

**Institution:** The University of Edinburgh

**Unit of Assessment:** B11 – Computer Science and Informatics

### A. Overview

Informatics is a multi-disciplinary science. Its reach extends beyond the confines of traditional computer science to embrace linguistics, cognitive science, neuroscience, medicine, biology, robotics, physics, and social sciences. Our uniqueness comes from our breadth of view. Research in the School of Informatics deepens the foundational theory connecting sub-disciplines. The School has six research institutes, which promote research in broad areas of investigation. The **Laboratory for Foundations of Computer Science (LFCS)** pursues research into the mathematical foundations of computation in all its forms. The **Institute for Computing Systems Architecture (ICSA)** innovates in programming language compilation, system architecture and engineering. Our **Institute for Perception, Action and Behaviour (IPAB)** relates the real world to computation through sensing and acting. The **Institute for Adaptive and Neural Computation (IANC)** studies learning and adaptation in natural and artificial systems. The **Institute for Language, Cognition and Computation (ILCC)** pursues basic and applied research on computational approaches to language, cognition, and communication. The **Centre for Intelligent Systems and their Applications (CISA)** innovates in automated reasoning based on formal knowledge representation languages. Institutes drive our core research strategy. We draw strength for core research from our activities which support interdisciplinary research and technology transfer. The applicability of our methods builds on our engagement with other sciences, engineering and industry.

### B. Research Strategy

#### B.1. Research Vision and Key Drivers

We focus on the processes and algorithms essential to informatics across sub-disciplines and have led on *natural and engineered systems that operate at large scale or exhibit great complexity*, understanding these through compact and elegant mathematical formalisms. We study how humans operate as information processing systems so we have *innovated in cognitive systems and experiments driven by data at all scales*, from micro-measurement of human behaviours through to massive, streamed social data. We investigate the informatics of physical and biological systems at all levels of organization from molecule to ecosystem *deepening our engagement with medicine, natural and physical sciences*, and building on the informatics expertise developed by our colleagues in these disciplines. We push the limits to scale and effectiveness of traditional engineered systems with *seminal work in the design of distributed, parallel and multicore systems*, benefiting from strong engagement from major industrial developers. We represent an international hub for cutting edge *robotics and autonomous systems* research through novel hardware platforms and algorithmic advances for their control.

Our strategy over the REF period was to promote the quality, scope and diversity of our research. Our headcount of academic staff numbers has increased since RAE 2008 (by 14) and a larger proportion of staff in our REF submission are senior academics than was the case in our RAE 2008 submission. Our research portfolio has increased by 69% since RAE 2008 to £87.54M. This consists of £63.44M of research funding complemented by £24.1M of additional translational research funding. Our PhD student body has grown by 84% (from 152 to a 280 student cohort). This growth in research portfolio and PhD student numbers has given us great energy to drive our academic research agenda. In tandem with this growth, the School has created a dedicated technology transfer operation that has produced more spinouts and startup companies than any other comparator UK institution. This combination of broad scientific reach and vigorous entrepreneurial activity is unique to Edinburgh.

Our RAE return followed a period of strong growth in preparation for merging our activities on a single site. This growth gave us the breadth and depth to pursue the vision of informatics research we reported in RAE 2008 (see B.2). In 2008, we argued that delivery of this vision also needed interdisciplinary research (reported in B.3) and translational research (reported in Section D).

**Environment template (REF5)****B.2. Research Achievements during the REF Period**

We present the School's key achievements by institute. For each institute we give: statistics indicating the extent of research activity; a summary of the strategic research focus; and highlight key elements of our success over the REF period.

**Laboratory for Foundations of Computer Science (LFCS)**

Key statistics: 31 academic staff, 12 research staff, 42 PhDs, £9.80M research portfolio.

Strategic research focus: We seek to uncover rich and profound insights into the semantic and mathematical foundations of computer science and informatics. Our key contributions are in data provenance, data management and scalable architectures for data processing, computational complexity, model checking, randomized algorithms, game theory, quantum computation, and automated verification. Our work finds applications through the design of programming languages for the Web and Internet, formal methods for software and system engineering, process modelling (discrete and continuous in computer and biological systems), dynamic modelling, and network modelling.

Evidence of success:

- Wadler's "Propositions as Sessions" paper (ICFP 2012) uncovers a new connection between session types and classical linear logic and underpins his £3.9M 2013 EPSRC programme grant "From Data Types to Session Types: A Basis for Concurrency and Distribution".
- Fan's paper "Towards certain fixes with editing rules and master data" provided an innovative new technique to repair inconsistencies in data with guaranteed accuracy. The paper won the VLDB best paper award in 2010. It was an output from EPSRC grant EP/E029213/1 (£481k).
- Mayr's paper "When Simulation Meets Antichains (On Checking Language Inclusion of Nondeterministic Finite (Tree) Automata)" won the EATCS Best Theory Paper Award at ETAPS 2010, and led to the award of a Royal Society Grant.

**Institute for Computing Systems Architecture (ICSA)**

Key statistics: 14 academic staff, 9 research staff, 33 PhDs, £7.07M research portfolio.

Strategic research focus: We seek to discover and develop new approaches to system design, programming and use that adapt to technology evolution. We develop and implement new architectures to tackle the challenges of providing systems that are both energy-aware and scalable from embedded devices to data centres. We do this by developing innovative approaches to exploiting large-scale parallelism and device heterogeneity. Our key contributions are in computer architecture and engineering, compilers, programming languages, ad-hoc networks, scalable protocols, and sensor networks.

Evidence of success:

- O'Boyle's paper "Partitioning Streaming Parallelism for Multi-cores: A Machine Learning Based Approach" is the first to use machine learning to predict an unbounded program graph structure rather than simple parametric optimisation settings. It won a best paper award at PACT 2010.
- Dubach's paper "Portable compiler optimisation across embedded programs and microarchitectures using machine learning" (MICRO 2009) was the first to offer performance portability across a microarchitecture design space, winning a HiPEAC paper award. In 2012, Dubach won an Intel Early-Career Faculty Honor Programme award for his work.
- Topham's paper "Cycle-accurate performance modelling in an ultra-fast just-in-time dynamic binary translation instruction set simulator" (SAMOS 2010) demonstrated a dramatic reduction in simulation time, which is critical for design-space exploration. It won the best paper award.

**Institute for Perception, Action and Behaviour (IPAB)**

Key statistics: 9 academic staff, 10 research staff, 33 PhDs, £7.42M research portfolio.

Strategic research focus: We develop anthropomorphic and mobile robotic systems that can adapt to and interact with humans, social settings and complex physical environments. We invent new methods of visual processing that take advantage of the massive data available in modern

**Environment template (REF5)**

images and video. Our key contributions are in statistical learning for motor control, biomimetic robotics, 3D vision, real-time mapping and navigation, learning visual concepts, and interactive animation and graphics. Application domains range from manufacturing to prosthetics, large-scale image retrieval, dynamic animation, and retargeting for movies.

## Evidence of success:

- Ferrari's paper "Segmentation Propagation in ImageNet" introduced a novel procedure to recursively segment large image databases without segmented training data. The paper won the best paper award at ECCV 2012 and led to the €1.4M ERC starting grant *VisCul*.
- Vijayakumar's paper "On Stochastic Optimal Control and Reinforcement Learning by Approximate Inference" was awarded the Best Paper runner-up at R:SS 2012. It presents a reformulation of the stochastic optimal control problem in terms of KL-divergence minimisation which simplifies the control problem. It contributed to a successful bid for a new robotics facility.
- Hermann's paper "Taming the Beast: Guided Self-organization of Behavior in Autonomous Robots" provided a practical solution to the exploration-exploitation dilemma for active agents and demonstrated the importance of autonomous learning. The paper won the best paper award at the "From Animals to Animats" conference in 2011.

**Institute for Adaptive and Neural Computation (IANC)**

Key statistics: 11 academic staff, 5 research staff, 71 PhDs, £7.70M research portfolio.

Strategic research focus: How natural and artificial systems can extract information from data. We develop the mathematical and statistical foundations of machine learning, with key contributions in Markov Chain Monte Carlo methods and probabilistic inference methods with particular emphasis on scalability and streaming data, Gaussian processes, and time-series modelling. We study adaptive processes in biological nervous systems, from genomic through neural to cognitive levels, with key contributions in modelling neural development and function, analysis of large-scale neural data, modelling gene expression dynamics, and systems biology modelling of the nervous system from the molecular to the cellular level. Our work engages with computational neuroscience, computational biology, drug discovery and proteomics/genomics.

## Evidence of success:

- Sanguinetti's paper "Point process modelling of the Afghan War Diary" (PNAS 2012) used statistical signal processing tools to fit a dynamical model to the Wikileaks Afghan conflict data. The work was awarded a Cozzarelli prize by the PNAS editorial board, an annual award to outstanding papers in the six main branches of the PNAS scope.
- Murray's paper "The neural autoregressive distribution estimator", which describes the first tractable high-dimensional binary distribution that gets state-of-the-art fits to data, achieved a Notable Paper Award at AISTATS in 2011. A CASE studentship from Novauris Technologies is continuing this work in deep learning.
- Bishop's paper "Beyond Atopy: Multiple Patterns of Sensitization in Relation to Asthma in a Birth Cohort Study" (AJRCCM 2010) applied Bayesian inference to understanding the factors influencing asthma, identifying the phenotype most strongly associated with its development.

**Institute for Language, Cognition and Computation (ILCC)**

Key statistics: 20 academic staff, 57 research staff, 64 PhDs, £25.31M research portfolio.

Strategic research focus: Advances in statistical methods have made possible new approaches to natural language and speech processing. Our computational models provide insights into human cognition and communication. Key contributions are in large-scale statistical machine translation; grammar-based robust natural language processing ("parsing the web"); mathematically well-founded models of summarization; and in speech synthesis and speech recognition that approaches human levels of reliability, adaptability and conversational richness. We have created widely-used open source toolkits (Festival, Moses, and NLTK) and contribute actively to projects started by others (HTS speech synthesis and Kaldi speech recognition). The clinical applications of speech technology include voice reconstruction for people with disordered speech (e.g. due to Motor Neurone Disease) and articulatory modelling for speech

**Environment template (REF5)**

therapy. Our work is effective at large scales including text mining and real-time processing of massive-scale social media data. We research in statistical models of human-computer dialogue, multimodal dialogue with robots, and negotiation with social cues. The practical applications of our work include cognitive assistants for older people, assistive technology for depression, and tutorial dialogue systems.

## Evidence of success:

- Osborne and Lavrenko's paper "Streaming first story detection with application to Twitter" paved the way to detecting breaking news in streams of data at scale. It led to the award of Osborne's 2012 EPSRC grant "CROSS: Real-time Story Detection across Multiple Massive Streams".
- Koehn's book on Statistical Machine Translation (CUP 2010) helped establish statistical machine translation as a mature research field. Its insights contributed to the widely-used Moses machine translation software.
- Yamagishi and Renals' paper "Robust speaker-adaptive HMM-based text-to-speech synthesis" led to the collaboration in the clinical applications of speech synthesis and was the key underpinning research behind the EPSRC Programme Grant "Natural Speech Technology" (2011–2016, coordinated by Renals).

**Centre for Intelligent Systems and their Applications (CISA)**

Key statistics: 9 academic staff, 24 research staff, 25 PhDs, £6.09M research portfolio.

Strategic research focus: Formal languages now provide the basis for massive, distributed systems of automated reasoning and for complex local problem solving. We develop new forms of language and styles of inference that are adapted to this scale and complexity, innovating in data architectures and social computation. We address the problems of coordination of agents in open and unpredictable environments, with applications in coalition formation and multi-agent learning. We explore the foundations of automated reasoning, with key contributions in automated theorem proving, logics for knowledge representation and geometric reasoning.

## Evidence of success:

- Bundy's paper "Scheme-Based Synthesis of Inductive Theories" (Best paper MICAL 2010) made a research breakthrough in the area of automatic conjecturing of interesting inductive theorems. This work was continued in the AI4FM grant; EPSRC Reference *EP/H024204/1*, value £513k.
- Atkinson's book "The DATA Bonanza: Improving Knowledge Discovery in Science, Engineering, and Business" (2013) describes the approach to the challenges of data-intensive computing developed around the language and architectures from our Data Intensive Research group.
- Robertson's paper "Programming the Social Computer" was the catalyst for the current surge of interest in the application of computational logic to social computation: now supported by the SociaM EPSRC Programme, Smart Societies IP, and many other related projects.

**B.3. Interdisciplinarity**

Informatics has expanded through interdisciplinary research. We have established the following major new interdisciplinary activities since 2008.

**Design Informatics** brings together researchers from the School of Informatics and the Edinburgh College of Art (now merged with the University) to explore design across the broad spectrum of our activities. The centre has attracted £1M of funding since its inception in 2011. It is the only design centre in the UK with a specific informatics focus.

**The Digital Healthcare Innovation Centre** is a £10M initiative to develop new systems of data-intensive healthcare across Scotland, in partnership with NHS24 and other Scottish universities. Edinburgh leads this initiative.

**The Farr Health Informatics Research Institute** is a £5M research activity funded by the MRC on medical informatics across Scotland, linking to the larger UK activity. Edinburgh is the research centre in Scotland (connecting to centres in London, Manchester and Swansea).

**The Data Science Innovation Centre** is a £10M activity run across the Scottish Informatics and Computing Science Alliance (SICSA) translating data-intensive research into industry and government. Edinburgh leads this initiative.

## Environment template (REF5)

**Robotarium** is a £6.1M robotic and autonomous systems research facility established jointly by the University of Edinburgh and Heriot-Watt University. It supports robotics research at the intersection between physical systems, collaborative interaction (human-robot and robot-robot), and smart spaces (see Section D for further information on this new facility).

We continue to grow our long-standing interdisciplinary activities. Principal amongst these are:

**Edinburgh Parallel Computing Centre**, a world-leading supercomputing centre. It manages the national facilities: HECToR (800 Tflop/s and 90,000 cores), DiRAC BlueGene/Q (6,000 nodes) and UK-RDF (8PB storage). It includes the Software Sustainability Institute, a national facility for users and developers of research software.

**SynthSys**, a centre for synthetic and systems biology led jointly by the Schools of Informatics, Biological Sciences, and Engineering (one of six Centres for Integrative and Systems Biology funded by BBSRC and EPSRC; £11.1M in 2007-11). SynthSys connects us to the Scottish Universities Life Sciences Alliance, SULSA, as a parallel to the Scottish Informatics and Computing Science Alliance, SICSA.

**Centre for Speech Technology Research**, linking Informatics with Linguistics and contributing to leading open-source toolkits: Festival, HTS (speech synthesis), and Kaldi (speech recognition).

**Neuroinformatics Doctoral Training Centre**, an £11M PhD programme for interdisciplinary research in neuroinformatics and computational neuroscience. Upon completion, it will have produced 111 PhD graduates, making it one of the largest programmes of its kind in the world.

### B.4. Future Priorities

Key research priorities by institute are:

**LFCS**: To develop the foundational theory necessary to analyse systems at large scale and in highly distributed settings.

**ICSA**: To innovate and develop architectures to satisfy future needs for energy-efficient and massively parallel systems, at granularities from processor to networked systems.

**IPAB**: To build a combined robotics and vision facility that is world-leading in adaptive and interactive robotics.

**IANC**: To extend models of learning from the neural and machine learning communities to meet the growing demand for adaptive systems from research, commerce and society.

**ILCC**: To construct foundational theories of cognition and natural language that are relevant to the expanding use of such theories in engineered systems.

**CISA**: To design new styles of automated reasoning and socially-derived inference built on rigorous mathematical principles.

In concert with these research priorities, we will grow our interdisciplinary engagement, building on the existing strong collaborations in biology, medicine, neuroscience and linguistics. We will grow our research training capacity through the establishment of three new Centres for Doctoral Training in the areas of **Data Science**, **Pervasive Parallelism** and, jointly with Heriot-Watt University, in **Autonomous Robotic Systems**. All will enroll their first students in 2014. We will foster the entrepreneurial culture responsible for our rapid rise in spin-out activity over the REF period and consolidate our position as a UK leader in this style of academic operation.

## C. People

### C.1. Staffing Strategy and Staff Development

At interview, we ask what attracts candidates to work in Informatics at Edinburgh. The most frequent answer is that they expect the School to substantially expand their horizons beyond their current sub-discipline. We encourage this attitude and foster it in our staff and students.

#### **Relationship between staffing strategy, research strategy and physical infrastructure:**

Our staffing strategy over the REF period has been to achieve controlled growth in our staff numbers while promoting the quality of our core informatics research and enhancing the scope and diversity of our interdisciplinary work. We have appointed 23 members of academic staff over the REF period to expand our expertise in:

- cognition, speech and language — Cohen, Lucas, Yamagishi;

## Environment template (REF5)

- complexity and algorithms — Diakonikolas, Sarkar;
- computer architecture — Grot, Lee;
- machine learning — Murray, Sanguinetti, Sutton;
- programming languages and XML — Cheney, Maneth;
- computer security and software engineering — Arapinis, Gordon, Rajan;
- neuroinformatics and systems biology — Henning, Thorne;
- robotics and vision — Chli, Fallon, Ferrari;
- computer systems, networks and compilers — Dubach, Leather, Patras.

In addition to the posts created during the REF period, we will benefit from the fact that the University is making 150 new five-year tenure-track Chancellors Fellows available in 2013-14. These posts are anticipated to lead to a permanent lectureship.

To accommodate expansion of its activities, the School of Informatics moved to two adjacent buildings on a single site at the heart of the University in 2008. **The Informatics Forum** brings the research staff submitted in this return under one roof. This £42M, 12,000m<sup>2</sup> new building was purpose-built to facilitate interaction and serendipitous encounters between members of the School. The excellence of the design of the Informatics Forum is acknowledged by awards from (amongst others) the Civic Trust, the Royal Incorporation of Architects in Scotland, the Royal Institute of British Architects, and the Carbon Trust. **The Appleton Tower** was refurbished at a cost of £6M, to provide an 1800m<sup>2</sup> centre for knowledge transfer and commercialisation, 1800m<sup>2</sup> of specialist teaching accommodation, and 300m<sup>2</sup> of electronic and mechanical workshops for the School.

**Support of staff career development:** The School supports all the principles of the Research Concordat via a tiered support strategy.

*Support for research assistants:* The Informatics Research Staff Society provides a forum for scientific, professional and social interaction, operating across research institutes. Elected members from this Society sit on our two most senior committees (Strategy Committee and Research Committee).

*Support for ECRs:* ECRs benefit from the supportive research community structure found within our six research institutes. Directors of institutes provide shape and direction to their research areas. Institutes deliver staff induction, provide access to senior role models, run popular seminar series, build links between experienced funding holders and first grant applicants, and will preferentially invest institute funding in PhD studentships to be supervised by ECRs. We support ECRs through pre-submission grant reviews and mock panels before fellowship interviews. Our average funding success rate runs at 61% as a result. The School chose to invest £2.2M of its own core funding into PhD studentships over the REF period including dedicated support for ECRs in building the research groups that sustain their careers. We work with the Institute for Academic Development (IAD) at university level.

*Support for established academic staff:* Our flexible working models support joint working with industry, or part-time appointments. We have a generous sabbatical policy as we aim for 10% of our staff to be on sabbatical each year. The allocation of teaching and administrative duties is transparent in the School. Internal promotions over the REF period provide strong evidence of successful support. Thirteen academic staff have been promoted to Professor (Anderson, Armstrong, Arvind, Cintra—now at Intel, Gilmore, Keller, Koehn, Lapata, Robertson, Simpson, Thompson, Vijayakumar, and Webb). Eleven academic staff have been promoted to Reader (Aspinall, Ferrari, Franke, Goldwater, Kashefi, Komura, Osborne, Sanguinetti, Santhanam, Storkey, and Viglas). Six academic staff have been promoted to Senior Lecturer (Bednar, Fleuriot, Jackson, Rovatsos, Smaill, and Stark).

All School of Informatics staff benefit from annual professional and development reviews. Nominated reviewers support our research staff by complementing line management within the School and the system of professional and development reviews through their focus on career development. We facilitate staff training uptake through provision of local courses, a notable recent example being recruitment law training for grant holders.

**Environment template (REF5)**

**Personal research fellowships won in open competition:** Our academic staff are successful at obtaining personal research fellowships. These are key to our research strategy as detailed below.

- Recognition of international excellence for established research programmes — Danos and Steedman have been awarded ERC Advanced Grants.
- Establishing new research programmes — Ferrari, Keller and Sanguinetti have been awarded ERC Starting Grants.
- Relating science to engineering — Dubach, Leather and Fensch won REng/EPSRC fellowships; Vijayakumar has been awarded a Microsoft Research/REng Senior Research Fellowship.
- Development of “rising stars” — Kashefi has been awarded an EPSRC Advanced fellowship; Yamagishi has been awarded a EPSRC Career Acceleration fellowship; Cheney won a Royal Society University Research Fellowship; Hennig has been awarded an MRC Career Development fellowship. Lapata completed her EPSRC Advanced Fellowship in 2011.

**International staff appointments:** We recruit internationally. Our priority in academic recruitment is to attract early-career staff with the highest potential. Of the academic staff appointed since 2008, fourteen have come to us from major centres of research outside the UK and fifteen have completed their PhD studies abroad. Recent appointments include Arapinis (Birmingham, PhD Paris XII); Cheney (Edinburgh, PhD Cornell); Chli (ETH Zurich, PhD Imperial), Cohen (Columbia, PhD Carnegie Mellon); Diakonikolas (UC Berkeley, PhD Columbia); Fallon (MIT, PhD Cambridge); Ferrari (Zurich, PhD ETH Zurich); Grot (EPFL, PhD U Texas at Austin); Lee (Purdue, PhD Purdue); Lucas (Carnegie Mellon, PhD UC Berkeley); Maneth (Oxford, PhD New South Wales); Murray (Toronto, PhD UCL); Nagarajan (UC Riverside, PhD Riverside); Patras (Maynooth, PhD Madrid); Rajan (Grenoble, PhD Minnesota); Sarkar (Berlin, PhD Stony Brook); Sutton (UC Berkeley, PhD UMass Amherst); Yamagishi (Edinburgh, PhD Tokyo Institute of Technology).

We generate many more scientists than we can retain. Outgoing staff have secured positions in the best international research organizations. Examples include Cintra, going to Intel Labs Braunschweig; Coja-Oghlan to Frankfurt/Main; Geerts to Antwerp; Harmeling to the Max Planck Institute; Momigliano to Milan; Pollock to Harvard; and Steel to ENS-Cachan, INRIA Paris.

**Visiting scholars:** Augmented by SICSA Distinguished Visiting Fellows, an active visitor programme enhances our research environment. The School provides private visitor offices and flexible hot-desk space. We have hosted in excess of 550 research visitors to our School over the REF period: durations range from a few days to a few months. Notable examples include Andrew Black (Portland State), Sharon Crook (Arizona State), Mari Ostendorf (U Washington), Prakash Panangaden (McGill), Phil Scott (Ottawa), Nitin Vaidja (U Illinois), Mihalis Yannakakis (Columbia).

**Equality and diversity:** The School of Informatics is one of only two schools in the subject area of computer science in the UK to hold a Silver Athena SWAN award, recognizing our commitment to advancing women's careers in our discipline. No school of computer science in the UK holds a higher Athena SWAN award. We are one of the most ethnically and culturally diverse schools in the University of Edinburgh: our staff returned for REF come from 20 nations. Our School Equality and Diversity committee engages with College and institutional structures. A member of our professoriate (Lapata) is our Equality and Diversity Coordinator. This role has a seat on our highest committee, Strategy Committee. A member of our support staff is one of a University-wide network of Dignity and Respect advisors. The School has invested staff development funding in the provision of external mentoring support for female research and academic staff via the Scottish Resource Centre for Women in Science, Engineering and Technology. Equality and diversity training is promoted to all staff and required for those on senior management committees and those responsible for recruitment.

**C.2. Research Students**

Our decision to make a strong investment of School funds in PhD studentship awards has contributed to a step change in the scale of our training and supervision of postgraduate research students since RAE 2008. The School has a large population of postgraduate research students (currently 268 PhDs and 12 MRes) who all belong to our Graduate School (responsible for their welfare and career development overall) and a research institute (providing a focus for their

activities amongst a cluster of more senior researchers). Our average number of research doctoral degrees awarded per year has increased by just over 60% since RAE 2008 (from 24.19 to 38.8), an 83% increase in awarded degree per Cat A FTE.

**Postgraduate research student recruitment:** Our applications for postgraduate research positions average 474 per academic year. The postgraduate selectors of each research institute direct applicants to potential supervisors. Each applicant's research proposal, CV and academic transcripts are reviewed at the institute level by at least two academics. Selected candidates are interviewed. The Head of Graduate School makes final decisions about academic eligibility, supervisor availability and funding awards in consultation with area-specific selectors. We average 90 offers per year of which an average of 70 students arrive — two-thirds from outside the UK.

**Training and support mechanisms:** The School of Informatics works closely with the University's Institute for Academic Development's postgraduate team to produce a programme of skills and research training tailored to its postgraduate research students. All first-year students take courses on PhD management, planning and thesis proposal authoring. Second-year students take courses on research presentation and on research paper writing. Third-year students take courses on thesis writing and viva preparation. We provide an informal mentoring scheme, where each student has a peer-supporter from amongst the cohort to support them. Courses are also provided in research planning, professional development, and communication. Research students are encouraged to attend appropriate technical courses from our taught MSc programme, with the choice of suitable courses made in consultation with their supervisor. Students typically take three or four such courses in their first or second years of study. Students are fully integrated into the School's research environment including seminar series, outreach programmes, commercialisation opportunities, computing support, and access to facilities. We provide opportunities for paid teaching, with tailored training in teaching given by the Institute of Academic Development. The School funds travel for postgraduate research students, enabling them to present their results at major national and international conferences, summer schools and workshops. A typical student attends three of these in the course of their studies, although this varies between students and research areas.

All of our postgraduate research students are also members of the SICSA Graduate Academy. SICSA is a collaboration of Scottish Universities whose goal is to develop and extend Scotland's position as a world leader in Informatics and Computer Science research and education. The Graduate Academy allows members to participate in SICSA activities without charge and alongside fundamental training in research, it organises a coordinated programme of research-related activities including transferable skills and entrepreneurship training. These include training courses, conferences and workshops for all computer science and informatics PhD students working in Scotland. SICSA also supports advanced, discipline-specific short courses delivered through summer schools – residential courses lasting one or two weeks. Examples include the summer school on affective computing and social signal processing (Edinburgh, 2010) and the summer school on multimodal systems for digital tourism (St Andrews, 2011). The majority of SICSA Distinguished Visiting Fellows give master-classes in their research specialisms.

**Progress monitoring:** All postgraduate research students have designated first and second supervisors. Each Institute holds an annual student review meeting to consider the progress of their students, with each discussed and minutes submitted to the Graduate School. Students and supervisors are required to submit annual review forms (internal to the institute) before this meeting, summarising their progress, identifying any issues and confirming that a review meeting for the student has taken place with a panel consisting of both supervisors and an additional independent staff member. Subsequent to the meeting, supervisors are required to make a formal annual report for each of their students. The deputy director of the Graduate School monitors the completion of these annual reports. Systems specific to our research institutes augment the progress monitoring which the School undertakes. For example, the Institute for Computer Systems Architecture (ICSA) provides one year's postdoctoral research extension for selected PhD students within its institute. Senior members of ICSA provide additional support at this stage to help these young researchers actively manage their careers. This gives a natural transition period

## Environment template (REF5)

that allows researchers to develop publications from their PhD studies, to travel, and to write grant applications. It provides time to build collaborations with industry and broader academic networks. This led to four PhD students subsequently obtaining prestigious five-year RAEng fellowships over the REF period.

**D. Income, infrastructure and facilities**

**Research funding portfolio:** The School maintains a portfolio of research awards of £63.44M. This is a 22% increase on the £51.9M portfolio reported in RAE 2008. We have achieved this by diversifying our funding sources, receiving funding from all the major UK funding councils and the major European funding bodies. We have established centres of excellence: see Section B.3 above. We have secured large-scale awards nationally, including EPSRC Programme grants. We have secured large-scale awards internationally, including EU Framework 7 IPs and ERC Advanced grants. We have invested in ten dedicated research administration staff across the School. The School has ambitious future targets for both the scope of our research agenda and the scale of our future research portfolio, as seen in our infrastructure and technology transfer plans.

**Specialist infrastructure and facilities:** The University has an integrated network providing wired and wireless connectivity across its estate, with secure remote access to UoE computing, data and library resources. The School extends this with its own security, network, storage and backup infrastructure, and a managed platform that provides a consistent Linux installation across more than 1,000 machines including over 200 servers. Our managed Linux platform builds on Scientific Linux, augmented with over 1000 additional software packages to support research and teaching. A staff of 26 dedicated computing/technical personnel support the School's computing facilities. Our mechanical and electrical workshops design, develop and maintain equipment for staff and student research projects.

For those with compute-intensive or data-intensive needs, the University's **Edinburgh Compute and Data Facility (ECDF)** provides around 3,000 CPU cores and 1PB of storage. Institutional investment in ECDF totals £1.6M, additionally Informatics invested £0.87M to purchase CPU hours for our research staff and students enabling free access irrespective of personal research funding position. The Edinburgh Parallel Computing Centre leads the **HPCX** consortium, and directs and operates **HECToR** (High End Computing Terascale Resources) at the University's **Advanced Computing Facility (ACF)**. **HECToR** was funded through a £115M, 6-year contract from EPSRC. This resource provides over 90,000 cores and a peak processing capacity of over 800 Tflop/s.

The School houses a suite of laboratories for research in speech, language, and cognitive science. We have four **sound studios** for general audio recording, audio recording in hemi-anechoic conditions, or audio recording in conjunction with eye-tracking or electromagnetic articulography. We have two **perception labs** with facilities for reaction-time measurements and audio presentation and recording. We have an **eye-tracking lab** with two head-mounted eye-trackers to be used individually for reading or perception experiments, or jointly for studying dialogue, interaction and collaboration. We have a **shape-capture lab** with a high-resolution static shape-capture system for 3D modelling. We have an **instrumented meeting room** for recording of multi-party meetings and videoconferences, using synchronised multiple audio and video channels as well as capture of hand writing, and data projector and whiteboard use. Over and above their initial set-up costs, these facilities have been supported through direct School investment of £0.3M over the REF period and through the establishment of a transparent recharge mechanism allowing access cost inclusion on funding awards. This supports sustainable spending on running costs and periodic replacement of equipment.

The School provides specialist infrastructure for vision and robotics. We have **anthropomorphic robotic platforms** including 15 Nao mini humanoids, a KUKA LWR robotic arm equipped with a Schunk Tactile hand and KUKA YouBot omnidirectional mobile platform. We have a large number of special **prototype robotic actuator** platforms developed completely in-house including novel variable impedance actuators. These are used for research on bipedal locomotion, to test limb prosthetics, to realise dynamic interactive tasks such as manipulation and brachiation, and to

**Environment template (REF5)**

develop bio-inspired hexapod all-terrain locomotion platforms. To support experiments in collaborative and aerial robotics we provide a dedicated **field robotics laboratory** with facilities including a robot soccer pitch, double-height facility for experiments with AUVs and real-time motion capture facility. In addition, we have a **sensorimotor lab** with equipment for 16-channel EMG measurement, polyhymous motion tracking, and a robotic manipulandum to conduct human psychophysics research and experiments involving reach and grasp. We have specialist **vision and range sensors** including: sensors capable of scanning surfaces of mechanical objects at a resolution of 5nm; sensors designed for heads, acquiring 7MP colour head images at 0.1mm resolution; a 25fps, 1.2MP 3D colour video system; and a 500fps, 1.2MP 3D monochrome system.

This is the core of our future **Robotarium**, a Centre of Excellence in robotics in Edinburgh funded by a UK government £6.1M award. The facility will house humanoid robots including the £2M DARPA Atlas robot, PR2, KUKA and Baxter humanoids as well as a complete suite of sensorised smart spaces for monitoring human/robot interaction. Our current sensorimotor lab will be augmented with a split-treadmill gait analysis laboratory and a prosthetics development facility in collaboration with Touch Bionics. Our **3D printing** and **laser LPKF systems** for fabricating variable modulus structures for microbots will strengthen our actuator design capabilities.

The School also provides unique **translational facilities**. The Informatics Forum has purpose-built **flexible event space**, with a lecture hall, atrium, café area and breakout space on the ground floor. The Forum is a popular event venue that has hosted prestigious conferences (Federated Logic Conference 2010, International Conference on Machine Learning 2012, Annual Conference on Learning Theory 2012, Parallel Architectures and Compilation Techniques 2013). **Public engagement** in science is supported via InSpace, a laboratory that explores the cultural significance of informatics and new media practice. The School has invested £0.28M in InSpace over the REF period. InSpace averages 12,500 visitors per annum and is now an established venue for the Edinburgh International Science Festival (EISF). A noteworthy example is LateLab, sited in InSpace, “a creative, social and participatory programme for adults where they can make, do and interact with art, science and technology.” (1,000 people over twelve nights, EISF 2012).

**Translational research, technology transfer, consultancies and professional services:**

Our technology transfer and innovation activities are summarised in REF3a. Since RAE 2008, we have developed a leading UK entrepreneurial initiative and run it alongside a leading UK research operation. Our current translational funding portfolio is £24.1M. Informatics has generated approximately one third of all spinouts created by the University over the past decade (44 of these since Jan 2008 alone). According to an independent survey by SpinoutsUK, the University of Edinburgh is the leading university in the UK for spinout creation over the past decade. The spinouts from Informatics alone would be enough to place the University of Edinburgh fourth in the UK. In addition to this activity, we have consultancies (£3.02M over the period) and licensing arrangements with major international companies that follow research through many stages into market. We have licensed our EnCore chip technology created by Topham to Synopsys (Nasdaq:SNPS), a \$4bn Silicon Valley electronic design automation company. Synopsys are now developing EnCore commercially and Topham and his group are working with them to bring EnCore to market. Patented ‘sense’ feedback technology developed with Touch Bionics will within the next five years be featured in their prosthetic hands currently fitted to over 500 amputees. This translational activity is unusually large for a School where academic excellence is the priority. We achieve this through dedicated School-level business development staff; bespoke entrepreneurial and startup training; access to professional support for consultancy; and University-level support for translational research via Edinburgh Research and Innovation.

**E. Collaboration or contribution to the discipline or research base****E.1. Research collaborations**

All of the School’s research institutes are underpinned by funded collaborations with major academic and industrial research centres internationally.

**Personal, funded research activities:** anchored by individuals in the School with others at major informatics centres nationally and internationally – examples include: with Oxford (Robertson,

## Environment template (REF5)

Libkin, Stevens), Cambridge (Armstrong, Stark, Sutton, Topham, Osborne, Renals, Yamagishi), UCL (Storkey, Bednar), DFKI-Germany (KoeHN), Trento (Rovatsos), Harvard (Murray, Sanguinetti), MIT (Keller), Toronto (Santhanam), McGill (Danos, Plotkin), Haifa (Armstrong), IIT Madras (Nagarajan), Beihang (Fan), Tokyo (Vijayakumar), Penn (Buneman).

**Research pooling:** we are one of the three core universities in the Scottish Informatics and Computing Science Alliance (covering 14 Scottish universities). SICSA is a research pool established in the previous RAE census period with the aim of promoting shared expertise and excellence in research across computing departments in Scotland's universities. The School became the entrepreneurial hub for SICSA in 2012, with an investment of funding from Scottish government of £1.17M.

**Major national/international programmes:** collaborations nationally involving multiple sites – for example through involvement in EPSRC Programme grants (we currently have five of these; held by Robertson, Stark, O'Boyle, Wadler (co-ordinator), and Renals (co-ordinator)) and internationally through a significant European portfolio; we currently have 34 awards (of which we co-ordinate 8) with a cumulative value of £21.9M.

**Industrial research collaborations:** where we have undertaken funded research with key labs in industry – for example AstraZeneca (Armstrong), BT (Goddard), Google NYC (Sutton), NIH (Bednar), Novauris (Murray), ARM (O'Boyle), IBM Haifa (Leather), IBM Watson (Dubach), INRIA (Leather, O'Boyle), Intel (Dubach, Viglas, Nagarajan), Samsung (Franke, Cole), Synopsys Advanced Processors Group (Topham), BBC (Renals), Cisco (Renals), Toshiba (Renals, Yamagishi), AT&T (Fan), Bell Labs (Etessami, Fan, Libkin), Google (Ferrari, Wadler), Microsoft Bangalore (Gordon), Microsoft Redmond (Jackson, Wadler), Microsoft Research (Sutton, Gordon, Hillston, Plotkin, Viglas, Vijayakumar, Wadler), Praxis (Jackson), Yahoo (Fan), Honda (Vijayakumar), Disney (Komura), Oracle (Dubach, Viglas).

## E.2. Interdisciplinary research

Interdisciplinary research cuts across all of the School's strategic research areas and has led to exciting new insights and discoveries. We provide funding for networking events, staff travel and research sabbaticals to promote interdisciplinary research. Every institute in the School engages with vigour in interdisciplinary research. We present highlights below by institute.

*LFCS [Biology] — Biological process modelling.* Our collaboration with systems biologists in Edinburgh led to the development and application of languages such as Bio-PEPA. Using static analysis through the computation of model invariants in Bio-PEPA, we have been able to discover previously unknown errors in published biological models and in Bio-Models, the premier database of curated biological models. [Key reference: Hillston's paper, "Bio-PEPA: A Framework for the Modelling and Analysis of Biological Systems", TCS 2009.]

*LFCS [Physics] — Quantum information theory.* Our collaboration with physicists in the UK, France, Austria and Canada led to the formulation of a new cryptographic protocol: universal blind quantum computing. This protocol demonstrated for the first time the possibility of preserving the privacy of computation using quantum properties. We achieved the first experimental demonstration of blind computing with Zeilinger's group in Vienna, arguably the best quantum optics group in the world. [Key reference: Kashefi's paper "Demonstration of Blind Quantum Computing", Science 2012.]

*ICSA [Engineering] — Processor design and visible light communications.* Our collaboration with electronics engineers in Edinburgh led to new processor designs for communication systems in the visible light spectrum. [Key reference: Topham's paper "Design-Space Exploration of Resource-Sharing Solutions for Custom Instruction Set Extensions", TCAD 2009.]

*IPAB [Healthcare] — Prosthetics.* Our collaboration with healthcare professionals and Touch Bionics Inc. in Edinburgh on augmented sensory feedback for improving dexterity of upper limb prosthetic use led to the introduction of novel 'pulsed' control modes in the iLIMB prosthetic hand. [Key reference: Vijayakumar's paper "The role of feed-forward and feedback processes for closed-loop prosthesis control", Journal of Neuro-engineering and Rehabilitation, 2011.]

*IPAB [Medicine] — Skin cancer classification.* Our collaboration with dermatologists at the University of Edinburgh led to the identification of inconsistent judgements in public self-detection and primary clinical assessment of potential melanomas. [Key reference: Fisher's paper "Novice identification of melanoma: not quite as straightforward as the ABCDs", Acta

## Environment template (REF5)

Dermato-Venereologica, 2011.]

*IANC [Physiology] — Computational models of the primary visual cortex.* Our collaboration with computational neuroscientists showed that a wide range of observations could be traced back to a small set of basic principles, suggesting an underlying simplicity behind complex adult cortical function. [Key reference: Bednar's paper "Building a mechanistic model of the development and function of the primary visual cortex", Journal of Physiology, 2012.]

*IANC [Neuroscience] — Neural networks from MRI scans.* Our collaboration with psychiatrists at Edinburgh demonstrated that intracortical similarities could provide a robust statistical description of the gray matter morphology of individual brains [Key reference: Serié and Willshaw's paper "Similarity-Based Extraction of Individual Networks from Gray Matter MRI Scans", Cerebral Cortex, 2011.]

*ILCC [Psychology] — Computational aspects of human language acquisition.* Our collaboration with linguists and psychologists in Stanford, Berkeley, Macquarie, and Georgia Tech led to the application of novel machine learning methods to research in cognitive science and language acquisition [Key reference: Goldwater and Steedman's paper "A Probabilistic Model of Syntactic and Semantic Acquisition from Child-Directed Utterances and their Meanings", EACL 2012].

*CISA [Physics] — Ontology evolution.* Our GALILEO system leads automated discovery of new Physics concepts through proof-system-based resolution between theories and experimental evidence. [Key reference: Bundy's paper "A Higher-Order Approach to Ontology Evolution in Physics", Journal on Data Semantics 2012.]

*ILCC [Healthcare] — Real time tongue tracking for speech therapy.* Our collaboration with speech therapists at Queen Margaret University led to the development of ultrasound technology to provide visual feedback in speech therapy for children. [Key reference: Richmond and Renals, "Ultrax: An animated midsagittal vocal tract display for speech therapy", Interspeech 2012].

### E.3. Exemplars of leadership

The School's leadership has been acknowledged throughout the REF period by a series of prestigious national and international awards.

*Major fellowships:* Fellowship of the Association for Computational Linguistics (Steedman, Webber), Fellowship of the Association for Computing Machinery (Fan, Libkin), Fellowship of the Royal Society (Buneman, Bundy), Fellowship of the Royal Society of Edinburgh (Fourman, Fan, Libkin, Vijayakumar).

*Prestigious awards and prizes:* ACM SIGPLAN Programming Languages Achievement Award (Plotkin), BCS Roger Needham Award (Fan), European Academy of Sciences, Blaise Pascal Medal in Computational Sciences and Information (Plotkin), PNAS Cozzarelli Prize for Engineering and Applied Sciences (Sanguinetti), Royal Society Milner Award (Plotkin), Royal Society-Wolfson Research Merit Award (Wadler; Plotkin).

*National honours for computing science:* Bundy was appointed Commander of the Order of the British Empire (CBE) in the 2012 New Year Honours for his services to computing science. Buneman was appointed Member of the Order of the British Empire (MBE) in the 2013 New Year Honours for his services to data systems and computing.

Our staff fulfill **national advisory roles**. Bundy was the first Vice President for the BCS Academy of Computing, its learned society for advancing the academic discipline of computer science. Renals serves on the Scottish Enterprise advisory board on Informatics and Computing. Willshaw is an advisor to the Wellcome Trust and MRC Senior Fellowship Committees and HM Government Scientific Representative to the International Neuroinformatics Coordinating Facility. Bundy and Oberlander serve on the Scottish Science Advisory Committee. The School engages with the broader UK discussion on the shape of our discipline via UKCRC, the BCS, and in other fora. Robertson is a member of the UKCRC executive. We fulfill **international advisory roles** including: Scientific Board, IST, Austria (Plotkin); Scientific Advisory Committee, Institute of Logic, Language and Information, University of Amsterdam (Plotkin); Advisory Board for the Computer Science Department of Harbin Institute of Technology (Sannella); Digiteo Scientific Council (Stevens); Scientific Advisory Board of the Excellence Cluster on Multimodal Computing and Interaction in Saarbruecken (Gordon); BCS representative to IFIP TC1 (Sannella); Associate member of the Canadian Institute for Advanced Research program on Neural Computation and Adaptive

## Environment template (REF5)

Perception (Williams); review board member for Germany's Bernstein Centres in Computational Neuroscience and in Neurotechnology (Willshaw); Scientific Advisory Board member, Goettingen Computational Neuroscience Centre (Willshaw).

Within **learned societies, funding bodies and professional organisations** we serve on:

Sectional Committee 1 of the Royal Society, (Plotkin 2008–2010), Sectional Committee B4 of the Royal Society of Edinburgh (Hillston and Webber 2008–2012, Fan 2012–date, Vijayakumar 2013–date), the Computer Science (PE6) Advanced Grants panel of the European Research Council (Hillston 2008, 2010, 2012), the EPSRC/GCHQ Cyber Research Institute panel (Gordon 2012), BCS/CPHC Distinguished Dissertation Awards Committee (Hillston 2006–2009, chair 2008–2009), ACM PODS Executive Committee (Libkin 2008–2009), W3C Provenance Incubator Group and Interchange Working Group (Cheney), ACM SIGPLAN (Wadler chair 2009–2012 and past chair 2012–2015), ISCA Advisory Council (Renals 2011–2014), IEEE Speech and Language Technical Committee (Richmond 2012–2014), ACL executive committee (Steedman), UKCRC (Bundy, Buneman, Fourman, Fan, Hillston, Libkin, Moore, Plotkin, Robertson, Wadler).

**Our invited keynote lectures** included in 2009 the ACL Annual Meeting (Webber), DATE (Topham), FASE (Gilmore), IJCNLP (Webber), PODS (Libkin), and Quantum Information and Technology (Kashefi). In 2010, CiE (Santhanam), ETAPS (Wadler), FOSSACS (Stirling), the Oxford University Strachey Lecture (Hillston), LICS (Danos and Plotkin), and NAACL/HLT (Renals). In 2011, CALCO (Danos), CiE (Kashefi), DATE (Topham), and QDB (Fan). In 2012, EDBT/ICDT (Fan), ICIAR (Fisher), LCC (Libkin), London Colloquia on Combinatorics (Cryan), POPL (Gordon), Princeton Turing Centennial (Wadler), and SOS/EXPRESS (Stirling). In 2013, Behaviour (Webb), ICIV (Webb), STACS (Etessami), and TGC (Hillston).

**Our best paper and book awards** include: the Test-of-Time Award ACM PODS (Fan 2010); VLDB (Fan, 2010); most influential ACM POPL 2000 Paper (Gordon, 2010); EATCS best theory paper (Mayr, 2010); ICDT (Libkin, 2011); ICSE ACM Distinguished Paper award (Rajan, 2008); HiPEAC award (Franke); PACT (Nagarajan, O'Boyle, 2010); ICSAMOS (Topham, 2010); Itakura Prize (Yamagishi, 2010); CC (O'Boyle, 2011); PLDI HiPEAC award (Topham, 2011); DeGroot Prize of the International Society for Bayesian Analysis (Williams, 2009); EC-TEL (Moore, 2010); AISTATS (Murray, 2011); SAB (Hermann, 2010); ECCV (Ferrari, 2012); QEST (Sanguinetti, 2013).

**Our research students** lead in their peer groups: O'Boyle's PhD student, Christophe Dubach, won the 2009 BCS distinguished dissertation award; Keller's PhD student, Vera Demberg-Winterfors, won a Robert J. Glushko dissertation prize in 2011 for an outstanding dissertation in cognitive science; Osborne's student Abby Levenberg won the European Association of Machine Translation PhD thesis award in 2011; Simpson's PhD student, Matteo Mio won the European Association for Computer Science Logic outstanding dissertation Ackermann Award in 2013; and Lapata and Sutton's PhD student Daniel Renshaw was one of only 39 PhD students worldwide who were awarded a Google PhD Fellowship in 2013.

We provide **editors in chief** for journals such as Theoretical Computer Science (Sannella) and ACM Transactions on Speech and Language Processing (Renals). We serve on **editorial boards** including Theoretical Computer Science (Hillston, Fan, Stevens), Mathematical Structures in Computer Science (Danos, Stirling), Transactions of the ACL (Goldwater, Koehn, Lapata, Steedman), Computational Linguistics (Koehn, Goldwater, Keller), IEEE Transactions on Audio, Speech and Language Processing (Renals), IEEE Signal Processing Letters (Renals), Journal of Natural Language Engineering (Webber), Speech Communication (Moore), ACM Transactions on Speech and Language Processing (Lapata), International Journal of Computer Vision (Williams), Journal of Machine Learning Research (Williams), Proceedings of the Royal Society A (Williams), IEEE Transactions in Pattern Analysis and Machine Intelligence (Storkey).

Finally, we have provided **general conference chairs** for, among others, NIPS'10, MiG'11, ACM/Eurographics'12, PACT'13, and PLDI'14 and **programme chairs** for ACL'08, EMNLP'08, EMNLP'09, HiPEAC'09, NIPS'09, AMDO'12, VR'10, MIG'10, CGO'11, ConLL'11, ACL'12, EACL'12, CC'12, LCTES'13, and PODS'13.