

Philosophy of Computer Science: A Course Outline

William J. Rapaport

**Department of Computer Science and Engineering,
Department of Philosophy, and Center for Cognitive Science
State University of New York at Buffalo, Buffalo, NY 14260-2000**

rapaport@cse.buffalo.edu
<http://www.cse.buffalo.edu/~rapaport/>

February 8, 2005

Abstract

There are many branches of philosophy called “the philosophy of X ”, where X = disciplines ranging from history to physics. The philosophy of artificial intelligence has a long history, and there are many courses and texts with that title. Surprisingly, the philosophy of computer science is not nearly as well-developed. This article proposes topics that might constitute the philosophy of computer science and describes a course covering those topics, along with suggested readings and assignments.

1 Introduction

During the Spring 2004 semester, I created and taught a course on the Philosophy of Computer Science. The course was both dual-listed at the upper-level undergraduate and first-year graduate levels and cross-listed in the Department of Computer Science and Engineering (CSE) (where I am an Associate Professor) and the Department of Philosophy (where I have a courtesy appointment as an Adjunct Professor) at State University of New York at Buffalo (“UB”).

The philosophy of computer science is not the philosophy of artificial intelligence (AI); it includes the philosophy of AI, of course, but extends far beyond it in scope. There seem to be less than a handful of such broader courses that have been taught: A Web search turned up some 3 or 4 that were similar to my

course in both title and content.¹ There are several more courses with that title, but their content is more accurately described as covering the philosophy of AI. The philosophy of computer science deserves more exposure at the university level. The UB course was popular (with an enrollment of just under 50), and the students found it valuable, not only for its coverage of topics in the philosophy of computer science, but also for the critical-thinking skills they learned (see §4.1). This article presents my ideas on what a course in the philosophy of computer science might look like.

Why teach philosophy of computer science? And why teach it in a computer science department rather than a philosophy department? As a professor of computer science with a Ph.D. in philosophy (and a previous career as a philosophy professor), I've long been interested in philosophical issues in computer science in general and artificial intelligence in particular. My colleague Stuart C. Shapiro in the UB CSE department had urged me to develop some philosophy courses for our students. Initially, I had resisted this, not being sure that such courses would be acceptable to my department or—more importantly—taken by enough students. Moreover, my colleague Randall R. Dipert in our philosophy department regularly offered an undergraduate course in the philosophy of AI, with which I didn't want to compete.

However, there were many metaphysical, epistemological, and ethical issues that I thought were of interest in the non-AI part of computer science, many of which have only recently begun to be examined in detail by philosophers and philosophically-oriented computer scientists, and many of which shed new light on classical topics in philosophy. This article surveys them and offers some interesting readings that deserve to be more well known. Moreover, a course such as this can serve as an introduction to philosophy for computer science students, an introduction to issues in computer science for philosophy students, a capstone course for senior undergraduate computer science students, or perhaps an overview course for beginning computer-science graduate students.

¹In particular, CD5650, Swedish National Course on Philosophy of Computer Science, at Mälardalen University (Sweden), coordinated by Gordana Dodig-Crnkovic [<http://www.idt.mdh.se/~gdc/PI-network-course.htm>]; Selected Topics in the Philosophy of Computer Science, at Tel Aviv University (Israel), taught by Eli Dresner [<http://www.tau.ac.il/humanities/digicult/english.htm>]; and PHI 319, Philosophy of Computing, at Arizona State University, taught by Bernard W. Kobes [http://www.asu.edu/clas/philosophy/course_descriptors.htm].

2 Syllabus

The course syllabus was organized around a set of questions whose various answers we examined during the semester:²

1. **What is philosophy?** In particular, what is “the philosophy of X ” (where X = things like: science, psychology, history, etc.)? [These questions are especially important to discuss in a course primarily aimed at computer science students, who might have misleading ideas of what philosophy is all about—or no idea at all.]
2. **What is computer science?** [Although the “final” answer to this question may simply be the extensional “whatever computer scientists do”, this is a reasonable issue to discuss, even if there is no intensional answer. The following subquestions indicate some of the interesting issues that this main question raises.]
 - (a) What is science? What is engineering?
 - (b) Is computer science a science? Or is it a branch of engineering?
 - (c) If it is a science, what is it a science of?
 - (d) Is it a science of computers (as some authors say)?
 - (e) What, then, is a computer?
 - (f) Or is computer science a science of computation (as other authors say)?
 - (g) What, then, is computation?
 - (h) What is an algorithm? Is an algorithm different from a procedure? Many authors say that an algorithm is (like) a recipe; is it, or are there important differences?
 - (i) What are Church’s and Turing’s “theses”?
 - (j) Some authors claim that there are forms of computation—often lumped together under the rubric “hypercomputation”—that, in some sense, go “beyond” Turing-machine (TM) computation: What is “hypercomputation”?

²I am grateful to Timothy Colburn, Randall R. Dipert, Eli Dresner, James H. Fetzer, Luciano Floridi, Bipin Indurkha, James Moor, Robert Stainton, and Chris Viger for (email) discussions on the questions that such a course might focus on.

3. What is a computer program?

- (a) What is the relation of a program to that which it models or simulates? What is simulation?
- (b) Are programs (scientific) theories?
- (c) What is an implementation?
- (d) What is software? How does it relate to hardware?
- (e) Can (or should) computer programs be copyrighted, or patented?
- (f) Can computer programs be verified?

4. What is the philosophy of artificial intelligence?

- (a) What is AI?
- (b) What is the relation of computation to cognition?
- (c) Can computers think?
- (d) What are the Turing Test and the Chinese Room Argument?

5. What is computer ethics? [This, like the philosophy of AI, is a vast question, deserving of its own course and having many textbooks devoted solely to it. For my purposes, I decided to focus on questions that don't seem to be the typical ones asked in such a course.]

- (a) Should we trust decisions made by computers?
- (b) Should we build "intelligent" computers?

The remainder of this paper surveys these topics, recommends suggested readings, discusses the sorts of assignments I gave, and presents some student reactions.³

3 Textbooks

Unfortunately, there is no textbook that exactly overlaps the above topics. Three possibilities were offered to the students as recommended texts: Luciano Floridi's *Philosophy and Computing* (1999), Timothy Colburn's *Philosophy and Computer Science* (2000), and Floridi's *Blackwell Guide to the Philosophy of Computing and Information* (2004). The first two are monographs offering the authors' points of view; there is nothing wrong with this, of course, but I preferred a more neutral

³The homepage for the course, with links to the complete syllabus, assignments, and other course webpages, is at [<http://www.cse.buffalo.edu/~rapaport/philcs.html>].

approach for the sort of course that I had in mind. Moreover, the topics covered in each of these had a relatively small intersection with my topics. The third book is an anthology, but—again—there was only a small overlap with my topics, and, in any case, I preferred that my students read primary sources rather than overviews. There are other sources, of course: A special issue of the philosophy journal *The Monist* (Vol. 82, No. 1, 1999) was devoted to the philosophy of computer science. The journal *Minds and Machines: Journal for Artificial Intelligence, Philosophy and Cognitive Science* is almost entirely devoted to philosophy of computer science broadly construed. And about half of the articles in the *Journal of Experimental and Theoretical Artificial Intelligence* are on philosophy of computer science. Finally, an excellent website, “Computational Philosophy”, is moderated by John Taylor [<http://www.crumpled.com/cp/>].⁴ In the sections that follow and in Appendix A, I recommend appropriate readings for the topics that we covered.

4 Topics and Readings

4.1 What is philosophy?

A typical advanced philosophy course in a philosophy department normally does not need to address the question of what philosophy is, but I felt that a course whose principal audience was computer-science students needed to. I suspect that many such students feel that philosophy is a “soft” subject where there are no answers, so everyone’s opinion is equally good.⁵ In contrast, I hoped to present to the students a view of philosophy as an analytical and critical discipline that could be of value to them.⁶

I began with a brief history of western philosophy, beginning with Socrates’s

⁴Pointers to these and other sources are at my course webpage “What is Philosophy of Computer Science?” [<http://www.cse.buffalo.edu/~rapaport/510/whatisphilcs.html>].

⁵This is based on a not unreasonable assumption that computer-science students tend to be “Dualists” who see (and fear?) philosophy as being a “Multiplistic” discipline. These are terms from William Perry’s (1970, 1981) “scheme” of intellectual and ethical development. For a quick online glimpse of Perry’s views, see my website, “William Perry’s Scheme of Intellectual and Ethical Development” [<http://www.cse.buffalo.edu/~rapaport/perry.positions.html>]. Roughly, “Dualists” believe that all questions have correct answers and that the student’s job is to learn these answers, whereas “Multiplists” believe that most questions have no known answers and, consequently, that everyone’s opinion is equally good. However, those are vast oversimplifications, and the interested reader is urged to consult Perry’s writings, or any of the other sources listed on my website. On whether there can be answers to philosophical questions and, thus, real progress in philosophy, see Rapaport 1982.

⁶In Perry’s terminology, philosophy is a “Contextually Relativistic” discipline, i.e., one that critically evaluates claims on the basis of evidence (the truth-value of a claim is “relative” to its evidential “context”).

and Plato's view of the philosopher as "gadfly" challenging others' assumptions. I offered my own definition of philosophy as the search for truth in any field by rational means (which might be limited to deductive logic, or might be extended to include empirical scientific investigation). And we defined the "philosophy of X" as the study of the fundamental assumptions and main goals of any discipline X.

I briefly covered some of the basic principles of critical thinking and informal argument analysis, including the following notions:

1. "argument" (a set of premises and a conclusion)
2. "premise" (a Boolean proposition used to support a conclusion)
3. "conclusion" (a Boolean proposition that someone tries to convince you of by means of a logical argument)
4. "valid argument" (an argument is valid iff it is impossible for the premises all to be true yet for the conclusion to be false; this semantic notion can also be supplemented with a syntactic one: an argument is (syntactically) valid iff it has the form of any of a given standard set of argument forms that are (semantically) valid, such as Modus Ponens)
5. "factual argument" (this is a non-standard, but useful, notion:⁷ an argument is factual iff all of its premises are true)
6. "sound" (an argument is sound iff it is factual and valid).⁸

I will have more to say about this in §5.2, where I discuss the course assignments, but I should point out that Computing Curricula 2001's "Social and Professional Issues" knowledge area includes the item "Methods and Tools of Analysis" (SP3), which covers precisely these sorts of argument-analysis techniques [<http://www.computer.org/education/cc2001/final/sp.htm#SP-MethodsAndTools>].

As a reading assignment, I asked the students to read at least one of a variety of brief introductions to philosophy (see Appendix A), and I listed Mark B. Woodhouse's *Preface to Philosophy* (2003) as another recommended textbook for the course.

⁷Learned from my former philosophy colleague, Kenneth G. Lucey.

⁸There are many excellent textbooks on critical thinking and informal logic, and, of course, it is the subject of many full courses on its own. A useful short introduction for a course such as this is Longview Community College's website "Critical Thinking Across the Curriculum Project" [<http://www.kcmetro.cc.mo.us/longview/ctac/toc.htm>].

4.2 What is computer science?

We began the first major section of the course by discussing the reasons one might ask what a discipline is: There are, of course, philosophical—primarily ontological—reasons. But there are also political reasons, especially in the case of a discipline such as computer science, which can be found both in (arts-and-)science faculties as well as in engineering faculties (sometimes in both at the same institution!), or even in its own faculty (either accompanied by other departments in, say, an informatics faculty or else by itself). Then, too, there is the question of the relationship between computer *science* and computer *engineering*.

We surveyed the following answers that have been given to the question “What is computer science?”:

- It is a *science of computers and surrounding phenomena* (such as algorithms, etc.) (Newell et al. 1967).
- It is the *study* (N.B.: not “science”) of *algorithms and surrounding phenomena* (such as the computers they run on, etc.) (Knuth 1974).
- It is the empirical study (“*artificial science*”) of the phenomena surrounding computers (Newell & Simon 1976; cf. Simon 1996).
- It is a *natural science*, not of computers or algorithms, but of *procedures* (Shapiro 2001).
- It is *not* a science, but a branch of *engineering* (Brooks 1996).
- It is the body of knowledge dealing with information-transforming *processes* (Denning 1985).
- It is the study of *information* itself (Hartmanis & Lin 1992).

Note that several of these (especially the first two) might be “extensionally equivalent” but approach the question from very different perspectives: Some emphasize the computer (hardware); others emphasize algorithms, processes, procedures, etc. (software), or even something more abstract (e.g., information). An orthogonal dimension focuses on whether computer science is a science or perhaps something else (a “study”, a “body of knowledge”, an engineering discipline, etc.). And, of course, the name itself varies (computer science, computing science, informatics, etc.), often for political, not philosophical, reasons.⁹

⁹Another (more recent) view is that computer science is the study of *virtual* phenomena (Crowcroft 2005).

4.2.1 Is “computer science” science or engineering?

The question of whether computer science is really a science or else is really a branch of engineering has been the subject of several essays. It has special relevance at UB ever since our former Department of Computer Science, housed in the Faculty of Natural Sciences and Mathematics, merged with several computer engineers¹⁰ from our former Department of Electrical and Computer Engineering to form a new Department of Computer Science and Engineering housed in the School of Engineering and Applied Sciences. This is not only confusing to read about, but has given rise to a certain identity crisis for both our students and faculty, and I thought it would provide interesting local color to an investigation of the nature of science and of engineering.

We first turned to the question “What is science?”, discussing both its goals (should it merely *describe* the world—as Ernst Mach thought (cf. Alexander 1967: 118f)—or *explain* it?) as well as the nature of its theories (are they merely instrumentalist, or realist?). We looked at debates over scientific method (is it experimental and cumulative, or does it proceed by paradigm and revolution?) and its branches (is mathematics a science?). The primary readings on science were selections from David Papineau’s “Philosophy of Science” (1996) and chapters from John G. Kemeny’s *A Philosopher Looks at Science* (1959).

We next looked at the history of engineering (Michael Davis’s *Thinking Like an Engineer* (1998), is especially useful), discussing engineering as applied science, as defined in terms of professional education, and as a design activity (Petroski 2003). And we looked at a definition of computer science as a new kind of engineering that studies the theory, design, analysis, and implementation of information-processing algorithms (Loui 1987, 1995).

4.3 What is a computer?—Part I

Insofar as computer science is the science (or study) primarily of computers, the next reasonable question is: What is a computer? This is a large topic, and I divided it into two parts.

The first part was a survey of the history of computers. I presented this in terms of two parallel goals: the goal of building a computing machine, and the goal of providing a foundation for mathematics. As I see it, these were two more-or-less independent goals that converged in the first half of the 20th century. (Whether or not this is a historically accurate way of looking at the matter is itself an interesting question; in any case, it is certainly a convenient way to organize the topic pedagogically.) Our discussion of the first goal involved the contributions

¹⁰Some of whom had their doctorates from departments of computer *science*!

of Babbage, Aiken, Atanasoff and Berry, Turing, and Eckert and Mauchly. The contributions of Leibniz, Boole, Frege, Hilbert, Turing, Church, and Gödel made up the overview of the second goal.

The history of computers is a large topic, and we did not spend much time on it. Consequently, the assigned readings were intended only to give the students a flavor of the main events. I prepared a website, “A Very Brief History of Computers”,¹¹ based on the IEEE’s “Timeline of Computing History”¹² and containing links for further information, and I asked the students to read O’Connor & Robertson 1998 (on Babbage), Simon & Newell 1958 (pp. 1–3 are also on Babbage), and Ensmenger 2004 (on the controversy over who deserved the US patent for the first computer).

4.4 What is an algorithm?—Part I

The other main answer to the question of what computer science studies is: algorithms. So, what is an algorithm? We began our two-part investigation of this by first considering what computation is. One informal, CS-1–style explanation proceeds as follows: A function f (viewed as a set of ordered pairs, or “inputs” and “outputs”) is *computable* means by definition that there is an “algorithm” that computes f , i.e., there is an algorithm A such that for all input i , $A(i) = f(i)$, and A specifies how f ’s inputs and outputs are related (or how f ’s outputs are produced by its inputs). Then an *algorithm for a problem P* can be characterized as a finite procedure (i.e., a finite set of instructions) for solving P that is:

1. unambiguous for the computer or human who will execute it; i.e., all steps of the procedure must be clear and well-defined for the executor, and
2. effective; i.e., it must eventually halt, and it must output a correct solution to P .¹³

It became an interesting exercise as we went through the semester to compare the different (informal) explications of ‘algorithm’, no two of which seem to be equivalent. This makes Turing’s accomplishment all the more interesting!

With this informal exposition in mind, we then turned to a careful reading of Turing’s *magnum opus*, “On Computable Numbers” (1936). There are several versions on the Web, though the most trustworthy is the one reprinted in Davis 1965. When I said “careful reading”, I meant it: We spent an entire 80-minute class doing a slow, “active”, line-by-line reading of as much of it as we

¹¹[\[http://www.cse.buffalo.edu/~rapaport/510/history.html\]](http://www.cse.buffalo.edu/~rapaport/510/history.html)

¹²[\[http://www.computer.org/computer/timeline/\]](http://www.computer.org/computer/timeline/)

¹³This is an adaptation of Stuart C. Shapiro’s informal characterization; personal communication.

could.¹⁴ I strongly recommend that all computer-science students (as well as computationally-oriented philosophy students, of course) do this at least once in their lives. In addition to being one of the most significant scientific papers of the 20th century, it is also fascinating, well-written, and contains many interesting philosophical insights. The students told me afterwards that this slow reading was one of the highlights of the course.

We also discussed the history of the mathematical investigation of the concept “computable”, and discussed the relationship of (1) Turing’s thesis that a function is (informally) computable if and only if it is TM-computable to (2) Church’s thesis that a function is (informally) computable if and only if it is lambda-definable (which is logically equivalent to being recursive and, of course, to being TM-computable).

Besides Turing 1936, I also asked the students to read Leon Henkin’s “Are Logic and Mathematics Identical?” (1962), which has a good discussion of the history of logic and the foundations of mathematics that led up to Turing’s analysis, and Gabor Herman’s “Theory of Algorithms” (1983), which discusses the informal notions of “algorithm” and “effective computability” and provides a good background for Turing 1936. I also especially recommend (to instructors, if not to students) Robert I. Soare’s “Computability and Recursion” (1996) for the clarity it brings to the history of the competing analyses of ‘computable’ (e.g., how Turing’s Thesis differs from Church’s Thesis).

4.5 What is a computer?—Part II

Armed with this background, we turned to philosophical questions surrounding the nature of computers. John Searle’s “Is the Brain a Digital Computer?” (1990) argues that *everything* is a digital computer (which seems to trivialize the question), and Patrick Hayes’s “What Is a Computer?” (1997) is a symposium that responds to Searle. Hayes’s own view is that a computer is a machine that can take, as input, patterns that describe changes to themselves and other patterns, and that causes the described changes to occur. (A related definition—a computer is a device that “change[s] variable assignments”—is offered in Thomason 2003: 328.) It turns out that it is surprisingly difficult to give a precise characterization of what a computer is.

A closely related topic for which a relevant reading did not appear till after the semester was over is the question of whether the universe itself is a computer (or whether parts of the universe compute; e.g., does the solar system compute Kepler’s laws?). On this, see Seth Lloyd & Y. Jack Ng’s “Black Hole Computers”

¹⁴On this sort of Talmudic, slow-but-active reading style, see Rapaport 2004, §5.

(2004). This issue also concerns the nature of simulation (see Rapaport 1998, Perruchet & Vinter 2002 (esp. §1.3.4), and §4.8.2, below).

4.6 What is an algorithm?—Part II

As hard as it is to define ‘computer’, the notion of “algorithm” is even murkier, despite the accomplishments of Church, Gödel, Kleene, Markov, Turing, Post, et al. Introductions to computer science often liken algorithms to recipes, and, indeed, there are clear similarities. But the differences are even more illuminating, given the informality with which recipes are usually presented. An interesting unpublished paper by Beth Preston (2000) suggests that recipes are more like specifications than they are like algorithms. And Carol Cleland has written a series of papers (1993, 1995, 2001, 2002) that explores the relationships between algorithms, recipes, and procedures, introducing a notion of “mundane” procedures (causal processes, including recipes), which are effective procedures that (she argues) are not TM-computable, since their effectiveness depends on the external world.

4.7 What is hypercomputation?

“Hypercomputation” is a name given by the philosopher Jack Copeland (2002) to the computation of functions that can’t be TM-computed. We briefly investigated Turing’s “oracle” machines, Putnam’s and Gold’s “trial & error” machines (Turing machines where it is the *last* answer that counts, not the first answer), Boolos & Jeffrey’s infinitely-accelerating “Zeus” machines, and Wegner’s “interaction” machines (such as automatic-teller machines or airline-reservation systems) (see Copeland 2002 for citations and other models). We also looked at Kugel’s (2002) thesis that Putnam-Gold machines may be needed for AI to succeed.

4.8 What is a computer program?

We focused on five aspects of this question: the nature of implementation, whether programs are theories, the nature of software (vs. hardware), whether software can or should be copyrighted or patented, and whether programs can be verified. Each is discussed briefly, below, with a digression on course evaluation.

4.8.1 What is implementation?

“Implementation” is a ubiquitous notion in computer science, but one that is rarely defined, and thus crying out for philosophical analysis. We say that programs implement algorithms, yet high-level programs can be implemented in machine

language. We say that particular data structures (e.g., arrays) can implement abstract data types (ADTs) (e.g., stacks), yet some ADTs (e.g., stacks) can be implemented in other ADTs (e.g., linked lists). Is implementation a relation between an abstraction and something “concrete”, or can it (also) be a relation between two abstractions? Is it an isomorphism, or a homomorphism? In rebuttal to Searle’s argument that everything is a computer (see §4.5, above), David Chalmers (1994) develops a notion of implementation as isomorphism. I have urged that implementation is best viewed as the semantic interpretation of an abstract formal system (Rapaport 1999). These issues were all touched upon, and I also used this opportunity to carefully develop the notions of syntax, semantics, and formal systems.

4.8.2 Are programs scientific theories?

Some computational cognitive scientists (e.g., Pylyshyn 1984: 76, Johnson-Laird 1988: 52) have claimed that cognitive theories are best expressed, not in the languages of statistics or mathematics, or even in natural language, but in computer programs. These programs, being simultaneously theories and models (or implementations of the theories), can then be executed, in order to test whether the theory is a good model of cognition. It has also been argued, of course, that such a program is more than merely a model or simulation of the cognitive phenomenon under investigation; some have argued that it actually exhibits the cognitive ability. As background, we also discussed the relationships between theories and models, simulations and “the real thing”, and simulations and emulations; philosophical theories of scientific explanation; and philosophical theories of scientific models. Relevant readings here also include Joseph Weizenbaum’s *Computer Power and Human Reason* (1976; Chs. 5 and 6 are on models and theories) and Herbert Simon’s *Sciences of the Artificial* (1996; Ch. 1, which discusses scientific theories, is also good reading for the question of whether computer science is a science).

4.8.3 What is software?

Introductory computer science courses often assume that the distinction between software and hardware is clear. Computer scientists (and philosophers) know otherwise. James Moor’s “Three Myths of Computer Science” (1978) points out the inadequacies of the usual “abstract” software vs. “concrete” hardware distinction, arguing that software is a computer program that is changeable by a person. This allows for the “software” to be “hard”wired, as long as it can be changed. The “software is abstract” point of view is well argued by Peter Suber (1988), who considers it to be “syntactic form” (and this ties in nicely with the

discussion of syntax vs. semantics in the section on implementation). Finally, Colburn (1999) views software as a “concrete abstraction”: It has a “medium of description” insofar as it is a text in a formal language (which is an abstraction), and it has a “medium of execution” insofar as it is implemented in circuits and semiconductors (which are concrete).

4.8.4 Interlude: Midsemester Course Evaluation and Course Correction

The previous topic brought us more or less to the midsemester point in the course. Borrowing an idea from my colleague Stuart C. Shapiro, I traditionally give a midsemester course evaluation. I strongly recommend this for any course: It is far more useful than an end-of-course evaluation that is not seen until the course is over and hence is of no use in improving the course that just ended. For this course, I asked two simple, open-ended questions: What aspects of the course would you like to see changed? and What aspects of the course do you especially like? The answers let me know what needed to be fixed and what was going well. I summarized the answers and posted a response to the course newsgroup.

For this course, the major complaint was the amount of reading. I told the students that I would try to comply with their request for less reading, but that there were just so many exciting things that I wanted them to read that I would compromise: From then on, I only assigned one (sometimes two) required readings for each of the remaining topics, per class session, but I recommended (sometimes strongly) other things to look at—if not now, then at their leisure after the semester was over. Thus, for example, instead of requiring the students to read Moor 1978 *and* Suber 1988 (which is a very long paper) *and* Colburn 1999 (which is philosophically challenging), I only *required* them to read Moor 1978 (which is well-written and also discusses other important topics), *strongly recommended* Suber 1988 (which is wide-ranging and has lots of things to think about), and recommended Colburn 1999. In lecture, however, I discussed all three.

I hasten to add that there were many compliments, too! Students were pleased with the topics and organization, and especially liked the writing assignments, which I discuss further in §5.2.

4.8.5 Can software be patented? Or should it be copyrighted?

The topic of whether computer programs are copyrightable entities or patentable entities¹⁵ is a fascinating one, because it combines legal, social, and metaphysical issues. We concentrated on the last of these, since it flows nicely from the previous topic of what software is.

¹⁵There is a third possibility: that they are trademarkable entities; we did not consider this option.

Here is the fundamental paradox: If a computer program is viewed as a written text, then it is, by definition, copyrightable. But the very “same” program, engraved on a CD-ROM and, hence, executable on a computer, can be viewed as a machine that is, by definition, patentable (as well as subject to legal limitations on exportation to foreign countries; see Colburn 1999). Yet, also by definition, nothing is both copyrightable and patentable. (Whether one *should* copyright or patent a program vs. whether programs should be “open source” is one of the interesting social issues that we did not have time to consider.)

We looked at the legal definitions of copyright and patent (available from various US government websites)¹⁶ and read a fascinating—and little known—essay by computer scientist Allen Newell (“The Models are Broken, the Models are Broken”) that appeared as part of a symposium on this topic in the *University of Pittsburgh Law Review* (1985-1986). Newell argues that we computer scientists need to devise better models—i.e., better ontological theories—of such computer-science entities as algorithms, programs, etc. In contrast, some legal scholars (e.g., Koepsell 2000) have argued that lawyers need to devise better methods of legal protection that better match the unique natures of computer software and hardware. The point in both cases is that there is a mismatch between computer-science entities, on the one hand, and legal forms of protection, on the other (or between computational ontology and legal ontology); something’s got to give.

4.8.6 Can programs be verified?

We ended our investigations into the nature of computer programs with an inquiry into whether they can be formally verified. There is a subfield of computer science and software engineering that looks into formal methods for proving program correctness (see, e.g., Gries 1981 for a classic treatment). Two philosophers have written essays that critique this approach. I am a firm believer in the value of such formal proofs (despite some very real limitations), and I have several times taught our department’s course on program verification. Consequently, I spent some time introducing some aspects of formal verification before turning to the criticisms.

Brian Cantwell Smith’s (1985) “Limits of Correctness in Computers” is, in my opinion, one of the most significant papers on all aspects—moral, legal, semantic, ontological, etc.—of the philosophy of computer science, and should be required reading for all computer science majors. Among other things, he argues that there is a gap between the world and our models of it and that computers are doubly

¹⁶For ‘copyright’, see the US Copyright Office Circular 1 at [<http://www.copyright.gov/circls/circ1.html#wci>]; for ‘patent’, see the US Patent and Trademark Office Glossary at [<http://www.uspto.gov/main/glossary/index.html#p>].

removed, relying on models of the models, yet must act in the real world.

The other critique is James Fetzer's explosive essay, "Program Verification: The Very Idea", that appeared in the *Communications of the ACM* in 1988 and that launched a vicious public debate on the pros and cons of verification. Briefly, Fetzer argues that *programs* can't be verified because you can't logically prove that causal systems won't fail; at best, you can verify an *algorithm*. Note that, in order to properly evaluate Fetzer's argument, you must have a firm grasp of the relationship of algorithm to program, which, by this time, my students were well-prepared for.

4.9 Philosophy of AI: *Could we build artificial intelligences?*

As I indicated above, the philosophy of AI deserves a full course to itself (see, e.g., Moulton & Voytek 1979, Rapaport 1986), and one of my motivations for creating a course in the philosophy of computer science (and not merely the philosophy of AI) was that there were many non-AI philosophical issues of interest. Nevertheless, the philosophy of AI is a proper part of the philosophy of computer science, it is my own area of expertise, and the students intensely wanted to discuss it.

I limited myself to two topics: the Turing Test and the Chinese-Room Argument. A case can be made that an excellent course on the philosophy of computer science could consist solely of close readings of Turing's two major essays: his 1936 paper on computability and his 1950 paper on whether computers can think. So, for this topic, we read Turing's "Computing Machinery and Intelligence" (1950) as well as the current (and probably perennially most popular) reply: John Searle's Chinese-Room Argument ("Minds, Brains, and Programs", 1980).

Turing 1950, as is well known, argued that a computer will be said to be able to think if we cannot distinguish its linguistic (hence cognitive) behavior from a human's. Searle 1980 proposed a now-classic counterexample that alleges that a computer could pass a Turing Test without really being able to think.¹⁷ (A good source for both of these, and related, papers is Shieber 2004; cf. Rapaport, forthcoming.) We closed this topic with my own attempt at a rebuttal of Searle (Rapaport 2000), arguing that syntactic symbol manipulation of the sort that computers do can suffice for semantic interpretation of the kind needed for computational cognition.

¹⁷Mention should be made that a very early version of Searle's thought experiment appears as a way of explicating Turing machines in Rogers 1959 (Part I, reprinted as Rogers 1969: 131, 133; based on a 1957 lecture).

4.10 Computer ethics

Our final topic was computer ethics. As noted above, and as with philosophy of AI, this is often the topic of full courses by itself and is the subject of numerous texts and anthologies. I gave a brief overview of (computer) ethics, based on Moor's "What Is Computer Ethics?" (1985). We focused on his claim that we need to have metaphysical and ontological theories of computers (in particular, their "logical malleability") and related phenomena in order to answer ethical and social questions about their nature and use.

I chose to concentrate on two issues that are not often covered in such courses or books: Are there decisions that computers should never make? and Should we build artificial intelligences?

We turned to Moor's "Are There Decisions Computers Should Never Make?" (1979). One of his main points is that there are no decisions computers shouldn't make, at least as long as their track record is better than that of humans, but it's up to us to accept or reject their decisions. An interesting contrasting opinion is that of Friedman & Kahn's "People Are Responsible, Computers Are Not" (1992), which argues that there *are* decisions that computers should not make, because only humans are capable of being moral agents. But "to err is human", and we looked at a recent case of an airline crash caused by following a human's decision instead of a computer's (as reported in George Johnson's "To Err Is Human", 2002).

On ethical issues in AI, we read Michael R. LaChat's "Artificial Intelligence and Ethics: An Exercise in the Moral Imagination" (1986). First, I outlined the plot of Stanislaw Lem's "Non Serviam" (1971)—which should be required reading for all researchers in artificial life!—in which what we would today call an ALife researcher is forced to pull the plug on his creations when his research grant ends. LaChat considers whether such research shouldn't even *begin*, but that, nevertheless, considering the possibilities enables us to deal with important issues such as: What is a person? Would an AI with personhood have rights? Could it be moral?

4.11 Philosophy of computer science: A summary and a unifying theme

In closing the semester, I asked the students to read two recent overviews of issues in the philosophy of computer science, as a way to gauge what they had learned: Matthias Scheutz's "Philosophical Issues about Computation" (2002) and Smith's "The Foundations of Computing" (2002), and we reviewed the semester's readings and discussion, with an eye towards themes that connected the several topics.

One such theme that the students and I became aware of as the semester

progressed is the relation of an abstract computation to the real world. This theme is addressed explicitly in some of the papers we read, and is implicit in many others. It emerges in Cleland’s discussion of the causal nature of “mundane” procedures, which produce some actual product or physically affect the real world in some way. This is also one of Smith’s themes in his “Limits of Computation” essay, as well as an underlying reason of Fetzer’s arguments against program verification. It is, of course, the subject matter of implementation, and underlies the paradoxical nature of software vs. hardware, and hence the issue of whether software is copyrightable or patentable. I recommend an exploration of this theme as a unifying idea for a future course in philosophy of computer science.

5 Assignments

5.1 A difficulty

I wanted the students to do a lot of reading and thinking. Thinking is best done by active reading (Rapaport 2004), discussion, and writing—lots of writing. There is a well-known drawback to assigning a lot of writing to students: The instructor has to read it all and, ideally, comment on it. When this course was first advertised, I expected about 10–15 students, in a small seminar-like setting. The first preliminary enrollment report said that 30 students had signed up. Thinking that they thought that this might be a “gut” course (“A philosophy course in a computer science department? Oh, this’ll be easy to ace!”), I posted a note to the undergraduate newsgroup spelling out the large quantities of writing that I would expect. Enrollment doubled to 60! It finally settled down at just under 50 students.¹⁸ Still, 50 ten-page term papers plus frequent short writing assignments during the semester was not a prospect that I looked forward to.

Nor could I rely on help from graduate teaching assistants or recitation sections (a problem I was familiar with from my days teaching at a primarily undergraduate institution). No recitation sections had been assigned to the course, since I had not

¹⁸The breakdown of student majors was as follows:

	CSE	PHI	Other
undergrads	72%	10%	12%
grads	75%	15%	10%
total	73%	12%	15%

CSE = Computer Science & Engineering majors; PHI = Philosophy majors; Other = students majoring in Biology, Economics, Electrical Engineering, Management, Management & Information Science, and Mathematics.

expected such a large enrollment. They would have been useful for discussion purposes, but that was not to be. I did have an excellent graduate teaching assistant, but he was a computer-science grad student, not a philosophy grad student (although he did have some undergraduate philosophy experience and was philosophically sophisticated), and, in any case, he had no recitation sections to lead. Consequently, he was of most use to me in keeping records, though he did hold office hours and students felt comfortable going to him for advice on writing.

But how to have students write a lot without having to read it all? And how to have discussions without special time allotted for them? Of course, faculty at undergraduate institutions face this problem all the time, unlike we faculty at research universities. And so I drew upon my experiences as a philosophy professor at an undergraduate institution with no TAs and no recitation sections.

5.2 A solution: Required, short position papers . . .

I assigned the students five 1-page position papers throughout the semester, roughly one every 2 or 3 weeks. A first draft of each assignment was due 1 week after it was announced. The day it was due we set aside for “peer editing” (adapted from techniques used in English composition classes; cf. Cho & Schunn 2004): Each student was asked to bring 5 copies of their position paper, one for me, one for themselves, and one each for 3 other students. I put the students into small groups of three or four “peers”, each of whom had written a response to the same assignment. I asked them to spend about 10–15 minutes on each paper, reading it, critiquing it, and making suggestions for improvement. The students were then given another week to revise their papers to be handed in for credit. To ease my burden of grading, I read and put copious comments on only about 40% of the papers for each of the 5 assignments; each student received at least 2 papers fully critiqued by me (the other 3 papers were recorded as being handed in).

Peer editing accomplished several goals simultaneously: The students had plenty of opportunities to discuss the material with each other. In fact, probably more students participated in these small groups than would have ordinarily spoken out in a large classroom setting (though such full-class discussions were encouraged, as well). Moreover, all students got multiple feedback on each paper, in addition to my feedback on a subset of their papers. Another advantage of peer editing in class is that I had the freedom (and responsibility) to walk around the room, listening to the student discussions and occasionally facilitating one or answering a question on a one-on-one basis.

The position papers were designed to supplement the lectures and readings, as well as to foster the students’ critical-thinking skills. In particular, the topics always involved an argument that the students were asked to evaluate in terms of

factuality (i.e., truth value of the premises) and validity (i.e., reasoning). Students were encouraged to present their opinions and to support them with reasons. As one example, Appendix B contains the instructions for Position Paper 1, on “What is computer science?”, and Appendix C contains the peer-editing guidelines for it.

As one student observed later, the argument-analysis format of the position papers made them somewhat easier to grade than an ordinary essay would have been. Since the students were required to examine a rigid structure of an argument, they had fewer “degrees of freedom” in writing their responses. Thus, grading such papers can be closer to grading a mathematical problem set than a typical essay. It also made grading somewhat more impartial and somewhat less controversial.¹⁹

5.3 ... And two optional assignments

In addition to the required position papers, there was an optional term paper, whose topic had to be approved by me in advance. I supplied a list of some possible topics, but I encouraged the students to explore areas of interest to them. As a default topic, a student could write a summary of the philosophy of computer science in the style of an encyclopedia article or else present his or her own answers to the syllabus questions (see §2).

An exclusive-alternative option was a final exam (students could do the exam or the term paper, but not both). This was a take-home, short-answer, essay-style exam, asking for analytic and evaluative summaries of the possible answers to the topic-questions.

5.4 A required reading journal

In addition, in order to provide evidence that the students were really reading the material, as well as to encourage them to read slowly and actively, I required them to keep a “reading journal”. For each essay they read, they were to copy interesting passages (or at least full references to them) and—most importantly—to provide their own comments on them and on the issues raised in each item read. (Suggestions on how to do this can be found in Rapaport 2004.) I collected these Journals at the end of the semester, and included them in the grade calculation.

Students who attended almost all classes and turned in a Reading Journal could get a C; students who did that plus all five position papers could get a B; and students who did all of that plus either the term paper or final exam could get an A. All but one student wrote the position papers. Over 80% of the students chose the exam/paper option, with about 70% choosing the exam option.

¹⁹For more thoughts on grading, and a “triage” theory of grading, see my website, “How I Grade” [<http://www.cse.buffalo.edu/~rapaport/howigrade.html>].

6 What the students did and didn't like

The students' favorite part of the course was the writing, peer-editing, and revising of the 1-page position papers: They enjoyed the discussions, the ability to revise (including an option to re-revise for a higher grade), and—most importantly—the skills they learned, and the practice they got, in critically analyzing and evaluating informal arguments.

They also appreciated the course website (see Appendix D). I began each new section of the course by putting up a webpage containing recommended readings for that topic. I then gave a quick overview in lecture about each of the readings. Students informed me that this was very useful because it provided a summary of what was to come, including the different positions that have been taken on each issue.

Here is what one student said, in an unsolicited email message I received after the course was over:

I'd like to thank you for putting together such a great course this semester. I'll admit, knowing very little about it, I never had much respect for philosophy in the past—but this course has provided me with an entirely new perspective. In fact, I'd say that I learned as much in your course as any other I've taken in my graduate career at UB (not to mention the fact that the skills I learned in [it] are far more transferable than the skills of the more esoteric CS courses). . . . I urge [you] to offer this course again in the future. It offers exactly the kind of breadth of education that the department needs to stress, and with its CS flavor, it can tap the interest of students who would otherwise blow it off. Thanks again for a great semester, and please consider making Philosophy of CS a regular offering :)

Another student observed that “I don't think there was a single student in the class whose critical thinking/writing/reading skills didn't improve as a result of taking this course.”

As noted above, the students' least favorite part of the course was the amount of reading. Of course, this is something that students almost always complain about, but, in this case, the complaint really was about the quantity, not the quality: By and large, they found all of the readings to be interesting and useful; their complaint was that they didn't have enough time to read them all as carefully as they (and I) would have liked. Fortunately, on the basis of the midsemester course evaluation, I found this out early enough to be able to do something about it. As discussed above, subsequent reading assignments were limited to at most two required readings, with suggestions for recommended (but optional) follow-up readings.

7 Conclusions

I believe this to have been a worthwhile course, both for me and—more importantly—for the students. It gave many of the computer science majors an option to think about many issues that they either hadn't thought of before or had thought about but had no venue for discussing. It also gave them an opportunity to (learn how to) think critically, and to find out what philosophy could be like. The philosophy majors, in addition, had the opportunity to learn some things about computers, computing, and computer science that they probably would not have come across in more traditional philosophy courses, as well as the opportunity to apply some of their philosophical skills and knowledge to a different domain.²⁰

²⁰I am grateful to my students Dima Dligach and Albert Goldfain and to my colleagues Peter D. Scott and Stuart C. Shapiro for comments on earlier drafts.

Appendix A: Recommended Readings

Note: For a complete list of suggested readings, of which the following form only a small subset, see the links on the course website at

[<http://www.cse.buffalo.edu/~rapaport/directory.html>].

Many of the readings are publicly available online; others may be available online by subscription or affiliation with a university library. This appendix is adapted from the course webpage on Reading Assignments [<http://www.cse.buffalo.edu/~rapaport/510/readings.html>].

Topic: What is philosophy? Read at least one of the following:

1. Woodhouse 2003, Chs. I–III (pp. 1–45). (A good introduction to what philosophy is all about.)
2. Colburn 2000, Chs. 3–4 (pp. 19–50). (A good survey of some of the history of philosophy that is relevant to CS.)
3. Plato, *The Apology*; various versions on-line. (Plato’s explanation of what Socrates thought philosophy was all about; a good introduction to the skeptical, questioning nature of philosophy.)
4. Audi, Robert (2001), “Philosophy: A Brief Guide for Undergraduates” (American Philosophical Association). (A good brief intro to what philosophy is and what its branches are.)

Topic: What is computer science? Read all of the following (preferably before the next lecture, since I plan on giving you more to read then:-), slowly (i.e., one sentence at a time) and actively (i.e., think about each sentence; make notes in your Reading Journal). Keep in mind that your main goal in reading these is to look for each author’s answer to our question; consequently, at least on a first reading, you don’t have to read the “irrelevant” parts quite as carefully. Below, I indicate which parts I think are of central importance for our present purposes and which parts you can just skim (i.e., read quickly):

- Newell, Perlis, & Simon (1967)
- Knuth 1974, at least §1, optionally §§2–4, also. (The more mathematically inclined may wish to read the whole thing :-)
- Newell & Simon 1976, pp. 113–116, 120, & “Conclusion” (pp. 125–126). (Skim the rest; we may read it later. For now, concentrate just on what they have to say about what CS is.)

- Denning et al. (1989). (Study pp. 9-12, 16ff; skim the rest (you can also just skim pp. 17ff).)
- Hartmanis & Lin 1992, “Computer Science & Engineering” (pp. 163–168) and ”Abstractions in Computer Systems” (pp. 168–174); skim the rest.
- Brooks 1996, pp. 61–64; skim the rest.
- Shapiro 2001

Topic: What is science?

Required:

1. Papineau 1996: esp. pp. 290–294, 298–308, 319–320; skim the rest.
2. Kemeny 1959: Intro., Ch.5 (“The Method”), Ch. 10 (“What Is Science?”)

Recommended: Popper 1962, Ch. 1 (esp. pp. 33–59); Hempel 1966, Ch. 1; Kyburg 1968, Ch. 1 (esp. pp. 1–7); Ziman 1968; Salmon 1984, Ch. 1; Rosenberg 2000, Ch. 1.

Topic: What is engineering?

Required: Davis 1998: Ch. 1 (pp. 3–17), pp. 25–28, pp. 31–37; skim the rest.

Recommended: Loui 1987; Brooks 1996; Petroski 2003.

Topic: What is a computer?

1. “A Very Brief History of Computers”; browse the linked websites. [<http://www.cse.buffalo.edu/~rapaport/510/history.html>];
2. O’Connor & Robertson (1998)
3. Simon & Newell 1958, pp. 1–3 (on Babbage); skim the rest.
4. Ensmenger 2004

Topic: What is an algorithm?

Required:

1. Henkin 1962 (Discusses the history of logic and the foundations of math that led up to Turing’s analysis.)
2. Herman 1983 (Discusses the informal notions of “algorithm” and “effective computability”; good background for Turing 1936.)
3. Turing 1936

- Concentrate on the informal expository parts; the technical parts are, of course, of interest, but are rather difficult to follow and incorrect in many parts, and can be skimmed.
- In particular, concentrate on §§1–6 (study the simple examples of Turing machines carefully; skim the complex ones) and §9, part I (which elaborates on what it is that a human computer does).
- §7 describes the universal Turing machine; §8 describes the Halting Problem. You can skim these sections.
- If you get lost, try Suber 1997, Copeland 2004, or Copeland & Gordon’s AlanTuring.net website for “gentler” expositions.

Recommended:

1. Browse through the “Examples of Algorithms”
[<http://www.cse.buffalo.edu/~rapaport/510/whatisanalg.html>];
enjoy the cartoons :-)
2. Boehm & Jacopini 1966; Haugeland 1981
3. Soare 1996 (tough going in spots (you can skim those spots), but the rest is a good discussion and history of the competing analyses of “computable” (e.g., how “Turing’s Thesis” is different from “Church’s Thesis”).)

Topic: What is an algorithm?—Part II: Soare 1996, §§1–3, 4.5–5 (skim the rest)

Topic: What is a computer?—Part II: Searle 1990, Hayes 1997

Topic: What is a procedure? Preston 2000 (skim §1–2; read §§3–4; skim the rest); Cleland 1993

Topic: What is hypercomputation? Copeland 2002 (“Hypercomputation”), Kugel 2002

Topic: What is a computer program/What is implementation? Chalmers 1993a or 1993b, Rapaport 1999

Topic: What is a computer program/Are programs theories?

Required:

1. Read the quotes from Johnson-Laird 1981, Pylyshyn 1984, and Johnson-Laird 1988 carefully
[<http://www.cse.buffalo.edu/~rapaport/510/whatisacomprog.html>].
2. Weizenbaum 1976, Chs. 5, 6.
3. Simon 1996, Ch. 1.

Recommended:

1. Wilks 1990 has a confusing, but useful, overview of the many meanings of “theory” and “model”.
2. Daubert v. Merrell 1993 has interesting observations on the nature of scientific theories and expertise.
3. Green 2004, Ch. 3, has a useful survey of different views of scientific explanation and scientific models, embedded in a discussion of connectionism.

Topic: What is software?

Required: Moor 1978

Strongly Recommended: Suber 1988

Recommended: Colburn 1999

Topic: Can programs be copyrighted or patented?

Required: Newell 1985-1986

Recommended: Samuelson 1990, Koepsell & Rapaport 1995

Topic: Can programs be verified?

Required: Smith 1985

Strongly recommended: Fetzer 1988

Recommended background for Fetzer 1988: De Millo et al. 1979, Ardis et al. 1989

Topic: Philosophy of AI Turing 1950, Searle 1980.

Topic: Computer Ethics

Required:

1. Moor 1979

Highly recommended: Johnson 2002

Recommended: Friedman & Kahn 1992

2. LaChat 1986

Highly recommended: Lem 1971

Recommended: Moor 1985

Topic: Philosophy of Computer Science: Scheutz 2002, Smith 2002

Appendix B: Position Paper 1—What is computer science?

The purpose of this position paper is to give you an opportunity to clarify *your* beliefs about what computer science is, so that as we continue to discuss the topic in class and as you continue to read about it, you'll know where *you* stand—what *your* beliefs are. Later, when your beliefs have been “contaminated” by further readings and by our discussions, you may wish to *revise* your beliefs. But you can't revise a belief that you don't have (you can only acquire new beliefs). So, here I am forcing you to *discover, clarify, and defend* the beliefs that you now have, by turning them into words and putting them on paper.

Imagine that you are the newly-appointed Dean of the School of Science at the University of Aix (pronounced like the letter ‘X’). In an attempt to build up the rival School of Engineering, the newly-appointed Dean of Engineering has proposed to the Provost (the Deans' boss) that the Department of Computer Science be moved—lock, stock, and computer,²¹ so to speak—to Engineering, on the following grounds:

1. Science
is the systematic observation, description, experimental investigation, and theoretical explanation of natural phenomena.
2. Computer science is the study of computers and related phenomena.
3. Therefore, computer science is not a science.

(The Dean of Engineering has not yet argued that computer science is an engineering discipline; that may come later.)

You may agree with this argument; then again, you may not agree with it. You should ignore political considerations: You may suppose that the move from Science to Engineering involves no loss of money, prestige, or anything else, and it is to be done, if at all, only on strictly intellectual grounds. How might you respond to the Dean of Engineering's argument? The Provost is eagerly awaiting your reply, and will abide by your decision . . . if, that is, you give a well-argued defense of your position.

²¹[<http://www.quinion.com/words/qa/qa-loc1.htm>]

There are several possible responses that you might have:

Response 1: You might **disagree** for any of 3 reasons:

- (a) You believe that premise (1) is false.
- (b) You believe that premise (2) is false.
- (c) You believe that (1) and (2) are true but that conclusion (3) does not follow from them. (E.g., you might believe that “computers and related phenomena” *are* “natural phenomena”.)

... or you might believe some combination of these.

Please explain to the Provost *why* you disagree, by explaining which of (a), (b), and/or (c) you *do* believe.

Response 2: You might **agree** for any of the following reasons:

(a) You believe (1) and (2) and that (3) follows from them. If so, please explain to the Provost why you believe (1) and (2), and how (3) follows from them. (E.g., you might believe that “computers and related phenomena” are *not* “natural phenomena”, or you might believe that computer science doesn’t study them “systematically”, or ...)

(b) You don’t believe (1) or you don’t believe (2), but you *do* believe (3), whether or not it follows from (1) and (2). If so, please explain why you don’t believe (1) and/or (2) and what *other* reasons you have for believing (3).

Response 3: You might **neither** agree **nor** disagree with (3); alternatively, you might **both** agree **and** disagree with it. For example, you might believe that computer science is *both* a science *and* an engineering discipline (or, alternatively, that it is neither). If so, then please give your reasons for this.

Other responses: You might not agree with any of these responses. However, I believe that any other response can, perhaps with a bit of force, be seen to fall under one of the above Responses. But if you really feel that your position is not exactly characterized by any of the above Responses, then please say what your position is, why you believe it, and why you think it is not one of the above.

Ground Rules:

1. If you resort to a dictionary, textbook, article, website, etc., be sure to say which one.
2. Your answer should honestly reflect *your* beliefs (not what you think the fictional Provost or Dean of Engineering want to hear!).
3. Your position paper should be approximately **1 typed page and double-spaced (i.e., about 250 words)**.
4. Please bring **5 copies** to class on the due date.
5. This paper only needs the title “Position Paper #1”, your name, and the date at the top of the page.
6. For general assistance with writing (including my preferred method of paper preparation and format, as well as advice on grammar), see my website “How to Write”.²²

DUE AT THE BEGINNING OF LECTURE, 1 WEEK FROM TODAY

²²[<http://www.cse.buffalo.edu/~rapaport/howtowrite.html>]

Appendix C: Suggestions and Guidelines for Peer-Group Editing of Position Paper #1

1. When you get into your small groups, introduce yourselves, and share copies of your papers with each other.
2. Choose one paper to discuss first. (Suggestion: Go in alphabetical order by family name.)
3. The other people in the group might find it useful to imagine themselves as members of a committee set up by the Provost to make a recommendation. Their purpose is to try to help the author clarify his or her beliefs and arguments, so that they will be able to make a recommendation to the Provost on purely logical grounds (again: ignore politics!).
4. Start by asking the author to state (or read) his or her beliefs about whether computer science is a science, giving his or her reasons for those beliefs.
5. Any time you have a question, ask it. Here are some suggestions:
 - Why did you say _____ rather than _____?
 - What did you mean when you said _____?
 - Can you give me an example of _____?
 - Can you give me more details about _____?
 - Do you think that _____ is always true?
 - Why? (This is always a good question to ask.)
 - How?
6. The author should not get defensive. The committee members are friendly. Critical, but friendly.
7. Keep a written record of the questions and replies. This will be useful to the author, for revision.
8. After spending at least 10 minutes on the first paper, move on to the next, going back to step (2) above, changing roles. Spend about 15 minutes per paper.
9. At home, over the next week, please revise your paper to take into consideration the comments made by your fellow students (i.e., your “peers”): Perhaps defend your claims better, or clarify statements that were misunderstood, etc. For help, see the TA or me.

1-2 PAGE (250-500 WORD) REVISION, 1 COPY, TYPED, IS DUE AT THE BEGINNING OF LECTURE, ONE WEEK FROM TODAY. NO LATE PAPERS WILL BE ACCEPTED!

Appendix D: Website

The website for the course, at [<http://www.cse.buffalo.edu/~rapaport/philcs.html>], has links to the syllabus and a directory of documents that, in turn, has a large bibliography, links to other relevant websites, and links to the assignments, position papers, term-paper topics, and final exam.

References

- Alexander, Peter (1967), "Mach, Ernst", in Paul Edwards (ed.), *Encyclopedia of Philosophy* 5: 115–119.
- Ardis, Mark; Basili, Victor; Gerhart, Susan; Good, Donald; Gries, David; Kemmerer, Richard; Leveson, Nancy; Musser, David; Neumann, Peter; & von Henke, Friedrich (1989), "Editorial Process Verification" (letter to the editor, with replies by James H. Fetzer and Peter J. Denning), *Communications of the ACM* 32(3) (March): 287–290.
- Audi, Robert (2001), "Philosophy: A Brief Guide for Undergraduates"
[<http://www.apa.udel.edu/apa/publications/texts/briefgd.html>].
- Boehm, C., & Jacopini, G. (1966), "Flow Diagrams, Turing Machines, and Languages with only Two Formation Rules", *Communications of the ACM* 9(5): 366–371.
- Brooks, Frederick P., Jr. (1996), "The Computer Scientist as Toolsmith II", *Communications of the ACM* 39(3) (March): 61–68.
- Chalmers, David J. (1993a), "A Computational Foundation for the Study of Cognition"
[<http://www.u.arizona.edu/~chalmers/papers/computation.html>].
- Chalmers, David J. (1993b), "Does a Rock Implement Every Finite-State Automaton?", *Synthese* 108 (1996): 309–333.
- Chalmers, David J. (1994), "On Implementing a Computation", *Minds and Machines* 4 (1994): 391–402.
- Cho, Kwangsu; & Schunn, Christian D. (2004), "You Write Better When You Get Feedback from Multiple Peers than an Expert", *Proceedings of the 20th Annual Conference of the Cognitive Science Society* (Mahwah, NJ: Lawrence Erlbaum Associates, 2005).
- Cleland, Carol E. (1993), "Is the Church-Turing Thesis True?", *Minds and Machines* 3(3) (August): 283–312.
- Cleland, Carol E. (1995), "Effective Procedures and Computable Functions", *Minds and Machines* 5(1): 9–23.
- Cleland, Carol E. (2001), "Recipes, Algorithms, and Programs", *Minds and Machines* 11(2) (May): 219–237.
- Cleland, Carol E. (2002), "On Effective Procedures", *Minds and Machines* 12(2) (May): 159–179.
- Colburn, Timothy R. (1999), "Software, Abstraction, and Ontology", *The Monist* 82(1): 3–19; reprinted (in slightly different form) in Colburn 2000, Ch. 12.
- Colburn, Timothy R. (2000), *Philosophy and Computer Science* (Armonk, NY: M.E. Sharpe).
- Copeland, B. Jack (2002), "Hypercomputation", *Minds and Machines* 12(4): 461–502.
- Copeland, B. Jack (2004), "Computation", in Floridi 2004.
- Crowcroft, Jon (2005), "On the Nature of Computing", *Communications of the ACM* 48(2) (February): 19–20.
- Daubert v. Merrell Dow Pharmaceuticals (92–102), 509 U.S. 579 (1993)
[<http://supct.law.cornell.edu:8080/supct/html/92-102.ZS.html>].
- Davis, Martin (ed.) (1965), *The Undecidable: Basic Papers on Undecidable Propositions, Unsolvable Problems and Computable Functions* (New York: Raven Press).
- Davis, Michael (1998), *Thinking Like an Engineer: Studies in the Ethics of a Profession*

- (New York: Oxford University Press).
- De Millo, Richard A.; Lipton, Richard J.; & Perlis, Alan J. (1979), “Social Processes and Proofs of Theorems and Programs”, *Communications of the ACM* 22(5): 271–280.
- Denning, Peter J. (1985), “What Is Computer Science?”, *American Scientist* 73 (January–February): 16–19.
- Denning, Peter J.; Comer, Douglas E.; Gries, David; Mulder, Michael C.; Tucker, Allen; Turner, A. Joe; & Young, Paul R. (1989), “Computing as a Discipline”, *Communications of the ACM* 32(1) (January): 9–23.
- Ensmenger, Nathan (2004), “Bits of History: Review of A.R. Burks’s *Who Invented the Computer? The Legal Battle that Changed Computing History*”, in *American Scientist* 91 (September–October): 467–468.
- Fetzer, James H. (1988), “Program Verification: The Very Idea”, *Communications of the ACM* 31(9) (September): 1048–1063; reprinted in Timothy R. Colburn, James H. Fetzer, & Terry L. Rankin (eds.), *Program Verification: Fundamental Issues in Computer Science* (Dordrecht, Holland: Kluwer Academic Publishers, 1993): 321–358.
- Floridi, Luciano (1999), *Philosophy and Computing: An Introduction* (London: Routledge).
- Floridi, Luciano (2004), *The Blackwell Guide to the Philosophy of Computing and Information* (Malden, MA: Blackwell).
- Friedman, Batya; & Kahn, Peter H., Jr. (1992), “People Are Responsible, Computers Are Not”, excerpt from their “Human Agency and Responsible Computing: Implications for Computer System Design”, *Journal of Systems Software* (1992): 7–14; excerpt reprinted in M. David Ermann, Mary B. Williams, & Michele S. Shauf (eds.) *Computers, Ethics, and Society, Second Edition* (New York: Oxford University Press, 1997): 303–314.
- Green, Christopher D. (2004), “(How) Do Connectionist Networks Model Cognition?” (unpublished Ph.D. dissertation, Department of Philosophy, University of Toronto).
- Gries, David (1981), *The Science of Programming* (New York: Springer-Verlag).
- Hartmanis, Juris, & Lin, Herbert (1992), “What Is Computer Science and Engineering?”, in Juris Hartmanis & Herbert Lin (eds.), *Computing the Future: A Broader Agenda for Computer Science and Engineering* (Washington, DC: National Academy Press), Ch. 6, pp. 163–216.
- Haugeland, John (1981), “Semantic Engines: An Introduction to Mind Design”, in John Haugeland (ed.), *Mind Design: Philosophy, Psychology, Artificial Intelligence* (Cambridge, MA: MIT Press): 95–128.
- Hayes, Patrick J. (1997), “What Is a Computer? An Electronic Discussion”, *The Monist* 80(3).
- Hempel, Carl G. (1966), *Philosophy of Natural Science* (Englewood Cliffs, NJ: Prentice-Hall).
- Henkin, Leon (1962), “Are Logic and Mathematics Identical?”, *Science* 138(3542) (November 16): 788–794.
- Herman, Gabor T. (1983), “Algorithms, Theory of”, in Anthony S. Ralston (ed.), *Encyclopedia of Computer Science and Engineering, 2nd edition* (New York: Van Nostrand Reinhold): 57–59.

- Johnson, George (2002), "To Err Is Human", *New York Times* (14 July).
- Johnson-Laird, Philip N. (1988), *The Computer and the Mind: An Introduction to Cognitive Science* (Cambridge, MA: Harvard University Press).
- Kemeny, John G. (1959), *A Philosopher Looks at Science* (Princeton: D. van Nostrand).
- Koepsell, David R. (2000), *The Ontology of Cyberspace: Philosophy, Law, and the Future of Intellectual Property* (Chicago: Open Court).
- Koepsell, David R.; & Rapaport, William J. (1995), "The Ontology of Cyberspace: Questions and Comments", *Technical Report 95-25* (Buffalo: SUNY Buffalo Department of Computer Science) and *Technical Report 95-09* (Buffalo: SUNY Buffalo Center for Cognitive Science).
- Knuth, Donald (1974), "Computer Science and Its Relation to Mathematics", *American Mathematical Monthly* 81(4) (April): 323–343.
- Kugel, Peter (2002), "Computing Machines Can't Be Intelligent (... and Turing Said So)", *Minds and Machines* 12(4): 563–579.
- Kyburg, Henry E., Jr. (1968), *Philosophy of Science: A Formal Approach* (New York: Macmillan).
- LaChat, Michael R. (1986), "Artificial Intelligence and Ethics: An Exercise in the Moral Imagination", *AI Magazine* 7(2): 70–79.
- Lem, Stanislaw (1971), "Non Serviam", in S. Lem, *A Perfect Vacuum*, trans. by Michael Kandel (New York: Harcourt Brace Jovanovich, 1979).
- Lloyd, Seth, & Ng, Y. Jack (2004), "Black Hole Computers", *Scientific American* 291(5) (November): 52–61.
- Loui, Michael C. (1987), "Computer Science Is an Engineering Discipline", *Engineering Education* 78(3): 175–178.
- Loui, Michael C. (1995), "Computer Science Is a New Engineering Discipline", *ACM Computing Surveys* 27(1) (March): 31–32.
- Moor, James H. (1978), "Three Myths of Computer Science", *British Journal for the Philosophy of Science* 29(3) (September): 213–222.
- Moor, James H. (1979), "Are There Decisions Computers Should Never Make?", *Nature and System* 1: 217–229.
- Moor, James H. (1985), "What Is Computer Ethics?", *Metaphilosophy* 16(4) (October): 266–275.
- Moulton, Janice, & Voytek, Jane (1979), "Can Humans Think Machines Think?", *Teaching Philosophy* 3(2): 153–167.
- Newell, Allen (1985-1986), "Response: The Models Are Broken, the Models Are Broken", *University of Pittsburgh Law Review* 47: 1023–1031.
- Newell, Allen; Perlis, Alan J.; & Simon, Herbert A. (1967), "Computer Science", *Science* 157(3795) (22 September): 1373–1374.
- Newell, Allen, & Simon, Herbert A. (1976), "Computer Science as Empirical Inquiry: Symbols and Search", *Communications of the ACM* 19(3) (March): 113–126.
- O'Connor, J.J., & Robertson, E.F. (1998), "Charles Babbage"
[<http://www-gap.dcs.st-and.ac.uk/~history/Mathematicians/Babbage.html>].
- Papineau, David (1996), "Philosophy of Science", in Nicholas Bunnin & E.P. Tsui-James (eds.), *The Blackwell Companion to Philosophy* (Oxford: Blackwell): 290–324.
- Perry, William G., Jr. (1970), *Forms of Intellectual and Ethical Development in the College*

- Years: A Scheme* (New York: Holt, Rinehart, and Winston).
- Perry, William G., Jr. (1981), "Cognitive and Ethical Growth: The Making of Meaning", in Arthur W. Chickering and Associates, *The Modern American College* (San Francisco: Jossey-Bass): 76–116.
- Perruchet, Pierre, & Vinter, Annie (2002), "The Self-Organizing Consciousness", *Behavioral and Brain Sciences* 25(3) (June): 297–388.
- Petroski, Henry (2003), "Early Education", *American Scientist* 91 (May-June): 206–209.
- Popper, Karl R. (1962), *Conjectures and Refutations: The Growth of Scientific Knowledge* (New York: Harper & Row).
- Preston, Beth (2000), "Recipes and Songs: Towards a Theory of Production" (unpublished ms.)
- Pylyshyn, Zenon W. (1984), *Computation and Cognition: Toward a Foundation for Cognitive Science* (Cambridge, MA: MIT Press).
- Rapaport, William J. (1982), "Unsolvable Problems and Philosophical Progress", *American Philosophical Quarterly* 19: 289–298.
- Rapaport, William J. (1986), "Philosophy of Artificial Intelligence: A Course Outline", *Teaching Philosophy* 9: 103–120.
- Rapaport, William J. (1998), "How Minds Can Be Computational Systems", *Journal of Experimental and Theoretical Artificial Intelligence* 10: 403–419.
- Rapaport, William J. (1999), "Implementation Is Semantic Interpretation", *The Monist* 82(1): 109-130.
- Rapaport, William J. (2004), "How to Study"
[<http://www.cse.buffalo.edu/~rapaport/howtostudy.html>].
- Rapaport, William J. (forthcoming), Review of Shieber 2004, *Computational Linguistics*.
- Rogers, Hartley, Jr. (1959), "The Present Theory of Turing Machine Computability", *Journal of the Society for Industrial and Applied Mathematics* 7: 114–130; reprinted in Jaakko Hintikka (ed.), *The Philosophy of Mathematics* (London: Oxford University Press, 1969): 130–146.
- Rosenberg, Alex (2000), *Philosophy of Science: A Contemporary Introduction* (London: Routledge).
- Salmon, Wesley C. (1984), *Scientific Explanation and the Causal Structure of the World* (Princeton: Princeton University Press).
- Samuelson, Pamela (1990), "Should Program Algorithms Be Patented?", *Communications of the ACM* 33(8) (August): 23–27.
- Scheutz, Matthias (2002), "Philosophical Issues about Computation", *Encyclopedia of Cognitive Science* (London: Macmillan).
- Searle, John R. (1980), "Minds, Brains, and Programs", *Behavioral and Brain Sciences* 3: 417–457.
- Searle, John R. (1990), "Is the Brain a Digital Computer?", *Proceedings and Addresses of the American Philosophical Association* 64: 21–37.
- Shapiro, Stuart C. (2001), "Computer Science: The Study of Procedures"
[<http://www.cse.buffalo.edu/~shapiro/Papers/whatiscs.ps>].
- Shieber, Stuart M. (2004), *The Turing Test: Verbal Behavior as the Hallmark of Intelligence* (Cambridge, MA: MIT Press).
- Simon, Herbert A. (1996), *The Sciences of the Artificial, Third Edition* (Cambridge, MA:

- MIT Press).
- Simon, Herbert A., & Newell, Allen (1958), "Heuristic Problem Solving: The Next Advance in Operations Research", *Operations Research* 6(1) (January–February): 1–10.
- Smith, Brian Cantwell (1985), "Limits of Correctness in Computers", *Technical Report CSLI-85-36* (Stanford, CA: Center for the Study of Language and Information); first published in Charles Dunlop & Rob Kling (eds.), *Computerization and Controversy* (San Diego: Academic Press, 1991): 632–646; reprinted in Timothy R. Colburn, James H. Fetzer, & Terry L. Rankin (eds.), *Program Verification: Fundamental Issues in Computer Science* (Dordrecht, Holland: Kluwer Academic Publishers, 1993): 275–293.
- Smith, Brian Cantwell (2002), "The Foundations of Computing", in Scheutz, Matthias (ed.), *Computationalism: New Directions* (Cambridge, MA: MIT Press): 23–58.
- Soare, Robert I. (1996), "Computability and Recursion", *Bulletin of Symbolic Logic* 2(3) (September): 284–321.
- Suber, Peter (1988), "What Is Software?", *Journal of Speculative Philosophy* 2(2): 89–119.
- Suber, Peter (1997), "Turing Machines"
[<http://www.earlham.edu/~peters/courses/logsys/turing.htm>].
- Thomason, Richmond H. (2003), "Dynamic Contextual Intensional Logic: Logical Foundations and an Application", in P. Blackburn et al. (eds.), *CONTEXT 2003, Lecture Notes in Artificial Intelligence 2680* (Berlin: Springer-Verlag): 328–341.
- Turing, Alan M. (1936), "On Computable Numbers, with an Application to the Entscheidungsproblem", *Proceedings of the London Mathematical Society*, Ser. 2, Vol. 42: 230–265.
- Turing, Alan M. (1950), "Computing Machinery and Intelligence", *Mind* 59: 433–460; reprinted in Shieber 2004.
- Weizenbaum, Joseph (1976), *Computer Power and Human Reason: From Judgment to Calculation* (New York: W.H. Freeman).
- Wilks, Yorick (1990), "One Small Head: Models and Theories", in Partridge, Derek; & Wilks, Yorick (eds.), *The Foundations of Artificial Intelligence: A Sourcebook* (Cambridge, UK: Cambridge University Press): 121–134.
- Woodhouse, Mark B. (2003), *A Preface to Philosophy, 7th edition* (Belmont, CA: Wadsworth/Thomson Learning).
- Ziman, John M. (1968), "What Is Science?", in Michalos, Alex C. (ed.) (1974), *Philosophical Problems of Science and Technology* (Boston: Allyn & Bacon): 1–27.