

Philosophy 395 Computer Science 228 Symbolic Systems 210	Lecture 7 — Connectionism and the Rise of Concepts	Philosophy of AI Stanford University Winter Quarter, 1989–90
File path:	Koyaanisqatsi : Philosophy of AI : Phai — 7 — Connectionism	
Copyright:	© 1990 Brian Cantwell Smith	
Last edited	Monday, March 5, 1990 — 1:33:31 PM	

a. Preamble

- a. Today start on 3rd part of course: Foundational issues in AI

a. Outline

- Concepts
 - What they are: big, unsolved issue in cog sci
 - Logic, model-theory, etc.
 - work: on inter-conceptual relations (proofs, analogies, etc.)
 - definitions: Fodor & Katz, etc.; \Rightarrow no definitions
 - conclusion: don't take them on
 - font-change semantics
 - Tarski's conventionT \Leftarrow certain humility
 - Fodor: this is all that allows ST to be possible
 - Related to language
 - Properties of public discourse
 - Properties of language
 - Properties of scientific theory itself
 - Counter tradition, though, that this isn't everything
 - Two simple examples:
 - Evans: sound coming from over your shoulder
 - Perry: shopping basket: essential indexicality
 - AI
 - need to connect with action, perception, etc.
 - brittleness, etc. (cf. Winograd)
 - Against this background, connectionism has emerged
 - Something to do with sub-, non-, intra-, ... conceptual analysis and/or modeling
- Connectionism
 - Review: fixed stock of units, measurable signals, quasi-linear, etc.
 - Can show various kinds of behaviour
 - Lots of traditional questions:
 - Parse of input: whose?

- Variables, etc.: non-local dependencies
- Quantifiers, disjunction, etc.
- Relation to neurology
- Most important question: relation to concepts
 - Specifically: are concepts emergent, or implemented?
 - What are the semantic aspects (if any) of connectionist level
 - Relation to perspectival, action, etc.

a. Notes to be included

- Concepts
 - whose (theorist's or agent's)
 - what does it take to have them
 - relation to representation
 - featural definition (Fodor & Katz)
 - intra-, rather than inter-, conceptual structure
- Implementation
 - what properties cross boundaries
 - cf. "complete, formal, precise"
 - standard line: implementation is explanatorily irrelevant to higher level. What Smolensky & Cussins are challenging is this claim (do I believe standard version? no!)
- Cussins: non-conceptual is at the level of experience (including consciousness). Smolensky: below the level of experience.
- Questions:
 - Does a "connectionist" level exist?
 - Is it intentional?
 - If so, what notion of semantics (or content)
 - Smolensky: relies on a theory of representation (even more so: goals)
 - Does it implement conceptual?
 - What properties cross boundaries?
 - Relation between "conceptual" and:
 - objective (world as something independent of agent)
 - non-perspectival (world in a way that is independent)
 - consciousness
 - Relation between theorist's and agent's conceptual (registrational) schemes
 - assumed = on conceptualist line
 - once ≠, opens up a possibility of non-conceptual content (cf. Cussins)
 - What properties are necessary for these splits?

a. General

-

a. Possible exam questions

- Compare and contrast:
 - Smolensky's "incomplete, informal, imprecise" analysis of the conceptual level and Haugeland's "second-order analog"
 - Cussin's "nonconceptual" and Smolensky's "subconceptual"
- How is Smolensky not merely an "implementationalist" wrt symbolic level. Or, rather, what is his notion of implementation, such that he isn't, and is that viable (true)?
- What properties of a computational system would be required to meet Smolensky's criteria. Cussin's criteria? (See Cussins p. 52 ¶ 2 sentence -1: "It may be that PTC needs C_3 more than C_3 needs PTC.").

—end of file —◆◆