The Ring

The journal of the Cambridge Computer Lab Ring

Issue XLVIII — May 2018

Who's Who

Adam Cohen–Rose (CL BA95) has joined Tesco Labs as a software development engineer.

William Denman (PhD14) has been appointed Quantitative Developer at Orbis Investments in Vancouver.

Simon Fothergill (PhD13) has recently joined PROWLER.io as a senior data scientist.

Andy Harter (F BA83 CC PhD90) CBE, FREng, CEng, FIET, FBCS, CITP, FLCM, FRSA, has been appointed the High Sheriff of Cambridgeshire.

Andy Harter CBE

There is a High Sheriff in every county in England and Wales, and their history and tradition goes back before the Norman Conquest. The Office is the oldest Royal appointment.

Bjarki Holm (M MA05, PhD10) is working for Lucidworks as VP of Solutions.

Laura James (CC MA00 PhD05) is Entrepreneur in Residence at the Department of Computer Science and Technology, University of Cambridge

John Lusty (CHU MA98) has joined FiveAI as VP Simulation.

Sean Moran (CTH BA05) is working at Huawei where he is a Senior Research Scientist in Machine Learning.

Barney Pell (T PhD93) has joined the Board of Directors of Ecoation Innovative Solutions Inc, The Ecoation artificial intelligence platform predicts the type, location and level of crop stress using plant signals, and communicates this intelligence daily to growers on their phone or computer.

David Singleton (JN BA02) recently joined Stripe from Google, where he was director of Android Wear. David will lead Stripe's global engineering efforts.

Bjarne Stroustrup (CHU PhD79) has been named the recipient of the 2017 Faraday Medal, the most prestigious of the Institution of Engineering and Technology (IET) Achievement Medals.

Bjarne has also been awarded the 2018 Charles Stark Draper Prize for Engineering "for conceptualizing and developing the C++ programming language."The $500,000 annual award is given to engineers whose accomplishments have significantly benefited society.

Stephen Thornhill (CHU MA98) is Senior Principal Engineer at Menlo Security Inc.

Andy Twigg (K PhD06) is Entrepreneur in Residence at Milliways Ventures in the San Francisco Bay Area.

Paul Ward (Q MA94) works for RBS in London where he is Data Innovation Strategy Lead.

David Young (JN BA90) has joined Wates Group as IT roadmap programme manager.

Ali Mustafa Zaidi (ED PhD14) has been appointed Research Architect at ARM.

Events calendar

2018

May

Tuesday 8th, 6.30pm
Bay Area Ringlet Bar
Southern Pacific Brewing, 620 Treat Ave, San Francisco

June

Wednesday 6th, 6.30pm
London Ringlet Bar
Venue to be confirmed

August

Thursday 2nd, 6.30pm
London Ringlet Bar
Venue to be confirmed

October

Wednesday 3rd, 6.30pm
London Ringlet Bar
Venue to be confirmed

December

Tuesday 4th, 6.30pm
London Ringlet Bar
Venue to be confirmed

Hall of fame news

Bromium

Bromium was named Cambridge Computer Lab Ring Company of the Year 2018 at the Hall of Fame Awards, held at Queens' College on April 11th 2018.

Bromium also received two Government Security News™ (GSN) Homeland Security Awards for the second consecutive year. It won in the Best Anti-Malware Solution and Best Email Security/Loss Management Solution categories, and was also named a finalist for Best Application Security Solution.

DroneDeploy

DroneDeploy, the leading cloud software platform for commercial drone, has launched a drone solution for the roofing industry that automatically generates accurate and comprehensive roof analytics. For the first time, roofing, solar, and insurance companies can easily and affordably inspect and measure roofs with industry–leading accuracy.

DroneDeploy's roofing solution integrates an autonomous 3D Flight App with a roof–specific flight mode and an application that generates PDF roof reports and DXF files. The 100% cloud–based mapping and roof inspection platform allows companies to capture information from multiple houses, and start generating insights immediately.

Factmata

Factmata has closed its seed funding round.

The investors included Biz Stone, co–founder of Twitter and Craig Newmark, founder of Craigslist.

GeoSpock

GeoSpock raised an additional $6.6M in closing its Series A funding round bringing the total raised to $13.2M. The round was led by Cambridge Innovation Capital (CIC) with existing investors Parkwalk Advisors and Sir Michael Marshall. In addition, Japanese strategic investors Global Brain and 31 Ventures also participated.

The company's data engine analyses extreme amounts of contextual data in sub–second response times. It has the capability to significantly improve artificial intelligence (AI) speed, enabling the use of extreme data to train new systems, and can integrate and correlate information from a wide variety of data sources at a moment's notice. The platform is used by customers in the mobility, smart city, ad tech, financial services, telematics, and telco sectors, who benefit from interactive analytics at scale.

Masabi

Masabi, the global leader in mobile ticketing and Software–as–a–Service (SaaS) based fare collection for public transit, has announced a new strategic partnership with Uber to add public transit mobile ticketing into the Uber app. Once an agreement is reached with a transit agency, Uber users will have the option to book and display Masabi customers' transit tickets in the app, allowing for seamless transfers from ride–sharing to public transit services for convenient multimodal journeys. The two companies will also be working together to offer the integration to new cities around the globe.

Uber's new ticketing option will be powered by Masabi's Justride SDK, the first and only mobile ticketing SDK for public transit. The SDK allows 3rd party applications to request fare types, make payments, and deliver visual and barcode mobile tickets to a passenger through a secure ticket wallet.

PetaGene

PetaGene won Product of the Year Award 2018 for PetaSuite, at the Cambridge Computer Lab Ring Year Hall of Fame Awards held on April 11th 2018.

PetaGene is the market leader for genomic data compression, making genetic data smaller, faster and better, to enable a new era of Precision Medicine.

PetaSuite by PetaGene is the market leading solution for genomic data compression. Thanks to the rapidly declining cost of DNA Sequencing, the promise of Precision Medicine for cancer, rare diseases and diagnostics is moving out of research labs and into the clinic. Unfortunately, Precision Medicine budgets are rapidly becoming dominated by IT infrastructure costs. PetaSuite enables organisations to dramatically slash the cost of storing and collaborating with genomic data, by up to a factor of six times. Since winning Best of Show at Bio–IT World beating EMC, IBM, Dell and other top–tier vendors, it has gone on to provide its solution to leading organisations around the world, including AstraZeneca, the NHS, DNAnexus and many research institutions.

Questionmark

Questionmark, global provider of assessment technologies and portal solutions, has been awarded a "Best Advance in Assessment and Survey Technology" award by leading independent research and analyst firm, Brandon Hall Group as part of their Excellence in Technology Program.

Re:infer

Re:infer, has raised $3.5m in a round led by Touchstone Innovations.

The London–based artificial intelligence startup delivering cognitive automation also drew support from Crane Ventures alongside existing backers Seedcamp and Dr Jason Kingdon, an AI entrepreneur.

Re:infer's deep learning tech automates the interpretation of communications data and bridges the gap between humans and information technology systems.

Sophos

Sophos, a global leader in network and endpoint security, has been recognized at the PLC Awards 2017 as the winner of the Innovation in Technology award. Sophos is the first company in the PLC award's 32–year history to win this award in successive years. The award acknowledges Sophos' continued focus on innovation in next–generation endpoint and network security to drive growth and market success, since its initial public listing in 2015.

The PLC Awards for 2017 were extended to offer recognition to all publicly quoted companies on the main market, from fledging indices through to FTSE 100, for their success and outstanding achievements throughout the year.

Speechmatics

Speechmatics has launched Global English, an accent–agnostic language pack for speech–to–text transcription.

Global English (GE) was trained on thousands of hours of spoken data from over 40 countries and tens of billions of words drawn from global sources, making it one of the most comprehensive and accurate accent–agnostic transcription solutions on the market.

Job listing

April 2018

StarLeaf
- *Senior software engineer*
- *Software development engineer in Test*
- *Product manager*

Outschool
- *Principal software engineer*

Graphmasters
- *Software engineer*

Klydo
- *Machine learning engineer*

March 2018

Leanplum
- *Various (based in San Francisco, New York and Sofia, Bulgaria)*

Tech Marionette
- Sotware developer

InferStat
- *Data scientist (part–time contracting)*
- *Full stack developer (part–time contracting*

February 2018

Flagfast
- *Full stack developer (Python)*

Ensoft
- *Software engineer*

If you have a job advert that you would like included in the weekly listing, please send the details (as a word doc) to cam–ring@cst.cam.ac.uk

Part IB Group Projects

Yifan Bai reports on Team Oscar's project Virtual World Generator

Introduction

Digitalisation and virtualisation are happening almost everywhere in people's daily life — taking a photo after the algorithm lecturer's notes spanned the blackboards, or holding a phone turning several rounds of 360 degrees to contribute to the Google Street View at a scenic spot. In the age when Virtual Reality is becoming more affordable and attracting public attention, it is natural to associate VR with the digitalisation of the real world, and envisage the day physical objects being populated with sensible size in the reconstructed virtual scene, where everything is virtual but sensed real.

More thoughts therefore arise — if the virtual scene can later become interactive with the semantics of objects, the Virtual World Generator can then be applied to the efficient production of numerous and various indoor environments for virtual training, gaming, and virtual tourism. For example, a SWAT officer is trained to investigate and clear the rooms in a building, which are all generated from real world and the trainee will only need to act on a VR motion platform to complete the tasks, for instance, picking up a ballistic shield, turning the door handle and pushing the door.

The ideal Virtual World Generator will allow the users to scan the scene with a commonly available device, e.g. a mobile phone, using the techniques of Computer Vision and Object Recognition, create a virtual world resembling the real one with all the objects appearing replaced with their virtual entities retrieved from the library. Additionally, an agent will be created, performing actions in the virtual scene with all the objects according to their semantics.

Due to the limitations on time and resources, the Project aimed to develop the Virtual World Generator with the problem restricted to a smaller context — focusing on the positions and sizes of objects placed on a table and developing the semantics of knives, apples, bananas, and laptops.

Investigated Work

VisualSFM: Visual Structure–from–motion [1] is a library with the input as a sequence of images taken by a monocular camera, generating the point cloud reconstructed. It utilises SIFT (Scale–Invariant Feature Transform) [2] to perform full pairwise image matching, from which it extracts the camera movement and rotation, thereby determining the 3D positions of the key points. As it is running full pairwise matching, the method is computationally expensive. A quick test on 283 images run with a modern machine took 86 minutes, which is not a feasible approach of the given problem.

LSD-SLAM: Large–scale direct monocular SLAM [3] is a real–time algorithm taking a new image at each time step, with its featureless tracking method applied to two images in adjacent time steps. After the tracking, the map (i.e. the 3D point cloud of the captured scene) is updated with the keyframe created/merged from the current image.

Deeper Depth Prediction with Fully Convolutional Residual Networks: the fully convolutional architecture [4] is based on Residual Network [5], with its novel use of up-projection layers, which achieved better performance both in terms of time and accuracy metrics on the NYU Depth v2 depth image dataset [6]. A quick integration test including the reconstruction step was done by Team Oscar, in which the generated 3D point cloud from its depth maps showed unsatisfactory 3D structures of the tested scene.

Semi-global Matching: This stereo–images–based method [7] matches two images with pixel–wise approach as well as an approximate global cost function as a constraint, thereby producing a disparity map for further operations to generate depth map.

Methodology

Team Oscar proposed a solution — STereo Image Capture Kit (STICK) to generate virtual objects from real objects placed on a table with a sequence of stereo images taken by two phones of identical model. The process consists of four stages:

1. Using the 3D–printed twin–phone case (Fig. 1) for two identical phones to stay in a stable relative position.

Figure 1: 3D–printed twin–phone case

2. Running the web application set up on Team Oscar's server, two cameras take photos simultaneously upon the message from the application and upload images to the server for processing in the next stage.

3. The list of stereo images is supplied to the pipeline, which runs between two adjacent frames in the time domain (illustrated in Fig. 2).

a. First, the key points will be marked by the SIFT [2] algorithm, among which a matching algorithm FLANN [8] runs to identify the matched key points in both images. The matched points are regarded as the projections of the single 3D point in the real world captured. The disparity due to the stereo images of the matched points is calculated, and the depth is calculated based on:

$$depth = \frac{baseline \times f}{disparity}$$

where baseline is the distance between two camera and f is the focal length.

b. Second, key points from the left image of the frame one time step before are taken to perform a three–way matching, after which the direct image alignment algorithm [9] is used to generate a vector describing the camera motion.

c. From the output vector consisting of the information of rotational and translational camera motion, the camera property matrix from the previous frame is updated by multiplying it by the matrices of translation and rotation.

d. The object recognition part happens in step d, where the object recognition library developed by Tensorflow Team [10] is applied and it outputs a list of tuples, each consisting of bounding rectangle positions, object class and recognition score. The list of recognised object is sorted in a decreasing order of the score of recognition, from which the objects above a predefined threshold are selected as the input to the next stage.

e. This step makes use of the object recognition boxes retrieved from step d, and considers them as the 2D projections of the objects in the 3D world to calculate the estimated world coordinate of the centres of the objects, given the information of depths in step a and camera positions in step c.

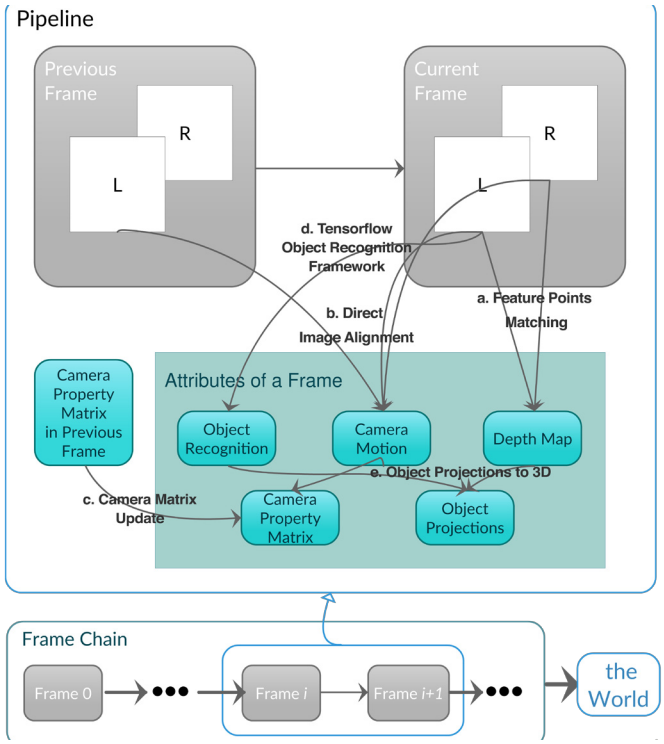

Figure 2: The Pipeline

4. The final step is the clustering of the 3D positions of object projections obtained in step 3.e. All object projections are first grouped according to their object classes (e.g. apple, laptop, etc.), forming groups of projections with each group containing only one class. The Bayesian Gaussian Mixture Model implemented by Scikit–Learn [11] is applied to the projections in each group, making clusters on their estimated centres in world coordinates, thereby deciding the number, positions and sizes of objects in the real world that the projections correspond to. The mean position and size of the projections are used in constructing the virtual counterpart of each identified real world

object. All the generated objects form a list, which is the complete list of the objects in the virtual world, and is sent to the Unity frontend as a JSON file. The Unity application runs on the client computer to populate the objects with the list, and the user can control the player on the client to interact with the virtual world.

Experiment Results

The pipeline described above is run with two iPhone 6 as the cameras, a desktop equipped with GeForce GTX 1060 graphics card for object recognition as the server to collect frames and perform the pipeline in step 3, and a laptop as the client running Unity application. The number of frames to capture was set to 40 and the capture rate set to one at every 250ms. From the end of capture to the successful population of objects in the Unity scene is around 40s on average, which to a large extent is the benefit of GPU support.

Figure 3 shows one of the best experiment results of STICK on a group of objects on a desk in the Intel Lab, where the relative sizes of the objects reflect the truth, and the relative position of the bowl and banana are also convincing, while the position of the laptop and the apple populated were different to actual ones in real world. However, STICK performs less satisfactory in the aspects of duplicate generation of the same object occasionally

Figure 3: The group of objects under experiment (l to the r: a laptop, a banana, a bowl, an apple, a spoon, a fork, a knife and a mug) and the result on the screen (Blurred areas are irrelevant)

The semantics of objects are implemented to a limited extent as shown in Figure 4 — objects can be picked up and put down to a different place; they can be held with the primary mouse button held down; an apple is cut into halves after it is touched by a knife and correspondingly a banana is peeled after the cuts.

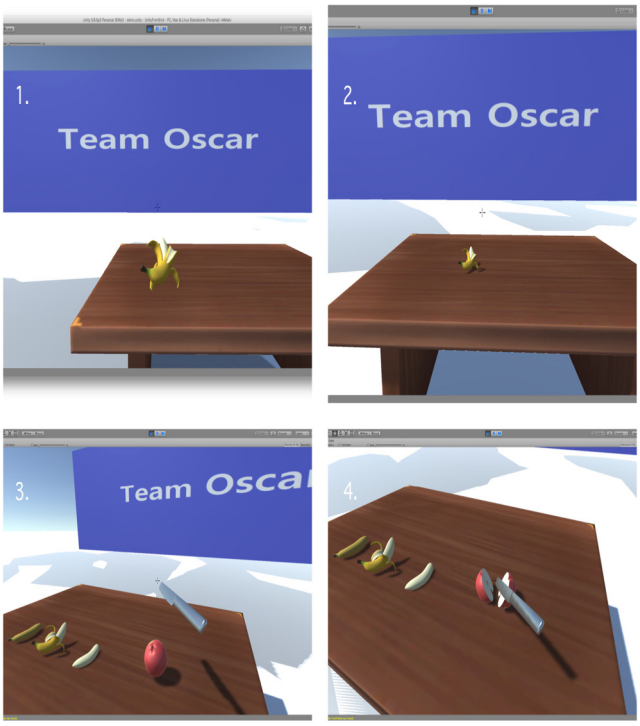

Figure 4: Semantics implemented: 1. Objects can be picked up. (A half-peeled banana.) 2. Objects can be put down. 3. Cutting a banana once makes it half–peeled and cutting a half-peeled one makes it fully–peeled. The agent is holding a knife, ready to cut the apple. 4. After the knife touched the apple, the apple became two halves.

Conclusion

The solution STICK is a proof of concept that using the principles and techniques of Computer Vision and Object Recognition, a cloud–based Virtual World Generator can be built with a twin–phone case, a pipeline calculating the depth map and object positions in world coordinate, and a Unity–based client visualising the real world.

Future improvements can be made to the objects clustering part to achieve better stability and to the Unity application to support the dynamic retrieval of assets and even their semantics as currently the effects of the knife and fruits are hard–coded.

References

[1] C. Wu. VisualSFM : A Visual Structure from Motion System. In ccwu.me/vsfm/doc.html.

[2] D. G. Lowe. Distinctive image features from scale-invariant keypoints. In International Journal of Computer Vision 60.2, pages 91–110, 2004.

[3] J. Engel, T. Schöps, and D. Cremers. LSD-SLAM: Large-scale direct monocular SLAM. In European Conference on Computer Vision, pages 834–849, 2014.

[4] I. Laina, C. Rupprecht, V. Belagiannis, F. Tombari and N. Navab. Deeper Depth Prediction with Fully Convolutional Residual Networks. In 2016 Fourth International Conference on 3D Vision (3DV), 2016.

[5] K. He, X. Zhang, S. Ren and J. Sun. Deep residual learning for image recognition. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 770–778, 2016.

[6] N. Silberman, D. Hoiem, P. Kohli and R. Fergus. Indoor segmentation and support inference from RGBD images. In European Conference on Computer Vision, pages 746–760, 2012.

[7] H. Hirschmuller. Stereo Processing by Semiglobal Matching and Mutual Information. In IEEE Transactions on Pattern Analysis and Machine Intelligence, volume 30, no. 2, pages 328–341, 2008.

[8] M. Muja and D.G. Lowe. Fast approximate nearest neighbors with automatic algorithm configuration. In VISAPP, 2009.

[9] J. Engel, J. Stückler, and D. Cremers. Large–scale direct SLAM with stereo cameras. In IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 1935–1942, 2015.

[10] J. Huang, V. Rathod, C. Sun, M. Zhu, A. Korattikara, A. Fathi, I. Fischer, Z. Wojna, Y. Song, S. Guadarrama and K. Murphy. Speed/accuracy trade–offs for modern convolutional object detectors. In IEEE CVPR, 2017

[11] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot and E. Duchesnay. Scikit-learn: Machine Learning in Python, Journal of Machine Learning Research, volume 12, pages 2825–2830, 2011.

Candidates for Part IB of the Computer Science Tripos are required to undertake a group design project as part of their practical work.

Team Oscar comprised Yifan Bai (T), Mansur Pasha (K), Sechyun Woo (M), Joe Yan (M), Charles Yoon (K) and Almos Zarandy (CHU).

Research Skills course

Richard Ngo: Review of 'Neural Acceleraton for General–Purpose Approximate Programs'

Programming is inherently far more precise than using natural language — as it has to be, because computers require such rigid instructions. We usually write software that tries to produce reliable and exact outcomes: loading a certain website, saving certain data in a spreadsheet, or autopiloting a plane or car. In each of these cases, even small errors might result in disproportionately large problems — one mistaken character might lead to an entirely different website or a budget which doesn't balance. But there are plenty of applications where precision isn't so important. For example, those which work with noisy inputs or outputs, or those for which there are a range of acceptable answers. The ever–increasing amount of data available to be analysed means that processor speed and energy efficiency have become bottlenecks in our ability to extract conclusions from such data. The field of approximate computing explores a trade–off which can be made in response to this problem: the sacrifice of precision for speed and efficiency.

In their paper 'Neural Acceleration for General-Purpose Approximate Programs', Esmaeilzadeh et al describe a way to speed up software by allowing programmers to label some regions of code as "approximable". Under their approach, the compiler then replaces these approximable regions (after an initial training period) with a neural processing unit (NPU), a dedicated chip which implements a basic neural net. This substitution is called a Parrot Algorithmic Transformation. The NPU will produce imprecise outputs based on observations from the training period. The authors demonstrate that Parrot transformations are able, in some cases, to produce significant decreases in running time and energy use, with decreases in accuracy of less than 10%.

These benefits can be attributed to two properties of NPUs: parallelism and simplicity. The former means that significant amounts of original code can be approximated quickly. The second means that NPUs can be run on dedicated hardware which consumes very little energy. Of course, these results depend on appropriate choice of approximable regions — in particular, they should be frequently–run code segments which use clearly defined inputs to produce clearly defined outputs without side effects. The authors note that automatic identification by the compiler of these promising regions could be a useful follow–up project.

One issue with approximate computing is that it makes programs more difficult to analyse rigorously (for instance, using formal verification methods). However, in the vast majority of cases such analysis doesn't occur anyway. Given how much data we currently have access to, the prospect of being able to analyse it more quickly makes approximate computing, and this paper in particular, promising areas of research.

The best essays from the Research Skills module of the MPhil in Advanced Computer Science course 2017/2018 are being published in 'The Ring'. This is the second of these essays.

Computer Laboratory news

Appointment of new Head of Department

The Department of Computer Science and Technology is delighted to announce the appointment of Professor Ann Copestake as new Head of Department. She has taken over from Professor Andy Hopper CBE, FIET, FREng, FRS, who held the post for almost 14 years.

Ann Copestake is Professor of Computational Linguistics. 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. Apart from DELPH–IN, her current research mainly concerns semantics.

Professor Ann Copestake

She has worked on a variety of application areas including scientific text processing, information extraction, augmentative and alternative communication (AAC), machine translation, Natural Language Interfaces, lexical acquisition and on tools for lexicographers.

Funding Successes

ERC Advanced Grant
October 1st2018 for five years, €2,473,844
Professor Peter Sewell: 'ELVER'

Engineering with Logic and Verification: Mathematically Rigorous Engineering for Safe and Secure Computer Systems

EPSRC Early Career Fellowship
September 1st 2018 for five years
Dr Hatice Gunes: 'Adaptive Robotic EQ for well-being (ARoEQ)'

The fellowship will equip humanoid robots with novel socio–economic intelligence and adaptation capabilities, grounded in the state of the art in affective computing, social signal processing, computer vision and machine learning. It will also create a novel, dynamic and adaptive socio–emotionally intelligent robotic framework for fostering human well–being, and generate a re–usable knowledge base of datasets, algorithms and evaluation metrics to enable benchmarking and standardisation in social HRI.

EPSRC TIPS (Trust, Identity, Privacy and Security) (with Aberdeen and Oxford).
October 1st 2018 for 28 months, £250K
Dr Jat Singh: 'Realising Accountable Intelligent Systems (RAInS)'
Dr Singh will explore issues of transparency and accountability in automated systems/IoT. The focus is on the means for better describing, auditing and controlling distributed/composite systems.

Innovate UK (collaboration with All Street)
April 1st 2018 to October 1st 2019, £200K to Cambridge
Professor Ted Briscoe: 'Virtual Investment Researcher'

The partners have previously developed a prototype virtual investment research system, SEWA (Structured Editing and Workflow Automation), which uses ML and AI to help human analysts to write investment reports. This project will move SEWA from prototype to viable product, by deploying a range of ML models embedded within an active learning infrastructure to improve SEWA's speed and predictive capabilities.

EPSRC Human–like Computing feasibility grant
(with University of Sussex)
1 April 1st 2018 – October 1st 2019, £303K
Dr Mateja Jamnik: 'How to (re)represent it'

Unlike humans, machines in general have fixed representations and do not have the understanding of the user. For example, sat nav systems will only give directions with elementary spatial commands or route planning functions, whereas humans give directions in many forms, for instance in terms of landmarks or other geographic features that are based on shared knowledge.

We want to model in computational systems the inherently human ability to choose or change appropriate representations, and make machines do the same. We want to find out what are the cognitive processes that humans use to select representations, what criteria they use to choose them, and how we can model this ability on machines. Our hypothesis is that when humans choose a representation of a problem, they use cognitive and formal properties of the problem and its representation to make their choice.

DARPA Microelectronics Technology Office (MTO)
Approx $829K to Cambridge over 36 months.
Dr Robert Watson, Professor Simon Moore and Professor Peter Sewell: 'CHERI–ISA Formal Verification (CIFV)'

In this joint project with SRI International, we will deploy formal ISA modelling and verification techniques developed by Peter Sewell's group to model and verify key security properties of the capability–based CHERI Instruction–Set Architecture (ISA). We believe that this approach is essential to the development, transition, and long–term maintenance of architectural security techniques, which rely on the ubiquitous enforcement of protections throughout an ISA as a foundation for software security.

DARPA
Approx $3M to Cambridge for 39 months starting December 2017
Dr Robert Watson and Professor Simon Moore: 'Extending the CHERI Architecture for Trustworthiness in SSITY (ECATS) Phase 1'

In this joint project with SRI International and ARM Research, we will explore how the capability–based CHERI protection model can be applied to full System–on–Chip (SoC) designs, blending novel hardware and software security approaches. We will deploy CHERI within 32–bit microcontroller architectures (using MIPS, RISC–V, and ARM–M), within heterogenous devices such as DMA engines and other compute types, and explore how capability–based protections can limit cross–SoC and I/O–based attacks. We will also explore development–tool approaches to rapidly deploying both new hardware protections within SoC designs, and also use and management of these features by SoC–facing software stacks.

EPSRC
November 1st 2017 – October 31st 2022, £512K
Dr Robert Watson, Professor Simon Moore and Dr Theo Markettos: 'Protection and memory safety for input/output security (IOSEC)'

Part of the new EPSRC/GCHQ Virtual Research Institute in Hardware Security, IOSEC is exploring novel hardware and software adversarial techniques and defensive approaches to attackers originating over I/O buses — especially those involving untrustworthy heterogeneous devices and peripherals with DMA access (e.g., via USB–C / Thunderbolt 3).

RAEng Frontiers of Development
February 1st 2018 – February 1st 2019, £15K
Dr Eiko Yoneki: BUGALERT: 'Pest and Disease monitoring in greenhouses with Raspberry-Pi Network'

Joint project with a Kenyan start–up to develop a smart greenhouse monitoring solution using image recognition to alert farmers to plant pests and diseases, using networked Raspberry Pis for data collection, and machine learning techniques to interrogate the images.

EPSRC (with QMUL)
January 1st 2018 – December 31st 2020, £840K to Cambridge
Dr Andrew Moore: EARL: 'sdn EnAbled measuRement for all'

This will use IXPs to demonstrate how SDN–enabled measurement–based network management can permit new Internet–wide services, through the creation of EARLnet, an operational, research–centred autonomous system. EARLnet will provide a real–world environment for real–time monitoring of network status and SDN–oriented research, and will also serve as a test–bed to develop and evaluate novel reactive network management solutions.

EPSRC GCRF Institutional Grant
£31.5K
Dr Jat Singh and Professor Simon Moore: 'People vs. bots — detecting social media manipulation in the developing world'

Using machine learning to identify sources of influence, bias and manipulation in the social media of developing countries.

EPSRC IAA Knowledge Transfer Fellowship
18 January 18th 2018 – January 17th 2019, £60K
Dr Richard Mortier and Dr Andres Arcia–Moret: 'IoT Spectrum Monitoring in Smart Cambridge'

To build a distributed, scalable spectrum data collection and visualisation platform, providing a suitable control plane to monitor the IoT network deployment in Smart Cambridge.

EPSRC IAA Knowledge Transfer Fellowship
February 1st 2018 – December 1st 2018, £60K
Professor Jon Crowcroft and Dr Carlos Molina: 'Tools for enforcement of Smart Contracts (TESCON)'

This project aims to evaluate the suitability of three smart contracts tools, after their alignments with a leading smart contract platform, for the enforcement of contracts used in data–trading applications developed by the Hat Community Foundation.

Cambridge Computer Laboratory hosts Inter-ACE cyber security challenge

More than 130 students representing 18 of the UK's top cybersecurity universities battled it out at the Inter–ACE 2018 cyber security challenge, hosted at the Computer Laboratory at the University of Cambridge.

The competition, supported by GCHQ's National Cyber Security Centre, and designed to attract the next generation of cyber security talent took place over two days on the 16th and 17th of March 2018.

The victorious team from the University of Edinburgh won the top prize of £6,000, with second place going to the University of Southampton and Imperial College London taking home bronze.

The winners will now compete with the best of the USA at C2C —'Cambridge2-Cambridge', a transatlantic contest jointly organised by the Massachusetts Institute of Technology (MIT) and the University of Cambridge to be held between the 29th of June and 1st of July 2018 at MIT's Computer Science and Artificial Intelligence Laboratory, Cambridge, MA, USA.

Now in its third year, Inter–ACE was established to help resolve the vast and growing cyber security skills gap, with an estimated shortfall of 1.8m workers worldwide by 2022[1]. Inter–ACE aims to inspire young tech enthusiasts into the cyber security sector, while also honing the skills of those who already have a strong aptitude for ethical hacking and helping them meet like-minded individuals and potential employers.

Professor Frank Stajano, Founder of Inter-ACE and Professor of Security and Privacy at the Department of Computer Science and Technology, University of Cambridge, said: 'It's no secret that the cyber security industry is suffering from a large and growing skills gap. We must do more to attract a more diverse pool of talent into the field. This is about demonstrating that careers in cyber security not only help to keep your country, your friends and your family safe, but are varied, valued and most of all fun.

There is still much more to be achieved, but I have been delighted over the last three years to be welcoming a growing number of female participants and contestants from increasingly diverse backgrounds to the two-day competition. We had 18 women competing this year, as opposed to just two when we started! It's working. There is no set profile for a cyber security professional and Inter-ACE contributes to reaching more people with that important message'.

Inter–ACE 2018 involved a number of different scenarios, including preventing a hack on a UK city's infrastructure and a tap on an undersea communications cable. Connected devices such as a children's toy were also used to demonstrate the impact of hacking techniques. The two–day event featured over 20 challenges in total, set by experts from the University of Cambridge and sponsors including Context IS and Palo Alto Networks.

Established through the UK's National Cyber Security Strategy and supported by GCHQ's National Cyber Security Centre, Inter–ACE is sponsored by Microsoft, BT, Palo Alto and Context IS.

The 18 universities that participated in this year's Inter–ACE were Queen's University Belfast, the University of Birmingham, the University of Cambridge, Cardiff University, De Montfort University, the University of Edinburgh, Edinburgh Napier University, Imperial College London, the University of Kent, Lancaster University, Newcastle University, the University of Oxford, Royal Holloway, University of London, the University of Southampton, the University of Surrey, University College London, the University of Warwick and the University of York.

[1] Frost & Sullivan "2017 Global Information Security Workforce Study Benchmarking Workforce Capacity and Response to Cyber Risk"

The Ring, Issue XLVIII May 2018
Web: http://www.cst.cam.ac.uk/ring/
E-mail: cam–ring@cst.cam.ac.uk
Tel: +44 1223 763585
Post: William Gates Building, Cambridge CB3 0FD
Published three times a year. Copy deadline for the September 2018 issue is August 10th 2018

The Ring is the journal of the Computer Lab Ring, which is the graduate association of the Department of Computer Science and Technology, University of Cambridge.