

Philosophy 395 Computer Science 228 Symbolic Systems 210	<h1>Syllabus</h1>	Philosophy of AI Stanford University Winter Quarter, 1989–90
File path: Koyaanisqatsi : Philosophy of AI : Phai — Syllabus Copyright: © 1990 Brian Cantwell Smith Last edited Monday, January 22, 1990 — 1:30 p.m.		

A. Administrative details

1. Class: Philosophy of AI (Philosophy 395, Computer Science 228, Symbolic Systems 210)
Time: Mondays 3:15 – 5:00 p.m.
Place: Cordura Main Seminar Room (1st floor)
2. Instructor: Brian C. Smith
Office: Xerox PARC, 3333 Coyote Hill Road, Palo Alto, CA 94304
Net mail: briansmith.pa@xerox.com
Phone: 415-494-4336
3. Office hours: Monday 5:00 – 7:00 p.m., Cordura 202 (phone: 723-0488)

B. Lecture schedule

A. *Lay of the Land*

1. Introduction (January 22, 1990)
 - Delineate the general project of artificial intelligence; discriminate various general approaches (such as strong & weak AI), clarify primary methodological point: that computation is as empirical an inquiry as psychology or cognitive science.
2. The intentional challenge (January 29, 1990)
 - Identify various general characteristics of intentionality (a general term assumed to subsume issues of symbols, language, meaning, semantics, representation, etc.), especially: *meaning* or *registration* (the fact that any intentional structure presents the world “as being a certain way”), and *reference* or *semantic reach* (the fact that the intentional elements of the world bear a non-effective relationship to whatever they refer to or are about).

B. *Philosophy of Computation*

3. Computation in the wild (February 5, 1990)
 - Assess computational practice; characterise various terms of art (‘program’, ‘process’, ‘interpreter’, ‘compiler’, ‘architecture’, ‘implementation’, etc.); highlight issues relevant to an adequate semantics of a computational system.
4. Digital state machines (February 12, 1990)
 - Examine the notion of a “digital” or discrete system; contrast it with a range of “analog”, continuous, and indefinite alternatives; consider the proposal that digitality is an essential — perhaps even defining — characteristic of what it is to be a computer.
5. Formal symbol manipulation (Tuesday, February 20, 7:00 – 9:00 p.m. ← **note date change**)
 - Scrutinise the widely-accepted but rarely-examined hypothesis that computation is “formal symbol manipulation”, with a particular eye on what the claim means, why it might be important, and whether it is true.
6. Effective computability (February 26, 1990)
 - Exhume the traditional notion of effective computability; distinguish Turing machines (qua devices) from the familiar theory of them; and suggest that, though relevant to

intentionality, and typically analysed under interpretation, the notion of effective computability is in fact independent of intentional considerations.

C. Foundational Issues in AI

7. The mind-body problem for machines (March 5, 1990)
 - Consider the hypothesis that computers “solve” the age-old wonder about the relationship of mind to body; provoke an analysis of the same question in the machine case; conclude that, current proposals notwithstanding, an adequate resolution remains elusive.
8. Connectionism and the rise of concepts (March 12, 1990)
 - Outline the intellectual issues triggered by the prospect of neurally-inspired connectionist networks; distinguish two broad classes of AI theorist: those that accept a conceptual scheme as a methodological given, and those that take the rise of concepts to be an essential phenomenon to be explained.
9. An Alternative Account (Optional session: March 12, 1990, 7:00 – 9:00 p.m. **← note this is the evening of the date of lecture #8**).
 - Sketch the instructor’s own theory of what computation is, how it addresses the various issues raised in the course, and what it entails for a theory of mind.

— end of file — □□