Personal Statement

Until Caltech, the underlying theme of my life had been movement – a state of flux born of multiple changes of scenery before my 18th birthday. From the snowy winters of Rhode Island to the rejuvenating monsoons of India to the unrelenting heat of the Texas sun, each new move has made me independent, adaptable, and unafraid to try new things. I have tried not to let the prospect of readjusting to a new life deter me from speaking up and finding my niche, and I continue to imbue this spirit and determination in each new challenge I face. Fleeting memories of a brief 6-month stint in London as a toddler and a longer 2-year sojourn in Virginia sometimes remind me that my formative years were full of oddly disjointed instruction, but I have taken this eclectic mix of a background and made it my own.

Contrary to previous experience, my time at Caltech has been characterized by routine. I follow a weekly schedule of lectures and problem sets, meticulously optimized in the first fortnight of each term and diligently attended to for the remainder. I take classes with a set group of friends, eat dinner at a specific table most nights, and plan and participate in House activities with familiar faces. I love the steadiness and rhythm that I have come to attribute to this campus, and the knowledge of the unique space I occupy in our small undergraduate community. However, this routine has also led me to become entrenched in the "Caltech bubble". I want to take study abroad as an opportunity to reclaim my sense of independence and regain a willingness to take risks – whether it be by signing up for classes that I have little background in or by joining clubs and intramural sports teams that I have never considered before. I look forward to experiencing a different lifestyle from the one I have been used to in suburbia – both a college town campus and a city campus afford many avenues for exploration. The stunning architecture, history, and culture of both Cambridge and Edinburgh add a new dimension to the learning experience – I plan to take full advantage of my time abroad by acquainting myself as well as possible with the places I live in and visit. Lastly, and perhaps most importantly, is the university community; with populations many times the size of Caltech's own and a much wider variety of majors and club activities, I hope to step out of my comfort zone as a Caltech CS major and revel in each unfamiliar interaction.

Academically, I am looking forward to being exposed to both new topics in Computer Science and fresh perspectives on the ones I've already learnt about at Caltech. The sheer range of courses available from which to choose is astounding, and gives me the chance to learn specialized topics that I certainly would not find elsewhere. In addition, Caltech's CS curriculum has traditionally leaned towards being theoretical, with an emphasis on underlying principles and mathematical backgrounds. I am eager to apply the academic rigor I have learnt here to more practical, hands-on courses in order to be fully prepared for any direction I may choose to go in with my career. I also think that the inevitable presence of common coursework across Caltech and European universities will lend itself to very interesting discussions about both the fields of study and their applications – I am eager to gain an understanding of how faculty and students abroad might view current topics at the forefront of the American CS industry, from the machine learning boom to user data privacy issues. I also hope that the diverse academics on study abroad university campuses allow me to interact with and learn from students beyond CS. I am very interested in international politics and economy (and am attempting to add the BEM major at Caltech!), and aim to broaden my perspective in these fields during my time abroad.

The thought of living on my own in an unfamiliar location is daunting, yet exhilarating. It comes with the chance to explore new ideas, diverse cultures, and unfamiliar routines, and maybe even the opportunity to come back with a new routine of my own.

Cambridge Michaelmas Program Fit

As a Computer Science major, the Cambridge Michaelmas program has some of the most pertinent elective classes I could hope to find for my interests, besides providing a uniquely individualized learning experience in the form of supervisions. Each of the courses I am planning to take enhances an aspect of my Caltech coursework or introduces me to new topics. I am currently completing Caltech's Machine Learning sequence, and taking Cambridge's Natural Language Processing class will be a very interesting related field of study, besides being one I have always wanted to learn more about. The Principles of Communications class covers topics that I have never been exposed to before, and should be a very appropriate precursor to CS144 (Networking) at Caltech, which I plan to take the term after studying abroad. Two of the classes I have listed – Concurrent and Distributed Systems, and Cloud Computing – are electives that I chose to learn more about. I am looking forward in particular to the Cloud Computing course, since this is a field that is fast becoming essential in the computer science industry as well as any research areas dealing with large amounts of data. My final course is Business Studies, which should complement to the practical courses above and provide a valuable perspective on what goes into turning code and theory into a product. It is also especially pertinent in my case since I am also attempting to complete the Business, Economics, and Management major at Caltech, and so this class should be a valuable interdisciplinary perspective. Finally, apart from the coursework, I believe the Cambridge supervision system will be one in which I thrive, since I have always learned better by conversing with professors and TAs about their thoughts on the material in addition to regular lectures. I hope that attending in the Michaelmas term will allow me to approach the time abroad with a refreshed mindset after the summer, eager to interact with and learn from professors and upperclassmen with very different experiences from those found at Caltech. It will also give me the chance to adjust well to Cambridge life through Freshers' Week, find friends who are as new to the location as me to explore with, and join active clubs related to my various interests (as well as be introduced to things I haven't tried before!).

Cambridge Michaelmas Proposed Course List

Total CIT Units for term abroad: 36
Course by Correspondence/Units: 9 (BEM 105)

1. Natural Language Processing

Tripos: Computer Science
Part: II
Term: Michaelmas
Number of lectures: 12
Lecture times, if available: M, W, F 10 am
Caltech units: 6
Caltech evaluator: Thomas Vidick
Type of Caltech credit (option, general, etc.): Option
State CIT equivalent course, if applicable: N/A
Course description (paste in paragraph):

This course introduces the fundamental techniques of natural language processing. It aims to explain the potential and the main limitations of these techniques. Some current research issues are introduced and some current and potential applications discussed and evaluated. Students will also be introduced to practical experimentation in natural language processing.

By the end of the course students should:

- be able to discuss the current and likely future performance of several NLP applications;
- be able to describe briefly a fundamental technique for processing language for several subtasks, such as morphological processing, parsing, word sense disambiguation etc.;
- understand how these techniques draw on and relate to other areas of computer science.

2. Principles of Communications

Tripos: Computer Science
Part: II
Term: Michaelmas
Number of lectures: 16
Lecture times, if available: M, W, F 11 am
Caltech units: 9
Caltech evaluator: Thomas Vidick
Type of Caltech credit (option, general, etc.): Option
State CIT equivalent course, if applicable: N/A
Course description (paste in paragraph):

This course aims to provide a detailed understanding of the underlying principles for how communications systems operate. Practical examples (from wired and wireless communications, the Internet, and other communications systems) are used to illustrate the principles.

At the end of the course students should be able to explain the underlying design and behaviour of protocols and networks, including capacity, topology, control and use. Several specific mathematical approaches are covered (control theory, optimisation).

3. Business Studies

Tripos: Computer Science
Part: II
Term: Michaelmas
Number of lectures: 8
Lecture times, if available: M, W, F 12 pm
Caltech units: 3
Caltech evaluator: Thomas Vidick
Type of Caltech credit (option, general, etc.): Option
State CIT equivalent course, if applicable: N/A
Course description (paste in paragraph):

How to start and run a computer company; the aims of this course are to introduce students to all the things that go to making a successful project or product other than just the programming. The course will survey some of the issues that students are likely to encounter in the world of commerce and that need to be considered when setting up a new computer company.

At the end of the course students should:

- be able to write and analyse a business plan;
- know how to construct PERT and GANTT diagrams and perform critical path analysis;
- appreciate the differences between profitability and cash flow, and have some notion of budget estimation;
- have an outline view of company formation, share structure, capital raising, growth and exit routes;
- have been introduced to concepts of team formation and management;
- know about quality documentation and productization processes;
- understand the rudiments of marketing and the sales process.

4. Concurrent and Distributed Systems

Tripos: Computer Science
Part: II
Term: Michaelmas
Number of lectures: 16
Lecture times, if available: Tu, Th 10 am
Caltech units: 9
Caltech evaluator: Thomas Vidick
Type of Caltech credit (option, general, etc.): Option
State CIT equivalent course, if applicable: N/A
Course description (paste in paragraph):

This course considers two closely related topics, Concurrent Systems and Distributed Systems, over 16 lectures. The aim of the first half of the course is to introduce concurrency control concepts and their implications for system design and implementation. The aims of the latter half of the course are to study the fundamental characteristics of distributed systems, including their models and architectures; the implications for software design; some of the techniques that have been used to build them; and the resulting details of good distributed algorithms and applications.

At the end of Concurrent Systems portion of the course, students should:

- understand the need for concurrency control in operating systems and applications, both mutual exclusion and condition synchronisation;
- understand how multi-threading can be supported and the implications of different approaches;
- be familiar with the support offered by various programming languages for concurrency control and be able to judge the scope, performance implications and possible applications of the various approaches;
- be aware that dynamic resource allocation can lead to deadlock;
- understand the concept of transaction; the properties of transactions, how they can be implemented, and how their performance can be optimised based on optimistic assumptions;
- understand how the persistence properties of transactions are addressed through logging; and
- have a high-level understanding of the evolution of software use of concurrency in the operating-system kernel case study.

At the end of the Distributed Systems portion of the course, students should:

- understand the difference between simple concurrent systems and distributed systems;
- understand the fundamental properties of distributed systems and their implications for system design;
- understand notions of time synchronisation, including logical clocks, vector clocks, and physical time;
- be familiar with various approaches to data and service replication, as well as the concept of data consistency;
- understand the effects of large scale on the provision of fundamental services and the tradeoffs arising from scale;
- appreciate the implications of individual node and network communications failures on distributed computation;
- be aware of a variety of tools used by distributed-system creators, such as RPC and object-oriented middleware (OOM);
- be familiar with a range of distributed algorithms;
- be familiar with a number of case studies in distributed-system design including: the Network File System (NFS), the Network Time Protocol (NTP), Java Remote Method Invocation (RMI), CORBA, the AFS and Coda filesystems, Network-Attached Secure Disks (NASD), and Google's MapReduce, BigTable, and Spanner systems.

5. Cloud Computing

Tripos: Computer Science
Part: II

Term: Michaelmas
Number of lectures: 16
Lecture times, if available: Tu, Th 2 pm
Caltech units: 9
Caltech evaluator: Thomas Vidick
Type of Caltech credit (option, general, etc.): Option
State CIT equivalent course, if applicable: N/A
Course description (paste in paragraph):

This module aims to teach students the fundamentals of Cloud Computing covering topics such as virtualization, data centres, cloud resource management, cloud storage and popular cloud applications including batch and data stream processing. Emphasis is given on the different backend technologies to build and run efficient clouds and the way clouds are used by applications to realise computing on demand. The course will include practical tutorials on different cloud infrastructure technologies. Students will be assessed via a Cloud-based coursework project.

By the end of the course students should:

- understand how modern clouds operate and provide computing on demand;
- understand about cloud availability, performance, scalability and cost;
- know about cloud infrastructure technologies including virtualization, data centres, resource management and storage;
- know how popular applications such as batch and data stream processing run efficiently on clouds;
- know how to build and operate a testbed cloud.

6. Course by Correspondence: Options (BEM105)

Tripos: N/A
Part: N/A
Term: Michaelmas
Number of lectures: N/A
Lecture times, if available: N/A
Caltech units: 9
Caltech evaluator: Jaksa Cvitanic
Type of Caltech credit (option, general, etc.): Additional Social Science electives
State CIT equivalent course, if applicable: N/A
Course description (paste in paragraph):

An introduction to option pricing theory and risk management in the discrete-time, binomial tree model, and the continuous-time Black-Scholes-Merton framework. Both the partial differential equations approach and the martingale approach (risk-neutral pricing by expected values) will be developed. The course will cover the basics of Stochastic, Ito Calculus. Since 2015, the course is offered in the flipped format: the students are required to watch lectures online, while problem solving and case and paper presentations are done in class.

Please note: This course by correspondence request has been discussed with and approved by the Caltech instructor.

7. Quantum Computing (ALTERNATE to Business Studies)

Tripos: Computer Science
Part: II
Term: Michaelmas
Number of lectures: 8
Lecture times, if available: M, W, F 11 am
Caltech units: 3
Caltech evaluator: Thomas Vidick
Type of Caltech credit (option, general, etc.): Option
State CIT equivalent course, if applicable: N/A
Course description (paste in paragraph):

The aims of the course are to introduce students to the basics of the quantum model of computation. The model will be used to study algorithms for searching and factorisation. Issues in the complexity of computation will also be explored.

At the end of the course students should:

- understand the quantum model of computation and the basic principles of quantum mechanics;
- be familiar with basic quantum algorithms and their analysis;
- be familiar with basic quantum protocols such as teleportation and superdense coding;
- see how the quantum model relates to classical models of deterministic and probabilistic computation.

Schedule of courses for Computer Science candidates:

#	Course Title	Time	Lectures, Supervisions	Units	Part
1	Natural Language Processing	MWF10	12, 3	6	II
3	Principles of Communications	MWF11	16, 4	9	II
4	Business Studies	MWF12	8, 2	3	II
5	Concurrent and Distributed Systems	TR10	16, 4	9	IB
6	Cloud Computing	TR2	16, 0	9	II

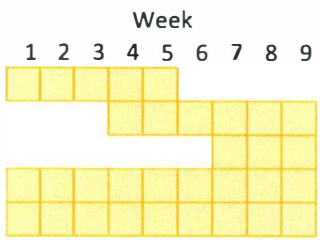

(ALTERNATE) **Schedule of courses for Computer Science candidates:**

#	Course Title	Time	Lectures, Supervisions	Units	Part
1	Natural Language Processing	MWF10	12, 3	6	II
3	Principles of Communications	MWF11	16, 4	9	II
4	Quantum Computing	MWF11	8, 2	3	II
5	Concurrent and Distributed Systems	TR10	16, 4	9	IB
6	Cloud Computing	TR2	16, 0	9	II