COGS 50.09:
Computational models of social cognition
Winter Term 2024

Instructor:	Dae Houlihan
Course Meetings:	Tuesday 2:25–4:15pm Thursday 2:25–4:15pm
Location:	Anonymous Hall, Room 202
PPL Exercises:	https://comosoco.daeh.info
Email:	dae.houlihan@dartmouth.edu
Office:	Moore Hall 233

Course Description

This course introduces the philosophy and practice of building computational models of social cognition. The course will follow three themes. The topic-theme, *social cognition*, will explore how people infer beliefs and preferences, plan interactions, and reason about emotions. Other topics in cognitive science will be incorporated to build a conceptual foundation for thinking about how people represent each other. The theoretical-theme, *models as epistemological expressions*, will emphasize thinking deeply about why and how models are made and used. We will consider how scientists' views on the mind and the world shape their approaches to building models of the mind, models of people's mental models of the world, and models of people's mental models of other minds. We will examine aspects of modeling that are often implicit or deemphasized, including what alternative models were not chosen, the philosophical traditions behind modeling assumptions, how information is learned and represented by models, how data is measured and used, and the various goals and functions of modeling. The methodological-theme, *probabilistic programming*, will give you a skillset for building formal cognitive models. You will be introduced to a probabilistic programming language (WebPPL) through guided exercises and problem sets. It will be helpful to have prior experience with probability (e.g. Bayes rule, conditioning, inference) and programming in any language (prior experience with WebPPL is not expected).

This course will involve a good deal of in-class discussion. Sessions will typically involve a lecture followed by group work. For some assigned assigned readings you will be asked to respond to questions about the research papers before the session. Lab sessions will focus on probabilistic

programing methods and provide the opportunity for you to get help on problem sets. The course will culminate with a final project, which may be an original model of data you collect, an extension of an existing model using open source data, or a critical paper applying ideas from the course to published research.

Assignments and Grading

40% Probabilistic programing exercises
25% Written responses (pre-class questions, in-class responses)
5% Final project proposal
30% Final project

Attendance

Since discussion is a major part of the course, attendance is expected.

Policies on late assignments

The scores of assignments turned in after the deadline will be scaled by max(0, 1 - days/10), e.g. an assignment turned in 2 days late (24<hours<=48 after the deadline) can be worth a maximum of 80% of the original score.

Week 1		
Thurs Jan 4	Goals of the course Modeling as philosophy	
Week 2		
Tues Jan 9	What makes a model good? What is a model, anyway? Introduction to probabilistic programing	Exercise: - Generative models Reading: - Marr Ch. 1 (1.1-1.2) - Varela Ch. 1 - Tenenbaum JB, Kemp C, Griffiths TL & Goodman ND. (2011). How to Grow a Mind: Statistics, Structure, and Abstraction. Science.
Thurs Jan 11	Generative models	Exercise: - Conditioning Reading: - Cusimano M, Hewitt LB & McDermott JH. (2023). Bayesian auditory scene synthesis explains human perception of illusions and everyday sounds [Preprint]. - Poldrack RA. (2021). The physics of representation. Synthese.

Week 3		
Tues Jan 16	Causal models, Intuitive theories	Exercise: - Causal and statistical dependence Reading: - Lake BM, Ullman TD, Tenenbaum JB & Gershman SJ. (2017). Building machines that learn and think like people. Behavioral and Brain Sciences. - Gerstenberg T & Tenenbaum JB. (2017). Intuitive Theories. In M. Waldmannn (Ed.), Oxford handbook of causal reasoning (pp. 515–548).
Thurs Jan 18	Hierarchical models, Explanation	Exercise: - Conditional dependence Reading: - McElreath, Chapter 6. The Haunted DAG & The Causal Terror - Pearl J. (2021). Causal and Counterfactual Inference. In M. Knauff & W. Spohn (Eds.), The Handbook of Rationality (pp. 427–438). The MIT Press.
Week 4		
Tues Jan 23	Your epistemology is showing	Exercise: - Hierarchical Models Reading: - McElreath, Chapter 13. Models With Memory - Brooke-Wilson T. (2022). How Is Perception Tractable? - Berke MD, Walter-Terrill R, Jara-Ettinger J & Scholl BJ. (2022). Flexible Goals Require that Inflexible Perceptual Systems Produce Veridical Representations: Implications for Realism as Revealed by Evolutionary Simulations. Cognitive Science.

Thurs Jan 25	Theory of Mind	Exercise: - Social cognition Reading: - Saxe R. (2005). Against simulation: the argument from error. Trends in Cognitive Sciences. - Phillips J, Buckwalter W, Cushman F, Friedman O, Martin A, Turri J, Santos L & Knobe J. (2021). Knowledge before belief. Behavioral and Brain Sciences. - Sap M, LeBras R, Fried D & Choi Y. (2023). Neural Theory-of-Mind? On the Limits of Social Intelligence in Large LMs (arXiv:2210.13312). arXiv.
Week 5		
Tues Jan 30	Inference	Exercise: - Algorithms for inference - Bayesian data analysis Reading: - McElreath, Chapter 9. Markov Chain Monte Carlo
Thurs Feb 1	Inverse planning	Reading: - Baker CL, Jara-Ettinger J, Saxe R & Tenenbaum JB. (2017). Rational quantitative attribution of beliefs, desires and percepts in human mentalizing. Nature Human Behaviour. - Rabinowitz N, Perbet F, Song F, Zhang C, Eslami SMA & Botvinick M. (2018). Machine theory of mind. In J. Dy & A. Krause (Eds.), Proceedings of the 35th international conference on machine learning (Vol. 80, pp. 4218–4227). PMLR.
Week 6		
Tues Feb 6	Emotion Reasoning and Emotion Recognition	Reading: - Houlihan SD, Ong DC, Cusimano M & Saxe R. (2023). Causal inference over an intuitive theory of emotion. - Ong DC, Zaki J & Goodman ND. (2015). Affective cognition: Exploring lay theories of emotion. Cognition. - Cowen AS & Keltner D. (2020). What the face displays: Mapping 28 emotions conveyed by naturalistic expression. American Psychologist.

Thurs Feb 8	Emotion Prediction	Reading: - Houlihan SD, Kleiman-Weiner M, Hewitt LB, Tenenbaum JB & Saxe R. (2023). Emotion prediction as computation over a generative theory of mind. Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences. - Thornton MA & Tamir DI. (2017). Mental models accurately predict emotion transitions. Proceedings of the National Academy of Sciences.
Week 7		**Project proposal Due**
Tues Feb 13	Agent models learning to act	Reading: - Tsividis PA, Loula J, Burga J, Foss N, Campero A, Pouncy T, Gershman SJ & Tenenbaum JB. (2021). Human-Level Reinforcement Learning through Theory-Based Modeling, Exploration, and Planning (arXiv:2107.12544). arXiv. - Gan C, Zhou S, Schwartz J, Alter S, Bhandwaldar A, Gutfreund D, Yamins DLK, DiCarlo JJ, McDermott J, Torralba A & Tenenbaum JB. (2022). The ThreeDWorld Transport Challenge: A Visually Guided Task-and-Motion Planning Benchmark Towards Physically Realistic Embodied AI. 2022 International Conference on Robotics and Automation (ICRA).
Thurs Feb 15	Final projects	Lab
Week 8		
Tues Feb 20	Probabilistic program learning	Reading: - Lake BM, Salakhutdinov R & Tenenbaum JB. (2015). Human-level concept learning through probabilistic program induction. Science. - Gershman S & Goodman N. (2014). Amortized inference in probabilistic reasoning. Proceedings of the 36th Annual Conference of the Cognitive Science Society.
Thurs Feb 22	Final projects	Lab
Week 9		

Tues Feb 27	Neuro-symbolic probabilistic program synthesis	Reading: - Hewitt LB, Anh Le T & Tenenbaum JB. (2020). Learning to learn generative programs with Memoised Wake-Sleep. In J. Peters & D. Sontag (Eds.), Proceedings of the 36th conference on uncertainty in artificial intelligence (UAI) (Vol. 124, pp. 1278–1287). PMLR. - Ellis K, Wong L, Nye M, Sablé-Meyer M, Cary L, Anaya Pozo L, Hewitt L, Solar-Lezama A & Tenenbaum JB. (2023). DreamCoder: growing generalizable, interpretable knowledge with wake–sleep Bayesian program learning. Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences.
Wed Feb 28 (NB X-period. No class on Thurs	Probabilistic language of thought	Reading: - Wong L, Grand G, Lew AK, Goodman ND, Mansinghka VK, Andreas J & Tenenbaum JB. (2023). From Word Models to World Models: Translating from Natural Language to the Probabilistic Language of Thought (arXiv:2306.12672). arXiv. - Grand G, Wong L, Bowers M, Olausson TX, Liu M, Tenenbaum JB & Andreas J. (2023). LILO: Learning Interpretable Libraries by Compressing and Documenting Code (arXiv:2310.19791). arXiv.
Week 10		
Tues March 5	Project presentations	
		Final projects due

Academic Honor

The faculty, administration, and students of Dartmouth College acknowledge the responsibility to maintain and perpetuate the principle of academic honor, and recognize that any instance of academic dishonesty is considered a violation of the Academic Honor Principle <https://students.dartmouth.edu/community-standards/policy/academic-honor-principle>.

Title IX

At Dartmouth, we value integrity, responsibility, and respect for the rights and interests of others, all central to our Principles of Community. We are dedicated to establishing and maintaining a safe and inclusive campus where all community members have equal access to

Dartmouth's educational and employment opportunities. We strive to promote an environment of sexual respect, safety, and well-being. Through the Sexual and Gender-Based Misconduct Policy (SMP), Dartmouth demonstrates that sex and gender-based discrimination, sex and gender-based harassment, sexual assault, dating violence, domestic violence, stalking, etc., are not tolerated in our community.

For more information regarding Title IX and to access helpful resources, visit Title IX's website (sexual-respect.dartmouth.edu). As a faculty member, I am required to share disclosures of sexual or gender-based misconduct with the Title IX office.

If you have any questions or want to explore support and assistance, please contact the Title IX office at 603-646-0922 or TitleIX@dartmouth.edu. Speaking to Title IX does not automatically initiate a college resolution. Instead, much of their work is around providing supportive measures to ensure you can continue to engage in Dartmouth's programs and activities.

The Research Center for Writing, and Information Technology (RWiT)

The Student Center for Research, Writing, and Information Technology (RWiT) is a place where you can meet with an undergraduate tutor to discuss a paper, research project, or multi-media assignment. The RWit tutors are trained to help you at any phase of your process. Whether you are brainstorming or planning, drafting or structuring, tweaking or polishing, the RWiT tutors can provide feedback that will help you to create final products of which you can be proud. (http://www.dartmouth.edu/~rwit/)

Religious Observances

Dartmouth has a deep commitment to support students' religious observances and diverse faith practices. Some students may wish to take part in religious observances that occur during this academic term. If you have a religious observance that conflicts with your participation in the course, please meet with me as soon as possible—before the end of the second week of the term at the latest—to discuss appropriate course adjustments.

Student Accessibility and Accommodations

Students requesting disability-related accommodations and services for this course are required to register with Student Accessibility Services (SAS; Apply for Services webpage; student.accessibility.services@dartmouth.edu; 1-603-646-9900) and to request that an accommodation email be sent to me in advance of the need for an accommodation. Then, students should schedule a follow-up meeting with me to determine relevant details such as what role SAS or its Testing Center may play in accommodation implementation. This process works best for everyone when completed as early in the quarter as possible. If students have questions about whether they are eligible for accommodations or have concerns about the implementation of their accommodations, they should contact the SAS office. All inquiries and discussions will remain confidential.

Mental Health and Wellness

The Committee on Student Life, with the approval and consideration of the COI, recommends including the following in your syllabus:

The academic environment is challenging, our terms are intensive, and classes are not the only demanding part of your life. There are a number of resources available to you on campus to support your wellness, including: the Counseling Center which allows you to book triage appointments online, the Student Wellness Center which offers wellness check-ins, and your undergraduate dean. The student-led Dartmouth Student Mental Health Union and their peer support program may be helpful if you would like to speak to a trained fellow student support listener. If you need immediate assistance, please contact the counselor on-call at (603) 646-9442 at any time. Please make me aware of anything that will hinder your success in this course.

Consent to Record

Particularly during remote teaching and learning terms, it is recommended that faculty include the following language in their course syllabus or Canvas site. You may also want to review this policy with your students, pointing in particular to the section about the instructor's copyright of materials and about the prohibition of recording one-on-one meetings.

(1) Consent to recording of course meetings and office hours that are open to multiple students.

By enrolling in this course,

a) I affirm my understanding that the instructor may record meetings of this course and any associated meetings open to multiple students and the instructor, including but not limited to scheduled and ad hoc office hours and other consultations, within any digital platform, including those used to offer remote instruction for this course.
b) I further affirm that the instructor owns the copyright to their instructional materials, of which these recordings constitute a part, and my distribution of any of these recordings in whole or in part to any person or entity other than other members of the class without prior written consent of the instructor may be subject to discipline by Dartmouth up to and including separation from Dartmouth.

(2) Requirement of consent to one-on-one recordings

By enrolling in this course, I hereby affirm that I will not make a recording in any medium of any one-on-one meeting with the instructor or another member of the class or group of members of the class without obtaining the prior written consent of all those participating, and I understand that if I violate this prohibition, I will be subject to discipline by Dartmouth up to and including separation from Dartmouth, as well as any other civil or criminal penalties under applicable law. I understand that an exception to this consent applies to accommodations approved by SAS for a student's disability, and that one or more students in a class may record class lectures, discussions, lab sessions, and review sessions and take pictures of essential information, and/or be provided class notes for personal study use only.

If you have questions, please contact the Office of the Dean of the Faculty of Arts and Sciences.