The University of Cambridge Department of Computer Science and Technology in 2018

This report describes the Department in September 2018, particularly its research and the wider environment in which it operates.

It was originally prepared for the University's September 2018 Strategic Research Review of the Department, and this version has been very lightly edited, primarily to remove details of the review process and confidential information. The Strategic Research Reviews provide expert external advice to support the development of the University's research environment.

The University of Cambridge

The University

The University comprises six Schools, each of which forms an administrative grouping of Faculties and other institutions. The Schools are: Arts and Humanities, Biological Sciences, Clinical Medicine, Humanities and Social Sciences, Physical Sciences, and Technology. The Schools are each represented on the University's General Board. In addition to carrying out research, the Faculties are responsible for delivering core course content and for the setting and marking of examinations (for undergraduate and postgraduate taught students) and for supervision (of postgraduate research students).

The University also hosts several Strategic Research Initiatives and Networks, and Interdisciplinary Research Centres which bring together research expertise from across the Schools with the aims: to address large-scale multi-disciplinary research challenges; to strengthen research collaborations and knowledge transfer across disciplines; to increase research capacity and profile by providing a platform for large-scale funding applications, recruitments and international research partnerships; and to enhance our ability to influence national and international research, policy and funding agendas.

The Colleges

The University of Cambridge has 31 constituent Colleges, each of which is an independent institution. Members of University Faculties may be invited to hold College Fellowships. Such membership of a College's academic community might include an associated expectation that the post-holder contribute to College life through College-based teaching and other duties which may be administrative and/or pastoral in nature. Academics may also be admitted to a College's fellowship without holding a University (i.e. Faculty) post, though may develop associations with a faculty through teaching and/or research activity.

For students, Colleges provide facilities equivalent to halls of residence, such as accommodation, catering and social space. In addition, they provide integral academic support to undergraduate students, through the co-ordination of small-group teaching sessions called 'supervisions', which supplement core Faculty teaching. This co-ordination is carried out by a Director of Studies, who may or may not be a College Fellow.

The Research Strategy Office

The Research Strategy Office (RSO) supports the work of the Pro-Vice-Chancellor for Research, the Heads of the Schools and the Research Policy Committee in devising and implementing strategies and policies to maintain and increase the University's research funding; shape the University's response to a changing research landscape and the requirements of the funding partners; and enhance Cambridge's standing as a world-leading institution.

The Research Strategy Office team also works with researchers and research groups to help develop collaborative cross-disciplinary research activity, facilitate the development of institutional relationships with major research and industrial funders, support the University's submission to the REF, oversee certain strategically important cross-School grants, and coordinate the University's restricted calls process.

The Research Operations Office

The Research Operations Office (ROO) is a central office that works with departments to negotiate research contracts, check applications for research funding and provide guidance on the management of research projects. The Research Operations Office is organised into School teams, with each Department having a named research grants advisor and contracts manager. Whilst the grant holding authority is centralised with ROO much of the responsibility and process of research funding is distributed into Departments to facilitate greater agility.

Chapter 1: THE DEPARTMENT OF COMPUTER SCIENCE AND TECHNOLOGY - AN OVERVIEW

Founded in the 1930s as the 'Mathematical Laboratory' and subsequently known as the Computer Laboratory (until October 2017), the Department of Computer Science and Technology (CST) is one of the oldest computer science departments in the world. It is located in the purpose-built William Gates Building (completed 2001), on the University's West Cambridge site. Within the organisational structure of the University, it is part of the School of Technology (SoT) together with the Departments of Engineering and Chemical Engineering, the Judge Business School and the Cambridge Institute for Sustainability Leadership.

Leadership and Management

The current Head of Department is Professor Ann Copestake. She has been in position since May 2018, and a member of the Department since October 2000. Prior to her appointment, Professor Andy Hopper served as Head of Department from 2004.

Professor Copestake is supported in the management of the Department by two Deputy Heads of Department, a Departmental Administrator and a Research Facilitator. There is also a senior member of staff who manages the Finance and Research Grants team. A team of computing staff is led by a Senior Computer Officer (a full organisational chart can be found in Appendix B). The senior team play an active role in the overall development of the Department; shaping and implementing policies, responding to new challenges, and maintaining its international reputation. A priority is to ensure that the Department sustains a breadth of active research across computer science, and at the same time, encourages the development of new areas and directions.

The Department aims to provide an open environment where all academics, research staff and students feel they can contribute to debate and discussion around topics affecting research in the department. Whilst the Head of Department, the Deputy Head (Research) and the Research Facilitator have overall responsibility for the strategic development of research, there are weekly staff meetings during term time where relevant research-related matters are discussed. These are complemented by occasional structured meetings, which allow a particular topic or issue to be explored in more detail, and by a yearly offsite. Informal interaction with all staff in the wider Department is facilitated via a termly afternoon tea to which everyone working in the building is invited.

Governance

The governing body for the Department is the Faculty Board, which is responsible for overseeing the work of the Faculty, for ensuring the provision of appropriate research facilities, for overseeing the teaching programme, and for ensuring that the teaching given is of a high standard. The Chairman of the Faculty Board is a University of Cambridge Professor (external to the Department), with the Departmental Administrator acting as its Secretary. There are 22 Faculty Board members including the Head of Department, four Professors, two members

appointed by the University Council, four further members of the department, two co-opted members, the Departmental Secretary and three student members.

Faculty Board members also form the membership of a Degree Committee, which oversees the admission of Graduate Students, the supervision of their work, and the award of degrees and certificates in respect of graduate study. The Degree Committee also considers matters relating to the MPhil in Advanced Computer Science (ACS). Three additional members who have broad understanding of the MPhil are co-opted onto the Degree Committee. These will ordinarily be the Graduate Education Manager, who acts as the Secretary of the Committee, and the Chief Examiner and Course Director for the MPhil.

The Tripos Management Committee (TMC) are responsible for organising all teaching activities related to the undergraduate Computer Science Tripos and for implementing and monitoring relevant policies as determined by the Faculty Board. The Graduate Education and Advanced Taught Courses Committee has responsibility for the ACS and for the PhD learning related activities.

The Head of Department delegates other areas of the work of the Department to the following Committees, listed in alphabetical order, who all report to the Faculty Board:

- Buildings and Environment
- Directors of Studies
- Ethics
- Equality and Diversity (also responsible for Athena SWAN)
- Graduate Students' Forum
- Health & Safety
- Outreach
- Research Staff Forum
- Staff-Student Consultative Forum

Academic Staff

The Department currently has 45 members of permanent academic staff (otherwise known as University Teaching Officers or UTOs). Tragically, Richard Gibbens, who was a member of the Systems Research Group, died during the period in which this report was being compiled. We have not changed the description of the group where his research contribution is discussed. The overall distribution of staff according to title is shown Figure 1.

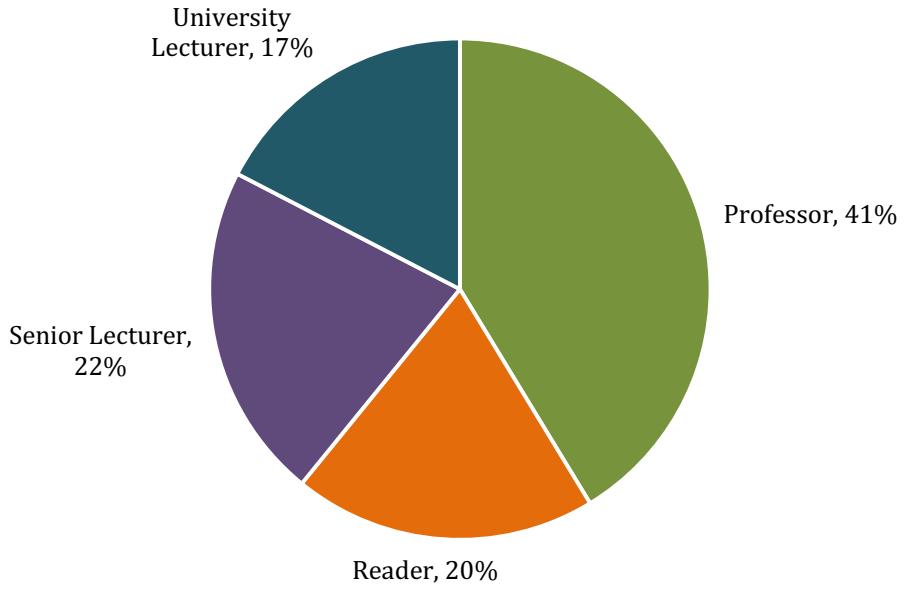

Figure 1: University Teaching Officers (academic staff) title distribution.

The formal approach to appointments at central University level is complex and is changing for 2018/19, so only a brief outline follows. Appointments to UTO positions at the level of Lecturer, Senior Lecturer and Reader are made by an Appointments Committee, which forms the selection and interview panels. The Head of Department and Departmental Administrator sit on all Appointment Committees, with other members proposed by the Head of Department's senior management team and approved by Faculty Board. Normally, at least one member is external to the Department and there may be members external to the University. Professorial appointments are undertaken by a Board of Electors appointed and convened at central University Level, with the Head of Department and Head of School sitting on all such Boards. Boards of Electors usually have three members external to the University.

The University's annual Senior Academic Promotion (SAP) scheme is the process under which academic staff apply for promotion to USL, Reader or Professor. Within this scheme, there is also a route for contract research staff to apply for promotion to the level of Principal Research Associate or Director of Research (subject to availability of funding). The Head of Department is responsible for identifying staff eligible to apply for promotion and discusses their applications with them. The Head of Department is by default a member of the Departmental level SAP Committee, which then makes recommendations to School and University-level committees. The Departmental Administrator is Secretary to the Committee, and the Faculty Board approves its membership. There is a bi-annual scheme for Professorial pay review which allows for application to progress within and between bands 1-4 of Grade 12. Progression is not automatic and requires applicants to give evidence of their sustained and improving contribution to teaching, administration, and research. The next exercise is due to be announced in Michaelmas 2018.

Post-doctoral Researchers

The Department hosts approximately 75 post-doctoral research staff and 25 research assistants (some of whom are also enrolled in our PhD programme), and the numbers continue to grow

in line with the increase in the number of academic staff in the Department. Academic staff are responsible for the recruitment of post-doctoral staff, supported by the Departmental Administrator and a team of HR staff. Subsequent employment matters are managed by the Departmental Administrator, who consults the Head of Department and Human Resources Division where appropriate. Post-doctoral researchers make a substantial contribution to the research in the Department and we work hard to support their ongoing development.

Research Staff have the opportunity to discuss issues relevant to their employment in the Department and broader University via the Research Staff Forum. Meetings occur termly, and are chaired by a member of research staff, with all research groups represented. The Departmental Secretary, the Research Facilitator and a member of academic staff also attend. Formal minutes are included in Faculty Board agendas and the Departmental Secretary also provides informal feedback to the Head of Department regarding requests, ideas and discussion points that arise at these meetings. The Chair of the Committee is invited to be an ex-officio member of Faculty Board and is encouraged to attend in that capacity.

Some post-doctoral staff, especially Senior Research Associates, are involved in University teaching, both lecturing and Masters project supervision. One of our aims is to make this opportunity more evenly available and the related processes more transparent, so that all post-docs who wish to develop their teaching experience are supported to do so. However, we are mindful that post-docs should never feel under any pressure to teach, whilst ensuring that those who do so are properly rewarded and supported.

Research Assistants are generally either researchers in the process of finishing their PhD in our or another University, and likely continuing as post-doctoral researchers (they would gain this status once the PhD is awarded), or they are full time PhD students: the University allows for research assistants to be enrolled as full time PhD students provided the topic of their PhD studies largely overlaps with the project they are employed on. This mechanism has been employed by the Department to increase PhD student funding.

Chapter 2: RESEARCH IN THE DEPARTMENT: OVERVIEW

Research Strategy

The Department's strategy for many years has been to "get good people and let them do what they want". This sounds like a lack of strategy (and perhaps is better described as a philosophy rather than a strategy), but it has been an active commitment on our part, and we firmly believe it has underpinned our considerable and ongoing success in research, as well as in recruiting and retaining staff. However, it actually has complex ramifications, and can be difficult to implement well. Some questions it raises include:

1. What areas do we hire new people in? How do we ensure that good people wish to join the Department?
2. How do we support people (recent recruits and long-established members of staff) to be effective in research? This includes supporting them with funding applications, helping them to build up their own teams of post-docs and PhD students and facilitating interactions with other groups inside and outside the University. It also includes insulating them from unnecessary bureaucratic demands as much as possible.
3. How do we retain staff, especially given that salaries are uncompetitive?
4. How do we anticipate what support people in the Department may need and obtain the funding for those resources?
5. How do we allow people as much time and freedom as possible for research, while also meeting the Department's teaching responsibilities and administrative needs and yet avoiding staff overload?
6. How do we ensure that the Department has the flexibility required to adapt to developments in research, funding, University policy and other external events?
7. How do we communicate this strategy, and its implications, to people outside the Department and people joining it? How do we persuade the University and others that we do have an effective strategy, even if it is expressed in terms that do not seem to be particularly 'strategic' at first glance?

Some of these can be addressed more easily than others. Our approach to departmental structure and recruitment is discussed below. The other questions will be returned to at the end of this report.

Our deliberately light-touch management strategy is guided by a number of underlying considerations. We aim to be excellent in the research that we do and that requires adequate depth. We also aim to maintain breadth: we wish to continue to offer an excellent educational experience from undergraduate level upwards, and we think it important that most courses are taught by people who have research expertise in the topic. Teaching is therefore an important guide to recruitment. Clearly, however, we cannot cover every possible area of research in depth.

Part of the strategy – often not explicitly articulated – is that the Department does not just aim to let people do what they want, but to let them collaborate with whomever they want. This has led to a quite distinctive and extremely successful approach to industrial collaboration, something that the previous Head of Department, Professor Andy Hopper nurtured with great success. The Department under Hopper aimed to minimise barriers to the formation of new

companies while aligning incentives for staff and students, avoiding IP issues, mentoring, being helpful in every possible way, and not picking winners. Furthermore, this has been a cradle-to-grave approach, ranging from undergraduate lectures given by entrepreneurs and investors (often alumni) to the maintenance of an industrial business club beyond the Department. Increasingly, this also involves various initiatives for interdisciplinary collaboration, which are discussed further in subsequent sections of this document. The results speak for themselves, with 270 companies formed by staff and students to date,[1] of which 50% are active with turnover of $1Bn, and 18% were sold for over $40Bn.

As stated above, the Department has consistently maintained an extremely lightweight approach to IP, with an emphasis on disruptive innovation and being first-to-market, rather than on incremental advances to existing technologies. Most academic staff will not sign NDAs, and there is a very strong tradition of releasing software as open source. We are also making moves into open source hardware development, not least with Mullins' involvement in the lowRISC project[2]. Many staff have experience of successful commercialisation and/or company founding, and willingly provide support or mentorship to those interested in doing so for the first time. Although none of this is mandatory, with individual academics remaining free to choose the most appropriate route 'out' for their particular project(s), we are firmly of the opinion that our overall approach has been central to our ongoing academic and commercial success, not least in enabling us to attract and retain entrepreneurially-minded staff, and in enabling us to provide the most up-to-date and high quality industrial exposure for our students.

The Department's organisational structure is primarily aimed at ensuring that communication channels are open and transparent, resources are appropriately directed, and that individuals (whether academic or support staff) do not end up overloaded. Research is loosely organised into eight groups, each of which have, on average, six academic staff, though there is considerable variation. Individual groups do not have rigid boundaries and many staff belong to more than one group. The group structure is thus a means of facilitating interaction and informal collaboration between those working in cognate areas, and of enabling potential funders, collaborators and students to identify more easily specific areas of interest.

Academic Staffing

Recruitment of permanent academic staff is the primary means by which the Department expresses its strategy insofar as it relates to areas of research. The Head of Department makes the decision as to which positions we seek to fill, after consultation with the Deputy Heads and other members of the Department. The case for support is then taken to the School's Needs Committee for discussion.

The overall goal is to recruit excellent people with relatively broad interests and expertise who will have the flexibility to adapt – we are, after all, recruiting people who may well be in the Department for 40 years. The ability to collaborate (within academia and with industry or other non-academic organisations) is extremely important and we now put considerable emphasis on collaborations external to the Department. A high proportion of our recent appointments have been interdisciplinary in some respect, and we expect that trend to continue. We naturally need to fill some gaps which arise when staff leave or retire, but this cannot be the main driver, given

1 www.cst.cam.ac.uk/ring/halloffame

2 www.lowrisc.org/

the rapidity with which research in CST moves. The choice of areas to expand into is driven partly by our perception of these areas' potential contribution to the field, but also by consideration of how the appointment could fit with the existing groups within the Department. Recruiting people who are at the intersection of two or more of the existing groups gives them a better basis for developing their careers than those who do not have this grounding. Indeed, we have recently taken the decision that we will seek to recruit at more senior levels (Reader or Professor), when it would be relatively difficult to make strong connections to existing groups.

The staff who have been appointed over the last three years are: Nada Amin (adding to our Formal Methods and Programming group), Paula Buttery (Computational Linguistics/Natural Language Processing, transferred from the Department of Theoretical and Applied Linguistics), Hatice Gunes (Affective Computing), Timothy Jones (Hardware), Eva Kalyvianaki (Distributed Systems), Neelakantan Krishnaswami (Programming Language Semantics and Theory), Rafal Mantiuk (Graphics), Amanda Prorok (expanding the Department's expertise into the area of Mobile Robotics), and Damon Wischik (Mathematical Modelling and Data Science). Some of these appointments replaced retiring members of staff, and some were proleptic,[3] while others represent a real expansion in numbers.

Three new members of staff will join us in October 2018. Two will join the NLP group (Ryan Cotterell, previously a PhD student at John Hopkins University and Andreas Vlachos, previously a Lecturer at the University of Sheffield). These appointments reflect the importance of NLP in its own right and also its close relationship with machine learning. They are additionally motivated by current and future vacancies: Stephen Clark has joined DeepMind (he formally leaves the Department in October 2018 but has primarily been working at DeepMind since October 2016), and two members of the group (Ann Copestake and Ted Briscoe) will retire in the future. The third new appointment is Alice Hutchings who has been appointed to the Security group (she is currently a post-doctoral researcher in the group, specializing in cybercrime). The appointment expands the very successful group, and underlines its importance to the Department. It was advertised as a general Security vacancy, but Dr Hutchings' background in Criminology fits well with the Department's increased emphasis on interdisciplinarity. These appointments will take us to 48 members of academic staff.

We expect to recruit four members of staff in the coming academic year (aiming for appointment in October 2019). We will soon advertise a Chair in Machine Learning, which has been endowed by DeepMind. The Department currently has far too little capacity in Machine Learning to meet demands from students and Industry and we also see huge potential for interdisciplinary collaboration in this area. Given the importance of this area, and our current limited expertise in Machine Learning, the ability to recruit at Professorial level is very welcome. We will also advertise a Chair in Computer Science (Robert Sansom Chair, left vacant by Prof Ian Leslie's appointment as Director of University Information Systems). A Readership in Environmental Data Science will be an expansion into an area of great importance to the University, as well as more externally. Ideally, as well as being interdisciplinary in its own right, it will enhance the potential for other members of the Department to collaborate with scientists across the University. The fourth vacancy is a

[3] Although we use the term `proleptic' in this document, this is a descriptive convenience rather than a strict statement since we do not have a fixed headcount as such.

lectureship in Computer Graphics, with a focus on 3D Modelling and Rendering. This is complementary to the research carried out by Rafal Mantiuk into Display Technologies and Perception. It meets a current teaching need (an Affiliated Lecturer has been teaching Graphics courses) and relates to interest from students and from Industry (particularly computer gaming).

Research Staff

In addition to our permanent academic staff, we have a number of prestigious fellowship holders within the Department, and details of these are given within the relevant Group sections. We make a distinction between Senior Research Fellows (SRFs), who may hold grants as a PI and supervise PhD students, and holders of more junior positions, such a college Junior Research Fellows (JRFs). College JRFs are hosted by the Department rather than being formally employed. All SRFs in the Department are regarded as full members of academic staff, helping to lead and direct the intellectual life of the research group of which they are a part, and being encouraged to apply for further research grants in their own right.

In terms of professional development, the University is a signatory to the UK Concordat to support the Career Development of Researchers, and Senior Fellowship holders are offered the opportunity to supervise relevant students and have access to the University's comprehensive professional development programme. As with newly appointed members of academic staff, they are assigned an existing member of permanent academic staff as their mentor. Mentors are able to monitor and advise on their programme of research through routine informal contact and more formal annual appraisal. JRFs may be involved in the supervision of PhD and Masters students. Research Fellows are not required to teach, but many choose to do so, and SRFs sometimes fill administrative positions in the Department.

Unlike the practice in many universities, the Department cannot make automatic offers of permanent posts to applicants who are successful in obtaining Senior Research Fellowships. This has sometimes led to us losing potential applicants for fellowships or fellowship holders to departments in other universities. So far this has not been regarded as a serious problem but it may require a rethink at University/School level, should more fellowships follow the recent UKRI pattern and require confirmation from the host institution that they will make applicants a firm offer of a permanent position at the end of their fellowship. Fellows are encouraged to apply for academic departmental posts in competition with external candidates. There is usually an option in the advertisement for a position that allows candidates from any area of computer science to apply, so fellowship applicants do not always have to wait for a vacancy in their area, although, if an offer is made under such circumstances, the usual process for approval of positions (the Needs Committee) has to be undertaken. If applicants are successful, they typically retain their research fellowship until it ends and then take up their academic post. In the last four years, one of our new academic appointees was previously a Research Fellow in the Department (Dr T. Jones).

Research Support Functions

Computing support within the Department is provided by a team of Computer Officers. Together they provide support for most aspects of research provision ranging from purchase through to installation and maintenance of servers, storage and networking. The team is geared

towards supporting research users and most of its staff have the strong technical background necessary to support this function in a diverse research environment.

Administrative support for funding applications and industrial sponsorship is provided within the Department by a Research Grants Administrator (provides costings and cross-checks funder regulations), the Finance Manager (maintains financial oversight of all applications and active research grants), the Departmental Administrator (provides advice and support on recruitment processes and other HR matters) and the Research Facilitator (provides assistance with all non-financial aspects of grant preparation).

The Department also has access to a specialist contracts negotiation team from the University's Research Operations Office, who undertake negotiations with commercial funders on our behalf. The University's Development and Alumni Relations office (CUDAR) provides support for the negotiation and management of philanthropic donations.

Research Resources

The Department has its own experimental data centre, a cluster of 80 computers interconnected by a highly reconfigurable 10/40Gb Ethernet network. Users may customise or replace the installed OS, and have the ability to intercept, capture and monitor network flows, and monitor power usage in real time. We also receive an annual resource allocation equivalent to $20K from the Microsoft Azure platform, which is primarily used to support undergraduate and postgraduate student use. Several academic staff also receive separate allocations from Azure, for their own use or for their students. Individual groups have additional specialised research equipment pertinent to them, usually purchased through research grants. The University Information Services' High Performance Computing clusters are also used.

Research Funding

As stated in the above section on Research Strategy, the Department gives academic staff absolute freedom in terms of their research direction and the funding sources that they approach and our funding portfolio is highly diverse as a consequence. Figure 2 shows the breakdown of research funding by funder type over the last five years per Principal Investigators, and it is worth elaborating further on a couple of points. Regarding European Commission funding, over the last five years, eight of our staff have been successful in applications to the European Research Council (ERC), comprising four Starting Grants (Clark, Mullins, Sauerwald, Stajano), one Consolidator Grant (Mantiuk), and three Advanced Grants (Paulson, Sewell, Winskel). We would obviously hope to maintain this level of success going forwards, but it remains to be seen if UK-based applicants remain eligible to apply to the ERC following Brexit. On a more positive note, the proportion of 'overseas government/educational' funding saw a considerable increase in the past full year which we have figures for, a trend which we expect to continue, given that much of this is due to a strong relationship between the Computer Architecture Group (Moore and Watson in particular) and SRI and their success in obtaining significant DARPA funding (see Chapter 4 for more information). On a similarly positive note, the volume of philanthropic funding has been increasing over the last three to four full financial years; this funding is of a higher value than its monetary value per se, as it generally enables flexibility in terms of being able to spend it where there is most need at a particular point in time – e.g. providing further (or bridging) funding for talented post-docs coming to the end of

their project, enabling us to fund small-scale exploratory work that would be too risky for a 'normal' grant application and supporting a not insignificant number of PhD students. We note that Figure 2 does not include funding coming to the Department when investigators are co-investigators of grants or where the principal investigator is in a different institution or department, and it is therefore a lower bound on the grant income.

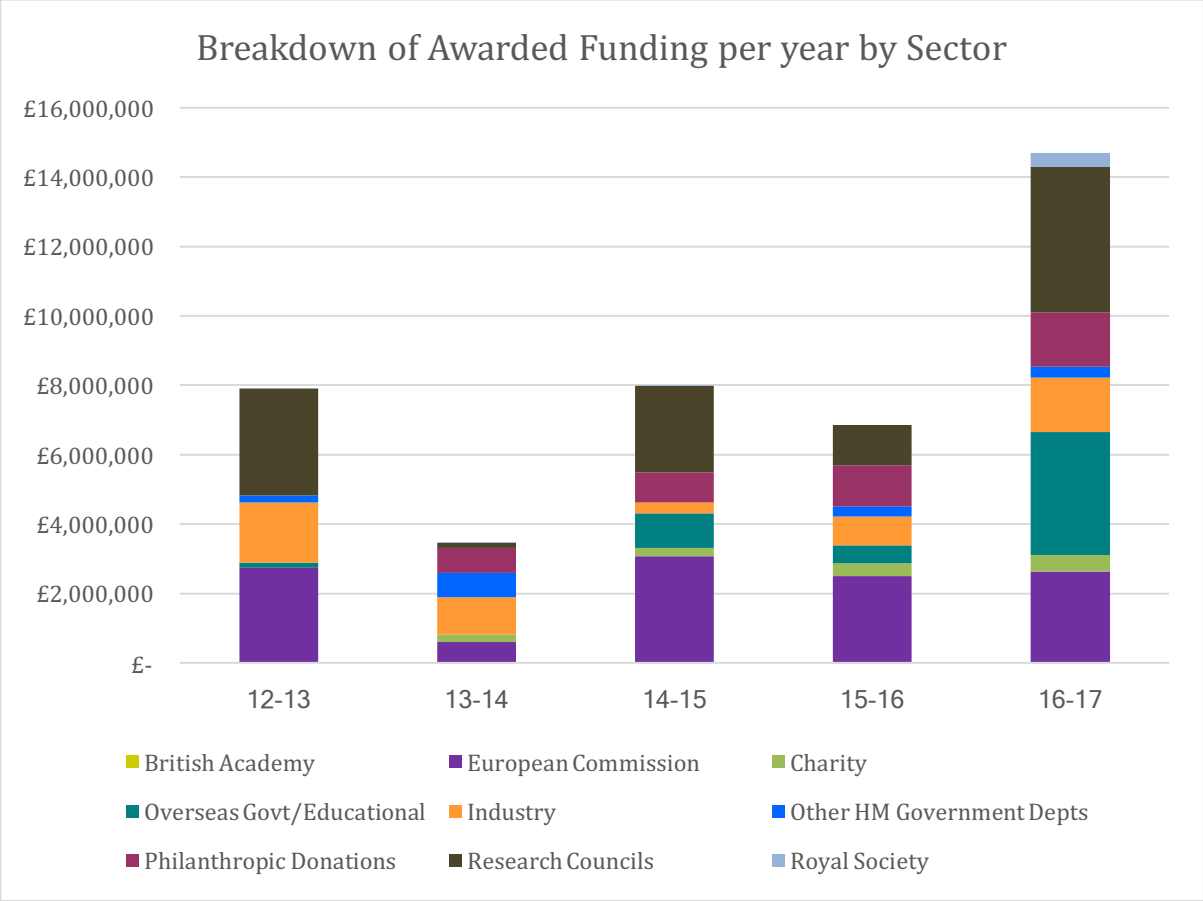

Figure 2: Breakdown of funding sources by year.

In terms of support from our industrial partners, recent years have seen the receipt of a number of large grants to support major initiatives in the Department. Illustrative highlights include:

- **OCaml Labs**[4] (2011), supported by a series of donations now in excess of £5.3M from Jane Street Trading with an additional £2.5M from other partners. The Director is Dr Anil Madhavapeddy and the centre concentrates on pushing OCaml and functional programming forward as a platform which can be used widely both industry and academia.

- **The Alta Institute**[5] (2013). The Institute for Automated Language Teaching and Assessment (ALTA) is funded by Cambridge English Language Assessment to conduct research into improving linguistic proficiency of learners and into automated forms of assessment of language performance. Professor Ted Briscoe from the Department is the Director of the Institute and coordinates the academic programme.

4 http://ocamllabs.io/about/

5 http://alta.cambridgeenglish.org/

- **The Centre for Mobile Wearable Systems and Augmented Intelligence**[6] (2018), funded by a £1.75M donation from Nokia Bell Labs over five years. The Centre will explore how to augment and improve the human experience in a digitally connected world, with areas of focus including predictive and personalised medicine; digital, physical, mental and social well-being; and sensory human communication experiences beyond visual and audio. The Centre will be co-directed by Professor Cecilia Mascolo and Dr Alastair Beresford and support roughly one post-doc and nine PhD studentships.

Almost all of our UK Research Council funding comes from Engineering and Physical Science Research Council (EPSRC), with a recent addition of Medical Research Council (MRC) funding, mainly secured by Professor Lio' and Professor Mascolo. Figure 3 shows the volume of EPSRC funding awarded per year over the last five full years. It shows when funds were awarded and does not reflect the duration of the grant. Where grants are active, awarded funds to date are captured. So we would therefore expect a 'successful' year (e.g. 2012-13 below) to be followed by a relatively unsuccessful year (according to the data available at least), as academics who have gotten funding are therefore not likely to be applying for further EPSRC funding for another two to four years.

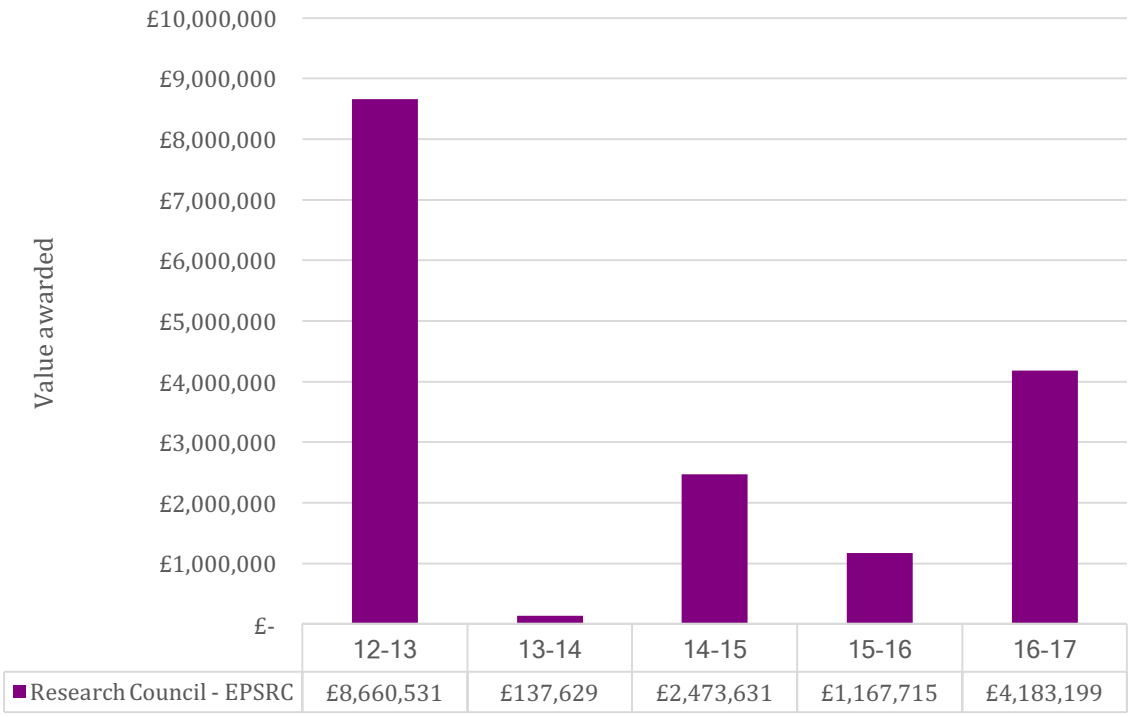

Figure 3: EPSRC funding awarded 2012/13 – 2016/17.

Whilst EPSRC grants are obviously a very important part of our overall funding portfolio, any internal strategic decision to increase such funding would be subject to considerable funder-driven constraints. In recent years, EPSRC (in common with the other UK Research Councils) has allocated an increasing proportion of its overall budget to funding calls that either have a particular focus or for single large centres in a particular area, at the same time as individual

6 https://mobicentre.cst.cam.ac.uk

universities have been increasing staff numbers and therefore increasing the pool of eligible applicants for such funding. Whilst we would obviously hope to remain highly competitive in terms of EPSRC funding applications, certain factors pertaining to success rates are therefore outside our control and demonstrate the importance of maintaining (and increasing) diversity within our funding portfolio in order to sustain our absolute research excellence.

Chapter 3: RESEARCH AREAS

Overview of Research Groups

The Department of Computer Science and Technology has over 250 research workers: academic staff, research associates, and PhD students. Figure 4 depicts the breakdown of the various roles across the research workers.

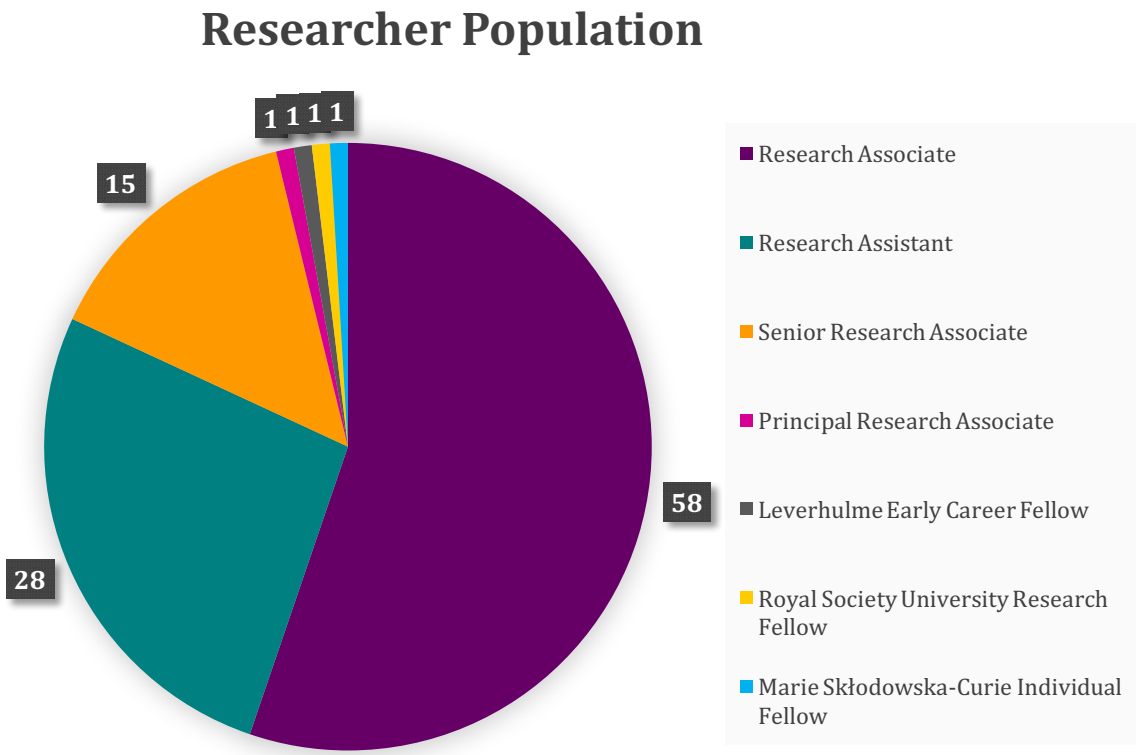

Figure 4: Breakdown of Research Workers' Roles 2017/18.

Figure 5 shows the temporal trends of the Department in terms of staff. The increase in academic staff is partly responsible for the increase in research staff.[7]

[7] There is a slight discrepancy in the number of academics mentioned earlier in the text and in Figure 5 due to data collection at different times and capturing differently retirements/appointments.

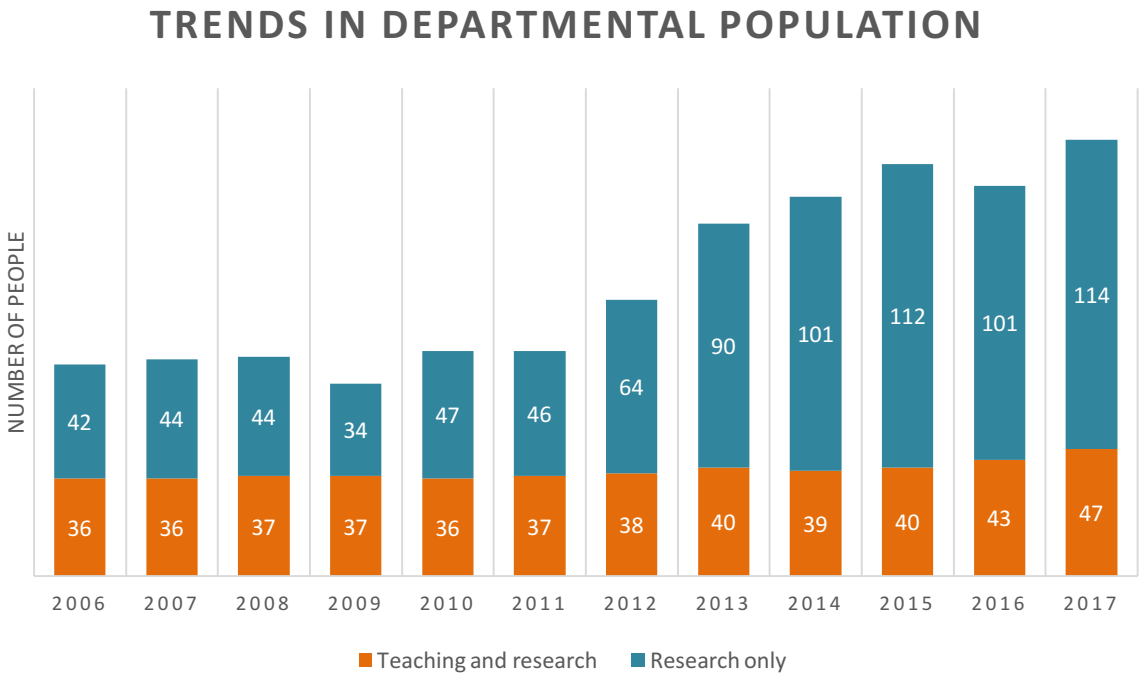

Figure 5: Temporal Trend in numbers of Academic and Research Staff.

Research is carried out across a broad range of subjects within Computer Science and is conducted primarily within small research groups:

- Artificial Intelligence Group
- Digital Technology Group
- Computer Architecture Group (Rainbow)
- Graphics and Interaction Research Group
- Natural Language Processing and Information Processing Group
- Programming, Logic, and Semantics Group
- Security Group
- Systems Research Group

We now describe the research and directions of each group. A full list of group membership and member descriptions can be found in Appendix A.

Artificial Intelligence Group

Research by the Artificial Intelligence Group is multi-disciplinary, spanning genomics and bio-informatics, machine learning, stochastic algorithms, game theory, automated theorem proving, computer vision, and human-like computation. A unifying theme is to understand multi-scale pattern recognition problems, seeking powerful statistical algorithms for modelling and solving them, and for learning from data.

Research Highlights
The AI Group seeks to continue to find synergies amongst ideas based in statistics, machine learning and reasoning, cognitive science, biology and engineering, and to develop practical and beneficial applications from them. For example, John Daugman's research has led to the invention and employment of automatic iris recognition. As of early 2018, at least 1.5 billion persons worldwide have had their iris patterns mathematically encoded using the Daugman algorithms for enrolment in national ID or entitlements programmes.

Research developing AI techniques for human-like computing has been led by Mateja Jamnik. She recently served as a Specialist Adviser to the House of Lords Select Committee on Artificial Intelligence, helping the UK government in policy direction, priority and focus in relation to the impact of AI on society. Recently Jamnik has started to apply AI and reasoning techniques to medical data to advance personalised cancer medicine.

Pietro Lio's research uses bioinformatics, computational biology models and machine learning to integrate various types of data across different spatial and temporal scales of biological complexity. By integrating different layers of evidence, predictive models will improve the accuracy of diagnosis of complex diseases in the presence of other chronic and acute conditions and will be able to identify effective markers for disease trajectory and suggest the possible composition of treatments.

Pietro Lio' and Mateja Jamnik recently started a new project on cancer data integration via machine learning and logic reasoning. This forms part of a large project on Integrated Cancer Medicine at The Mark Foundation Institute for Integrated Cancer Medicine at Cambridge University and is a collaboration with Oncology, Imaging, Mathematics and Physics. The project will aim to develop and use novel AI methods, in particular machine learning algorithms and also logical approaches to integrate diverse types of patient data (e.g., genomic, molecular, imaging, clinical) in order to provide better and more personalised diagnostic and treatment plans. In particular, we will develop techniques from supervised and unsupervised machine learning methods, including random forests and various flavours of deep learning (for example, graphs, attention, agents based, autoencoders) and combine them with network science (multilayers) and biomedical knowledge to integrate heterogeneous data for individual breast cancer patients, such as their genomic, cell and clinical data. In parallel, we propose to use these predictive models to extract a generative logic model with reasoning rules that provides causal explanations of the decisions made by the model. These rules will be used to build medical decision support systems for personalised patient prognosis, therapy and drug optimisation.

Altogether the AI group hosts 5 Academics, 9 Research Associates and 16 PhD students.

Digital Technology Group

The DTG has a very wide area of expertise, its research scope ranging from system design, analysis and implementation at the physical level to development of novel devices and applications – an experience gained from a long history of applied research. The group also studies control and planning algorithms for autonomous systems, with theoretical guarantees on their predicted performance. As a distinctive feature, most projects within the group are also implemented and tested in real-world environments. Approaching computing problems by defining a pioneering and non-standard vision allows new technical solutions to emerge which

have application way beyond the original framework. This has happened with assured computing where our provenance research produced early implementations of blockchain (broadly defined) which give engineering choices between energy-proportional computation and unbounded computation when dealing with falsification and circumvention in computer systems. In addition to continuing work on various classification tasks, the group will be extending work on machine learning to the design and optimisation of wireless networks in order to improve the capacity and throughput, particularly addressing aspects such as resource allocation, quality of service and physical layer throughput. Doing so is vital given the ever increasing demands being placed upon wireless networks. Thus, the underlying aim will be to address large-scale unsupervised learning problems, where longer-term open questions will be raised about appropriate objectives for learning good representations for such large scale problems.

Research Highlights

Examples of the group's research include the Trve Data project[8], run by Dr Beresford, which is developing new and enhanced security and privacy features for collaborative applications. For example, we are devising and implementing the fundamental distributed computing algorithms needed to replace systems such as Google Docs, Evernote and Wunderlist with solutions which do not require you to trust service providers with the contents of your shared documents, to-do lists, calendar appointments or notes. Trve Data is funded by The Boeing Company.

The Device Analyzer[9] is another project managed by Dr Beresford and Dr Rice. This project has collected statistical usage data from over 30,000 Android mobile devices across the world. Device Analyzer pre-processes data on the mobile device to remove direct personal identifiers and reduce the privacy risk of sharing the data with the university. If our study participants agree, we share a subset of the data with over 80 partner research groups. We also use the data ourselves for research, for example, by looking at the state of Android security. Device Analyzer is funded by donation from Google.

The Isaac Platform[10] is a project that uses recent developments in web technology and computer-based educational methods to improve physics and maths teaching in English schools, working in partnership with teachers and educators. Dr Beresford and Dr Rice lead the technical team and work in close collaboration with Dr Lisa Jardine-Wright who manages the overall project as well as the content and outreach teams in the Cavendish Laboratory. The platform has over 120,000 registered users who make tens of thousands of question attempts on our platform every day and is funded by the Department for Education.

The group will continue with large-scale system deployments along the lines of Device Analyzer and The Isaac Platform. The aim of such systems is not only to understand behaviour in the chosen domains, but also improve our knowledge about how such systems are best designed, implemented and maintained. Such projects also feed into our teaching programme, ensuring that our undergraduate and graduate course reflect contemporary practice. Several of our projects raise significant privacy and security issues and this is another area which we intend to explore further in the coming years.

---

8 www.trvedata.org

9 https://deviceanalyzer.cl.cam.ac.uk

10 https://isaacphysics.org

The CamFort project[11] has developed an open-source software verification and refactoring tool for Fortran. The project, run by Dr Rice, works closely with scientists and the work has been presented nationally and overseas at Fortran Modernisation Workshops. CamFort will shortly be integrated into the software quality pipeline at the MetOffice for the Unified Weather Model, which is used for forecasting both in the UK and under license in many other countries. CamFort is funded by the EPSRC. Dr Rice is intending to continue with the CamFort project with the goal of developing new techniques to help scientists who are writing computational models. The emerging area of combining machine learning techniques with program analysis is a particularly interesting avenue for investigation. Rice will be leading a course on Machine Learning for Programming on the MPhil in Advanced Computer Science next year.

Research centred on human movement has been led by Dr Harle, with a focus on the use of consumer smart devices, including smartphones and watches. Recent projects have considered movement at a coarse level, estimating where someone is and where they go, particularly indoors. Most recently his work in this area has focused on developing automated techniques for the construction radio maps from crowd-sourced smart device data. In parallel with this, he has been investigating the use of smart devices to monitor more fine-grained movements, which are considering how people move. Initially applied to unobtrusive limb-tracking in sports, the most recent project involves partnering with Addenbrooke's Hospital and Google to track the progression of degenerative disease using only the signals captures from personal phones and watches.

Further wide-reaching research includes Professor Hopper and Dr Rice's work on how digital technologies can support a more sustainable planet. They developed a framework consisting of four components: "green computing", "computing for green", "assured computing", and "wealth in cyberspace". This perspective has been aimed at industry and a broad audience rather than solely academic computer science. Over the decade since it was published quite a number of the green computing goals have been achieved and industry is well on the way with the remaining parts of the vision. However there are a number of problematic area as well, not least to do with the dilemmas of privacy and surveillance.

Research contributions rooted in the domain of multi-robot systems are led by Dr Amanda Prorok. Her group's goals are to develop inter-robot coordination algorithms for robotic systems that operate in uncertain and adversarial settings. A recent research highlight addressed the problem of robot assignment on transport graphs, where origin-to-goal travel times are uncertain. This research proposed a novel concept that allows to curb the expected performance loss: robot redundancy. Although they show that the nominal redundant assignment problem is strongly NP-hard, they have developed a near-optimal, efficient assignment algorithm based on supermodular optimization. The novel assignment algorithm proved to significantly reduce average robot travel times in cases where uncertainty is present in robot origin locations, robot goal locations, as well as robot travel times. This result is transformative for a number of applications, such as multi-vehicle systems, automated transport, and robot swarms. Prorok will continue her line of work in coordination algorithms for robotic systems that operate in uncertain and adversarial settings. In the near-term, she and her collaborators will extend some preliminary results on using Graph Neural Networks (GNNs) for learning multi-robot coordination; their focus will be on relaxing some of the underlying assumptions that GNNs build on, to make them more amenable to systems that operate in real environments (e.g., lossy and asynchronous communications). They will also focus on hybrid control mechanisms that

[11] http://camfort.github.io

are capable of sliding between distinct operational modes (cooperative and centralised planning, and independent decentralised planning), as a function of real-time factors that define the level of cooperation in multi-robot or multi-agent systems.

The group includes 6 Academics, 10 Research Associates, 2 Research Assistants and 13 PhD students.

Computer Architecture Group

The Computer Architecture research group works on a wide range of topics in Computer Architecture, spanning low-level circuit design through to complete computer systems. Current work spans a range of activities including alternative massively parallel architectures (e.g. Loki and POETS projects), secure processors (e.g. CHERI processor in the CTSRD project), automatic parallelisation and vectorisation (e.g. HELIX), binary translation, open-source system-on-chip design (lowRISC project), the design of accelerators for machine learning and network design (e.g. NetFPGA).

Research Highlights
Dr Mullins has begun a three-year research collaboration with Samsung to develop scalable accelerators for neural networks capable of on-device learning. The project is exploring techniques to compress the model itself and techniques to ensure that the resulting model can be executed efficiently in hardware. LowRISC[12] is a not-for-profit organisation and research project. The goal is to create a fully open-source, Linux capable, RISC-V based SoC (System-on-Chip). Mullins is a founder and Director and the project is supported by Google and other industrial and private donors. The project also leads the RISC LLVM compiler effort. Mullins will be focusing on the development of the LowRISC project which is set to grow over the next year with significant new industry funding. The goal is to help further develop the open-source hardware ecosystem to the point where numerous commercial companies are producing volume silicon using open-source designs. This will be achieved with help from both academic and commercial partners.

Mullins is also about to complete a five-year project funded by an ERC Starting Grant. The project has explored the design of a flexible tiled many-core processor that incorporates support for message-passing at the instruction-set level. A 128-core prototype has recently been sent for fabrication in TSMC's 40nm process (June 2018).

As part of his EPSRC Early Career Fellowship (Sept 2013 - Sept 2018), Tim Jones has been developing new strategies for data prefetching, especially focusing on graph, database and HPC application areas. His group have developed hardware and software schemes to bring data into the processor's caches just before being needed, so as to hide the costly trip to main memory to fetch it. An automatic compiler pass (published at CGO 2017) adds support for common patterns that target existing processors; a novel heterogeneous architecture (ASPLOS 2018) uses an event-driven programming model to guide an array of highly-efficient microcontroller-sized cores to perform the prefetching. HELIX is an automatic parallelising compiler for generic sequential applications, originally developed by Jones and colleagues at Harvard University. First published at CGO 2012, the research considered pure software techniques as well as hardware extensions to support low-cost communication between processor cores

[12] www.lowrisc.org

(ISCA 2014, CACM 2017). Lessons learnt from this prototype fed into a PhD dissertation on automatic parallelisation of application binaries (Zhou, viva Nov 2017) and an EPSRC grant on the same topic (July 2017 - Dec 2020).

Further research highlights include the CHERI secure processor work (Watson and Moore in collaboration with SRI International), funded largely by DARPA, starting in 2010 with the CTSRD project that originally ran for four years but, through various extensions, ran for nearly eight years. More recently they have received additional contracts to (a) transition CHERI to industry (CSST project); (b) undertake formal verification of key security properties (CIFV project, Sewell, Moore, Watson); and most recently (c) to explore complete system-on-chip security (ECATS project, with ARM and SRI as partners). Publications cover a number of communities from architecture (ISCA 2014, IEEE Micro 2016, ICCD 2017) through languages (ASPLOS 2015 and 2017, PLDI 2016), security (IEEE Symposium on Security and Privacy 2015, ACM CCS 2015), and formal methods (MEMOCODE 2015, FMCAD 2016). So far this has required 70+ person years of effort. Through the EPSRC, Watson and Moore have also begun the IOSEC project in 2018, exploring I/O security with a particular focus on exploits from malicious pluggable peripherals (e.g. over USB-C that also wraps Thunderbolt and PCIe).

The POETS project (Moore & Jones) is an EPSRC programme grant, started in 2016 in collaboration with Southampton, Imperial and Newcastle, to explore massively parallel and highly interconnected compute using an event triggered model. Cambridge is providing the hardware architecture (on many FPGAs) that supports this effort.

Simon Moore's future plans are to focus on security and rigorous engineering of computer architectures that underpin trustworthy technology. In particular, to commercialise the CHERI processor architecture by working with major players; to continue work with Robert Watson across the hardware and software security architectures; and work with Peter Sewell on verifying key security properties of the whole system. We are also applying for funding to explore vulnerabilities due to misuse of speculative execution mechanisms (e.g. like the Spectre attacks). Finally, Moore intends to continue work on massively parallel compute with a particular focus on communication as a first class design constraint.

The groups includes 3 Academics, 15 Research Associates and 13 PhD students.

Graphics and Interaction Research Group

The group has undertaken research in computer graphics since the mid-1960s. Over the decades, the focus has moved from considerations of display hardware and hardware architecture (1970s), through rendering algorithms (1980s), 3D display technologies (1990s) and modelling (2000s), to our current work on new display technologies (HDR, VR, multi-focal) and imaging. Graphics research is undertaken primarily by PhD students and post-docs working under the supervision of Rafał Mantiuk. The group has undertaken research in the field of human-computer interaction for many years. Much of that research has focused on novel interface techniques that put computing power into the hands of a wider range of users under the direction of Peter Robinson. This work has been complemented by research into applications of HCI theory, including cognitive science, social science and design practice, under the direction of Alan Blackwell. With rapid advances in key computing technologies and the heightened user expectation of computers, the development of socially and emotionally adept technologies is becoming a necessity. HCI research increasingly involves collaboration

with other parts of the university. The Crucible network is jointly directed by Alan Blackwell, and by David Good in the Department of Psychology. Crucible has undertaken over 180 interdisciplinary projects since 2000, and now leads the Cambridge agenda in Global Challenges.

Research Highlights

Two large grants have supported research in affective computing: ASC-Inclusion was a €1.9m collaboration funded by the EC to investigate the use of a 'serious game' as an intervention for children with high-functioning autism. The group's main contribution was a system to provide feedback from facial expressions. The face tracker developed as part of this has been made freely available and is widely used around the world. The educational music-programming language Sonic Pi, developed by Sam Aaron and integrating the group's research agendas in programming languages and music, has become the world's most widely used live coding language, with over a million current users. A series of other projects funded by sponsors including Boeing Research and Technology, BT, and the Health Foundation have resulted in the creation of interactive intelligent tools that demonstrate potential of mixed-initiative interaction for intelligent data analytics. This work has resulted in frequent invitations to contribute to the work of the Alan Turing Institute.

Alan Blackwell's work in human factors of programming languages has been recognised with both 10-year and 20-year most influential paper awards at the IEEE conference on Visual Languages and Human-Centric Computing. His HCI-driven work is increasingly influential on the mainstream programming language research community, with invited talks at conferences such as POPL and SPLASH. Blackwell is currently investigating intelligent user interfaces, including critical assessment of the economic and social implications of machine learning algorithms with colleagues in philosophy, politics, and science and technology studies. Working with Cambridge Global Challenges and partners across sub-Sarahan Africa, he is planning a programme of comparative education, entrepreneurship and policy research for global AI.

Over the last five years, Hatice Gunes has been part of the £2M EPSRC project 'Being There – Humans and Robots in Public Spaces'[13] (2013 – 2017) to investigate with Universities of Exeter, Oxford, Bath, and the Bristol Robotics Laboratory how to achieve greater social integration in public spaces, and how to increase access to public spaces in robot proxy forms. Dr Gunes' publications, live demos and the dataset stemming from this project have led to numerous invited talks/media coverage, new grants, and a Special Issue in *Frontiers in Robotics and AI* in the form of a Research Topic on 'Affective and Social Signals for HRI' which has been shortlisted for this year's Frontiers Spotlight Award. Gunes was also an academic partner in the £562k Innovate UK / Technology Strategy Board project 'Sensing Feeling'[14] (2016-2018) that aimed to enhance user experience in retail. The project developed a compact and portable hardware sensor with capabilities of deep learning-based automatic emotion sensing and interpretation based on real-time facial expression sensing and monitoring, data transmission, and aggregation. The product has been patented and is currently in the process of being commercialised by the industrial partner.

Dr Gunes' current research vision is to embrace the challenges present in the area of health and to empower the lives of people through technology by continuing her research with applications

[13] http://being-there.org.uk

[14] http://sensingfeeling.com

to social robotics and wellbeing. She has recently been awarded the prestigious Early Career Fellowship (2019-2024) by the Engineering and Physical Sciences Research Council UK (EPSRC) to investigate adaptive robotic emotional intelligence for wellbeing (£1M FEC). She is also taking part in the Small Business Research Initiative (SBRI) project led by industrial partners focusing on social training in virtual environments using home-based emotion-sensing. In July 2018, Dr Gunes has been named a Turing Fellow with the Alan Turing Institute, commencing 1 April 2019. She will be collaborating with various Turing Interest Groups on how to devise novel approaches to advance the interaction and adaptation capabilities of autonomous robots, and how to utilise the adaptive intelligent robotic frameworks for fostering human wellbeing.

Rafał Mantiuk's work on perceptual display algorithms has helped establish a start-up company (IRYStec). The company has now multi-million evaluation and was successful in integrating their solutions into automotive infotainment systems (sun-light and dim-display compensation). Mantiuk was awarded an ERC Consolidator grant (€1.8M) to work on perceptual coding for future display technologies (VR, AR, 3DS, multi-focal, HDR). The project aims at finding efficient, perception-limited representations for visual content intended for the display technologies that will support high dynamic range, changing view-point, binocular and focal depth cues. He has started collaboration with Human Vision scientists (Berkeley, Liverpool, Oxford, Aberdeen, Southampton), aimed at developing models and collecting data using prototype display devices of extended capabilities, such as a 15,000 cd/m2 HDR display, or a multi-focal HDR stereo display. Such collaboration is partly funded by the EC H2020 Innovative Training Network grant (€820k in Cambridge) and an EPSRC grant (£320k in Cambridge). The latter project aims at developing spatio-chromative vision models, which will be a foundation for efficient visual coding and adaptive display algorithms.

Other research activities include Endeavour (Enhancing Driver Experiences through Vision Research), which is a £570k collaboration with industry to investigate applications of computer vision in the automotive domain. This involves exploring ways in which computer vision can be used to enhance the experience of driving a car. Peter Robinson is continuing investigations of human interfaces for future transport, looking at ways that drivers interactions with the rapidly changing world of motoring can be adjusted to increase their confidence, comfort and safety. Damon Wischik is co-PI on a £650k 18-month project at the Alan Turing Institute, sponsored by the Toyota Mobility Foundation, on using AI to improve traffic flow. The project will investigate two themes: reinforcement learning for better traffic signal control and modelling the interaction between cities and fleets like Uber.

The group includes 5 Academics, 8 Research Associates and 8 PhD students and one College Junior Fellow.

Natural Language Processing and Information Processing Group

The Natural Language & Information Processing (NLIP) group conducts research into a broad range of topics involving the intersection of computing and language (natural/human languages as opposed to computer programming languages) and has a high international standing. Their research falls within two areas: Natural Language Processing (NLP) as a branch of Artificial Intelligence, the ultimate aim of which is for humans and computers to interact using natural language. Members of the NLIP group work on methods for automatically discovering the underlying structure of language; methods for representing and generating meaning; as well as

specific computer applications, including (but not limited to) document summarisation, automated marking of essays and automatic translation between languages. The NLIP group is also concerned with understanding language processing in the brain. Having recently founded the "Language, Minds and Machines" Special Interest Group (within the context of the Cambridge Language Sciences Interdisciplinary Research Centre) members of the NLIP group are also interested in how computational language processing methods may yield predictive models for language cognition.

Natural Language Processing is one of the most challenging modalities for modern Machine Learning. A recent paradigm shift in NLP finds the group working with and contributing to cutting-edge techniques for Machine Learning including Deep Learning. In this respect the group has demonstrated an impressive ability to adapt as the research environment rapidly evolved – they have reoriented themselves around big-data driven technological advances, whilst still grounding their work in linguistic expertise.

The NLIP group has an excellent record of collaboration within the Department, across the University and with other national and international research institutes. Current research projects see intra-departmental collaboration with the Security and Digital Technology Group; across the University with CRASSH, the Cambridge Africa Foundation, the Engineering Department, MML (Linguistics) and Psychology; nationally with the Alan Turing Institute; and internationally with Makerere University, Uganda and Tokyo Technological Institute, Japan. There are non-funded collaborations with individuals in many other universities, including Stanford, Washington and Trento. The group also has existing and growing ties with local and international industry including Amazon, DeepMind and Facebook. This ensures we are always working on cutting-edge problems. Group members are sought-after for their expertise on national and international panels: Ann Copestake in currently a member of the Royal Society Fellowship review panel and will be a member of the panel reviewing applications for the European Research Council from 2019. Reflecting our world-leading expertise in Language Assessment, several members of the group are currently running the Building Education Applications workshops (which are co-located with the field's largest conference).

The group also has a central role in research teaching at Masters level and contributes disproportionately to the MPhil in ACS/Part III: on average 25% of students do NLIP projects, and the combined student uptake of NLIP modules is regularly over 20% of the total, with the concomitant assessment load. There is such high demand for our students from industry at the moment that all the NLIP courses on the ACS are always oversubscribed. In the past some very good students have been lost to other universities because the NLIP options on the ACS are insufficiently broad.

The NLIP group was historically an outlier in the Department, partly because of the research area and partly because of its emphasis on Masters-level teaching, before the general ACS Masters course was introduced. That situation has now changed completely, with the increase in importance of machine learning across the Department, the decision to start teaching it from the first year of the undergraduate Tripos (led by members of the NLIP group), increased collaboration with other groups in the Department, and increased departmental emphasis on interdisciplinarity. Combined with the rapid changes to the field of NLP as a whole and the recent changes in group membership, the group is in a period of considerable flux. While we do not expect that the new DeepMind Chair will be an NLP researcher, the new appointment may well be very helpful to the group.

It is clear that there are great opportunities to collaborate and obtain funding, and many exciting research areas to pursue, with no shortage of students interested in working in the area at Masters level or PhD, though we are seeing promising PhD students being lost to companies at an increased rate. The descriptions of ongoing research below demonstrate that existing members of the group are exploring a range of new areas. In this landscape, the renewal of the ALTA funding is a very important stabilizing factor. The NLIP group's emphasis on interdisciplinarity and its involvement with Cambridge Language Sciences fits with the general Departmental strategy of emphasizing what differentiates university research from research in companies. The move into areas of more direct relevance to cognition is particularly important in this respect. However, while the group does not want to compete directly with commercial research, it is necessary to have a minimum level of compute resource in order to be able to run experiments on neural network architectures, and resourcing this is a challenge.

Research Highlights

The group has an excellent publishing record, without exception publishing multiple papers every year at all the major conferences in the field. The interdisciplinary nature of our work also sees members of the group regularly publishing in Linguistics and Education journals as well as contributing book chapters to titles in Linguistics. For example, Ann Copestake (together with PIs in Engineering and Linguistics) is leading a project funded via the Centre for Research in the Arts, Social Sciences and Humanities (CRASSH). The project —'Giving Voice to Digital Democracies: The Social Impact of Artificially Intelligent Communication Technology'— started with a pilot phase, with a number of panels bringing together researchers from companies and third sector organisations with academics in Computer Science, Social Sciences and Humanities. The main project will employ three post-doctoral researchers based in CRASSH.

Former doctoral student Guy Emerson has recently been awarded a Junior Research Fellowship to work on methods for automatically learning human interpretable computational representations of meaning. In his PhD thesis, *Semantic Functions: Learning Interpretable Representations of Meaning From Disparate Data Sources,* Guy developed a semantic framework which can be used to learn the meanings of words, based on a large body of text. Having built on well-understood modelling techniques, both from formal linguistic theory and from probabilistic machine learning, there are natural ways to extend his framework. During his fellowship, he plans to incorporate data sources beyond just text, including both grounded data (such as labelled images) and structured data (such as dictionaries and databases). Building a coherent semantic model that can use such disparate data sources would constitute a major step forward in computational linguistics.

The Cambridge Institute for Automated Language Teaching and Assessment (ALTA) is a virtual institute that brings together researchers from Cambridge Assessment and the Departments of Computer Science and Technology, Engineering, and Modern and Medieval Languages (Linguistics). Using techniques from Machine Learning and Natural Language Processing, ALTA's primary objective is to develop technology to enhance the experience of language learning and to develop cutting-edge approaches to automated assessment that will benefit learners and teachers worldwide. This interdisciplinary institute, financed by Cambridge Assessment, was initially funded for a five-year period and has recently been awarded continuation funding for a further three years (with the expectation of rolling funding for subsequent three-year periods). Under the inaugural Directorship of Ted Briscoe, ALTA has become world-leading in the broad and challenging application area of automated

assessment and also in providing feedback for further language learning. The Institute will continue under the Directorship of Paula Buttery and has extended its research focus into the areas of personalised/adaptive learning and the automatic generation of teaching materials. ALTA currently helps to support four University Teaching Officers and employs around 12 post-docs as well as one or two PhD students a year.

Over the last three years, Professor Ted Briscoe has had two Innovate UK grants awarded jointly with All Street,[15] a London-based start-up using analytics to generate investment reports for individual investors. The company has just launched its first product and is undertaking its first funding round.

The group has several current research awards to work with not for profit groups. Paula Buttery has a collaboration with Africa's Voices that will fund a Research Assistant in 2019. Africa's Voices is a non-profit research organisation spun out of the Department of Politics and International Studies, University of Cambridge. Africa's Voices converses with African citizens via SMS and social media on a range of health and education related issues. The foundation's goal is to understand attitudes and beliefs in the target populations in order to enrich the aid that support organisations can provide. Paula has also been awarded funding from Cambridge-Africa Alborada Research Fund and the Cambridge Africa Partnership for Research Excellence. These projects involve Computer-Assisted Language Learning (CALL) for the revitalization of endangered languages (in collaboration with Fridah Katushemererwe, Makerere University, Uganda).

In collaboration with Alice Hutchings in the Security group, Paula Buttery has been funded by the Defence and Security Programme at the Alan Turing Institute to investigate NLP approaches to understanding key actors on online hacking forums. Current NLP technology has severe limitations in processing important but non-standard writing (such as the posts on internet forums or on social media). This project looks at the language of grey hat forums (online communities that bring together individuals interested in hacking and illicit monetising techniques). By developing appropriate NLP technology in collaboration with criminologists and network analysts, it is possible to identify key players in these criminal domains.

From October 2018, the group expects to have 6 Academics, 5 Senior Research Associates, 3 Research Associates, 1 Research Assistant, and 17 PhD students. One of our current PhDs has recently won a Junior Research College Fellowship.

Programming, Logic, and Semantics Group

Nearly everything in the modern world relies upon computers, yet even tiny oversights can lead to crashes, wrong answers, massive leaks of sensitive information, or allowing malicious adversaries to seize control. The research of the Programming, Logic, and Semantics Group centres around programming languages, semantics and verification techniques to explore new ways of creating and analysing computing systems (hardware, software, network) that are easier, faster, safer and offer reliable guarantees. We investigate foundations grounded in logics and mathematical models such as category theory, and also collaborate with researchers and practitioners in systems and application domains to address pragmatic concerns and validate

[15] www.allstreet.org/allstreetplatform

our principles at scale. We also develop mathematical techniques for analysing the complexity and correctness of computational systems.

Research Highlights

The REMS project (2013-2020) is developing rigorous semantic methods for the engineering of mainstream systems, especially at the architecture and programming-language levels. Funded by an EPSRC Programme Grant, it involves PLS faculty Professor Sewell (PI), Professor Pitts and Dr Krishnaswami; computer architecture, systems, and security colleagues (Watson, Moore, Crowcroft, Madhavapeddy) in the CHERI project; and Systems Research Group, and colleagues in Edinburgh and Imperial (Stark and Gardner); it also has close links with ARM and IBM. Highlights include:

- Work with ARM to make their user-mode concurrency architecture well-defined, now incorporated into their ARMv8 architecture specification, and work with the emerging RISC-V architecture to do the same. The RMEM tool for exploring these models.
- Work with ARM to establish usable theorem-prover and emulation-capable versions of large parts of the full ARM 64-bit instruction-set architecture, including systems aspects, using our Sail ISA description language.
- Development of the Cerberus semantics for large parts of C, and ongoing engagement with the ISO WG14 standards committee and key industry figures to clarify the C memory object model. The Cerberus tool for exploring these models.
- Work with the CHERI project, developing research architecture and a software stack providing fine-grained protection and secure encapsulation, using a formal ISA model as a key design tool.

Peter Sewell plans to build on some of the work developed in REMS, especially on architectural and systems-programming-language semantics, and with the CHERI project, to develop better engineering techniques for real systems, focused especially on their fundamental security properties.

Glynn Winskel obtained an Advanced Grant from the European Research Council for his project ‘Events, Causality and Symmetry’ (ECSYM, 2011–2017), which aims to build a new foundation for the semantics of computation. Traditional semantic techniques have fallen short in several respects. Right from the start, a central aim to represent distributed interactive computation led to an emphasis on causal models and in particular on event structures in which events are related by their causal dependencies. The demands of compositionality (the feature of explaining the behaviour of complex systems from that of their parts) led to a new theory of concurrent games in which interactive processes are understood as distributed strategies and their types as distributed games. The discovery of concurrent games is beginning to have a broad influence as they are applied to areas outside the traditional range of semantics, for example, to a semantics of classical proofs as concurrent strategies, which in particular leads to a compositional proof of Herbrand’s theorem; probabilistic distributed strategies, and a form of “structural game theory” in which games and optimal strategies can be programmed; an extension to quantum strategies, leading to a fully abstract model of the quantum lambda calculus; to other difficulties within the semantics of computation generally, from weak memory models to probabilistic programming.

These lines of work continue in collaboration with and between colleagues at ENS Lyon, Paris, Imperial, UCL, Oxford, and recently at ARM; Winskel is keen to continue his collaborations

and work on concurrent games and strategies, exploiting the logical and quantitative features that they share. He plans to provide further value theorems for probabilistic concurrent games, the keystones of a structural game theory.

An ERC Advanced Grant has been awarded to Lawrence Paulson for his project 'Large-Scale Formal Proof for the Working Mathematician' (ALEXANDRIA, 2017–2022). Its overall goal is to make existing verification tools useful to research mathematicians, in addition to the original target group of computer technologists. The research programme begins with a pilot study in which mathematicians formalise advanced material using Isabelle/HOL. This has the twin aims of identifying issues and adding to Isabelle's libraries of mechanised mathematics. Other themes of the programme include intelligent search (so that users can easily locate relevant theorems and proof fragments) and computer algebra. When mechanising proofs, a primary objective is to retain the link between the formal proof and the original material, with the formal proof recognisable as mathematics in its own right. Paulson's ALEXANDRIA project will dominate his work for the next few years, so to a focus on the mechanisation of mathematics. However, he retains his interests in computer security (e.g. bitcoin-like protocols) and theorem proving technology.

Anuj Dawar's current research involves further investigating the limits of algorithms that respect natural symmetries in data (in particular, algorithms that are automatically generated from high-level specifications) as well as the complexity of detecting symmetries in data. Funded by an ERC starting grant for the project 'Dynamics of Multiple, Interacting and Concurrent Markov Chains', till 2021), Thomas Sauerwald aims to investigate the foundations of parallel random walks (e.g. the time it takes to explore a network) and use these insights to design more efficient algorithms for massive networks and data sets.

Andrew Pitts has recently begun to contribute to the area of Homotopy Type Theory, which lies at the interface between logic (constructive type theory), computer science (functional programming and proof assistants) and mathematics (algebraic topology); he plans to focus on its applications to the foundations of computer systems for constructing formal proofs in mathematics and software verification. Neel Krishnaswami is designing a new family of type inference algorithms, which are both easy to implement and easily handle advanced type-theoretic features (such as type dependency, modality, and sub-structural logic). The goal is to apply this research to the design of a new memory-safe low-level programming language which both offer programmers fine control over efficiency considerations such as data and stack layout, while still permitting the use of high-level reasoning principles and programming idioms.

Alan Mycroft is currently working with his three PhD students on two projects. The first works towards a better understanding of object-oriented vs functional programming styles under the moniker 'the expression problem', and investigating to what extent partial evaluation can automatically re-express programs. The second has to do with the idea of effect systems and their dual of co-effect systems. Recent work uses co-effects to give an effect system for call-by-need (lazy) languages.

The group includes 9 Academics, 1 Royal Society Fellow, 1 Marie Curie Fellow, about 10 Research Associates and 23 PhD students.

Security Group

Computer security has been among the Department's research interests for many years, along with related topics, such as cryptology, formal methods, hardware design, biometrics, and the robustness of distributed systems in general. Our research strategy, insofar as we have one, has always been to hire the brightest people we could find and let them do whatever turns them on. That said, there are two large research projects that are currently in full flow: CHERI and the Cambridge Cybercrime Centre. CHERI refers to Capability Hardware Enhanced RISC Instructions, an Instruction-Set Architecture (ISA) extension that implements a hybrid capability-system model providing fine-grained memory protection and scalable software compartmentalisation within processes. This has been supported for a number of years by DARPA with additional funding from Google and ARM; we have implemented a prototype processor and done extensive work on the toolchain to explore the possibilities. In addition to publishing a series of systems papers since 2012, we hope to see real-world deployment in the next two to three years, leading to major global impact on access control architectures. The Cambridge Cybercrime Centre was established as a global resource for cybercrime research, making a wide range of datasets available to academic researchers worldwide. We currently make data and tools available to over 20 researchers worldwide, covering a wide variety of online harms from spam and phish to malware and botnet command-and-control traffic. It is our ambition to become the hub of worldwide cybercrime research by providing the data and tools that academics require to do research on the same basis as Google, Facebook or national law-enforcement agencies.

More generally the group's research over the past 20 years has pushed out the boundaries of the subject by studying "security of X" for various applications X, and "security and Y" for various disciplines Y. Previous values of X include health records, payments, prepayment meters and industrial control systems, while Y has run through signal processing, economics and psychology. These interactions have led to a number of continuing strands of research interest including payment fraud, medical privacy, emission security and the economics of security and dependability – which in turn led to the Cambridge Cybercrime Centre. This tactic does not yet appear to be exhausted; for example, a new research student has X = autonomous vehicles and Y = AI/ML.

Research Highlights

Ross Anderson was awarded the Lovelace Medal (2015), the UK's top award in computing in recognition of his many contributions to building Security Engineering into a discipline. In 2016, Anderson, along with Alice Hutchings and Richard Clayton, produced 'Taking down websites to prevent crime' (eCrime 2016), a study for the Home Office reporting on what works and what does not when trying to take down illegal websites. More recently he has published work on cryptocurrencies, such as 'Bitcoin Redux' (WEIS 18). This presents better tools for tracing stolen bitcoin and analyses what we can learn about regulatory failures in the bitcoin ecosystem.

Robert Clayton and Anderson have produced 'Standardisation and certification of safety, security and privacy in the "Internet of Things"' (EU 2018). This report for the European Commission analyses what changes will be needed in device safety regulation once software in cars, medical devices and other safety-critical equipment starts to be updated regularly. Clayton has also been working on 'Bitter harvest: Systematically fingerprinting low- and medium-interaction honeypots at internet scale' (WOOT 18); this is a class break of honeypots

currently deployed on the Internet, showing how single-packet probes can distinguish them from the systems they emulate.

Within the group, The Tamper Lab (Anderson, Kuhn, Skorobogatov) studies existing security products, documents how they have been penetrated in the past, develops new attack techniques, and tries to forecast how newly available technologies will make it easier to bypass hardware security mechanisms. They then develop and evaluate new countermeasures and assist industrial designers in staying ahead of the game, most of all by giving them an advanced understanding of which attack techniques are most dangerous. They are especially interested in protection systems for mass-market applications, and in forensic applications. For example, Sergei Skorobogatov has shown how the iPhone 5c password guessing limit can be circumvented, contrary to claims made by the then Director of the FBI (2016). More recently, in 2017, Skorobogatov has been reporting the reverse engineering of the Insulet insulin pump to enable access for researchers trying to develop an artificial pancreas. Markus Kuhn's research (2018) has shown how to improve template power-analysis attacks against microcontrollers to the point where direct distinction of all 256 values of on-chip parallel 8-bit data interfaces becomes possible, even when targeting just a single load instruction, and across multiple devices. The Tamper Lab is a world centre of expertise in semiconductor reverse engineering and compromising emanations, where semi-invasive and optical side-channel attacks were pioneered and which continues to push and define the state of the art in reading out information from semiconductor devices.

Interdisciplinary research is also being carried out by a number of faculty (Hutchings, Clayton, Beresford) who have been exploring ethical issues in research, such as using datasets of illicit origin, including the ethics of doing research with data harvested from underground fora, bots and other criminal activity.

The group's largest single project was published in April 2017 as 'Capability Hardware Enhanced RISC Instructions: CHERI Instruction-Set Architecture (Version 6). Technical Report UCAM-CL-TR-907, Computer Laboratory, April 2017.

The group includes 5 Academics, 1 Leverhulme Early Career Fellow, 1 Principal Research Associate, 3 Senior Research Associates, 8 Research Associates and 17 PhD students.

Systems Research Group

The group undertakes teaching and research into topics encompassing computer systems in general and more specifically: design and analysis, including network monitoring and protocol design, mobile systems, operating systems design and networked-system practical distributed systems. We regularly incorporate cutting edge elements of other Computer Science disciplines, such as formal methods, programming languages, digital hardware, machine learning and security into our research. Our methodology as a systems group is to combine these research strands into an artefact-led scientific evaluation in order to shed new understanding in the interplay of complex computer systems in an interdisciplinary environment.

Research Highlights

Recent highlights include the five-year (to date) MCCRC programme of multi-disciplinary work in Cloud Computing and Law, focussing on privacy and localisation, analytics use of machine learning and transparency, blockchain and ownership models, and IoT and liability. This work is funded through the Microsoft Corporate HQ compliance/legal group, and is a collaboration with Queen Mary University of London Law Department. Related to this are two EPSRC funded projects: 'Towards a legally-compliant Internet of Things' and 'Realising Accountable Intelligent Systems (RAInS)' (a TIPS 2.0 in collaboration with University of Oxford and Aberdeen). Both of these are interdisciplinary projects that focus on technical measures for better aligning networked and mobile systems, data management, and data processing (ML) aspects with legal and accountability concerns.

The Hub-of-all-things (HAT) project which started life as digital economy research and is now a start-up which offers a platform for new business models and markets in personal data, with over £30M of partnership funding, and deployments ongoing in multiple countries. The work is with Warwick University and a number of companies. In a related project, Maru, which is funded through the Turing Institute, we are building platforms to exploit secure enclaves for processing sensitive personal data (e.g. healthcare, financial) in multi-tenant data centres, with improved security. The work is a collaboration with Imperial College and also relates to the use of Trustzone potentially for the Databox and HAT work. Several aspects of this work link with other group members work in verified systems, e.g. under the EPSRC REMS programme grant.

Three specific projects are currently underway that align with the Human-Data Interaction agenda (Mortier et al, IDF HCI Encyclopaedia 2015). The first, the EPSRC-funded Databox project is designing, building and evaluating a platform for mediating access to personal data. The core innovation is to consider how we can support movement of data processing code to a computing platform where the personal data is kept, rather than the today's approach where personal data is sent to the cloud with the resulting loss of legibility, agency and negotiability. The second, the soon-to-start EPSRC-funded DADA project will explore how to secure domestic environments from rogue IOT devices by making the network behaviour of such devices legible, and providing means to enforce control over such devices. The third, Owl/Actor, is a distributed machine learning framework in the OCaml language that targets building machine learning and other data-processing systems for small form-factor low-power computing platforms, distributed over the wide-area, in support of systems such as both Databox and DADA.

The group is also involved in larger Centres that have interests in the Digital Economy, the Internet of Things, and their effects on society. For example, the group is a spoke of Horizon Digital Economy Research based at the University of Nottingham, and is represented in the Investigator team of the Centre for Digital Built Britain, based in the Institute for Manufacturing here in Cambridge.

The group is a worldwide hub of knowledge in networked-hardware, mostly through the NetFPGA project, an open source platform for networking research and teaching. The platform has over 1200 users in more than 200 universities in 47 different countries. The NetFPGA-SUME platform, released by our team in 2015, has over 600 users and over 800 boards sold. The NetFPGA platform was used to develop a number of successful open-source projects, Emu (open source network as a service), P4-NetFPGA (open source programmable data planes) and the NDP switch, which won SIGCOMM 2017 best paper award. Alongside being a research

and teaching vehicle, the NetFPGA platform underpins the development of other tools such as the OSNT: an extensible open source network tester, the pcie-bench toolkit to aid understanding in PCIe-related performance for end host networking, and NRG, a hardware and software toolset, enabling reproducible large-scale networked-system experimentation in a small-scale environment and the study of network effects on application's performance.

The OCaml Labs[16] group is at the centre of modern functional programming developments. The group has spun out a successful consultancy wing, and is (along with Inria) responsible for the stewardship of the OCaml language and tooling. Our technology has been adopted by some big established industrial groups (Facebook, Microsoft, Jane Street) and up and coming start-ups (the Tezos cryptocurrency). The research and industry work all combine via open-source interaction, with millions of lines of code having been published since 2012. Research projects are incubated in this environment – the MirageOS unikernel, the Irmin distributed data store, and the multicore memory model are being picked up by industry in all sorts of unexpected deployments, as well as being published in top-tier conferences.

Cecilia Mascolo has recently become Director of the Centre of Mobile, Wearable and Augmented Intelligence[17], with Nokia Bell Labs as founding partner. She has published a range of papers on the use of local computational resources of mobile devices for enabling efficient and effective continuous audio sensing. Her work on Project EmotionSense was one of the first to adopt mobile sensing as digital biomarker for mental health diagnostics. In addition, the work using the microphone as a sensor for human emotion analytics on device (published in Ubicomp 2010) was pioneering the area of extracting features from sound to augment digital biomarking efficiently. Mascolo is also publishing work on the use of large urban data analytics to improve the understanding of how cities function and she has collaborations with major companies including Foursquare and Tencent for the analysis of the data for urban planning.

The group includes 8 Academics, 1 EPSRC Research Fellow, 1 Leverhulme Early Career Fellow, 4 Senior Research Associates, 29 Research Associates, and 28 PhD students.

[16] http://ocamllabs.io
[17] http://mobicentre.cst.cam.ac.uk

Chapter 4: RESEARCH COLLABORATION

Collaboration is a natural way of working in the Department, given our long history of generating substantial impact coupled with our relatively small scale. Whilst the majority of research projects focus exclusively on computer science, a growing proportion are interdisciplinary, and in recent years have come to include disciplines such as Sociology, Psychology, Economics, Architecture, Clinical Sciences, Law and Ethics as well as our long-standing relationships with Linguistics, Mathematics and Engineering.

Within the University

In addition to collaborations between the groups within the Department, there are many formal and informal collaborations with other departments across the University (this is particularly true for Engineering, specifically, Information Engineering). In June 2018, the two departments jointly received an award from the University's Academic Seed Fund to establish 'The Observatory for Human-Machine Collaboration', which will enable holistic study of the fundamentals of human-machine interaction by putting human and technological perspectives on an equal footing. Minute examination of these interactions means that researchers will be able to explore how collaborations between people and machines can be optimised, and create outcomes that are more effective, adaptable, reliable and socially acceptable as a result. The central feature will be a fully equipped laboratory containing a smart, multi-modal metrology system, alongside a Virtual and Augmented Reality system and a set of high-performance projectors and loudspeakers. The initial award was for £600K, and we plan jointly to submit a Capital Equipment bid to EPSRC for further funding shortly.

Earlier this year, we also entered into a pilot scheme with Engineering and the University of Michigan to seed-fund new collaborations between academics in the two institutions (Cambridge and Michigan) using our own internal funds. The first successful projects were announced in May 2018 and included a proposal from Nada Amin (CST, in collaboration with Kasikci, Michigan) on 'Specifying and Checking Performance Properties via Hardware-Software Co-Design'.

In terms of collaborations with University departments outside the School of Technology, Pietro Lio's work on machine learning approaches for the analysis of bio-medical data sets has led to collaborations with Clinical Neurosciences, Haematology, Oncology, Psychiatry, Radiology, and Veterinary Medicine. Sean Holden has been collaborating with researchers in the Department of Biochemistry since 2012, applying Bayesian inference, probabilistic programming and computational learning theory to drug design and proteomics. Cecilia Mascolo has a long history of working with different domains on the application of mobile sensor technology. She is currently collaborating with Clinical Neuroscience on early diagnostic methodologies for Alzheimer's disease, and with the Department of Health Psychology on behavioural interventions delivered via mobile devices. Outside of the clinical sphere, she is working with Departments of Architecture and Urbanism to study city functionality through digitally crowdsourced data.

Alan Blackwell, as Professor of Interdisciplinary Design, has specialised in the development of interdisciplinary approaches to technology development, and in particular to collaboration with the Arts, Humanities and Social Sciences. He created the Crucible network for research

in interdisciplinary design in 2000, as one of the initiatives of the Cambridge-MIT Institute. Since then, Crucible has facilitated more than 180 interdisciplinary projects, involving many members of this Department. Spin-offs from Crucible activity have included the Cambridge Digital Humanities initiative (now hosted within the University Library and CRASSH Centre for Research in Arts Social Science and Humanities), and the interdisciplinary approach of the Cambridge Global Challenges initiative. A key element of Crucible activity has been reflection on processes of interdisciplinarity, assisted by anthropologists and philosophers of science. Findings from the network have been more broadly influential in computer science, for example as a model for interdisciplinary research and teaching in HCI.

Members of the NLIP group have a long history of collaboration with linguists and psychologists in Cambridge and elsewhere, as well as with Information Engineering. The NLIP group was central to the formation of the Cambridge Language Sciences Strategic Research Initiative (SRI) in 2011, one of the first University SRIs. SRIs build on areas of existing research strength by bringing together a critical mass of expertise from across the University's six academic Schools, with four key aims: to address large-scale multidisciplinary research challenges; to strengthen research collaborations across disciplines; to increase research capacity and profile; and to influence research, policy and funding agendas. Research on Language Sciences is relevant to all Schools in the University and the introduction of the SRI allowed researchers from different disciplines to become aware of the related work that was going on across the University. The ALTA project, led by Ted Briscoe from the NLIP group, was a key demonstration of the importance of the multidisciplinary approach to language learning and assessment. The SRI was upgraded to an Interdisciplinary Research Centre (IRC)[18] in October 2017, which allows more independence from individual departments, thus (among other things) facilitating collaboration with external organisations. Ann Copestake currently co-directs the IRC, together with Professor Ianthi Tsimpli (Linguistics).

In 2017, the Department was awarded University funding to establish a Strategic Research Initiative (SRI) in Trust and Technology. Due to be formally launched on 20th September 2018, the Trust and Technology Initiative[19] will draw on departments across the University to "*explore the dynamics of trust and distrust in relation to internet technologies, society and power; to better inform trustworthy design and governance of next generation tech at the R&D stage; and to promote informed, critical and engaging voices supporting individuals, communities and institutions in light of technology's increasing pervasiveness in societies."* It is chaired by Professor Simon Moore and Dr Jat Singh (both CST).

We were also awarded funding to establish the Cambridge Global Challenges SRI.[20] Borne out of a recognition that Cambridge operates as a global university with responsibilities beyond the UK, and that global priorities are encoded in the UN Sustainable Development Goals, Cambridge Global Challenges is focused on promoting research that will advance the SDGs, working in collaboration with those living in low-income countries. The need for collaboration across disciplines, in partnership toward applied goals, has provided an opportunity to apply principles developed out of computer science expertise in interdisciplinary design.

[18] https://www.languagesciences.cam.ac.uk/
[19] www.trusttech.cam.ac.uk/
[20] www.gci.cam.ac.uk/

Other academic institutions and non-commercial organisations

The Department has collaboration with a wide range of other academic organisations, both nationally and internationally through either informal long standing interaction or formally through research grants. A couple of most prominent examples are:

Alan Turing Institute

The Alan Turing Institute is the UK's national institute for data science and artificial intelligence. The University of Cambridge was one of the five founder institutions. The Alan Turing Institute was formed three years ago by five universities and a number of strategic industry and government partners, to act as the national focus for data science, and more recently, after the House of Lords report on the topic, for AI. The Institute has its physical hub in the British Library in London (next to the Crick and, notably, Deepmind). Here, there are a core of PhD and post-doc research programs supported by around ten Research Software Engineers, as well as a business and events team for industry and public outreach. However, it also runs a highly decentralised operation, drawing on faculty from many Universities, who act as visiting fellows a few days a week. The technical work covers the underpinning core principles, tools and techniques of data science, including mathematics, statistics, algorithms, systems, NLP, and, uniquely (compared to other AI & big data organisations), social data science, ethics research, as well as hardware (for accelerating or securing ML).

The Cambridge Department of Computer Science and Technology was a founding member, and has several faculty spending time, there as well as advising Turing PhDs and running projects. Some of the key people, and their types of engagement include:

- Professor Crowcroft who chairs the Program Committee - the ATI's research strategy advisory group. He also has a project in collaboration with Imperial on providing confidential machine learning platforms in the cloud.
- Professor Gibbens and Dr Wischik work in transport, specifically in electrical vehicle charging networks, where the constraints for scheduling make an interesting challenge.
- Professor Mascolo has work on sensing for health and well-being with some seed funding from the Turing on use of mobile sensing data for early diagnosis of neurodegenerative diseases such as Alzheimer's.
- Professor Dawar is contributing to foundational work in logic for data science.
- Dr Buttery has a project in NLP in understanding social media and unrest.

Recent outputs from the Turing Institute include: Grammenos' (Mascolo's PhD student) two publications; one on automatic calibration of accelerometer data in mobile phones to produce data more suitable for medical studies and another for a new result for bounded latency in real time stream processing. Professor Crowcroft has published a paper on techniques for running secure (in the sense of confidential) computations in untrusted cloud platforms.

The Institute recently expanded with the addition of at least eight more universities, many of whom we are already in collaboration with. In addition, several more Cambridge Computer Scientists have become engaged with work there including: Dr Yoneki (SRA) on optimisation of complex data processing frameworks in Machine Learning; and Dr Singh (Fellow) on the policy/legal/ethical aspects of AI/ML. Finally, as part of the Turing's program for systems at scale, we expect to release new versions of Spark & Flink that exploit Intel's SGX facility to provide hardware encrypted memory.

Smart Cambridge

More locally, building on our interests in sensors and smart cities, we are directly supporting Smart Cambridge,[21] overseeing the design and construction of the urban digital infrastructure across the region. We have implemented regional infrastructure far more capable than would otherwise be the case, both benefiting people in the region and also fit-for-purpose for our own research. In addition, researchers in other departments in the University (including Architecture, Chemistry, Engineering and Land Economy) are using the data from our sensors in their own research projects and programmes.

SRI International

SRI International (previously Stanford Research Institute) has its headquarters in Menlo Park, California, and a number of additional outposts. For over 70 years, SRI has been an independent, non-profit research centre that takes the most advanced R&D from the laboratory to the marketplace. Serving government and industry, they collaborate across technical and scientific disciplines to invent solutions that solve challenging problems today and look ahead to the needs of the future. Their achievements range from Siri and online banking to medical ultrasound, cancer treatments, and much more.

Over a decade ago, the Department worked with SRI on speech technologies. More recently, the Department has had eight consecutive years of joint funding from DARPA (Defence Advanced Research Projects Agency) with SRI as the prime contractor, to research fundamentally more secure computer systems. In Cambridge, this work has been led by Dr Robert Watson and Professor Simon Moore. In collaboration with SRI, we have established a team of over twenty researchers with a broad range of computer science skills, enabling work on new processor architecture, hardware prototypes, compiler support, operating system porting, and benchmarking of applications. This level of funding has gone beyond anything available through standard EPSRC funding mechanisms (e.g. considerably more funding than is available on a Programme Grant) and has been critical to hire a team with the required breadth of skills to undertake the research and to compete with top universities in the USA.

Industrial Collaboration

The companies that the Department works with range from nascent local start-ups (often founded by our alumni) to large global businesses, and the ways in which we collaborate are many and varied. Often, collaborations are informal, with individuals on both sides sharing information on the latest developments in research, and industry examples that are pertinent to teaching. Formal collaborations with smaller companies often consist of a joint application for funding to organisations such as Innovate UK (for example, in the last three years, Briscoe has had two consecutive Innovate UK projects with All Street, who have just launched their first product, and the outputs of Gunes' recent project with 'Sensing Feeling' are currently being commercialised). Larger companies generally make either financial or in-kind contributions to funding applications or provide direct support for research projects and/or PhD studentships. Between October 2009 and April 2018, 18% of students (45/250) received some sort of support from industrial partners including the large donation to the School of Technology by Qualcomm. The figure seems to be increasing: between October 2014 and April 2018, the industrial support has been around 19.8% (22/111). Overall, funding from industry partners does seem to be on an upwards trajectory.

[21] https://smartcambridge.org/

We have had a high success rate for applications to the EPSRC Impact Acceleration Account held by the University, where the purpose of the funding is to aid the transfer of ideas developed in EPSRC-funded research grants into more mature prototypes which could be then transferred into the real world in various forms, including commercially. Over the last four years, we have received funding for a range of Knowledge Transfer Fellowships, which support a post-doctoral researcher to spend a year working with an industrial partner to apply the results of an EPSRC-funded research project to that business. Examples include TranscenData (2014), ARM (three in total during 2016-17), Smart Cambridge (2017) and the Met Office (2018). Some examples of larger and more strategic relationships are outlined in Chapter 2, under 'Research Funding'.

Our approach to industrial collaboration centres on finding real alignment between academic research interest and the 'problem' the company wants to solve, and in doing so from the outset. Often, an initial approach from a potential partner is made to the Research Facilitator, who will engage in a discussion with the company to identify the questions they want to address and how these might overlap with research interests in the Department, in addition to any resources that are available or that might need to be applied for. The company would then normally provide a summary of their requirements in writing, and the Research Facilitator will then share and discuss this with relevant academics and take a view on whether to proceed to face-to-face meetings. This careful triaging of approaches and managing of expectations alongside, aims to ensure that both academics and companies derive some benefit from any meetings that occur – even if in many instances this is more of a 'meeting of minds' rather than progression to a tangible research project (at least at that point).

Chapter 5: RESEARCH AND TEACHING

Undergraduate Teaching

The three-year undergraduate degree programme taught by the Department (the Computer Science Tripos, delivered to an annual intake of around 100-120 in each year, currently with over ten applicants for each place) is strongly oriented toward teaching recent research, and giving undergraduates practical research experience. Very few of the courses follow standard textbooks. Instead, the majority of the teaching is carried out by active researchers in the field, who generally draw on the latest research when introducing advanced elements of courses at every level.

At the time of the review, we are completing a total overhaul of the curriculum, and are introducing practical research exercises in computer science for first-year students as an alternative to the traditional option of a natural science subject, such as Physics or Chemistry. This reform has also introduced a third-year course that allows undergraduates to select two Masters-level papers in the final year of their BA. This option has proven extremely popular, and has been taken up by 75% of the first cohort.

A core component of the second year (Part Ib) in the Tripos is the Group Design Project, in which students work in teams to create a prototype of an innovative product, guided by a "client" recruited from the network of industrial supporters and research collaborators affiliated with the Department. Many of these projects attempt ambitious applications of recent research, recently including topics such as blockchain, deep learning, augmented reality, internet of things, computer vision and so on. These projects are often used as testing grounds for concepts with research collaborators, leading to research internships for the students, research career opportunities, or further exploration of specialist problems in further years of their degree. Many undergraduate students pursue research internships in the Department during summers, administered via the Cambridge UROP (Undergraduate Research Opportunities Programme) scheme.

The largest components in the third year (Part II) of the Tripos are an individual project and dissertation. Both of these are individually supervised by an active researcher in the Department. Although students are warned of the dangers of attempting an overly ambitious undergraduate research project and course guidance explicitly states that original research is not expected, many students and supervisors do use this opportunity to produce novel results, potentially leading to research publications or to further investigation in graduate research.

Students who complete the three-year Tripos with a first-class result have the option, instead of graduating with a BA, of continuing study for a fourth year (Part III) leading to the MEng. This fourth year follows, as closely as possible, the MPhil in Advanced Computer Science with its emphasis on research preparation, as discussed in the next section.

Masters Teaching

The MPhil in Advanced Computer Science, and closely related "Part III" (optional 4th year) of the Computer Science Tripos, are tightly focused as 'research preparation degrees'. Promotional material emphasises that the degree is designed to prepare students for the start of

a PhD (and subsequent research career). The core element is a research dissertation on a topic arranged with a personal supervisor, weighted as 7/12 of the final grade for MPhil students, and 4/9 for Part III students (reflecting shorter undergraduate terms).

MPhil and Part III students select an individual course of study composed of five modules selected from about 30 options spread across all the research groups of the Department. A wide range of study and assessment styles are provided, reflecting the different research styles and skills required in different subfields of computer science. These range from reading groups and seminars for qualitative topics such as computer security and systems research, to examined lecture courses for theoretical topics, and a substantial number of practical courses incorporating research mini-projects - individual investigations that often lead to a conference-style report (with coursework reports from the best students regularly accepted for conference presentation). MPhil students (and optionally Part III students) also attend an individually selected series of 'research skills' sessions that focus on practical skills necessary in an academic career, including publication, funding, collaboration, specific research methods, and even soldering, short-run manufacture and 3D prototyping for those expecting to carry out research with electronic or physical components.

Admission to the ACS degree is highly competitive, with 200-300 applications for around 50 places. The minimum standard for admission, as for continuation of undergraduates to the Part III year, is a Cambridge first-class degree or equivalent. Those admitted generally have undergraduate grades placing them in the top 5% of a cohort at a prestigious institution. It is not unusual for the ACS to be taken as a second Masters. Applicants reaching this standard are interviewed before admission, usually by the faculty member who is likely to supervise the research dissertation. Some load-balancing is necessary in response to greater popularity of fields (such as machine learning) where the proportion of applicants is very high in relation to research group size. The admissions panel sets higher standards in these areas, to ensure that research in popular areas continues to maintain high standards.

Results from the MPhil ACS and Part III are extremely high, with the majority of students achieving project results that could be published, with minimal modification, at a leading research venue in their sub-field of computer science. This is reflected in the high proportion of distinctions that are awarded – a fact that is almost always commented on, and enthusiastically supported in considering the standard of achievement, by the external examiners for these degrees.

In the past three years, the number of MPhil students has been 31, 32 and 58. Target size for the course is now between 40-50 students, and admission criteria for the next academic year have been adjusted accordingly. Over the past three years 46% of MPhil students have come from outside the EU, 30% from EU countries and 24% from the UK (noting that students from Cambridge are more likely to take the Part III route rather than applying for MPhil). Gender balance of the MPhil cohort is 79% male, 21% female. Distribution of research interests can be estimated by the group in which the dissertation supervisor is based. Over the past three years, 30 MPhil dissertations have been supervised in the NLIP group, 21 in the AI group, 19 in the Graphics and Interaction Group, 15 in the Security group, 13 in Programming Language and Semantics, ten in Systems Research, seven in Computer Architectures and six in the Digital Technology group.

PhD Training

PhD training in the Department follows the traditional "apprenticeship" model of PhDs in the UK, each student working closely with a personal supervisor, and contributing to the work of a research group and collaborating with other researchers as appropriate to the style of research undertaken by that supervisor. This traditional model provides the flexibility to accommodate a very wide range of approaches to computer science research, ranging from formal mathematics (often requiring intensive private study) to large-scale engineering development as a member of a larger team.

The diversity of research training in different subfields is supplemented by the programme of research skills development, managed across the Department, through which students must participate in a minimum number of focused skill sessions. Research skills training is offered to MPhil students, PhD students, and also post-doctoral research staff. PhD students in specialist fields are free to develop skills in other areas (including foreign language tuition) if they wish to do so.

The PhD programme includes formal assessment of progress at the end of each year. PhD students are not at first registered for a degree, and registration for the PhD is conditional on passing a viva examination based on a written first-year report including a report on progress in the first year and a research proposal. If the examiners judge that the student is unlikely to proceed successfully to PhD, an opportunity is provided for resubmission and re-examination. A small number of students who do not proceed to the PhD are awarded a certificate of postgraduate study, allowing a managed exit for those who find they are not suited to academic research. The second-year report is also formally examined, but is mainly intended as an opportunity for faculty members other than the supervisor to check the ambition and feasibility of plans for completion of the PhD. A third-year progress statement is required at the end of the ninth term which gives an overview of the student's work since submission of the second year report and a schedule for the next one to three terms' work, or, ideally, a draft of the dissertation. This report is considered by the Director of Graduate Education.

PhDs are examined following the standard UK model, with one internal and one external examiner conducting a viva examination. Provision for interdisciplinary research includes reasonably routine appointment of an internal examiner from another Cambridge department, and very few restrictions on the selection of the external examiner, who is almost always a recognised expert specialist in the topic.

These processes are overseen by a specialist Degree Committee, which reviews all admissions to the PhD, appoints examiners, receives and evaluates the reports of examiners, and formally recommends award of the PhD. The Degree Committee includes external members from other Cambridge Faculties, and also independent industry representatives, reflecting both research and technology product sectors. The high quality of PhD research in the Department is recognised by regular distinguished dissertation awards administered by the British Computer Society, and occasional other awards, such as the E.W. Beth Dissertation Prize.

There are presently 114 PhD students in the Department (28 first year pre-registration), and a further 21 arriving in the coming year. Seven of these are also Research Assistants. 41% of these are from outside the EU, 36% from EU countries and 23% from the UK. Gender balance is 78% male, 22% female. Distribution of students across the research groups in the Department is currently 17 students in the NLIP group, 28 in Systems Research, 13 in AI, 23 in

Programming Languages and Semantics, 13 in Digital Technology Group, 17 in Security, 13 in Computer Architectures, and eight in the Graphics and Interaction Group. Of the 114 current PhD students, six completed our Part III UG programme (representing 15% of the Part III classes over four years, and 5% of current PhD students) and 38 the MPhil (representing 28% of the MPhil classes over four years, 33% of current PhD students).

Opportunities for PhD study are largely driven by the availability of funding. Over the past eight years, we have admitted 257 students, who have been funded in the following ways: 85 from the UK Engineering and Physical Sciences Research Council (including 13 industrial CASE collaborations), 61 from Cambridge University resources (the Cambridge Trusts, Gates scholarships, College scholarships, and four from the Department's own funds), 68 from project grants or direct industry sponsorship, 24 from foreign governments and 19 self-funded.

Chapter 6: EQUALITY AND DIVERSITY

Gender Balance and Athena SWAN

The primary concern for Equality & Diversity (E&D) in Computer Science as a subject is the relatively low levels of representation of women. Of course, the University as a whole supports diversity in a number of ways, including a range of diversity networks for staff and students (https://www.equality.admin.cam.ac.uk/diversity-networks), particularly recruitment where diversity is explicitly monitored by HR. We are currently considering whether the Department could and should be playing a more active role in this regard. However, based on feedback to the Equality & Diversity Committee (EDC) through surveys and focus groups, the most critical problem that the Department faces, in common with other Computer Science departments, is still the low numbers of women. We have therefore focused on E&D matters through the vehicle of Athena SWAN as that emphasises women in STEM subjects.

The Department currently holds an Athena SWAN Bronze award and is now preparing its submission for the Silver Award, which is due in November 2018. The responsibility for preparing the submission falls to the Equality & Diversity Committee (EDC), a re-forming of the Self-Assessment Panel that was responsible for the initial Athena SWAN Bronze submission. The EDC meets at least four times annually, has representation from all staff groups and roles, as well as a range of individual circumstances (international and UK, male and female, differing caring responsibilities, and so on), plus representation from women@CL (see below). It has carried out regular staff and student surveys and focus groups to ensure a good flow of information about the challenges faced by and desires of Department members. The Department also provides a range of opportunities for formal and informal interactions between staff groups and students, including: termly social teas, a weekly Happy Hour (social gathering run by PhD students and post-doc staff and open to all), free small group running, Pilates classes, and research group lunches. These activities bring together all members of the Department, students, researchers and staff.

The EDC is particularly responsible for overseeing implementation of the Action Plan presented in the Bronze award, and for developing the next Action Plan as part of the Silver submission. Among the direct actions taken by the EDC in implementing the Bronze award Action Plan are:

- Rescheduling of weekly Department meetings to sit within core hours, to make it easier for those with caring responsibilities to attend.
- Introduction of Unconscious Bias training for all staff involved in recruitment and Department supervision.
- Offering a course for those in Colleges involved in undergraduate admissions, attempting to address some of the problems in recruiting women into computer science.
- Provision of mentoring for post-doctoral researchers through an appraisal process involving an academic other than the post-doc's PI.
- Increased prominence of women in recruitment materials, the Department's website, and at Department Open Days (from 21% in 2014 to 33% in 2018).
- Following a review of our Tripos, various studies which have shown that an excess of dry mathematical material in the early stages of computer science degree courses has the effect of turning women away. The Tripos was thus refreshed to provide a broader, better balanced introduction to the subject in the first year, and an increased number of options in later years. Given the current gender balance of our staff, this also allowed

us to increase the number of lecture-hours offered by women in the early stages of the Tripos, with the aim of addressing issues noted by the People Like Me campaign.

The Department supports the University's Graduated Return policy, actively promoting it to all staff, and providing for part-time/term-time working, job sharing, flexitime, and breastfeeding and refrigeration facilities. This has resulted in three successful applicants since its introduction in 2013 (one academic, two researchers). The Department also took the initiative to provide funding to support researchers whose grants have expired by the time of their return to work (four since 2013), which at least partially mitigates the impact on their careers. Since 2013, nine men have taken paternity leave of up to two weeks, one of whom took an extended period of unpaid leave. In this context it is also worth noting that the Department is flexible in how formally specified hours and days of work are handled: this flexibility accounts at least in part for the relatively small numbers of formal requests for flexible working. In practice, most computer scientists require little in the way of access to specialist equipment, so working from home and working remotely are both commonplace. While most staff therefore choose to work full-time, there has been a slow increase of the number of (male) staff choosing to work part-time (from 3% in 2014 to 18% in 2018). Of course, the Department is also cognisant of the need to ensure staff do take proper advantage of the various leaves and forms of flexible working offered and do not, for example, turn a period of maternity leave into an extended period of simply working from home.

Recruitment, Retention and Promotion

The Department is slowly but steadily improving its gender balance among both students and staff. While the Department has no direct control over undergraduate admissions, the actions of the EDC in raising awareness of the problem of gender balance in computer science, coupled with the Department's provision of admissions and unconscious bias training appears to have led to an improvement in proportion of female undergraduate applicants over the period 2013/14 to 2017/18 from 11% to 15% and of female undergraduate acceptances from 14% to 21%. We anticipate around 25% female students in 2018/19. The fee status of female applicants has also become more proportionate with respect to the overall cohort. Over the same period the total number of admitted students has grown from 86 to 105.

The numbers for the MPhil show a similar pattern, with a steady increase in the proportion of female offers and admissions from 2013/14 to 2016/17 (15% to 24% and 14% to 24%). These proportions fell back slightly in 2017/18 (to 20% for both offers and admissions) though that year was something of an outlier as applications doubled compared to 2016/17. The proportion of female PhD students has remained roughly constant at 18-22%, with absolute numbers remaining above 25. We continue to monitor this situation as we are placed below the HESA average (at 24% compared to HESA average of 27%).

The proportion of female research staff has also improved. In 2013/14 the proportions of female researchers and research fellows were 17% and 50% respectively, and in 2017/18 these had increased to 22% and 75% (Figure 6 shows the trends with respect to research staff overall). These numbers do fluctuate, however, due to certain subareas, particularly interdisciplinary areas, attracting significantly higher numbers of female applicants, meaning that availability or otherwise of targeted research funding can have a significant effect on numbers from one year to the next. PIs and Chairs of hiring committees are actively encouraged to be aware of these issues and to seek applications through as many channels as possible.

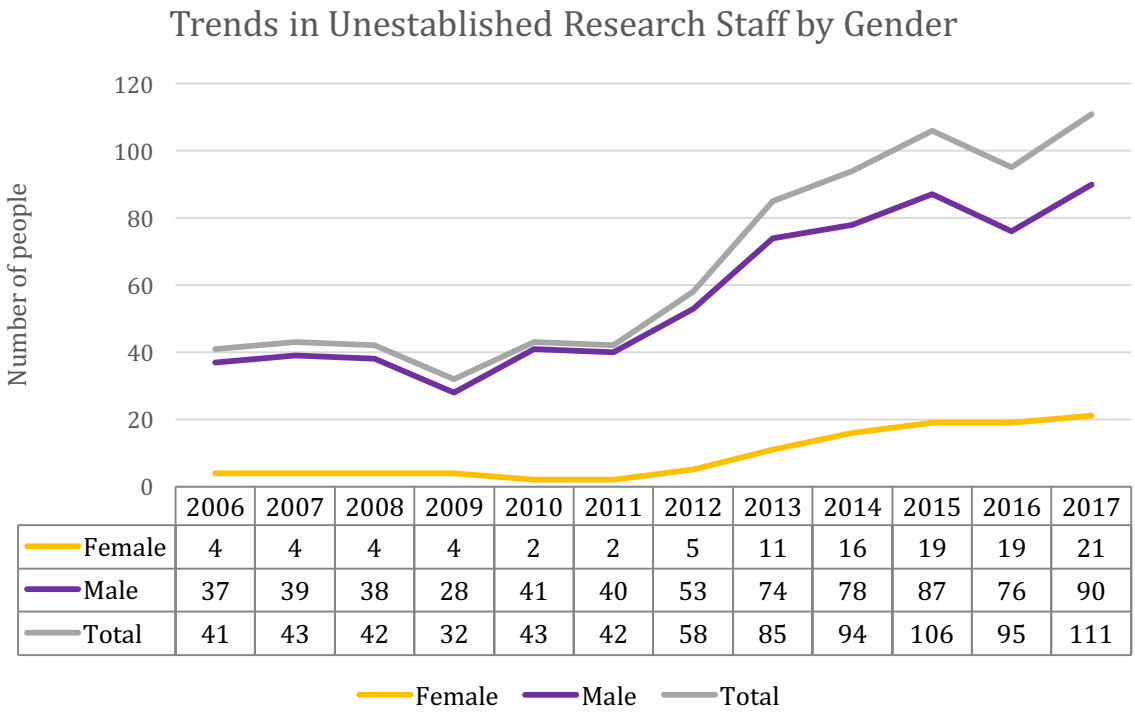

	2006	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017
Female	4	4	4	4	2	2	5	11	16	19	19	21
Male	37	39	38	28	41	40	53	74	78	87	76	90
Total	41	43	42	32	43	42	58	85	94	106	95	111

Figure 6: Trends in Research Staff by Gender.

Gender balance in academic staff in Computer Science is poor at most other UK institutions and we are unfortunately no exception. We thus committed to undertake several actions to improve the balance in our Department as part of our Bronze Action Plan, instituting pro-active searches for female applicants by the Chairs of selection panels. The result has been a significant improvement, from 11% female academic staff at the time of the Athena SWAN Bronze award submission in 2015, to 23% currently. Some improvement has taken place at all academic levels except Reader which has remained stable at 22% (Prof: 5% to 16%; SL: 10% to 20%; L: 0% to 29%). As we are expecting to continue hiring through to 2020, we will continue this practice which we hope will further improve the gender balance among academic staff. Figure 7 shows the temporal trend.

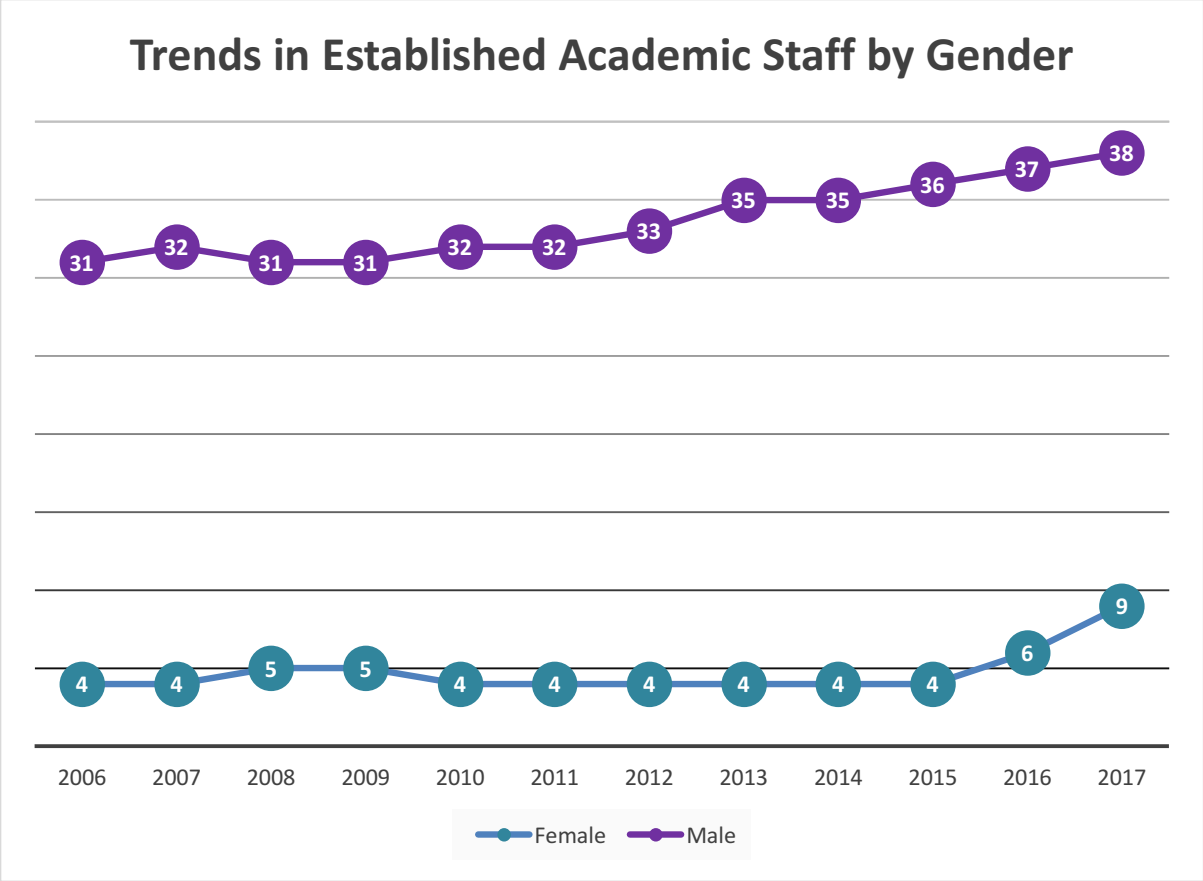

Figure 7: Temporal Trend for Academic Staff by Gender.

Staff retention is generally good, with one (male) academic departing in 2015/16 to return to his home country of New Zealand and another (male) academic leaving at the end of September 2018. The Department is generally highly successful in the Senior Academic Promotion process. From 2013-2017, 17 of 20 applications were successful, breaking down as 100% of five female applicants and 80% of 15 male applicants. All staff who are eligible to apply for promotion can discuss the opportunity at their appraisal, and if they wish to apply for promotion they are supported by the HoD and Departmental Secretary with their application. Our University Equality and Diversity Office offer a Senior Academic Promotions CV Scheme to encourage and support more female and male academics to apply for promotion within the University. Academics are made aware of this scheme at their Induction, probation and appraisal meetings. Almost all cases of promotion to Senior Research Associate are successful, with 21 successes in the period 2013/14 to 2017/18.

Unfortunately, the University continues to lack a mechanism for Professional and Support Staff promotion, although the Department Secretary does circulate details of the criteria and application process annually. All staff are encouraged to apply and support for staff contemplating applying for an increment is given through their Line Manager. Female staff have generally been more successful than male in their applications to this scheme (84% against 50% of a total of 39 applications).

women@CL

The women@CL initiative provides local, national and international activities for women engaged in computing research and academic leadership in the Department of Computer Science and Technology. It exists to help women aspire to leadership positions in both academia and in industry and to support them in their careers. The initiative is directed from the ground up by a staff and student executive committee, with the department providing administrative support and core funding. This year, the initiative has coordinated the following major activities:

- **Annual Conference:** Hosted the 5th Annual Oxbridge Women in Computer Science conference in the Lab, with sponsorship from Google. Attendance was robust from both universities, with a full day of research talks from across the seniority spectrum and a networking dinner in the evening.
- **Networking Events:** Several Cambridge dinners from industry sponsors for the purposes of recruitment and networking. This year ARM and Bloomberg continued their sponsorship of the dinners. There have also been termly coffee and cake gatherings in the Department for informal networking during the working day.
- **women@CL Talklet series:** There have been a series of talks from female academics on their research, hosted in the Department and advertised on talks.cam.ac.uk to encourage cross-departmental attendance where appropriate.

The women@CL budget has remained fairly static this year, with around £1,400 spent on the networking and talkets and the remainder funded directly by industrial sponsors. The executive committee is preparing the introduction of sponsorship tiers to formalise this direct funding and provide an expanded budget for future committees. This sponsorship aims to increase the working budget of the initiative to around £20-40,000 from external funds, in order to permit more ambitious projects in the future. Suggestions from the membership body include: providing some day care in the Department during school holidays, providing access to equality consultants to provide feedback on gendered text in job advertisements and group web pages, and a hardship fund for female students at all levels.

Chapter 7: IMPACT AND OUTREACH

The Department has a long tradition of generating commercial and social impacts from its research, and has been an important micro-engine of wealth creation in the UK. Staff and alumni have so far created over 270 companies, a list of which is maintained on the Departmental webpages. Recent examples of commercial and non-commercial impact include:

Unikernels

Unikernels are small, potentially transient computer modules specialised to undertake a single task at the point in time when it is needed. They are, in effect, the opposite of generic virtual machines - each one is designed to undertake a single task, and they are small, simple and quick, using just enough code to enable the relevant application or process to run (about 4% of a traditional operating system). They are likely to become increasingly used in applications where security and efficiency are vital, such as systems storing personal data and IoT applications. Researchers in the Department started restructuring VMs into flexible modular components in 2009, as part of the RCUK-funded MirageOS project. By the end of 2016, the unikernel technology arising from MirageOS was sufficiently advanced that the team, led by Dr. Anil Madhavapeddy, decided to found Unikernel Systems, which was acquired eight weeks later by San Francisco-based Docker Inc. to accelerate the development and broad adoption of the technology. The unikernel technology has subsequently been adopted in one of the fastest growing segments of the IT industry (containers), and ships in Docker for Mac and Windows that have tens of millions of users.

Raspberry Pi

In 2006, several members of the Department (Upton, Lang, Mullins, Mycroft) became concerned at the decline in numbers and skills of students applying for Computer Science. They foresaw how a tiny and affordable computer could potentially mitigate this, and started to produce what would become the Raspberry Pi. Several early prototypes were designed but were very limited by the high cost and low power processors for mobile devices at that time. In 2008, the team started a collaboration with Pete Lomas, MD of Norcott Technologies and David Braben, the co-author of the seminal BBC micro game Elite, and formed the Raspberry Pi Foundation[22]. Three years later, the Raspberry Pi Model B was born, which sold over two million units within two years of mass production. In early 2013 the organization split into two parts: the Raspberry Pi Foundation and Raspberry Pi (Trading) Ltd responsible for the engineering and trading activities. The Raspberry Pi Foundation is a UK registered charity with the objective of furthering the advancement of education of adults and children, particularly in the field of computers, computer science and related subjects. Raspberry Pi (Trading) Ltd is a wholly owned subsidiary of Foundation, with the money earned from sales of Raspberry Pi products being used to fund the Foundation's charitable work. Since the first Raspberry Pi went on sale on 29 February 2012, the Raspberry Pi Foundation has sold more than 15 million computers.

[22] https://www.raspberrypi.org

The HAT (Hub of All Things)

HAT started as a £1.2M multi-disciplinary project funded by RCUK's Digital Economy Programme (Crowcroft as Co-Investigator). It has evolved into the first ever personal data platform (HATPDP) created to trade and exchange individuals' own data for services in a standardised and structured manner. While there are many personal data lockers and repositories, the HAT has a schema that 'flattens' and 'liberates' vertical structures of data so that new mashups and new ways of putting together data for new services can be created to serve individual lives. Individuals can acquire their own data from internet-connected objects or services, and this acquired data is then transformed by the HATPDP to enable individuals to contextualise their own data, making it meaningful and useful for decision-making. With that data, individuals can buy apps that allow them to analyse, view, trade or make important decisions based on their own data for a smarter and more effective life. Perhaps most importantly, the HAT and its transformed data is owned by the individual.

The HAT is now run by two support organisations: The HAT Community Foundation (HCF) and the HAT Data Exchange (HATDeX). HCF is a non-profit organisation that promotes the use of private data accounts by individuals, start-ups, corporations, universities and governments. Recent figures from the HCF website (June 2018) state that there are 1100 HAT owners, 30 start-ups, 50 partners, ten Foundation Members and seven apps in the HAT ecosystem. HATDeX recently (June 2018) received £30M investment to fund at least 10m new HAT microservers, provisioning them for new apps and services, and growing the technology infrastructure.

Sonic Pi

Sonic Pi is a live coding environment that supports computing and music lessons in schools, and which can be used by professional musicians. It was originally borne out of a research project involving members of the Cambridge Computer Science Department, Sam Aaron, Dominic Orchard and Alan Blackwell, and the first release was in 2012.[23] In 2017, African and Finnish tech and education innovators collaborated to use Sonic Pi to deliver creative coding workshops engaging almost 2000 children in ten African countries, through the Codebus Africa initiative.[24]

The Isaac Physics Project

The Isaac Platform[25] has been mentioned earlier in this document. It uses recent developments in web technology and computer-based educational methods to improve physics and maths teaching in English schools, working in partnership with teachers and educators. Dr Beresford and Dr Rice from the Computer Lab lead the technical team and work in close collaboration with Dr Lisa Jardine-Wright who manages the overall project as well as the content and outreach teams in the Cavendish Laboratory. The platform has over 120,000 registered users who make tens of thousands of question attempts on our platform every day. The Isaac platform is funded by the Department for Education.

---

23 https://sonic-pi.net/

24 http://finland100africa.fi

25 https://isaacphysics.org

Consultancy and Professional Services

All academic staff are free to undertake consultancy work without limit, provided there is no negative impact on their academic commitments. Consultancy is regarded as a way of enhancing experience of industrial problems, building connections, and of transferring knowledge to industry and government. Of course, it also provides staff with a means of generating additional income, and the Department regards flexibility in this respect as a valuable retention tool. Some work is done pro bono (e.g. for HM Government), whilst paid consultancy may be done through Cambridge University Technical Services (CUTS) or privately, as the staff member chooses. The service provided by CUTS includes contract negotiation, invoicing, and professional liability insurance, in return for a small percentage fee. Some members of staff run their own companies to provide consultancy. Since 2008, Laboratory staff have carried out consultancy work for many companies, including (as indicative examples) Alcatel (Bell Labs), Apple, AOL, BBC, BT, CSR, France Telecom, Google, Intel, Juniper Networks, Microsoft, Netronome, and Orange.

In addition to consultancy, members of staff can take unpaid leave to work with companies. This is most common during sabbatical, but, with HoD approval, it may also be done in addition to sabbatical. This requires careful consideration, especially with respect to cover for teaching, PhD supervision and so on, but is highly important for staff retention and our impact agenda. Staff may also be employed by companies as well as in their academic positions, as long as they are able to carry out their academic job. This can be appropriate where the company employment is in a research lab, with the main condition being that a high percentage of the research can be published openly. Resources available in the commercial environment are normally a significant part of the attraction to academics, as well as the enhanced salary. Again, this situation requires careful monitoring, to ensure that other members of staff are not disadvantaged.

The Cambridge Ring

The Ring[26] is the Department's graduate association. It was founded in 2002 to enable graduates to get a lifetime benefit from their Cambridge degree, professionally, technically and socially. There are currently (Aug 2018) over 1000 members, the majority of whom work in industry, and who benefit from activities such as a mentoring scheme, careers advice, circulation of job adverts, formal and informal social events, and a service putting them in touch with others working in a particular sector of job function. They also receive three newsletters each year, providing updates on individual members, the Department, and articles from current staff and students. The Ring is also an important resource for entrepreneurial students, as it acts as an effective business club. At the annual Ring dinner, prizes are awarded to the company of the year and product of the year (eligibility, in both cases, requires a connection to the Department), and to the research publication of the year from Departmental staff and students. The five most recent "Company of the Year" awards went to Bromium, Improbable, Unikernel Systems, Swiftkey and DeepMind Technologies.

[26] https://www.cst.cam.ac.uk/ring

The Supporters' Club

The Supporters Club is a group of companies which actively support teaching or research in the Department, via networking and industrial opportunities and financial donations to the Department. These are a mix of local, national and international businesses, ranging from the very small, to major multinationals. Companies receive a slot at our annual recruitment fair, some give 'tech talks' to our staff and students, and can act as clients for our second year group projects. There is an annual dinner with academic staff and postgraduate students, and companies are able to post internships and job opportunities on a dedicated website. There are currently 82 member companies (August 2018), and the donations are used in direct support of our research students and other activities.

Outreach Activities

The Department's outreach activities focus on attracting the best students to apply to Cambridge, regardless of their backgrounds. The most visible events are the University's Open Days in early July where we provide talks, demonstrations and the opportunity to chat with faculty to several hundred potential students and their families. Beyond that, we focus on high-achieving potential students who come from disadvantaged backgrounds through our annual Sutton Trust Summer School[27] for Year 12 pupils from state-funded schools, which runs for a week in August. Here, students learn different kinds of algorithm, computer architecture and graphics, through lectures and backed up with significant amounts of practical experience in labs where they program on Raspberry Pis that they can then keep at the end of the school. This year we have expanded the work we do for this group of students by taking part in an Experience Cambridge day on a Saturday in July, providing taster sessions on similar topics. As well as this, we give ad hoc talks on Computer Science subjects to school pupils attending college-run events, such as 'How to build an iPhone' (Jones) for the Gonville and Caius' Norfolk Masterclass in March, or Downing's annual outreach residential in June, and a talk on affective computing (Gunes) to a Trinity Hall Women in STEM summer school in August. The Department's outreach activities are being reviewed so as to focus more on the groups from which we receive very low applicant numbers (e.g. women and black and ethnic minority students).

Links to the Local Economic Environment

Cambridge has a very high concentration of technology companies, known colloquially as "the Cambridge Phenomenon" or "Silicon Fen". Many of these were originally spin-outs from the Department's own staff or students, or from cognate departments in the University (e.g. Engineering, Physics), and ongoing engagement has generated very productive formal and informal collaborations. Examples include work with the Raspberry Pi Foundation (Mullins, Mycroft), ARM (Jones, Moore, Sewell), Bromium (Madhavapeddy), Cambridge Assessment (Briscoe), Cambridge Consultants (Blackwell), Linguamatics (Copestake), Microsoft Research (Blackwell), Nokia Bell Labs (Mascolo), RealVNC (Hopper), SolarFlare (Hopper, A Moore, Rice), and Ubisense (Hopper).

In addition, in recent years a number of industrial research laboratories have opened a site in Cambridge, starting with Microsoft Research, Toshiba and Philips Research and then, more

[27] https://summerschools.suttontrust.com/

recently, with Amazon Research Cambridge, Apple, Nokia Bell Labs and Samsung Research. Many undergraduate, postgraduate and PhD students as well as post-doctoral researchers have been recruited from the Department to work in these laboratories. Industrial research seminars are often advertised and attended by university researchers and departmental talks often see industrial researchers participating. Microsoft Research organises a Doctoral Students Summer schools for all their European funded students but also invites all the Department's first-year PhD students each year.

Open Source

Many members of the Department actively choose to release their research code as open source, in line with the Department's philosophy of enabling collaboration to be as free and flexible as possible. In the case of open source software, this is of great importance both for collaborations with companies and with other academic organisations. While the Department would never attempt to instruct an academic to release software as open source, it strongly defends the right of individual academics to do this and is prepared to push back on ill-informed objections that open source releases impede the University's attempts to monetise Intellectual Property. Xen is one notable example where the open source release changed the discipline. Other recent notable open source releases include the Cam3D corpus and CLM-Z face tracker (Robinson) which have both been made publicly available and are being used by research groups in other universities; the NetFPGA project (A Moore, with Stanford and others) which now has hundreds of users across institutions worldwide; OCaml Labs (Madhavapeddy), which is pushing forwards the development and uptake of OCaml and functional programming in general; and Databox,[28] which is developing a personal networked device to collate, curate and control access to personal data.

The Research Excellence Framework (REF)

The Department was one of 89 departments submitting to 'Unit of Assessment 11: Computer Science and Informatics' in the most recent Research Excellence Framework exercise (REF2014). Early in the preparation period, the Department took the active decision that REF assessment would not be allowed to influence departmental policy or management. In total, 56 staff submitted outputs to the REF, comprising 36 out of 39 established staff and a further 20 contract research staff. Two of the academic staff omitted worked mainly in teaching roles, the third had been seconded to the central University administration for several years, and no new staff were recruited to improve the Department's record. This contrasted with the approach adopted widely elsewhere and criticised in the Stern review of REF 2014.

37% of our submitted outputs (generally publications in conference proceedings or journals) were classified as world leading, 47% as internationally excellent, 15% as internationally recognised and 1% as nationally recognised. Six impact case studies were submitted, of which five were classified as world leading and the other as internationally excellent. 45% of the environment was classified as world leading and 55% as internationally excellent. The net effect was that the Department was ranked about sixth of 89 departments in its Unit of Assessment. However, the volume of our submission exceeded that of several larger

[28] www.databoxproject.uk

departments, and has resulted in substantial funding through the block grant to the University in the following years.

The goals articulated in our 2014 submission, and subsequent progress is as follows:

1. Expand the academic staff by four over the period 2013-2018 - We have recruited an additional twelve academic staff, in addition to three replacement appointments (Nada Amin to replace Mike Gordon, Andreas Vlachos to replace Stephen Clark and Rafal Mantiuk to replace Neil Dodgson). The new appointments both consolidate existing expertise and bridge the gaps between current disciplinary groups (more information is provided elsewhere in this document). All hires have substantial publication record which will support our REF2021 submission.

2. Expand support functions, particularly the finance office - We have made 1.5FTE additional appointments to the Finance Office, taking staff numbers to 4.5FTE, with the aim of ensuring sufficient support for costing grant applications as academic staff numbers increase (one of these posts is currently being appointed to, and exact roles and responsibilities are therefore being finalised).

3. Start planning towards a new building for 2020 onwards - plans have been drawn up based on an extension to our current building, and we are also considering alternatives such as a completely new building.

4. Smooth out the cluster of 2020 retirements by using reserves to recruit early to some of these posts - most of our additional hires have been funded as proleptic appointments. We have 10 scheduled academic retirements in the next five years.

5. Further diversify funding and increase industrial engagement – we have further increased the volume of funding from DARPA, and four additional staff are now in receipt of ERC awards (Paulson, Mantiuk, Sauerwald, Sewell). We have just announced a new collaboration with Nokia Bell Labs, who are funding a Centre for Mobile, Wearable Systems and Augmented Intelligence.[29] We have also recently announced the establishment of the DeepMind Professorship of Machine Learning, supported by an endowment from DeepMind. We are developing plans to pursue industrial funding more systematically - see the final section of the report.

The Department is now actively preparing for REF2021. A departmental committee has been formed consisting of six members of the academic staff (including two who served on the national panels in 2008 and 2014), the Research Facilitator, an external advisor (who served on the national panel in 2008), a University impact advisor and a part-time administrator. The committee will oversee preparation of the Department's submission, including assessment and selection of outputs and impact case studies. Initial indications (based on numbers of staff submitted) suggest that we will have to submit 160-170 research outputs and six impact case studies. Staff have already been asked to identify possible outputs for inclusion, resulting in a long list of 325 outputs for internal review. We have also begun identifying possible impact case studies, and have a current list of 18 for further exploration.

[29] http://www.cam.ac.uk/news/cambridge-and-nokia-bell-labs-establish-new-research-centre-to-advance-ai-supported-multi-sensory

An initial progress report was prepared in June 2018, a draft submission will be prepared for June 2019, and this will be revised in June 2020 for submission in November 2020. Our likely goals will be:

- To continue proleptic appointments for academic staff approaching retirement (including established Chairs), and making new appointments to increase the establishment to 50 academic staff
- To continue working towards an extension or new building elsewhere
- To manage the excessive workload of many staff
- To implement appropriate changes following the outcome of this Research Review

Chapter 8: HIGHLIGHTED CHALLENGES AND OPPORTUNITIES

Brexit

Brexit is likely to have a very negative effect on the Department, in common with most others in the UK. We have been very successful in obtaining ERC funding, and this will have to be replaced by other sources. Diversification of funding is important, but industry seems the most promising source, as discussed further below. EU students will be less likely to want to come to the UK, which could affect us severely, especially at PhD level. We will have to seek to make the Department more attractive to them. Staff recruitment will be affected, at all levels. We can seek to mitigate this by continuing a generous policy of funding visa expenses, but it may well become more difficult for EU staff to supplement their salary by consultancy, and it will be more difficult for partners to find work. Existing staff may leave if the UK climate continues to become more hostile and the local computer industry could decline. In principle, Cambridge is better placed than most UK universities to deal with this crisis, and could indeed use its resources to improve its competitive position relative to other universities. At Departmental level we see very little we can do other than to seek to enhance and better communicate what we see as our most positive aspects: collegiality, freedom to pursue individual goals and flexibility. Thus Brexit makes addressing other challenges more urgent and more difficult, but does not fundamentally change our approach.

Space

The Department had an extremely serious space problem in the 1980s and 1990s which is generally agreed to have stunted its development. Moving to the current building in 2001 allowed a rapid expansion, and considerable foresight was shown in creating a building that was much bigger than initially needed, though the Department's expansion into the additional space happened more rapidly than expected. For the last few years, we have been experiencing space shortages again, which have been partially alleviated by converting space into offices (parts of the library, alcoves and so on). However, we are approaching the limit of what can be reasonably done in that respect. An original plan for expansion when the building was designed was that the large lectures theatres could be converted to offices. The University initially proposed shared facilities, including large central lecture theatres on the West Cambridge site, but these did not materialise. In fact, the expected size of the 2018 undergraduate intake means that we are approaching the limits of the lecture theatre space and space for practicals (despite the latter being remodelled).

Extending the current building has been discussed over the last few years, and various outline plans drawn up. However, the University's Master plan for the West Cambridge site only allows a rather small extension to the north of the building, which would barely meet our expected needs in three years. Attempts to discuss further options have met with resistance from the central University presumably because they would require changing the outline planning permission. It might be possible to put another storey on the building, but this would be extremely expensive for the additional space provided and would also be very disruptive.

Our fundamental problem is that the outline plan for the site currently envisages the Department being hemmed in on all sides by buildings belonging to Engineering and Physics with no substantial space for expansion. While the aim that Information Engineering be located close to the Department is very welcome, the detailed status of these plans is currently very unclear. It is difficult to begin a serious attempt to raise funds for an extension in these circumstances. Even though we expect the Department to stop (or at least pause) its current expansion once it reaches 50 UTOs, many of our staff have ambitions to increase the size of their research teams and the numbers of PhD students, and that will not be possible without a substantial increase in space. We also do not want the Department's size to be limited in the long term, since it may well make sense to expand again in the future.

We additionally have the problem that the Department is located some distance from the main computer companies in Cambridge who we want to collaborate with, who are now predominantly located in the area around the station (CB1). While there are small amounts of space for start-ups on the West Cambridge site, and plans for more commercial space, none of those companies we have spoken to recently have much interest in being located here.

The current feeling of the HoD team is that we need a rethink, with a bolder and more ambitious plan than a simple extension. Ideally this would allow us to collaborate more easily with Information Engineering and with local companies. A site which allowed for the Department to be co-located with Information Engineering, small company research groups, start-ups and third-sector organisations would be very attractive and help fulfil the Department's philosophy of making interaction with industry as easy as possible. However, these ideas are currently in their infancy. Any extension or new building will require substantial funding which will presumably have to come from external sources.

Other Needs

Apart from space, the main need for additional funding that is felt across the whole of the Department is for PhD studentships. This need is felt by most departments in Cambridge. We currently have, on average, 2.5 PhD students for every UTO, although some of the students are actually supervised by Research Fellows, and the distribution of students is very uneven, both in terms of individual UTOs and research groups. Because some funders (notably EPSRC) do not fund PhD students on normal grants, a number of PhD students are actually employed as Research Assistants and pay their PhD fees (at the reduced rate for University employees) out of their salaries. This is somewhat unsatisfactory because they are subject to different conditions as employees from the other PhD students. Many good applicants cannot be accepted at all because of lack of funding, and this reduces the utility of the ACS to the Department, since some extremely good students end up accepting PhD places at other universities. The University is making PhD studentships a major focus of its new fund-raising campaign.

Funding From External Sources

The Department's record in obtaining research funding has been discussed extensively above. We have an extremely good track record with grant applications by individuals or small groups, and have diversified funding to a considerable extent, including attracting large grants from DARPA. We have been less successful at applications on an institutional scale. The

Department led two applications for EPSRC Centres for Doctoral Training this year, neither of which were successful. Two applications for CDTs in previous rounds also failed. It is unclear whether there is any systematic reason for this lack of success.

As previously discussed, we have recently obtained significant funding from donations from companies, including DeepMind and Nokia Bell Labs as well as longer-term supporters, such as ARM and Microsoft. Other companies have indicated their willingness to make substantial donations as part of a package where they are also funding specific projects on a contract basis. We intend to try and make this more systematic and we are investigating the possibility of supplementing the existing Supporters' Club with a more research-focused activity, targeted at a smaller number of large companies. One issue here is competition from other groups in the University who are seeking to do something similar.

There is also the potential for more funds to be raised from alumni. There is an agreement with colleges that individual alumni may not be approached by departments unless they are of high net worth, but it is possible to ask alumni collectively for donations, and this approach is being pursued by the Department of Chemistry, for instance, via their alumni newsletter. Anecdotally, some of our alumni would be very happy to be asked for funds to support the Department's activities, so we may investigate a very light touch approach to doing this. CUDAR is currently discussing the possibility of making approaches to suitable individual alumni.

University Funding

It has been a long-standing source of frustration to many in the University that very little of the income raised from overheads and course fees comes back to the departments. An example of this would be the fee income from ACS students (33 students in 2016-17, and 54 in the following year). The additional students were a considerable cost to the Department, directly in terms of additional payments to research staff who had to be employed to help with teaching, and indirectly in the additional time that was required by academic and support staff (which takes time away from grant applications and so on). However, almost none of the additional income came back into the Department. As far as the Department's own finances are concerned, the ACS is subsidised from other activities. The ACS income is, of course, a factor that can be used in making a case for additional academic or support staff, but the indirectness of the relationship means we cannot straightforwardly plan on that basis. Similar remarks apply to overheads on research income, very little of which comes into the Department.

Departmental Organisation

As described previously, the Department has several formal and semi-formal mechanisms to organise our activities. However, approaches which worked well for 25 UTOs are becoming substantially less effective as our numbers approach 50. Meetings are becoming too large for real discussion. The danger is that this creates a distance which means people start to perceive themselves as being managed by some third party, rather than being collectively involved in the activities of the Department. Load-balancing has become ineffective.

To a large extent, these problems are affecting teaching and administration more than research. The level of collaboration between research groups has been increasing because of decisions

by individuals rather than any top-down directive. Nevertheless, this does have implications for research activity if administration becomes ineffective or teaching takes up more time than it should. Over-emphasis on research groups as a means of general departmental organisation could inhibit expansion into new areas which must be consciously guarded against and research groups are not the right units for organisation of teaching.

Given the nature of the Department, an attempt to impose a new group structure top-down will not work. The current plan is to organise teaching by grouping courses rather than people, but this will require considerable discussion and consultation. Since this is not directly related to research, we do not discuss this further here.

Workload

Academics, especially those who combine teaching, research, research management and administrative roles, are juggling a range of complex demands. Many people feel they should excel at all of them. The pressures are constantly increasing and the chunks of time available to concentrate on any one activity are diminishing. Post-doctoral staff have fewer teaching and administrative responsibilities but are under acute and increasing pressure to publish, especially if they aim to obtain a permanent academic position. While the situation is somewhat different for academic-related staff, they also face increasing pressures and time-fragmentation. This is one of the most serious and insidious issues the Department faces and resists any simple solution. Better load-balancing would help, but can only alleviate the worst problems rather than address the trend. We want to make awareness of this central to decision-making, ensure that we think very carefully about actions that will increase workload, and seek to alleviate pressure from outside. This may mean not becoming involved in some initiatives and being careful not to chase funding which has very deleterious consequences for workload.

External Communication

There is a long-standing problem in explaining the nature of computer science, especially when discussing academic work. The situation may even be getting worse, with the public narrative becoming increasing negative in discussion of the giant computer technology companies and worries about AI. The new computer science curriculum in schools is also often seen negatively. That being the case, the Department has to try and present aspects of what we do that will interest and inspire outsiders, rather than seek to define CS as a whole.

Our website has been acknowledged to be a problem for years, and progress on replacing it has been slow. We are in the process of recruiting two members of staff to address the problem, one of whom will concentrate on technical issues and the other on content. Our aim is to produce a narrative about the Department that celebrates the fluidity and diversity of Computer Science. We also need to make the case more strongly for why university research is essential and what differentiates it from company research.

Recruitment Plans

Our recruitment plans for next year have been detailed above. If we are successful, this would take UTO numbers up to 48. We currently hope to recruit two additional UTOs in 2019-20

(one in machine learning or closely related area), which would leave us at the target of 50 UTOs. On this plan, subsequent recruitment would mainly be replacement of UTOs.

Competition with industry for staff is becoming intense, though it is probably affecting us more at the level of post-docs than UTOs (and even at PhD student level, as company research labs are now willing to hire MPhil students). This is an area we will have to look at. Although some of our post-docs do obtain faculty positions at other universities, there are far too few such positions for this to be a realistic prospect for most. Fewer people will be willing to accept a low salary in return for a small chance of a faculty position, especially when industrial labs are carrying out exciting research and paying vastly higher salaries. Most of the factors are outside our control as a Department, but we want to improve the working environment for post-doc staff, for instance by making sure that there is no pressure to help out with teaching, but that any teaching is properly rewarded, both financially and with formal recognition.

Appendix A: Academic Staff of Research Groups

Artificial Intelligence Group

John Daugman is Professor of Computer Vision and Pattern Recognition. Daugman received his degrees at Harvard University and then taught at Harvard before coming to Cambridge University. He has held the Johann Bernoulli Chair of Mathematics and Informatics at the University of Groningen, and the Toshiba Endowed Chair at the Tokyo Institute of Technology. His areas of research and teaching at Cambridge include Computer Vision, Information Theory, Neurocomputing and Statistical Pattern Recognition. Awards for his work in science and technology include the Information Technology Award and Medal of the British Computer Society, the "Time 100" Innovators Award, and the OBE, Order of the British Empire. He has been elected to Fellowships of: the Royal Academy of Engineering; the US National Academy of Inventors; the Institute of Mathematics and its Applications; the International Association for Pattern Recognition; and the British Computer Society. He was one of three finalists for the European Inventor of the Year Award, and he has been inducted into the US National Inventors Hall of Fame. He is the founder and benefactor of the Cambridge Chrysalis Trust.

Sean Holden is a Senior Lecturer in Machine Learning, and Fellow and Director of Studies in Computing at Trinity College. His research areas are Computational Learning Theory, Bayesian Inference, Machine Learning, and applications in Organelle Proteomics and Automated Theorem Proving. Sean obtained his PhD from the Signal Processing and Communications Group at the Cambridge University Engineering Department. He held post-doctoral positions in the Signals, Circuits and Systems Research Group at King's College London, and in the Speech, Vision and Robotics Group in Cambridge University Engineering Department. In his early work on computational learning theory, he was particularly interested in the sample complexity of error estimates such as cross-validation. He later moved on to more applied work, first in the application of machine learning to drug design, and more recently in the development of advanced data fusion methods for organelle proteomics. He is also currently involved in research aiming to incorporate machine learning into automated theorem proving systems. He has supervised numerous successful PhD projects covering a wide range of subject areas including Bayesian inference, drug design, kernel methods, and planning algorithms.

Mateja Jamnik is a Reader in Artificial Intelligence with a focus on developing AI techniques for human-like computing. She has recently served as a Specialist Adviser to the House of Lords Select Committee on Artificial Intelligence. Previously, she held an EPSRC Advanced Research Fellowship. Her work focuses on how people solve problems using informal techniques like diagrams, and she then models this type of reasoning on computers to enable machines to reason in a similar way to humans. She aims to develop research techniques and to build a research community in human-like computing in AI. Recently, she has begun to apply AI and reasoning techniques to medical data to advance personalised cancer medicine. At the start of the millennium, she was one of the founders of a new interdisciplinary research area and conference series "Diagrams", on the theory and application of diagrams. Jamnik's research bridges theoretical computer science (such as automated reasoning) and artificial intelligence, and has been supported by the UK Engineering and Physical Sciences Research Council (EPSRC), the Leverhulme Trust, the Mark Foundation, and the European Research Council. Mateja is an active supporter of women scientists and in 2003 founded a national network, women@CL, for women in computing research. In recognition of these contributions, Mateja was awarded the Athena Prize in 2016 by the Royal Society.

Pietro Lio' will be Professor in Computational Biology as of October 2018. He earned PhDs in Complex Systems and Non Linear Dynamics from the School of Informatics, Dept. of Engineering of the University of Firenze and (Theoretical) Genetics from the University of Pavia. His research focuses on using bioinformatics, computational biology models and machine learning to integrate various types of data (molecular and clinical, drugs, social and lifestyle) across different spatial and temporal scales of biological complexity. In the context of basic science, these approaches are effective in understanding the mechanisms and the dynamics of how biological elements build up properties, such as sensing the environment, conveying information, being programmable, communicating, and computing. In the context of biomedical fields, by integrating different layers of evidence, predictive models will improve the accuracy of diagnosis of complex diseases in the presence of other chronic and acute conditions; will identify effective markers for disease trajectory; and will suggest composition of treatments (drugs and lifestyle) before the manifestation of symptoms.

Thomas Sauerwald is a University Senior Lecturer (from 1 October 2018). He is a member of both the Artificial Intelligence group and the Programming, Language and Semantics group. Thomas obtained his PhD in Computer Science from the University of Paderborn, Germany. Before moving to Cambridge, he held post-doctoral positions at Berkeley, Vancouver and the Max Planck Institute for Informatics, where he was also appointed Senior Researcher and Junior Research Group Leader in 2012. In 2015, he was a visiting researcher at Microsoft Research in Cambridge. His research focuses on randomised algorithms and stochastic processes, especially random walks and Markov chains. Specific examples include distributed algorithms for load balancing, scheduling and information dissemination, streaming algorithms and graph partitioning. Currently his research is supported by an ERC Starting Grant, which is a five-year project from 2016 to 2021. The objectives of this project are to (i) improve our understanding of multiple random walks and (ii) apply these insights to design more efficient algorithms that can cope with static and dynamic graphs.

Digital Technology Group

Alastair Beresford is a Reader. His work examines the security and privacy of large-scale distributed computer systems, with a particular interest in the security and privacy of networked mobile and embedded devices. He looks at the security of the devices themselves as well as the security and privacy problems induced by the interaction between mobile devices and cloud-based Internet services. He takes a variety of approaches, including the critical evaluation of existing products, the design and construction of novel prototype technologies, and the measurement of human behaviour. Dr Beresford also work on large-scale online learning platforms, and in this context runs the Isaac Platform together with Dr Rice.

Robert Harle is a Senior Lecturer whose research domain is in cyber-physical systems, with a particular focus on sensing and data analysis. His research work typically involves mobile and wearable systems. He has published extensively in the fields of indoor positioning, context-awareness and ubiquitous computing.

Andy Hopper, Professor of Computer Technology, works in several connected areas spanning computer systems, sensor-driven and context-aware computing, and using computer technology to ensure sustainability of the planet. Professor Hopper has pursued academic and

industrial careers simultaneously (as mentioned in the Research Strategy statement above) and in the industrial context he has co-founded thirteen spin-outs and start-ups, three of which floated on stock markets, as well as working for multinational companies. In recent years the companies he co-founded have received five Queen's Awards for Enterprise. Previously Head of this Department from 2004, he is currently Treasurer and Vice-President of the Royal Society.

Andrew Moore (see Systems Research Group)

Amanda Prorok is a Lecturer. Her research focuses on robot networks and cyber-physical systems, including algorithms for coordination, control and planning, with applications to multi-vehicle systems, automated transport, and robot swarms. Prior to joining Cambridge in 2017, she was a post-doctoral Researcher in the General Robotics, Automation, Sensing and Perception (GRASP) Laboratory at the University of Pennsylvania, USA, where she worked with Professor Vijay Kumar on coordination and control algorithms for multi-robot systems. She completed her PhD at EPFL, Switzerland. Her dissertation was awarded the Asea Brown Boveri (ABB) award for the best thesis at EPFL in the fields of Computer Sciences, Automatics and Telecommunications. Further awards include Finalist for Best Multi-Robot Systems Paper at ICRA 2017, Best Paper at BICT 2015, and MIT Rising Stars 2015.

Andrew Rice, who is a Reader, works with Professor Hopper developing the Computing for the Future of the Planet framework which seeks to identify areas of computer science research that could benefit global issues. His research covers topics arising from this framework. His current research focus is into providing tools to support software developers arises from this. He is PI on the EPSRC-funded CamFort grant which seeks to develop techniques for scientists to use for improving the software quality of computational models written in Fortran. His previous work developed new techniques for measuring and understanding the power consumption of smartphones.

Ian Wassell is a Senior Lecturer with a long history of conducting research concerning various aspects of wireless communication systems, including radio propagation measurement and modelling, physical layer design and simulation, media access control for cognitive radio, and cooperative relay networks. He has also undertaken a number of successful projects in the area of wireless sensor networks (WSNs), primarily aimed at civil infrastructure monitoring. This work led to significant contributions to a new data sampling paradigm known as compressive sensing (CS) that eliminates the need for explicit source compression at the sensor node and lowering energy consumption. CS is based on the notion that most signals of interest are sparse, provided an appropriate dictionary is established. The concept of sparseness has also been applied to improving the performance of image and video coding, and has been extended to address image classification, e.g., face, task and activity recognition. Improving the performance of deep-learning-based machine learning via the use of sparse and other priors is a growing area of research activity. Dr Wassell is a Chartered Engineer and a member of the Institution of Engineering and Technology.

Computer Architecture Group

Timothy Jones is a Lecturer (Reader from 1 October 2018) at the Computer Laboratory. His research focuses on extracting the many different forms of parallelism from applications (e.g.

thread-level, data-level, memory-level) to increase performance and address energy-efficiency, reliability and security challenges within compilers, binary translators and microarchitectures. Jones plans to focus on security, reliability, and optimised binary translation in the future. For security, his group has already performed work into mitigations for the recent Spectre hardware vulnerabilities and he intends to continue to address other speculative side-channels within the processor, as well as providing hardware support for security monitoring. Reliability is an ongoing concern in many areas and his aim is to develop the architecture used for prefetching such that it can correct soft errors within the core, as well as providing software schemes that target existing hardware.

Andrew Moore (see Systems Research Group)

Simon Moore has been Professor of Computer Engineering since 2014, having been a faculty member since 1998. He has worked on the design of CPUs including architectures to provide hardware support for fine-grained access control by means of capabilities. He was a designer of the BERI experimental CPU used in capability research.

Robert Mullins is a Senior Lecturer. His research and teaching focuses on computer architecture and VLSI (chip) design. His current projects focus on the design of manycore processors, hardware accelerators for machine learning and open-source hardware.

Robert Watson (see Security Group)

Graphics and Interaction Research Group

Alan Blackwell, Professor of Interdisciplinary Design, approaches interaction and graphics from a design research perspective, with a particular emphasis on ensuring that technology research be properly situated within the arts, humanities and social sciences. This brings a critical concern with the social dynamics of interdisciplinary creation and application of new technologies, and also practice-based research in the creative arts, where software is viewed as a craft material. His technical background is in engineering, AI and cognitive neuroscience, but the work of his group has proven particularly influential in the design of novel programming languages (including visual, tangible, dataflow and grid languages) from a human-centred perspective, for example in end-user programming, arts and education.

Hatice Gunes is a Senior Lecturer. Her research develops novel computational methods for analysing and understanding human behaviour, social signals and affect from facial expressions, vocal nuances, body posture/ gesture, and physiological signals; and for modelling these phenomena for creating socio-emotionally intelligent games, assistive technologies, virtual agents and robotic systems. She is currently collaborating with colleagues from Cambridge Clinical Neurosciences, Autism Research Centre, and Department of Engineering's Machine Intelligence Laboratory as well as UK start-up companies in the context of InnovateUK and Small Business Research Initiative projects.

Rafał Mantiuk is a Senior Lecturer whose work investigates how the knowledge of human visual perception can be incorporated within computer graphics and imaging. His recent interests focus on designing imaging algorithms that adapt to human visual performance and viewing conditions in order to deliver the best images given limited resources, such as

bandwidth, computation time or display contrast. He contributed to early work on high dynamic range imaging, including quality metrics (HDR-VDP), video compression, tone-mapping and open source software (pfstools).

Peter Robinson, Professor of Computer Technology, works on new technologies to enhance communication between computers and their users, and new applications to exploit these technologies. The main focus for this is human-computer interaction, where he has been leading work for some years on the use of video and paper as part of the user interface. Recent work has included desk-size projected displays and tangible interfaces. He has led investigations of the inference of people's mental states from facial expressions, vocal nuances, body posture and gesture, and other physiological signals, and also considered the expression of emotions by robots and cartoon avatars. He has also pursued parallel research into healthcare applications including diagnostic and therapeutic systems, and general inclusive user interfaces.

Damon Wischik is a Lecturer with a research interest in the challenges of big data. His research explores how machine learning can change the way we model with data: how machine-learnt data descriptions can be used instead of classic equation-based modelling. This is promising for narrative data, e.g. travel records, where machine learning might discover a grammar of behaviour, which could be communicated to the investigator by automatically picking out representative "stories". He is also seeking to improve the experience of interactive data modelling, especially for narrative data which does not fit well into the tabular worldview of current tools.

Natural Language and Information Processing (Nlip) Research Group

Ted Briscoe is Professor of Computational Linguistics and the inaugural Director of the Alta Institute, which conducts research into automated language teaching and assessment. He has published over 100 research articles and three books in the areas of automated speech and language processing. He was senior editor of the journal Computer Speech and Language for 15 years, and is co-founder and CEO of iLexIR Ltd., a consultancy and technology provider specialising in language processing applications. In 2009, iLexIR spun-out Swiftkey, maker of the world's most popular predictive keyboard for smartphones. In 2013, he co-founded English Language iTutoring Ltd and is its chief scientist.

Paula Buttery is a Reader in Computing and Language with research interests in both computer applications (Natural Language Processing) and language cognition (Computational Psycholinguistics). Her research focuses on building Natural Language Processing tools that work with non-canonical forms of natural language (spoken language, learners, aphasics) and also with low resource languages (endangered languages, dialects). She is interested in both the automatic machine processing of non-canonical language and the cognitive processes underlying that language – understanding the cognitive aspects of language processing is essential if we are to build language tools that can be intuitive to users. She is the current Director of the Automated Language Teaching and Assessment (ALTA) project. Her work within ALTA focuses on the spoken language of learners of English; but she is also funded to work on under-resourced endangered languages as part of the Cambridge Africa programme which partners Cambridge researchers with colleagues in Uganda and Ghana. Other research interests include the computational modelling of first and second language acquisition and language evolution.

Stephen Clark is a Reader in Natural Language Processing. He has been on academic leave since October 2016 working at DeepMind and will formally leave the Department on September 30, 2018.

Ann Copestake is Professor of Computational Linguistics and also the current Head of Department and a co-director of the Cambridge Language Sciences Interdisciplinary Research Centre. Her research involves developing computer models of human languages (or, more precisely, models of some aspects of human languages). In conjunction with DELPH-IN, an informal international consortium, she has developed software which has been used to develop formal computational accounts of the syntax and compositional semantics of many different languages. Her current research mainly concerns the development of models of semantics which are compatible with broad-coverage computational processing (parsing and generation). She is also interested in the formal aspects of combining distributional semantics with model theoretic accounts and in utilising DELPH-IN technology to establish the performance of deep learning systems according to linguistic criteria.

Simone Teufel is Professor of Information and Language. Her research interests are the study of (human or automatic) text comprehension, in particular how observable surface linguistic indications contribute to this process. She tries to find explanatory models of how a human or machine might rationally integrate linguistic indications into their representation of the meaning of the entire text. She also develops models of discourse structure and argumentation in scientific text, and is interested in how "folk logic" connects to linguistic expressions, which is particularly of interest for argument mining.

Andreas Vlachos and Ryan Cotterell will join the NLIP group in October 2018 at Senior Lecturer and Lecturer level respectively. Both experts in machine learning, Andreas' main research focus is automated fact-checking and imitation learning; Ryan's work focuses on deep learning and statistical approaches to phonology, morphology, linguistic typology and low-resource languages. Both Andreas and Ryan have been internationally recognised by the field with awards at major conferences.

Programming, Logic, and Semantics Group

Nada Amin is a University Lecturer who is interested in programming languages, including type systems, generative programming, relational programming, domain-specific languages, neuro-symbolic reasoning, deriving compilers, and applications to machine learning, biology and music.

Anuj Dawar is a Professor of Logic and Algorithms who is interested in finite model theory and its connection to the study of computational complexity; the theory of databases; the complexity of games; and the expressive power of logical formalisms.

Marcelo Fiore, Professor of Mathematical Foundations of Computer Science, is interested in mathematical models of computation; semantics of programming languages and concurrent systems; metalanguages, type systems, and program logics; and applications of category theory to computer science.

Tim Griffin (see Systems Research Group)

Mateja Jamnik (see Artificial Intelligence Group)

Neel Krishnaswami is a University Lecturer and interested in the applications of proof theory, type theory, and denotational semantics to programming language design and program verification.

Alan Mycroft, Professor of Computing, is interested in programming languages, type systems, program analysis and compilation, especially techniques bridging the theory-systems divide. Mycroft was co-founder of the Raspberry Pi Foundation for which he was also a Trustee 2008-2015.

Larry Paulson (FRS, FACM) is Professor of Computational Logic. Interested in automated theorem proving and verification, cryptographic protocols, and the mechanisation of mathematics. He is the originator and one of the primary developers of the Isabelle theorem prover, as well as MetiTarsk, an automatic theorem prover based on a combination of resolution and a decision procedure for the theory of real closed fields.

Andrew Pitts (FACM), Professor of Theoretical Computer Science, is interested in the applications of logic, type theory and category theory to the foundations of programming languages and their semantics and to computer-aided proof.

Martin Richards is a University Senior Lecturer (retired). Interested in typeless languages; compact byte-stream interpretive codes; machine-independent operating systems for process control and real-time applications; and real-time analysis of digital musical sound for use in an automatic accompanist system.

Thomas Sauerwald (See Artificial Intelligence Group)

Peter Sewell, Professor of Computer Science, is interested in the development of mathematically rigorous techniques that let us put the development of real systems on more solid foundations, to make them more robust and secure. This applied semantics work involves programming languages, computer architecture, security, network protocols, and verification. Peter Sewell has been awarded an ERC Advanced Grant (ELVER, 2018-2023).

Glynn Winskel, Professor of Computer Science, is interested in theoretical computer science and related mathematics, especially applications of logic and category theory to computer science, over a variety of areas from the foundations of semantics to systems biology.

Security Group

Ross Anderson has been Professor of Security Engineering in the lab since 2002, having been a faculty member since 1995. He has worked on cryptographic algorithms and protocols, hardware security, information hiding, and applications ranging from payment systems to the safety and privacy of healthcare IT. He pioneered the study of the economics of information security and is principal investigator of the Cambridge Cybercrime Centre, started in 2015.

Alastair Beresford (See Digital Technology Group)

Alice Hutchings has been appointed Lecturer from October 2018, having been a Senior Research Associate since 2014. She studies cybercrime, and has a particular interest in understanding cybercrime offenders: how they get into crime, and how they stop. She has worked on travel fraud, accommodation fraud, and various aspects of the prevention and disruption of online crime. She maintains the Cambridge Computer Crime database.

Markus Kuhn has been Senior Lecturer since 2010, having been a faculty member since 2001. He has worked on the hardware and signal-processing aspects of security including physical tamper-resistance, emissions security and side-channel attacks; he has also worked on distance bounding protocols, data compression and on satellite applications from pay-TV security through the protection of GNSS navigation systems.

Andrew Moore (see Systems Research Group)

Simon Moore (see Computer Architecture Group)

Frank Stajano has been Professor of Security and Privacy in the lab since 2017, having been appointed a faculty member in 2000, originally from the Engineering Department. He wrote *Security for Ubiquitous Computing* (Wiley, 2002). He has worked on cryptographic protocols, sensor networks, ubiquitous computing and human factors in security. He invented Pico, an authentication system to replace passwords, founded the national Inter-ACE and international Cambridge2Cambridge cyber security competitions and heads the university's Academic Centre of Excellence in Cyber Security Research.

Robert Watson has been Senior Lecturer since 2017, having been a faculty member since 2014. He leads several research projects spanning computer architecture, compilers, operating systems, networking and security. To this end he is also a member of the Security Group. He has led development of the security-focused capability-based CHERI Instruction-Set Architecture (ISA) as PI and Co-I of four DARPA contracts since 2010. He is also PI on a DARPA project in local and distributed operating-system tracing. Prior to coming to Cambridge, he was a Principal Research Scientist and DARPA PI at a series of industrial research labs, developing the kernel access-control framework now used for sandboxing in FreeBSD, Mac OS X, iOS, Junos, and other systems. He is a director (and past president) of the FreeBSD Foundation, and co-author of the Design and Implementation of the FreeBSD Operating System (Second Edition).

Systems Research Group

Jon Crowcroft has been the Marconi Professor of Communications Systems in the Computer Laboratory since October 2001. He worked on Internet support for multimedia communications for two decades. Three main topics of interest have been scalable multicast routing, practical approaches to traffic management, and the design of useful end-to-end protocols. Nowadays, he works on mobile, opportunistic, social, low energy, privacy preserving systems. Crowcroft intends to continue pushing towards re-decentralisation of the internet, web, cloud as well as to try to characterise and counteract social media abuse by humans and bots. He is also interested in continuing scaling network control as just another use of distributed computing and use of typesafe language techniques to secure computation on

simpler new hardware. Finally he is interested in deployment of communication infrastructure in challenging areas of the world.

Richard Gibbens is a Professor in Network Modelling within the group. His research interests are in the modelling of flows in congested systems with applications to communication, road and most recently energy networks. He lectures courses on Data Science and Computer Systems Modelling as well as supervising a broad range of courses in mathematics and computer science. Gibbens is a Fellow of the Alan Turing Institute where he has a funded project modelling the distributed control of resources in a charging network for electric vehicles. In other collaborative work which began with an EPSRC-funded project he is researching the use of storage in electric grid networks to enhance robustness, demand response and balancing in the presence of renewable energy sources such as wind or solar power generation. Very sadly, Professor Gibbens died in August 2018.

David Greaves is a Senior University Lecturer undertaking research in system specification with emphasis on interconnection, networking and component assembly. He has worked on tools for hardware RTL synthesis using model checkers and automated provers and is now applying these techniques to more general software systems.

Tim Griffin is a Reader in the group. Prior to joining the Department in 2005, Tim had been a researcher with Intel Research, AT&T Research, and Bell Laboratories. Tim's research is currently focused on applying rigorous modelling and analysis methods to problems of network design and network protocol design, especially Internet routing protocols. Dr Griffin plans to take the lessons learned in the analysis of internet routing to the wider community working on other kinds of optimisation problems. In particular, his results suggest that there is new, largely unexplored, class of “constrained shortest-path” problems that can be solved efficiently. He anticipates applications that are far removed from network routing.

Eva Kalyvianaki has been a Senior University Lecturer in the group since October 2017. Her research work spans the areas of cloud computing, big data processing, autonomic computing, distributed systems and systems research in general. She is interested in the design and management of next generation, large-scale applications in the cloud. Her current work involves real-time resource management of data centres. She is also going to be an Alan Turing Fellow from October 2018. Dr Kalyvianaki’s future research plans involve bringing the gap between application design and development and run-time performance management. Additional projects include designing and managing high-performance big data processing infrastructures with cost-, delay- and data shedding-related considerations for flexible processing.

Anil Madhavapeddy is a University Lecturer. His research interests are at the intersection of operating systems and programming languages, and he runs the OCaml Labs initiative across the Systems Research Group and Programming Languages and Semantics group. He is an active open-source developer, most recently on the MirageOS unikernel, the Xen hypervisor, the OpenBSD operating system and the OCaml programming language. Dr Madhavapeddy is working towards a “universal” programming platform for programming heterogeneous execution environments – from FPGAs to unikernels to JavaScript. This uses OCaml as a basis, but is a next-generation language with a formal specification and rewritten type system at its core. Once bootstrapped, this robust core will be the basis for the next generation of distributed and parallel high performance computing, for example for AI and machine learning, but with

a strong specification to ensure that code can survive more generationally in the face of hardware and operating system evolution.

Cecilia Mascolo is Professor of Mobile Systems. Her research concentrates on devising techniques to make mobile devices more efficient in their job of understanding human behaviour. This includes research on use of local computational units for data inference as well as better data modelling. She has also devised novel analytics techniques for mobile sensing datasets, exhibiting spatio-temporal characteristics. Mascolo's focus on applications includes aspects of mobile health and urban computing with her research concentrating more and more on the efficient use of on device computing for data analytics and inference, with special focus on mobile health applications: her current projects include neuroscientists, cardiologists, chest consultants and health psychologists looking at how devices and sensing can improve diagnostics and monitoring of diseases. The challenges to unlock mobile and wearable systems technology in practical applications are an important focus of her research which will encompass embedded machine learning with a special focus on acoustic sensing.

Andrew Moore is Reader in Systems. Driven by a fascination in the behaviour of complex systems, significant past contributions have ranged from novel, high-performance, network measurement and monitoring, to the first use of machine-learning for classification of network-use. Utilizing a system-wide interest: from application-software down to physical-layer coding and packet-based networks, Andrew has contributed to a number of open-source projects and as current director of the NetFPGA project continues to steward this early success in open-source software/hardware projects NetFPGA work has led to project-outcomes including open-source testing tools built upon the NetFPGA platform, alongside new generations of NetFPGA hardware and software. Now with a focus on latency-metrics and latency-reduction, as director of the Low Latency Laboratory he will investigate this topic further.

Richard Mortier is a Reader in Computing and Human-Data Interaction. Prior to re-joining the SRG he worked at Sprint ATL, Microsoft Research and Horizon Digital Economy Research at the University of Nottingham, as well as founding a start-up. His interests are at the intersection of Computing and Human-Data Interaction – how do the needs of interaction affect the systems we design, and how do the systems we build affect the interactions they can support. At present his focus topics are personal data processing, distributed machine learning, and unikernels. Mortier's research is increasingly aligned to the Human-Data Interaction agenda, exploring ways to design and build data processing and communication systems that enable dynamic yet legible control over compute resources and interactions around data.

Robert Watson (See Security Group)

Appendix B: Organisational Chart (correct as of 5th September 2018)

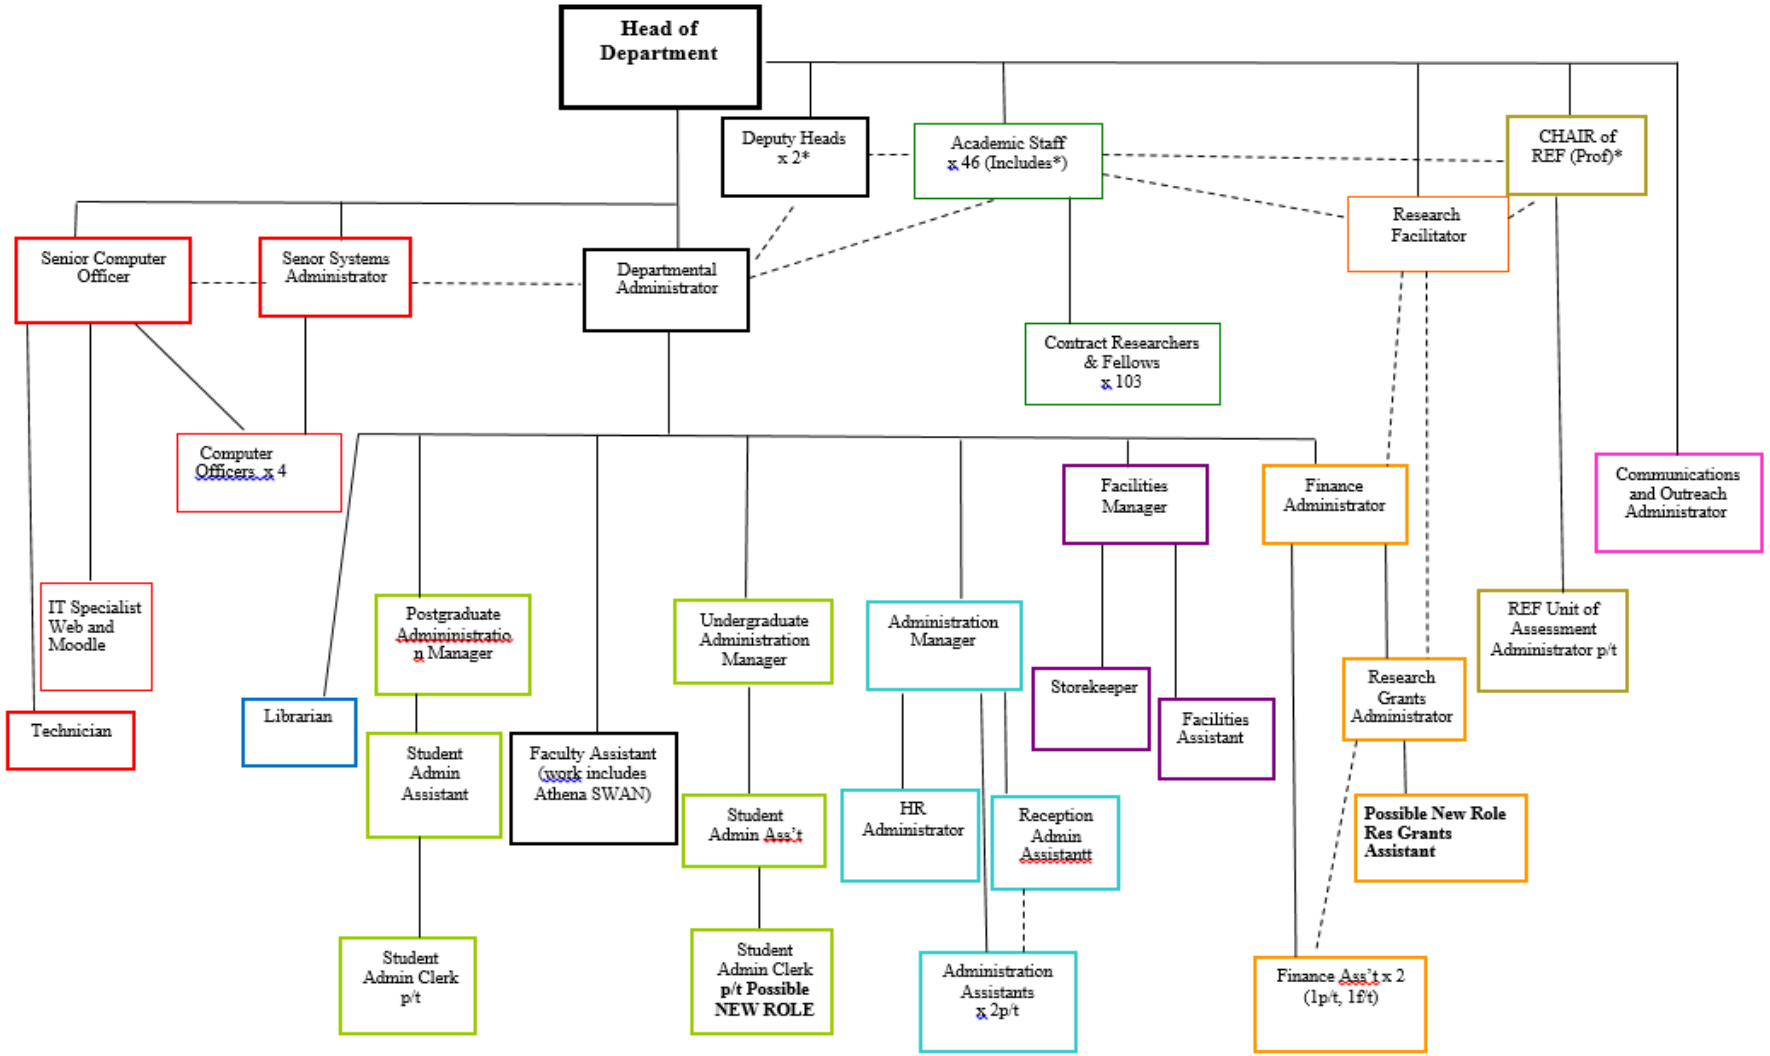

Appendix C: Research Income Distribution at the University of Cambridge

The University of Cambridge operates an income allocation model to distribute the funding received from external research sponsors between the Department/Faculty hosting the project and the University's central finances (the "Chest"). The Chest re-distributes the income retained centrally for research, via the University's six Schools, to Faculties and Departments and to central functions such as the University Library and central administration services. The distribution of Chest funds is based on the University's financial plans that are refreshed on an annual basis, drawing in turn upon the forward plans prepared by the Schools and the non-School institutions. The income allocation policy has recently been revised and the new policy is applied to all research grants activated after 1st August 2016.

Key principles of the Income Allocation Policy at the University of Cambridge are as follows:

- Funding received from the funder for directly incurred costs (e.g. salaries of research staff employed on the grant, consumables, travel and subsistence) are returned to the department in accordance with the funders' policies (e.g. 80% for UKRI, 100% for medical charities)

- Directly allocated costs provided by the funder to contribute to the salaries of investigators who are Chest-funded (e.g. academic Principal Investigator and Co-Investigator) are distributed between the Chest and department/faculty on a ratio of 80:20. There are two exceptions to this policy as follows: (i) the salaries of Chest-funded European Research Council Fellows and Principal Investigators of Horizon 2020 grants; and (ii) directly allocated investigator salaries recovered on industry-funded contracts are distributed between the Chest and department/faculty on a ratio of 50:50.

- The overhead recovered on research grants (e.g. funds received for estates and indirect costs on RCUK grants and the 25% of direct cost overhead recovered under Horizon 2020) are first used to make up any shortfall on direct costs, and the remainder is distributed between the Chest and faculty/department on a ratio of 80:20.

- The overhead recovered in excess of the direct costs of industry-funded contracts and agreements are distributed between the Chest and faculty/department on a ratio of 50:50. The policy also sets an expectation that all industry contracts should be expected to achieve an equivalent return of a minimum of 100% FEC.

The wider financial context for income allocation from research at Cambridge is relatively complex. The University is a collegiate institution in which the Colleges are independent educational charities that fall outside the financial planning of the University, other than as recipients of a share of the fees charged to undergraduate and postgraduate students. Nonetheless, the University benefits from the activities of the Colleges, for example through their investment in College Teaching Officers (CTOs) who support the work of the faculties, and from College support for research posts, including the cohort of Junior Research Fellows (JRFs). The

University also benefits from the income from its substantial endowment and the income streams available to the University from Cambridge University Press and Cambridge Assessment. However, there are also significant financial pressures on the University. In line with UK government policy the University of Cambridge is able to charge £9,250 for undergraduate fees from 2017/18[30], but the University's total costs for providing an undergraduate degree are significantly in excess of this.[31] The University is also constrained by the fact that it is only able to recover of the order of 85% of the full economic costs of the University's research activities from its funders. Combined, this wider financial context limits the flexibility of the University, both centrally and within Faculties, to invest in research activities.

[30] UK students comprise 87% of the University's undergraduate population (2015/16).

[31] The latest calculation for the average cost per student stands at £18,000/year (*Cambridge Reporter*, 13th July 2016)

Appendix D: The Research Excellence Framework 2021 – What We Know So Far (April 2018)

Initial Decisions on the Research Excellence Framework 2021[32]

The Research Excellence Framework is designed to secure the continuation of a world-class, dynamic and responsive research base across the full academic spectrum within UK higher education and:

- To provide accountability for public investment in research and produce evidence of the benefits of this investment
- To provide benchmarking information and establish reputational yardsticks, for use within the
  HE sector and for public information
- To inform the selective allocation of quality-related funding for research (for 2017-18, this equates to approximately £1.6 billion per year)
- To provide a rich evidence base to inform strategic decisions about national research priorities
- To create a strong performance incentive for HEIs and individual researchers
- To inform decisions on resource allocation by individual HEIs and other bodies

The model builds on previous similar exercises, including the Research Excellence Framework of 2014 and the Research Assessment Exercises which took place between 1986 and 2008.

Higher Education Institutions will be invited to return submissions in 34 broad discipline areas called Units of Assessment (UoAs). Returns will be submitted in late 2020. Expert sub-panels for each UoA will carry out the assessment, working under the leadership and guidance of four main panels. Results will likely be published in late 2021.

[32] Full text available at http://www.hefce.ac.uk/rsrch/ref2021/

For each submission, three elements will be assessed, with each element weighted to form a proportion of the 'overall excellence' assessment:

Element	Weighting	Key facts
Outputs	60%	• Assessment period for outputs will be January 2014 to December 2020. • Will be assessed through expert review, supported by citation data and other metrics where appropriate • Average number of outputs required per submitted researcher is to be determined • Precise conditions for researcher eligibility and for the portability (or otherwise) of outputs between institutions are to be determined
Impact	25%	• Includes the impact of research on the economy, policy, culture, the environment or quality of life outside the academic sphere, as well as certain types of impact within academia (such as impact on teaching) • Must be linked to research undertaken at the submitting institution between January 2000 and December 2020. The impact itself must have occurred between August 2013 and July 2020. • Will be assessed via a set of case studies, proportionate in number to the volume of research conducted in the submitting unit
Environment	15%	• Includes appraisal of academic impact, resourcing (including staffing, research income, infrastructure and facilities) and management (including strategy, staff development and training of postgraduate researchers) • Will be assessed through data submission (such as research funding awarded and doctoral degrees awarded) and through a statement (with sections on both the institutional and more localised environment) per Unit of Assessment. • An additional statement regarding research impact will also be required and will contribute to the score in this section.

Assessment Criteria and Level Definitions[33]

The ratings for each element are likely to be determined using the following criteria:

	Outputs	Impact	Environment
Criteria	Originality, significance and rigour	Reach and significance	Vitality and sustainability
4*	World-leading	Outstanding	Conducive to world-leading research
3*	Internationally excellent	Very considerable	Conducive to internationally excellent research
2*	Recognised internationally	Considerable	Conducive to internationally recognised research
1*	Recognised nationally	Recognised but modest	Conducive to nationally recognised research
U	Below these standards or not a research output	Little or no impact, or ineligible	Not conducive to research of these standards

[33] Sourced from *Panel Criteria and Working Methods* document (REF 2014)

Summary of Key Known Changes for Research Excellence Framework 2021 (c.f. REF 2014)

Element	Notes
Staff	• All staff with "significant responsibility for research" will be included in an institution's submission to a Unit of Assessment

Environment	• The environment statement template is to be more structured and will include more explicit sections on:
	  ○ The submitting unit's approach to interdisciplinary research
	  ○ Support for collaboration beyond the sector
	  ○ Open Research (including Open Access) strategy
	  ○ Equality and diversity

Impact	• Additional guidance will be provided on:
	  ○ The criteria which determine the "reach and significance" of impact
	  ○ Impact arising from public engagement
	• Eligibility of impact on teaching is to be broadened, to include impact within the submitting institution.
	• Impact which has arisen from a body of work (rather than from a single output) will be eligible.
	• The impact statement template will be assessed within the 'Environment' section

Units of Assessment	• 'Engineering' is to be a singular Unit of Assessment (having originally been sub-divided into four UoAs). • 'Geography and Environmental Studies' and 'Archaeology' are to be separate UoAs (having originally been combined). • 'Film and Screen Studies' has been added to the title of the 'Music, Drama, Dance and Performing Arts' UoA.

Pilots	• Institutional level assessment will be piloted, the outcomes of which will not be included in the results of the exercise. This will include:
	  ○ An overall assessment of each institution's environment. It will be made by a cross-disciplinary panel and will draw on institutional-level information (supplied within UoA-level environment statements)
	  ○ Assessment of institution-level impact case studies, to showcase impact arising from more interdisciplinary research and from institution-wide activity

Appendix E – Data Parameters and Caveats

These pertain to the graphs shown in the Papers

Funding Data

A dataset was supplied by the University's Planning and Resource Allocation Office (PRAO) in August 2018. It comprised a list of all research grant applications and secured grants which dated from financial year 2012/2013, i.e. 1 October 2012 onwards, to the present (derived from information held in the University's X5 and pFACT systems) and either:

- were attributed to an individual PI who was an employee of the Department of Department of Computer Science and Technology during that time up to 1 August 2018 (determined by the RSO and the Research and Partnerships Facilitator and the Departmental Administrator of the Department of Computer Science and Technology)

  or

- were attributed to the Department of Computer Science and Technology or, as it was previously named, the Computer Lab, during this period.

Philanthropic donations for the last five years were supplied in two datasets. Donations below £100,000 were provided by year by the Department. Donations above £100,000 were requested by the RSO from the Cambridge University Development and Alumni Relations Office (CUDAR). Both datasets were combined and used for funding charts in line with GDPR.

All datasets were checked manually for accuracy and edited to produce a clean dataset used in the analysis, as follows:

- Wrongly assigned and attributed grants were removed from the data.
- Grants held by PIs whose main affiliation is to another department were removed.
- Grants transferred to another institution without being held in the Department were removed.
- Abandoned duplicate applications to the same funder with matching grant numbers, dates and titles were removed.
- Donations mistakenly classified as research grants were removed and consolidated with the separately provided datasets of philanthropic donations.
- Applications for internal University of Cambridge funding were removed because an accurate and complete record of these bids cannot be guaranteed.
- Applications to Leverhulme Early Career Fellowships were combined with their corresponding applications for matched funding to the Isaac Newton Trust.
- Unsuccessful applications for phased projects were combined to comprise to whole project.
- Successful applications for phased projects the awarded money of which matches the combined total, or, if still active, the amount applied for to date were combined.
- Approved extensions to successful applications matching the Research Grant (RG) number, PI and money applied for were combined.

- Awards spanning several years with phased applications and/or extensions where awarded funds did not match the combined or even partial total applied for were checked in the RCO database and, when necessary, queried with the Research Operations Office (ROO) to find the correct amounts to the best of our knowledge.

HR Staff and Student Data

Records for all staff labelled as employed in the Department of Computer Science and Technology were obtained from the University's HR database and checked for accuracy by the Research Strategy Office and the Research and Partnerships Facilitator and Departmental Administrator at the Department of Computer Science and Technology. An updated spreadsheet was provided by the Department and both datasets were manually compared and consolidated to reach a high level of accuracy.

Staff data was correct on 31/07/2018.

Student data was supplied by the University's Planning and Resource Allocation Office (PRAO) in June 2018 and checked for accuracy by the Research Strategy Office. Corrections were made manually by comparing the dataset to official data published by the Cambridge Admissions Office. Further minor corrections to student numbers were made on 05/09/2018 after consulting the Research and Partnerships Facilitator and the Director of Research at the Department.

Student data was correct on 05/09/2018.