

Prof. Andrew Chin
Law 552.1
www.andrewchin.com

Spring 2024 Syllabus
University of North Carolina School of Law
Room 3031, Thursdays 9:50 – 11:50 (3 credit hours)

ARTIFICIAL INTELLIGENCE AND THE LAW

Description

Artificial intelligence comprises a set of computationally intensive information technologies intended to substitute for human intelligence in the performance of cognitive tasks. These technologies are being used in connection with an ever-increasing range of performance-enhancing and labor-saving commercial, industrial and research applications, and regulatory, law enforcement and military. Accompanying these developments are emerging concerns about the transformation of labor, including the practices of legal services and institutions; the rights, duties and liabilities attaching to the designs and activities of robotic and transactional AI agents; and the privacy, security, social and ethical impacts of AI and automation. Technological changes in AI, and responses to those changes in law and governance, are ongoing, rapid and often difficult to predict. This seminar will prepare students to be critical and informed observers of, and participants in, the legal developments, disputes and discourses surrounding AI technologies that will emerge over the course of their careers. We will engage with a broad range of cases, statutes, regulations, commentaries and other materials pertaining to AI and the law through a variety of discussions, presentations, and writing assignments.

Course Prerequisites

None. Comfort with reading technical literature written for a popular audience is helpful.

Learning Outcomes

Students shall be able to find, analyze, and use relevant legal materials (including, among others, statutes, cases, regulations, and other administrative materials) in order to identify and resolve problems and communicate legal analysis in a variety of written and oral formats. Students shall be able to exercise legal research and writing needed for well-rounded, competent, and ethical participation as members of the legal profession, including, but not limited to, applying legal and other scholarship to understand and affect legal policy.

Course Requirements

An academic-genre paper of at least 9,000 words; one guest speaker critique and one peer critique of at least 800 words each; short in-class final presentation. An example of a past guest speaker critique is appended.

Course Format

The course will begin with an **introductory guest lecture and in-class discussion** on artificial intelligence and the law, from which you should have a sense of the substantive scope of this course.

Our third and tenth meetings will be devoted to **meeting as a research group** in which you will have a designated partner who will be your primary peer mentor and reviewer.

Our eleventh, twelfth and thirteenth meetings will conclude the course with an opportunity to **present** the main ideas of your paper to the class in prepared talks of 10 to 15 minutes.

Alternatively, you may record a video presentation of 10 to 15 minutes to be viewed asynchronously.

In most other weeks, we will welcome a series of eight distinguished **guest speakers** on diverse topics, each presenting for one hour. Each of you will be responsible for **hosting** one of these speakers. You may express a preference for a particular speaker, but if other students express the same preference, only one of you will be randomly chosen as the host and the others will be assigned as needed for coverage. Most speakers will provide readings (e.g., an article or book

chapter) to accompany their presentation at least one week in advance. As host, you will write a **summary and critique** of the reading (1) for the benefit of any classmates who do not have time to do the full reading, (2) for the purpose of beginning the discussion and Q&A after the guest speaker's presentation and continuing the discussion into the second hour as warranted. Your summary and critique papers are due by **11:59 p.m.** the **Tuesday** preceding the presentation, and will be shared with your classmates before the class and with the guest speaker after the class. Even if you are not hosting, you should at least skim the reading and come to class prepared with a question to ask the speaker.

In the remaining time, we will also discuss the textbook and some related topics. I've already scheduled myself for a presentation of an AI/patents article, but if time permits, I also plan to present some of my other work in this field.

Readings

Most of the readings are linked from the Detailed Course Coverage table below. Other readings, presentations, and guest critique papers will be distributed via Sakai and/or Zoom.

While most of your grade will be based on written work, we will be discussing reading material every week. Reading all of the distributed materials will usually be necessary to participate fully in these discussions. You may occasionally read the critique paper in lieu of a guest speaker's article, but you should not make this your usual practice.

You will find reading for the purpose of discussion less intensive and detail-oriented than reading in preparation for an examination. The emerging field of artificial intelligence and the law by its nature is so diverse and pervasive that none of us is going to have all the answers.

Attendance and Participation

As a research group, our seminar requires the attendance, punctuality and participation of all members. Under the Law School's attendance policy, a student must attend at least 80% of the required meetings; i.e., 11 of our 13 meetings. A student who does not do so presumptively fails the course. To avoid failing the course, students can make up missed course meetings by watching the recordings (available via the Echo360 system). If you email me in advance, you may attend via Zoom on any week that we have an outside guest speaker.

Some absences are excused. Excused absences include COVID-related situations, personal/family illness or emergency, pregnancy or disability as approved by Accessibility and Resource Services, and up to two religious observances under the Law School's policy.

If you accumulate so many excused absences that you cannot meet the 80% minimum, you may make up absences to bring your number of attended classes up to 10. Making up an excused absence includes viewing the entire recorded class, turning all assignments and any other written work completed by the students in attendance, and working with your partner as necessary to ensure that both of you can timely complete all subsequent coursework. If possible, you should view the recording for a missed class session before the next class session.

Please send me a timely email notifying me of any absence, whether excused or unexcused. Please also send me an email once you have completed the necessary work to make up any excused absence(s).

Writing Requirements, Class Participation and Grading

To excel in this course, you will need to participate fully in both written and oral discussions throughout the semester. Your grade will be based on the following components. Assignments are due at 11:59 p.m. on the date indicated. One letter grade will be deducted for a late assignment.

- Final draft of your research paper (at least 9,000 words) 70%

- Topic description due Feb. 1
- Draft outline and partial draft due March 21
- Presentations April 4 and April 11 (10%)
- Final paper due May 6, 5 p.m. (60%)
- Critique papers (at least 800 words each) 20%
 - Guest critique (due 11:59 p.m. Tuesday before) (10%)
 - One peer critique due March 28 (10%)
- Discussion participation 10%

ChatGPT

I neither encourage nor discourage the use of ChatGPT as a general matter. After you have some experience in exploring the limits of its legal knowledge and analytical abilities, however, you should understand that it is inadvisable to rely on ChatGPT as the sole authority for any legal or factual proposition.

If you use ChatGPT in any way to assist you in researching or writing your paper, your paper must provide a statement (not included in the word count) fully and specifically describing the purpose and extent of any such use.

You should (re)familiarize yourself with the Law School's plagiarism policy at <https://law.unc.edu/academics/academic-policies/>. To avoid violating this policy, in your final paper you must attribute to ChatGPT 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, 2023 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)."

You should also provide an appendix (not included in the word 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.

Honor Code

The Honor Code is in effect in this class and all others at the University. I am committed to treating Honor Code violations seriously and encourage all students to become familiar with its terms set out at <http://instrument.unc.edu>. If you have questions, it is your responsibility to ask me about the Code's application. All written work must be submitted with a statement that you have complied with the requirements of the Honor Code in all aspects of the submitted work.

Office Hours

After class or via Zoom by appointment.

Detailed Course Coverage Table (tentative)

Revisions to this syllabus will be announced in class and distributed via Sakai.

<i>Date</i>	<i>Topic(s)/Title</i>	<i>Guest Presenter</i>	<i>Reading/Topic</i>
1/11	Introduction to AI and Machine Learning; legal issues surrounding ChatGPT; assign hosts		
1/18	Case discussion; assign hosts and partners	Ryan Abbott	Brauneis, <i>AI Litigation Database</i> ; Abbott, <i>The Reasonable Computer: Disrupting the Paradigm of Tort Liability</i>
1/25	AI and discrimination Topic consultations; assign partners		Katyal, <i>Private Accountability in the Age of Artificial Intelligence</i> (skim)
2/1	AI and rights of publicity	Mark Bartholomew	Bartholomew, <i>A Right to Be Left Dead</i>
2/8	AI and patent, privacy and/or antitrust law		Chin, <i>Surgically Precise But Kinematically Abstract Patents</i> ; Chin, <i>Differential Privacy: The Facebook Advertiser Case Study</i> (w/Dinerstein v. Google) or Chin, <i>Antitrust Analysis in Software Product Markets</i>
2/14	Topic consultations (via Zoom)		
2/15	Generative AI and copyright II	Carol Hayes	Lemley & Casey, <i>Fair Learning</i> ; Hayes, <i>Generative AI and Copyright: Both Sides of the Black Box</i>
2/22	AI and health care; generative AI and finding the law	Sara Gerke, Paul Callister	Gerke, <i>"Nutrition Facts Labels" for AI/ML Medical Devices</i> ; Callister, <i>Generative AI and Finding the Law</i>
2/29	Explainable AI and taxation	Anthony Niblett*	Alarie & Niblett, <i>Explainable AI and Taxation: A Real-Life Application</i>
3/5	Consultations		
3/21	Draft outline and partial draft due AI governance and regulation; or AI and legal personality	Simon Chesterman	Chesterman, <i>The Tragedy of AI Governance</i> or Chesterman, <i>Artificial Intelligence and the Limits of Legal Personality</i>
3/28	Peer critiques due Draft consultations	You and your partner	Your partner's draft
4/4	Presentations	Some of you	
4/11	Presentations	The rest of you	

An Example of a Guest Critique Paper

Introduction

Tejas Narechania's *Machine Learning as Natural Monopoly* piece concludes that, in at least some instances, machine learning applications can be considered natural monopolies that give rise to a need for increased regulation to protect against a myriad of possible harms, including intrusive data collection and algorithmic redlining. In his argument, Professor Narechania draws a number of parallels to the railway and telephone network monopolies of the past and indicates that, while the monopoly and the issues that arise as a result of the monopoly may have changed, the fundamental need to have regulatory supervision over rates and service standards has not. At bottom, Professor Narechania presents a compelling argument for the need to classify at least some machine learning applications as natural monopolies, however some parts of his use of analogy would benefit from greater clarification and discussion because they appear to stretch the analogy almost to the point of no longer being comparable.

Summary

Professor Narechania's contribution to the literature is by exploring the extent to which previous natural monopoly frameworks, such as the telephone and railway monopolies, can apply to more modern machine learning applications and how policymakers can similarly draw from that previous monopoly framework and apply it to this more modern technological space. In general, Professor Narechania defined a natural monopoly as one in which consumers are likely to face high prices, poor quality, or both and in which competition is unlikely to address these price- and quality-related problems. As he recounts, the goal of monopoly regulation is to avoid "wasteful" competition by offering a company a monopoly in exchange for a commitment to provide "reasonable service at reasonable rates."

To illustrate this notion, Professor Narechania offers two prime examples: telephone and railway monopolies. In both cases, the cost to enter a particular market was high. In the telephone context, there is substantial cost and effort required to physically run the telephone lines between each city to enable telephone service. In the railway context, there is a similarly high cost to run the tracks through each area that would benefit from rail service. After this initial outlay of capital and the buildout, both services become relatively cheaper to maintain and to sign up additional customers. In the telephone context, each new subscriber would only need, in many cases, the simple flip of a switch to enable service that has already been run to one's home or office. In the railway context, each passenger or freight company would only need to pay the requisite fare to utilize the service. There would be a high barrier to entry for a new telephone or railway company because of the substantial outlay of capital required. As a result, companies that have the infrastructure in place can more or less dictate the cost to the consumer with little competition from other companies. This led to the need for regulation to ensure that these high barrier to entry scenarios did not lead to out of control pricing and a decrease of service quality to maximize profits for the company.

With that background, Professor Narechania then draws the parallel to more modern machine learning applications. In at least some applications, the fixed costs of developing a given machine learning system can, like the telephone and railway examples, be exceedingly high. Not only can the physical development of the software be costly, but the costs of acquiring the specialized hardware and data to "train" the application can also cause a substantial increase in the startup costs. Conversely, the cost to distribute the system is quite low, amounting to nothing more than clicking on an internet site or, at most, distributing a physical copy of the developed software. Much like the telephone and railway examples, then, there is a possibility that a given company that has invested in developing a machine learning application can

similarly dictate the cost to the consumer with little competition from other companies. Like the telephone and railway examples, this can lead to exorbitant pricing and reduced quality. In the machine learning context, there is also the added risk that profits can be extracted from using and selling the information gathered from the machine learning system. This can also lead to privacy issues.

Given this rather clear parallel between historical natural monopolies and the more modern machine learning natural monopolies, Professor Narechania proposes a similar need for regulatory scrutiny and how that scrutiny would need to address machine learning monopoly problems of accuracy, autonomy, bias, and privacy. The exact details of this conclusion still need to be fleshed out in this preliminary draft of the piece.

Assessment

The main line of reasoning that Professor Narechania uses throughout his piece lies in examples and analogies to the telephone and railway natural monopolies. In the context of the need for increased regulation of a new industry, the analogies make a compelling argument because they show that previous precedent has adequately handled the frequent price increases and service declines inherent to natural monopolies.

On the other hand, some of the details of Professor Narechania's use of analogy are less compelling. By and large, comparing the cost of laying telephone cable or train track to the development of a new piece of software seems to stretch the comparison quite a bit. One of the generally understood concepts in computing is that the cost to entry to develop software is quite low. Typically, a few programmers can develop software and then release the beta version of the software. With the beta, end users can use it for free and offer feedback, as well as their sample data, to improve the software for the next iteration. The cost for all of this effort, on its face at least, appears to be substantially less than hiring thousands of employees or contractors to lay the cable or train track, handle the easements necessary with the municipalities including each local municipalities' regulations and requirements, and then have substantial exposure to liability for physical injury both during the build process and when the service is in operation. With this facially lower barrier to entry, it becomes difficult to see why a given machine learning market would not be subject to adequate competition to make monopoly regulation unnecessary. To be fair, Professor Narechania does acknowledge that the barrier to entry part of the analogy is slightly different in the machine learning context when he describes that expertise, specialized computing equipment, and vast volumes of data are costly, however it is less clear whether even this really amounts to anywhere near the cost of the physical labor and effort required to build a telephone or railway network.

Moreover, the end user pricing part of the analogy also seems quite stretched. Most consumer driven machine learning applications are offered for little or no money. As a result, it seems difficult to appreciate how a consumer could be subjected to exorbitant pricing, even when considering Professor Narechania's discussion of the value of the consumer's data that is used in lieu of monetary payment.

Conclusion

All together, Professor Narechania's piece offers a compelling parallel between a previous need for regulation and a more modern need for regulation, however some aspects of the parallel seem stretched and hard to apply in the modern context. It would be helpful to address these issues more directly and, specifically, describe how these might affect either the need for regulation altogether or adjust the need for regulation up or down.