

COGS 50.09

Computational models of social cognition

FALL 2024

Instructor:	Dae Houlihan
Course Meetings:	TBD
Location:	TBD
PPL Exercises:	https://comosoco.daeh.info
Email:	dae.houlihan@dartmouth.edu
Office:	Moore Hall 240

COURSE DESCRIPTION

This course introduces the philosophy and practice of building computational models of social cognition. The course will follow three themes. The methodological-theme, **probabilistic programming**, will give you a skillset for building formal cognitive models. You will be introduced to a probabilistic programming language (PPL) through guided exercises and problem sets. We will practice ‘thinking in code’ while covering core topics in causal inference, Bayesian data analysis, probability as extended logic, generative models, and symbolic AI. The topic-theme, **social cognition**, will explore how people infer others’ beliefs and preferences, plan interactions, and reason about emotions. We will see how the probabilistic programming readily translates to modeling human cognition, exploring the computational bases of mental representation, concepts, perception, learning, and reasoning. 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 model building.

Class sessions will include lectures, discussion, interactive demos, and guided exercises. Readings will include state-of-the-art research papers, seminal theoretical perspectives, textbook chapters, and methodical tutorials. Assignments will include problem sets, reading and discussion responses, and a final modeling project or paper.

COURSE GOALS

By the end of the course, you should have acquired a toolset enabling you to build probabilistic generative models of human cognition. You should also be well positioned to describe what epistemological commitments are implied by your modeling choices.

PREREQUISITES

COSC 1 or PSYC 10 or instructor permission.

It will be helpful to have prior experience with probability (e.g. Bayes rule, conditioning, inference). Prior experience with a PPL is not expected.

ASSIGNMENTS AND GRADING

- 40% Probabilistic programming exercises
- 20% Written responses (pre-class reading questions, in-class discussion responses)
- 10% In-class participation (discussions and attendance)
- 5% Final project proposal
- 25% Final project

The course culminates 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.

ATTENDANCE

Attendance is expected. The material builds on itself every session, and small group discussion are a major part of the course.

LATE ASSIGNMENTS

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

COURSE SCHEDULE

Date	Topic	Assignment
(1) Mon Sep 16	Goals of the course	What makes a model good? What is a model, anyway? Introduction to probabilistic programming
(2) Wed Sep 18	Modeling as philosophy	Reading: <ul style="list-style-type: none">• Marr D. (1982). <i>Vision: A Computational Investigation into the Human Representation and Processing of Visual Information</i>. Ch. 1 (1.1-1.2).• Jaynes ET, Bretthorst GL. (2003). <i>Probability Theory: The Logic of Science</i>. Preface and Ch. 1.• Varela FJ, Thompson E, Rosch E. (1991). <i>The Embodied Mind: Cognitive Science and Human Experience</i>. Ch. 1.
(3) Mon Sep 23	Generative models — thinking with causal programs (representation, abstraction, and symbols)	Reading: <ul style="list-style-type: none">• Tenenbaum JB, Kemp C, Griffiths TL, Goodman ND. (2011). “How to Grow a Mind: Statistics, Structure, and Abstraction.” <i>Science</i>.• Poldrack RA. (2021). “The Physics of Representation.” <i>Synthese</i>.
(4) Wed Sep 25	Generative models — from Bayesian statistics to epistemological theory	Exercise: <ul style="list-style-type: none">• Generative models

Continued on next page

(Continued)

(5) Mon Sep 30	Intuitive theories — what do we know?	Reading: <ul style="list-style-type: none">• Lake BM, Ullman TD, Tenenbaum JB, Gershman SJ. (2017). “Building Machines That Learn and Think like People.” <i>Behavioral and Brain Sciences</i>.• Gerstenberg T, Tenenbaum JB. (2017). “Intuitive Theories.” <i>Oxford Handbook of Causal Reasoning</i>.
(6) Wed Oct 2	Intuitive theories — inductive constraints	Exercise: <ul style="list-style-type: none">• Conditioning
(7) Mon Oct 7	Causal models — how do we know?	Reading: <ul style="list-style-type: none">• McElreath, Chapter 6. The Haunted DAG & The Causal Terror.• Pearl J. (2021). “Causal and Counterfactual Inference.” <i>The Handbook of Rationality</i>.
(8) Wed Oct 9	Causal motifs — conservation of belief, patterns of explanation	Exercise: <ul style="list-style-type: none">• Dependence
(9) Mon Oct 14	Causal analysis — backdoors, do-calculus, adjustment sets	Reading: <ul style="list-style-type: none">• Cinelli C, Forney A, Pearl J. (2022). “A Crash Course in Good and Bad Controls.” <i>Sociological Methods & Research</i>.
(10) Wed Oct 16	Intuitive causal reasoning — from BDA to mental models	Exercise: <ul style="list-style-type: none">• Conditional Dependence
(11) Mon Oct 21	Theory of Mind	Reading: <ul style="list-style-type: none">• Saxe R. (2005). “Against Simulation: The Argument from Error.” <i>Trends in Cognitive Sciences</i>.• Phillips J, Buckwalter W, Cushman F, Friedman O, Martin A, Turri J, Santos L, Knobe J. (2021). “Knowledge before Belief.” <i>Behavioral and Brain Sciences</i>.• Sap M, LeBras R, Fried D, Choi Y. (2023). <i>Neural Theory-of-Mind? On the Limits of Social Intelligence in Large LMs</i>. arxiv.
(12) Wed Oct 23	Inference about inference	Exercise: <ul style="list-style-type: none">• Social Cognition
(13) Mon Oct 28	Algorithms for inference	Reading: <ul style="list-style-type: none">• McElreath, Chapter 9. Markov Chain Monte Carlo• Harrison PMC, Marjeh R, Adolphi F, Rijn P, Anglada-Tort M, Tchernichovski O, Larrouy-Maestri P, Jacoby N. (2020). “Gibbs Sampling with People.” <i>Advances in Neural Information Processing Systems</i>.

Continued on next page

(Continued)

<p>(14) Wed Oct 30</p>	<p>Inverse planning</p>	<p>Reading:</p> <ul style="list-style-type: none">• Baker CL, Jara-Ettinger J, Saxe R, Tenenbaum JB. (2017). “Rational Quantitative Attribution of Beliefs, Desires and Percepts in Human Mentalizing.” <i>Nature Human Behaviour</i>.• Rabinowitz N, Perbet F, Song F, Zhang C, Eslami SMA, Botvinick M. (2018). “Machine Theory of Mind.” <i>Proceedings of the 35th International Conference on Machine Learning</i>.
<p>(15) Mon Nov 4</p>	<p>Emotion reasoning and emotion recognition</p>	<p>Reading:</p> <ul style="list-style-type: none">• Houlihan SD, Ong D, Cusimano M, Saxe R. (2022). “Reasoning about the Antecedents of Emotions: Bayesian Causal Inference over an Intuitive Theory of Mind.” <i>Proceedings of the 44th Annual Conference of the Cognitive Science Society</i>.• Ong DC, Zaki J, Goodman ND. (2015). “Affective Cognition: Exploring Lay Theories of Emotion.” <i>Cognition</i>.• Cowen AS, Keltner D. (2020). “What the Face Displays: Mapping 28 Emotions Conveyed by Naturalistic Expression.” <i>American Psychologist</i>.
<p>(16) Wed Nov 6</p>	<p>Emotion prediction</p>	<p>Reading:</p> <ul style="list-style-type: none">• Houlihan SD, Kleiman-Weiner M, Hewitt LB, Tenenbaum JB, Saxe R. (2023). “Emotion Prediction as Computation over a Generative Theory of Mind.” <i>Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences</i>.• Thornton MA, Tamir DI. (2017). “Mental Models Accurately Predict Emotion Transitions.” <i>Proceedings of the National Academy of Sciences</i>.
<p>(17) Mon Nov 11</p>	<p>Probabilistic language of thought</p>	<p>Reading:</p> <ul style="list-style-type: none">• Goodman ND, Gerstenberg T, Tenenbaum JB. (2023). “Probabilistic Programs as a Unifying Language of Thought.” <i>Bayesian Models of Cognition: Reverse Engineering the Mind</i>.
<p>(18) Wed Nov 13</p>	<p>Probabilistic program learning</p>	<p>Reading:</p> <ul style="list-style-type: none">• Lake BM, Salakhutdinov R, Tenenbaum JB. (2015). “Human-Level Concept Learning through Probabilistic Program Induction.” <i>Science</i>.• Hewitt LB, Anh Le T, Tenenbaum JB. (2020). “Learning to Learn Generative Programs with Memoised Wake-Sleep.” <i>Proceedings of the 36th Conference on Uncertainty in Artificial Intelligence (UAI)</i>.

Continued on next page

(Continued)

<p>(19) Mon Nov 18</p>	<p>Neuro-symbolic probabilistic program synthesis</p>	<p>Reading:</p> <ul style="list-style-type: none">• 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.” <i>Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences</i>.• Wong L, Grand G, Lew AK, Goodman ND, Mansinghka VK, Andreas J, Tenenbaum JB. (2023). <i>From Word Models to World Models: Translating from Natural Language to the Probabilistic Language of Thought</i>. arxiv.
----------------------------	---	--

LEARNING RESOURCES

THE RESEARCH CENTER FOR WRITING, AND INFORMATION TECHNOLOGY (RWIT)

The Student Center for Research, Writing, and Information Technology (RWIT; <https://writing.dartmouth.edu/support/writing-center>) 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.

ACADEMIC SKILLS CENTER (ASC)

Open to the entire Dartmouth community, the ASC assists students in achieving their academic goals through tutoring and learning skills trainings: <https://students.dartmouth.edu/academic-skills/>.

POLICY STATEMENTS

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.

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 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

(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.