

The Philosopher

Alumni shape the emerging industry of AI law

 Legal

 Tech

 Society

● By Susie Allen

In February 2018, Danny Tobey '03 presented a paper at the inaugural Conference on Artificial Intelligence, Ethics, and Society that made a bold claim. Tobey argued that artificial intelligence (AI) could potentially supplement or replace the work of doctors, engineers, architects, and yes, even lawyers. As a result, he suggested, software companies would be forced to contend with the complex liability issues these specialized professions have long faced.

“It was a controversial paper at the time,” recalled Tobey, who is now the global co-chair of DLA Piper’s AI and data analytics practice. But nearly 10 years on, the questions he raised are indeed playing out in boardrooms and courtrooms around the world. Tobey himself has been involved in some of the early generative AI litigation, and, he said,

in the Machine

YLS

“we really are seeing some of those trends that I was predicting, along with the defenses and mitigations.”

As AI moves from the periphery to the center of society and the legal profession, ideas that once sounded like science fiction — software that can diagnose disease, generate complex documents, and analyze reams of evidence — have become contemplatable (and, in some cases, already arrived). In the face of rapid change, alumni including Tobey; Cynthia Cwik '87; Rebecca Crootof '11 JD, '16 PHD; and David Robinson '12 are using their Yale Law School educations to shape how this transformative new technology can best be used to benefit society.



The speed of technological transformation has intrigued Tobey since his time at the Law School. As a student in the early 2000s, he witnessed the transition from paper notetaking to omnipresent laptops — a shift that motivated him to develop legal education software. He eventually traded technology entrepreneurship for the practice of technology law and has become a sought-after expert for companies seeking to mitigate the risks of new innovations.

Contending with the thorny ethical and legal questions posed by AI sits at the heart of Tobey’s work today. He’s involved in the first state attorney general settlement over generative AI accuracy — but even figuring out how to quantify accuracy for something that is “as open-ended and probabilistic as generative AI is an incredibly difficult legal and scientific and philosophical challenge,” he said.

Beyond litigation, Tobey’s team, which includes both lawyers and data scientists, is also developing new ways to test legal compliance in AI models. They’ve adapted the cybersecurity practice of “red teaming” — essentially, simulated hacking to uncover a system’s vulnerabilities — to identify legal risks in the responses provided to users. “It’s been a very successful endeavor,” Tobey said, and has helped companies build important new safeguards into their models. The issues at play have unfolded on a global scale, and recently, Tobey was named senior consultant to the United Nations on Parliamentary Engagement on Artificial Intelligence, helping countries address AI.

Tobey says his years at Yale Law School provided an essential foundation for his current work — though he didn’t know it at the time. “To be an AI lawyer, you really have to be a philosopher, because you can’t wrestle with the issues that AI raises without very quickly getting into pretty heavy questions of what is truth? How do we measure it? What is a product versus a service? What does it mean to have intelligence or judgment?” he said. “So the philosophical bent of Yale Law School prepared me to help extend the law into an area where some fundamental assumptions have to be examined and can’t be avoided.”

Cwik, too, is fascinated by the range of skill sets and approaches that technology law demands. Now an arbitrator

and mediator with JAMS, Cwik has spent much of her career focused on legal issues related to science and technology. That’s required a lot of on-the-job learning: One of her earliest cases, for example, involved understanding the nuances of sophisticated immunology tests.

Cwik traces her interest in science and technology back to her childhood in Pittsburgh, Pennsylvania, where she attended a public elementary school that had established a partnership with the University of Pittsburgh and provided her with early access to computers. “I remember one of my proud moments in elementary school was when I was able to beat the computer at tic-tac-toe,” she said. This early exposure to computers, she believes, helped her “not be afraid to experiment and see what they’re capable of.”

Cwik was able to continue to develop her knowledge about artificial intelligence in 2018. During a yearlong fellowship at Stanford, she took a class on AI and the law taught by Mariano-Florentino Cuéllar ’97, former associate justice of the Supreme Court of California. “It was the first time this



Danny Tobey

Danny Tobey '03:

“To be an AI lawyer, you really have to be a philosopher, because you can’t wrestle with the issues that AI raises without very quickly getting into pretty heavy questions of what is truth? How do we measure it? What is a product versus a service? What does it mean to have intelligence or judgment?”

Cynthia Cwik '87:

“I remember one of my proud moments in elementary school was when I was able to beat the computer at tic-tac-toe. [This early exposure to computers helped me] not be afraid to experiment and see what they’re capable of.”

class was offered, and it was fascinating to hear his perspective,” she recalled. She has continued studying and thinking about AI ever since, including serving on the Planning Committee for the American Bar Association (ABA) Artificial Intelligence and Robotics National Institute since its inception in 2020 and serving as the vice-chair of the ABA Task Force on Law and AI.

Cwik brought those years of expertise to “Artificial Intelligence: Legal Issues, Policy, and Practical Strategies,” a book she co-edited alongside Christopher A. Suarez ’11 and Lucy L. Thomson. “We had an interdisciplinary approach, which I think is important,” she said — contributors included not just lawyers but also voices from computer science, government, and the nonprofit sector. “We really tried to be comprehensive,” Cwik said, exploring how AI may shape areas including legal education, the judiciary, employment law, intellectual property, privacy, and national security.

Thinking about and studying these issues deeply has helped Cwik recognize both the peril and the potential of AI. While she believes it’s essential to recognize the limitations and risks of AI tools, she also sees powerful opportunities to help lawyers in their work — and even to increase access to justice for nonlawyers.



Cynthia Cwik

Cwik tries to share this cautiously optimistic perspective when she speaks to law students and early-career lawyers, many of whom are understandably anxious about how AI might affect their jobs. She encourages them to stay open-minded and continually experiment with AI, as well as other new tools that may emerge. After all, she says, the legal technology issues of the future might be very different from the ones that preoccupy us today, something she has learned firsthand: “I had no idea when I was at Yale Law School that I would take the direction I did.”

Crootof followed a similarly unexpected path into technology law. She arrived at Yale Law School planning to study civil rights law but shifted her focus after taking a seminar on emerging issues in international law taught by Professor Oona Hathaway ’97. Crootof was surprised to discover she found herself most drawn to questions related to technology: When does a cyber operation constitute an armed attack that justifies a responsive use of force? What is required to lawfully use drones in another country’s territory? These issues “were really exciting for me,” she said, “because it’s fascinating to see how law and technology foster each other’s evolution, and to think about when and how we can proactively direct it.”

Crootof’s interest deepened during her doctoral studies. Hathaway often encouraged Crootof to join her at high-profile events about issues at the forefront of human rights and national security. One of these was focused on autonomous weapon systems, which ended up becoming a major research interest of Crootof’s. The event was something “I would never have been invited to,” she said, but ultimately “set my scholarly agenda for the next 15 years.”

Crootof also credits her time at Yale Law School’s Information Society Project, where she was a fellow and executive director, with granting her a more nuanced understanding of the relationship between law, technology, and society. “Jack Balkin is truly the Godfather of Technology Law — he created a space where a community of techlaw scholars could connect and grow,” Crootof said. Along with BJ Ard ’10 JD, ’17 PHD, she is now publishing an open-access technology law coursebook.

Crootof has continued to explore the intersection of technology, national security, and international law as a law professor at the University of Richmond. Recently, she’s



Rebecca Crootof

continued her study of human/machine decision-making processes, such as the use of AI in gathering and processing intelligence. In some regards, it's an ideal match between technology and application — “one of the things AI is good at is processing superhuman amounts of data at superhuman speeds” — but also fraught. “Done well, it could be an incredibly useful augmentation tool,” Crootof said. “Done poorly, it could lead to a lot of accidents where nobody actually acts with criminal mens rea, but terrible, terrible things happen.” Crootof has long argued that autonomous weapon systems and other new military technologies have made this accountability gap in armed conflict more salient, and she has proposed a “war torts” legal regime that would establish a route to redress for harmed civilians.

Maximizing benefits while mitigating or preventing harm was the focus of Crootof's work last year as the inaugural Ethical, Legal, and Societal Implications (ELSI) Visiting Scholar at the Defense Advanced Research Projects Agency (DARPA). This small agency has had an outsized influence far beyond the military, creating everything from the original internet protocols to the first autonomous vehicles to the software that powers today's voice recognition systems. DARPA started the ELSI visiting scholar position because it is “aware

that the choices it makes in the military R&D context have a broader impact on society,” Crootof explained.

Her presence initially attracted some skepticism. “I showed up and everyone was like, are you the fun police?” she recalled. But she worked hard to convince her DARPA colleagues that considering the potential downstream effects of a new technology, good and bad, at the design stage would ultimately make that technology better — not just morally better, but practically and operationally better. “I had to win people over, and it helped that I had concrete examples of how thinking through uses and misuses and implications ahead of time could help improve the product itself, could make it more capable, or minimize accidents,” she said.

By the end of her year at DARPA, Crootof had designed an ELSI evaluation process that all new DARPA programs undergo — and found influential allies in DARPA's service liaisons, from each branch of the military. If people who have actually deployed believe incorporating ELSI matters, Crootof says, “I can take a lot of heart that this isn't just an academic exercise.”

Robinson, too, is focused intently on issues of technological safety. As part of OpenAI's safety systems team, he is



David Robinson

Rebecca Crootof '11 JD, '16 PHD:

“Done well, [AI] could be an incredibly useful augmentation tool. Done poorly, it could lead to a lot of accidents where nobody actually acts with criminal mens rea, but terrible, terrible things happen.”

David Robinson '12:

“Having strong views about AI in the future and what it should be...is healthy. This shouldn't be something that we, or companies like us, do alone.”

working to raise awareness of the safeguards built into the company's technology. Robinson is also tasked with helping design and communicate the company's efforts to mitigate catastrophic risks, such as the use of AI by malicious actors or AI models that go rogue.

With each new launch, OpenAI releases a technical report on its safety measures. Robinson's job, he explained, is to “make those documents as strong and useful as we can.” For example, OpenAI regularly tests its models to see how adept they are at hacking; the company then communicates the methods and results of those tests to users. Building greater understanding of how this safety testing works and what it uncovers, Robinson said, “is a foundation for trust.”

Facilitating trust in — and *trustworthy* — technology is a longstanding interest of Robinson's. While he was a student at Yale Law School, Robinson cofounded Upturn, a nonprofit that promotes equity and justice in the design, governance, and use of digital technology. He also wrote a book, “Voices in the Code: A Story About People, Their Values, and the Algorithm They Made,” that describes how a diverse group

of stakeholders came together to build a new transplant matching algorithm.

These experiences helped him see that what people want from a technology can change in meaningful ways as they learn more about it and interact with other users of it. Figuring out how to create opportunities for that kind of collaboration and consensus-building is “an important question as a human being and citizen and certainly as somebody working on AI,” he said. After all, “having strong views about AI in the future and what it should be ... is healthy. This shouldn't be something that we, or companies like us, do alone.”

Robinson sees his time at Yale Law School as essential to the work he is doing today. “One thing that Yale encourages people to do that has been useful is to think about institutional architecture,” Robinson said. That's been vital, he said, for someone whose work involves thinking about how to govern new things in new ways. At Yale Law School, he learned how to think about “not just, how do you make today's law work for the people who need it — which is vital — but how else could it be?”

AI and ChatGPT Push Students and Faculty to Take a Deeper Look at the Law

As AI continues to overtake almost every sector of modern society, Yale Law School is utilizing this moment to encourage its students to explore AI's implications and prepare them for an ever-changing technological world and legal landscape.

In the classroom, students in Knight Professor of Constitutional Law and the First Amendment Jack M. Balkin's reading group, “Technology Law: From Accidents to AI,” are examining how the introduction of different technologies have shaped the law and legal system throughout time.

Mikey McGovern and Isaac May '24, fellows in the Information Society Project, run the reading group under Balkin's supervision.

Shah Khan '27, a member of the reading group, said he came to the Law School to study how the law has historically adapted to what he calls “disruptive technologies”

to understand how the law might change to accommodate the rise of generative AI.

“YLS has been a great place to explore this topic because of the incredibly accomplished and thoughtful professors,” said Khan. “Last semester, I led a reading group on AI and the law, and I took a class with a visiting professor on privacy and data governance. In the Private Law Clinic, I've had the opportunity to think about how to help plaintiffs bring claims against certain AI companies. And, this semester, I'm participating in the technology law reading group, which traces how novel technologies have changed our legal system — exactly what I came here to study.”

These class discussions are bolstered through extracurricular programs at the Law School. During the fall 2025 semester, the Lillian Goldman Law Library

sponsored a Critical Legal AI Literacies Series, a speaker series that welcomed scholars from a variety of fields to the Law School to share their perspectives on generative AI and how they relate to law.

The series, held from Sept. 9 through Oct. 14, included University of Washington Professor of Linguistics Emily Bender speaking on “Large Language Models and the Lawyer's Search for Meaning”; University of Maryland Francis King Carey School of Law Professor of Law Chaz Arnett speaking on “Legal AI and the Permanence of Racism”; and American Library Association Office for Intellectual Freedom Deputy Director Sarah Lamdan speaking on “Data Cartels: The Companies that Control Legal AI.” The program was made possible with the support of the Law School's Oscar M. Ruebhausen Fund.