

Syllabus for CS491/PSCH 494
Special Topic: *Introduction to Cognitive Science*
Fall semester, 2015
Instructor Stellan Ohlsson

Background and Purpose

Cognitive Science emerged in the aftermath of World War II, driven by the invention of what we now recognize as information processing technologies: radar, sonar, bomb sights, cryptography, etc., as well as information-related theories and sciences: information theory, Boolean logic, automata theory, cybernetics, decision theory, etc. The most impactful of these developments was the invention of digital computers and the associated software tools. The concept of a stored program computer, plus the ideas of symbolic (as opposed to arithmetic) data, and algorithms for processing such data were key advances. These technical and conceptual innovations inspired new approaches to old problems, the formulation of new problems, and the development of new research techniques, both empirical and theoretical. These new concepts, problems, tools, and forms of inquiry combined to form a new discipline, Cognitive Science. Its core is the application of a computational perspective to the study of complicated systems. A key feature of this approach is its interdisciplinary character. Studies in artificial intelligence, psychology, neuroscience, linguistics, philosophy, and the social sciences illuminate each other when viewed through the lens of computation. The cognitive sciences are also characterized by the use of computational modeling as an important research practice.

The purpose of this course is two-fold. At a general level, it aims to transmit the computational perspective to students at the 400 and 500 levels. To this end, the course will briefly cover the history and emergence of Cognitive Science, with emphasis on what was new and unique about this perspective, and the ways in which it extended earlier forms of inquiry. The course will proceed to an examination of basic concepts like representation and process, with explicit links to programming. Simple programming exercises, some in the form of actual computer programs and some in the form of paper-and-pencil exercises, will instantiate the general concepts. Attention will be given to the concept and practice of computational modeling of systems of various kinds. This type of modeling will be set in the context of the place and role of modeling in science in general. At the general level, the unifying theme of the course is Cognitive Science as the study of general intelligence. Consequently, the course is focused on cognitive psychology and AI, with less attention on linguistics, philosophy, and the social sciences.

At a more specific level, the goal is to familiarize the students with particular lines of work on specific cognitive problems, and to teach basic Cognitive Science modeling skills. The readings and exercises include some historical papers, review papers, as well as peer reviewed research articles. There is no text book.

The course will end with a brief look forward. What developments are rising over the horizon that are likely to impact future work on the computational analysis of complicated systems?

Method of Instruction

The course is primarily intended for students from the Departments of Computer Science and Psychology. It might also be of interest to students from the Learning Sciences Program, the Information and Decision Sciences program, and the UIC School of Education. The course will be taught by Stellan Ohlsson, Professor of Psychology and Adjunct Professor of Computer Science. The class will meet once a week. Each class period will include a mixture of lecture, discussion of readings, and hands-on training in particular cognitive science practices. Students will be graded on the basis of their level of participation in class discussions and the quality of a semester-long class project.

To accommodate the expected diverse academic backgrounds of the students, there is a range of alternative prerequisites. At the 200-level, either CS 251 Data Structures or PSCH 242 Research Methods is required. At least one of CS 411 Artificial Intelligence I, CS 421 Natural Language Processing, PSCH 459 Cognitive Methods, PSCH 457 Cognitive Psychology of Skill and Knowledge Acquisition, is recommended.

Instructions for Student Projects

The purpose of the course project is to acquire hands-on experience of computational modeling. Each student will program a model and investigate its behavior. The ambition level will depend on past modeling experience and programming skills. Because this is an interdisciplinary course, the prior knowledge is expected to vary greatly from student to student. Lack of prior experience is not an obstacle to taking the course. The instructor will work with each student to specify a project that is suited to his or her skill level. The important point is that each student stretches his or her modeling skills in the course of the semester.

Goal of Project Modeling begins with a choice of something to model. The most straightforward target for cognitive science modeling is a person performing a task of some (perhaps very simple) sort. However, any system is potentially a target for modeling: brain modules, social networks, sand piles, anthills, economic markets, species, and so on. The requirements are that (a) the system is complex (and any scenario involving humans doing tasks counts as complex, no matter how simple the task); (b) there is empirical research on the system, so that there is easily accessible information about the system's behavior; (c) there should be something interesting about that behavior, like a regularity or other type of finding that can be an interesting target for modeling; (d) the behavior of the target system changes over time (think learning curves, stock market fluctuations, etc.), so that a part of the

modeling work is to think about possible change mechanisms, one of the themes of the course. If the student cannot think of an appropriate modeling target, the instructor will assign one.

Execution Each class period will be devoted, in part, to a discussion of the readings, but also to instruction in modeling. The instructor will break down model construction into a series of steps. Students with advanced modeling skills need not follow these exactly, but students with less prior skills will benefit from doing so. The steps/topics are listed along with the readings for each week. The steps will be explained in lecture format, and then illustrated or exemplified by the instructor building a simple model alongside the students.

The models will preferably be written in the Python programming language. Alternatives include Lisp and Prolog. For students who do not know these languages, or do not have *any* prior programming experience, the instructor will make special accommodations.

The purpose of a model is to reproduce the target system's behavior. When the program runs, it should behave in the same way as the targeted system, at some level of granularity. Qualitative comparisons (curve fitting etc.) are desirable, but they are not always possible. The nature of the comparison depends on the nature of phenomenon or the types of changes in behavior over time. Preferably, the model should mimic the target system in such a way that it becomes understandable why the system exhibits the phenomenon or the changes in question.

Criteria for Successful Completion At the end of the semester, the student should submit three files as attachments to an email to the instructor. The first file should contain the Python code for the model. There should be enough comments for the instructor to understand the code. This file should be in such a form that the instructor can run the model as is.

Second, the student should submit a file that documents of least two runs with the model, exhibiting different behaviors, either across conditions or over time. The format of this file can vary, depending on the type of behavior modeled. For example, the runs might be documented by submitting a graph of the model's behavior.

Finally, there should be a third file with a 2-page text that explains how the model works, and how the model clarifies the modeled phenomenon, regularity, or change over time.

The three files should be submitted to the instructor on or before 5 pm on the last day of classes.

Grading of Project The modeling project can be aligned with the student's own research and/or other coursework, but it has to require work that is unique to this

course. Submitting modeling work done in another context contradicts the idea that the course should provide additional modeling experience.

Students can work individually or in teams of two. Larger teams are not allowed, unless there is some specific reason for an exception.

The students are encouraged to talk to each other about their modeling projects, to help each other, and to seek advice and input from the instructor. A portion of the class time each week will be devoted to issues and questions that come up in the various modeling projects.

As the semester unfolds, the instructor will ask students to volunteer to give brief reports on what they are working on. Due to the number of students, not all students will be able to do this. To encourage very preliminary reports, the grade on the project is based solely on the quality of the final report.

Reading List

The readings will be posted on the Blackboard site for the course and can be downloaded from there. There will be 1-3 readings per week. Many of the articles are exceedingly long. The instructors will issue readings instructions that identify the essential parts of the texts, point out sections that can be skipped etc. The readings will be discussed in class.

BEGINNINGS

Week 1: Origin and history of Cognitive Science. What are the essential characteristics and long-term goals of research in the cognitive sciences, especially cognitive psychology and artificial intelligence? Preliminary review of some basic concepts. Overview of course, practical matters, etc.

Readings:

Miller, G. A. (2003). The cognitive revolution: a historical perspective. *TRENDS in the Cognitive Sciences*, vol. 7, pp. 141-144.

Wang, P. (2006). Artificial Intelligence: What it is, and what it should be. In *AAAI Spring Symposium: Between a Rock and a Hard Place: Cognitive Science Principles Meet AI-Hard Problems* (p. 97-102).

Week 2: Computational models of intelligence. Modeling as a research practice. Essential features of computational models. Basic steps in building a computational model. General intelligence: The first theory.

Peschl, M. F., & Scheutz, M. (2001). Explicating the epistemological role of simulation in the development of theories of cognition. In *Proceedings of the 7th International Colloquium on Cognitive Science (ICCS-01)* (pp. 274-281).

Simon, H. A., & Newell, A. (1971). Human problem solving: The state of the theory in 1970. *American Psychologist*, vol. 26, pp. 145-159.

INTELLIGENCE AS THE ABILITY TO PERFORM COMPLEX TASKS

Week 3: General intelligence through so-called weak methods. Problem spaces and search strategies. How does the 1970 theory hold up?

Simon, H. A., & Reed, S. K. (1976). Modeling strategy shifts in a problem-solving task. *Cognitive Psychology*, vol. 8, pp. 86-97.

Atwood, M. E., Masson, M. E. J., & Polson, P. G., (1980). Further explorations with a process model of water jug problems. *Memory & Cognition*, vol. 8, pp. 182-192.

Week 4: Bounded rationality. Cognition is constrained by computational capacity limits. Working memory as example. Computational complexity in AI.

Simon, H. A. (1990). Invariants of human behavior. *Annual Review of Psychology*, vol. 41, pp. 1-19.

Yeh, Y.-C., Tsai, J.-L, Hsu, W.-C, & Lin, C. F. (2014). A model of how working memory capacity influences insight problem solving in situations with multiple visual representations: An eye tracking analysis. *Thinking Skills and Creativity*, vol. 13, pp. 153-167.

Week 5: Knowledge and expertise. Rule-based systems. How much knowledge is enough?

Buchanan, B. G., & Smith, R. G. (1988). Fundamentals of expert systems. *Annual Review of Computer Science*, vol. 3, 23-58.

Pandit, M. (2013). Expert systems – A review article. *International Journal of Engineering Sciences & Research Technology*, vol. 2, pp. 1583-1585.

Chen, Y., Hsu, C.-Y., Liu, L., & Yang, S. (2012). Constructing a nutrition diagnosis expert system. *Expert Systems with Applications*, vol. 39, pp. 2132-2156.

Week 6: Learning by doing. The acquisition of procedural knowledge. The necessity of practice. Multiple mechanisms for learning by practicing.

Anderson, J. R. (1987). Skill acquisition: Compilation of weak-method solutions. *Psychological Review*, vol. 94, pp. 192-210.

Laird, J. D., Rosenbloom, P. S., & Newell, A. (1986). Chunking in SOAR: The anatomy of a general learning mechanism. *Machine Learning*, vol. 1, pp. 11-46.

Week 7: Pulling it all together; cognitive architectures as models of general intelligence. How are the proposed architectures similar or different? Can current architectures reach general intelligence? How to evaluate and choose among architectures?

Duch, W., Oentaryo, R., & Pasquier, M. (2008). Cognitive architectures: Where do we go from here? In P. Wang B. Goertzel, & S. Franklin, (Eds.) *Artificial General Intelligence 2008: Proceedings of the First AGI Conference* (vol. 171, pp. 121-136). Amsterdam, The Netherlands: IOS Press.

Langley, P., Choi, Dongku, & Rogers, S. (2009). Acquisition of hierarchical reactive skills in a unified cognitive architecture. *Cognitive Systems Research*, vol. 10, pp. 316-332.

INTELLIGENCE AS THE ABILITY TO UNDERSTAND AND REASON

Week 8: Semantic memory. Declarative knowledge. Language and semantic memory. Concepts and categories; meaning, and reference. Semantic networks;

Quillian, M. R. (1969, August). The teachable language comprehender: A simulation program and theory of language. *Communications of the ACM*, vol. 12(8), pp. 459-475.

Lenat, D. B., Guha, R. V., Pittman, K., Pratt, D., & Shepard, M. (1990, August). Cyc: Toward programs with common sense. *Communications of the ACM*, vol. 33, pp. 30-49.

Week 9: Semantic inference. Basic semantic inferences. Semantic retrieval.

Szymański, J., & Duch, W. (2012). Information retrieval with semantic memory model. *Cognitive Systems Research*, vol. 14, pp. 84-100.

Liu, H., & Singh, P. (2004, October). ConceptNet – a practical commonsense reasoning tool-kit. *BT Technology Journal*, vol. 22, pp. 211-226.

Week 10: Abduction.

Bridewell, W., & Langley, P. (2011). A computational account of everyday abductive inference. In *Proceedings of the Thirty-Third Annual Meeting of the Cognitive Science Society* (pp. 2289-2294).

Wang, H., Johnson, T. R., & Zhang, J. (2006). The order effect in human abductive reasoning: an empirical and computational study. *Journal of Experimental & Theoretical Artificial Intelligence*, vol. 18, pp. 215-247.

Week 11: Explanatory Coherence

Thagard, P. (1989). Explanatory coherence. *Brain and Behavioral Sciences*, vol. 12, pp. 435 – 502 (N. B: *Read only pp. 435-467.*)

Spellman, B. A., Ullman, J. B., & Holyoak, K. J. (1993). A coherence model of cognitive consistency: Dynamics of attitude change during the Persian Gulf War. *Journal of Social Issues*, vol. 49, pp. 147-165.

Week 12: Theory change in science. How do scientists do it?

Thagard, P., & Findlay, S. (2011). Changing minds about climate change: Belief revision, coherence, and emotion. In E. J. Olsson and S. Enqvist, (Eds.), *Belief Revision Meets Philosophy of Science* (Chap. 14, pp. 329-343). Springer Science.

Rose, D., & Langley, P. (1986). Chemical discovery as belief revision. *Machine Learning*, vol. 1., pp. 423-451.

INTELLIGENCE IN PHYSICALLY REALIZED AGENTS

Week 13: Embodied cognition. Living in a body; models of perception and motor action; implicit vs. explicit cognition. Brain in a vat? Cognitive architectures as models of brains.

Chemero, A. (2013). Radical embodied cognitive science. *Review of General Psychology*, vol.17, pp. 145-150.

Patsenko, E. G., & Altmann, E. (2010). How planful is routine behavior? A selective-attention model of performance in the Tower of Hanoi. *Journal of Experimental Psychology: General*, vol. 135, pp. 95-116.

Week 14: Robotics; models of autonomous agents; the contribution of cognitive science to robotics.

Bogue, R. (2014). The role of artificial intelligence in robotics. *Industrial Robot: An International Journal*, 41(2), 119-123.

Kurup, U., & Lebiere, C. (2012). What can cognitive architectures do for robotics? *Biologically Inspired Cognitive Architectures*, vol. 2, pp. 88-99.

Puigbo, J.-Y., Pumarola, A., Angulo, C., & Tellez, R. (2015). Using a cognitive architecture for general purpose service robot control. *Connection Science*, vol. 27, pp. 105-117.

INTELLIGENCE AS AN EMERGENT PROPERTY OF AGENT NETWORKS

Week 15: Multi-agent networks. Communication; social networks; models of social creativity and innovation. Small world phenomena.

Macal, C. M., & North, M. J. (2010). Tutorial on agent-based modeling and simulation. *Journal of Simulation*, vol. 4, pp. 151-162.

Lazer, D., & Friedman, A. (2007). The network structure of exploration and exploitation. *Administrative Science Quarterly*, vol. 52, pp. 667-694.

(End)