

Natural and Artificial Neural Networks

COGS GU4---

Instructors

John Morrison
Philosophy Department
Barnard College
jmorrison@barnard.edu

Christos Papadimitriou
Computer Science Department
Columbia University
christos@columbia.edu

Teaching Assistants

Clayton Sanford
Computer Science Department
Columbia University
clayton@cs.columbia.edu

Samuel Deng
Computer Science Department
Columbia University
samdeng@cs.columbia.edu

Course Description

Artificial neural networks can do amazing things. They can play chess, recognize faces, predict human behavior, learn language, and create art. Natural neural networks -- that is to say, brains -- can do many of the same things, often a little more clumsily. But, unlike artificial networks, they can switch seamlessly between two tasks, learn to perform them without supervision, and do not need to be told to -- actually, they can choose to refuse. Brains provided the initial inspiration for artificial networks, which is why we call them 'artificial neural networks.' But how deep are the similarities between the two? Do they share more than a few abilities, a similar structure, and a common nomenclature?

Learning Outcomes

- Develop an appreciation for the importance of both artificial and natural neural networks
- Learn about cutting-edge research and ongoing controversies in neuroscience, machine learning, and philosophy

- Encounter perspectives from outside one's home field of study
- Computer scientists and neuroscientists will think more critically and abstractly about natural and artificial networks, and philosophers will develop a more concrete, better informed understanding of their mechanics.

Assignments

1. Responses (30%)
 - Weekly responses to reading. Approximately 1 page.
2. Short Paper (10%)
 - One of the weekly responses should be expanded into a five-page paper.
 - Undergraduates must submit by the 7th week of course. Graduate students can submit at any time, or can choose to expand their final paper by five pages (in which case it will be worth 65%).
3. Final paper or project (55%)
 - A final paper or project chosen in consultation with at least one of the professors. It should build on existing expertise in computer science, neuroscience, or philosophy.
 - A final paper might explore the differences in how brains and recurrent neural networks store information in short-term memory, and take a stand on whether those differences are deep or fundamental. Another final paper might consider whether it's appropriate to apply traditional concepts from cognitive science, such as representation and inference, to convolutional neural networks, such as the networks trained to recognize objects. Yet another final paper might propose new ways of behaviorally testing whether artificial neural networks are brain-like, especially now that many of the standard benchmarks have been met.
 - A final project might involve training a recurrent and a non-recurrent neural network to identify the subject in a sentence using an existing database of sentences and then asking which performs better. Another final project might be to train a behavioral decoder on existing neural data and test the extent to which it can successfully predict the animal's behavior. Yet another final project might compare neural networks that are more and less neurally plausible on object recognition and compare their mistakes.
 - Papers should be 10-15 pages in length. Projects are more variable, but should still have a written component. The exact length can be set by the professor on a case-by-case basis.
4. Final paper or project proposal (5%)
 - By the 10th week of the course all students must have submitted a one page plan for their final paper or project and have met with one of the professors about it.

Schedule

(all readings are 30-50 pages)

Week 1 (January 18): Natural neural networks: Historical Introduction (CP)

- Chapter 1 of Papadimitriou's *How Does the Brain Beget the Mind*

Week 2 (January 25): Natural neural networks: Topical Introduction (JM)

- Kandell and Hudspeth, "The Brain and Behavior," Ch 1 of *Principles of Neural Science*

Week 3 (February 1): Classical views of computation (CP)

- Selection from Dasgupta et al.'s *Algorithms*

Week 4 (February 8): Classical views of cognition (JM)

- Pylyshyn (1999), "What's in your mind?"
- Stillings et al. (1995), "What is Cognitive Science?"

Week 5 (February 15): Artificial neural networks: Historical introduction (JM)

- Chapter 3 of Lindsay's *Models of Mind*

Week 6 (February 22): Artificial neural networks: Topical introduction (CP)

- Hinton (1993), "How Neural Networks Learn from Experience"

Week 7 (March 1): General comparison (JM)

- Marcus (2018), "Deep Learning: A critical Appraisal"
- Storrs and Kriegeskorte (2019), "Deep Learning for cognitive neuroscience" in Gazzaniga (ed.) *The Cognitive Neurosciences*

Week 8 (March 15): Memory (JM)

- Gallistel and King, *Memory and the Computational Brain*, Ch 9
- Fusi, "Memory capacity of neural network models" in *Human Memory*

Week 9 (March 22): Introduction to Linguistics (CP)

- Hauser et al. (2002), "The Language Faculty"

Week 10 (March 29): Natural Language Processing, and Language in the Brain (CP)

- Chapter 1 or Jurawski and Martin's *Speech and Language Processing*
- (possibly) A survey of bias in deep NLP Garrido-Munoz et al, Applied Science

Week 11 (April 5): Algorithms (JM & CP)

- Selection from Ch 8 of Clark's *Being There*

Week 12 (April 12): Representations (JM)

- Chapter 8 of Sterelny's *The Representational Theory of Mind*

Week 13 (April 19): Using artificial networks to understand the brain

Lecturer: Larry Abbott, <https://zuckermaninstitute.columbia.edu/larry-f-abbott-phd>

- Dicarlo & Yamins (2016), "Using goal-driven deep learning models to understand sensory cortex"
- Schrimpf et al. (2021), "Artificial Neural Networks Accurately Predict Language Processing in the Brain"
- Wang et al. (2021), "Evolving the Olfactory System with Machine Learning"

Week 14 (April 26): Ethics (CP)

- T Gebru et al. (2018), "Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification"

Honor Code

Approved by the student body in 1912 and updated in 2016, the Code states:

We, the students of Barnard College, resolve to uphold the honor of the College by engaging with integrity in all of our academic pursuits. We affirm that academic integrity is the honorable creation and presentation of our own work. We acknowledge that it is our responsibility to seek clarification of proper forms of collaboration and use of academic resources in all assignments or exams. We consider academic integrity to include the proper use and care for all print, electronic, or other academic resources. We will respect the rights of others to engage in pursuit of learning in order to uphold our commitment to honor. We pledge to do all that is in our power to create a spirit of honesty and honor for its own sake.

Wellness Statement

It is important for undergraduates to recognize and identify the different pressures, burdens, and stressors you may be facing, whether personal, emotional, physical, financial, mental, or academic. We as a community urge you to make yourself--your own health, sanity, and wellness--your priority throughout this term and your career here. Sleep, exercise, and eating well can all be a part of a healthy regimen to cope with stress. Resources exist to support you in several sectors of your life, and we encourage you to make use of them. Should you have any questions about navigating these resources, please visit these sites:

- <http://barnard.edu/primarycare>
- <https://barnard.edu/about-counseling>
- <http://barnard.edu/wellwoman/about>

Center for Accessibility Resources & Disability Services Statement

If you believe you may encounter barriers to the academic environment due to a documented disability or emerging health challenges, please feel free to contact me and/or the Center for Accessibility Resources & Disability Services (CARDS). Any student with approved academic accommodations is encouraged to contact me during office hours or via email. If you have questions regarding registering a disability or receiving accommodations for the semester, please contact CARDS at (212) 854-4634, cards@barnard.edu, or learn more at barnard.edu/disabilityservices. CARDS is located in 101 Altschul Hall.

Affordable Access to Course Texts and Materials Statement

All students deserve to be able to study and make use of course texts and materials regardless of cost. Barnard librarians have partnered with students, faculty, and staff to find ways to increase student access to textbooks. By the first day of advance registration for each term, faculty will have provided information about required texts for each course on CourseWorks (including ISBN or author, title, publisher, copyright date, and price), which can be viewed by students. A number of cost-free or low-cost methods for accessing some types of courses texts are detailed on the Barnard Library Textbook Affordability guide (library.barnard.edu/textbook-affordability). Undergraduate students who identify as first-generation and/or low-income students may check out items from the FLIP lending libraries in the Barnard Library (library.barnard.edu/flip) and in Butler Library for an entire semester. Students may also consult with their professors, the Dean of Studies, and the Financial Aid Office about additional affordable alternatives for having access to course texts. Visit the guide and talk to your professors and your librarian for more details