Computer Science and Philosophy

Outline of a Planned New Degree Programme at Oxford

Introduction

The computer is now ubiquitous, impacting throughout our social, cultural, and intellectual lives. Every academic field has been profoundly affected by it, many to such an extent that progress has become almost unthinkable without the assistance of computer technology. But we have so far had little more than a quarter century of the personal computer, yet the students we admit in 2012 (the centenary year of Alan Turing's birth) will be working past the three quarter century mark! The better they understand this computer-dominated world, the more effectively will they be able to benefit from, contribute to, and shape it.

All this would be reason enough for philosophers to take a serious interest in computers as a subject of study. But there are also far stronger reasons, deriving from the nature of this radically novel device. Philosophy, at least since Aristotle, has focused heavily on reasoning and its formalisation, taking particular interest in areas where "intuitive" human reason can be replaced by the rigorous application of rules. The computer brings to this task a new algorithmic paradigm, which holds the promise of combining the breadth of Philosophy – its ability to treat information of many different sorts from many different sources – with the solidity of Mathematics. Algorithmic approaches have thus become increasingly important to a wide range of fields that philosophers traditionally discuss, and in these fields, many of the sorts of questions raised (e.g. about how relevant information can be most faithfully represented, or appropriately reasoned with, or the potential scope and limits of such representation and reasoning) are very closely akin to those traditionally asked by philosophers.

Since Philosophy and Computing are such close intellectual cousins, it is not surprising that those with an interest and aptitude for one of them are frequently attracted to the other. Very many graduates of Philosophy have moved into employment in Computing, and several notable contributors to Computer Science have been philosophers at some stage of their career. The two subjects also nicely complement each other, in that Philosophy appeals to potential students as being creative and intellectually exciting, while Computing holds more prudential promise for the future.

One particular virtue of this combination is that *both* Computing and Philosophy can be studied at Oxford without requiring any previous qualifications in either subject. Students who like the idea of doing Philosophy with a broadly scientific focus can realistically apply, even if they have never previously studied Philosophy, Computing or *any* of the traditional sciences. ***The only prerequisite qualification is a good A-Level (or equivalent) in Mathematics.***

Links between Computing and Philosophy

As already sketched, the intellectual links between Computing and Philosophy are broad and deep. The development of the computer has provided philosophers with novel methods of exchanging and storing information, processing and analysing texts, and representing logical structures and arguments (e.g. graphically or dynamically). It has also opened some entirely new philosophical areas, such as:

- Philosophy of Artificial Intelligence;
- Philosophy of Artificial Life;
- Philosophy of Computation;
- Philosophy of Information.

Each of these includes a range of important and fundamental topics (e.g. distinguishing between simulation and performance, the frame problem, the possibility of non-biological life, the philosophical relevance of complexity, classification of algorithms, induction and verification, data and information,

the data/symbol grounding problem). But quite apart from these specialist concerns, computational methods and results are increasingly impacting on work in traditional areas of Philosophy, including:

- Epistemology (e.g. computational epistemology, procedural epistemology);
- Ethics and Aesthetics (e.g. computer and information ethics: privacy, intellectual property, constraints on autonomous/robotic systems etc.; models of cooperation, digital art, creativity);
- Logic (e.g. automated reasoning systems and logics, logic programming and representation);
- Metaphysics (e.g. digital or information metaphysics, emergence, causation, formal ontologies);
- Philosophy of Action (e.g. the nature of agency and autonomy as applied to machines);
- Philosophy of Language (e.g. much of computational linguistics / natural language processing);
- Philosophy of Mathematics (e.g. the status of machine-assisted proofs, computability);
- Philosophy of Mind (e.g. the computational theory of mind, computational models of reasoning, concept recognition, representation of cognitive states, neural representation, virtual reality);
- Philosophy of Science (e.g. the status and role of computer models, automated induction);
- Political theory and Economics (e.g. computational game theory, agent models of behaviour).

The last three of these areas all connect closely with the increasing use of computer modelling across the entire spectrum of scientific investigation, in all of the Physical, Engineering, Medical, Biological and Social Sciences. Understanding of such models, their status and behaviour, is therefore becoming ever more important for philosophers who wish to investigate a wide range of topics, from the creation of the Universe to the likely future impact of global warming, and from the nature of biological evolutionary processes to the prospects for new evolutionary, game-theoretic, or agent-centred approaches to Economics, Politics, and Morals. In some of these areas (e.g. reconstruction of the early Universe or global warming), the production of such models will be confined to specialists with highly detailed knowledge of the relevant domain. But in many areas that have the greatest traditional interest to philosophers – where the questions are broader and the patterns less specific (e.g. the general understanding of evolutionary or agent-centred processes) – it is entirely feasible for appropriately skilled philosophers themselves to contribute to the relevant modelling. Thus Computing can provide philosophers of the 21st century with a wonderful tool for massively enhancing a very traditional philosophical activity: trying out thought experiments that would be far too complex to assess unaided "on paper" or "in the armchair". This also imposes a discipline which, again, is entirely in the spirit of traditional philosophical endeavour. Implementation of ideas "*in silico*" requires explicit statement of all assumptions, and removes any risk that results will be swayed by wishful thinking or vagueness. In Dennett's words, philosophers equipped with computing abilities have available to them

> "the creation of prosthetically controlled thought experiments of indefinite complexity … kept honest by requirements that could never be imposed on the naked mind of a human thinker acting alone. … This is a great way of confirming or disconfirming many of the intuitions or hunches that otherwise have to pass as data … Philosophers who see this opportunity will want to leap into the field, at whatever level of abstraction suits their interests, and gird their conceptual loins with the simulational virtuosity of computers."

Thus computers and Computer Science provide not only a rich source of new topics for the philosopher to study, but also a potentially invaluable new tool. For more on all this, see www.philocomp.net.

Proposed Outline of the Degree Programme

The proposed degree programme is largely modelled – with appropriate modifications – on the Computer Science half of the "Mathematics and Computer Science" programme and the Philosophy half of the "Mathematics and Philosophy" programme. Hence readers wishing to find out more detail can investigate appropriate pages on the Oxford University website.

However a major innovation is a new course for first years, bringing together the two halves in a stimulating and conceptually exciting combination: "Alan Turing on Computability and Intelligence". This course will also provide an opportunity for the course leader (Dr Peter Millican) to get to know all the students well, and for students from different colleges to get to know each other. Various social events are planned during the year, in a similar spirit.

First Year

Students will spend roughly half of their time studying each discipline, and will be examined in both at the end of the year. The two halves are constituted as follows:

Computing Functional Programming, Design and Analysis of Algorithms, Imperative Programming 1 and 2, Discrete Mathematics, Probability

Philosophy General Philosophy, Introduction to Logic, Elements of Deductive Logic, Alan Turing on Computability and Intelligence

Second Year

The workload is again divided roughly in half, but only the Computer Science half gets examined at the end of the year (these are the "Computer Science Part A" examinations):

Computing Models of Computation
3 options chosen by the student [1]

Philosophy 2 Philosophy papers, chosen by the student, to be examined at the end of the Third Year

Third Year

Here students can choose whether to focus more on Computer Science or on Philosophy (with a 75%-25% split in either direction), or to continue to spend half their time on each:

Computing 2, 4, or 6 options chosen by the student [2]

Philosophy 3, 2, or 1 papers (respectively), chosen by the student (note that a Philosophy paper is double the weight of a Computer Science option)

At the end of the third year, students are examined on the Computer Science options done during the year (the "Computer Science Part B" examinations), and all of the Philosophy papers taken so far.[3] Assuming they pass, they have now qualified for the B.A. in Computer Science and Philosophy.

Fourth Year *(optional)*

Students can leave with a B.A. degree after the third year, but may choose to stay on one more year for a Masters degree: the MCompPhil in Computer Science and Philosophy. The fourth year is even more flexible than the third, enabling students to concentrate, if they wish, on just one of the disciplines. Students select any three "units", where a "unit" is either a Computing project, two advanced courses in Computing,[4] a Philosophy thesis, or a Philosophy paper (together with an extended essay).

---

[1] Likely second-year options include Advanced Data Structures and Algorithms, Compilers, Concurrency, Concurrent Programming, Formal Program Design, Object Oriented Programming, and Principles of Programming Languages.

[2] Up to two of these options can be taken from the second-year list. Likely third-year options include Computational Complexity, Computer Security, Computers in Society, Databases, Intelligent Systems I, Intelligent Systems II, Lambda Calculus and Types, Logic of Multi-Agent Information Flow.

[3] There is a very wide choice of Philosophy papers, the only significant constraint being that over the second and third years combined, at least two of them must be from the following list: Formal Logic, History of Philosophy from Descartes to Kant, Knowledge and Reality, Philosophy of Logic and Language, Philosophy of Mathematics, Philosophy of Mind, and *either* Philosophy of Science *or* Philosophy of Science and Philosophy of Psychology and Neuroscience. Virtually all of the many other Philosophy papers offered at Oxford will be made available to students of Computer Science and Philosophy.

[4] Likely fourth-year options include Automata, Logic and Games, Categories, Proofs and Processes, Computational Linguistics, Computer Aided Formal Verification, Game Semantics, Information Retrieval, Probabilistic Model Checking, Program Analysis, Randomised Algorithms, and Theory of Data and Knowledge Bases.