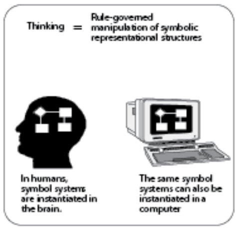

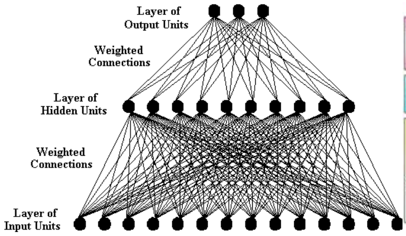

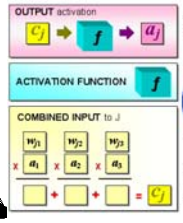

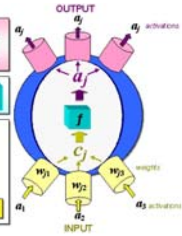

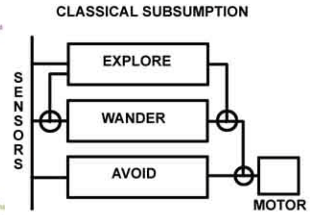

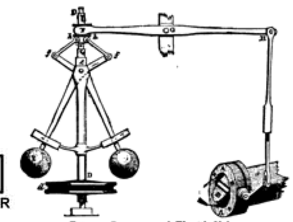
FIG. 4.—Governor and Throttle-Valve.

Cognitive Modeling, Intelligence, & Intelligent Systems

Phil 237 (Fall 2014): TR 11:00 – 12:20

Professor William Seeley, 315 Hedge Hall
Office Hours: TR: 12:30-1:30 & by appointment
wseeley@bates.edu

Course Description:
Artificial Intelligence is an interdisciplinary research field dedicated to the study of intelligence and intelligent systems. The field draws on materials from computer science, psychology, neuroscience, and philosophy. Its history is closely associated with the development of cognitive science. The course provides a historical introduction and overview of theories and methods within the field with a strong focus on the development of computational modeling and simulation as research methods. Coursework will include hands-on modeling and simulation exercises to enable students to explore the practical applications and limits of different models for intelligent behavior. No prior programming experience is required.

Course Goals:
The goals of this course are threefold. We will try to come to an understanding of what artificial intelligence is, and more importantly what it is not, as a research program in psychology and computer science. In this context we will evaluate the validity of several philosophical problems associated with artificial intelligence. This debate will be used to evaluate the traditional symbol system model for AI and introduce contemporary dynamic systems and behavior-based alternatives. Along the way we will evaluate, challenge, and develop our own common sense assumptions about the nature of intelligence.

** This seminar will be of interest to students with some background in cognitive science (e.g. students who have taken courses like Philosophy of Mind, Cognitive Psychology, Sensation and Perception, Brain Matters, Animal Learning, Physiological Psychology, or Formal Logic). However, the course does not presuppose any prior specialized knowledge of philosophy, psychology, or computer science.

Texts:
- *Mind Design II*, ed. John Haugeland (Cambridge, MA: MIT Press, 1997). ***(MD)***
- *Associative Engines*, Andy Clark (Cambridge, MA: MIT Press, 1993). ***(AE)***
- Electronic resources: online resources & pdf files on Lyceum ***(L)***.

Requirements:
- ***Class participation*** (including attendance) is worth +/- (5%) of your grade.
- **Cognitive modeling** and **artificial neural network** exercises(15%).
- A ***4-5 page analysis paper*** (15%).
- A ***6 page synthesis paper*** (25%).
- A ***12 page final paper*** on a topic of your choosing (40%). Topics must be cleared by week 10.

SCHEDULE OF READINGS:

Lecture 1, 09/04. Introduction

The goal of this session is to preview the semester by convening a discussion of our folk intuitions about intelligence.

Lecture 2, 09/09. A Model for Cognitive Science

The goal of this session is to introduce students to the basic framework of research in cognitive science developed by David Marr and discuss its implications for research in artificial intelligence.

David Marr (1982). *Vision* (New York: W. H. Freeman & Co., 1982), 8-37. ***(L)***

Lecture 3, 09/11. The Symbol System Hypothesis I: A Computational Theory of Mind

The goal of this session is to introduce students to some fundamental concepts in AI research and the basic framework of newell & Simon's *physical symbol system hypothesis.*

John Haugeland (1997). What is Mind Design? In ed. John Haugeland, *Mind Design II* (Cambridge, MA: MIT Press, 1997), 1-28. ***(MD)***

Lecture 4, 09/16. Automatic Formal Systems

The goal of this session is to introduce the concept of an *automatic symbol system* and open discussion about *Turing Machines*, computer architecture, and intelligence.

Alan Turing (1950). Computing Machinery & Intelligence. Reprinted in ed. John Haugeland, *Mind Design II* (Cambridge, MA: MIT Press, 1997), 29-56. ***(MD)***

John Haugeland (1996). Automatic Formal Systems. In *Artificial Intelligence: The Very Idea*. Cambridge, MA: MIT Press, 1996), 47-86. (***L)***

Lecture 5, 09/18. The Symbol System Hypothesis II: The Architecture of Intelligence

The goal of this session is to continue the discussion of *Turing Machines*, computer architecture, and computer intelligence.

Jack Copeland (2001). The Symbol System Hypothesis. In *Artificial Intelligence: A Philosophical Introduction* (Malden, MA: Blackwell Publishers, 2001), 58-82. ***(L)***

John Haugeland (1996). Computer Architecture. In *Artificial Intelligence: The Very Idea* (Cambridge, MA: MIT Press, 1996), 125-166. ***(L)***

The Mindproject, Turing Machine Module: Machine Tables & Turing Machines:
http://www.mind.ilstu.edu/curriculum/modOverview.php?modGUI=240

Lecture 6, 09/23. The Symbol System Hypothesis III: Strong AI

The goal of this session is to discuss a range of early implementations of the symbol system hypothesis.

Allen Newell & Herbert A. Simon (1976). Computer Science as Empirical Enquiry: Symbols & Search. In ed. John Haugeland, *Mind Design II* (Cambridge, MA: MIT Press, 1997), 81-110.***(MD)***

John Haugeland (1996). Real Machines. In *Artificial Intelligence: The Very Idea*. Cambridge, MA: MIT Press, 1996), 176-195. ***(L)***

Jack Copeland (2001). Some Dazzling Exhibits. In *Artificial Intelligence: A Philosophical Introduction* (Malden, MA: Blackwell Publishers, 2001), 11-26. ***(L)***

Jack Copeland (2001). General Problem Solver ii. In *Artificial Intelligence: A Philosophical Introduction* (Malden, MA: Blackwell Publishers, 2001), 87-91. ***(L)***

Lecture 7, 09/25. Heuristic Search Strategies

The goal of this session is to introduce knowledge representation as a potentially sticky issue for the symbol system hypothesis.

The Tower of Hanoi Problem ***(Onl)***
http://www-rci.rutgers.edu/~cfs/472_html/ProblemSolving/TOH_Strategies/TOH_TOC.html (broken)

Cannibals & Missionaries ***(ONL)***
http://www.aiai.ed.ac.uk/~gwickler/missionaries.html

The following supplementary resources are also useful to thumb through

Cannibals & Missionaries (John McCarthy):
http://www-formal.stanford.edu/jmc/elaboration/node2.html

General Problem Solver ***(Onl):***
http://ai-su13.artifice.cc/gps.html

Larry Learner Program:
http://www.mind.ilstu.edu/curriculum/modOverview.php?modGUI=209

A more detailed introduction to the general computational issues surrounding heuristic search strategies:
http://www.rci.rutgers.edu/~cfs/472_html/CogArch/CogArchToc.html (broken)

Lecture 8. 09/30. Knowledge Representation I: The Approach

The goal of this session is to introduce knowledge representation as a potentially sticky issue for the symbol system hypothesis.

Marvin Minsky (1975). A Framework for Representing Knowledge. In ed. John Haugeland, *Mind Design II* (Cambridge, MA: MIT Press, 1997), 111-142. ***(MD)***

John Haugeland (1996). Real Machines. In *Artificial Intelligence: The Very Idea*. Cambridge, MA: MIT Press, 1996), 195-203. ***(L)***

Jack Copeland (2001). Some Dazzling Exhibits. In *Artificial Intelligence: A Philosophical Introduction* (Malden, MA: Blackwell Publishers, 2001), 26-32. ***(L)***

Lecture 9, 10/02. The Chinese Room Problem

The goal of this session is to introduce and evaluate Searle's *Chinese room problem* as an objection to the strong symbol system hypothesis.

John Searle (1980). Minds, Brains, & Programs. In ed. John Haugeland, *Mind Design II* (Cambridge, MA: MIT Press, 1997), 183-204. ***(MD)***

Jack Copeland (2001). The Chinese Room from a Logical Point of View. In *Artificial Intelligence: A Philosophical Introduction* (Malden, MA: Blackwell Publishers, 2001), 109-122. ***(L)***

The Mind Project: The Chinese Room ***(Onl)***
http://www.mind.ilstu.edu/curriculum/searle_chinese_room/searle_chinese_room.php

Lecture 10, 10/07. Knowledge Representation II: The Frame Problem

The goal of this session is to introduce and evaluate *the frame problem* as an objection to the strong symbol system hypothesis.

Hubert L. Dreyfus (1979). From Micro Worlds to Knowledge: AI at an Impasse. In ed. John Haugeland, *Mind Design II* (Cambridge, MA: MIT Press, 1997), 143-182. ***(MD)***

John Haugeland (1996). Real Machines. In *Artificial Intelligence: The Very Idea*. Cambridge, MA: MIT Press, 1996), 203-212. ***(L)***

Jack Copeland (2001). The Knowledge Representation Problem. In *Artificial Intelligence: A Philosophical Introduction* (Malden, MA: Blackwell Publishers, 2001), 91-95. ***(L)***

Margaret Boden, *Mind as Machine*, Volume 2 (New York: Oxford University Press, 2006), 772-774. ***(L)***

Lecture 11, 10/09. Connectionism I

The goal of this session is to introduce *connectionist modeling* as an alternative to symbol system approaches.

David E. Rumelhart (1989). The Architecture of Mind: A Connectionist Approach. In ed. John Haugeland, *Mind Design II* (Cambridge, MA: MIT Press, 1997), 205-232. ***(MD)***

Paul Smolensky (1989). Connectionist Modeling. In ed. John Haugeland, *Mind Design II* (Cambridge, MA: MIT Press, 1997), 233-250. ***(MD)***

The Mind Project: Connectionism 1-3. ***(Onl)***
http://www.mind.ilstu.edu/curriculum/connectionism_intro/connectionism_1.php

Lecture 12, 10/14. Connectionism II

The goal of this session is to open discussion about *connectionist modeling* as an alternative to the symbol system hypothesis.

Andy Clark (1993). Computational Models, Syntax, and the Folk Solids. In *Associative Engines: Connectionism, Concepts, and Representational Change* (Cambridge, MA: MIT Press, 1993), 3-15. ***(AE)***

Andy Clark (1993). Connectionism, Code, and Context. In *Associative Engines: Connectionism, Concepts, and Representational Change* (Cambridge, MA: MIT Press, 1993), 17-40. ***(AE)***

FALL BREAK: 10/15 – 10/19

Lecture 13, 10/21. Connectionism III

The goal of this session is to evaluate what *connectionist models* might be said to know and not know, or more accurately information processing constraints on interpreting these simulations as models of intelligent behavior.

Andy Clark (1993). What Networks Know. In *Associative Engines: Connectionism, Concepts, and Representational Change* (Cambridge, MA: MIT Press, 1993), 41-68. ***(AE)***

Andy Clark (1993). What Networks Don't Know. In *Associative Engines: Connectionism, Concepts, and Representational Change* (Cambridge, MA: MIT Press, 1993), 69-86. ***(AE)***

Lecture 14, 10/23. Connectionism IV

The goal of this session is to continue discussion about the potential of *connectionist architectures*.

Andy Clark (1993). Concept, Category, Prototype. In *Associative Engines: Connectionism, Concepts, and Representational Change* (Cambridge, MA: MIT Press, 1993), 87-114. ***(AE)***

Lecture 15, 10/28. Connectionism V

The goal of this session is to evaluate *the systematicity argument* against connectionist models in artificial intelligence.

Jerry A. Fodor & Zenon W. Pylyshyn (1988). Connectionism and Cognitive Architecture: A Critical Analysis. In ed. John Haugeland, *Mind Design II* (Cambridge, MA: MIT Press, 1997), 309-350. ***(MD)***

Lecture 16, 10/30. Connectionism VI

The goal of this session is to discuss how knowledge effects and long term memory might be modeled in connectionist systems.

Andy Clark (1993). The Presence of a Symbol. In *Associative Engines: Connectionism, Concepts, and Representational Change* (Cambridge, MA: MIT Press, 1993), 115-130. ***(AE)***

Andy Clark (1993). The Role of Representational Trajectories. In *Associative Engines: Connectionism, Concepts, and Representational Change* (Cambridge, MA: MIT Press, 1993), 151-170 ***(AE)***

Lecture 17, 11/04. Connectionism, & Artificial Neural Networks: Text & Exercises I

The goal of these sessions is to acquire hands-on experience with connectionist and artificial neural network modeling in order to gain a better practical understanding of the promise and pitfalls of this approach to understanding cognition and modeling intelligent behavior. We will work through some of the exercises from each of the resources below.

Using A Simple Neural Network, Jeff Heaton
http://www.heatonresearch.com/online/introduction-neural-networks-java-edition-2/chapter-1/page4.html

The Mind Project: Connectionism 1-3. ***(Onl)***
http://www.mind.ilstu.edu/curriculum/connectionism_intro/connectionism_1.php

Kim Plunkett and Jeffrey Elman, *Exercises in Rethinking Innateness* (Cambridge, MA: MIT Press, 1997).
(see Tlearn folders on the Classes course page: \\netapp\etna\Classes\Fall2014\Group\Tlearn)

Supplemental: Other Simulators and Exercises

Michael R. W. Dawson (2008). *Connectionism: A Hands-On Approach*. Malden, MA: Wiley-Blackwell. (Chapters 1- 8, pages 1-47 + exercises)
(http://www.bcp.psych.ualberta.ca/~mike/Book3/Exercises/index.html)

Distributed Memory, Hebbian Learning, & The Delta Rule
(http://www.bcp.psych.ualberta.ca/~mike/Book3/index.html)

Dawson: Rosenblatt, & Rumelhart (software & exercises)
(http://www.bcp.psych.ualberta.ca/~mike/Software/)

Fred Cummins, *BasicProp*:
(http://basicprop.wordpress.com/)

Jay McClelland: *The PDP lab*
(http://web.stanford.edu/group/pdplab/resources.html#pdptool)

Lecture 18, 11/06. Connectionism, & Artificial Neural Networks: Text & Exercises II

The goal of these sessions is to acquire hands-on experience with connectionist and artificial neural network modeling in order to gain a better practical understanding of the promise and pitfalls of this approach to understanding cognition and modeling intelligent behavior.

Kim Plunkett and Jeffrey Elman, *Exercises in Rethinking Innateness* (Cambridge, MA: MIT Press, 1997).
Chapter 1, Introductions and Overview, pp. 1-18. ***(L & Reserve)***
Chapter 2, The Method of Simulations, pp. 19-30. ***(L & Reserve)***

Lecture 19, 11/11. Connectionism, & Artificial Neural Networks: Text & Exercises III

The goal of these sessions is to acquire hands-on experience with connectionist and artificial neural network modeling in order to gain a better practical understanding of the promise and pitfalls of this approach to understanding cognition and modeling intelligent behavior.

Kim Plunkett and Jeffrey Elman, *Exercises in Rethinking Innateness* (Cambridge, MA: MIT Press, 1997).
Chapter 3, Learning to Use the Simulator, pp. 31-74. ***(L & Reserve)***
Chapter 4, Learning Internal Representations, pp. 75-99. ***(L & Reserve)***

Lecture 20, 11/13. Connectionism, & Artificial Neural Networks: Text & Exercises IV

The goal of these sessions is to acquire hands-on experience with connectionist and artificial neural network modeling in order to gain a better practical understanding of the promise and pitfalls of this approach to understanding cognition and modeling intelligent behavior.

Kim Plunkett and Jeffrey Elman, *Exercises in Rethinking Innateness* (Cambridge, MA: MIT Press, 1997).
Chapter 5, Autoassociation, pp. 99-114. ***(L & Reserve)***
Chapter 6, Generalization, pp. 115-137. ***(L & Reserve)***

Lecture 21, 11/18. Connectionism, & Artificial Neural Networks: Text & Exercises VI

The goal of these session is to acquire hands-on experience with connectionist and artificial neural network modeling in order to gain a better practical understanding of the promise and pitfalls of this approach to understanding cognition and modeling intelligent behavior.

Kim Plunkett and Jeffrey Elman, *Exercises in Rethinking Innateness* (Cambridge, MA: MIT Press, 1997).
Chapter 8, Simple Recurrent Networks, pp. 151-163. ***(L** & **Reserve)***
Chapter 10, Modeling Stages in Cognitive Development, pp. 179-203. ***(L** & **Reserve)***

Lecture 22, 11/18. Connectionism, & Artificial Neural Networks: Text & Exercises VI

The goal of these session is to acquire hands-on experience with connectionist and artificial neural network modeling in order to gain a better practical understanding of the promise and pitfalls of this approach to understanding cognition and modeling intelligent behavior.

Kim Plunkett and Jeffrey Elman, *Exercises in Rethinking Innateness* (Cambridge, MA: MIT Press, 1997).
Chapter 12, The Importance of Starting Small, pp. 233-259. ***(L** & **Reserve)***

THANKSGIVING BREAK!!!!!!!

Lecture 23, 11/20. Behavior-Based Robotics & Biorobotics Modeling

The goal of this session is to introduce *behavior-based robotics* as an alternative to GOFAI and Connectionist approaches to artificial intelligence.

Randall Beer, Hillel Chiel, & Leon Sterling (1992). A Biological Perspective on Autonomous Agent Design. In Patti Maes, *Designing Autonomous Agents: Theory and Practice from Biology to Engineering and Back*. Cambridge, MA: MIT Press, 1992, 169-186. ***(L)***

Barbara Webb (2001). A Spiking Neuron Controller for Robot Phonotaxis. In eds. Barbara Webb and Thomas R. Consi, *Biorobotics: methods and Applications*. Cambridge, MA: MIT Press, 2001, 3-20. ***(L)***

Lecture 24, 11/20. Dynamic Systems Approaches

The goal of this session is to introduce a *dynamic systems* approach to modeling intelligent behavior as an alternative to approaches to artificial intelligence.

Randall Beer (2000). Dynamical Approaches to Cognitive Science. *TRENDS in Cognitive Sciences*, 4(3): 91-99. ***(L)***
Andy Clark (2014). Dynamics. *Mindware*, 2nd Edition. Cambridge, MA: MIT Press, 2014, 140-141; 145-156. ***(L)***

ASSIGNMENTS: All assignments must be handed in both in hard copy and electronically via the dropbox for that assignment on LYCEUM

Analysis Paper (1400 words): The goal of this paper is to critically evaluate the argument presented in syllabus reading. Focus should be placed on critically evaluating the reasoning and evidence presented in support of the argument in the reading.

Due Date: October 5, 2014

Synthesis Paper (1800 words): The purpose of this paper is to demonstrate that you can identify & evaluate a standard argument in the literature and synthesize the diverse range of material covered in the first half of the semester into a coherent position.

Due date: November 21, 2014

Final Paper (1800 words): The purpose of this paper is to demonstrate that you can identify & evaluate a standard argument in the literature, synthesize the diverse range of material covered on the syllabus into a coherent position, and incorporate comments on revisions to a draft. Please write on a topic of your choosing.

Topic Due Date: Week 12
Paper Due Date: The scheduled exam date (there is no final exam), December 11, 2014.

Computational Modeling Exercises:

I will assign several exercises in October exploring the nature of algorithms, symbol systems, and heuristic search for students to work through on their own (see schedule of readings for dates and links to software & exercises).

We will meet for three weeks in November in the computer classroon at The Imaging Center to run through a range of connectionist / artificial neural network modeling exercises (see schedule of readings for dates and links to software & exercises).

Turing Machines: (for Lecture 5)	Work through the *Machine Tables* & *Turing Machine* exercises in the "Introduction to Symbolic Models" module on The Mindproject website on your own. We will discuss these execises in class (see syllabus above).
Heuristic Search: (for Lecture 7)	Work through the online materials for the *Tower of Hanoi* and *Cannibals and Missionaries* problems.
Excel modeling exercises: (due Lecture 18)	Please use *Excel* to replicate the simple neural networks for the logical functions AND, OR, XOR introduced in the *Uusing a Simple Neural Network* materials online (see modeling exercises sheet below)
Tlearn modeling exercises: (due dates –see below)	AND, OR, XOR, Autoassociation, Generalization, Recurrent Networks, Phased Learning (see chapters from Plunkett & Elman *Exercises in Rethinking Innateness* listed on the syllabus above)
** The James Program:	Distributed Memory; Hebbian Learning; The Delta Rule (see supplemental materials from Dawson, *Connictionism: A Hands-On Approach* listed on the syllabus above)

**** Due Dates for Tlearn exercises: Sundays at midnight at the end of weeks 9, 10, & 11.

Some Miscellaneous Notes and Guidelines:

Moral behavior is the grounds for, and the framework of, a healthy society. In this regard it is each of our responsibility as an individual within the community of our classroom to act responsibly. This includes following the rules and guidelines set out by Bates College for academic behavior. Plagiarism is a serious matter. It goes without saying that each of you is expected to do his or her own work and to cite EVERY text that is used to prepare a paper for this class.

Please familiarize yourself with the guidelines for academic integrity posted on the Bates Website:
http://www.bates.edu/entering/policy/judicial-affairs/code-of-student-conduct/academic-misconduct/

This is a seminar. This means that the content of the course, and our progress through the syllabus, should ideally be student driven. I have designed the course to allow us some flexibility so that we can spend more time on issues of interest to the class. I reserve the right to make changes to the syllabus as we go along in order to accommodate our interests as they emerge in class discussions. I will also occasionally upload supplementary materials to *Lyceum* for students interested in pursuing particular issues beyond class discussion.

SCHEDULE OF READINGS		
The reading schedule that follows is a loose guideline for our progress through the syllabus. It is open to change at the Professor's discretion contingent on the pace of the class and evolving interests of the group.		
Date	**Readings**	**Assignments (due Fridays: 5pm)**
1: 09/03	No Readings	
2: 09/09	Marr, *Vision* (excerpt): 8-37. ***(L)***	
3: 09/11	Haugeland, What is Mind Design?: 8-21. ***(MD)***	
4: 09/16	Turing, Computing Machinery…: 29-56. ***(MD)*** Haugeland, Automatic Formal Systems*:* 47-86*.* **(L)**	**Ruring Machine Exercises**
5: 09/18	Copeland, The Symbol System Hypothesis: 58-82. ***(L)*** Haugeland, Computer Architecture: 125-166. ***(L)***	
6: 09/23	Newell & Simon, Computer Science as…:81-110. ***(MD)*** Haugeland, Real Machines: 176-195. ***(L)*** Copeland, Some Dazzling Exhibits: 11-26. ***(L)***	**Heuristic Search Exercises**
7: 09/25	Minsky, A Framework…: 111-142. ***(MD)*** Haugeland, Real Machines: 195-203. ***(L)*** Copeland, Some Dazzling Exhibits: 26-32. ***(L)***	
8: 09/30	***NO CLASS!!***	
9: 10/02	Searle, Minds, Brains, Programs: 183-204. ***(MD)*** Copeland, The Chinese Room from…: 109-122. ***(L)*** The Mind Project: The Chinese Room. ***(Onl)***	
10: 10/07	Dreyfus, AI at an Impasse: 143-182. ***(MD)*** Haugeland, Real Machines: 203-212. ***(L)*** Copeland, knowledge representation: 91-95. ***(L)*** Boden, the frame problem: 772-774. ***(L)***	
11: 10/09	Rumelhart, The Architecture of…: 205-232. ***(MD)*** Smolensky, Connectionist Modeling: 233-250. ***(MD)*** The Mind Project, Connectionism. ***(Onl)***	
12: 10/14	Clark, Computational Models: 3-16. ***(AE)*** Clark, Connectionism, Code, & Context: 17-40. ***(AE)***	**Analysis Paper: Tuesday evening**
Fall Recess: 10/15 – 10/19		
13: 10/21	Clark, What Networks Know: 41-68. ***(AE)*** Clark, What Networks Don't Know: 69-86. ***(AE)***	
14: 10/23	Fodor & Pylyshyn, Connectionism and…: 309-350. ***(MD)***	
15: 10/28	Clark, The Presence of a…: 115-130. ***(AE)*** Clark, The Role of…: 151-179. ***(AE)***	
16: 10/30	Clark, The Presence of a…: 115-130. ***(AE)*** Clark, The Role of…: 151-179. ***(AE)***	

17: 11/04	Using a Simple Neural Netork ***(Onl)*** Artificial Neural Networks ***(Onl)***	
18: 11/06	Plunkett & Elman, Introduction & Overview : 1-18. ***(L)*** Plunkett & Elman, The Methodology of Simulations: 19-31. ***(L)***	**Excel Modeling Exercises** **P&E, Chapter 1 Exercises**
19: 11/11	Plunkett & Elman, Learning to Use the Simulator: 32- 74. ***(L)*** Plunkett & Elman, Learning Intenal Representations: 75-98. ***(L)***.	
20: 11/13	Plunkett & Elman , Autoassociation: 99-114. ***(L)*** Plunkett & Elman, Generalization: 115-136. ***(L)***	
21: 11/18	Plunkett & Elman, Simple Recurrent networks: 151-162. ***(L)*** Plunkett & Elman, Modeling Stages of Cognitive Development: 179-202. ***(L)*** *supplemental advanced exercises*: Dawson, *Connectionism*: 1-29. ***(L)***	
22: 11/20	Plunkett & Elman , The Importance of Starting Small: 233-258. ***(L)*** *supplemental advanced exercises*: Dawson, *Connectionism*: 30-47. ***(L)***	**Synthesis Paper: Friday @ 4**
THANKSGIVING BREAK!!!!!!!!!		
23: 12/02	Beer, Chiel, & Sterling, A Biological perspective on…: 169-186. ***(L)*** Webb, A Spiking Neuron Controller…: 3-20. ***(L)***	
24: 12/04	Beer, Dynamical Approaches to…: 91-99. ***(L)*** Clark, Dynamics: 140-165. ***(L)***	**Final Paper Topics Due**
Final 12/12	**NO FINAL EXAM**	**Final paper**

Computational Modeling Exercises: *Tlearn*, *Exercises in Rethinking Innateness*, class exercises
Computational Modeling, Intelligence, and Intelligent Systems (Phil 237)
Professor Seeley, Fall 2014

Tlearn = software for Plunkett & Elman, *Exercises in Rethinking Innateness*

We will work through the exercises in chapters 1-6 in Plunkett & Elman's *Exercises in Rethinking Innateness* (Cambridge, MA: MIT Press, 1997). These chapters have been loaded to LYCEUM and I have updated the syllabus to match our expected progress through the material. If we have time we will work through the exercises in chapters 8, 10, & 12 as we move through them as well.

I will arrange for four 2-hour workshop sessions in the Coram computer classroom. These are tentatvely scheduled (pending your schedules) for W/Th November 19 & 20 and W/Th December 3 & 4 from 2-4 pm.

Tlearn is a neural network simulator designed for the exercises in the Plunkett & Elman textbook. Work through the exercises in each chapter using the *Tlearn* simulator. Answers for the exercises are provided at the end of each chapter. Please:

- complete the exercises
- correct your work, providing an explanation of your understanding of any errors you might have made in italicized text following your original answer. Your goal in these cases is to demonstrate an understanding of the material by providing an explanation of your errors
- type your answers and corrections up in a word file (or write them very neatly and scan them as a pdf file) and turn them in through the dropbox for each chapter.

I have also constructed a set of *Excel modeling exercise* and a larger *Tlearn* exercise that combines all of the logic exercises together (see below). We'll talk about these exercises in class, but have a look and be thinking how to solve it as you learn to use the simulator.

I have uploaded the software and project files to our class folder on the *etna server*. click on the Windows icon in the lower left corner of your desktop and search for - \\netapp\etna\Classes\Fall2014 - Phil237a\ This should bring up a list of classes. Our class is *Phil 237A*.

Everyone has a folder in the "Students" folder. Go to the "Groups" folder, click on the "Tlearn" folder, and copy the "wintlrn1.0.3+exercises" folder into your "Students" folder. When you are running exercises or experimenting with the simulator you will need to copy this folder onto your desktop. The program is old and a little buggy and I have found this is the most successful strategy. However, if the program still crashes, remember that we had to go a step further in class and copy the *.prj files (e.g. *and*; *and.cf*; *and.dat*; *and.teach*) to the desktop to run the simulator.

Last thing. Please set up a folder in your student folder where you can save copies of the wintlrn1.0.3+exercises, any project files, weight files, or etc. you might write while you are experimenting with the simulator (either via exercises or on your own). It is good to keep a record of your work this way and can be helpful to refer back to.

Also, it is critical that you do not make any changes to the wintlrn1.0.3+exercises folder in the Groups folder – if you do you will have altered the master file everyone is using to run their exercises. Instead make a copy for yourself and store it in your Student folder.

Good Luck!

Have Fun!

Bill

Connectionist / Artificial Neural Network Modeling Exercises

The general rule in all the following exercises is that tinkering is good, failing is better (when it leads to more tinkering), and although succeeding is best…it's better if it's built on bumps in the road (and more tinkering).

Excel modeling exercises:

Please use Excel to replicate the neural networks for AND, OR, and XOR that we introduced in class from Jeff Heaton's *Using Simple Neural Networks* < http://www.heatonresearch.com/online/introduction-neural-networks-java-edition-2/chapter-1/page4.html>

The example we ran through in class is loaded to our Groups folder on the *etna server* to help you get started.

Due date: end of week 9, Sunday November 9 @ midnight.

Tlearn chapter exercises:

Please work through the exercises in chapters 1-4 of *Rethinking Innateness* (on LYCEUM) and type them up in a word file to hand in on LYCEUM.

After you complete the exercises please use the answer keys in the back of the chapters to correct them.

Please use *italic font* to write up explanations of any errors or questions that you could not answer. Include your corrections in the text of your file just following your answers.

Due dates:

Chapter 1: end of week 9 / Sunday November 9 @ midnight.

Chapters 3-6: end of week 10 / Sunday November 16 @ midnight.

* Chapters 8 & 10: end of week 11 / Friday, November 21 @ 5pm.

* Chapter 12: end of week 12 / Friday, December 5 @5pm.

Tlearn class exercises:

a) Combine the data sets and teacher files from the Tlearn exercises in Chapters 3 & 4 of *Exercises in Rethinking Innateness* and generate a single network that "knows" these three basic logic functions.

What would the input vectors have to look like to accomplish this?

How many input nodes would you need to add if you wanted to explicitly label each input in the training set appropriately (match inputs to target outputs in the teacher file)? How would this alter your network architecture and configuration file? Do you need to do this or would the *.teach file suffice to label the input vectors in the training set?

Does your network work?

b) There is one more basic logical function that has been left out of the Tlearn exercises: if..then. Conditional statements are only false when their antecedent is true and their consequent false.

1	0 0
0	1 0
1	0 1
1	0 1

Can you use this data set and teacher file with the others above to construct a network that successfully learns this logical function?

Can you combine this data set and teacher file with the others to construct a single network that "knows" all four basic logic functions?

c) Can you replicate Elman's "phased learning" strategy to incrementally train the network to learn all four logic functions (or just the first 3?)?

Is there any difference in the learning behavior (is one faster than the other)?

Due Date: end of week 12 / Friday, December 5 @5pm.

Supplementary Resources:

Stanford, Neural Networks: http://cs.stanford.edu/people/eroberts/courses/soco/projects/neural-networks/Neuron/index.html

Rocks & Mines Artificial Neural Network

R. Paul Gorman and Terrence J. Sejnowski (1988). Analysis of Hidden Units in a Layered Network Trained to Classify Sonar Targets. *Neural Networks* 1: 75-89. ***(L)***

Rocks & Mines Data Set. University of California Irvine Machine Learning Repository https://archive.ics.uci.edu/ml/datasets/Connectionist+Bench+(Sonar,+Mines+vs.+Rocks)

One potential advanced project is to try to replicate Gorman & Sejnowski's project using the archived data set form the University of california, Irvine website.