

680-001: AI and Legal Philosophy

Jack Kieffaber - jkieffab@gmu.edu

Course Overview: The legal academy is *just starting* to teach lawyers how to use AI. That's a question of where AI is now. But the largest untapped reservoir in legal scholarship today is on where AI is *going* — and where it's going is a radically new form of governance, both public and private. All major legal questions culminate in the debate over how that government should look; interpretive questions, natural law questions, structural questions, everything plays at this new frontier. This class is designed to place students at the cutting edge of that debate and personally acquaint them with its major players — so they can help write the Federalist Papers of the next American founding.

Introduction

Week 1 — The Algorithmic View of Governance

- An experimental paradigm that views the law as nothing more than a giant algorithm whereby every possible factual occurrence is connected to a legal outcome. We'll discuss how the American system goes about building this algorithm, what the algorithm's purpose should be, how algorithms can be assessed relative to a purpose, and how AI could be employed to produce a more complete algorithm.
 1. Kieffaber: *A Paradigm for Algorithmic Governance* (handout)
 2. *Riggs v. Palmer*, 115 N.Y. 506, 22 N.E. 188 (1889)
 3. **OPTIONAL:** Federalist 9, 10, 51, 62, 73

Judge.ai

Week 2 — Formalism: Sowing the Seeds of Automation?

- Is formalism a kind of contrived automation? If so, is that a good thing? If not, how is it distinct from other methodologies? And does it have any higher end than predictability? We'll explore the various candidates for formalism's lodestar and debate whether human morality is one of them.
 1. Kieffaber, [Predictability, AI, and Judicial Futurism: Why Robots Will Run the Law and Textualists Will Like It](#) (pp. 8–29)
 2. Newsom & Frederick, [Snails, Trains, and Pragmatist Claims](#)
 3. Fuller, [The Case of the Speluncean Explorers](#)
 4. Hart, [Positivism and the Separation of Law and Morals](#)
 5. Scalia, [A Matter of Interpretation](#)
 6. Michael Showalter, [Why Arbitrator.ai Is Desirable](#) (pp. 7–9, 12–15)
 7. Rao, [Textualism's Political Morality](#)
 8. Matey, [The Fearful and The Fainthearted](#)
 9. **OPTIONAL:** *Bostock v. Clayton Cnty., Georgia*, 590 U.S. 644, (2020); *Thomas v. Reeves*, 961 F.3d 800, 825 (5th Cir. 2020); Hayek, *The Constitution of Liberty*, 30–36 (1960); Hume, *An Inquiry Concerning Human Understanding*, 1–20.
- [Available guests:
 1. Hon. Paul B. Matey (US Federal Court of Appeals for the Third Circuit)
 2. Michael Showalter (author of several AI articles, rebutted my predictability thesis in the Notre Dame Law Review. Frequent debate partner of mine at law school events)]

Week 3 — Playing it Safe: AI as Interpreter

- Let's start small and measured, like most thinkers in this space have: What if AI was just a tool, like the dictionary, that judges could use to determine the meanings of words? On what grounds could a formalist object? What technological shortcomings might arise? We'll discuss the rise of AI interpretation and discuss (attack) its closest competitor, corpus linguistics.
 1. Arbel & Hoffman, [Generative Interpretation](#)
 2. Engel & McAdams, [Asking GPT for the Ordinary Meaning of Statutory Terms](#)
 3. *Snell v. United Specialty Insurance Co.*, 102 F.4th 1208, 1221 (11th Cir. 2024) (Newsom concurrence)
 4. *United States v. Deleon*, 116 F.4th 1260 (11th Cir. 2024) (Newsom concurrence)
 5. [Ross v. United States](#), 331 A.3d 220 (D.C. 2025) (Howard and Deahl concurrences)
 6. Lee & Eggbert, [Artificial Meaning?](#)
 7. **OPTIONAL:** Kieffaber & Mayesh, [The Dying Gasp of Corpus Linguistics](#)
- [Available guests:
 1. Austin Peters (Bigelow Fellow at UChicago and Newsom clerk during 2024, the *Snell* and *DeLeon* term)
 3. Shlomo Klapper (founder and CEO, of Learned Hand, an AI law clerk)]

Week 4 — Getting Dangerous: AI as Judge

- It's the obvious question: If AI can help the judge, why can't AI *be* the judge? And if AI can be the judge, shouldn't it just be the entire judiciary? What would that sort of government even look like? Is it lawful? Is it a Due Process violation? What laws would have to change? We'll discuss several models of AI judging — with most emphasis on the Judge.ai model of talking statutes.
 1. Kieffaber, [Predictability, AI, and Judicial Futurism: Why Robots Will Run the Law and Textualists Will Like It](#) (pp. 3–8, 29–53)
 2. Lawrence Solum, [Artificially Intelligent Law](#)
 3. Eugene Volokh, [Chief Justice Robots](#)
 4. Richard Re, [Developing Artificially Intelligent Justice](#)
 5. Adam Unikowsky, *In AI We Trust* pts. I and II.
 6. Posner & Saran, [Judge AI: Assessing Large Language Models in Judicial Decision-Making](#)
 7. **OPTIONAL:** Showalter, [The Honorable Robot: Why AI Adjudicators May Be Judiciously Virtuous](#)
- [Available guests:
 1. Adam Unikowsky (Jenner and Block partner, author of several influential pieces touting AI's potential as a jurist)
 2. Jameson Payne (Graduate Researcher, Hillsdale College, author of exciting new work on advisory opinions in the Judge.ai context)

Week 5 — Rabbit Hole: Judge.ai and Unwritten Law

- Would Judge.ai leave room for equity? For the common law and general law? Or would it codify all of this unwritten law into a fixed statute? And is such codification necessary for formalism to effectively function? We'll dive into the work of the modern general lawyers to decide the fate of *ex post* lawmaking under Judge.ai.
 1. Kieffaber, [Predictability, AI, and Judicial Futurism: Why Robots Will Run the Law and Textualists Will Like It](#) (pp. 38–45)
 2. [Erie R. Co. v. Tompkins](#), 304 U.S. 64, 58 S. Ct. 817, 82 L. Ed. 1188 (1938)
 3. Baude et al., [General Law and the Fourteenth Amendment](#)
 4. Baude & Sachs, [The Official Story of The Law](#)
 5. Sachs, [Finding Law](#)
 6. Campbell, [Tradition, Originalism, and General Fundamental Law](#)
- [Available guests:
 1. Jud Campbell, professor of law, Stanford University)]

Week 6 — Technical and Moral Objections to Judge.ai

- Judge.ai is a fun thought experiment. But can today's AI technology *actually do the job*? And even if it can, is there something morally impracticable about removing humans from the judicial process? We'll examine the strongest counterarguments and best rebuttals on both the normative and positive sides of the ball.
 1. Choi, [Off-the-Shelf Large Language Models Are Unreliable Judges](#)
 2. Grimmelman, [Generative Misinterpretation](#)
 3. Waldon et al., [Large Language Models for Legal Interpretation? Don't Take Their Word for It](#)
 4. Kieffaber et. al., [LLMs Are Bad Judges. So Use Our Classifier Instead.](#)
 5. Ted Folkman, [Against Whatever This Is](#)
 6. Kieffaber, [Predictability, AI, and Judicial Futurism: Why Robots Will Run the Law and Textualists Will Like It](#) (pp. 34–78)
- [Available guests:
 1. Kenny McLaren (CTO and chief software engineer, Fortuna Insights)
 2. Colin Rule (CEO of Online Dispute Resolution)]

Week 7 — Judge.ai in the Private Sphere: A Hands-On Demonstration

- My company, Fortuna Insights, has actually built a private version of Judge.ai — it's called Arbitrus, and it arbitrates contract disputes at rapid speeds following roughly the same processes we've discussed. This class session will be interactive; Fortuna's founder and CEO Kimo Gandall will present Arbitrus to the class, and students will submit briefs and litigate in front of it. Discussion will follow as to what the students liked and disliked about the process — and whether the role of the lawyer can be diminished to that of a mere prompter.
 1. Gandall et. al., [We Built Judge.ai. And You Should Buy It.](#)
- Guests:
 1. Kimo Gandall (Founder and CEO of Fortuna Insights, creator of Arbitrus)
 2. Kenny McLaren (CTO and chief software engineer, Fortuna Insights)
 3. Brian Potts (partner, Husch Blackwell)

Week 8 — Closing the Loop: The Future of the Legal Profession

- Products like Arbitrus suggest that we might already have the technology to create Judge.ai. But it will almost certainly be a while before America bends the knee and replaces the Article III judiciary. That leaves one final question: What will practice look like in the interim, when the tech exists but isn't institutionalized? As the formalist movement makes judges more homogenous and predictable, what does law practice look like when robots can accurately forecast winners and losers? When fully-functional judge bots exist parallel to federal and state courts in the private sphere? When judges and lawyers alike are asking legal questions to the same machines? We'll discuss the mercantilist (and feudalist) implications for biglaw, litigation finance, and business more generally — as well as professional ethics.
 1. Gandall et. al., [Predicting Precedent: A Psycholinguistic Artificial Intelligence in the Supreme Court](#)
 2. Kieffaber & Maxwell, [The Last Gold Rush: AI and the New Mercantilism](#)
 3. Croly, *The Promise of American Life*, (pp. 25–35)
 4. Kieffaber, [Spontaneous Disorder: The Protected Sphere and the Coming Web 3.0 Age](#), Part III
 5. *United States v. El Paso Natural Gas*, 376 U.S. 651, 656 (1964)
 6. *Anderson v. City of Bessemer City, N.C.*, 470 U.S. 564, 572 (1985)
 7. Rosenthal & Yoon, [Judicial Ghostwriting: Authorship on the Supreme Court](#)
- [Available guests:
 1. James Burnham (former general counsel of DOGE, founder of the AI Innovation Council, President of Vallecito Capital and King Street Legal)
 2. Pierce Babirak (General counsel and Executive Vice President, Vallecito Capital. Forward-thinking litigation financier)
 3. Paul Healy, (Director of Legal Engineering at Norm AI)
 4. Patricia Geli (Managing Partner, C10 Labs, an AI-specific VC fund)]

Leg.ai

Week 9 — Going Off the Deep End: AI as Lawmaker?

- Judge.ai is a statute that talks — it's self-executing. But what if the statute was also *self-creating*? A smart statute that evolved over time? How would that look in practice? And how could/should it be coded to make law? Should it isolate the democratic median, or strive for objective truth? We'll discuss a few paradigms for how Leg.ai might operate — each stemming from either a moralist or a nihilist view — and debate which view should animate the legislative process.
 1. Christiansen, [Congress.AI](#)
 2. Kieffaber et. al., [AI and The Ordinance of Reason](#)
 3. The Holy Bible
 - Genesis 1:26-27
 - Wisdom 6:1-6
 - Daniel 2:20-21, 3:1-30
 - John 1:1-4; 14
 - Romans 2:14-16, 13:1-7
 - Acts 5:29
 - 1 Peter 2:15-16
 - Galatians 5:1
 4. Nietzsche, [On the Genealogy of Morals](#), First Treatise
 5. Keller, [The Reason for Living](#) (audio)
 6. Lawrence Solum, [Artificial Meaning](#)
 7. Federalist 51
- [Available guests:
 1. Prof. Jeremy Christiansen (Professor of Law at Regent University, author of *Congress.ai*.)
 2. Garrett Snedeker (Deputy Director, James Wilson Institute)
 3. Hon. Paul B. Matey, US Federal Court of Appeals for the Third Circuit)]

Week 10 — God.ai, GoldenCalf.ai, and the Enumeration Paradox: Can the Natural Law be Written Down?

- A deeper dive into last week's concepts; things that can be written down are things that AI can digest and master. Is the natural law one of those things? If so, why not make your legislature GodAI? If not, is there even a natural law in the first place? Must a thing be enumerable to be knowable? We'll discuss the legislative and moral paradoxes of our enlightenment culture of enumeration.
 1. Aquinas, *The Unjust "Law"*, ST, I-II, Q 96, art. 4, 5, 6; ST, II-II, Q. 60, art. 1, 2, 3, 4, 5, 6
 2. Aquinas, *Natural Law*, ST, I-II, Q. 94, art. 1, 2, 3, 4, 5, 6
 3. Vermeule, *Common Good Constitutionalism* 1–51
 4. Arkes, *Mere Natural Law*, Ch. I–II.
 5. Finnis, *Natural Law and Natural Rights*, I.1, II.1, II.2, II.8
 6. Vermeule, *Law Skepticism on The New Right*
 7. **Optional:** Aristotle, *Nicomachean Ethics*, *Nicomachean Ethics*, Book I, Ch. IV - Ch. V (end of third paragraph); Book II, Ch. I – VI; Book V
- [Available guests:
 1. Prof. Adrian Vermeule (professor at Harvard Law School, prominent natural law theorist and author of *Common Good Constitutionalism*)
 2. Prof. Hadley Arkes (distinguished leader of the James Wilson Institute, a natural law institution)
 3. Prof. Gerard Bradley (co-director of the James Wilson Institute, professor emeritus at Notre Dame Law School)
 4. Conor Casey (lecturer at the University of Surrey, frequent Vermeule co-author and managing editor of *The New Digest*)]

Week 11 — A Demonstration: How to Code the Natural Law

- The next frontier of ethics and technology will be enumerating moral truth; the ethicists must decide what is to be coded, and the technologists must learn how to code it. Shaping the internal “conscience” of these bots will be hugely important as their role in our lives expands. Accordingly, we’re joined this week by Prof. Hanseok Ko — the Newton-Bennet Chair Professor of Electrical Engineering and Computer Science at Catholic University and the director of Catholic’s nascent *Institute for AI and Morality*. He’ll discuss the technical methods of enumerating the kinds of moral principles we’ve studied thus far, and we’ll get a technical demonstration — prompting, critiquing and comparing bots trained on various moral precepts.
 1. Kieffaber, *God.ai, GoldenCalf.ai, and The Enumeration Paradox*
- Guests:
 1. Prof. Hanseok Ko (Professor of Electrical Engineering and Computer Science at Catholic University who has spent a career studying and designing ethical AI models)

Exec.ai

Week 12 — The End of The World: AI as Executive

- If an AI can adjudicate the law... and AI can make the law... what can’t AI *physically execute* the laws? Would this entail some presidential apparatus? Or would Exec.ai simply be the legal code’s servant, carrying out the outcomes that accompany the various factual inputs? Is this simply a question of robotics, such that the AI can physically punish infractions in the real world? Or does Exec.ai more practically manifest in digital spaces, where a centralized entity already controls the relevant physics such that actors cannot physically violate the law in the first place? Is that a utopia? And where will Exec.ai manifest — publicly or (perhaps more terrifyingly) *privately*? We’ll take some wild swings in this session, indulging in doomy futurism and theorizing how close the contrived physics of law might get to the relative omnipresence and omnipotence of actual physics.
 1. The Holy Bible, [Revelations](#)
 2. Gandall et. al., *We Built Judge.ai. And You Should Buy It*, (pp. 82–88)
 3. Kieffaber, [Spontaneous Disorder: The Protected Sphere and the Coming Web 3.0 Age](#), Part IV.
 4. Kokotajlo, [Why the AI Race Ends in Disaster](#) (video)
- [Available guests:
 1. Daniel Kokotajlo (formerly of OpenAI, currently heads the AI Futures Project. Has garnered substantial media attention for his doomsday superintelligence forecasts)]

Week 13 — Back to Reality: AI and The Administrative State

- Matrix-style Exec.ai is probably unrealistic. But an automated administrative state is already under construction — and, since the modern administrative state tends to consolidate judging, legislating, and executing into one entity, it's the perfect subject with which to sum up the course. Are automated ALJ's feasible? Constitutional? If yes to both, how slippery is this slope? And if an AI can adjudicate, can it rulemake? Indeed, could it perform both functions in an Article I capacity to remove all nondelegation worries from our system? Would that satisfy the administrative state's detractors? We'll close with *Marbury v. Madison* and ponder how the separation of powers might be credibly rearranged to facilitate the algorithmic model.

1. *Gundy v. United States*, 588 U.S. 128 (2019) (Gorsuch dissent)
2. Friendly, [Some Kind of Hearing](#)
3. *Heckler v. Campbell* 461 U.S. 458, 460–61(1983)
4. *Mathews v. Eldridge*, 424 U.S. 319, 348 (1976)
5. *Bridges v. Wixon*, 326 U.S. 135 (1945)
6. *Accardi v. Shaughnessy*, 347 U.S. 260 (1954)
7. *Morgan v. United States*, 298 U.S. 468, 482
8. Gandall et. al, *Does Due Process Require A Human Adjudicator?* (handout, excerpt from Fortuna Insights whitepaper)
9. Gorsuch, *Over Ruled*, (chs. I, III)
10. *Marbury v. Madison*, 5 U.S. 137, 2 L. Ed. 60 (1803)
11. **OPTIONAL:** Vermeule, [What is “The Unitary Executive”?](#)

- [Available guests:

1. Austin Raynor (Former General Counsel, DOGE)
2. Eli Nachmany (Covington & Burling associate, rising star administrative law scholar)]