

April Success Stories from the School of Informatics

Two informatics PhD students received 2019 Google Fellowships: Nicolai Oswald in Systems and Networking and Reinald Kim Amplayo in Natural Language Processing (including Information Retrieval and Extraction). Google PhD Fellowships include support for students as well as a Google Research Mentor. The Program was created to recognize outstanding graduate students doing exceptional work in computer science and related research areas.

Lexi Birch and colleagues were awarded £730000 from Horizon 2020 to work on their GoURMET (Global Under-Resourced MEdia Translation) project. GoURMET is an EU project co-ordinated by the University of Edinburgh in partnership with the University of Amsterdam, the University of Alicante, the BBC and Deutsche Welle. It runs from Jan 2019 to the end of 2021. Machine translation (MT) is an increasingly important technology for supporting communication in a globalised world. Recent advances in neural machine translation (NMT) have resulted in significant interest in industry and have led to very rapid adoption of the new paradigm. However these models are data intensive and require parallel corpora of many millions of humanly translated sentences for training. Neural machine translation is currently not able to deliver usable translations for the vast majority of language pairs in the world. This is especially problematic for our user partners, the BBC and DW who need access to fast and accurate translation for languages with very few resources. The aim of GoURMET is to significantly improve the robustness and applicability of neural machine translation for low-resource language pairs and domains. The outputs of the project will be field-tested at partners BBC and DW, and the platform will be further validated through innovation intensives such as the workshops centred around our user group and BBC NewsHacks.

<https://gourmet-project.eu/>

Subramanian Ramamoorthy received funding of £263000 from the Turing Institute to work on his project Safety Critical Artificial Intelligence over two years (2019 – 20). AI is more and more often deployed in safety-critical applications involving physical interaction between humans and machines. This raises several new challenges that must necessarily be addressed if these technologies are to realize their potential benefits. The first challenge is the need for robust decision making despite noisy sensing, dynamic environments and unforeseen events. The second challenge pertains to the complexity of specifications required to capture the nature of planning and control in complex domains. This project addresses these challenges using a suite of techniques including interactively learning task specifications from demonstration data, formally analysing and introspecting about the properties of the learned models to establish safety properties, and then to synthesise correct by construction policies for robot motion. These methods will be applied in the domain of surgical assistance. Working with Co-I, Dr Paul Brennan from UoE Clinical Brain Sciences and NHS Lothian, researchers will apply these methods to a set of problems arising in the operating theatre, wherein robotic systems will provide assistance to the surgical team.

Boris Grot received \$73000 from a Google Faculty Research Award to work on his Accelerating Address Translation via a Learned Page Table Index Project. Address translation is an established performance bottleneck in workloads operating on large datasets due to frequent TLB misses and subsequent page table walks that often require multiple memory accesses to resolve. Inspired by recent research at Google on Learned Index Structures, this project will explore a radical alternative to traditional page table walks based on learned models using neural networks. A key question explored in this research will be how to microarchitect a hardware-friendly learned page table indexing scheme with respect to latency, area, and accuracy.

Catherine Lai was awarded £81,000 to continue the project 'Spoken Language Processing for Robot Companions' funded by Toyota Motor Europe. Access to speech-based interfaces can be crucial to navigate the modern world, especially for people with reduced physical mobility. The goal of this project is to improve speech technologies for assistive robots to enable more satisfying, personalized interactions which can help build long-term human-robot rapport. In particular, this project focuses on developing speech technologies to detect both what people are talking about and how they feel about specific topics in conversational speech (i.e. automatic topic and speaker stance detection), in order to develop useful and engaging interaction strategies. To achieve this, this work explores machine learning techniques for leveraging the richness of spoken language and interaction beyond text transcription.

Paul Patras was awarded £70000 by the University of Glasgow to work on his project 'Automatic threat detection and anomaly counteraction in home IoT'. The number of Internet-connected devices is expected to reach 1 trillion by 2035 and a large fraction of these devices will become an integral part of households. This can improve the productivity and quality of life of their users, but also exposes them to new cyber security and privacy risks. Supported by the Centre of Excellence for Sensing and Imaging Systems and Internet of Things Technologies (CENSIS) and Arm, this project will design intelligent algorithms that can detect and counteract cyber threats originating from or targeting Internet of Things (IoT) devices in user homes, without requiring manual intervention. Novel mechanisms will also be devised to allow users to operate safely with potentially compromised gadgets, while ensuring such devices will not damage the networking infrastructure or the operation of other equipment in homes.

Lexi Birch and Mirella Lapata have been listed in top 10 of most influential scholars in Natural Language Processing by AMiner, a free online service for academic social network analysis and mining. As of 2018, the system has collected information on over 136 million researchers, 230 million publication papers, and 368,402 venues. The system has been in operation on the Internet since 2006 and has been visited by nearly 8.32 million independent IP accesses. It provides various search/mining services for publishers, NSFC, and research venues such as ACM/IEEE Transactions, ACM SIGKDD, ACM WSDM, and IEEE ICDM.

<https://www.aminer.cn/ai10/nlp>