

COGS/PHIL 3750: Philosophy of Artificial Intelligence

Fall Semester 2018-2019

Instructor

Muhammad Ali Khalidi

Email

khalidi@yorku.ca

Office Location & Hours

Ross S 438-Thursdays 11:30-12:30 & by appointment

Course Description

This course will address two sets of philosophical issues raised by Artificial Intelligence (AI). The first set has to do with the nature of intelligence and the very possibility of creating a machine that is intelligent. What would it take to design an intelligent computer? Is this even a feasible goal, or are there principled obstacles to creating artificial intelligence? If it is possible in principle, what kind of features would such a machine have? The second set has to do with the ethical, social, and political implications of artificial intelligence. On the assumption that it is possible to create artificial intelligence systems, what is the likelihood of developing AI systems that are superior to humans in intelligence, and can we ensure that such systems do not cause harm to the human race? Can such systems be programmed in such a way as to avoid any adverse moral consequences? And what would the future look like if many, if not most, current human occupations are taken over by AI systems?

Course Requirements

First in-class quiz (20%): This quiz is designed to make sure that you are doing the reading and following class discussion. There will be some multiple-choice questions and some questions requiring short answers (roughly, one paragraph each). More details and sample questions will be distributed two weeks prior to the quiz.

Second in-class quiz (20%): Same as first in-class quiz.

Class Attendance, Participation, and Weekly Posts (15%) Roughly each week I will post a question on the readings for that week. You will be expected to post an answer to at least five questions during the course of the semester; if you post more, your best five will be counted. Your answer should be no longer than 100 words. Comments have to be posted by the announced deadline; late comments will not be accepted. Attendance and participation will be assessed in the usual ways.

Analytic Essay (20%): Essay topics will be distributed around halfway through the semester. You will be asked to write an essay about an article that we didn't read in class on a topic that we covered in class. Further details and specifications will be announced with essay topics.

Take-Home Final Exam (25%): Take-home final exam will consist of short essays on assigned topics. It will be assigned before classes end and will be due during the final exam period.

Deadlines and Makeups

All deadlines and quiz dates are firm. If, due to a serious medical or other emergency, you need to miss a deadline or quiz date, you must email me *in advance* to say so and then produce *proper documentation* (e.g. a doctor's report) within a week. In such cases, a makeup will be given (but note that makeups may not offer a choice of questions) or late penalties will be waived.

Course Schedule

Topic	Date	Reading
-------	------	---------

All readings listed below are **required** and will be freely available via Moodle. Other material (films, podcasts, news articles) marked with an asterisk (*) is **optional** but may provide important background. Please make sure you do the required reading before the class in question. There is no substitute for reading the assigned material and coming to class prepared to discuss it. (Bibliographic details will be posted on Moodle.)

Introduction:	Thurs Sept 6	No assigned reading
What is super-intelligence?	Tues Sept 11	Bostrom, "When Machines Outsmart Humans" *Podcast: Chalmers, "The Singularity"
	Thurs Sept 13	Muehlhauser and Bostrom, "Why We Need Friendly AI" *Video: Bostrom, "What Happens When Our Computers Get Smarter than We Are?"
Algorithms and Turing Machines	Tues Sept 18	Crane, "Computers and Thought" *Podcast: "The Turing Problem"
	Thurs Sept 20	Haugeland, "Computer Architecture" *Video: "Algorithms: The Secret Rules of Modern Living"
The Turing Test and the Chinese Room	Tues Sept 25	Turing, "Computing Machinery and Intelligence" *Video: "How the Most Human Human Passed the Turing Test"
	Thurs Sept 27	Searle, "Can Computers Think?" *Video: Interview with Searle
	Tues Oct 2	Churchland & Churchland, "Could a Machine Think?" *Podcast: "Clever Bots"
	Thurs Oct 4	Dennett, "Can Machines Think?" *Podcast: "More or Less Human"
Reading Week	Tues Oct 9	No Class
	Thurs Oct 11	No Class
	Tues Oct 16	First in-class quiz

Topic	Date	Reading
Symbolic AI, Neural Networks, and Deep Learning	Thurs Oct 18	Newell & Simon, "Computer Science and Empirical Inquiry"
	Tues Oct 23	Garson, "Connectionism" *Video: But what *is* a Neural Network?
	Thurs Oct 25	Garson, "Connectionism" (continued)
	Tues Oct 30	Gopnik, "An AI that Knows the World like Children Do" *Video: AlphaGo
Robotics and Situated Intelligence	Thurs Nov 1	Brooks, "Intelligence without Representation" *Video: "Making Robots More Like Children"
	Tues Nov 6	Clark & Toribio, "Doing without Representing?"
	Thurs Nov 8	Dennett, "The Practical Requirements for Making a Conscious Robot"
	Tues Nov 13	Second in-class quiz
The Singularity and Ethical Implications	Thurs Nov 15	Rini, "Raising Good Robots" *Video: Rini and Hughes
	Tues Nov 20	Schwitzgebel and Garza, "Designing AI with Rights, Consciousness, Self-Respect, and Freedom"
	Thurs Nov 22	Wynsberghe & Robbins, "Critiquing the Reasons for Making Artificial Moral Agents"
Social and Political Implications of AI	Tues Nov 27	Totschnig, "The problem of superintelligence: political, not technological"
	Thurs Nov 29	Danaher, "Will life be worth living in a world without work? Technological unemployment and the meaning of life"
	Tues Dec 4	Wrap-up and review

Additional Information and Resources

Academic Honesty

All students are expected to abide strictly by standards of academic honesty. Please familiarize yourselves with the [University Senate policy on academic dishonesty](#). If you have any questions concerning what constitutes cheating or plagiarism, please consult with me.

Accessibility

Students with health-related, learning, physical, psychiatric, or sensory disabilities who require reasonable accommodations in teaching style or evaluation methods should discuss their concerns with me *as soon as possible* so that appropriate arrangements can be made.

Moodle and Email

All registered students will be automatically added to the course's Moodle site, so please make sure that you can access Moodle and if you have any problems, please contact technical support (askit@yorku.ca). All required texts for this course will be accessible via Moodle, so it's crucial that you be able to access it as soon as possible. Important announcements concerning the course and course policies will also be posted on Moodle, so you are responsible for checking the Moodle site on a regular basis. I will also post all lecture slides on Moodle (but only *after* they have been presented in class). Email is a good way of communicating with me, but if you have a substantive question about course content, I strongly recommend that you arrange to discuss it in a face-to-face meeting during office hours rather than try to resolve it over email. It's generally difficult to answer philosophical questions adequately in an email message. Please note that I will only respond to emails if you identify yourself by **including your complete name** (both first and last names) in the body of the message!

Use of Electronic Devices

As long as it's not disruptive to those around you, I'm willing to let you use laptops, tablets, and other electronic devices in class to take notes or to look up passages in the readings (in case you decide not to print them out and prefer to read them on a screen). But I would strongly encourage you to *turn your wifi off or put your device in airplane mode*. That way, there will be fewer distractions from incoming emails and other notifications and you can concentrate on what's going on in class. There has been some recent research that suggests that Internet use among students in university classes is correlated with worse performance. As the researchers put it: "nonacademic Internet use was common among students who brought laptops to class and was inversely related to class performance. This relationship was upheld after we accounted for motivation, interest, and intelligence. Class-related Internet use was not associated with a benefit to classroom performance."¹ Other research even suggests that taking notes longhand (using pen and paper, the old-fashioned way) is more conducive to learning than taking notes on a laptop computer.² But since I tend to take notes mostly on a laptop, I can't very well discourage it in others. However, if I find that people are doing things other than taking notes and looking up the readings on their devices, I may have to revise this policy of tolerating the use of electronic devices. Needless to say, you should also keep your phones turned off or silent in class. Cellphones should be used for text messaging only in the case of emergency.

¹ Ravizza, S. M., Uitvlugt, M. G., & Fenn, K. M. (2017). Logged in and zoned out: How laptop internet use relates to classroom learning. *Psychological science*, 28(2), 171-180.

² Mueller, P. A., & Oppenheimer, D. M. (2014). The pen is mightier than the keyboard: Advantages of longhand over laptop note taking. *Psychological science*, 25(6), 1159-1168.