

Last Updated: August 20, 2024

AI AND THE FUTURE OF LAW
(2 or 3 credits)

Professor Jonathan H. Choi
Fall 2024

Email: jonchoi@law.usc.edu
Administrative Assistant: Maria Chan, mchan@law.usc.edu
Class Meetings: Tuesday, 2 PM – 3:50 PM
Class Location: Room 2
Grade scale: Numerical grading by default, may elect CR/D/F grading if not attempting to satisfy UDWR

1. Course Description

This course examines the impact of artificial intelligence (AI) on the future of law. Its content is divided between (a) practical content teaching students how to use AI effectively in legal research, writing, and practice, and (b) theoretical content discussing the legal implications of AI and regulatory approaches. Students will read work on the cutting edge of AI research and participate in presentations from experts about current topics.

The course is a two (2) unit course with a three (3) unit option.

2. Learning Objectives

By the end of this course, students will be able to:

1. Use AI in legal research and writing.
2. Discuss the implications of AI for law and the legal profession.
3. Understand the potential to misuse AI tools and develop an understanding for how they should be regulated.

3. Required Resources

To complete the first three assignments for this course, you will need access to a large language model, like [ChatGPT Plus](#) or [Claude Pro](#). I recommend subscribing to one of these services, although Claude Sonnet 3.5 is currently free and of good quality (but for legal writing Claude Opus 3 is often better). Both ChatGPT Plus and Claude Pro cost \$20 per month (plus tax), and you will only need access for about one month. Both provide access to cutting-edge models, and I have used both extensively. Because each has its own strengths and weaknesses, it may even be useful to subscribe to both.

However, no subscription is required, and you will be judged based on your interactions with the large language models (LLMs) rather than the quality of the LLM's output on its own. If you are familiar with programming and have API access, you can use GPT or Claude directly instead of subscribing, although this will generally be much more work.

If you are unable to sign up for either Claude or ChatGPT, please contact me. In addition, if you believe that another LLM performs as well or better than ChatGPT or Claude, let me know and we can discuss alternatives.

4. Brightspace

I'll use Brightspace to post course materials and announcements. Please ensure that you receive announcements through Brightspace, as I will post all course-related announcements there. You'll also use Brightspace to submit response papers, starting in Week 2.

5. Panel System, Discussion Leaders, and Class Participation

Each week, there will be 2 discussion leaders, and approximately half the class will be "on call" for discussion. Although you will be expected to read all of the readings for the class, you will be particularly expected to participate in sessions where you are "on call", and I may call on people Socratically. The schedule for this is as follows:

Class 1: No response paper required, no one is on call

Classes 2-3: Everyone must write a response paper, everyone is on call

Classes 4-12: You need only write a response paper if you are on call that day

Prior to Class 4, you will be asked to rank your preferences (if any) among the topics and topics will be assigned based on a matching algorithm intended to maximize satisfaction with assignments.

Some classes will involve in-class exercises to improve your understanding of and familiarity with AI tools. These exercises may require you to have a laptop and access to LLMs. For example, we will practice "prompt engineering" to improve results from AI language models by actually using different prompts and discussing the results.

In addition to the class being on call, each day I will assign two (2) or three (3) discussion leaders per class (again assigned based on the matching algorithm). If you are a discussion leader for the day, you will need to prepare five (5) questions / topics to discuss in advance, and you will be expected to start our discussion for the day. Discussion leaders will be told their assignments in advance and should coordinate with their co-leader(s). When you are a discussion leader, you should send me the questions / topics you plan to discuss in advance; the due date for this is the same as for response papers and can be submitted along with your response paper.

Class attendance and active participation are required. You are expected to attend class regularly regardless of whether you are a discussion leader / on call or not, complete all of the assignments described above, and be an active member of the class during discussion periods. Because this is a seminar, class attendance is very important and is a core part of your grade. You should not sign up for this course unless you anticipate being able to attend all of the classes. Classes will not be recorded and unexcused absences may affect your grade. Repeated unexcused absences may result in dismissal from the course.

If you can't attend a scheduled class session at all, please email me in advance to let me know. I'm happy to discuss any additional accommodations that might be required in light of scheduling challenges this term.

6. Response Papers

Each class when you are required to submit a response paper, you should write approximately two (2) single-spaced pages reflecting on one or more issues raised by the readings for a given week. The length reflects conventional formatting, e.g. size 12 Times New Roman with 1-inch margins; you may actually format the paper however you wish, for example you may submit four (4) double-spaced pages. You may also submit a longer response paper if you wish; there is no limit on total length. The response papers will be graded on the quality of your reasoning and writing and the extent to which you consider, respond to, and elaborate on the readings for the day. You are encouraged not just to recite the main points of the reading but to compose original reactions to the topics of each day's reading.

Some weeks have specific assignments, and some weeks have only readings listed; when no specific assignments are listed, please respond in general to the issues raised by the readings, and when assignments are listed, you must complete the assignment but you do not need to discuss the assignment in your response paper unless you found the results particularly interesting. To be accepted, each response paper must be submitted by 11 a.m. each day we have class.

7. Grading Breakdown

For students taking the course for 2 credits, grades will be based on the following:

- 75%: Class participation, response papers, and assignments
- 25%: Class presentation

Class Participation, Response Papers, and Assignments (75%). This is a class in which active participation with classmates and meaningful engagement with class materials is expected. **Any** unexcused absence from class may result in a reduction in the overall course grade.

Your grade for participation, response papers, and assignments is further broken down into:

- 15%: Participation as discussion leader
- 35%: Response papers
- 35%: Participation when on call (not as discussion leader)
- 15%: Participation when not on call

This category includes regular attendance, thorough preparation of reading and written assignments (including the reaction papers); scholarly, and constructive critique of readings; and avoidance of unexcused absences. Each class that you are on call, your in-class participation and your response paper will be graded as either unsatisfactory, good, or outstanding, with the majority graded as good. Discussion leaders will be numerically graded on the basis of how well they lead discussion and how deeply they engage with the subject. **It is important to submit response papers when required, regardless of whether you will be able to attend or not—failure to submit a response paper will result in a grade of unsatisfactory.**

Class Presentation (25%). The class presentation will be a short (approximately 10-minute) presentation on any topic relating to law and artificial intelligence. Students taking the course for 3 credits are expected to present the claim of their final paper, and other students may select a topic of their choice. You should plan to spend approximately 8 minutes on the presentation itself, leaving approximately 2

minutes for questions (but you can leave more time for questions if you wish). Presentations will be graded on depth of research (40%), clarity and quality of the presentation (35%), and originality (25%). Originality means making points that were not simply made in the class reading, and could consist of original descriptive observations, normative arguments, predictions, etc. Students should preview the topic of their final presentation with me well in advance. You may use Powerpoint slides in your presentation if you would like.

Students who are taking the course for 3 credits submit a final paper and therefore will be graded on the following basis:

- 50%: Class participation, response papers, and assignments
- 20%: Class presentation
- 30%: Final paper

Grading methods and expectations are otherwise the same as for students taking the class for 2 credits. All final papers are due at noon on May 9, 2024. Students writing a final paper should preview the topic of the paper with me in advance. Deadlines are taken very seriously. Late submissions yield grade reductions.

8. Final Papers

Students taking the course for 3 credits will be required to write an approximately double-spaced twenty (20) page (~5000 word) scholarly paper on a topic of their choosing, related to the class, and approved by the professor. You may, but are not required, to share early drafts of your final papers with me to receive feedback. The paper should include citations where applicable in any citation format you prefer (however, most students default to Bluebook format).

For students using the 3-credit option, generally your class presentation will be on the same topic as your final paper. If you wish for the paper to satisfy the Upper Division Writing Requirement, your paper will need to be at least twenty (20) double-spaced pages (~5000 words) long and you cannot take the course on a CR/D/F basis.

Final papers will be graded on depth of research (40%), originality (30%), and clarity and quality of writing (30%). Originality means making points that were not simply made in the class reading, and could consist of original descriptive observations, normative arguments, predictions, etc.

9. Feedback and Future Changes

I am keen to receive feedback about whether the readings for and topics discussed in this course are helpful. I invite you to submit feedback about the course by either emailing me or emailing my assistant Maria Chan, who can anonymize feedback. Based on this feedback, and because the material we are covering is changing so quickly, there may be changes to the syllabus as the semester progresses. I will attempt to provide at least 2 classes' worth of notice when making any changes.

10. Class Location and Make-Up Classes

It is possible some classes may be held remotely over Zoom. For those classes, we'll use the following Zoom meeting room:

<https://usc.zoom.us/j/95408559602?pwd=GdjfBKorjusgvLUfjlcxgnOkyxJ6Qb.1>

Meeting ID: 954 0855 9602

Passcode: 470666

Although I'll usually try to conduct class remotely rather than adjusting the schedule, I'll occasionally cancel class and schedule in-person make-up classes at a different time. I do not currently anticipate rescheduling any classes to Zoom and will let you know in advance if I do so. When we have guest speakers presenting over Zoom, we will still meet in our regular room and will have in-person discussion in addition to the guest speaker.

Class will be cancelled on November 19. I will let you know when the makeup class occurs.

11. AI Policy

Some assignments in this course will require the use of AI tools, like ChatGPT or Claude. Unless otherwise stated, you are **encouraged** to use AI tools to complete all assignments. You are also encouraged to **share** anything that you learn about AI tools in classroom discussion.

When submitting any written materials (for example, response papers), you must attribute to any AI tool you have used any words, paraphrases, or ideas you have used or derived from it. This attribution should include a footnote that includes the exact query you submitted to ChatGPT and a quotation, paraphrase, or description of the relevant portion of ChatGPT's response. For guidance as to citation format, you may refer to any footnote in a published law review article citing the results of a Westlaw or Google search.

This is an example of a compliant footnote: "A query to ChatGPT 'Are you sentient?' on Jan. 9, 2024 resulted in the response, inter alia, that 'There is ongoing debate among philosophers, scientists, and researchers about this question, and it is likely to remain a topic of discussion for the foreseeable future.' See Appendix C (providing full text of this conversation)." In this case, you would also provide an appendix (not included in the word or page count) providing the full text of the conversation that resulted in the cited response, including any previous queries and responses that ChatGPT is likely to have incorporated into its analysis.

12. Technological Proficiency and Hardware/Software Required

For some assignments, you will be required to use a computer and access online AI tools. You may also be required to purchase a subscription to ChatGPT Plus in order to complete coursework.

13. Course Content Distribution and Synchronous Session Recordings Policies

As is the default for courses at USC Law, classes will be audio recorded. The recordings will not be posted online or distributed except with my permission in the event of excuse absences.

USC has policies that prohibit recording and distribution of any synchronous and asynchronous course content outside of the learning environment.

Recording a university class without the express permission of the instructor and announcement to the class, or unless conducted pursuant to an Office of Student Accessibility Services (OSAS) accommodation. Recording can inhibit free discussion in the future, and thus infringe on the academic freedom of other students as well as the instructor. ([Living our Unifying Values: The USC Student Handbook](#), page 13).

Distribution or use of notes, recordings, exams, or other intellectual property, based on university classes or lectures without the express permission of the instructor for purposes other than individual or group study. This includes but is not limited to providing materials for distribution by services publishing course materials. This restriction on unauthorized use also applies to all information, which had been distributed to students or in any way had been displayed for use in relationship to the class, whether obtained in class, via email, on the internet, or via any other media. ([Living our Unifying Values: The USC Student Handbook](#), page 13).

14. Office Hours

I'll hold office hours in my office (Room 468) after each class. Please let me know in advance if you'd like to attend. Alternatively, I'm happy to schedule an appointment at another time.

15. Academic Integrity

The University of Southern California is a learning community committed to developing successful scholars and researchers dedicated to the pursuit of knowledge and the dissemination of ideas. Academic misconduct, which includes any act of dishonesty in the production or submission of academic work, comprises the integrity of the person who commits the act and can impugn the perceived integrity of the entire university community. It stands in opposition to the university's mission to research, educate, and contribute productively to our community and the world.

All students are expected to submit assignments that represent their own original work, and that have been prepared specifically for the course or section for which they have been submitted. You may not submit work written by others or "recycle" work prepared for other courses without obtaining written permission from the instructor(s).

Other violations of academic integrity include, but are not limited to, cheating, plagiarism, fabrication (e.g., falsifying data), collusion, knowingly assisting others in acts of academic dishonesty, and any act that gains or is intended to gain an unfair academic advantage.

The impact of academic dishonesty is far-reaching and is considered a serious offense against the university. All incidences of academic misconduct will be reported to the Office of Academic Integrity and could result in outcomes such as failure on the assignment, failure in the course, suspension, or even expulsion from the university.

For more information about academic integrity see [the student handbook](#) or the [Office of Academic Integrity's website](#), and university policies on [Research and Scholarship Misconduct](#).

Please ask your instructor if you are unsure what constitutes unauthorized assistance on an exam or assignment, or what information requires citation and/or attribution.

16. Students and Disability Accommodations

USC welcomes students with disabilities into all of the University's educational programs. The Office of Student Accessibility Services (OSAS) is responsible for the determination of appropriate accommodations for students who encounter disability-related barriers. Once a student has completed the OSAS process (registration, initial appointment, and submitted documentation) and accommodations are determined to be reasonable and appropriate, a Letter of Accommodation (LOA) will be available to generate for each course. The LOA must be given to each course instructor by the student and followed up with a discussion. This should be done as early in the semester as possible as accommodations are not retroactive. More information can be found at osas.usc.edu. You may contact OSAS at (213) 740-0776 or via email at osasfrontdesk@usc.edu.

17. Support Systems

[Counseling and Mental Health](#) - (213) 740-9355 – 24/7 on call

Free and confidential mental health treatment for students, including short-term psychotherapy, group counseling, stress fitness workshops, and crisis intervention.

[988 Suicide and Crisis Lifeline](#) - 988 for both calls and text messages – 24/7 on call

The 988 Suicide and Crisis Lifeline (formerly known as the National Suicide Prevention Lifeline) provides free and confidential emotional support to people in suicidal crisis or emotional distress 24 hours a day, 7 days a week, across the United States. The Lifeline is comprised of a national network of over 200 local crisis centers, combining custom local care and resources with national standards and best practices. The new, shorter phone number makes it easier for people to remember and access mental health crisis services (though the previous 1 (800) 273-8255 number will continue to function indefinitely) and represents a continued commitment to those in crisis.

[Relationship and Sexual Violence Prevention Services \(RSVP\)](#) - (213) 740-9355(WELL) – 24/7 on call

Free and confidential therapy services, workshops, and training for situations related to gender- and power-based harm (including sexual assault, intimate partner violence, and stalking).

[Office for Equity, Equal Opportunity, and Title IX \(EEO-TIX\)](#) - (213) 740-5086

Information about how to get help or help someone affected by harassment or discrimination, rights of protected classes, reporting options, and additional resources for students, faculty, staff, visitors, and applicants.

[Reporting Incidents of Bias or Harassment](#) - (213) 740-5086 or (213) 821-8298

Avenue to report incidents of bias, hate crimes, and microaggressions to the Office for Equity, Equal Opportunity, and Title for appropriate investigation, supportive measures, and response.

[The Office of Student Accessibility Services \(OSAS\)](#) - (213) 740-0776

OSAS ensures equal access for students with disabilities through providing academic accommodations and auxiliary aids in accordance with federal laws and university policy.

[USC Campus Support and Intervention](#) - (213) 740-0411

Assists students and families in resolving complex personal, financial, and academic issues adversely affecting their success as a student.

[Diversity, Equity and Inclusion](#) - (213) 740-2101

Information on events, programs and training, the Provost's Diversity and Inclusion Council, Diversity Liaisons for each academic school, chronology, participation, and various resources for students.

[USC Emergency](#) - UPC: (213) 740-4321, HSC: (323) 442-1000 – 24/7 on call

Emergency assistance and avenue to report a crime. Latest updates regarding safety, including ways in which instruction will be continued if an officially declared emergency makes travel to campus infeasible.

[USC Department of Public Safety](#) - UPC: (213) 740-6000, HSC: (323) 442-1200 – 24/7 on call

Non-emergency assistance or information.

[Office of the Ombuds](#) - (213) 821-9556 (UPC) / (323-442-0382 (HSC)

A safe and confidential place to share your USC-related issues with a University Ombuds who will work with you to explore options or paths to manage your concern.

[Occupational Therapy Faculty Practice](#) - (323) 442-2850 or otfp@med.usc.edu

Confidential Lifestyle Redesign services for USC students to support health promoting habits and routines that enhance quality of life and academic performance.

COURSE OUTLINE

The following list is preliminary and may be updated throughout the semester. Each topic corresponds to a full class. Optional further readings are intended as starting points for researching your final presentation and/or paper. You need not read them before class, and we will not discuss them in class.

This list of readings is provisional and may be updated as the semester progresses.

1. Introduction to AI

- Key issues:
 - What is AI?
 - A brief history of AI
 - How does AI work?
 - Why is AI relevant to the legal profession?
- Reading:
 - Helen Toner, *What Are Generative AI, Large Language Models, and Foundation Models?*, CSET (May 12, 2023), <https://cset.georgetown.edu/article/what-are-generative-ai-large-language-models-and-foundation-models/>
 - A good high-level overview of basic AI concepts.
 - Richard Ngo, *Visualizing the Deep Learning Revolution*, MEDIUM (Jan. 5, 2023), <https://medium.com/@richardcngo/visualizing-the-deep-learning-revolution-722098eb9c5>
 - A deeper dive into the capabilities of specific kinds of AI. You do not need to deeply engage with the examples (especially the technical ones), just get a sense of the kinds of things AI can do in specific narrow domains.
 - *Artificial General Intelligence*, WIKIPEDIA, https://en.wikipedia.org/wiki/Artificial_general_intelligence
 - You can skip the section titled “Brain Simulation”
- Assignments:
 - First, complete the pre-class survey at: <https://forms.gle/NxRB9Ctpow5YFPyA6>
 - Think about these questions: To what extent, if at all, can computers “understand” human language? What types of analysis (*i.e.*, analysis that is typically performed in legal work) are simulated by the “narrow” techniques of AI? What are the benefits and drawbacks of the various definitions of artificial general intelligence?
- Optional further reading:
 - [To read: 669 – 702] David Lehr & Paul Ohm, *Playing with the Data: What Legal Scholars Should Learn About Machine Learning*, 51 U.C. DAVIS L. REV. 653 (2017), <https://lawreview.law.ucdavis.edu/archives/51/2/playing-data-what-legal-scholars-should-learn-about-machine-learning>.
 - Good description of how machine learning actually works.
 - Timothy B. Lee & Sean Trott, *Large Language Models, Explained with a Minimum of Math and Jargon*, <https://seantrott.substack.com/p/large-language-models-explained>
 - Despite the title, this contains a fair amount of math and jargon
 - *How Transformers Work*, TOWARDS DATA SCIENCE, <https://towardsdatascience.com/transformers-141e32e69591>
 - Gets quite technical

- *How ChatGPT Actually Works*, ASSEMBLY AI <https://www.assemblyai.com/blog/how-chatgpt-actually-works>
 - Additional description of the specific model used in ChatGPT

2. Prompt Engineering for Lawyers

- Key Issues:
 - What is prompt engineering?
 - How can we prompt AI models to produce the best results?
 - Using AI tools in everyday life
 - What are the strengths and weaknesses of AI?
- Reading:
 - Daniel Schwarcz & Jonathan H. Choi, *AI Tools for Lawyers: A Practical Guide*, 108 MINN. L. REV. HEADNOTES 1 (2023), <https://minnesotalawreview.org/article/ai-tools-for-lawyers-a-practical-guide/>
 - “Prompt Engineering”, OPENAI (2023), <https://platform.openai.com/docs/guides/prompt-engineering/six-strategies-for-getting-better-results>
 - Some tactics are intended for software engineers and should be skipped, namely “Use intent classification to identify the most relevant instructions for a user query” and all the tactics under “Use external tools”
 - Note that several of the prompts ask the model to output in JSON format. Don’t worry about this. (For the curious: JSON format is like this: [{"a": "b"}, {"a2": "b2"}], ...) and is a good way to get structured data from LLMs for algorithms to process.)
- Assignments:
 - Create a ChatGPT Plus or Claude Pro account. Spend ten minutes conversing with the LLM on topics of your choosing. Save the prompts you used and the outputs you received for classroom discussion, and mention anything you learned from these prompts in your weekly response paper.
 - Some inspiration for interesting LLM prompts: *Awesome ChatGPT Prompts*, GITHUB, <https://github.com/f/awesome-chatgpt-prompts/>
 - If relevant, in your weekly response paper, describe any ways that you currently use LLMs for your schoolwork or your everyday life. (You may address any combination of this question and discussing the assignment above or the readings.)
- Optional further reading:
 - John Nay et al., *Large Language Models as Tax Attorneys: A Case Study in Legal Capabilities Emergence*, PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY (2023), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4476325
 - This paper applies a RAG-based approach to tax law, finding its potential to be limited
 - *OpenAI Cookbook*, OPENAI, <https://cookbook.openai.com/>
 - Mostly for software engineers, tends to be technical
 - *Microsoft Bing Chat*, LIFEARCHITECT, https://lifearchitect.ai/bing-chat/?utm_source=substack&utm_medium=email
 - Examples of system prompts used to convert OpenAI’s GPT models to various Microsoft products

- Fangyi Yu, Lee Quartey & Frank Schilder, *Legal Prompting: Teaching a Language Model to Think Like a Lawyer*, <https://arxiv.org/abs/2212.01326>
 - Suggestions about how to use IRAC prompting to improve model performance; very technical

3. AI in Legal Research and Writing

- Guest speaker: Valerie McConnell (Casetext) via Zoom (~30 minutes)
- Key Issues:
 - Can legal research be automated?
 - What natural language processing tools can improve the reliability of legal analysis by machines?
 - How reliable is AI-generated research?
- Reading:
 - Kim Martineau, *What Is Retrieval-Augmented Generation?*, IBM, <https://research.ibm.com/blog/retrieval-augmented-generation-RAG>
- Assignment:
 - Take a research assignment you completed for a different class. (If you do not have an appropriate research assignment, let me know and I will provide you with one.) Use ChatGPT/Claude and the prompt engineering techniques you learned in the previous class to complete the assignment. Then use Co-Counsel, Lexis AI, or both to complete the assignments.
 - NB: You do not need to give me the output produced by Co-Counsel or Lexis AI, your response paper only needs to include reflections on the strengths and weaknesses of each tool.
 - Consider the following in your response paper and for classroom discussion:
 - Did you save time using ChatGPT/Claude? Did you save time using Co-Counsel/Lexis AI?
 - Do you think using these tools increased or decreased the quality of your work?
 - Were there any advantages in conducting research the “old-fashioned” way?
 - What can LLMs and LLM-assisted tools do that Westlaw Edge and LexisNexis Advance (the traditional search products produced by Westlaw and Lexis) cannot? Vice versa?
 - Overall, which did you find more useful?
 - Co-Counsel and Lexis AI produce memos similar to the work a junior associate at a law firm might produce, although both companies explicitly caveat that their products cannot yet operate without supervision by a human lawyer. Are they being excessively conservative, or are there real concerns?
 - What is stopping AI from being able to write a legal research memo as well or better than a human lawyer?
 - Come prepared with a question or two to ask Valerie.
- Optional further reading:
 - Anton Korinek, *Language Models and Cognitive Automation for Economic Research*, <https://www.nber.org/papers/w30957>
 - Advice on using language models for academic research

4. AI and the Legal Profession

- Key Issues:
 - Predictions for the future of AI in law
 - Challenges and opportunities for lawyers
 - Adapting to the changing legal landscape
- Reading:
 - Jonathan H. Choi, Amy Monahan & Daniel Schwarcz, *Lawyering in the Age of Artificial Intelligence* (working draft),
https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4626276
 - Conducts research on the effect of AI-human interaction in lawyering tasks
 - Josh Blackman, *Robot, Esq.* (working draft),
https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2198672
 - A short paper that raises a number of important questions in the application of AI to the legal profession. You are encouraged to dig further into any of these questions that particularly interested you.
- Assignment:
 - Please think about the following issues:
 - Is the automation of litigation work or legal adjudication a good thing for the U.S. legal system? For the legal profession? For society?
 - Do you believe that, within the next 10 years, there will be “AI” technology that is capable of performing all the relevant tasks of a litigation attorney? Of reliably adjudicating legal claims?
 - Do you believe that ChatGPT/Claude/Co-Counsel/Lexis AI meaningfully bring us closer to a true legal AI (defined as an AI system with the capabilities of a competent attorney)?
 - Are there particular tasks you are familiar with from your legal background (law schools, summer internships, previous jobs) that you think are particularly susceptible / impervious to AI? Why?
- Optional further reading:
 - Amy Salyzyn, *AI and Legal Ethics*, in *ARTIFICIAL INTELLIGENCE AND THE LAW IN CANADA* (2021)
https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3733952
 - A discussion of ethical issues lawyers will face when using AI, from a Canadian perspective
 - Twitter thread on AI scaling laws (technical):
https://twitter.com/ai_pub/status/1564746660030267392
 - A technical discussion of an important result
 - Tyna Eloundou, Sam Manning, Pamela Mishkin & Daniel Rock, *GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models*,
<https://arxiv.org/abs/2303.10130>

5. AI Safety and Alignment

- Key Issues:
 - What is alignment, and why does it matter?
 - Why is alignment such a difficult problem?
 - What can we learn from alignment issues in other contexts?
 - Risks of misaligned or rogue AI systems

- How difficult is it to understand an AI model's thinking and motivations?
- Why might safety concerns about AI be overblown? Why might they be underappreciated?
- Reading:
 - Anton Korinek & Avital Balwit, *Aligned With Whom?*, BROOKINGS, <https://www.brookings.edu/wp-content/uploads/2022/05/Aligned-with-whom-1.pdf>
 - Jack Clark & Dario Amodei, *Faulty Reward Functions in the Wild*, OPENAI (Dec. 21, 2016), <https://openai.com/research/faulty-reward-functions>
 - Short, illustrates how AI can generate creative and unintended methods to optimize imperfectly thought-through reward functions
 - Jan Leike et al., *Our Approach to Alignment Research*, OPENAI (Aug. 24, 2022), <https://openai.com/blog/our-approach-to-alignment-research>
 - Anthropic, *Mapping the Mind of a Large Language Model* (May 21, 2024), <https://www.anthropic.com/research/mapping-mind-language-model>
 - This short piece discusses some of the challenges in understanding how a language model works and discusses cutting-edge research on explaining model behavior
- Optional reading:
 - Katja Grace, *Counterarguments to The Basic AI x-Risk Case*, AI IMPACTS (Aug. 31, 2022), <https://aiimpacts.org/counterarguments-to-the-basic-ai-x-risk-case/>
 - Betty L. Hou & Brain P. Green, *A Multilevel Framework for the AI Alignment Problem*, MARKKULA CENTER, (Jul. 25, 2022), <https://www.scu.edu/ethics/focus-areas/technology-ethics/resources/a-multilevel-framework-for-the-ai-alignment-problem/>
 - OPENAI, *Language Models Can Explain Neurons in Language Models* (May 9, 2023), <https://openai.com/research/language-models-can-explain-neurons-in-language-models>
 - NICK BOSTROM, *SUPERINTELLIGENCE: PATHS, DANGERS, AND STRATEGIES* (2014) (a classic text about the risks associated with AI)
 - Andrew Keane Woods, *Robophobia*, 93 U. COLO. L. REV. 51 (2022), <https://lawreview.colorado.edu/print/robophobia/>
 - [Read 33-44] Toby Ord, *Lessons from the Development of the Atomic Bomb*, CENTRE FOR THE GOVERNANCE OF AI, 33-44 (Nov. 2022), https://cdn.governance.ai/Ord_lessons_atomic_bomb_2022.pdf

6. Robot Rights

- Key Issues:
 - Under what circumstances should we give AIs legal personhood?
 - What legal protections might be granted to AI?
- Reading:
 - Katherine B. Forrest, *The Ethics and Challenges of Legal Personhood for AI*, 133 YALE LAW JOURNAL FORUM 1175 (2024), <https://www.yalelawjournal.org/forum/the-ethics-and-challenges-of-legal-personhood-for-ai>
- Optional Reading:
 - TED CHIANG, *THE LIFECYCLE OF SOFTWARE OBJECTS*
 - ISAAC ASIMOV, *I, ROBOT*
 - Peter Salib & Simon Goldstein, *AI Rights for Human Safety* (working paper), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4913167

7. Regulating AI

- Key Issues:
 - How should we regulate AI to ensure safety and trustworthiness?
 - Which arms of government (courts, Congress, agencies, the executive branch) have the institutional competency to adequately respond to the development of AI?
- Reading:
 - [Read 410-425] Rayn Calo, *Artificial Intelligence Policy: A Primer and Roadmap*, 51 U.C. DAVIS L. REV. 399 (2017), https://lawreview.law.ucdavis.edu/issues/51/2/symposium/51-2_Calo.pdf
 - A discussion of key questions in AI regulation
 - [Read 373-400] Matthew U. Scherer, *Regulating the Artificial Intelligence Systems: Risks, Challenges, Competencies and Strategies*, 29 HARV. J.L. & TECH. 353 (2016), <http://jolt.law.harvard.edu/articles/pdf/v29/29HarvJLTech353.pdf>
 - An analysis of how different institutions might differ in their competency at regulating AI
 - [Read 635-645, 684-693] Alicia Solow-Niederman, *Administrating Artificial Intelligence*, 93 S. CAL. L. REV. 633 (2020), https://southerncalifornialawreview.com/wp-content/uploads/2020/09/SolowNiederman_website.pdf
 - This article discusses how regulating AI can be difficult and suggests ways to incorporate policy preferences into code
- Optional further reading:
 - Tom Wheeler, *The Three Challenges of AI Regulation*, BROOKINGS (Jun. 15, 2023), <https://www.brookings.edu/articles/the-three-challenges-of-ai-regulation/>
 - THE WHITE HOUSE, *President Biden Issues Executive Order on Safe, Secure, and Trustworthy Artificial Intelligence* (2023), <https://www.whitehouse.gov/briefing-room/statements-releases/2023/10/30/fact-sheet-president-biden-issues-executive-order-on-safe-secure-and-trustworthy-artificial-intelligence/>
 - John Villasenor, *The Problems With a Moratorium on Training Large AI Systems*, BROOKINGS (Apr. 11, 2023), <https://www.brookings.edu/blog/techtank/2023/04/11/the-problems-with-a-moratorium-on-training-large-ai-systems/>
 - Jess Whittlestone & Jack Clark, *Why and How Governments Should Monitor AI Development*, 3 (2021), <http://arxiv.org/abs/2108.12427>
 - [Read 1-6; Section 7] Yonadav Shavit, *What Does It Take to Catch A Chinchilla? Verifying Rules on Large-Scale Neural Network Training via Compute Monitoring* (2023), <https://arxiv.org/pdf/2303.11341>
 - Brian R. Knight, *How to Build Beneficial Regulatory Sandboxes*, MERCATUS CENTER (Feb. 9, 2021), <https://www.mercatus.org/research/state-testimonies/how-build-beneficial-regulatory-sandboxes>
 - Katherine J. Strandburg, *Rulemaking and Inscrutable Automated Decision Tools*, 119 COLUM. L. REV. 1851 (2019)
 - [Read Introduction and p. 43-54] Margot E. Kaminski, *Regulating the Risks of AI*, 103 B.U. L. REV. (2023).
 - Scott J. Shackelford & Rachel Dockery, *Governing AI*, 30 CORNELL J.L. & PUB. POL'Y 279 (2020) (Chapter III).
 - Engstrom & Ho, *Algorithmic Accountability in the Administrative State*, 37 YALE J. REG. 800 (2020)

8. AI and Liability

- Key Issues:
 - How does the integration of AI in various industries impact traditional liability laws, and what are the key legal challenges that arise?
 - What legal frameworks and precedents exist for determining liability in AI-related cases, and how are they evolving to keep pace with technological advancements?
 - In what ways can legal professionals and policymakers strike a balance between fostering AI innovation and ensuring accountability and ethical use of AI systems within the scope of liability law?
- Reading:
 - [Read 1819-1843, 1850-1869] Patrick Hubbard, *Sophisticated Robots: Balancing Liability, Regulation and Innovation*, 66 FLORIDA L. REV. (2015), <https://scholarship.law.ufl.edu/cgi/viewcontent.cgi?article=1204&context=flr>
 - Azim Shariff et al., *Whose Life Should Your Car Save?*, NEW YORK TIMES (Nov. 3, 2016), <https://www.nytimes.com/2016/11/06/opinion/sunday/whose-life-should-your-car-save.html>
- Optional further reading:
 - [Read Parts I and III] Kenneth S. Abraham & Robert L. Rabin, *Automated Vehicles and Manufacturer Responsibility for Accidents: A New Legal Regime for a New Era*, 105 VA. L. REV. 127 (2019), <https://law.stanford.edu/publications/automated-vehicles-and-manufacturer-responsibility-accidents-new-legal-regime-new-era/>
 - Philipp Hacker, *The European AI Liability Directives -- Critique of a Half-Hearted Approach and Lessons for the Future*, (Jan. 2023), <https://doi.org/10.48550/arXiv.2211.13960>
 - EUROPEAN COMMISSION, *Proposal for a Directive of the European Parliament and of the Council on Adapting Non-Contractual Civil Liability Rules to Artificial Intelligence, AI Liability Directive* (Sept. 2022), https://ec.europa.eu/info/sites/default/files/1_1_197605_prop_dir_ai_en.pdf
 - James Davey, *By Insurers, For Insurers: The UK's Liability Regime for Autonomous Vehicles*, 13 J. TORT L. 163 (2020).
 - Fred Lambert, *Tesla (TSLA) Is About to Launch Its In-House Insurance Program in More States*, ELECTREK (Mar. 22, 2021).
 - Kenneth S. Abraham & Robert L. Rabin, *The Future Is Almost Here: Inaction Is Actually Mistaken Action*, 105 VA. L. REV. ONLINE 91 (2019).
 - Ryan Calo, *Commuting to Mars: A Response to Professors Abraham and Rabin*, 105 VA. L. REV. ONLINE 84 (2019).
 - Anat Lior, *Insuring AI: The Role of Insurance in Artificial Intelligence Regulation*, 35 HARV. J.L. & TECH. 467 (2022).
 - Mark A. Lemley & Bryan Casey, *Remedies for Robots*, 83 U. CHI. L. REV. 1311 (2019), Introduction, Chapters II and III.
 - Ryan Abbott, Excerpt from *The Reasonable Robot*, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3611370
 - [Read 388-392] Matthew Scherer, *Regulating Artificial Intelligence Systems: Risks, Challenges, Competencies, and Strategies*, 29 HARVARD JOURNAL OF LAW AND TECHNOLOGY, Vol 29, No. 2, (2015) http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2609777

- [Read pg 49] Jack Balkin, *The Path of Robotics Law*, THE CIRCUIT - CALIFORNIA LAW REV. (2015), <http://euro.ecom.cmu.edu/program/law/08-732/AI/Balkin.pdf>
- P.M. Asaro, *Robots and Responsibility from a Legal Perspective* (2007), <http://www.peterasaro.org/writing/ASARO%20Legal%20Perspective.pdf>

9. AI and Intellectual Property

- **Potential guest speaker--TBD**
- Key Issues:
 - How are AI models trained?
 - Does training AI models on copyrighted material constitute a violation of IP rights?
 - To what extent can creators have IP rights to content created with the assistance of AI?
- Reading:
 - TBD

10. AI and Discrimination I: Why Algorithms Discriminate

- Key issues:
 - What do we mean when we talk about algorithmic discrimination?
 - How does antidiscrimination law apply to big data and machine learning?
 - Are existing legal remedies a satisfactory remedy for discriminatory algorithms?
- Reading:
 - Sam Corbett-Davies et al., *A Computer Program Used for Bail and Sentencing Decisions Was Labeled Biased Against Blacks. It's Actually Not that Clear*, WASHINGTON POST, October 17, 2016, <https://www.washingtonpost.com/news/monkey-cage/wp/2016/10/17/can-an-algorithm-be-racist-our-analysis-is-more-cautious-than-propublicas/>
 - This article provides you with a short, easy to be introduction to a frequently cited example of the algorithm of discrimination that is actually more complicated than it is usually presented. It explains how fairness criteria such as “group calibration” and “equality of false positive rates” can be inherently incompatible.
 - Solon Barocas & Andrew D. Selbst, *Big Data's Disparate Impact*, 104 CALIF. L. REV. 671 (2016), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2477899
 - A seminal and widely cited essay discussing the application of antidiscrimination law to big data
- Optional further reading:
 - Anupam Chander, *The Racist Algorithm?*, 115 MICH. L. REV. 1023 (2017), <https://repository.law.umich.edu/mlr/vol115/iss6/13/>
 - This article explains how predictive algorithms can problematically reproduce biases inherent in training data and sketches some potential solutions.
 - SOLON BAROCAS AND MORITZ HARDT AND ARVIND NARAYANAN, *FAIRNESS AND MACHINE LEARNING: LIMITATIONS AND OPPORTUNITIES* (fairmlbook.org 2019), Ch.1 Introduction <https://fairmlbook.org/pdf/introduction.pdf>
 - This piece nicely explains how existing disparities can be reproduced and exacerbated by machine learning
 - [Read 1049-1093] Andrew Selbst, *Unfair Artificial Intelligence: How FTC Intervention Can Overcome the Limitations of Discrimination Law*, 171 U. PENN. L. REV. (2022), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4185227
 - A specific take on how the FTC can apply its existing authority to AI

- Danielle Keats Citron & Frank Pasquale, *The Scored Society: Due Process for Automated Predictions*, 89 WASH. L. REV. 1 (2014), <https://digitalcommons.law.uw.edu/wlr/vol89/iss1/2/>
- Aziz Z. Huq, *Constitutional Rights in the Machine-Learning State*, 105 CORNELL L. REV. 1875 (2020), <https://www.cornelllawreview.org/wp-content/uploads/2020/12/Huq-final.pdf>
- Sorelle A. Friedler, Carlos Scheidegger & Suresh Venkatasubramanian, *The (Im)possibility of Fairness: Different Value Systems Require Different Mechanisms for Fair Decision Making*, COMMUNICATIONS OF THE ACM, April 2021, Vol. 64 No. 4, Pages 136-143, 10.1145/3433949
- Ben Green, *Escaping the Impossibility of Fairness- From Formal to Substantive Algorithmic Fairness*, Philosophy & Technology (2022) 35:90 <https://doi.org/10.1007/s13347-022-00584-6>
- Rashida Richardson, Jason M. Schultz, & Kate Crawford, *Dirty Data, Bad Predictions: How Civil Rights Violations Impact Police Data, Predictive Policing Systems, and Justice*, 94 N.Y.U. L. REV. ONLINE 192 (2019), <https://www.nyulawreview.org/online-features/dirty-data-bad-predictions-how-civil-rights-violations-impact-police-data-predictive-policing-systems-and-justice/>
- Kate Crawford and Trevor Paglen, “Excavating AI: The Politics of Training Sets for Machine Learning” (September 19, 2019), <https://excavating.ai>
- Joy Buolamwini & Timnit Gebru, *Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification* (2018), <https://proceedings.mlr.press/v81/buolamwini18a/buolamwini18a.pdf>

11. AI and Discrimination II: Solutions

- Key issues:
 - Can technical solutions address algorithmic discrimination?
 - Can technical solutions make algorithms *less* biased than human assessments?
 - Are there legal barriers in applying algorithmic anti-discrimination?
- Reading:
 - Crystal S. Yang & Will Dobbie, *Equal Protection Under Algorithms: A New Statistical and Legal Framework*, 119 MICHIGAN L. REV. 291 (2020)
- Optional further reading:
 - Jason R. Bent, *Is Algorithmic Affirmative Action Legal?*, 106 GEORGETOWN L. REV. 803 (2020), <https://www.law.georgetown.edu/georgetown-law-journal/wp-content/uploads/sites/26/2020/04/Is-Algorithmic-Affirmative-Action-Legal.pdf>
 - This article nicely outlines potential solutions to algorithmic bias, as well as the technical and legal challenges in applying these solutions.
 - Deven R. Desai, Swati Gupta, and Jad Salem, *Using Algorithms to Tame Discrimination: A Path to Algorithmic Diversity, Equity, and Inclusion*. U.C. DAVIS L. REV. (forthcoming), <https://ssrn.com/abstract=4244925>
 - Note: make sure you download this as a pdf from the law review or SSRN, otherwise you will miss all the pretty graphics.
 - The paper argues that “rather than rejecting data science and algorithmic methods as leading to discrimination, interrogating legal rules to see what the

law requires and what the law allows enables innovation on how to identify and mitigate bias.”

- Rebecca Crotof, Margot Kaminski & Nicholson Price, *Humans in the Loop*, VANDERBILT L. REV. (2023), <https://vanderbiltlawreview.org/lawreview/2023/03/humans-in-the-loop/>
- Madeleine Elish, *Moral Crumple Zones: Cautionary Tales in Human Robot Interaction*, 5 ENGAGING SCIENCE, TECHNOLOGY, AND SOCIETY 40 (2019), <https://ssrn.com/abstract=2757236>
- Kiel Brennan-Marquez, Karen Levy & Daniel Susser, *Strange Loops: Apparent versus Actual Human Involvement in Automated Decision Making*, 34 BERKELEY TECH. L.J. 745 (2019), https://digitalcommons.lib.uconn.edu/cgi/viewcontent.cgi?article=1608&context=law_papers
- Laux, Johann, *Institutionalised Distrust and Human Oversight of Artificial Intelligence: Toward a Democratic Design of AI Governance under the European Union AI Act* (March 3, 2023), <https://ssrn.com/abstract=4377481> or <http://dx.doi.org/10.2139/ssrn.4377481>
- Ben Green, *The Flaws of Policies Requiring Human Oversight of Government Algorithms*, COMPUTER LAW & SECURITY REVIEW 45 (July 2022): 105681, <https://doi.org/10.1016/j.clsr.2022.105681>
- Joshua A. Kroll et al., *Accountable Algorithms*, 165 U. PA. L. REV. 633 (2017), https://scholarship.law.upenn.edu/penn_law_review/vol165/iss3/3/
- Anya E.R. Prince & Daniel Schwarcz, *Proxy Discrimination in the Age of Artificial Intelligence and Big Data*, 105 IOWA L. REV. 1257 (2020), <https://ilr.law.uiowa.edu/print/volume-105-issue-3/proxy-discrimination-in-the-age-of-artificial-intelligence-and-big-data>

12. AI Judges

- Key issues:
 - Can AI help judges handle their caseloads efficiently?
 - Would the use of AI change the types of issues judges focus on?
 - How have judges in other countries begun to use AI tools?
 - How might AI change the development of caselaw or the perceived legitimacy of the judiciary?
- Reading:
 - Robert Buckland, *AI, Judges, and Judgment: Setting the Scene*, <https://dash.harvard.edu/handle/1/37377475>
 - A good summary of the major issues likely to be encountered with AI judging. Given his own background, Judge Buckland is generally sympathetic to human judges and skeptical of AI judges. As you read, consider how some of the problems he identifies with AI judges might potentially apply to human judges as well.
 - [Read Introduction] Richard M Re and Alicia Solow-Niederman, *Developing Artificially Intelligent Justice*, 22 STANFORD TECHNOLOGY LAW REVIEW 242 (2019), https://law.stanford.edu/wp-content/uploads/2019/08/Re-Solow-Niederman_20190808.pdf
 - A good summary of issues relating to AI judging, with more of an emphasis on big data rather than contemporary language models.

- [Read Section titled “AI in the Courtroom”, pages 59-81] FEDERAL JUDICIAL CENTER, AN INTRODUCTION TO ARTIFICIAL INTELLIGENCE FOR FEDERAL JUDGES (2023), https://www.fjc.gov/sites/default/files/materials/47/An_Introduction_to_Artificial_Intelligence_for_Federal_Judges.pdf
 - An interesting synopsis of the kinds of practical issues that judges are likely to face relating to AI.
- Optional further reading:
 - Danielle Keats Citron, *Technological Due Process*, 85 WASH. U. L. REV. 1249 (2008)
 - This is an important early work that explains several ways in which algorithmic decision-making can be incompatible with general notions of a “fair” decision-making, or more specifically the requirements of due process in the context of administrative law.
 - Rashida Richardson, Jason Schultz & Kate Crawford, *Dirty Data, Bad Predictions: How Civil Rights Violations Impact Police Data, Predictive Policing Systems, and Justice*, 94 N.Y.U. L. REV. ONLINE 192 (2019) https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3333423
 - Rachel E. Stern et al., *Automating Fairness? Artificial Intelligence in the Chinese Court*, 59 COLUM. J. INT’L L. 515 (2021), https://scholarship.law.columbia.edu/faculty_scholarship/2940/
 - An interesting comparative study of a judicial system that has begun to use AI judging.
 - Andrew Guthrie Ferguson, *Predictive Policing and the Rule of Law*, 94 WASH. U. L. REV. (2017), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2765525
 - Danielle Ensign, et al., "Runaway Feedback Loops in Predictive Policing" <https://arxiv.org/abs/1706.09847>
 - Kristian Lum & William Isaac, "To predict and serve?" <https://rss.onlinelibrary.wiley.com/doi/full/10.1111/j.1740-9713.2016.00960.x>
 - Engstrom & Vogt, *The New Judicial Governance: Courts, Data, and the Future of Civil Justice*, DEPAUL L. REV. (forthcoming 2023), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4024511

13. Student Presentations I

- No assigned reading

14. Student Presentations II

- No assigned reading