

UNIVERSITY OF CAMBRIDGE COMPUTER LABORATORY

A Proposal for Graduate Studentships



UNIVERSITY OF
CAMBRIDGE

A global competition

The competition for the best graduate students is global and fierce. The reality is that Cambridge loses many of its best applicants to peer institutions around the world.

A 2013 survey of those who were made an offer for graduate admission, yet withdrew, identified that 65% did so because they received a funded offer at another institution or could not fulfil the financial condition of their offer at Cambridge.

These are great minds of the future who are missing out on a Cambridge education. These are students who we need at Cambridge in order to stay at the cutting-edge of 21st century research.

The Computer Lab is seeking to build a studentship fund in order to provide the financial support which enables us to welcome the brightest students into our community and apply their enthusiasm and energy to the world-leading research we produce.

The Lab's success to date has largely depended on the pioneering discoveries of its graduate students. Today our students across eight research groups continue to re-define our understanding of Computer Science. It is our duty to make sure that their circumstances pose no barrier to them doing so.

Cambridge's Computer Laboratory has accumulated an impressive list of successes since its establishment in 1937. It was at the Mathematical Laboratory as it was then called, under Professor Maurice Wilkes' direction, that pioneering work led to the commissioning of the earliest complete computer EDSAC in 1949. It was also here that the world's first course in computing was delivered from 1953; that networked computing and the webcam were born; and that the creators of the web search engine, the programming language C++ and the Raspberry Pi learned their trade as graduate students.


Today, the Computer Lab continues to build on its long and distinguished history, and to uphold its reputation, by producing research which is unrivalled in terms of breadth, depth and quality. The Lab is also known for its entrepreneurial spirit, and for turning this innovative research into successful start-ups and spin-outs. Two hundred of the Cambridge cluster's 2,100 hi-tech companies have been started by Computer Lab staff and students. We have additional links and partnerships with a number of successful, influential companies, including Microsoft, Google, IBM, Alcatel-Lucent Bell Labs and Cambridge Assessment.

One thing which I am particularly proud of, and which these companies capitalise upon, is the calibre of our graduate students. At present there are around 50 MPhil students and a further 120 PhD candidates at the Lab. They are at the heart of our success, as drivers of the new and innovative research which we produce, and are integral to my vision for the coming years. On the following pages, you will see the ways in which the Lab's students are mindful of the real-world application of their research, and how they are using computer science as a tool for creating broad and deep impact.

And indeed, the societal value of Computer Science becomes more evident with each day that goes by. I think it is fair to say that when, in 1977, Ken Olsen – the founder and Chairman of Digital Equipment Corporation – said "there is no reason anyone would want a computer in their home", he had not foreseen the way that computers would transform our lives!

As a completely unique department – there is no other in the world quite like us - we must continue to attract the most intellectually capable, hard-working and determined graduate students to join us and exploit the societal purpose of computing. We must also ensure that financial or social circumstance poses no barrier to their ability to push forwards the boundaries of knowledge in Computer Science.

Professor Andy Hopper, CBE, FIET, FRS, FEng
Head of Department



In a short space of time, Cambridge's graduate students must move from learning information at undergraduate level to advancing knowledge and understanding at the cutting-edge of their field.

This is a significant challenge, and often a major test of character. But it is also why graduate education at Cambridge is the best education there is.

Through research, graduate students learn to challenge received wisdom and take the initiative. They must think creatively and laterally, show self-reliance and attention to detail. Most importantly, they must learn to back, to test and then to follow their intellectual convictions. All of these skills equip them to lead in every major profession and career.

Whilst Cambridge offers graduate students an unrivalled education and experience they are, in turn, essential to our mission. For that reason, we have committed to increasing the number we admit by 2% each year across the University, around 130 places. This is a major challenge, and will not be possible without major philanthropic investment.

We need to be able to offer competitive funding, early in the application cycle, to attract the best global talent. Programmes such as the Gates Cambridge Scholarship make a major difference, providing 95 fully-funded scholarships per year, but the need is much greater.

At the same time, more and more careers require specialist expertise that can only be acquired through graduate study. Whilst Cambridge can guarantee need-blind admission at undergraduate level, we cannot do so at graduate level.



We need to make sure that our graduate programmes, and in particular our taught masters programmes are open to all applicants, regardless of their financial circumstances.


Graduate Students at the Computer Lab

Cambridge has a culture of radical thinking, of risk taking and pursuing answers to the big unanswered questions. Nowhere is this truer than at the Computer Lab. This approach is a key enabler of our track record of making world-changing discoveries in Computer Science, and it is to this endeavour that we expect our graduate students to contribute.

The Computer Lab's influential partner companies are all keen to employ our graduates. More generally, the demand for Cambridge Computer Science and Technology graduates is intense – the annual recruitment fair at the Lab attracts dozens of companies, with many more wanting to attend but not able to due to a lack of space.

Graduate students play an integral role across the breadth of the Computer Lab's eight research groups – Artificial Intelligence; Computer Architecture; Digital Technology; Graphics and Interaction; Natural Language and Information Processing; Programming, Logic and Semantics; Security and Systems Research. Some of these students past and present tell us their stories and guide us through their pioneering research, on the following pages.





Clockwise from bottom left: Rana El Kaliouby speaks at TEDWomen2015; Eben Upton holding the Raspberry Pi and Steffen Lösch

remarked how the challenges of training computers to read faces mirrored the difficulties that his autistic brother had recognising emotions.

She set about teaching a computer to learn from facial expression. By the time she completed her PhD, she had built MindReader, a program that could track several complex emotions in relatively unstructured settings. She subsequently co-founded MIT's Media Lab spin-out Affectiva, and in 2015 was awarded the *Smithsonian* magazine's American Ingenuity Award for Technology, and was named as one of seven "Women to Watch" in 2014 by *Entrepreneur Magazine*.

Steffen Lösch was awarded his PhD by the Computer Lab in 2014, and his dissertation nominated for the 2015 BCS (the Chartered Institute for IT) Distinguished Dissertation Award.

The Computer Lab's Hall of Fame – or list of companies started by graduates and staff – currently stands at 246.

Each year, the Lab celebrates the success of these companies and recognises particularly promising examples. In 2013, Raspberry Pi was awarded the 'Company of the Year'. The company was founded by a team comprising PhD candidate Eben Upton, Professor Alan Mycroft, Senior Lecturer Rob Mullins, and entrepreneur and business angel Jack Lang.

The year-on-year decline in the number and skill level of students applying to read Computer Science at Cambridge served as the inspiration for the development of the credit card sized computer. Its intended market in the first instance was school children, in the hope that they would be taught core and hobbyist programming skills at an early age, and inspired to pursue programming later on.

By 2015, and having reached five million sales, the device had become the best-selling British computer, and is being embraced by a far wider audience, including members of the hacking, or security testing, community and the developing world.

It was Rana El Kaliouby's experience as a PhD student in affective computing at Cambridge which similarly inspired her to specialise in human-human communication for individuals on the autism spectrum. As she presented her PhD research on human-computer interaction, a fellow student

Steffen was a Gates Scholar and explains: *"this [scholarship] was crucial for me, my parents are small-scale farmers in Germany, and I couldn't have afforded to do my studies in Cambridge otherwise. I also received additional support after my son was born, for which I will always be grateful.*

My PhD extended a common verification technique to new programming features. It provides the basis for improving tools that automatically find programming errors, used in companies like Microsoft or Facebook."

I was drawn to Cambridge for my Masters by the reputation and spectrum of the course. I stayed on to do my PhD because of the quality of the research and the interdisciplinary contact in the colleges. My time here really gave me the opportunity to engage with smart and interesting people, and so broadened my horizons both within and outside my field. The skills I gained and the reputation of Cambridge meant that I have gained a stimulating position in industry - I now work as a Software Engineer for Charles Simonyi's Intentional Software Corporation in Seattle.

Scholarships are absolutely fundamental for attracting the best students to the Computer Lab. The students accepted to study at the CL in all honesty could choose any number of other institutions, many of which will provide some sort of compensation. Without scholarships, the CL simply couldn't compete – in my case I probably would have gone to university in Germany where I did my undergrad, and where most PhD students receive a salary."

Yoli Shavit

3rd year PhD candidate in the Artificial Intelligence Group, supervised by Dr Pietro Liò.



I was studying medicine as an undergraduate in Tel Aviv, when one evening I saw a friend doing his computer science homework. I was hooked - I'd known very little about the field previously, but soon switched to do a double major in life sciences and computer science.

After doing a Masters in Bioinformatics in the UK, and having really enjoyed how efficient and condensed teaching within the higher education system here is, I wanted to study for a PhD. And so I applied to Cambridge's Computer Lab, encouraged by its reputation as one of the best, and was fortunate to receive a place.

The major hurdle for me was money - I am a recipient of a Cambridge International Scholarship, but without it (and being subject to pay international fees), I would not have been able to come here to do my PhD. When I found out that I had been given a bursary the summer before I was due to commence my studies, I did three lightning-speed laps of the lawn of our Kibbutz in celebration! The neighbours may have thought I was crazy but I was overjoyed.

So now I'm in the final year of my PhD. My research is to interpret big genomic data. Up until recently, genomes were mostly considered in their linear form and spatial information was obtained with imaging microscopy only for few genes at a time. Today, new molecular data are becoming available so that we can map entire genomes in 3D and enable spatial and temporal navigation. This involves many computational challenges!

A large part of my work is developing algorithms to condense and translate huge volumes of genomic

data into information I can do something with - for example, spatial information. If we use a Google Maps analogy - this means that we can zoom in and out of a genome, and understand the position of genes in relation to others.

This technology has the potential to revolutionise personalised medicine. Imagine if a physician was able to look at your 'map' and develop a treatment just for you. This in turn could enable us to design better targeted drugs, which know exactly where to go; serve a function in the diagnosis of diseases such as leukaemia; or help us to better understand the process by which an 'open' genome integrates a virus.

Working within the Computer Lab has been great - as a Department it encourages independence, especially in the first year of a PhD, where students are given space and freedom to really think about what they want to do and achieve - to navigate your way around your research. This freedom encouraged me to really form my research questions and ideas, and more broadly independence.

Another benefit to studying in Cambridge is the concentration of expertise in the area - I did an internship at Microsoft Research here in Cambridge one summer, and throughout my studies I have been able to hop on my bike and visit or learn about the work of researchers in the Wellcome Trust Sanger Institute (genomics and genetics research institute); the Babraham Institute (life sciences research).

I'm also the Chair of Women at CL, which draws out and brings together the inspiring female academics and researchers in the Department, though only 15% of graduate students here are female. We hope to turn preconceptions many women seem to have of Computer Science as an esoteric or irrelevant discipline on their head. In truth, almost every industry and the companies within it will be touched by advances in Computer Science in the near future. Computer Science is a field which will make a difference, and we need women at the helm!

Matic Horvat

2nd year PhD candidate in the Natural Language Group, co-supervised by Professor Ann Copestake, and Professor Bill Byrne in the Department of Engineering.



Whilst studying for my undergraduate degree I became aware of the field of Machine Learning and Artificial Intelligence, but soon discovered that I wasn't able to pursue Natural Language Processing, the specific area which I was particularly keen to explore, at home in Slovenia. I got a scholarship to do an MPhil in Cambridge, and thanks to support and drive from the Department to secure funding, have been able to stay on to do a PhD.

My research draws on Machine Learning expertise within both the Computer Lab and the Department of Engineering. I particularly enjoy undertaking my research at both Departments, and combining the different approaches offered by each into a single one. In the context of Machine Learning, the Computer Lab takes a more theoretical approach, with a focus on understanding; whereas the Information Engineering division at the Department of Engineering takes a more pragmatic, problem-fixing approach.

I look at Natural Language Processing specifically – that is, making computers understand human language. At present we have speech recognition systems and speech synthesis, and question-answering systems are in their infancy. But machine translation remains a fundamental problem in Artificial Intelligence. We started to look at it in the 1940s and 50s, and it has evolved substantially – for example Google Translate – but there is still a lot of work to be done. I take a statistical approach to machine translation, in which computers take, and learn from, huge volumes of data. This approach has been refined over time - in the beginning, a programme would be able to translate a single 'unit

of meaning' (word) into a target language. Then the units widened, so that we were able to translate phrases; and then again to hierarchical-based translation, or nested phrases.

The next step was to add in linguistic approaches. Syntax was embedded, and what I am now doing is adding in another layer of semantics – so teaching computers to deal with the meaning of a sentence. For example – does 'bank' in a passage mean a riverbed or a financial institution? If the computer can understand the meaning and structure of a sentence, then there is the potential to vastly increase the accuracy of machine translation. The approach I'm taking hasn't really been tried before.

The goal of my research is to make translation better. Not all translations are equally difficult – for example translating from French and English is simpler than between syntactically distant languages such as Japanese and English. For this reason, I am currently looking at, and challenging computers to understand and translate Japanese and German, which offers some complicated sentence structures for computers to understand!

The translation industry is big business – we need only look at the amount the EU must spend on translation to get an idea of how an innovation such as this could have an impact. One day we may even be able to talk into our phones and have it translate our words into another language, but in our own voice.

Flora Ponjou Tasse

3rd year PhD candidate in the Graphics and Interaction (Rainbow) Group, supervised by Professor Neil Dodgson.



The reason I wanted to do [Computer Science] was because I really like movies! I remember being fascinated with the simulations in Jurassic Park and telling my father that I wanted to learn to do that.

When I finished my MS in Cape Town in 2011 I was accepted to Cambridge, but was unsure whether I'd be able to come, having struggled to find funding – but thankfully Cambridge offered me the International Scholarship in July. Then last year I was encouraged by my supervisor to submit an application for a Google Doctoral Scholarship, available to two people across the whole University. I was astonished that I had even been singled out to apply, leave alone to receive the award, which covers my tuition, a monthly stipend and a fund for travel and conferences. It has given me space and flexibility to allow me to focus on my research.

My research is into 3D shape retrieval from partial queries. My motivation for doing this is to open up the world of 3D printing, set to be a \$20bn industry by 2020, to the mainstream. Whereas 3D printers were very expensive when they first gained prominence, they are now becoming cheaper as they are mass produced. This removes one hurdle to making them accessible to everyone, but another exists – that it is difficult to find 3D designs upon which to base your 3D printing project.

My research aims to bridge this gap by enabling people to take a picture of something they want to reproduce – on their phones, tablets, cameras – then to upload it to a search engine, which will recognise the object based on its depth and 3D properties, and return 'hits' which are 3D designs enabling them to print it themselves.

So the ultimate goal is to extend the application of 3D printing to a much broader audience than it is currently available to, and to those who don't have the high-level expertise to use these technologies.

The innovation could enable us to replace broken parts or objects; or in commercial terms it could enable consumer companies to advertise and offer recommendations to customers based on what they have been uploading and searching recently. Over the course of my PhD I have seen more and more researchers in academic institutions change their minds and become interested in the research implications of my work.

The research also fulfils my artistic desire, which is becoming more prominent as I come to understand more and more about what goes into designing and modelling an object. As I code, I come to understand how shapes and designs are created. Perhaps my work will enable people to engage with graphics at home. In the future I would like to explore how artists think and the processes they go through as they design a product or piece of work.

Next I might pursue a PostDoc position, and explore ideas I have for start-ups on the side. Last week I attended the Cambridge Judge Business School's EnterpriseWISE programme for female entrepreneurs – I pitched my idea to a panel and they seemed to think it was a good one, so watch this space!

“My scholarship has given me space and flexibility to allow me to focus on my research.”

Advait Sarkar

2nd year PhD candidate in the Graphics and Interaction (Rainbow) Group, supervised by Professor Alan Blackwell and Dr Mateja Jamnik.



I'm a Cambridge-bred researcher – I did my undergrad and my Masters in the Computer Lab, and now I'm in the second year of my PhD here. Sometimes I get 'interdisciplinarity envy' when I look at my colleagues who have studied music or dance and have subsequently transferred into the Lab - but then I suppose you can't fault my commitment to Computer Science!

I chose Cambridge's Computer Lab because it's the best in the country for my field of work. I come from a particularly technologically-minded household in Bangalore, the centre of India's high-tech industry, so you could say I was exposed from an early age.

I've been fortunate to receive a Computer Laboratory Premium Studentship, and also an EPSRC Scholarship with an iCASE award – one in which a business contributes additional funding. In my case it's BT, and this funding enables me to purchase any specialist equipment, attend conferences and travel when needed. It also means that I benefit by gaining access to the company's datasets and contacts, and that my research – being based on this real-world data - will be applicable, relevant and will stand up to the real requirements of industry.

So – on to what my research is. In a nutshell, I look to produce simple interfaces for complex analytics. A simple interface can be as simple as a spreadsheet. And analytics too can be simple – we do analytics all day, every day without even realising. For example, when a teacher thinks about whether their students are performing well; when we are comparing loans online and deciding which to take; or when we are evaluating the viability of a home business and whether it's likely to go under.

But as analytics have evolved beyond averages and arithmetic to become something altogether

more advanced (involving statistical modelling and machine learning), so too must our ability and methods to perform them. At the moment, the interfaces too are complex, being based on programming languages (eg. R, SPSS, Python or MATLAB), and so are not accessible to the non-expert. I am looking to simplify the interface so that this can change and complex analytics can be interpreted and manipulated by non-experts.

I compare it to car maintenance at the turn of the 20th century, when car ownership entailed understanding the structure and mechanics of the car. Whereas nowadays we have discovered what is safe and useful to abstract - so that the complexity has been removed and all the driver need do is drive. I'm trying to see if the same can be done for analytics.

Among the visual analytics tools I have developed is 'Gatherminer', which compactly displays large volumes of time-series data – for example the data generated by an army of BT engineers testing cables and devices, hubs, or routers on a given network.

Capturing the data is only the first, and a small part of, the problem – next we need to identify patterns and explanations, which up until now has been a process which requires a lot of time, expertise, and the guidance of an expert analyst. Gatherminer first uses computing power to gather together similar datapoints, a process which makes certain types of patterns visually stand out. Next, the user can select these groups of points and ask the system to 'explain' what they have in common, which it does by using an algorithm to dig deep into the data.

Now we have intuitive visual representation we can work with and understand something from. We are now planning to implement this tool for use by BT – it's gone from paper to desk faster than any prototype I know of.

We recently had a group of users with no training in computing or statistics whatsoever participate in experiments evaluating my software. One involved predicting and deciding whether they should 'buy or sell' a used car using a historic dataset of good 'buy/sell' decisions. They were very able to perform the complex analytics, which I think proves that the simple interfaces I have set out to develop are fit for purpose!

Raoul-Gabriel Urma

Final year PhD candidate in the Programming, Logic and Semantics Group, supervised by Professor Alan Mycroft.



There are three parts to my involvement with the Computer Lab – not only my PhD, but also a book I have published, and the Coding Academy which I founded. It's been a busy few years!

My research looks at the evolution of programming language. Programming languages change as natural languages (for example French, Spanish and English) do, except that the timescale is rather different. Rather than evolving over the course of centuries, programming languages change constantly and annually.

Programmers help us to make use of these, but you can imagine that it places huge pressure on them to keep up. Older versions become obsolete and unsupported – if there's a bug then it won't be fixable since no-one will have the skills to do so anymore.

If we look at Google as an example, which uses half a billion lines of code across its five main programming languages - hundreds of staff members they really have to maintain or update all the time to keep it secure. My PhD is about documenting this issue, and finding techniques to help developers do this. And this involves taking on a lot of the subtleties and problems which exist when updating code.

My book was a collaboration with my supervisor. I would like to emphasise that I am eternally grateful to him for everything. You certainly form strong relationships when completing a PhD, and Professor Alan Mycroft has become one of the most important people in my life! Anyway, the book is about Java 8. Java is the most popular programming language in the world, and version 8 is the latest. The book aims to help us use the latest features and therefore to write code more effectively. I presented the work at the Oracle (who wrote the Java code) Conference in

San Francisco – book signings and cocktail-drinking were not things I had foreseen my PhD leading to!

The unique value of studying at Cambridge for me was being able to work, from day one, on whatever I wanted to. Usually you would be expected to work on an existing project, but I didn't have to do that here, which is really fantastic. There are also great opportunities here - for example the Cambridge Judge Business School's Start-Ups Accelerators and Incubators, which are all free to access. I participated in the Enterprisers programme recently. The [Cambridge] network is also massive, so you have a chance to make lots of connections. You can accelerate ideas like crazy! If you want to do it, it's there for you.

The final strand of my work is as co-founder of the Cambridge Coding Academy. This is our vehicle for empowering the next generation to make their ideas a reality. We all have ideas, but not the acumen or the expertise to realise them. The Academy teaches people in three 'segments' – 14-18 year olds; university students and early-career academics; and people with a basic level of programming experience to coding skills and to go on to prototype tangible products.

It does this by hosting workshops and mentors – I love teaching and host many of the workshops myself, and come up with new courses – for example on Machine Learning or Data Analysis. At the moment we are looking at developing adaptive learning patterns, so that we can extend our workshops to anyone, online.

Fitting it all in isn't a problem – if you like what you do, it's easy. I'd rather spend time doing things I enjoy rather than 'relaxing' or doing nothing. I finish my PhD this summer – and after that I will focus on the Coding Academy.

My goal is to turn it into a sustainable business, which reaches millions of people in the next three years.

We need your help in order to attract the most brilliant young graduate students from across the globe to Cambridge's Computer Lab. Their research will underpin our ability to transform understanding both of Computer Science and, in turn, the world. We are seeking partners to invest in students so that they can pursue their passion for further study, regardless of their social or financial circumstance, and as the next generation of leaders and innovators in Computer Science.

We are building a **flexible studentships fund** which will enable us to attract and invest in our students, and allocate funding where it is needed most.

As an indicator, £800k will endow one PhD studentship in perpetuity. It is our ambition to establish one endowed PhD studentship to accompany each Professorship within the Department. This will ensure that research emerging from Cambridge in strategically important fields of Computer Science will continue far into the future.

Alternatively, £150k funds one student through three years of PhD research.

For more information, please contact:

Professor Andy Hopper CBE, FIET, FRS, FEng

Head of Department
The Computer Laboratory

t: +44 (0)1223 334607
e: ah12@cam.ac.uk

Chris Chaney

Head of Major Gifts - Schools Based

t: +44 (0)1223 764482
e: chris.chaney@admin.cam.ac.uk

