

Meetings: Tuesday and Thursday 11:00 AM - 12:20 PM – 202 HL*

Instructor: Prof. Robert Van Gulick
RNVANGUL@syr.edu

Office Hours: Wed 1:00 PM - 2:30 PM and other times by appointment - all on Zoom.

* Course is currently scheduled to be 100% in-person. Depending on covid situation that may change for some parts of the semester.

Course Topic: An investigation of philosophical issues arising within and with respect to work in Artificial Intelligence and the Computational Theory of Mind. Much of the course will be devoted to trying to understand and answer the question, "Can machines think?" This will require us to define what we mean by a *machine* and figuring out what is involved in being able to *think*. Can machines reason, understand, be conscious, be self-aware, learn, be creative, have emotions, and use natural language? We will ask these questions both about man-made computers, and also about the hypothesis that the mental properties of the human mind are best understood by treating the brain as a kind of computer. Most readings will be by philosophers, but some will be by workers in Artificial Intelligence.

Course Readings: (there are no texts required for purchase):

- All readings will be posted on **Blackboard** at least one week before reading date.
- You will be asked to **watch several sci-fi movies** about AI available on commercial sites for a **small rental fee**. (Total costs should be less than \$25.)

Course Requirements - All parts of the course must be completed to receive course credit.

1. **READING** the assigned readings (*before* the class at which they are to be discussed.)
2. Regular **Attendance & Participation in discussions** including small group discussions*.
3. **MIDTERM EXAM: In-class Thurs, March 10, (OR depending on covid *may be changed to take-home exam* - if so, it will be distributed March 10 and due March 24).**
4. A **take-home FINAL EXAM Monday May 10, 5:15pm-7:15pm** (depending on covid *may be changed to take-home exam – if so, dates tba*)
5. An **ESSAY** (5 pages) on one of a variety of topics that will be distributed during the second half of the course. There also be a **peer-commentary component**.
6. Regular **SHORT WRITINGS** (1-page or less) will be assigned - weekly or biweekly.
7. **THINKING** about the issues and questions both in class and out of class (very important - *this class is about thinking and **requires active thinking.***)

Basis of course grade:

Midterm Exam = 25%,

Final Exam = 25%,

Essay = 25%,

Participation & Short Assignments (including peer commentary of essay) = 25%

*Midterm exam, final exam and essay will get normal quality-based letter grades A-F.

Participation and short assignment component will receive A grade if all parts are satisfactorily completed.

Part I of Course (up to Midterm)

Week of 1/25 & 1/27 Can Machines Think? - Defining the question.

"In the Age of AI", video from PBS Frontline (November 2019). Can be watched at <https://www.pbs.org/video/in-the-age-of-ai-zwfwzb/>

Week of 2/1 & 2/3 Where did the question come from? Can animals think?

Descartes, Rene. three selections: excerpts from the *The Discourse on Method*, *letter to Henry More*, and *letter to the Marquis of Newcastle*.

Haugeland, John. "The saga of the modern mind" from *Mind Design*.

Week of 2/8 & 2/10 - The Turing Test - One possible answer to the question.

Turing, Alan. "Computing machinery and intelligence"

Moor, James. "Turing Test"

Dennett, Daniel. "Can machines think?"

Week of 2/15 & 2/17 Intentional Systems - Another behavioral answer.

Dennett, Daniel. "True believers, the intentional stance and why it works"

Week of 2/22 & 2/24 Intrinsic Meaning & the Introspective challenge

Searle, John. "Minds, brains and programs"

Week of 3/1 & 3/3 Brain-like computing - Nets, Nodes and Networks

Churchland, Patricia and Churchland, Paul. "Could a machine think?"

Searle, John. "Is the brain's mind a computer program?"

****Week of 3/8 & 3/10 Semantic Grounding - Where does meaning come from?**

3/8 Van Gulick, Robert. "Consciousness, intrinsic intentionality and self-understanding machines"

➡ **3/10, Thurs: Midterm Exam in-class (OR take-home exam distributed)****

END of Part I of course

Part II of Course

*List of likely readings - subject to change. Dates for readings tba.**

Possible Solutions (?) to the Semantics Problem

- Selections from John Haugeland, *AI The Very Idea*: “Semantics” offering possible solutions to the Paradox of Mechanical Reason, Mystery of Original Meaning, and Disaster of the Homuncular Regress.

Could a machine/computer be conscious?

- Armstrong, David, “What is consciousness?”
- McGinn, Colin, “Could a machine be conscious?”
- Nagel, Thomas “What is it like to be a bat?”
- Van Gulick, Robert “*Consciousness Reconsidered* Reconsidered”
- Zelazney, Roger “For a breath I tarry”

Could a computer/robot be a person?

- Dennett, Daniel “Conditions of Personhood”
- Asimov, Isaac “The bicentennial man”

Could machines be moral?

- Colin Allen, “Why Machine Ethics?”
- Merel Noorman “Computing and Moral Responsibility” (*Stanford Encyclopedia of Philosophy*) <https://plato.stanford.edu/entries/computing-responsibility/>
- *NYTimes* “UN panel urges moratorium on LARS” (Lethal Autonomous Robotic Systems)
- Abney, Keith. “Autonomous Robots and the Future of Just War Theory.” In *Routledge Handbook of Ethics and War: Just War Theory in the 21st Century*, edited by Fritz Allhoff, Nicholas G. Evans, and Adam Henschke, 338–51. Routledge, 2013.
- Dennett, Daniel C. “When HAL Kills, Who’s to Blame?” In *HAL’s Legacy: 2001’s Computer As Dream and Reality*, edited by David G. Stork. MIT Press, 1998.

The Social Hazards of AI and Algorithms

- Nguyen, C. Thi, “Echo Chambers and Epistemic Bubbles”
- Johnson, Gabbrielle M., Algorithmic bias: on the implicit biases of social technology.
- “Coded Bias”, PBS *Independent Lens* program from 3/22/21 (on Netflix). Also PBS Passport <https://www.pbs.org/independentlens/films/coded-bias/>

The Singularity & the Simulation

- David Chalmers, “The Singularity: A Philosophical Analysis”
- Stanislaus Lem, “Non servian” (also called “The experiment”)

AI in Science Fiction Films. during second half of the term you will also be asked to watch and think about a number of scientific films about AI including:

- *Blade Runner* (classic film directed by Ridley Scott, based on Philip K Dick's *Do Androids Dream of Electric Sheep?*, about artificial (bio-engineered) human-like beings).
- *Ex Machina* (film about testing an android to determine if it conscious, 2015).
- *Her* (Joaquin Phoenix, Scarlett Johanson film about a sentient operating system, 2013).
- *A.I. Artificial Intelligence*. Spielberg film about android ('mecha') programmed to love.

**Part II Readings subject to some changes. Other readings and topics may be added as our conversation develops. You will given be a retrospective list of everything we have covered before exams*