THE UNIVERSITY *of* EDINBURGH
informatics

SCHOOL OF INFORMATICS

School Strategy and Plan 2016/17 to 2018/19

Academic Recruitment Priorities

SCHOOL PLAN 2016/17 TO 2018/19

Hiring priorities

Academic appointments

In 2014/15 we recruited to Lecturer posts in Cyber Security and Privacy and in Data Science, in addition to a newly established chair in Cyber Security and Privacy.

More recently, we have recruited an additional Lecturer in Robotics (as an indirect replacement for one of our Chancellor's Fellows who resigned in order to relocate to Switzerland).

Other academic posts currently (December 2015) being advertised or approved are:

- A further Lecturer/Reader in Robotics
- Lecturer/Reader in Databases
- Lecturer/Reader in Machine Learning
- A vacant established chair in Computer Science with a specialism in Algorithms and Complexity
- Lecturer/Reader in Algorithms and Complexity

We will continue to seek to replace staff who leave and to backfill staff for whom a substantial proportion of their salary is covered by Fellowships and similar awards. Because of the rapidly developing nature of our discipline, whilst replacements are likely to be in related fields they will not necessarily be directly like-for-like.

We will seek to replace two further staff who have recently resigned and cases are in preparation for these posts. In addition, we will be bringing forward a case for backfill for seconded posts. These will comprise the following three posts:

- Lecturer/Reader in Computer Vision/Graphics
- Lecturer/Reader in Multimedia Information Retrieval
- Lecturer/Reader in Design Informatics and Visualisation

In addition to the above, there is an urgent need to make appointments in fields related to The Alan Turing Institute, in order that the School is in a position to take advantage of secondment opportunities to the ATI and, at the same time, to ensure resource to continue the normal business of the School. We require four posts in order to provide the necessary capacity and need to proceed with these imminently if we are to take advantage of the current recruiting round and have the capacity in place, in time, in order to allow the University to take advantage of the opportunity represented by the ATI. These posts may be in a range of areas covered by the ATI. The four proposed below are illustrative and not definitive.

Four Lecturer/Reader posts in Data Science supporting TATI secondment: A process for seconding data science academics to the Alan Turing Institute is underway and will intensify in 2016/17. For the University to take full advantage of this opportunity, the School needs to make multiple proleptic appointments so that we can both second data science staff (whether extant or new hires) and at the same time grow our teaching in accord with the planned increase in student numbers. Two types of posts are needed: 1) with a focus on development of principled methodologies for data science, inspired and proven in diverse applications, and 2) with a focus on a key application area, applying and extending data science methods in the process. We give four examples of key posts.

- Methodology: Deep Learning, Time Series and Multistream: Application areas such as speech recognition, computer vision and machine translation are now dominated by use of these methods, with several high-profile examples (e.g., autonomous vehicles). This post supports the intense interest in the formal underpinnings of these methods, with a focus on working with real-world, complex data, and making connections with other Institutes/Schools as well as external bodies. We have run highly successful and oversubscribed national deep learning workshops for the last two years, the most recent associated with a Turing Institute workshop we hosted on the same topic. Industry interest in this area is evident at our annual Data Science Research Day event (cf Google's acquisition of Deep Mind). Student interest is keen: our new deep learning practical course enrolled 140 students (two-thirds MSc) in its first iteration in 2015/16.

- **Methodology**: **Inference and Prediction in Heterogeneous Data**: Applications of machine-learning are increasingly in areas where multiple sources and types of large-scale data need to be integrated (e.g., smart cities, security, learning-analytics). Models of such complex systems usually involve networks of interacting elements, often hierarchically arranged, which poses interesting and challenging theoretical questions; the solution of which will have wide-ranging application. The post will focus on principled methods for addressing the challenges in data integration in complex models, with an outward looking stance to use applications to drive and test theoretical advance, and theoretical advances to enable new applications. This post will build our existing strength in inference into new areas and applications.

- Application: Systems Architecture for Data Science: This is a rapidly growing research area, with high industrial relevance and very strong potential for impact. A unique combination of existing research strength (see REF 2014 impact cases), industrial links (including Intel involvement at ATI) and colocation with leading Data Science research makes the School of Informatics ideally placed to benefit. The computing technologies which underpin Data Science have characteristics and constraints which pose distinctive and contrasting challenges to computer system researchers, unified by the need to combine high performance with energy efficiency. This is true on all scales, from the data centres which provide the massive computing power required by modern data analytics (for example "Cloud Computing"), in the individual devices which pervade the environment, gathering data and requesting services in the form of sensor networks and mobile devices, and in the networking technologies which glue everything together ("Internet of Things"). The market for this technology is huge enough to make the development of devices and systems software customised to its use cases economically viable and desirable. For example, the design of processors and networks specialized to the characteristics of data centre workloads offers substantial energy savings over more traditional mass market designs. Preliminary exploration reveals that an advertised position would attract a very strong field, featuring candidates who already have well developed industrial links.

- Application: Biomedical/Clinical Data Science: this post develops principled machine learning approaches to complex data in areas of fundamental biomedical and clinical relevance. New disruptive technologies are leading to a step change in data production, both at the fundamental and clinical ends of biomedical research, creating a need for novel computational thinking on how to link the existing modelling strengths of Neuroinformatics and systems biology with a data-driven, machine-learning approach. Through the DTC we have built up lasting links with basic biomedical areas at Edinburgh, providing a strong platform for extending collaborative activities. Pre-clinical/clinical research is viewed as a

particular strength of Edinburgh, with opportunities arising from data intensive translational research now conducted in several centres (e.g. MRC-CRM, MRC-CCACE, IGMM, CCBS, CCNS, CIP, Roslin and further groupings such as the Patrick Wild Centre).

All of the above should be cost neutral.

Additional academic recruitment priorities, in order of priority, are as follows:

A1. Lecturer/Reader in Software Engineering: Our Industrial Advisory Board rated software engineering as one of their top priorities for skills in graduates, but SE students make up only 10% of our current UG cohort and our capacity in this area is limited. This post is expected to link to existing and planned strength in data science and cyber-security to help meet this demand. This post would enable us to target gaps in our existing coverage (for example, requirements engineering) as well as emerging areas, such as cloud-based and search-based software engineering. There is significant unsatisfied student demand in this area, therefore a strong business case will be made on the basis of increase student fee income in addition to linkages into existing and developing areas of research.

A2. Lecture/Reader in Human-like Computing (new EPSRC priority): As a leading centre for AI, Informatics is already benefitting enormously from renewed interest in the field, mostly due to significant advances of data-driven methods in robotics, vision, speech and language processing, and machine learning. In areas that investigate how human knowledge and reasoning can be emulated by machines in ways that are similar to and understandable by humans, the School is in danger of losing capacity as only four of its nine academic staff are not nearing (or already beyond) retirement age at the Centre for Intelligent Systems and their Applications (CISA), where such research has been traditionally anchored.

Within AI, it is widely recognised that more human-like AI techniques will be needed to go beyond the solution of isolated intelligence problems toward integrating several of these facilities in systems that exhibit broader common sense reasoning and decision-making abilities, and to enable these systems to explain and justify their decisions toward humans. Research at CISA on topics related to this challenge, such as the integration and evolution of different representations, collective intelligence systems that involve large numbers of humans and machines, computational creativity, and AI-supported human collaboration systems is currently attracting over £5 million in external research funding, including major UK and EU grants.

Recently, EPSRC has declared “Human-Like Computing” a priority area in its strategic plan, and members of CISA are closely involved in the process of shaping the Council’s strategy in the area. It is vital to recruit new staff in this area now to benefit from funding programmes that are expected to be announced over the next six months. Similar opportunities exist in the context of initiatives such as “Industry 4.0” that aim at seamless collaboration between humans and intelligent machines in next-generation manufacturing.

The requested post would enable us to tap into these opportunities at an early stage of this emergent area through increased research capacity, and help address the sustainability issues arising from a lack in personnel in research topics that already produce a healthy income for the School. It can play a significant role in integrating work from other research groups (e.g. in robotics, machine learning, language processing), and help maintain the School's teaching portfolio in foundations of AI. Depending on the profile of the individual recruited, it can also help fill existing gaps in our human-computer interaction and software systems engineering teaching capability.

A3. Senior Data Scientist: There is a large demand from many areas of the university to carry out research based on data science methods. Much of this work can be carried out with standard methods, which makes it unattractive to academic staff working in the methodology of data science; however, it is crucial for progress in these application areas that state-of-the-art methods are used. Our vision to address this demand is to establish a team of practicing data scientists, who are able to advise and train research staff in the various domain areas about appropriate methods, supporting the much-promised data-driven translational work across Schools and Colleges. The SDS will be the first member of this team. The SDS will also organise meetings to share best practice among the research staff working across different areas in data science methods, providing valuable opportunities for cross-fertilisation across areas. The SDS position would be able to have much higher valency than academic staff, and would also be able to pass on non-standard problems to academic staff, providing a very useful triage stage. This position and further team members will be funded from core funding (preferably cross-School) initially, but will be expected to grow funding through inclusion in funding applications which utilise the service. Existence of this service will improve the ability of Informatics and other Schools to attract funding and the best academics, who are looking for this kind of support.

A similar initiative has been undertaken by our competitor institutions. For example, UCL has invested in groups of research software developers, who complement academic strength by providing in-depth development support. The rapid growth and success of these investments is reflected in the Environment ratings secured by these institutions in assessments such as REF.

A4. Lecturer/Reader in Security and Privacy: The recent appointment of a Chair in Cyber Security and Privacy is the beginning of an expansion in this key area, in which we seek to widen and deepen our existing activities within Informatics, as well as with collaborators across the University. The new Chair will lead further expansion into areas that reinforce and/or complement our current strengths. One obvious example of an important area in which we are currently lacking is network security. Most threats to security involve access to network-accessible resources or devices, and network security concerns the prevention and monitoring of such unauthorised access.