Computer Science and Technology
Research Students' Symposium

19-21 May 2021

Programme

Wednesday 19 May

2:00PM–2:20PM · Stage
Welcome and Introduction
Professor Ann Copestake, Head of Department, will welcome the attendees

2:20PM–2:40PM · Stage
Dr Andreas Vlachos
Keynote talk: Fact-checking as a conversation

2:40PM–3:00PM · Stage
Dr Dong Ma
Keynote talk: OESense: Employing Occlusion Effect for In-ear Human Sensing

3:00PM–3:15PM · Stage
Seng Ah (Michelle) Lee
Context-conscious algorithmic fairness in machine learning

3:15PM–3:30PM · Stage
Aamir Mustafa
Image-to-Image Translation Using Less Data and Perceptual Loss Function

3:30PM–3:45PM · Stage
Tea break

3:45PM–4:00PM · Stage
Param Hanji
Inverse problems in graphics and vision

4:00PM–4:15PM · Stage
Dmitry Kazhdan
Concept-based Explanation Approaches for Deep Neural Networks

4:15PM–4:30PM · Stage
Eric Meissner
Building Better, More Fair AI Systems with AutoAI

4:30PM–4:45PM · Stage
Diana Robinson
Probabilistic Modelling of Postoperative Bleeding Decisions

4:45PM–5:00PM · Stage
Sami Alabed
Role of Computer Systems Probabilistic Graphical Models in Optimization

5:00PM–5:45PM · Sessions
- Human-Centred Computing theme meeting
- Security theme meeting
- Algorithms and Complexity theme meeting

5:45PM–6:00PM · Stage
End of day 1
See you tomorrow

Michelle Seng Ah Lee *supervised by Dr J. Singh*

Context-conscious algorithmic fairness in machine learning

Algorithms are increasingly used in high-impact domains, but there is little consensus on how to define and mitigate unfair bias. Because the topic is cross-disciplinary, it is challenging to agree on terminology and approach. My work aims to introduce an end-to-end context-conscious methodology to govern the risk of unfair bias in ML. First, I will share the three challenges I have addressed so far in my papers: 1) the gap between how fairness is considered in computer science literature vs. in ethical philosophy and welfare economics, 2) barriers to implementation due to gaps in open source toolkits, and 3) unintended biases introduced through people and process rather than technology and data. Then, I will present the three key questions driving the next phase of my research: 1) why are some models better at triangulating sensitive features, 2) how do we separate the proxies of risk from proxies of sensitive features, and 3) how do we define fairness under uncertainty in a dynamic system?

Aamir Mustafa *supervised by Dr R. Mantiuk*

Image-to-Image Translation Using Less Data and Perceptual Loss Function

Although deep neural networks have achieved immense success in a wide range of computer vision applications, they owe much of their success to large collections of labelled datasets. In real-world scenarios creating these extensive datasets is expensive requiring time-consuming human labelling, e.g. expert annotators, as in case of medical predictions and artistic reconstructions. Scarcity of labelled data has motivated the development of semi-supervised learning methods, which learn from large portions of unlabelled data alongside a few labelled samples. The primary objective of my work is to come up with image-to-image (I2I) translation models that require minimum amount of labelled data for training.

As part of my initial work, I propose a semi-supervised training scheme for I2I settings called Transformation Consistency Regularization, which introduces a diverse set of geometric transformations and enforces the model's predictions for unlabelled data to be invariant to those transformations. The method is significantly data efficient, requiring only around 10 - 20% of labelled samples to achieve similar image reconstructions to its fully supervised counterpart.

Param Hanji *supervised by Dr R. Mantiuk*

Inverse problems in graphics and vision

Several problems in computer graphics and vision can be modeled as invertible functions. Commonly encountered problems also tend to be easy in the forward direction, but their inverses are often ill-posed. Estimation of a set of hidden parameters from multiple noisy observations is one such problem, that formed the focus of my research during my first year.

Dmitry Kazhdan *supervised by Prof. M. Jamnik and Prof. P. Liò*

Concept-based Explanation Approaches for Deep Neural Networks

Deep Neural Networks (DNNs) have achieved remarkable performance on a range of tasks. A key step to further empowering DNN-based approaches is improving their explainability and interpretability. Our work explores and expands on a novel class of explanation approaches - referred to as 'concept-based explanations'. Concept-based explanation approaches provide model explanations in terms of human-understandable units, rather than individual features, pixels, or characters (e.g., the concepts of a wheel and a door are important for the detection of cars). This class of approaches primarily focuses on (i) explaining trained DNN models in a post-hoc fashion (analysing which concepts they rely on), (ii) introducing concept-based models which are interpretable by design (explicitly forcing a model to rely on particular concepts). In our work we analyse, discuss and expand on these approaches, emphasising their primary strengths, limitations, and relations to other well-known subfields, such as disentanglement learning.

Eric Meissner *supervised by Prof N. D. Lawrence*

Building Better, More Fair AI Systems with AutoAI

Fairness research in machine learning has received a lot in the last few years, and has converged on the idea that any definition of fairness is inherently subjective and contextual. This implies that what we mean by fair differs across people and contexts, so there is no "solution" to the problem. So every problem needs to be evaluated from a variety of stakeholder perspectives (the individuals the system impacts, the owners and maintainers of the system, government, etc.) to guide the decision making process in implementing fair AI systems. However currently, it is difficult to explain what complex AI systems are doing to ourselves, much less to the general public. In order to move the field towards this participatory, socially inclusive future we first need to understand what is happening. This leads us to the AutoAI project, which is focused around building FIT (fair, interpretable, transparent) AI systems in the real world, through a combination of data-oriented architectures, shadow system emulation, and counterfactual simulations. These properties will form the foundation of future work, and provide the necessary tools for those outside the techno-sphere to have input on how AI systems are built and operate.

Diana Robinson *supervised by Prof. A. F. Blackwell*

Probabilistic Modelling of Postoperative Bleeding Decisions

Diana's research focus is on representing and modeling uncertainty in decisions around postoperative bleeding following cardiac surgery. Her aim is to see whether a tool that is built on probabilistic programming could be useful in incorporating clinical assumptions, modeling the uncertainty explicitly and reasoning through to likelihood estimates of the next best course of action. One of her interests is in how quantifying uncertainty might impact the social process of decision making, particularly the coordination and handovers between multidisciplinary teams. By developing this from an end user perspective, and co-designing with clinicians, her aim is for clinicians themselves to eventually be able to build and adapt their own models to inform their decisions.

Sami Alabed *supervised by Dr E. Yoneki*

Role of Computer Systems Probabilistic Graphical Models in Optimization

Injecting expert knowledge in machine learning optimization methods such as Bayesian Optimization has many advantages from robustness to speed of convergence. I am looking at automating the discovery of probabilistic graphical models of computer systems.