brand name "NotCoMilk". This product has been designed using artificial intelligence to inform a plant-based alternative beverage to better mimic the taste, nutrition and experience of dairy milk. The PhD candidate will focus on understanding critical areas such as the molecular composition, sensory elements and the key nutritional factors of leading dairy alternative products in the market and how these compare to that of dairy. They will also generate fundamental understanding of the compositional and sensory properties of the main ingredients used to produce plant-based dairy milk alternatives to help avoid the less appealing attributes like sourness, bitterness, and sulphur notes that are common in these products. They will also scope other non-traditional ingredients that might help achieve these aims. that the student will utilise the developed knowledge of composition and sensory properties to explore their use in dairy products, spoonables and plant-based cheese. This will require optimisation of critical functionality (i.e., gelation properties); sensory attributes, such as mouthfeel and flavour profile. By linking the molecular composition with functional and flavour data, the PhD candidate can work with Data61 scientists to simulate and visualise key protein structures and their interactions with small molecules (e.g., aroma molecules). All this information will help create the basis for an ingredient database that will include AI elements for ingredient selection and process development that will be explored in subsequent projects to deliver the precise molecular recipe responsible for an enhanced alternative dairy experience, "Like Milk". This project would represent a collaborative partnership between CSIRO, Deakin and Noumi bringing together complementary research capabilities and an industry partner who can translate the research to market as optimised plant-based dairy beverages

Relevant publications

- Paul, A.; Kumar, S.; Kumar, V.; Sharma, R. 2019. Milk Analog: Plant based alternatives to conventional milk, production, potential and health concerns. Critical Reviews in Food Science and Nutrition, 60, 305-3023. https://www.tandfonline.com/doi/full/10.1080/10408398.2019.1674243

- McClements, D.J.; Newman, E.; McClements, I.F. 2019. Plant-based Milks: A Review of the Science Underpinning Their Design, Fabrication, and Performance. Comprehensive Reviews in Food Science and Food Safety, 18, 2047-2067. https://ift.onlinelibrary.wiley.com/doi/10.1111/1541-4337.12505

- Mavani, N.R., Ali, J.M., Othman, S. et al. 2022. Application of Artificial Intelligence in Food Industry—a Guideline. Food Eng Rev 14, 134–175. https://link.springer.com/article/10.1007/s12393-021-09290-z

Skillset required

Degree in chemistry or food science with an understanding of basic biochemistry principles and lab-based experience. Experience in data analysis and project Management

Contact Details

Dr Joanna M. Gambetta (Joanna.gambetta@csiro.au)

Prof Russell Keast, Deakin University (russell.keast@deakin.edu.au)

AI4M-PHD-14_Rapid, high quality, clinical grade, 3D surface reconstruction of the human body from consumer-based videos

Description

Wound care in Australian residential aged care facilities has several challenges. The care quality depends on staff experience requiring health specialists, resulting in resident mobility. To avoid the risk and costs associated with resident mobility, telehealth solutions can be adopted. However, the current telehealth solutions for wound care follow a time-consuming off-line assessment approach leading to a lack of timely and coordinated care that reduces its effectiveness.

To overcome these issues, we propose to develop an advanced 3D reconstruction and analysis technology of wounds from images taken with a consumer-grade camera. This new technology will enable a breakthrough in this area by providing objective and precise clinical data (e.g., 3D wound measurements, tissue composition, and wound classification) to health professionals allowing them to provide more effective care through a telehealth apparatus.

Current challenges in human body reconstruction from consumer grade cameras include uncontrolled illumination, self-similarity of the texture on the skin surface, deformable surfaces, and computational requirements. Further constraints in clinical scenarios include uncontrolled motion of the patient during image acquisition which leads to blurry images, inaccurate reconstruction, and significant errors in the clinical measurements.

Hence, our goal is to conceive a technology to efficiently reconstruct a high-resolution, high quality 3D model of the body from consumer-based videos.

Proposed approach: The student will start by investigating classical methods like structure-from-motion (SfM) and non-rigid SfM, as well as very recent AI developments like Neural Radiance Fields (NeRF) and Neural Implicit Surfaces (NeuS), see references below, reasoning about their limitations and potential. In our proposed full-body reconstruction scenario, there will be issues stemming from poor image quality due to illumination and motion blur. In addition, classical methods are likely to have a limited accuracy on the self-similar human skin surface, while the AI-based approaches would also require significant compute resources. Hence, the student will explore new ideas to overcome these issues contributing to the 3D reconstruction field of computer vision as well as advancing new approaches suitable for AI deployments on edge devices within a clinical health setting.

External collaboration: In addition to research skills, the student will develop transferable skills such as project management and communication as a result of the university-industry interaction. This project is done in collaboration with Coviu (https://www.coviu.com/en-au/), a leading telehealth company, under the Medical Research Future Fund program. This offers a unique opportunity for the student to work on cutting edge computer vision and deformable object modelling, to transform aged care facilities through telehealth.

Relevant publications

- Mildenhall, B., Srinivasan, P. P., Tancik, M., Barron, J. T., Ramamoorthi, R., & Ng, R. (2020, August). Nerf: Representing scenes as neural radiance fields for view synthesis. In *European conference on computer vision* (pp. 405-421). Springer, Cham. [Paper][Preprint][Webpage]
- Wang, P., Liu, L., Liu, Y., Theobalt, C., Komura, T., & Wang, W. (2021). NeuS: Learning Neural Implicit Surfaces by Volume Rendering for Multi-view Reconstruction. *Advances in Neural Information Processing Systems*, 34. [Paper][Preprint][Webpage]

- Zoppo, G., Marrone, F., Pittarello, M., Farina, M., Uberti, A., Demarchi, D., ... & Ricci, E. (2020). AI technology for remote clinical assessment and monitoring. *Journal of wound care*, *29*(12), 692-706. [Paper]

Skillset required

- Basic knowledge of machine learning algorithms and deep learning models. Knowledge in AI-based 3D reconstruction is highly desirable, but not essential.
- Basic knowledge of 3D computer vision and multiple view geometry. Previous experience with 3D reconstruction algorithms is also highly desirable, but not essential.
- Excellent programming skills in Python and C++ using Linux environments.

Contact Details

Dr Rodrigo Santa Cruz, CSIRO (rodrigo.santacruz@csiro.au)

Prof Clinton Fookes, Queensland University of Technology (c.fookes@qut.edu.au)

AI4M-PHD-15_Reliable deep model validation using cheap data for next-generation phenotyping

Description

Deep learning algorithms have achieved great successes in a wide variety of applications and have been widely deployed in many daily life sectors. Deep models require a large amount of data to train and update. However, it is expensive, time-consuming, or even infeasible to accurately label large-scale datasets. To reduce the annotation cost, cheap methods such as crowdsourcing labelling and automated annotation methods have been widely used. For example, with CSIRO video phenomics project, oats spikelets were manually labelled in video frames. Because the oat spikelets are very dense (> 6,000 per video frame), it is very laborious, and mislabelling cannot be avoided. How to use the noisy labels to train and validate a reliable deep learning model has been a challenging task.

Although cheap data with noisy labels has now been widely exploited to train deep models, little study has been done on selecting reliable models using noisy data. Note that training models with noisy labels is very different from selecting models with noisy data. The former is to prevent the complex model from overfitting/memorising the label noise while the latter is to select/validate the reliability of a given model. As deep algorithms are usually non-convex and very complex, a lot of intermediate models are generated during the training procedure. Usually, a clean validation set will be used to select a reliable model from the intermediate ones for deployment. However, if the validation set itself contains label noise, the selected model will not be optimal, which causes performance degeneration and raises significant concerns on its reliability, especially for decision-critical applications. Current methods developed for learning with noisy labels cannot be directly employed for model validation.

In this project, we will study model validation using noisy data by modelling the label noise, using the relationship between the noisy and clean validation data. One intuitive way is to make use of the training data by assuming that the label noise pattern is invariant from the training data to validation data. Then, utilising the noisy validation data and the relationship between the noisy and clean data learned from the noisy training data to select a reliable model. We will further investigate the more challenging case when the label noise pattern varies from training and validation sets and the distribution of the clean validation data varies. We will use the techniques developed for transfer learning to resolve the issue.

The research outcome is expected to be published at top machine learning venues (such as NeurIPS, ICML, ICLR, AISTATS, UAI). The results obtained will be sufficient for the student to deliver a high-quality PhD thesis on the topic: Reliable deep model validation using cheap data with noisy labels.

Student activities include:

- Conducting literature review of existing methods for learning with label noise;
    - Developing general methodologies for model selection without using clean data; conducting experiments by coding using Python;
    - Analysing the statistical property, e.g., consistency, of the designed model selection method;
    - Conducting AI-based simulation studies, case studies/analysis of real phenotyping data, paper writing for journals and conferences.

Relevant publications

(1) Natarajan, Nagarajan, Inderjit S. Dhillon, Pradeep K. Ravikumar, and Ambuj Tewari. "Learning with noisy labels." Advances in neural information processing systems, 2013.

(2) Chen, Pengfei, Ben Ben Liao, Guangyong Chen, and Shengyu Zhang. "Understanding and utilizing deep neural networks trained with noisy labels." In International Conference on Machine Learning, 2019.

(3) Northcutt, Curtis G., Anish Athalye, and Jonas Mueller. "Pervasive label errors in test sets destabilize machine learning benchmarks." Advances in neural information processing systems, 2021.

Skillset required

- Strong mathematical background, knowledgeable in statistical analysis method
- Strong experience and skills in programming (e.g., Python)
- Knowledgeable in machine learning-based modelling.

Contact Details

Dr Dadong Wang, CSIRO Data61 (Dadong.Wang@csiro.au)

Dr Tongliang Liu, The University of Sydney (tongliang.liu@sydney.edu.au)

AI4M-PHD-16_Towards Enabling Real-time and Secure Geospatial Data Analytics

Description

Advances of the sensor and GPS techniques have motivated the proliferation of geospatial data. According to a report by Research and Markets, the global geospatial analytics market is expected to grow from $38.65 billion in 2017 to reach $174.65 billion in 2027 with a CAGR of 18.2%. Capabilities in effective real-time geospatial data analytics will empower business and organizations to identify markets and revenue sources, deliver timely and tailored marketing campaigns, and make informed decisions in business planning.

Despite the enormous benefits of geospatial data analytics, the large volume, multi-modality, dynamic updates, and the timely requirements of analytics tasks all pose great challenges. In this project, we will investigate two fundamental tasks over geospatial data, large-scale trajectory data analytics and path-based query processing. Task 1: Representation Learning for Privacy Preserving Trajectory Analytics at Scale. A trajectory records the human behaviour traces by a sequence of geospatial locations. Effective trajectory analytics extract movement patterns, revealing insightful information for downstream tasks such as location-based products or stores recommendation. The high computation cost of trajectory analytics measures such as clustering and similarity computation has become the de facto bottleneck for trajectory analytics at scale.

Moreover, the trajectory data is usually sensitive and the data privacy should be guaranteed. In this task, we will propose an AI-based model for trajectory analysis which will effectively reduce the high computation costs based on the state-of-the-art techniques including attentive neural networks, transformers, and contrastive learning. We will also design computing methods to allow parallel or distributed computation on high performance computing hardware to improve the efficiency and scalability of the model. To offer better data privacy, we will explore privacy-preserving machine learning techniques such as concentrated differential privacy, private aggregation of teacher ensembles and federated learning for trajectory representation learning. Furthermore, we will exploit reciprocal learning and augment the simulation training data to increase our model's robustness against potential adversarial attacks. Task 2: Efficient Path-based Query over Multi-modal Geospatial Data Geo-locations often accompany rich semantic information acquired from numerous online sources including business directories such as Google My Business, location-based social networks such as Foursquare, rate and review services such as TripAdvisor, and spatial check-in information on social platforms such as Facebook. The proliferation of such semantic data greatly enriches and mandates more intuitive integration with geospatial data. In this task we investigate path-based query over geo-spatial data enriched with semantic information. We model real-world road networks as time-dependent networks where an edge is associated with a time-dependent weight function for modelling speed during different time of the day and different day of the week (e.g., weekday vs weekend, public holiday). The first challenge originates from high computational costs over time-dependent road networks. A simple path query (e.g., s-t path with a source and destination) adopts A* or Dijkstra style approaches will incur prohibitively high computation cost in dozens of seconds over large time-dependent road networks. To address this challenge, we will explore index-based solutions which support fast route queries by employing 2-hop labelling and tree-decomposition techniques. The second challenge is to effectively fuse multi-modal semantic data, enabling more intuitive queries such as "find the beach within 45 mins walking distance, with good sushi restaurants and cafe". We will model the surrounding environment as interconnected entities in a graph database and explore their relationship. Novel interactive searching strategies utilizing graph neural networks will be developed to retrieve real-time and tailored recommendations

Relevant publications

- Yang, Peilun, et al. "T3S: effective representation learning for trajectory similarity computation."ICDE 2021. This paper presents state-of-the-art techniques for trajectory representation learning which effectively reduces the time complexity for trajectory similarity computation.
- Sun, Hao, Zhiqun Zhao, and Zhihai He. "Reciprocal learning networks for human trajectory prediction."CVPR2020.
- Yuli Jiang, Yu Rong, Hong Cheng, Xin Huang, Kangfei Zhao, Junzhou Huang, "Query Driven-Graph Neural Networks for Community Search: From Non-Attributed, Attributed to Interactive Attributed", PVLDB2022.

A novel methods for query processing over attributed graphs using graph neural networks.

Skillset required

The project requires capabilities in machine learning, algorithms, and graph theory.

Contact Details

Wei Ni, CSIRO Data61 (Wei.Ni@csiro.au)

Wenjie Zhang, University of New South Wales (wenjie.zhang at unsw.edu.au)