SYMSYS 103: Modern Research Methods in Cognitive Science

Summer Quarter: June 22 - August 13, 2026

Instructor: Noah Hermalin (he)
Class Time: Two 2-hour class sessions per week (Day/Time TBA)

Course Description

Please note that the syllabus as presented here is tentative and subject to change.

The purpose of this course is twofold: First, to give students a foundational understanding of new and innovative methods in cognitive science research, particularly those that make use of artificial intelligence/machine learning; and second, to explore the current state of the literature in cognitive science by reading recent reviews and research papers. Among other topics, we will see examples of how current methods are being used to address challenges pertaining to experiment design; data collection, cleaning, and synthesis; and modeling of cognitive processes. By the end of the course, students should have a firm understanding not only of the concepts that underlie modern AI-based research methods, but also how these methodological advances help us understand human minds and behavior. In addition, we will consider the advantages that more traditional methods may still have during this period of rapid technology-driven innovations, as well as risks and ethical considerations associated with AI-based methods.

Prerequisites and Administrative Considerations

This course officially has no listed requirements. However, be aware that we will be covering a range of topics in cognitive science (and associated fields such as psychology and linguistics), and it is recommended that students have at least some background in such areas. Students who are new to cognitive science should be prepared to put in a bit of extra work, especially early on, to make sure that they are able to fully understand readings, lectures, and discussions. In addition, while this course will not go into fine-grained detail on the mathematical foundations of modern AI architecture, students are expected to be comfortable with reading and understanding mathematical notations and formulae. If you have any questions or concerns about background, please let the instructor know!

Course Requirements

Grading

The grading breakdown for this course is as follows:

Discussion and Participation	20%
Exams and quizzes	30%
Reading Assignments	30%
Project Proposal	10%
Additional Assignments	10%

Project Proposal

During the latter half of the course, each student will be expected to write a project proposal for an original research project. This proposal should integrate ideas and methods that we have learned about in the course. Note that only a project *proposal* is required, and students are not expected to carry out a research project for the purpose of this class.

Exams and Quizzes

This course will include a final exam. In addition, there will be regular quizzes. These quizzes will usually focus on new content learned in a given week. The quizzes are intended as low-stakes opportunities for students to demonstrate their newly acquired knowledge, as well as a way for the instructor to keep track of how well students are understanding the material. Each student's two lowest-scoring quizzes will be dropped.

Reading Assignments

The course is fairly reading-intensive, with most day's content requiring a familiarity with that day's reading. As such, students are expected to finish the readings before class. We'll discuss some strategies for how to read academic publications during the first week of class. If you find yourself struggling with the readings, don't hesitate to get in touch, we can further discuss that might make the reading portion of the course go more smoothly.

In addition to doing the readings, most readings will also have a corresponding homework assignment. These reading assignments will vary, and may include notes/annotations, short question/problem sets, or similar tasks.

Participation

In-class discussion and participation constitute a significant portions of the course. This has two consequences: first, attendance is **mandatory** (unless otherwise excused); and second, part of each student's grade will be based on their participation in in-class discussions and activities. Keep in mind that what it means to be an active participant can vary from person to person and from context to context. As long as you attend class and actively (and respectfully) engage in the daily group discussions, you will have no problem scoring a perfect participation grade.

While attendance is required, sometimes things occur that prevent attendance, and missing class for a valid reason will never count against your participation grade. If you are unable to attend class (e.g. because of illness, family emergency, a wedding, etc.), please try to let me know as early as possible. In addition, each person's first unexcused absence will not be factored into their grade.

Ethics, Community, Communication, and Additional Information

Interactions and discussions: Learning requires a respectful atmosphere. I humbly implore that you always be respectful in class to the best of your ability. Hurtful or harmful statements and actions will not be tolerated. If at any point there's anything I can do to make class time more welcoming and conducive towards your education, please do not hesitate to let me know.

People should always feel welcome to agree or disagree with each other or the readings, but be mindful of the difference between respectful disagreement/feedback, which is ok, and being disrespectful, condescending, or dismissive, which is not ok. It's my hope that we'll together be able to create and maintain a classroom environment that is open-minded, conducive towards fruitful discussion and learning, and is free from judgment.

Some of the readings and studies we will see in this course may touch on potentially sensitive or out-dated topics. I'll do my best to forecast such topics ahead of time and we'll discuss how best to approach and think about such topics.

Technology in the classroom: Class time each day will have designated periods where no electronics (phones, tablets, laptops, etc.) are permitted, unless needed to satisfy an accommodation. At other times, electronics may be permitted as long as they are solely being used for class-related purposes.

Technology outside of the classroom: The use of generative AI tools to complete any assignments for this course is not allowed unless you have received explicit permission from the instructor, or if the assignment specifically directs you to use AI tools. Examples of generative AI usage which are **not** permitted include, but are not limited to:

- Composing text that is submitted (in an altered or unaltered state) to the course for any written assignment
- Generating summaries of readings which are read instead of reading the actual assigned readings
- Generating notes on readings
- Brainstorming ideas

Violation of this policy will be treated as academic dishonesty (i.e., cheating) and will be addressed as such.

Schedule

The following schedule offers a general overview of the topic of focus for each week. This schedule is tentative, or at least somewhat flexible, and may be adjusted as the quarter continues.

Week 1: Introduction, Technical and Conceptual Basics

Summary: The first week will summarize the goals and structure of the course. We'll cover some basic concepts and information which will help us better understand the course content moving forward. In addition, we'll go over strategies for how best to tackle the assignments for this course, including academic literacy.

Week 2: Theoretical Foundations

Summary: This week, we will continue to build our foundational understanding of the two major 'pieces' of this course: the practical and theoretical basics of contemporary AI; and how the current era of cognitive science fits in with the history of the field.

Week 3: Case Study: Language Acquisition

Summary: This week will focus on a specific topic in cognitive science, namely language acquisition. We will read and discuss recent research publications on language acquisition that make use of AI-based methods in various ways.

Week 4: Case Study: Visual Cognition

Summary: This week will focus on a specific topic in cognitive science, namely vision and visual cognition. We will read and discuss recent research publications on visual cognition that make use of AI-based methods in various ways.

Week 5: Synthetic Data

Summary: This week will focus on the topic of synthetic data: data which is generated as a proxy for naturally occurring data (for example, generating sample texts rather than using extant written documents when training a model). We'll consider the pros and cons, validity, and potential applications of using synthetic data for cognitive science research.

Week 6: Advantages and Disadvantages of AI Methods and Non-AI Methods

Summary: This week, we'll review what we've learned in the course thus far and consider what advantages, and disadvantages, recent AI-based methodological innovations may have

over more 'traditional' methods in cognitive science research. Are there instances where AI-based methods offer little benefit, or are even detrimental, when compared with other approaches? What theoretical and/or ethical shortcomings might call the validity of AI-based methods into question?

Week 7: Students' Choice and CogSci 2026

Summary: This week will focus on topics chosen by you, the students. In addition, we will survey some of the new research being presented at the 2026 meeting of the annual Cognitive Science Society conference, which will be taking place around this time. (Note that students are not expected to attend the conference.)

Week 8: Conclusion, Project Presentations, and Exam Preparation

Summary: In the final week of the course, we will wrap up and review what we've learned in preparation for the final exam. In addition, students will present on their project proposals.

List of Readings

The precise list of readings for the course is still subject to change, both because assignments may be adjusted to better suit students' interests and because the course's focus on ongoing innovations means that papers newly published leading up to the start of the quarter are valid candidates for inclusion. To provide a sense of what kinds of readings to expect, some sample readings are included here.

Abdurahman, S., Atari, M., Karimi-Malekabadi, F., Xue, M. J., Trager, J., Park, P. S., Golazizian, P., Omrani, A., & Dehghani, M. (2024). Perils and opportunities in using large language models in psychological research. *PNAS nexus*, 3(7).

Aher, G. V., Arriaga, R. I., & Kalai, A. T. (2023). Using large language models to simulate multiple humans and replicate human subject studies. In *International conference on machine learning* (pp. 337-371).

Alzahrani, A. (2025). The acceptability and validity of AI-generated psycholinguistic stimuli. *Heliyon*, 11(2).

Cuskley, C., Woods, R., & Flaherty, M. (2024). The limitations of large language models for understanding human language and cognition. *Open Mind*, 8, 1058-1083.

Dillion, D., Tandon, N., Gu, Y., & Gray, K. (2023). Can AI language models replace human participants?. *Trends in Cognitive Sciences*, 27(7), 597-600.

Frank, M. C., & Goodman, N. D. (2025). Cognitive modeling using artificial intelligence. *Annual Review of Psychology*, 77.

Futrell, R., & Mahowald, K. (2025). How linguistics learned to stop worrying and love the language models. *arXiv preprint arXiv:2501.17047.*

Goyal, M., & Mahmoud, Q. H. (2024). A systematic review of synthetic data generation techniques using generative AI. *Electronics*, 13(17), 3509.

Hosseini, E. A., Schrimpf, M., Zhang, Y., Bowman, S., Zaslavsky, N., & Fedorenko, E. (2024). Artificial neural network language models predict human brain responses to language even after a developmentally realistic amount of training. *Neurobiology of Language*, 5(1), 43-63.

Ke, L., Tong, S., Cheng, P., & Peng, K. (2025). Exploring the frontiers of llms in psychological applications: A comprehensive review. *Artificial Intelligence Review*, 58(10), 305.

Ong, D. C. (2024). GPT-ology, Computational Models, Silicon Sampling: How should we think about LLMs in Cognitive Science? *arXiv preprint arXiv:2406.09464.*

Orhan, A. E., & Lake, B. M. (2024). Learning high-level visual representations from a child's perspective without strong inductive biases. *Nature Machine Intelligence*, 6(3), 271-283.

Schulze Buschoff, L. M., Akata, E., Bethge, M., & Schulz, E. (2025). Visual cognition in multimodal large language models. *Nature Machine Intelligence*, 7(1), 96-106.

Van Rooij, I., Guest, O., Adolfi, F., de Haan, R., Kolokolova, A., & Rich, P. (2024). Reclaiming AI as a theoretical tool for cognitive science. *Computational Brain & Behavior*, 7(4), 616-636.

Xu, T., Kuribayashi, T., Oseki, Y., Cotterell, R., & Warstadt, A. (2025). Can Language Models Learn Typologically Implausible Languages? *arXiv preprint arXiv:2502.12317.*