

PHL342H: Minds & Machines

Lectures: Mondays, 3pm-5pm

Location: RW 110 (first two weeks on Zoom)

TAs:

Liang Zhou Koh, liangzhou.koh@mail.utoronto.ca

Philip Rajewicz, philip.rajewicz@mail.utoronto.ca

Instructor: Daniel Munro

Email: daniel.munro@mail.utoronto.ca

Office hours: Wednesdays @ 3pm on Zoom (link in first Quercus module)

Course Description

Cognitive science studies the mind by applying methods from a variety of disciplines, including psychology, philosophy, neuroscience, linguistics, and computer science. Artificial intelligence research applies insights from cognitive science to investigate the relationship between minds and computers. In this course, we'll investigate the philosophical foundations and applications of these disciplines.

The course will be divided into three parts:

1. **What is the Mind?:** Foundational philosophical issues about the nature of the mind.
2. **Could Machines Have Minds?:** Philosophical considerations about artificial intelligence.
3. **Epistemological, Ethical, and Social Dimensions of AI:** What are the epistemological, ethical, and social implications of AI research and development?

Course Delivery

Our first two lectures for the course (Sept. 13 & 20) will be hosted on Zoom, given the Faculty of Arts and Science requirement that students may attend classes online until September 23. You can find the Zoom link for these lectures in the first Quercus module. Beginning September 27, lectures will be in person.

Tutorials begin September 20; the first tutorial will be hosted on Zoom, after which tutorials will be in person. Your TA will post a link to your Zoom tutorial in a Quercus announcement. Make sure you register in a tutorial section if you haven't yet done so.

Readings for each Monday lecture will be posted on Quercus at least a week in advance. **You should do the readings and be prepared to discuss them prior to each Monday lecture**, given that tutorials take place immediately after lectures. Lecture slides will be posted on Quercus after each lecture.

Course Evaluation

Tutorial attendance, participation, and reading responses

Worth 15% of course grade. Your TA will evaluate you based on an equal mix of overall attendance, overall participation, and two reading responses which will help guide tutorial discussion. You'll sign up to submit a reading response on the Quercus discussion board twice this term (**due the Friday before the tutorials you sign up for**), and these student posts will help guide discussion. To get full credit for this assignment,

you must attend and participate in the tutorials you sign up for. See more details on the Quercus discussion board.

Topical Reflection

Worth 10% of course grade. At some point during the term, find a current event, recent news article, popular science article, film, artwork, etc. that's in some way related to our course themes. Post it on the Quercus discussion board with a 300-word explanation of its relevance to the course and why it's philosophically interesting. **Due on Quercus discussion board by December 8**, but you're encouraged to submit it sometime earlier throughout the term. Find more details on the Quercus discussion board.

Papers

Short paper #1: 500 words, worth 15% of course grade. Explain an argument or debate from a course reading (more detailed instructions posted on Quercus before due date). **Due Oct. 11.**

Short paper #2: 1,000 words, worth 25% of course grade. Explain an argument from a course reading and raise an original objection to it (more detailed instructions on Quercus before due date). **Due Nov. 8.**

Term paper or project

The final assignment for the course will be worth 35%. You can choose one of two options for this assignment. **Due Dec. 13.**

Option 1: term paper

Complete a term paper of 2,000 words. Will ask you to take an original stance on some debate we've covered in the course (more detailed instructions posted on Quercus before due date).

Option 2: public philosophy project

The majority of assignments you complete in university courses will be seen by you plus a TA or instructor, after which no one will ever look at them again. This option will instead allow you to prepare a piece of "public philosophy," philosophical work that is instead intended for consumption by a wider audience, thereby broadening the potential impact of your work in the course. The exact design of this project is relatively open-ended: I encourage to exercise your creativity in coming up with an idea that personally excites you (if you like, you can base it on the same topic as your "topical reflection" assignment—see above). The main criterion is that the project directly relates to (some aspect of) this course's philosophical content. Here are some examples of how this could look, though you are by no means limited to these ideas:

- Create a YouTube video or podcast.
- Propose a substantive edit to a Wikipedia article or propose an entirely new Wikipedia article.
- Write a philosophical op-ed or blog post.
- Conduct a philosophically substantive interview with someone whose work is related to course content (philosopher, academic, artist, journalist, etc.).
- Utilize another online medium or social media platform (Reddit, Twitter, Facebook, etc.), perhaps by designing some way to engage non-course participants in philosophical activity.

If you'd like to take advantage of this option, you should **email me by November 29th with a short description of your project idea**. As needed, I'll work with you to make sure the project is substantive enough to satisfy the assignment criteria while feasible to complete without taking up significantly more time and effort than a term paper would (but note that the open-endedness of this project may nevertheless

result in slightly more work than a term paper). Once project details are figured out, I'll also make it clear how your project will be submitted and evaluated (since this may vary a bit by student).

Reading & Lecture Schedule

It's recommended that you do the readings in the order they appear here. All readings will be posted on Quercus at least a week prior to each lecture. Reading list subject to slight modification, with advanced notice, as the course progresses.

Part 1: What is the Mind?

Sept. 13: Dualism vs. Physicalism

- Rene Descartes, *Meditations* (Meditations II & VI required)
- J.J.C. Smart, "Sensations and Brain Processes"

Sept. 20: Functionalism

- Hilary Putnam, "The Nature of Mental states"
- Jerry Fodor, "The Mind-Body Problem"

Sept. 27: Qualia and Consciousness: Dualism's Revenge?

- Frank Jackson, "Epiphenomenal Qualia"
- David Chalmers, *The Conscious Mind* (excerpt)

Oct. 4: Intentionality

- Alex Byrne, "Intentionality"
- Part 1 wrap-up & looking ahead

Oct. 11: No class (Thanksgiving)

- **Short paper #1 due**

Part 2: Could Computers Have Minds?

Oct. 18: The Turing Test

- Alan Turing, "Computing Machinery and Intelligence"

Oct. 25: Functionalism, Intentionality, and Artificial Intelligence

- John Searle, "Minds, Brains, and Programs" (excerpt)
- Dan Dennett, "Can Machines Think?" (excerpt)
- McClelland et al., "Extending Machine Language Models toward Human-Level Language Understanding"

Nov. 1: Consciousness and Artificial Intelligence

- Ned Block, “Comparing the Major Theories of Consciousness”
- Dehaene et al., “What is Consciousness, and Could Machines Have It?”

Nov. 8: No class (reading week)

- **Short paper #2 due**

Nov. 15: The Deep Learning Revolution

- Cameron Buckner, “Deep Learning: A Philosophical Introduction”
- Gary Marcus, “Deep Learning: A Critical Appraisal”

Part 3: Epistemological, Ethical, and Social Dimensions of AI

Nov. 22: Ethics of Robots and Autonomous Vehicles

- Bartneck et al., “Autonomous Vehicles”
- Regina Rini, “Raising Good Robots”

Nov. 29: AI in Criminal Justice

- *Hi-Phi Nation*, “The Precrime Unit” (podcast)
- *Hi-Phi Nation*, “Risky Business” (podcast)

Dec. 6: Deepfakes

- Regina Rini, “Deepfakes and the Epistemic Backstop”

Term paper due Dec. 13

Other Course Business

Course-related communications

All emails to Daniel or your TA should include “PHL342” in the subject line and be sent from a UofT email address. You can expect a reply within 48 hours (so try not to send questions about assignments less than 48 hours before due dates). If you don’t hear back by then, send a follow up. Email is appropriate for short questions about course business not already answered on the syllabus or assignment sheets. More philosophically substantive discussion is better reserved for office hours.

Late and missed work

Late assignments will be assessed a penalty of 1/3 letter grade per day.

Requests for extensions

If extenuating circumstances prevent you from completing a paper on time, please get in touch with your TA as far in advance as possible to request an extension.

Requests for regrades

If you feel a regrade on an assignment is warranted, email Daniel with your request. Include a valid explanation of why you think a re-grade is warranted. (Note: “I need a better grade in this course to [bring up my average/get into a program/apply to grad school]” is not an example of a valid explanation).

Plagiarism detection

Normally, students will be required to submit their course essays to the University’s plagiarism detection tool for a review of textual similarity and detection of possible plagiarism. In doing so, students will allow their essays to be included as source documents in the tool’s reference database, where they will be used solely for the purpose of detecting plagiarism. The terms that apply to the University’s use of this tool are described on the Centre for Teaching Support & Innovation web site (<https://uoft.me/pdt-faq>). Use of this tool is optional, but please email me well before the due date of an assignment if you’d like to make alternative submission arrangements.

Accessibility and accommodations

If you have any disability or health related concerns that require accommodation, please get in touch as soon as possible to figure out how we can best make this course a fruitful learning experience for you. Please also get in touch with [Accessibility Services](#).

If you require accommodation for course conflicts with religious holidays, please also get in touch as soon as possible to make arrangements.

You may also be interested in resources for [mental health support](#), in the [philosophy essay clinic](#) if you’re looking for help with philosophical writing, or in UofT’s [English language learning supports](#).

Academic integrity

It’s your responsibility to familiarize yourself with the behaviours that constitute academic dishonesty. If you’re feeling overwhelmed with coursework, worried about meeting a deadline, or having trouble understanding what’s expected of you on an assignment, please get in touch with me instead of resorting to academic dishonesty. For information about academic integrity at UofT, visit [the University’s academic integrity page](#).