COGS 223, EECS 273
Computational Cognitive Neuroscience

University of California, Merced

Spring, 2018

Instructor

David C. Noelle, Ph.D.
dnoelle@ucmerced.edu
262A Social Sciences & Management Building
(209) 228-4101
Office Hours: Tuesdays, 4:30 P.M. – 5:30 P.M., and by appointment

Teaching Assistants

Jacob Rafati & Tim Shea
jrafatiheravi@ucmerced.edu & tshea2@ucmerced.edu
220 Social Sciences & Management Building
Office Hours: by appointment

Class Meetings

Tuesdays & Thursdays, 3:00 P.M. – 4:15 P.M., in 120 Student Services Building
Thursdays, 4:30 P.M. – 7:15 P.M., in 154 Social Sciences & Management Building
Thursdays, 7:30 P.M. – 10:15 P.M., in 154 Social Sciences & Management Building

Overview

Cognitive neuroscience is concerned with the question of how the mass of interconnected neural cells that make up the brain can give rise to complex mental processes and intelligent behavior. Computer simulations of cognitive processes, informed and constrained by our knowledge of brain anatomy and physiology, can play a central role in this research endeavor by testing the feasibility of theories of brain function, uncovering hidden implications of such theories, and demonstrating how relatively simple biological mechanisms can work together to form powerful information processing systems which exhibit human-like patterns of performance. The generation, analysis, and testing of such computer simulations is the focus of *computational cognitive neuroscience*.

This course will introduce a collection of computer simulation techniques useful for investigating a variety of cognitive phenomena involving perception, action, learning, and memory. This introduction will include discussions of how such methods can leverage our growing neuroscientific knowledge.

This course is formatted to support interdisciplinary inquiry, with the backgrounds of students expected to vary broadly across the range of such disciplines as computer science, cognitive science, psychology, and neuroscience, as well as other related fields. The learning of both classic and contemporary methods for cognitive modeling will be facilitated by readings, presentations by the teaching team, computer simulation exercises, quizzes, a term project, and ample interaction and discussion between attendees.

Learning Outcomes

Students who successfully complete this course will have acquired an introductory understanding of the basic principles of *computational cognitive neuroscience (CCN)* modeling (including extensions to *parallel distributed processing (PDP)*, *connectionist*, and *artificial neural network* models). By the end of the term, students will be able to relate the formal properties of such models to known biological mechanisms as well as to behavioral phenomena, and they will possess intellectual tools for modifying such models in the light of new psychological and neuroscientific findings. The knowledge acquired by the successful student will encompass a variety of broad topics, including: the mathematical and computational properties of CCN models, a survey of common modeling techniques, methods for extending traditional PDP techniques to incorporate relevant biological details, the use of CCN simulation software, methods for studying cognition through computational modeling and analysis, CCN approaches to perception and motor control, CCN models of language use and language learning, neural and cognitive development from a CCN perspective, and the role of CCN models in exploring the impact of brain damage and neurodegenerative disease.

These *student learning outcomes* are expected to be valuable to students pursuing a variety of graduate education programs. With regard to the *Cognitive and Information Sciences (CIS)* Ph.D. program, these *student learning outcomes* contribute to a number of the CIS *program learning outcomes*, including increasing student understanding of foundational CIS concepts, development of skills in foundational CIS methods, practice of scientific communication skills, and exposure to the interdisciplinary nature of CIS. This course satisfies a CIS Ph.D. program graduate course requirement, and it provides training relevant to the required *integrative review papers* and the program's *qualifying examination*, as well as student research projects. Students enrolled in a variety of graduate education programs should find this course useful for meeting their program's objectives.

Resources

Meetings

Course participants will meet every Tuesday and Thursday of this semester according to the schedule outlined at the end of this document. Lecture and discussion meetings will take place from 3:00 P.M. to 4:15 P.M. in Room 120 of the Student Services Building. These class meetings will consist of reviews of material available in course readings, introductions to new material, timed quizzes, and discussions of related topics. The intended structure of these meetings includes substantial flexibility, allowing them to adapt to the interests and concerns of the participants.

In addition to these lecture and discussion meetings, course participants are expected to attend one laboratory session each week of the semester. Two such sessions are scheduled on a weekly basis, with both being held in Room 154 of the Social Sciences & Management Building. Laboratory sections are scheduled for 4:30 P.M. to 7:15 P.M. and for 7:30 P.M. to 10:15 P.M. on Thursdays. These laboratory sessions will allow students to execute and explore computer simulations of human behavior and brain function in the company of a teaching assistant and fellow students. Computer simulation exercises will be conducted by individual participants during these laboratory meetings, and these exercises act as a central component of the learning process.

Aggregated student responses to questions will be collected during class meetings using the *Top Hat* classroom response system. This system allows the intructor to present a question to the class, and it allows students to respond to the question using computers, mobile devices (iOS® or Android®), and SMS text messaging enabled phones. This system will be used to check on the understanding of readings and lecture material, and it will be used to monitor class attendance. In most cases, the correctness of student responses will *not* be scored, but student participation will be recorded. About once per week, this same system will be used to conduct a time-limited quiz concerning laboratory exercises, asssigned readings, and lecture content. Responses to these weekly quizzes will play a role in the evaluation of student understanding of the course content.

The *Top Hat* classroom response system is an online subscription service to which each student must suscribe. Such a subscription costs about $26, but there are more expensive subscription plans that allow for the use of *Top Hat* in other classes during later semesters. (Note that the "secure test" option will not be used in this class.) Each enrolled student is expected to receive an electronic mail message directly from *Top Hat*, explaining subscription options. The *Student Quick Start Guide* may be found at:

`http://tinyurl.com/THStudentRegistration`

If a subscription invitation electronic mail message is not received, the instructor should be promptly informed of this fact. Also, information about registering specifically for this class may be found at:

`http://app.tophat.com/e/659015`

Subscribing to *Top Hat* and registering for this course within that system should be completed as soon as possible. Further information about this classroom response system may be found at "`www.tophat.com`".

Web Site

Online materials for this class will be disseminated through the `CatCourses` learning management system (LMS). This system may be accessed through its central portal:

`https://catcourses.ucmerced.edu/`

Enrolled students should be provided with access to a section of `CatCourses` specifically reserved for this class. In particular, a shared section called "`S18-COGS 123 01/COGS 223 01/CSE 173 01/EECS 273 01`" should be available to all enrolled students. The instructor should be promptly informed if such access is not appropriately granted.

This web site will be used to announce updates to the class schedule, as well as to distribute class materials. It will be used for some quizzes, and exercise submissions will be made through this site. Students are required to obtain regular access to this resource, and they are strongly advised to consult it frequently (e.g., daily).

Readings

The primary source of expository readings for this class will be:

CCN : O'Reilly, R. C., Munakata, Y., Frank, M. J., Hazy, T. E., and Contributors (2016) *Computational Cognitive Neuroscience.* Wiki Book, 3rd (partial) Edition. URL: `http://ccnbook.colorado.edu`

In addition to providing explanatory text, this book will also be used to guide the sequencing of topics discussed in this class, and computer simulation exercises will be drawn almost exclusively from its pages.

This online textbook is an updated and reduced version of the following previously published book:

CECN : O'Reilly, R. C. & Munakata, Y. (2000) *Computational Explorations in Cognitive Neuroscience: Understanding the Mind by Simulating the Brain.* MIT Press: Cambridge, MA.

While somewhat older, this second textbook contains more extensive material on some topics. Electronic copies of the chapters in this book are available from the online *CogNet* service. While access to *CogNet* typically requires a subscription to the service, it is available to you, free of charge, through the university library web site (`http://library.ucmerced.edu/`).

While these two books provide rather comprehensive coverage of topics relevant to this course, they do not always offer thorough explanations of alternative or competing approaches to the modeling enterprise. Participants interested in obtaining a slightly wider view of the range of research methods employed in this field are encouraged to supplement their reading with other sources, including the following recommendations.

There is a collection of three seminal tomes in the world of PDP modeling, produced by a team of researchers at the University of California, San Diego in the early 1980s. Even after about 30 years, the PDP volumes remain an excellent introduction to the use of connectionist techniques to understand psychological phenomena:

PDP : Rumelhart, D. E., McClelland, J. L., & the PDP Research Group (1986) *Parallel Distributed Processing: Explorations in the Microstructure of Cognition. Volume 1: Foundations.* MIT Press: Cambridge, MA.

McClelland, J. L., Rumelhart, D. E., & the PDP Research Group (1986) *Parallel Distributed Processing: Explorations in the Microstructure of Cognition. Volume 2: Psychological and Biological Models.* MIT Press: Cambridge, MA.

Handbook : McClelland, J. L., Rumelhart, D. E. (1988) *Explorations in Parallel Distributed Processing: A Handbook of Models, Programs, and Exercises.* MIT Press: Cambridge, MA.

These books are still in print at the MIT Press (with the exception of the third volume, which contains exercises). Also, electronic copies of the chapters of these volumes are offered by *CogNet*.

A great collection of classic papers in the field of artificial neural network modeling may be found in:

Neurocomputing : Anderson, J. A. & Rosenfeld, E. (Eds.) (1988) *Neurocomputing: Foundations of Research*. MIT Press: Cambridge, MA.

This is also an MIT Press book, and an electronic version is also provided by the *CogNet* service.

An excellent, more recently published, text on computational neuroscience, focusing a bit more on neurobiological phenomena than behavioral phenomena, is:

Theoretical : Dayan, P. & Abbott, L. F. (2001) *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*. MIT Press: Cambridge, MA.

This book also appears in electronic format on the *CogNet* service.

Software

Reading about computational modeling mechanisms is rarely sufficient to develop a deep understanding of them. In hopes of fostering a deeper understanding, class participants will conduct small simulation experiments throughout the semester using the `emergent` software package. This software system provides tools for the design, execution, and analysis of CCN models utilizing a wide variety of computational mechanisms and methods. The guiding text of this course provides extensive examples and exercises using `emergent`, and versions of these exercises will act as critical tools for communicating key concepts.

The `emergent` system is an open source software package available free of charge from the following World Wide Web site:

`http://grey.colorado.edu/emergent/`

While `emergent` is fully extendible through the incorporation of user written `C++` code, no knowledge of computer programming is needed in order to make basic use of the system. A graphical user interface, involving pull-down menus, buttons, and the like, is provided, and this is the standard manner in which modelers interact with the `emergent` tools.

The `emergent` programs were originally developed for computers running Unix®-like operating systems. The package has been ported, however, to Microsoft® Windows®. There has also been a port to Mac® OS, and much of the recent development on `emergent` has been performed on the Mac® platform.

Computer simulation exercises using `emergent` draw extensively from the main text to be used in this class. While all of these exercises offer opportunities for learning, the instructor will identify a subset as particularly relevant. These exercises will help the student investigate general computational modeling mechanisms using this particular package of simulation tools. The main text offers much guidance for conducting these simulations, as does online documentation for the exercises, but there are additional sources of information on the fruitful use of the `emergent` software. Primary among these is the online documentation available from the `emergent` World Wide Web site. Further information on obtaining, installing, and using `emergent` may be found on the the official `emergent` web site.

Laboratory

The `emergent` software system has been installed on the computers that will be available during class laboratory sessions. During early meetings in this laboratory, students will be guided in the use of `emergent` in this operating environment.

While scheduled laboratory sessions are expected to be of sufficient length to allow course participants to complete the regular exercises that will be assigned during the semester, additional access to `emergent` may be needed for students to complete their term projects and to conduct other self-guided investigations into computational cognitive neuroscience. Thus, students are encouraged to install `emergent` on computers located in their usual work spaces, either in their offices or in their homes, allowing them to explore CCN modeling methods in the comfort of their standard work environments.

Expectations & Evaluations

Background

Course participants are expected to have some background knowledge in cognitive science, computer science, psychology, or neuroscience and a basic understanding of the most fundamental concepts of differential calculus, linear algebra, and statistics. Computer programming skills may be useful to some students as they conduct their term projects, but such skills are certainly *not* required. As previously mentioned, the `emergent` software system which will be used in this course allows for the design and control of simulations through a rich "point and click" interface. Students are only expected to be sufficiently familiar with an appropriate computing environment so as to be able to perform such tasks as editing and managing files and manipulating objects in a window-based graphical environment.

The University of California, Merced, is committed to ensuring equal educational opportunities for students with disabilities. An integral part of this commitment is the coordination of specialized academic support services through the Disability Services office. Students with a physical or learning disability may ask Disability Services to assist in communicating this fact to the instructor so that appropriate accommodation may be provided. Absent notification, the instructor may assume that no such accommodation is sought.

Participation

Studying the neural basis of cognition using computational models is a challenging interdisciplinary endeavor requiring familiarity with notions from artificial intelligence, statistics, psychology, and neuroscience. Thus, students may find this class both rewarding and demanding. Mastery of the course materials will require extensive reading, puzzling through unfamiliar concepts, active participation in classroom discussions, learning new mathematical formalisms, many hours of hands-on experience building and analyzing computer simulations, and a willingness to view human cognition in new ways.

Specific readings will be suggested as appropriate for each class meeting, and participants will be expected to have studied those readings *prior* to gathering, so as to promote thoughtful questions and knowledgeable discussion. Students will be expected to contribute constructively to discussions, bringing to bear both insights into the material at hand and relevant knowledge acquired in other contexts.

This class is structured so as to leverage the variety and wealth of knowledge that the participants will bring to this learning enterprise. Interaction during class meetings should be seen as an important goal. Participants are strongly encouraged to investigate connections between course topics and their own fields of study and to bring insights from such explorations with them to class meetings. Indeed, such contributions will be expected.

Computational methods for cognitive neuroscience are best learned through active experimentation. Thus, a number of exercises will be assigned to help students acquire an understanding of CCN concepts and techniques and to aid in the evaluation of their understanding. These exercises will require substantial time and effort to complete, and they will involve the use of the `emergent` simulation software. Student solutions to these assignments will be evaluated by the teaching team, and appropriate feedback will be given.

Approximately once per week, a small amount of time at the beginning of a class meeting will be devoted to a time-restricted quiz, designed to test student understanding of recently completed simulation exercises, as well as comprehension and memory of all material covered in assigned readings and class meetings up to the time of the quiz.

Near the end of the term, students are expected to complete original term projects. Each student will be expected to prepare a short research report (about 10–15 pages) describing the student's use of computational modeling techniques to advance our understanding of some phenomenon of the cognitive neurosciences. This final paper should provide each student with an opportunity to extend the course material in a direction of personal interest and should demonstrate an understanding of key CCN concepts in some interesting way. Minimally, term projects should extend a computational model presented during the semester or analyze a presented model in a new way. Note, however, that term projects are generally expected to contain some modicum of genuine research value, providing original insights or results through the generation and analysis of novel computer simulations. In addition to producing a written report describing the work conducted as part of a term project, each student will also be required to share their findings with the other course participants through brief presentations to be held during the course's final meetings.

Class participants are expected to embrace the course material with earnest effort, to contribute constructively to the learning of other students, and to always behave ethically and with civic concern. Students should come to every class meeting prepared to discuss relevant topics. Exercises are to be completed by their respective due dates. The ideas and contributions of others should be apprpriately cited. (This includes ideas and contributions garnered from readings, online resources, presentations, conversations, and any other source. This specifically includes help from other students and from members of the teaching team.) Students are expected to bring educational obstacles to the instructor's attention as early as possible, so that such problems may be promptly resolved.

Learning can be greatly facilitated by interactions between class participants, and these interactions are encouraged. Students should feel free to discuss lecture topics, readings, project ideas, and even exercise assignments with each other. The actual completion of exercises, however, should be conducted on an individual basis, as should responses to quiz questions, and all work on term projects should be completed by the individual authors of those projects. All assignments submitted for evaluation should reflect the understanding and effort of the individual participant. If there is ever any doubt concerning the propriety of a given interaction, it is the student's responsibility to approach the instructor and clarify the situation *prior* to the submission of work results. Also, helpful conversations with fellow students, or any other person (including members of the teaching team), should be explicitly mentioned in submitted assignments. Failure to appropriately cite sources is one form of *plagiarism*, and it will not be tolerated!

Evaluation

The teaching team will provide comments on assigned work submitted in a timely manner for evaluation. Those students who are to receive grades for this course will have their work assessed roughly as follows:

Exercises #1 – #11	2%	each
Weekly Cumulative Quizzes	3%	each
Project Oral Report	5%	
Project Written Report	25%	
Class Attendance & Participation	15%	

Student performance will be evaluated in comparison to that of other students, both past and present. Class participation, including responding to questions and actively facilitating class discussions, will be closely monitored throughout the term.

Exercises will be assigned each Thursday and will be due one week from that Thursday, at 3:00 P.M. on that day. Late assignments which arrive in the instructor's hands before 3:00 P.M. on the day after a due date (Friday) will be evaluated and will receive 90% of the credit for the assignment. Late assignments which arrive in the instructor's hands before 3:00 P.M. on the subsequent day (Saturday) will be evaluated and will receive 80% of the credit for the assignment. Assignments which are submitted later than this will not be evaluated, and no credit will be given.

Final written project reports will be due to the instructor by 9:30 P.M. on Tuesday, May 8th. Late written project reports will *not* be accepted.

Schedule

In the schedule that appears on the following pages, a check mark ($\surd$) identifies an assigned reading for the date in question. Other readings may be considered supplementary and optional. Those marked with a diamond ($\diamondsuit$) may be found in the *CECN* textbook, in the *PDP* volumes, in *Neurocomputing*, in *Theoretical Neuroscience*, or in online `emergent` documentation. Other papers, marked with a heart ($\heartsuit$) do not appear in any of these sources and will not be provided by the instructor.

Introduction

January 16 : Introduction

- *Syllabus Distributed*

January 18 : Historical Background

- *Initial Use of Top Hat Classroom Response System*

√ Syllabus

√ *CCN*, Chapter 1.

√ "`Documentation: Getting Started`", *Emergent Wiki*.

♢ *CECN*, Forward & Chapter 1.

♢ *CECN*, Appendix A.

♢ McClelland, J. L., Rumelhart, D. E., & Hinton, G. E. (1986) The appeal of parallel distributed processing. *PDP*, Chapter 1.

♢ McCulloch, W. S. & Pitts, W. (1943) A logical calculus of the ideas immanent in nervous activity. *Neurocomputing*, Chapter 2.

♢ von Neumann, J. (1958) The computer and the brain. *Neurocomputing*, Chapter 7.

Neural Activity

January 23 : Individual Neurons — The Membrane Potential

√ *CCN*, § 2.1–2.4.1.

♢ *CECN*, Chapter 2.

♢ *Theoretical*, § 5.1–5.5.

♢ Rumelhart, D. E., McClelland, J. L., & Hinton, G. E. (1986) A general framework for parallel distributed processing. *PDP*, Chapter 2.

♢ McClelland, J. L. & Rumelhart, D. E. (1988) Introduction. *Handbook*, Chapter 1.

♡ Siegelbaum, S. A. & Koester, J. (1991) Ion channels. Chapter 5 of E. R. Kandel, J. H. Schwartz, & T. M. Jessell (Eds.), *Principles of Neural Science: Third Edition*. Appleton & Lange: Norwalk, Connecticut.

♡ Koester, J. (1991) Membrane potential. Chapter 6 of E. R. Kandel, J. H. Schwartz, & T. M. Jessell (Eds.), *Principles of Neural Science: Third Edition*. Appleton & Lange: Norwalk, Connecticut.

♡ Koester, J. (1991) Passive membrane properties of the neuron. Chapter 7 of E. R. Kandel, J. H. Schwartz, & T. M. Jessell (Eds.), *Principles of Neural Science: Third Edition*. Appleton & Lange: Norwalk, Connecticut.

January 25 : Individual Neurons — The Action Potential

- *Exercise #1 Specification Distributed*

√ *CCN*, § 2.4.2–2.4.3, 2.5–2.6 (§ 2.7 is optional).

♢ *Theoretical*, § 7.1–7.4.

♡ Koester, J. (1991) Voltage-gated ion channels and the generation of the action potential. Chapter 8 of E. R. Kandel, J. H. Schwartz, & T. M. Jessell (Eds.), *Principles of Neural Science: Third Edition*. Appleton & Lange: Norwalk, Connecticut.

January 30 : Distributed Representations

√ *CCN*, § 3.1–3.2.

♢ *CECN*, § 3.1–3.4.

♢ Hinton, G. E., McClelland, J. L., & Rumelhart, D. E. (1986) Distributed representations. *PDP*, Chapter 3.

February 01 : Inhibition & Attractor Dynamics

- *Quiz #1 At 3:00 P.M.*
- *Exercise #1 Due At 3:00 P.M.*
- *Exercise #2 Specification Distributed*

√ *CCN*, § 3.3–3.5.

♢ *CECN*, § 3.5–3.8.

♢ *Theoretical*, § 7.5.

♢ Smolensky, P. (1986) Information processing in dynamical systems: Foundations of harmony theory. *PDP*, Chapter 6.

♢ McClelland, J. L. & Rumelhart, D. E. (1988) Interactive activation and competition. *Handbook*, Chapter 2, pp. 11–17.

♢ McClelland, J. L. & Rumelhart, D. E. (1981) An interactive activation model of context effects in letter perception: Part 1. An account of basic findings. *Neurocomputing*, Chapter 25.

♢ Rumelhart, D. E., Smolensky, P., McClelland, J. L., & Hinton, G. E. (1986) Schemata and sequential thought processes in PDP models. *PDP*, Chapter 14.

♢ McClelland, J. L. & Rumelhart, D. E. (1988) Constraint satisfaction in PDP systems. *Handbook*, Chapter 3, pp. 49–54, 68–73, 75–81.

♢ McClelland, J. L. & Elman, J. L. (1986) Interactive processes in speech perception: The TRACE model. *PDP*, Chapter 15.

♢ Selfridge, O. G. (1958) Pandemonium: A paradigm for learning. *Neurocomputing*, Chapter 9.

♢ Hopfield, J. J. (1982) Neural networks and physical systems with emergent collective computational abilities. *Neurocomputing*, Chapter 27.

♢ Geman, S. & Geman, D. (1984) Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *Neurocomputing*, Chapter 37.

February 06 : Mathematics Review

- ♢ *Theoretical*, § A.1, A.3, & A.5.
- ♢ *Theoretical*, § 8.1–8.2.
- ♢ Jordan, M. I. (1986) An introduction to linear algebra in parallel distributed processing. *PDP*, Chapter 9.

Synaptic Plasticity

February 08 : Unsupervised Learning & Self-Organization

- *Quiz #2 At 3:00 P.M.*
- *Exercise #2 Due At 3:00 P.M.*
- *Exercise #3 Specification Distributed*
- √ *CCN*, § 4.1–4.2.2.
- √ *CCN*, § 4.5 (Hebbian Learning).
- √ *CCN*, § 4.5 (STDP).
- ♢ *CECN*, Chapter 4.
- ♢ *Theoretical*, § 10.1–10.2.
- ♢ Rumelhart, D. E. & Zipser, D. (1986) Feature discovery by competitive learning. *PDP*, Chapter 5.
- ♢ McClelland, J. L. & Rumelhart, D. E. (1988) Learning in PDP models: The pattern associator. *Handbook*, Chapter 4, pp. 83–89.
- ♢ Hebb, D. O. (1949) The first stage of perception: growth of the assembly. *Neurocomputing*, chapter 4.
- ♢ Bienenstock, E. L., Cooper, L. N., & Munro, P. W. (1982) Theory for the development of neuron selectivity: Orientation specificity and binocular interaction in visual cortex. *Neurocomputing*, Chapter 26.
- ♢ Munro, P. W. (1986) State-dependent factors influencing neural plasticity: A partial account of the critical period. *PDP*, Chapter 24.
- ♢ Kohonen, T. (1982) Self-organized formation of topologically correct feature maps. *Neurocomputing*, Chapter 30.
- ♢ Hinton, G. E. & Sejnowski, T. J. (1986) Learning and relearning in Boltzmann machines. *PDP*, Chapter 7.
- ♢ Ackley, D. H., Hinton, G. E., & Sejnowski, T. J. (1985) A learning algorithm for Boltzmann machines. *Neurocomputing*, Chapter 38.
- ♢ Anderson, J. A., Silverstein, J. W., Ritz, S. A., & Jones, R. S. (1977) Distinctive features, categorical perception, and probability learning: Some applications of a neural model. *Neurocomputing*, Chapter 22.
- ♢ Kirkpatrick, S., Gelatt, Jr., C. D., Vecchi, M. P. (1983) Optimization by simulated annealing. *Neurocomputing*, Chapter 33.

February 13 : The Delta Rule

- √ *CCN*, § 4.2.3–4.2.3.1.
- ♢ *CECN*, § 5.1–5.5.
- ♢ Rosenblatt, F. (1958) The perceptron: A probabilistic model for information storage and organization in the brain. *Neurocomputing*, Chapter 8.
- ♢ Widrow, B. & Hoff, M. E. (1960) Adaptive switching circuits. *Neurocomputing*, Chapter 10.
- ♢ Minsky, M. & Papert, S. (1969) Perceptrons. *Neurocomputing*, Chapter 13.
- ♢ McClelland, J. L. & Rumelhart, D. E. (1988) Learning in PDP models: The pattern associator. *Handbook*, Chapter 4, pp. 89–99.
- ♢ Stone, G. O. (1986) An analysis of the delta rule and the learning of statistical associations. *PDP*, Chapter 11.
- ♡ Gluck, M. A. & Bower, G. H. (1988) From conditioning to category learning: An adaptive network model. *Journal of Experimental Psychology: General*, 117, pp. 227–247.
- ♡ Sutton, R. S. & Barto, A. G. (1981) Toward a modern theory of adaptive networks: Expectation and prediction. *Psychological Review*, 88, pp. 135–170.

February 15 : The Generalized Delta Rule

- *Quiz #3 At 3:00 P.M.*
- *Exercise #3 Due At 3:00 P.M.*
- *Exercise #4 Specification Distributed*
- √ *CCN*, § 4.5 (Backpropagation).
- ♢ *CECN*, § 5.6.
- ♢ Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986) Learning internal representations by error propagation. *PDP*, Chapter 8.
- ♢ *Theoretical*, § 8.4.
- ♡ Hinton, G. E. (1989) Connectionist learning procedures. *Artificial Intelligence*, 40, pp. 185–234.
- ♡ Rumelhart, D. E., Durbin, R., Golden, R., & Chauvin, Y. (1995) Backpropagation: The basic theory. Chapter 1 in Y. Chauvin & D. E. Rumelhart (Eds.), *Backpropagation: Theory, Architectures, and Applications*. Lawrence Erlbaum: Hillsdale, NJ.
- ♡ Rumelhart, D. E., Durbin, R., Golden, R., & Chauvin, Y. (1996) Backpropagation: The basic theory. Chapter 15 in P. Smolensky, M. C. Mozer, & D. E. Rumelhart (Eds.), *Mathematical Perspectives on Neural Networks*. Lawrence Erlbaum: Hillsdale, NJ.

February 20 : Special Network Architectures

- √ Elman, J. L. (1990) Finding structure in time. *Cognitive Science*, 14, pp. 179–211. (available from instructor)
- ◇ *CECN*, § 6.5–6.6.
- ◇ McClelland, J. L. & Rumelhart, D. E. (1988) Training hidden units: The generalized delta rule. *Handbook*, Chapter 5, pp. 121–137.
- ◇ Williams, R. J. (1986) The logic of activation functions. *PDP*, Chapter 10.
- ◇ McClelland, J. L. & Rumelhart, D. E. (1988) Training hidden units: The generalized delta rule. *Handbook*, Chapter 5, pp. 155–158.
- ♡ Morgan, N. & Bourlard, H. (1990) Generalization and parameter estimation in feedforward nets: Some experiments. In D. S. Touretzky (Ed.), *Advances in Neural Information Processing Systems 2*, pp. 630–637. Morgan Kaufmann: San Mateo, CA.
- ♡ le Cun, Y., Denker, J. S., & Solla, S. A. (1990) Optimal brain damage. In D. S. Touretzky (Ed.), *Advances in Neural Information Processing Systems 2*, pp. 598–605. Morgan Kaufmann: San Mateo, CA.
- ♡ Weigand, A. S., Rumelhart, D. E., & Huberman, B. A. (1991) Generalization by weight-elimination with application to forcasting. In R. P. Lippmann, J. E. Moody, & D. S. Touretzky (Eds.), *Advances in Neural Information Processing Systems 3*, pp. 875–882. Morgan Kaufmann: San Mateo, CA.
- ♡ Cottrell, G. W., Munro, P., & Zipser, D. (1988) Image compression by back propagation: An example of extensional programming. In N. E. Sharkey (Ed.), *Advances in Cognitive Science*, Volume 3. Ablex: Norwood, NJ.
- ♡ Pollack, J. B. (1990) Recursive distributed representations. *Artificial Intelligence*, 46, pp. 77–105.
- ♡ Jordan, M. I. (1986) Attractor dynamics and parallelism in a connectionist sequential machine. In *Proceedings of the 8th Annual Conference of the Cognitive Science Society*, pp. 531–546. Lawrence Erlbaum: Hillsdale, NJ.
- ♡ Williams, R. J. & Zipser, D. (1995) Gradient-based learning algorithms for recurrent networks and their computational complexity. Chapter 13 in Y. Chauvin & D. E. Rumelhart (Eds.), *Backpropagation: Theory, Architectures, and Applications*. Lawrence Erlbaum: Hillsdale, NJ.
- ♡ Pearlmutter, B. (1989) Learning state space trajectories in recurrent neural networks. *Neural Computation*, 1, pp. 263–269.

February 22 : Leabra Learning

- *Quiz #4 At 3:00 P.M.*
- *Exercise #4 Due At 3:00 P.M.*
- *Exercise #5 Specification Distributed*
- √ *CCN*, § 4.2.3.2–4.4.
- ◇ *CECN*, § 5.7–5.11.
- ◇ *CECN*, § 6.1–6.4.

Large-Scale Functional Brain Organization

February 27 : Cognitive Architecture

- √ *CCN*, Chapter 5.
- ♢ *CECN*, Chapter 7.

March 01 : Designing Simulation Projects

- *Quiz #5 At 3:00 P.M.*
- *Exercise #5 Due At 3:00 P.M.*
- *Exercise #6 Specification Distributed*
- ♢ *CECN*, Appendix B.

Perception & Action

March 06 : Visual Perception

- √ *CCN*, § 6.1–6.3.
- ♢ *CECN*, § 8.1–8.4.
- ♢ von der Malsburg, C. (1973) Self-organization of orientation sensitive cells in the striate cortex. *Neurocomputing*, Chapter 17.
- ♢ Marr, D. & Poggio, T. (1976) Cooperative computation of stereo disparity. *Neurocomputing*, Chapter 20.
- ♢ Marr, D. (1982) Vision. *Neurocomputing*, Chapter 28.
- ♢ Fukushima, K., Miyake, S., & Ito, T. (1983) Neocognitron: A neural network model for a mechanism of visual pattern recognition. *Neurocomputing*, Chapter 31.

March 08 : Visual Attention

- *Quiz #6 At 3:00 P.M.*
- *Exercise #6 Due At 3:00 P.M.*
- *Exercise #7 Specification Distributed*
- √ *CCN*, § 6.4.
- ♢ *CECN*, § 8.5–8.8.
- ♢ Crick, F. (1984) Function of the thalamic reticular complex: the searchlight hypothesis. *Neurocomputing*, Chapter 34.
- ♡ Mozer, M. C. & Sitton, M. (1998) Computational modeling of spatial attention. In H. Pashler (Ed.), *Attention*, pp. 341–393. UCL Press: London.

March 13 : Reinforcement Learning

- √ *CCN*, § 7.1–7.2.
- ♢ *CECN*, § 6.7–6.9.
- ♢ *Theoretical*, Chapter 9.
- ♢ Barto, A. G., Sutton, R. S., & Anderson, C. W. (1983) Neuronlike adaptive elements that can solve difficult learning control problems. *Neurocomputing*, Chapter 32.
- ♡ Sutton, R. S. & Barto, A. G. (1998) *Reinforcement Learning: An Introduction*. MIT Press: Cambridge, MA.
- ♡ Sutton, R. S. (1988) Learning to predict by the method of temporal differences. *Machine Learning*, 3, pp. 9–44.
- ♡ Barto, A. G. (1994) Adaptive critics and the basal ganglia. Chapter 11 in J. C. Houk, J. L. Davis, & D. G. Beiser (Eds.), *Models of Information Processing in the Basal Ganglia*. MIT Press: Cambridge, MA.

March 15 : Motor Control

- *Quiz #7 At 3:00 P.M.*
- *Exercise #7 Due At 3:00 P.M.*
- *Exercise #8 Specification Distributed*
- √ *CCN*, § 7.3.
- ♢ Jordan, M. I. & Rumelhart, D. E. (1992) Forward models: Supervised learning with a distal teacher. *Cognitive Science*, 16, pp. 307–354.
- ♢ Wolpert, D. M., Ghahramani, Z., & Jordan, M. I. (1995) An internal model for sensorimotor integration. *Science*, 269, pp. 1880–1882.
- ♡ Kawato, M. (1999) Internal models for motor control and trajectory planning. *Current Opinion in Neurobiology*, 9, pp. 718–727.
- ♡ Ohyama, T., Medina, J. F., Nores, W. L., & Mauk, M. D. (2002) Trying to understand the cerebellum well enough to build one. *Annals of the New York Academy of Sciences*, 978, pp. 425–438.

Memory

March 20 : Episodic Memory

- √ *CCN*, § 8.1.
- ♢ *CECN*, § 9.3.
- ♢ McClelland, J. L. & Rumelhart, D. E. (1986) Amnesia and distributed memory. *PDP*, Chapter 25.
- ♡ McCloskey, M. & Cohen, N. J. (1989) Catastrophic interference in connectionist networks: The sequential learning problem. *The Psychology of Learning and Motivation*, 24, pp. 109–165.
- ♡ McClelland, J. L., McNaughton, B. L., and O'Reilly, R. C. (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, pp. 419–457.
- ♡ Norman, K. A. & O'Reilly, R. C. (2003) Modeling hippocampal and neocortical contributions to recognition memory: A complementary-learning-systems approach. *Psychological Review*, 110, pp. 611–646.

March 22 : Priming

- *Quiz #8 At 3:00 P.M.*
- *Exercise #8 Due At 3:00 P.M.*
- *Exercise #9 Specification Distributed*

√ *CCN*, § 8.2–8.3.

♢ *CECN*, § 9.1–9.2.

March 27 : Spring Break

- *No Meeting*

March 29 : Spring Break

- *No Meeting*

Language

April 03 : Reading, Morphology, & Mental Rules

√ *CCN*, § 9.1–9.3.

♢ *CECN*, § 10.1–10.5.

♢ Sejnowski, T. J. & Rosenberg, C. R. (1986) NETtalk: A parallel network that learns to read aloud. *Neurocomputing*, Chapter 40.

♢ McClelland, J. L. (1986) The programmable blackboard model of reading. *PDP*, Chapter 16.

♢ Rumelhart, D. E. & McClelland, J. L. (1986) On learning the past tenses of english verbs. *PDP*, Chapter 18.

♡ Seidenberg, M. S. & McClelland, J. L. (1989) A distributed developmental model of word recognition and naming. *Psychological Review*, 96, pp. 523–568.

♡ Plaut, D. C., McClelland, J. L., Seidenberg, M. S., & Patterson, K. (1996) Understanding normal and impaired word reading: Computational principles in quasi-regular domains. *Psychological Review*, 103, pp. 56–115.

♡ Harm, M.W. & Seidenberg, M.S. (1999) Phonology, reading acquisition, and dyslexia: Insights from connectionist models. *Psychological Review*, 106, pp. 491–528.

♡ Pinker, S. & Prince, A. (1988) On language and connectionism: Analysis of a parallel distributed processing model of language acquisition. *Cognition*, 28, pp. 73–193.

♡ Plunkett, K. & Marchman, V. (1991) U-shaped learning and frequency effects in a multi-layered perceptron: Implications for child language acquisition. *Cognition*, 38, pp. 43–102.

April 05 : Semantics, Syntax, & Sentence Processing

- *Quiz #9 At 3:00 P.M.*
- *Exercise #9 Due At 3:00 P.M.*
- *Exercise #10 Specification Distributed*

√ *CCN*, § 9.4–9.6.

◇ *CECN*, § 10.6–10.9.

◇ McClelland, J. L. & Kawamoto, A. H. (1986) Mechanisms of sentence processing: Assigning roles to constituents. *PDP*, Chapter 19.

♡ St. John, M. F. & McClelland, J. L. (1990) Learning and applying contextual constraints in sentence comprehension. *Artificial Intelligence*, 46, pp. 217–257.

♡ Elman, J. L. (1993) Learning and development in neural networks: The importance of starting small. *Cognition*, 48, pp. 71–99.

Higher-Level Cognition

April 10 : Cognitive Control

√ *CCN*, § 10.3–10.4.

◇ *CECN*, § 9.4–9.9.

◇ *CECN*, § 11.1–11.3.

♡ Cohen, J. D., Dunbar, K., & McClelland, J. L. (1990) On the control of automatic processes: A parallel distributed processing model of the Stroop effect. *Psychological Review*, 97, pp. 332–361.

♡ Cohen, J. D. & Servan-Schreiber, D. (1992) Context, cortex, and dopamine: A connectionist approach to behavior and biology in schizophrenia. *Psychological Review*, 99, pp. 45–77.

April 12 : Cognitive Flexibility & Working Memory

- *Quiz #10 At 3:00 P.M.*
- *Exercise #10 Due At 3:00 P.M.*
- *Exercise #11 Specification Distributed*

√ *CCN*, § 10.1–10.2, 10.5–10.11.

◇ *CECN*, § 11.4–11.8.

♡ 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. Cambridge University Press.

♡ Durstewitz, D., Seamans, J. K., Sejnowski, T. J. (2000) Neurocomputational models of working memory. *Nature Neuroscience*, 3 (supplement), pp. 1184–1191.

♡ Braver, T. S. & Cohen, J. D. (2000) On the control of control: The role of dopamine in regulating prefrontal function and working memory. In S. Monsell & J. Driver (Eds.), *Control of Cognitive Processes: Attention and Performance XVIII*, pp. 713–738. MIT Press: Cambridge, MA.

♡ O'Reilly, R. C., Noelle, D. C., Braver, T. S., & Cohen, J. D. (2002) Prefrontal cortex and dynamic categorization tasks: Representational organization and neuromodulatory control. *Cerebral Cortex*, 12, pp. 246–257.

♡ Rougier, N. P., Noelle, D. C., 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 Science*, 102(20), pp. 7338–7343.

♡ O'Reilly, R. C. & Frank, M. J. (2006) Making working memory work: A computational model of learning in the frontal cortex and basal ganglia. *Neural Computation*, 18, pp. 283–328.

April 17 : Cognitive Development

√ Elman, J. L., Bates, E. A., Johnson, M. H., Karmiloff-Smith, A., Parisi, D., & Plunkett, K. (1996) New perspectives on development. *Rethinking Innateness: A Connectionist Perspective on Development*, Chapter 1. (available from instructor)

♡ Munakata, Y., McClelland, J. L., Johnson, Mark H., Siegler, R. S. (1997) Rethinking infant knowledge: Toward an adaptive process account of successes and failures in object permanence tasks. *Psychological Review*, 104, pp. 686–713.

April 19 : Computational Modeling Challenges

- *Quiz #11 At 3:00 P.M.*
- *Exercise #11 Due At 3:00 P.M.*

◇ *CECN*, Chapter 12.

◇ Rumelhart, D. E. & McClelland, J. L. (1986) PDP models and general issues in cognitive science. *PDP*, Chapter 4.

◇ Norman, D. A. (1986) Reflections on cognition and parallel distributed processing. *PDP*, Chapter 26.

◇ Rumelhart, D. E. & McClelland, J. L. (1986) Future directions. *PDP*, last chapter.

◇ Sejnowski, T. J. (1986) Open questions about computation in cerebral cortex. *PDP*, Chapter 21.

◇ Smolensky, P. (1986) Neural and conceptual interpretation of PDP models. *PDP*, Chapter 22.

♡ Fodor, J. A. & Pylyshyn, Z. W. (1988) Connectionism and cognitive architecture: A critical analysis. *Cognition*, 28, pp. 3–71.

♡ Smolensky, P. (1988) On the proper treatment of connectionism. *Behavioral and Brain Sciences*, 11, pp. 1–74.

April 24 : Project Oral Reports

April 26 : Project Oral Reports

May 01 : Project Oral Reports

May 03 : Project Oral Reports

May 07 : Project Oral Reports (Monday, 6:30 P.M. to 9:30 P.M.)

May 08 : Written Project Reports Due

- *Reports must be in the instructor's hands by 9:30 P.M. on this day.*