

# Psychology 878

## Human Learning II: Memory

Course Syllabus, Spring 2007

- Course:** PSYCH 878  
**Call number:** 17986-3  
**Credits:** 4  
**Dates:** Jan 4–March 8, 2007  
**Times:** Mondays 09:30–11:48 a.m.  
**Room:** Psychology Bldg, Room 0117  
**Prerequisites:** No formal prerequisites. Some knowledge of cognitive psychology, experimental design, neuroscience, and college-level math is assumed.  
**Websites:** <https://carmen.osu.edu> and <http://alexpetrov.com/teach/ccn/>  
**Textbook:** Randall O'Reilly & Yuko Munakata (2000). *Computational Explorations in Cognitive Neuroscience: Understanding the Mind by Simulating the Brain*. Cambridge, MA: MIT Press.  
**Instructor:** Dr. Alexander Petrov  
(614) 247-2734 [ccncourse@orion.psy.ohio-state.edu](mailto:ccncourse@orion.psy.ohio-state.edu)  
200B Lazenby Hall Office hours: By appointment

### Course Overview

How does cognition emerge from the brain? This course introduces you to the new and exciting field of Computational Cognitive Neuroscience (CCN) that provides important pieces of the answer to this question. We focus on simulations of cognitive and perceptual processes, using neural network models that bridge the gap between biology and behavior. We adopt the *Leabra* framework of Randy O'Reilly and Yuko Munakata, and use their 2000 book as the main text for the course. We first consider briefly the basic biological and computational properties of individual neurons and networks of neurons, as well as their idealized *Leabra* counterparts. We discuss their role in basic processing mechanisms such as spreading activation, inhibition, and multiple constraint satisfaction. We then discuss learning mechanisms that allow networks of neurons to build internal models of their environments and perform complex tasks. Models illustrating these ideas will be demonstrated in class and explored in homework assignments. We complement the simple demos with a few case studies of full-blown models of various aspects of perception, language, and memory. Finally, we turn to big-picture issues and present (an outline of) a comprehensive connectionist proposal of a cognitive architecture. We discuss how different brain systems (e.g., hippocampus, parietal cortex, frontal cortex) specialize to solve difficult computational tradeoffs. The course introduces the PDP++ neural network simulator, which will enable you to explore the variety of models that come with the book.

## **Textbook**

The main textbook is *Computational Explorations in Cognitive Neuroscience* (O'Reilly & Munakata, 2000). Various resources are available from Randy's web site at [http://psych.colorado.edu/~oreilly/comp\\_ex\\_cog\\_neuro.html](http://psych.colorado.edu/~oreilly/comp_ex_cog_neuro.html) Additional readings are listed in the lecture plan below.

## **Prerequisites**

There are no formal prerequisites for this course, other than graduate standing. Some knowledge of cognitive psychology, experimental design, neuroscience, and college-level math is assumed. Obviously, prior exposure to neural network models (e.g., Psych 617) and/or mathematical models (e.g., Psych 609) will be extremely useful, although the present course has little overlap with these two courses. While the models we will be using are mathematically based, only algebra and some simple calculus-level concepts are involved. The focus will be on intuitive and practical applications. We rely primarily on computer simulations to explore the models and develop intuitions about how they work. Computer programming experience is not required because the models are accessible via a graphical interface. Still, you will be expected to install the PDP++ simulator on your own and be able to run a few "canned" exercises with selected models.

## **Teaching Method**

Classes will be part lecture and part discussion/tutorial. There will be many in-class demonstrations. The course will require preparation prior to each class: reading chapters from the textbook, additional readings, and hands-on explorations with the PDP++ software. Short summaries of the readings will be collected each week and form 15% of the final grade.

## **PDP++ Software**

Each student should have access to the PDP++ simulator that comes with the book. It is open-source and available on all major platforms: Linux, Mac OS X (preferred), and Windows. Precompiled binaries (and C sources) are available for download from [http://psych.colorado.edu/~oreilly/cecn\\_download.html](http://psych.colorado.edu/~oreilly/cecn_download.html) It is assumed that all graduate students have access to a personal computer and are willing to install the software. If these assumptions are not met, access to a computer in Lazenby 015 will be provided. Please notify the instructor at least a week in advance to activate this option.

Hands-on experience with actual running models is an important part of the course. It is a unique exploratory learning opportunity. The difficulty level will be matched to the average ability of the students in the class. We will begin with very simple exercises and progress to more complex ones depending on your interest and skill. The simulator gives complete control over all aspects of the network and its training environment. It also provides dynamic, colorful visualizations that are an indispensable tool for developing intuitions about how complex cognition can emerge from a network of neurons.

## Accommodations for Students with Special Needs

The policy of The Ohio State University is to provide every reasonable, appropriate, and necessary accommodation to qualified disabled students. The University's colleges and academic centers evaluate and judge applications on an individual basis and no categories of disabled individuals are automatically barred from admission. The privacy rights of each disabled person are honored to the fullest extent possible. The University's interest in a student's disabilities are only for the purpose of accommodating his/her specific disability, thereby providing an academically qualified disabled student access to programs and activities accorded all other qualified students. Whenever generally accessible facilities do not adequately accommodate a specific disability, the University makes every reasonable accommodation and program or facility adjustment to assure individual access. These policies are fully supported and practiced in this class.

If you have a disability documented with the Office of Disability Services (<http://www.ods.ohio-state.edu>, 150 Pomerene Hall, 614-292-3307), please contact Dr. Petrov privately ([petrov.11@osu.edu](mailto:petrov.11@osu.edu), 200B Lazenby Hall, 614-247-2734) by the end of the second week of classes (4/06/2007) so that any accommodations can be made.

## Evaluation

Your grade will depend on four components in the following proportions:

Reading reactions	15%
Model explorations	35%
Class participation	15%
Final paper	35%

Grades are based on absolute cutoffs: A=90-100%, B=80-89%, C=70-79%, D=60-69%.

**Reading reactions:** For each chapter, you will be asked to write a few sentences about the topic you found most interesting in the chapter and why. Reading reactions (RRs) must be uploaded to the "Reading Reactions" dropbox on the Carmen website, prior to the class meeting they are due. Acceptable formats are MS Word (.doc), plain text (.txt), PDF, HTML, or RTF format. RRs in other formats, submitted by email, slipped under the door, etc. will not be graded and do not bring any points.

**Model explorations:** The textbook comes with a large number of pre-built neural network models that illustrate key principles and phenomena. Every week, you will explore these pre-built models and document your explorations by answering questions from the textbook. I will tell you which exercises to do for each chapter. If you start doing the exercises before you receive the assignment, you risk doing some exercises that are not in the assignment (which is not the end of the world, but you have been warned). You should write up all assigned simulation exercises for each chapter and upload them to the corresponding Carmen dropbox. The acceptable formats are the same as for the reading reactions. Note that the model exploration report for each chapter must be in a separate file and dropped in a separate dropbox, prior to the date specified in the Calendar

section of the syllabus. Late submissions will be penalized 2% for each day after the due date. Reports submitted by email, slipped under the door, etc. will not be graded and do not bring any points. Teamwork on these explorations is encouraged, but you must write your reports *individually*. We want to see that each individual person understands the material, and so this should be evident in your writing. It is best to write down results and first drafts of answers as you work through the simulations; they can sometimes take a while and you do not want to have to run them repeatedly.

**Class participation:** Productive participation in class discussion is encouraged to help you get the most out of this course. You are expected to read the text chapters the week they are assigned and to come to class prepared to discuss the issues and answer questions. Attendance is required: one percentage point will be deducted from your participation score for each absence. Besides, if you are absent, you forego the opportunity to score positive points for participating in the discussion.

**Final paper:** 35 percent of your grade are based on a final paper. It must be uploaded to the “Final Paper” dropbox on the Carmen website by 11:59 pm on Monday, June 4. It must be 8-12 pages long (double-spaced, 12-pt font, excluding figures) and be submitted in one of the acceptable file formats. Late penalty 5% per day. For the final paper, you can either (a) use simulations to examine some psychological phenomenon of interest to you or (b) do a literature search and write a review that compares a neural network model to another model (neural network or otherwise) in light of some behavioral, neuropsychological and/or neurophysiological data. I expect (and recommend) that most of you will choose option (b), but the more challenging option (a) is available too. If you choose (a), *do not be overly ambitious* – relatively simple but thoughtful work is preferred to a complicated half-baked mess. Do not be misled by the relative simplicity of running the canned exercises in the book. Developing a simulation from scratch takes a very long time.

Literature review (option b) will require more reading and integration of scientific articles. While this does not involve you in actual simulations, it does require in-depth understanding of publishable research on connectionist modeling. You will have to compare at least two published models and how they differentially account for your phenomenon of interest. One (or both) of these models must be a neural network model. The minimum review that can still bring you 35% towards your grade would take some of the models in the textbook and compare it against the original publication of the same model. For example, the Leabra implementation of the triangular model presented in Section 10.3 in the book (Figure 10.5) can be compared against the original backprop implementation of Plaut & Shallice (1993). Your review paper should contain a concise introduction to the psychological issue or phenomenon and outline the principle(s) instantiated by the two models. It should contain methods, results, and a concluding discussion of the significance of the results, how the models can be improved, etc.

By Sunday, May 06, you should upload on Carmen a one-page proposal for your final paper. This proposal should specify whether you plan to do (a) a simulation or (b) a literature review. It should also contain a summary of your question of interest, your proposed approach to explore this question through simulation (a) or the two models you

propose to compare (b). This one-page proposal is worth 5% toward the 35% allocation for the final paper. Your proposal must be approved by the instructor as a prerequisite for the final paper. Unapproved papers will not be graded and do not bring any points.

## Academic Ethics

All students enrolled in OSU courses are bound by the Code of Student Conduct ([http://studentaffairs.osu.edu/resource\\_csc.asp](http://studentaffairs.osu.edu/resource_csc.asp)). Suspected violations of the Code will be dealt with according to the procedures detailed in the Code. Specifically, any alleged cases of misconduct will be referred to the Committee on Academic Misconduct.

## Course Calendar

### 1. March 26: Introduction and overview. Neurons (I).

Getting started. The appeal of Parallel Distributed Processing (PDP). Overview of the course. Overview of *Leabra*. Overview of the PDP++ simulator. A network in action: constraint satisfaction. Jets-and-Sharks model (aka Cats-and-Dogs model, Sect. 3.6.4). Necker cube model (Sect. 3.6.5). The neuron as a detector.

Readings:

- Chapter 1 in O'Reilly & Munakata (2000, henceforth O&M).
- O'Reilly, R. C. (1998). Six principles for biologically-based computational models of cortical cognition. *Trends in Cognitive Sciences*, 2, 455-462. <http://psych.colorado.edu/~oreilly/pubs-abstr.html#OReilly98>
- McClelland, J. L., Rumelhart, D. E., & Hinton, G. E. (1986). The appeal of parallel distributed processing. In D. Rumelhart & J. McClelland (Eds.) *Parallel Distributed Processing, Vol. 1: Foundations*. Cambridge, MA, MIT Press. (Available on Carmen in read-only PDF.)
- Skim through O&M Chapter 2. Must read only the summary (Sect.2.10).

Deliverables: Reading reactions for McClelland et al. (1986), **due 3/30**, 11:59 pm.

### 2. April 02: Neurons (II) and Networks

Point-neuron activation function (Box 2.2 in O&M, p. 49). Biology of the neuron. Cortical networks. Localist vs. distributed representations. Feedforward and recurrent networks. NetTalk demo. Pattern completion. Inhibition and its k-winner-takes-all (kWTA) approximation. Constraint satisfaction. Jets-and-Sharks and Necker-cube models revisited.

Readings:

- O'Reilly, R. C. (1998), *TiCS*, 2, 455-462 – see full citation above
- Skim through O&M Chapters 2 and 3. Read all summaries and boxes carefully. You will need Sections 2.6.3, 3.6.4, and 3.6.5 for the assigned model explorations.
- [Optional] Sejnowski, T. J. & Rosenberg, C. R. (1987). Parallel networks that learn to pronounce English text. *Complex Systems*, 1, 145-168. (PDF reprint and audio file [nettalk-audio-demo.mp3](#) available on Carmen.)

Deliverables: Reading reactions for O'Reilly (1998), **due 4/01**, 11:59 pm.

Model explorations (5%): Exercises 2.7, 2.9, 3.15, and 3.16, **due 4/08**, 11:59 pm.

### 3. April 09: Hebbian Model Learning

Learning in neural networks. Hebbian vs. error-driven learning: what is the difference? Biological mechanisms of learning. Computational objectives of learning. Conditional PCA. Self-organizing learning. Explorations of Hebbian (Sect. 4.6) and self-organizing learning (Sect. 4.8.1).

Readings:

- O&M Chapter 4. Feel free to skip all mathematical derivations, as well as the technicalities of renormalization and contrast enhancement (Sect. 4.7).
- [Completely optional] Kandel, E. (2001). The molecular biology of memory storage: A dialogue between genes and synapses. *Science*, 294, 1030-1038. (Based on Eric Kandel's acceptance speech for the 2000 Nobel Prize for Medicine. PDF available on Carmen.)

Deliverables: Reading reactions for O&M Chapter 4, **due 4/08**, 11:59 pm.

Model explorations: Nothing is due this week but you should do Exercises 4.3 and 4.6 in preparation for next week's assignment.

### 4. April 16: Applications of Hebbian Learning

Learning self-organizing maps in primary visual cortex (Sect. 8.3). Learning semantic representations from word co-occurrences (Sect. 10.6).

Readings:

- O&M Sections 4.5, 4.8, 8.3, and 10.6.
- [Optional] Olshausen, B. A. & Field, D. J. (1996). Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381, 607-609. (PDF reprint available on Carmen.)

Deliverables: RRs for Sect. 8.3 if you chose to explore the model in Sect. 10.6; RRs for Sect. 10.6 if you chose to explore that in Sect. 8.3. **Due 4/15**, 11:59 pm. Model explorations (5%): Choose one of the following sets: (a) Ex. 4.3, 4.6, 8.1, and 8.2; (b) Ex. 4.3, 4.6, 10.11, and 10.12. **Due 4/22**, 11:59 pm. For 2% extra credit, you are welcome to do both sets of exercises.

### 5. April 23: Error-Driven Task Learning. Combined Hebbian+Task Learning

Learning of arbitrary input-output mappings. XOR problem. Delta rule and its generalizations. The GeneRec learning algorithm computes error from activation differences. Functional pros and cons of Hebbian and error-driven learning. The advantages of combining them. Inductive biases. Generalization.

Readings:

- Skim through O&M Chapters 5 and 6, up to Sect. 6.4 inclusive. Feel free to skip all mathematical derivations. Make sure to read Sections 5.8, 5.10, 6.2, 6.8, and Box 6.1.

Deliverables: RRs for the "make-sure" sections listed above, **due 4/22**, 11:59 pm.

Model explorations (5%): Exercises 5.2, 5.5a, 5.5c, and 6.2, **due 4/29**, 11:59 pm.

### 6. April 30: Invariant Object Recognition and Selective Attention

Dorsal and ventral visual streams. Invariant object recognition: the problem. Structural vs. image-based approaches. The binding problem. A simple model of

object recognition (Sect. 8.4). A simple model of spatial attention (Sect. 8.5). An integrated model (Sect. 8.6).

Readings:

- O&M Chapter 8.
- [Optional] Riesenhuber, M. & Poggio, T. (1999). Hierarchical Models of Object Recognition in Cortex. *Nature Neuroscience*, 2, 1019-1025. (PDF available on Carmen.)
- [Optional] Mozer, M. C. & Sitton, M. (1998). Computational Modeling of Spatial Attention. In H. Pashler (Ed.), *Attention* (pp. 341-393). Philadelphia, PA: Psychology Press. (Contact Dr. Petrov for a hard copy.)
- [Optional] Hummel, J. E. (2001). Complementary solutions to the binding problem in vision: Implications for shape perception and object recognition. *Visual Cognition*, 8, 489-517. (Special issue on binding. PDF available on Carmen.)

Deliverables: RRs for Chapter 8, **due 4/29**, 11:59 pm.

Model explorations (5%): Exercises 8.4, 8.6, 8.10, and 8.11. **Due 5/06**, 11:59 pm.

### **7. May 07: English Past Tense. Triangular Model**

Dyslexia. Triangular model (Sect. 10.3). Orthography-to-phonology mapping (Sect. 10.4). Overregularization in English past-tense inflections (Sect. 10.5).

Readings:

- O&M Chapter 10, up to Section 10.5 inclusive.
- [Optional] The original publications cited throughout the chapter (e.g., Plaut & Shallice, 1993; Seidenberg & McClelland, 1989; Rumelhart & McClelland, 1986) and/or their various follow-ups.

Deliverables: RRs for Chapter 10, **due 5/06**, 11:59 pm.

One-page proposal for final paper, **due 5/06**, 11:59 pm.

Model explorations (5%): Exercises 10.2a, 10.4a, and 10.7, **due 5/14**, 11:59 pm.

### **8. May 14: Complementary Learning Systems**

Memory is not unitary. Weight-based vs. activation-based memories. AB-AC learning. Catastrophic interference. Role of the medial temporal lobe. Amnesia. Functional tradeoffs and the need for complementary learning systems.

Readings:

- O&M Chapter 9.
- [Optional] McClelland, J. L., McNaughton, B. L., & O'Reilly, R. (1995). Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, 102, 419-457. <http://psych.colorado.edu/~oreilly/pubs-abstr.html>  
[#McClellandMcNaughtonOReilly95](#)

Deliverables: RRs for Chapter 9, **due 5/13**, 11:59 pm.

Model explorations (5%): Exercises 9.2, 9.4 (competition for extra credit!), and 9.7. **Due 5/20**, 11:59 pm.



## 9. May 21: Tripartite Cognitive Architecture

Cognitive architecture: A big idea of science. Tripartite functional organization of the brain. Prefrontal cortex. Working memory (Sect. 9.5). Stroop model (Sect. 11.3). Basal ganglia: reinforcement learning, gating.

Readings:

- O&M Chapter 7 and Sections 9.5, 11.1, 11.2, 11.3, and 11.7.
- [Optional] O'Reilly, R.C., Braver, T. S. & Cohen, J. D. (1999). A Biologically Based Computational Model of Working Memory. In A. Miyake & P. Shah (Eds.) *Models of Working Memory: Mechanisms of Active Maintenance and Executive Control* (pp. 375-411). New York: Cambridge University Press. <http://psych.colorado.edu/~oreilly/pubs-abstr.html#OReillyBraverCohen99>
- [Optional] Anderson, J. R. & Lebiere, C. L. (2003). The Newell test for a theory of cognition. *Behavioral & Brain Sciences*, 26, 587-637. <http://act-r.psy.cmu.edu/publications/pubinfo.php?id=507>

Deliverables: RRs for Chapter 7, **due 5/20**, 11:59 pm.

Model explorations (5%): Exercises 9.9, 9.11, 11.2, and 11.3. **Due 5/27**, 11:59 pm

## 10. TBA (in lieu of 5/28, Memorial Day): Advanced Topics. Conclusions

Higher-level cognition in neural networks. Final discussion.

Readings:

- O&M Chapter 12 and Sections 11.5, 11.6, and 11.7.
- O'Reilly, R. C. (2006). Biologically Based Computational Models of High-Level Cognition. *Science*, 314, 91-94. <http://psych.colorado.edu/~oreilly/pubs-abstr.html#OReilly06>
- [Optional] Rougier, N. P., Noelle, D., Braver, T. S., Cohen, J. D. & O'Reilly, R. C. (2005). Prefrontal Cortex and the Flexibility of Cognitive Control: Rules Without Symbols. *Proceedings of the National Academy of Sciences*, 102, 7338-7343. <http://psych.colorado.edu/~oreilly/pubs-abstr.html#RougierEtAl05>

Deliverables: RRs for O'Reilly (2006), **due 5/27**, 11:59 pm.

## F. June 04: Final Paper Due

The above calendar is subject to change at the discretion of the instructor, depending on the rate of progress through the material, student interest in alternative topics, and/or scheduling constraints. The model exploration assignments are also subject to change; you will be notified at least five days in advance for any changes in the assignments.

Finally, welcome to the course. I hope that you will enjoy the class and learn valuable information and skills. I look forward to seeing you on March 26.

Alex Petrov

syllabus878-sp07.doc, last updated 24 March 2007