

September Success Stories from the School of Informatics

Professor Gordon Plotkin has been selected by the British Computer Society as the 2018 winner of the Lovelace Medal. Gordon, a Professor in the School of Informatics, University of Edinburgh, is a theoretical computer scientist best known for his structural approach to computer programming languages, which describes how the individual steps of a computer-based calculation takes place. Gordon has also contributed to the fields of artificial intelligence, logic and linguistics. The BCS Lovelace Medal is the top award in computing in the UK and is awarded by BCS, The Chartered Institute for IT. The award is presented annually to individuals who, in the opinion of BCS, have made a significant contribution to the advancement of Information Systems.

Professor Mark Steedman was awarded 2018 Lifetime Achievement Award by the Association for Computational Linguistics (ACL). The ceremony took place during its 56th annual meeting in Melbourne, Australia (July 15 – July 20, 2018). Mark Steedman is Professor of Cognitive Science in the School of Informatics at the University of Edinburgh. He has made outstanding contributions to Computational Linguistics in particular in the areas of syntax and semantics. He is the main developer of Combinatory Categorial Grammar (CCG), a mildly context-sensitive grammar formalism that induces semantic representations in typed lambda calculus. Through the development of efficient and accurate broad-coverage parsers, CCG has become one of the most influential grammar formalisms in the field. The Association for Computational Linguistics (ACL) is the premier international scientific and professional society for people working on computational problems involving human language, a field often referred to as either computational linguistics or natural language processing (NLP).

Dr Paul Patras has been elevated to the grade of Senior Member of Institute of Electrical and Electronics Engineers (IEEE). Senior Members is the highest professional grade of IEEE for which a member may apply. It requires extensive experience and reflects professional accomplishment and maturity. Only 10% out of more than 400,000 members have achieved this level. IEEE is the world's largest technical professional organization dedicated to advancing technology for the benefit of humanity.

The University Defence Research onsoritum (UDRC) Phase 3 "Signal Processing in the Information Age" project is funding a consortium of Universities to develop underpinning signal processing and machine learning methodologies that are likely to be valuable in defence applications. Partners include the University of Edinburgh, Heriot-Watt University, University of Strathclyde and Queen's University Belfast. The School of Informatics with Tim Hospedales as a PI is leading the sub-project on Verifiable Deep Learning. Deep Neural Networks (DNNs) are deployed in increasingly many and increasingly mission-critical applications due to their impressive performance in a variety of areas. It is thus increasingly important to be confident that their outputs can be relied upon for decision-making. However as relatively "black box" systems, they have so far been hard to analyse in order to provide any kind of guarantee that they will always function as expected, particularly when exposed to new, unexpected, and potentially adversarial situations and inputs. This project will develop methods to certify and verify that DNNs are fit for purpose, even when extrapolating in the presence of novel inputs. This will involve developing methods that can provide limited

mathematical guarantees on DNNs' behaviour, despite their black box nature; and also methods for explaining their reasoning process so that their decisions can be manually checked for validity. This research should allow us to benefit from the applications being enabled by DNN's high performance, while relaxing in the confidence that they are doing the "right thing".

Professor Wenfei Fan received a Royal Society Wolfson Merit Award for his Querying Big Data with Limited Resources project. Big data is the next frontier for innovation, competition and productivity. However, big data analytics is often prohibitively costly. It can take days to combine information from tables with 10 million tuples. In other words, computation problems that are traditionally considered "tractable" in the classical theory of computation may nevertheless be impossible in practice. Wenfei's project will develop a new query evaluation paradigm that starts from the observation that with right auxiliary structures, many queries can be answered from a small subset of the data -- one whose size can be bounded in advance of evaluation. Researchers can hence reduce such "boundedly evaluable" queries on big data to computations on small data. For queries that are not bounded, they will introduce a scheme to compute approximate answers with accuracy guarantees, again by accessing a bounded amount of data. The characterisations of boundedly evaluable queries and the new approximation scheme extend the theory of scalable querying of big data that was proposed earlier. Big data analysis is currently the province of organisations that have massive computational resources at their disposal. This project has a potential to make it feasible for small businesses and researchers.

Pavlos Petoumenos secured funding for his project "Deep Learning For Easier Compiler Analysis and Optimisation". The project will use deep learning to understand and analyse computer code to enable novel optimisations for high performing software. Computing platforms are becoming more powerful but also more complex and harder to program. Producing code that uses the hardware efficiently requires time and expertise. Most developers lack either or both of them, so applications are inefficient, wasting processing power and electricity. Through deep neural networks, this project will remove human expertise from this process, making it easy, fast, and low cost to generate programs adapted to their hardware.

Lexi Birch secured funding for her project MTStretch: Low-resource Machine Translation. Neural machine translation (NMT) has recently made major advances in translation quality, however high performing neural models require many millions of human translated sentences for training. For many real-world applications, there is not enough data to build useful MT systems. In this project researchers plan to stretch the resources and capabilities that they have, in order to develop robust MT technologies which can be deployed for low-resource language pairs and for highly specialised low-resource domains. They will investigate making translation significantly more robust by using the intuition that translated (or parallel) corpora contain enormous redundancies, and are an inefficient way to learn to translate. Inspired by human learning, they will study Bayesian models which build up meaning compositionally and are able to learn to learn, thus creating models which only need a few training examples. They will also develop machine learning techniques, such as transfer learning and data augmentation, to extract knowledge from monolingual and parallel resources from other languages and domains. These advances will allow their partners, BBC World Service and BBC Monitoring, to cover under-resourced languages.

Hakan Bilen secured funding for his project Joint Inference Multi-task Networks. When we look at an image that contains a number of desks, chairs and a whiteboard, we can effortlessly predict that it is a classroom. Surprisingly even this simple problem requires us to coherently solve multiple tasks such as categorising the objects (e.g. chair, desk, whiteboard), estimating the depth and geometry of the room, spatial relations between the objects. During this process, our perception benefits from interactions among these tasks such that solving one would render solving another one easier (e.g. knowing the floor plane can help to find chairs, knowing the chair locations can conversely help to detect the floor). Machines would divide the same problem into smaller bits, but solve each subproblem disjointly and ignore their interactions. This paradigm prevents machines to develop a holistic understanding of visual world and thus to produce accurate predictions. This project will address such limitations by introducing collaborative problem solving strategies for machine perception and boost the performance of machine vision.