

**LAW 6930 AI Foundations & Legal Practice [12472]**  
**UNIVERSITY OF FLORIDA LEVIN COLLEGE OF LAW**  
**SPRING SYLLABUS – LAW 6930 – 2 CREDITS**

Professor Thinh Nguyen  
BG 105A  
Office Phone: 352-273-0642  
Email: nguyenthinh@ufl.edu  
Office Hours: Mondays 1 - 3 PM

**MEETING TIME:** Monday and Wednesday, 3:00 - 3:55 PM

**LOCATION:** HH - 355D

**COURSE DESCRIPTION AND OBJECTIVES:**

This course explores the technical, legal, ethical, and governance landscape of artificial intelligence from its historical foundations to its contemporary applications. We will develop an in-depth understanding of the principles of artificial intelligence, its capabilities, limitations, and use cases. We will use this foundation as the basis for understanding approaches to AI governance at the local, state, national, and international levels, including its implications for legal domains like employment law, consumer protection, privacy/cybersecurity, advertising, election law, criminal law, product liability, technology transactions, intellectual property, and international trade. We will also seek to understand how regimes of laws, norms, and technology shape the development of AI, and are in turn shaped by it.

**Prerequisite:** Students taking this course are required to take Professor Nguyen's "Python for Machine Learning and AI" as a spring compressed class. Students who previously took this compressed class with a passing grade already meet this requirement. Students may seek a waiver of this requirement with the written approval of Professor Nguyen by demonstrating adequate knowledge and proficiency with Python.

**STUDENT LEARNING OUTCOMES:**

At the end of this course, students should be able to:

- understand the foundations and principles of AI and distinguish between key approaches to AI including expert knowledge systems and deep learning,
- analyze the evolving local, state, national, and international AI governance mechanisms, including U.S. and international regulatory regimes, and understand how to design AI compliance programs,
- identify the normative content of AI systems and understand how they inform moral, ethical, and legal issues in the field of AI,
- explain practical applications of AI in law and other domains, including RAG, semantic search, structured text extraction, and agents,
- advise clients on specific legal issues raised by AI applications in areas like intellectual property, privacy, cybersecurity, employment, consumer protection, regulation of professions, torts, free speech, and international law.

**REQUIRED READING MATERIALS:**

We will use this textbook for the technical introduction to artificial intelligence:

Simon J.D. Prince, *Understanding Deep Learning* (MIT Press 2023).  
ISBN: 978-0-262-04864-4

This book is available for sale in hardcover but it is also available online as a free downloadable PDF licensed under the Creative Commons BY-NC-ND license:

<https://anthology-of-data.science/resources/prince2023udl.pdf>

Other reading assignments are listed in the syllabus below. I may supplement the reading as needed based on class discussions. All readings are available online for free download or viewing. **You do not need to purchase any textbooks or other materials for this class unless you want a hard copy of the Prince textbook above.**

Please be sure to register for the Canvas course and have any required materials with you in print or easily accessible electronic form in class. You are responsible for checking your Canvas page and the e-mail connected to the page on a regular basis for any class announcements or adjustments.

### **COURSE EXPECTATIONS AND GRADING EVALUATION:**

Students are expected to be prepared for class by completing the assigned reading prior to class. Students will be assigned to discussion panels for certain classes. Students serving on panels are expected to be prepared and to participate actively in the discussions. Towards the end of the class, we will have a case simulation. Students will be assigned to teams and are expected to prepare for the simulation both individually and as team members.

Students are expected to be able to read and write basic Python (see course requirements above) in order to experience hands-on learning of AI through homework assignments and a group project. Homework assignments will be posted on Canvas by the morning after class and must be completed and turned in prior to the beginning of the next class.

Students will be evaluated based upon participation in class discussion panels (20%), homework (30%), and a final exam (50%).

### **CLASS ATTENDANCE POLICY:**

Attendance in class is required by both the ABA and the Law School. Attendance will be taken at each class meeting. Students are allowed three 3 absences during the course of the semester. Students are responsible for ensuring that they are not recorded as absent if they come in late. A student who fails to meet the attendance requirement will be dropped from the course. The law school's policy on attendance can be found [here](#).

### **COMPLIANCE WITH UF HONOR CODE:**

Academic honesty and integrity are fundamental values of the University community. Students should be sure that they understand the UF Law Honor Code located [here](#). The UF Law Honor Code also prohibits use of artificial intelligence, including, but not limited to, ChatGPT and Harvey, to assist in completing quizzes, exams, papers, or other assessments unless expressly authorized by the professor to do so.

### **INFORMATION ON UF LAW GRADING POLICIES:**

The Levin College of Law's mean and mandatory distributions are posted on the College's website and this class adheres to that posted grading policy. The following chart describes the specific letter grade/grade point equivalent in place:

Letter Grade	Point Equivalent	Letter Grade	Point Equivalent
A (Excellent)	4.0	C (Satisfactory)	2.0
A-	3.67	C-	1.67
B+	3.33	D+	1.33
B	3.0	D (Poor)	1.0
B-	2.67	D-	0.67
C+	2.33	E (Failure)	0.0

The law school grading policy is available [here](#).

### **OBSERVANCE OF RELIGIOUS HOLIDAYS:**

UF Law respects students' [observance of religious holidays](#).

- Students, upon prior notification to their instructors, shall be excused from class or other scheduled academic activity to observe a religious holy day of their faith.
- Students shall be permitted a reasonable amount of time to make up the material or activities covered in their absence.
- Students shall not be penalized due to absence from class or other scheduled academic activity because of religious observances.

### **EXAM DELAYS AND ACCOMMODATIONS:**

The law school policy on exam delays and accommodations can be found [here](#).

### **STATEMENT RELATED TO ACCOMODATIONS FOR STUDENTS WITH DISABILITIES**

Students requesting accommodations for disabilities must first register with the Disability Resource Center (<https://disability.ufl.edu/>). Once registered, students will receive an accommodation letter, which must be presented to the Assistant Dean for Student Affairs (Assistant Dean Brian Mitchell). Students with disabilities should follow this procedure as early as possible in the semester. It is important for students to share their accommodation letter with their instructor and discuss their access needs as early as possible in the semester. Students may access information about various resources on the UF Law Student Resources Canvas page, available at <https://ufl.instructure.com/courses/427635>.

### **STUDENT COURSE EVALUATIONS**

Students are expected to provide professional and respectful feedback on the quality of instruction in this course by completing course evaluations online via GatorEvals. Click [here](#) for guidance on how to give feedback in a professional and respectful manner. Students will be notified when the evaluation period opens and may complete evaluations through the email they receive from GatorEvals, in their Canvas course menu under GatorEvals, or via <https://ufl.bluer.com/ufl/>. Summaries of course evaluation results are available to students [here](#).

### **RECORDINGS OF CLASS**

Students are allowed to record video or audio of class lectures. However, the purposes for which these recordings may be used are strictly controlled. The only allowable purposes are (1) for personal educational use, (2) in connection with a complaint to the university, or (3) as evidence in, or in preparation for, a criminal or civil proceeding. All other purposes are prohibited. Specifically, students may not publish recorded lectures without the written consent of the instructor. A "class lecture" is an educational presentation intended to inform or teach enrolled students about a particular subject, including

any instructor-led discussions that form part of the presentation, and delivered by any instructor hired or appointed by the University, or by a guest instructor, as part of a University of Florida course. A class lecture does not include lab sessions, student presentations, clinical presentations such as patient history, academic exercises involving solely student participation, assessments (quizzes, tests, exams), field trips, private conversations between students in the class or between a student and the faculty or guest lecturer during a class session. Publication without permission of the instructor is prohibited. To “publish” means to share, transmit, circulate, distribute, or provide access to a recording, regardless of format or medium, to another person (or persons), including but not limited to another student within the same class section. Additionally, a recording, or transcript of a recording, is considered published if it is posted on or uploaded to, in whole or in part, any media platform, including but not limited to social media, book, magazine, newspaper, leaflet, or third-party note/tutoring services. A student who publishes a recording without written consent may be subject to a civil cause of action instituted by a person injured by the publication and/or discipline under UF Regulation 4.040 Student Honor and Student Conduct Code.

**ABA OUT-OF-CLASS HOURS REQUIREMENTS:** ABA Standard 310 requires that students devote 120 minutes to out-of-class preparation for every “classroom hour” of in-class instruction. Each weekly class is approximately 2 hours in length, requiring at least **4 hours of preparation** outside of class including reading the assigned materials, homework assignments, preparing for panels and in-class presentations as part of case simulations.

### **COURSE SCHEDULE OF TOPICS AND ASSIGNMENTS**

This syllabus is offered as a guide to the direction of the course. Our pace will depend in part on the level of interest and the level of difficulty of each section and is subject to change.

### **USE OF ARTIFICIAL INTELLIGENCE**

Use of artificial intelligence resources is permitted for all class assignments unless otherwise stated in the assignment instructions. Use of artificial intelligence on the final exam is prohibited.

### **UF LEVIN COLLEGE OF LAW STANDARD SYLLABUS POLICIES:**

Other information about UF Levin College of Law policies, including compliance with the UF Honor Code, Grading, Accommodations, Class Recordings, and Course Evaluations can be found at [this link](#).

### **UF ACADEMIC POLICIES AND RESOURCES:**

Other information about UF academic policies and resources can be found at [this link](#).

---

## **Week 1 - Move 37: Machine Intelligence**

We will discuss the history of artificial intelligence, from the classical age to modern large language models, including the key milestones along the way such as the invention of the McCulloch-Pitts Neuron, the Perceptron, AlexNet, AlphaGo, AlphaFold, and ChatGPT. We will examine the two basic approaches to AI systems: expert knowledge systems and machine learning and see how they have complemented, and competed with, each other over the course of history. We will explore the reasons for the rise of deep learning systems in the 1980s and why they came to dominate the modern AI landscape.

### **Reading:**

*Prince*, Chapters 1, 2

**Watch:**

Google DeepMind, *AlphaGo*, Youtube (March 13, 2020),  
[https://youtu.be/WXuK6gekU1Y?si=HbFRl4wJ\\_x9X4RoO](https://youtu.be/WXuK6gekU1Y?si=HbFRl4wJ_x9X4RoO)

---

**Week 2 - The Perceptron: The Origins of Artificial Neural Networks**

We will discuss the earliest neural networks, starting from the McCulloch-Pitts Neuron (MPN). We will explore the mathematical properties of the MPN and its application in three fundamental tasks performed by AI: association, regression, and classification. Then, we will see how the Perceptron, invented in the 1950s, adapted MPNs to perform complex classification tasks, but we will also see the limits of the first Perceptron model (the XOR problem identified by Minsky, et al), which led to the first "AI Winter."

**Readings:**

*Prince*, Chapters 3, 4

**Watch**

3Blue1Brown, *But What is a Neural Network?*, Youtube,,  
<https://youtu.be/aircAruvnKk?si=5MamkBVxEaEBTCLd>

---

**Week 3 - Descent: Training Artificial Neural Networks**

We motivate the search for the solution to the XOR problem, namely, multilayer Perceptrons with activation functions. We will see how by adding layers and activation functions to our Perceptron networks, we can build a network with the mathematical properties of a Universal Function Approximator. We will see a demonstration of how multilayer Perceptrons can be used to fit arbitrarily complex data by adjusting weights and biases, starting in two-dimensions and then in higher dimensions. Then, we will learn how we can train our Multilayer Perceptrons using gradient descent and backpropagation. We will also explore the fundamental challenges of training, including overfitting, underfitting, local minima, the trade-off between precision and recall, and how researchers use techniques like regularization, drop-out, hyperparameter search, and loss validation to mitigate these problems.

**Readings:**

*Prince*, Chapters 5, 6

**Watch:**

3Blue1Brown, *Gradient Descent, How Neural Networks Learn*, Youtube (October 16, 2017), <https://youtu.be/IHZwWFHwa-w?si=Tw6ZwwdiZevXzWsm>

3Blue1Brown, *Backpropagation, Intuitively*, Youtube (November 3, 2017),  
<https://youtu.be/Ilg3gGewQ5U?si=DqFDPanrs-5JWUVW>

---

**Week 4 - AI Winter, AI Spring: Rise of Deep Learning**

To solve more complicated problems with our Multilayer Perceptron requires additional techniques that involve feature extraction. One of these, extracting shapes from images using

convolutional layers, led to the breakthrough that ended the AI winter: AlexNet. We will explore why AlexNet was a significant conceptual breakthrough in deep learning, along with recurrent neural networks and encoder / decoders, where we introduce the concept of latent embeddings. All of these became the basis for the invention of the modern transformer, the breakthrough that enabled large language models (LLMs) like ChatGPT.

**Readings:**

*Prince*, Chapter 10, 11

3Blue1Brown, *Transformers, the Tech Behind LLMs*, Youtube,  
<[https://youtu.be/wjZofJX0v4M?si=LOqX\\_eJD8ydMp9XG](https://youtu.be/wjZofJX0v4M?si=LOqX_eJD8ydMp9XG)>

---

**Week 5 - Attention is All You Need: The Transformer Architecture**

We will continue our exploration of the transformer architecture by studying the concepts of tokenization, semantic embeddings, and attention. Paired with multilayer Perceptrons, the attention mechanism gave rise to the transformer networks (BERT and GPT) that became the basis for large language models. We will understand how these can be used to solve general learning problems like language translation, text generation, image understanding, playing games like Go, predicting the weather, and even solving one of the greatest scientific challenges in history: the protein folding problem.

**Readings:**

*Prince*, Chapter 12

3Blue1Brown, *Attention in Transformers*, Youtube,  
<<https://youtu.be/eMlx5fFNoYc?si=uMJ4jkqPEc7QvHFJ>>

---

**Week 6 - Machines of Loving Grace: Practical Applications of LLMs**

This class will demonstrate practical applications of LLMs using the University of Florida's Navigator Toolkit and the OpenAI API. We will demonstrate how to use the API to perform chat completions, chat with memory, cache-augmented generation (CAG), structured text generation, text summarization, semantic search, retrieval-augmented generation (RAG), and agentic applications (with tool calling). We will also demonstrate multi-modal LLM capabilities, including transcription, data extraction, and image-based and multi-modal reasoning. We will also look at techniques for customizing LLMs like fine-tuning, prompt and context engineering, and few-shot training. We will discuss how AI is transforming the legal profession (e.g., legal research tools, contract analysis, predictive analytics in litigation) and issues of professional responsibility (e.g., competence, confidentiality) in the context of AI tools.

**Readings:**

*Prince*, Chapters 20, 21

Robert J. Couture, *The Impact of Artificial Intelligence on Law Firms' Business Models*, Center on the Legal Profession, Harvard Law School (Feb. 25, 2025),  
<https://clp.law.harvard.edu/knowledge-hub/insights/the-impact-of-artificial-intelligence-on-law-law-firms-business-models/>

Annabel V. Teiling, *Will AI Render Lawyers Obsolete?*, N.Y. St. Bar Ass'n (Oct. 7, 2025), [https://nysba.org/will-ai-render-lawyers-obsolete/?srsltid=AfmBOorwTGIDbj8-n9P4XyTVY-Mnrqt\\_-EdBZfXaLjegoIX3t1Al2qxX](https://nysba.org/will-ai-render-lawyers-obsolete/?srsltid=AfmBOorwTGIDbj8-n9P4XyTVY-Mnrqt_-EdBZfXaLjegoIX3t1Al2qxX)

---

### **Week 7 - Ghost in the Machine: Normative Dimensions of AI**

Creating AI systems and deploying them in the world involves a sequence of value choices, which are often opaque and hidden from the end user. We will apply everything we've learned about AI training to help us unpack the normative content of these decisions so we can understand how those choices, such as how to collect and label data, how to deal with overfitting vs. underfitting, how to balance precision vs. recall, how to align the model to human objectives, and how to deploy, monitor, and maintain these systems can have profound downstream implications for bias and discrimination, fairness, equality, well-being, and moral uncertainty.

#### **Readings:**

Daniel J. Solove, *Artificial Intelligence and Privacy*, 77 Fla. L. Rev. 1 (2025), <https://www.floridalawreview.com/api/v1/articles/129976-artificial-intelligence-and-privacy.pdf>

Harvard Law Review, *Resetting Antidiscrimination Law in the Age of AI*, 138 Harv. L. Rev. 1562, <https://harvardlawreview.org/print/vol-138/resetting-antidiscrimination-law-in-the-age-of-ai/>

Mirko Bagaric, et al., *The Solution to the Pervasive Bias and Discrimination in the Criminal Justice System: Transparent and Fair Artificial Intelligence*, 59 Am. Crim. L. Rev. 95 (2022), [https://www.law.georgetown.edu/american-criminal-law-review/wp-content/uploads/sites/15/2022/02/59-1-Bagaric-Solution\\_to\\_Pervasive\\_Bias\\_in\\_Criminal\\_Justice\\_System.pdf](https://www.law.georgetown.edu/american-criminal-law-review/wp-content/uploads/sites/15/2022/02/59-1-Bagaric-Solution_to_Pervasive_Bias_in_Criminal_Justice_System.pdf)

---

### **Week 8 - The Prisoner's Dilemma: Competition and Cooperation in the Age of AI**

Competition over AI supremacy can lead to a type of game recognized in game theory as the "prisoner's dilemma". Here, we may end up with a socially suboptimal outcome, even though all agents act rationally in their self-interest. AI creates a nested set of such collective action problems extending from individual and firm behavior all the way to grand strategy. We will look at the AI competition between the US, EU, and China and the attempts to address these collective action problems through negotiations, treaties, and norms. But we will also see the effects of the Prisoner's Dilemma in areas such as corporate investments, R&D, international trade, export controls, sanctions, antitrust, technology embargoes, cybersecurity, and economic espionage.

#### **Readings:**

James Tobin, *The Prisoner's Dilemma*, Heritage Project, (2025) <https://heritage.umich.edu/stories/the-prisoners-dilemma/>

Michael Froman, *China, the United States, and the AI Race*, COUNCIL ON FOREIGN RELATIONS (2025), <https://www.cfr.org/article/china-united-states-and-ai-race>

Colin H. Kahl & Jim Mitre, *The Real AI Race*, Foreign Affs. (July 9, 2025), <https://www.foreignaffairs.com/united-states/china-real-artificial-intelligence-race-innovation>

---

## **Week 9 - Move 78: Legal and Regulatory Responses to AI**

This class will be an overview of current regulatory schemas that are starting to grapple with the issues raised by AI, focusing on how lawyers and corporations see these problems in practical terms. We will look at corporate, local, state, national and international governance mechanisms for AI. We will focus on how corporate forms, such as non-profit vs. for-profit forms, may influence the development of AI, examining OpenAI's attempt to restructure its business. Then, we will look at whether these frameworks are effective at combating AI-based discrimination, bias, privacy loss, misinformation, deepfakes, and other potential harms. We will explore the implications of the "black box" problem of AI for the coherence of doctrinal frameworks that rely on legal explanation. We will compare these efforts in the U.S. to European regulatory frameworks, specifically, the EU AI Act. We will also discuss how businesses and law firms approach advising, designing, and maintaining compliance programs centered around AI risks in areas such as employment, housing, healthcare, finance, autonomous vehicles, insurance, consumer protection, regulation of professions, negligence, product liability, privacy, free speech, cybersecurity, and more.

### **Readings:**

Filippo Lancieri, Laura Edelson & Stefan Bechtold, *AI Regulation: Competition, Arbitrage & Regulatory Capture* (2024) (forthcoming in Theor. Inq. L.), <https://scholarship.law.georgetown.edu/facpub/2647>

Sandra Wachter, *Limitations and Loopholes in the EU AI Act and AI Liability Directives: What This Means for the European Union, the United States, and Beyond*, 26 Yale J.L. & Tech. 671 (2024), [https://yjolt.org/sites/default/files/wachter\\_26yalejltech671.pdf](https://yjolt.org/sites/default/files/wachter_26yalejltech671.pdf)

Jennifer Wang, et al., *Assessing the Implementation of Federal AI Leadership and Compliance Mandates*, Stanford Univ. Human-Centered Artificial Intelligence (Jan. 2025), <https://hai.stanford.edu/assets/files/hai-reglab-white-paper-federal-ai-leadership-and-compliance.pdf>

---

## **Week 10 - The Tragedy of the Commons: Intellectual Property in the Age of AI**

We will discuss the intellectual property issues raised by AI. Starting with cases involving use of copyrighted works in training, we will understand the elements of infringement involved in training AI and the defenses raised by companies, including the Fair Use defense. We will discuss the policy reasons for and against protecting AI training under Fair Use and see how the courts have applied these factors in the most recent cases to emerge. We will also examine the issues involved with claiming ownership over AI-generated content, including cases involving the copyrightability of AI content and the patentability of AI-assisted invention. Finally, we discuss how data use agreements, intellectual property and privacy rights combine to create a new "data economy".

### **Readings:**



Kadrey v. Meta Platforms, Inc. (N.D. Cal. June 25, 2025),  
<https://law.justia.com/cases/federal/district-courts/california/candce/3:2023cv03417/415175/598/>

Thomson Reuters v. Ross Intelligence Inc. (D. Del. Feb. 11, 2025),  
[https://www.ded.uscourts.gov/sites/ded/files/opinions/20-613\\_5.pdf](https://www.ded.uscourts.gov/sites/ded/files/opinions/20-613_5.pdf)

Bartz v. Anthropic PBC (N.D. Cal. June 23, 2025) (ECF No. 231),  
<https://www.courtlistener.com/docket/69058235/231/bartz-v-anthropic-pbc/>

---

### **Week 11 - The Tragedy of the AntiCommons: Conceptions of Open Source AI**

AI has deep roots within the open source community. In this class we explore concepts of free software, open source software, open access, and open culture, and how they have shaped the development of artificial intelligence, and are in turn being shaped by it. We will examine the definitions of open source, including the OSI Open Source Definition, and the challenges and controversies faced by the open source community in adapting it to define Open Source AI. We will examine the impact that DeepSeek has had on the open source community and different strategies employed by China, the EU, and the US in the open source AI space. We will discuss how the struggle for dominance in open source AI may have deep and enduring implications for the future of AI development.

#### **Readings:**

Michael Heller, *The Tragedy of the Anticommons*, *The Wealth of the Commons*,  
<https://wealthofthecommons.org/essay/tragedy-anticommons>

Alek Tarkowski, *Data Governance in Open Source AI*, OPEN SOURCE INITIATIVE  
(Feb. 3, 2025), <https://opensource.org/data-governance-open-source-ai>.

The Harvard Law Review, *Co-governance and the Future of AI Regulation*, 138 Harv. L. Rev. 1609 (2025),  
<https://harvardlawreview.org/print/vol-138/co-governance-and-the-future-of-ai-regulation/>

---

### **Week 12 - The Imitation Game: Agents, Alignment, and Security**

We will examine the special challenges posed by AI agents, artificial intelligence systems that possess both agency and the power to act on behalf of their human or corporate principals. We will examine the issues related to single-agent and multi-agent safety and alignment, focusing on monitoring and model transparency, reward hacking, reinforcement learning with human feedback, deception, adversarial AI, and other alignment risks. We will look at the principles and practices of AI safety engineering and how regulation and legal liability regimes may influence the adoption and evolution of such practices (e.g., OECD AI Principles, UNESCO Recommendation on the Ethics of AI), including multi-stakeholder and local governance models. We will examine the challenges of governing adversarial AI agents in the context of cybersecurity, national security, biosecurity, and the human/AI future.

#### **Readings:**

Kate Crawford & Jason Schultz, *AI Systems as State Actors*, 119 Colum. L. Rev. 1941 (2019), <https://www.columbialawreview.org/content/ai-systems-as-state-actors/>

Ian Ayres & Jack M. Balkin, *The Law of AI Is the Law of Risky Agents Without Intentions*, U. Chi. L. Rev. Online (Nov. 27, 2024), <https://lawreview.uchicago.edu/online-archive/law-ai-law-risky-agents-without-intentions>

Dan Hendrycks, Mantas Mazeika, & Thomas Woodside, *An Overview of Catastrophic AI Risks*, arXiv:2306.12001 (June 21, 2023), <https://doi.org/10.48550/arXiv.2306.12001>

---

### **Week 13 - Emergence: Artificial General Intelligence and Beyond**

As AI systems become more complex, scaling laws predict improvements in AI capabilities as a result of ever larger models and larger training data sets. These scaling laws can be non-linear: potentially giving rise to unexpected abilities, called emergent properties, which are not observed with smaller networks. Emerging capabilities such as artificial systems for AI R&D automation (ASARA) may lead to artificial general intelligence (AGI) or even artificial superintelligence (ASI). We will examine different attitudes towards these risks, from doomers to accelerationists, and the legal, regulatory, and governance responses that attempt to prepare society for the potential rise of AGI or ASI.

#### **Readings:**

This week's classes will involve a simulation of future AI policy scenarios, where you will work in teams. Reading assignments will be based on your team assignment.