

Artificial Intelligence and Technology Syllabus

PHIL ### Syllabus

CRN: ###

Instructor: Jacob Huls

Email: jacob.huls@slu.edu

HOW TO CONTACT: Contact via email any time. I will try to get back to you as quickly as I can.

Office hours: By appointment. There will be office hours scheduled around tests. I will notify you about these closer to those respective assignments.

10:00 am–10:50 MW

Hall 214

Fall or Spring 20##

1. Course Description

This is a sample AI and technology syllabus. It is intended to introduce students to contemporary debates on artificial intelligence, machine learning, and technology generally. It is (to some degree) geared towards computer science undergrads in addition to higher-level philosophy undergrads. We start with some fundamental metaphysical questions: Do contemporary AI systems deliver things with content and meaning? If so, how is that possible? We move on to consider the famous question about whether strong AI is possible: that is, can there be conscious artificial intelligence? We move on to consider some epistemological questions that have arisen out of recent developments in AI and technology generally. Are some AI systems evidence of empiricism? Is it likely that we are living in a simulation? Finally, we consider a number of pragmatic and ethical questions: Is it feasible that AI could eventually destroy us? Should we try to instill values into AI? If so, what values and in what way should we attempt to instill them?

2. Schedule and Readings

REQUIRED TEXTS:

- [*Making AI Intelligible*](#) by Herman Cappelen and Josh Dever (FREE OPEN ACCESS PDF AT LINK)
- *Reality +: Virtual Worlds and the Problems of Philosophy* by David Chalmers
ISBN: 0393635805
- *Superintelligence: Paths, Dangers, Strategies* by Nick Bostrom
ISBN: 0143442465

REQUIRED MEDIA

- *Star Trek: The Next Generation*, Season 2, Episode 9 (if you don't have Paramount+, you can watch this for free with a Paramount+ free trial. Otherwise, you can buy just this episode digitally on Amazon)

FREE RESOURCES:

- [1000 Word Philosophy](#) – A collection of very short (approximately 1000-word) entries on many philosophical topics.
- [The Internet Encyclopedia of Philosophy](#)
- [The Stanford Encyclopedia of Philosophy](#) – this is generally less accessible than the Internet Encyclopedia of Philosophy for beginners

REQUIRED MATERIALS: Please bring a pen or pencil and plenty of paper to class. At the very least, make sure you bring plenty of paper and a pen/pencil to the final test.

READINGS: All readings either come from our required texts or are posted on Canvas.

The difficulty of the required readings vary. Some of the readings are relatively easy, while others are relatively difficult. You are not expected to fully grasp the readings. When you find that you are not understanding a particular part of the reading, keep going and prepare to ask questions in class.

Readings may be removed, added, or changed (although it is unlikely this will happen). But I will notify you about any changes at least two weeks before the readings are due.

SCHEDULE

Week 1: What are Neural Nets, ML, and AI? Intro to the Question of Content

Mon #/#: Syllabus, What is AI?

Reading: No reading, but please do the following online module: [Neural Networks from Scratch](#). It may be a little complicated for some, but it will help you get the gist of what neural networks, machine learning, and artificial intelligence are.

Wed #/#: Intro: Content and AI systems

Reading: *Making AI Intelligible* (henceforth MAII) pages 1–30

Week 2: Skepticism about Content

Mon #/#: Dialogue w/ Alfred the Skeptic

Reading: MAII pages 31–50

Wed: #/#: Some philosophy of language

Reading: MAII pages 51–80

Week 3: Devising a theory of Content for AI

Mon #/#: Picking out properties
Reading: MAII pages 81–102

Wed #/#: Reference and AI
Reading: MAII pages 103–138

Week 4: Finishing the content question; beginning the strong AI debate

Mon #/#: Conclusion of MAII
Reading: MAII pages 139–166

Wed #/#: The Chinese Room Argument
WATCH: *Star Trek: The Next Generation*, Season 2, Episode 9: “The Measure of a Man”
Reading: [The Chinese Room SEP Article](#)

Week 5: Gödel and Strong AI

Mon #/#: What are Gödel’s theorems and do they show strong AI is impossible?
Reading: [“Minds, Machines, and Gödel”](#) by J. R. Lucas (PDF on Canvas in addition to link)
Instructor Handout on Gödel’s incompleteness theorems

Wed #/#: Putnam’s response
Reading: [“Minds and Machines”](#) by Hilary Putnam

Week 6: The Simulation Hypothesis I

Mon #/#: Introduction to the simulation hypothesis
Reading: *Reality* + (henceforth R+) pages xi–19

Wed #/#: Intro to the simulation hypothesis continued
Reading: R+ pages 20–42

Week 7: The Simulation Hypothesis II

- Mon #/#: Some basic epistemology
Reading: R+ pages 43–80 (**NOTE:** This is a longer reading. But, if you're a philosophy student you can probably skip a lot [especially if you've had an epistemology class before this]. If you aren't a philosophy student, I recommend focusing on pages 43–60. Then, if you have time, read 61–80. It is also worth noting that this reading is very easy).
- Wed #/#: The simulation argument
Reading: R+ pages 81–104

Week 8: Philosophical Questions about Simulated and Virtual Reality I

- Mon #/#: What is reality? Is simulated reality real?
Reading: R+ pages 105–123
- Wed #/#: The simulation argument for God
Reading: R+ pages 124–146

Week 9: Philosophical Questions about Simulated and Virtual Reality II

- Mon #/#: Is the universe made of information?
Reading: R+ pages 145–166
- Wed #/#: Can there be consciousness in a digital world?
Reading: R+ pages 274–293

Week 10: Practical & Ethical Questions I

- Mon #/#: Deepfakes
Reading: “Deepfakes, Deep Harms” by Regina Rini & Leah Cohen (PDF on Canvas)
R+ pages 238–254
(Longer reading again, focus on the first reading)
- Wed #/#: Do simulated lives matter?
Reading: R+ pages 331–349

Week 11: Practical & Ethical Questions II: Could Superintelligence Take Over

- Mon #/#: What could superintelligence do?
Reading: *Superintelligence* by Nick Bostrom pages 91–114
- Wed #/#: Doom and avoiding doom
Reading: *Superintelligence* pages 115–144

Week 12: Practical & Ethical Questions III: The Control Problem

Mon #/#: Castes of supertintelligence?
Reading: Superintelligence pages 145–184

Wed #/#: Values
Reading: Superintelligence pages 185–207

Week 13: Practical & Ethical Questions IV: The Control Problem Continued

Mon #/#: Choosing the Criteria for Choosing
Reading: Superintelligence pages 209–227

Wed #/#: The strategic picture
Reading: Superintelligence 228–253

Week 14: Conclusion & Review

Mon #/#: Conclusion of *Superintelligence* & review
Reading: Superintelligence pages 255–260
ALSO: Prepare by looking at the study guide for the final test

Wed #/#: Review & Q&A
How to prepare: look at study guide and prepare questions

4. Assignments and Grading

There will be three categories of assignments you will be graded on:

a. ATTENDANCE & PARTICIPATION (30% of total grade)

Attendance will be taken at the beginning of class every session. Attendance + participation constitutes 15% of your total grade. Participation will be determined by the instructor's impression of how often you ask questions or answer questions (the Q&A and review sessions will be a good time to make up for participation). Participation counts for exactly 7.5% of your grade and attendance accounts for exactly 7.5% of your grade. You are granted *two free* absences without excuse. After you have missed two classes without excuse, you must contact the professor to request an absence.

TARDINESS: Being significantly late to class will generally count as an absence. If you show up a few minutes late and miss attendance, talk to the instructor after class to have your attendance recorded.

b. WEEKLY QUESTIONS (10% of total grade)

Each week you will be asked to choose one of the readings to submit questions about. Questions will be due the night before class at midnight. For more about this assignment, see the “Substantive Questions” handout on Canvas.

c. TESTS (60% of total grade)

There will be two tests for this class: one midterm and one final. The midterm test will be composed of a T/F section and a multiple choice section, and it will be taken in the testing center. The final test will have one T/F section, one multiple choice section, and one long-answer section. The final test will be taken during finals week. Both tests will be accompanied by study guides. We will have one week of in-class review before the final test.

GRADING SCALE: This course will use the following grading scale.

A (4.0):	94-100
A- (3.7):	90-94
B+(3.3):	87-90
B (3.0):	83-87
B- (2.7):	80-83
C+ (2.3):	77-80
C (2.0):	73-77
C- (1.7):	70-73
D (1.0):	60-70
F (0.0):	0-6

LATE SUBMISSIONS: Late assignments will, generally, be considered with a penalty to be determined by the instructor. The later the assignment is submitted, generally, the greater the penalty will be.

RE-DOs: Case-by-case basis. Talk to the professor about redoing assignments.

CHATGPT and AI: The use of ChatGPT or generative AI is strictly prohibited. As such, all assignments (**with the exception of weekly questions**) will be written in class or in the testing center. The use of ChatGPT or AI in class will result in a grade of zero on the assignment it was used for.

EXTRA CREDIT: There will be a number of extra credit options available throughout the semester. There will be one extra credit question at the end of both tests. Students who regularly ask questions will likely receive some extra credit at the end of the course.

5. Academic Honesty/Integrity

Academic integrity is honest, truthful and responsible conduct in all academic endeavors. The mission of Saint Louis University is “the pursuit of truth for the greater glory of God and for the service of humanity.” Accordingly, all acts of falsehood demean and compromise the corporate endeavors of teaching, research, health care, and community service through which SLU fulfills its mission. The University strives to prepare students for lives of personal and professional integrity, and therefore regards all breaches of academic integrity as matters of serious concern. The full University-level Academic Integrity Policy can be found on the Provost's Office website at: <https://www.slu.edu/provost/policies/academic-and-course/academic-integrity-policy.pdf>

Additionally, each SLU College, School, and Center has its own academic integrity policies, available on their respective websites.

6. Disability Accommodations

Students with a documented disability who wish to request academic accommodations must formally register their disability with the University. Once successfully registered, students also must notify their course instructor that they wish to use their approved accommodations in the course.

Please contact the Center for Accessibility and Disability Resources (CADR) to schedule an appointment to discuss accommodation requests and eligibility requirements. Most students on the St. Louis campus will contact CADR, located in the Student Success Center and available by email at accessibility_disability@slu.edu or by phone at 314.977.3484. Once approved, information about a student's eligibility for academic accommodations will be shared with course instructors by email from CADR and within the instructor's official course roster. Students who do not have a documented disability but who think they may have one also are encouraged to contact to CADR. Confidentiality will be observed in all inquiries.

Note: due to accreditation requirements, regulatory differences, and/or location-specific resources, the School of Law, the School of Medicine, and SLU Madrid have their own standard language for syllabus statements related to disability accommodations. Faculty in those units should seek guidance for syllabus requirements from their dean's office.

7. Title IX

Saint Louis University and its faculty are committed to supporting our students and seeking an environment that is free of bias, discrimination, and harassment. If you have encountered any form of discrimination on the basis of sex, including sexual harassment, sexual assault, stalking, domestic or dating violence, we encourage you to report this to the University. Discrimination on the basis of sex includes discrimination on the basis of assigned sex at birth, sex characteristics, pregnancy and pregnancy related conditions, sexual orientation and gender identity. If you speak with a faculty member about an incident that involves a Title IX matter, that faculty member must notify SLU's Title IX Coordinator that you shared an experience relating to Title IX. This is true even if you ask the faculty member not to disclose the incident. The Title IX Coordinator will then be available to assist you in understanding all of your options and in connecting you with all possible resources on and off campus.

If you are pregnant or experiencing a pregnancy related condition, the Title IX Coordinator can assist you in understanding your rights and options as well as provide supportive measures.

Anna Kratky is the Title IX Coordinator at Saint Louis University (DuBourg Hall, room 36; anna.kratky@slu.edu; 314-977-3886). If you wish to speak with a confidential source, you may contact the counselors at the University Counseling Center at 314-977-TALK or make an anonymous report through SLU's Integrity Hotline by calling 1-877-525-5669 or online at <http://www.lighthouse-services.com/slu>. To view SLU's policies, and for resources, please visit the following web addresses: <https://www.slu.edu/about/safety/sexual-assault-resources/index.php>.

Note: due to accreditation requirements, regulatory differences, and/or location-specific resources, the School of Law, the School of Medicine, and SLU Madrid have their own standard language for syllabus statements related to Title IX. Faculty in those units should seek guidance for syllabus requirements from their dean's office.

8. Outcomes

The purpose of this course is to provide an introduction to some of the most historically important philosophical texts, ideas, and thinkers as well as to the distinctive activity of philosophy itself. Over the course of the semester, students will:

1. Learn what great thinkers have to say in response to enduring philosophical questions and how their answers represent philosophically and dialectically distinct answers to such questions. Students will consider a range of philosophical voices and perspectives--including classic texts from Plato and Aristotle, texts from thinkers in the Catholic tradition, and texts by thinkers historically underrepresented in the western canon. In reading and discussing such thinkers and texts students will be encouraged to consider and assess both the philosophical merits and costs of the various views they encounter.
2. Learn to *explain, analyze, and evaluate* arguments for philosophical answers to ultimate questions. Students will learn to give a close reading of philosophical texts with a view toward being able to explain and analyze the arguments they contain. By the end of the semester, students should be able to (i) identify the key premises, assumptions, and conclusions present in the texts they encounter and (ii) explain the content of the arguments premises in a way that (iii) illuminates how these premises support an inference to the conclusion. Finally, students will learn to evaluate the arguments they encounter by, among other things, assessing them according to logical notions of validity and soundness.
3. Begin to shape their own answers to ultimate questions by applying distinctively logical and philosophical ways of thinking to them. To this end, students will learn to raise philosophical questions, to identify unexamined presuppositions, and to reflect on their own assumptions and commitments. Students will also work at philosophically articulating their own views by learning how to offer evidence, arguments, and responses to objections to their own answers to ultimate questions.

4. Learn to articulate their own worldview and practices in dialogue with different answers to ultimate questions. Students will be provided with opportunities to reflect on the way in which their own commitments, values, and practices may overlap with or diverge from the perspectives they encounter in their readings, in lectures, and in dialogue with one another.

Ultimate Questions: Philosophy is one of 19 Core Components. The University Core SLO(s) that this component is designed to intentionally advance are listed below:

University Core Student Learning Outcomes

The Core SLO(s) that this component is intentionally designed to advance are:

SLO 1: Examine their actions and vocations in dialogue with the Catholic, Jesuit tradition

SLO 2: Integrate knowledge from multiple disciplines to address complex questions

SLO 3: Assess evidence and draw reasoned conclusions

Additionally, the Core Component-level Student Learning Outcomes are listed below:

Component-level Student Learning Outcomes

Students who complete this course will be able to:

Explain and evaluate arguments for philosophical answers, including those in the Catholic tradition, to ultimate questions concerning the nature of humanity, reality and God, knowledge, and/or the good life

Apply an array of analytic methods, conceptual tools, logical principles, and other resources to their own inquiry into ultimate questions

Articulate and reflect on their own worldview and practices in dialogue with different answers to ultimate questions, including answers in the Catholic tradition