 Last edited Monday, January 8, 2001

Syllabus

I. Administrative Details

1. *Course:*	Computer Science B607 [1392] • Philosophy of Computation (3 hours)		
Time:	Tuesdays and Thursdays 9:30 – 11:00 a.m.		
Place:	Lindley Hall • 101		
2. *Instructor:*	Brian Cantwell Smith	*Office hours:*	Tuesdays 2:00 – 4:00 p.m.
Office:	Sycamore • 126 / Lindley • 228		(Sycamore 126 office)
Email:	smithbc@indiana.edu	*Phone:*	812·855·3788
3. *Web site:*	**http://www.ageofsig.org/courses/b607**		

II. Description

Can a computer be conscious? What would a continuous programming language be like? What kind of "materiality" exists on the Web? What does 'virtual' mean? How will computation affect the future of science? Will quantum or DNA computers change our conception of what is and what s not computable? Is it ethical to give computers power over human life?

Addressing such questions requires knowing what computation is, and what computers are—to a depth, it will be argued in this course, beyond the reach of current theories. The course undertakes a critical examination of the conceptual foundations of computing—in an attempt to figure out what we know, what we don't know, and what a more adequate theory would look like. Overall, the motivation is to understand:

1. The models and metaphors in terms of which we understand computing, from programs to processes, architecture to abstraction, parameterization to parallelism; and
2. The use of computational concepts in adjacent fields, from cognitive science to physics, economics to art.

Specific attention will be given to six traditional views of computation: formal symbol manipulation, recursive function theory, effective computability & computational complexity, digital state machines, information processing, and Newell and Simon's notion of a physical symbol system. Some non-standard views will also be (briefly) considered, including connectionism, non-linear dynamics, and artificial life. Each view will be judged by, among other things, its ability to do justice to practice.

In conclusion, some consideration will be given to the wider role of computational concepts in intellectual life—including their affect on our human self-conception.

III. Content

1. *Reading:* There is no explicit text. Primary reading will consist of:
   a. Lecture notes, available (in PDF, for reading and printing) on the web site.
   b. Background papers by (among others): Chalmers, Dretske, Dreyfus, Fodor, Goodman, Haugeland, Hayes, Kleene, Minsky, Newell, Penrose, Putnam, Searle, Shannon, Simon, Turing, Webb, and others. Copies of these papers will be available (i) in a class reader, and (ii) for copying, in the computer science department copier room (Lindley Hall 210).
   c. Selections from the instructor's forthcoming series of books on the philosophy of computing: *The Age of Significance: An Essay in the Foundations of Computation and Intentionality* (also available on the course web site).

2. *Prerequisites:* Students should have substantial computational expertise, typically from a combination of programming and instruction, plus at least one (ideally both) of the following:
   a. Prior course work in logic or the theory of computation (Turing machines, computability theory, complexity theory, etc.); or
   b. Prior course work in philosophy.

   Students in philosophy or history & philosophy of science with a background in logic may get by without substantial prior computational expertise, but are liable to want to work on problem sets (see below) in conjunction with students with more computational experience. (For questions see the instructor.)

3. *Problem sets:* Course work will consist of 14 "problem sets," one per week (the first due in the second week). Each problem set will consist of a structured, essay question, often with multiple parts.

   The notion of a "right" or "wrong" answer is considered inappropriate. Grades will be based on the clarity, imagination, and depth of answer. Experience suggests that it will take a typical student a minimum of 4–5 hours to complete each problem set to a grade-A level.

   Note: See the discussion of collaboration, below.

4. *Format:* Problem sets will be distributed in class, and made available online. Responses should be submitted to the class web page, using the **Annotate** program (developed by Ruth Eberle of the IU Cognitive Science Program). Responses will be graded, commented, cross-referenced, and made available to other class members to read and reference.

   To submit responses, connect to **http://www.indiana.edu/~annotate** using (a recent version of) the Netscape[1] browser.

   Note: Although responses will be made available to other class members, they remain password protected, unavailable to anyone outside the class.

[1] At the moment, Annotate does not work with Microsoft Explorer (a new, Java-based version is being developed, which will do so, but it is unlikely to be available for use this semester).

5. *Collaboration:* Students are encouraged to work on problem sets in groups (of two or more), debating and discussing questions and responses with other class members. This will be especially useful for students lacking a thorough background in both computer science and philosophy, but is recommended for everyone.

   If worked on in groups, responses may be submitted: (i) individually, by each group member; or (ii) if it is acceptable to all parties, as a single document written by the group as a whole.[2] In the former case, each person should write the text of their own reply, listing the other group members at the beginning, and giving acknowledgment, on a point-by-point basis (in the way that would be appropriate in a professional paper), of the insights, contributions, and/or suggestions of others. In the latter case, the response should be posted under the name of one group member, begin with a list of all group members, and indicate, throughout the response, contributions of each member, as appropriate.[3]

6. *Sections:* No formal sections.

7. *Grading:* No exams.
   98% — for problem sets (7% each)
   15% — for class participation
   (If anyone achieves more than 100% they will definitely get an A+)

IV. Schedule (provisional)

1. The following schedule, specifying 15 weeks of lectures, is provisional. If we cover material faster than expected, we will devote remaining time at the end of the semester to questions of special interest to class members.
2. The secondary readings listed below are intended to open out into a deeper analysis of each topic; they are illustrative readings, rather than comprehensive bibliographies. How many of these (or other) secondary papers are read is left to the student's discretion; you should probe deeply enough to feel confident about your problem set replies. (The listed papers are also a bit out of date; I may update them as the course proceeds.)

   A partially-annotated web-based bibliography on the philosophy of computing is under development; as soon as possible a pointer to it, with revised readings and links, will be posted on the class web page.
3. Only authors and titles are given below; full references (together with short annotations/commentaries) can be found in the class Reading List—distributed on the first day of class, and available on the class web site.

A. **Part I — Introduction** (2 weeks)

1. Summary
   a. Smith, Brian Cantwell, "The Foundations of Computing"
2. Primary

---

[2] Or by a subgroup of the whole group, if appropriate.

[3] Obviously, exact credit will be impossible to assign. What is expected are comments such as: "The basic idea was Randy's; but details were developed in discussion with Pat, Hilary, and Llewelyn"; or "Overall, this part of the response was generated collaboratively, but it was Jedediah's idea to frame it in topological terms." Or whatever.

   a. AOS·I (Introduction) — Chapter 1: “Project”
   b. AOS·I (Introduction) — Chapter 2: “Construals”
   c. AOS·I (Introduction) — Chapter 3: “State of the Art”
3. Secondary:
   a. Haugeland, John, “Semantic Engines”

B. Part II — Formal Symbol Manipulation (4 weeks)

1. Primary: AOS·II (Formal Symbol Manipulation) Chapters 1–4
2. Secondary
   a. Background
      i. Hunter, Geoffrey, Part I §s 1–7 of *Metalogic: An Introduction to the Metatheory of Standard First Order Logic*.
   b. Computation as formal symbol manipulation
      i. Hayes, Patrick J., “Computation and Deduction”
      ii. Kowalski, Robert, “Algorithm = Logic + Control”
      iii. Newell, Alan and Simon, Herbert A., “Computer Science as Empirical Inquiry”
      iv. Newell, Alan, “Physical Symbol Systems”
   c. Analysis, discussion, and critique
      i. Fodor, Jerry A., “Methodological Solipsism Considered as a Research Strategy in Cognitive Psychology”
      ii. Fodor, Jerry A, and Pylyshyn, Zenon, “Connectionism and Cognitive Architecture: A Critical Analysis”
      iii. Dretske, Fred I., “Machines and the Mental”
      iv. Searle, John “Minds, Brains, and Programs”
      v. Haugeland, John, “Syntax, Semantics, Physics”

C. Part III — Effective Computability and Recursion Theory (5 weeks)

1. Primary
   a. AOS·III (Effective Computability) Chapters 1–4
   b. AOS·I (Introduction) — Chapter 6: “Synopsis • I”
2. Secondary
   a. For Turing machines themselves
      i. Minsky, Marvin, Chapters 5–8 of *Finite & Infinite Machines*
      ii. Turing, Alan M., “On Computable Numbers, with an application to the Entscheidungsproblem”
      iii. Turing, Alan M., “Computing machinery and intelligence”
      iv. Kleene, Stephen C., “Turing's Analysis of Computability, and Major Appl’ns of It”
   b. For discussion (historical)
      i. Gandy, Robin, “The Confluence of Ideas in 1936”
      ii. Davis, Martin, “Mathematical Logic & the Origin of the Modern Computer”
      iii. Webb, Judson, Introduction & Chapter 1 of *Mechanism, Mentalism, and Metamathematics: An Essay on Finitism*
      iv. Gandy, Robin, “Church’s Thesis and Principles for Mechanisms”
   c. Analysis, discussion, and critique (recent)
      i. Searle, John, chapter 9 from *The Rediscovery of the Mind*; or “Is the Brain a Digital

Computer?"
- ii. Putnam, Hilary, appendix to *Representation and Reality*
- iii. Chalmers, Dave, "Does a Rock Implement Every Finite-State Automaton?"
- iv. Copeland, B. J. (1996) "What is Computation?"
- v. Scheutz, Matthias, chapters 2–4 from his dissertation "The Missing Link: Implementation and Realization of Computations in Computer and Cognitive Science"

D. Part IV — Digital State Machines (3 weeks)

1. Primary: AOS·I (Introduction) — Chapter 7: "Synopsis • II"
2. Secondary
   - a. For the notion of a digital state machine
     - i. Minsky, Marvin, Chapters 1 & 2 of *Finite & Infinite Machines*
   - b. For the notion of digitality
     - i. Haugeland, John, Chapter 2 of *Artificial Intelligence: The Very Idea*
     - ii. Goodman, Nelson, Chapter 4 of *Languages of Art*
     - iii. Lewis, David, "Analog and Digital"
     - iv. Haugeland, John, "Analog and Analog"
     - v. Dretske, Fred, "Sensation & Perception" (*Knowledge & the Flow of Information*, ch. 6)
     - vi. Fodor, Jerry A. & Ned J. Block, "Cognitivism & the Analog/Digital Distinction"

E. Part V — Information Processing (to be skipped)

1. Primary: AOS·I (Introduction) — Chapter 7: "Synopsis • II"
2. Secondary
   - a. For the syntactic notion
     - i. Weaver, Warren, "Recent Contributions to the Math. Theory of Communication"
     - ii. Shannon, Claude E., Part I of "The Mathematical Theory of Communication"
     - iii. Singh, Jagjit, Ch. 1–9 of *Great Ideas in Information Theory, Language, and Cybernetics*
   - b. For the semantic notion
     - i. Dretske, Fred I., "Précis of *Knowledge and the Flow of Information*"
     - ii. Dretske, Fred I. Chapter 3 of *Knowledge and the Flow of Information*
     - iii. Israel, David and John Perry, "What is Information?"
   - c. For application of the semantic notion to AI and computer science (respectively)
     - i. Rosenschein, Stanley J., "Formal theories of Knowledge in AI and Robotics"
     - ii. Halpern, Joseph, "Using Reasoning about Knowledge to Analyze Distributed Systems"

F. Part VI — Conclusion (The Age of Significance) (1 week)

1. Primary
   - a. [AOS·I (Introduction) — Chapter 6: "Synopsis • I"]
   - b. [AOS·I (Introduction) — Chapter 7: "Synopsis • II"]
   - c. AOS·I (Introduction) — Chapter 8: "The Middle Distance"

——*end of file* ——