



Annual Report

2025

INSTITUTE
FOR LAW & AI

CONTENTS

3	Executive Summary
4	Message from the Director
6	Research
7	US Law and Policy
10	EU Law
12	Legal Frontiers
15	Fellowships
16	Summer Research Fellowship
19	Events
20	Summer Institute on Law and AI
21	Workshop on Law-Following AI
22	Cambridge Forum on Law and AI
23	Finances
24	Team
28	Looking Ahead

The Institute for Law & AI (LawAI) is an independent legal research organization focused on artificial intelligence. We aim to produce sound legal analysis that promotes security, welfare, and the rule of law in the age of AI.

EXECUTIVE SUMMARY

The Institute for Law & AI (LawAI) engages in three core activities:

- Independent legal research and analysis
- Consulting for governments, academia, and industry
- Field-building through fellowships and events

In 2025, our activities grew substantially. Our core team grew from 20 to 28 staff. Applications to our Seasonal Research Fellowship were 9.5x higher than in 2024. We also launched three flagship events, featuring 200+ participants and 50+ experts and moderators.

Meanwhile, our Director, Christoph Winter, served in a legal advisory role to the European AI Office on the General-Purpose AI Code of Practice. He also testified before the European Parliament, offering recommendations on law-following AI and institutional capacity-building at the European AI Office.

Our US Law & Policy team produced 7 new publications, including a report on mobilizing technical talent into government for AI-related national security crises, and discussed implications of the Executive Order on AI preemption with CNN, the Financial Times, Politico, NBC, and The Verge.

Our Legal Frontiers team published a pioneering research article on law-following AI, and has begun to catalyze a field of technical, legal, and policy research to develop the idea more fully, including through our own subsequent research and our Workshop on Law-Following AI.

In 2026, our work will emphasize “radical optionality”—research and policy approaches designed to preserve flexibility amid uncertainty about AI’s trajectory. Our EU Law team will continue to publish new chapters of the *Cambridge Commentary on EU General-Purpose AI Law*, while our Legal Frontiers team will expand its new Automated Governance workstream. Our fellowship programs are expanding across all tracks with our inaugural Winter Research Fellowship.

30+

Publications total in 2025

This includes our research article introducing law-following AI (*Fordham Law Review*) and a book on AI governance architectures (Oxford University Press).

40%

Growth in full-time core staff

Our total headcount grew from 20 to 28 staff during 2025. We expect to grow our staff by a further 50% during 2026.

9.5x

Increase in fellowship applications

Our Seasonal Research Fellowship brings students and professionals to the leading edge of AI, law, and policy. In 2025, we received 6,089 applications (vs. 638 in 2024).

9.6

Average likelihood to recommend (EU SRF)

Our EU Law Summer Research Fellowship received an average score of 9.6 for the likelihood to recommend the program to colleagues and friends.

Message from the Director



Uncertainty defined the past year in AI, law, and policy. Above all, this uncertainty reflects the pace and unpredictability of AI progress itself. Rapid advances have validated long-standing concerns about various risks while compressing the time available for careful intervention. Policymakers face a sharpening dilemma: act too early and risk stifling innovation, or act too late and risk losing control.

This uncertainty has been compounded by institutional and political strain. Major governance bodies remain understaffed and under-resourced. To illustrate what the current resource gap looks like: Recent reports suggest Meta offered one AI researcher a salary package of €190 million. By contrast, three of the most prominent public institutions tasked with overseeing the industry – the US Center for AI Standards and Innovation (“US CAISI”), the European AI Office, and the UK AI Security Institute – each operate on annual budgets smaller than that. (US CAISI operates on a budget of roughly \$10 million, or twenty times less than the compensation reportedly offered to a single researcher.) Meanwhile, jurisdictions struggle to coordinate regulatory approaches, and strategic competition between major powers continues to accelerate AI development. Together, these pressures make effective law and policy more difficult and more urgent.

At the same time, it is an age of profound opportunity: 2025 has shown that there is significant demand for timely, trustworthy, and sound legal analysis to promote security, welfare, and the rule of law. We aim to fill this demand through three complementary efforts. First, we conduct research and advisory work that helps policymakers, regulators, and industry make decisions that are legally sound, technically informed, and institutionally realistic. Second, we develop talent through fellowship programs that address acute expertise shortages across many institutions. Third, we convene scholars and practitioners across disciplines and political perspectives to identify and tackle the most pressing issues at the intersection of AI, law, and policy.

In 2025, our research priorities were shaped by both rapidly evolving policy debates and the need for independent, foundational legal analysis. In the US, our team engaged closely with policymakers, briefing Congressional staff on the proposed federal AI moratorium and contributing to public debate on the Executive Order on AI preemption. In the EU, I served in a legal advisory role to the European AI Office on the General-Purpose AI Code of Practice – the world’s first comprehensive legal framework for general-purpose AI model compliance.

As LawAI enters its eighth year, I have been encouraged by how much the field, and our role within it, has developed since our early discussions in 2019.

I also testified before the European Parliament, offering recommendations that it strengthen the European AI Office and require AI systems to be law-following.

This latter recommendation – that AI systems ought to be architecturally constrained to refuse illegal actions – emerged from research by our Legal Frontiers team, which continues to develop new legal paradigms for an era of increasingly capable AI systems. This included the publication of its article introducing law-following AI in the *Fordham Law Review*, and our book on the architectures of global AI governance with Oxford University Press. One common thread runs across these efforts: translating fast-moving technical developments into legally grounded, institutionally realistic policy options.

Alongside our research and advisory work, we continued to invest in talent. Demand for our Seasonal Research Fellowships reached record levels in 2025, attracting nearly ten times as many applications compared with 2024. Alumni of our programs now work across leading public institutions, AI labs, academia, and think tanks. We also convened hundreds of experts through our events, including the Workshop on Law-Following AI, the Cambridge Forum on Law and AI, and the Summer Institute on Law and AI.

Looking ahead, uncertainty is unlikely to abate. To escape the dilemma of regulating too early or waiting too long, LawAI's work in 2026 will emphasize a third path: *radical optionality*. As AI systems continue to evolve, we will focus our research and consulting on approaches designed to preserve flexibility across technical, institutional, and regulatory choices. Our US Law & Policy and EU Law teams will advance this model in their respective jurisdictions. In parallel, our EU Law team has begun to publish a major commentary on general-purpose AI regulation under the EU AI Act, in collaboration with the University of Cambridge. Our Legal Frontiers team will continue exporting the results of its new Automated Governance workstream. Institutionally, we are preparing for a period of growth to support this expanding agenda, with plans to add capacity across research, programs, and operations. We expect to grow our staff by 1.5× annually over the next two years.

As LawAI enters its eighth year, I have been encouraged by how much the field, and our role within it, has developed since our early discussions in 2019. The legal and policy decisions made in the coming years will plausibly shape humanity's trajectory for decades to come. I believe LawAI's work is only becoming more urgent. I am deeply grateful to our donors,

partners, and colleagues, whose support makes our independent, evidence-based work possible. We remain committed to promoting security, welfare, and the rule of law in the age of AI – and to meeting uncertainty with seriousness, rigour, and care.



Cristoph Winter

Director & Founder, Institute for Law & AI
Assistant Professor of Law and AI,
University of Cambridge

Research

Through our research and consulting, LawAI is creating legal infrastructure for transformative AI across three specialized teams:

US Law and Policy

EU Law

Legal Frontiers

This structure – combining teams focused on immediate policy needs with one dedicated to breakthrough research – enables us to both respond to today’s urgent challenges and build the frameworks needed for tomorrow’s AI systems.

Through our applied research, we help policymakers, regulators, think tanks, industry, and advocacy organizations to think about and craft legal frameworks and policy. We work with organizations and people across political lines, and we are not affiliated with any political party.

US Law and Policy

Our US Law and Policy team provides practical legal research to improve the quality and clarity of current AI policy efforts in the United States.

The team currently focuses on three workstreams:



Mackenzie Arnold
Director of US Policy

Institutions and Procedures:

This workstream aims to produce insights into the institutional and procedural factors that impact the effectiveness of AI regulation. The capabilities that advanced AI systems will possess in the future, and the regulatory challenges that those capabilities will pose, are difficult to predict. But even when the best course of action is unclear in the present, we can equip regulators with the tools, procedures, skills, information, and authority to develop better answers over time. Relevant questions in this workstream include: What authorities do regulators already have to enact regulations, gather information, monitor for compliance, and enforce the law? What new authorities would they benefit from, and how could they be created?

Law and Compute:

Advanced AI models use a large amount of specialized computer hardware resources (“compute”) during development and deployment. Compared to the other primary inputs to AI systems – data and algorithms – compute is uniquely detectable, excludable, and quantifiable. Additionally, the supply chains that facilitate the manufacture or acquisition of compute are unusually concentrated. Compute is therefore a promising policy target for AI governance, and this workstream aims

to help policymakers leverage compute governance in an effective and balanced way. Specific research questions include: What role can compute thresholds play in AI governance efforts? How can we design measures of compute usage that are robust to attempts at circumvention?

Liability and Insurance Law:

Whereas many legislative or regulatory approaches create ex-ante requirements, liability and insurance impose ex-post penalties after harm has occurred. There are benefits to this ex-post approach: by waiting for harm to occur, liability and insurance can tailor relief to the specific circumstances that develop. But this focus on retroactive relief is only part of the story. Liability and insurance also set forward-looking incentives, such as encouraging the development of professional standards and safety procedures. This workstream aims to address questions like: How will existing liability law and insurance markets operate by default? What will they do well, and what gaps will emerge?

Our real-time analysis and strategic advice has proved central to [key debates](#) in AI policy, and made us a go-to resource for policymakers, industry, and civil society.

2025 Activities

Promoting adaptive governance

Uncertainty around AI's trajectory invites a dilemma for those working on policy and governance more broadly: regulate too early and stifle innovation, or wait too long and risk losing control entirely. In 2025, we worked on developing a third option: adaptive governance. By focusing our research and consulting on work that deliberately preserves flexibility across technical, institutional, and policy choices, we hope to make it easier for governance to rapidly adapt to future circumstances.

This includes work that helps the US government learn, coordinate, reason, and respond, such as our [research](#) on preparing the US government to respond to AI-related national security crises. It also includes our engagement with the Draft Report of the Joint California Policy Working Group on AI Frontier Models. In April, we submitted feedback to the Working Group in relation to its [scoping and definitions](#), [whistleblower protections](#), and [liability and insurance](#). We recommended substantive changes, and we were satisfied with the content of the [final report](#), published in June, which should more effectively inform California policy going forward.

By [increasing option value](#) and addressing uncertainty, we put the government in position to make informed decisions over time as we've expounded upon in our recent essay on radical optionality.

Providing expert advice to policymakers

We continued to provide consulting services to state and federal legislatures working at the intersection of AI, law, and policy. Our expertise in US and international law has helped legislatures make their work more informed, concrete, and feasible in the face of rapid political and technological change. This year, our services included drafting model legislation and regulation, reviewing proposed policy, and briefing legislative staff on discrete legal questions.

Weighing in on the proposed federal moratorium

The US continues to grapple with a debate over preempting state AI laws. There is, however, considerable confusion amongst policymakers and advocates, especially with respect to the federal moratorium on state AI regulation proposed in the "Big Beautiful Bill." We were pleased to help resolve some of these uncertainties via our [commentary](#), which concluded that the proposal raised significant legal issues under the Senate's Byrd Rule. In June, we found ourselves becoming the "go-to advisor" on everything from BEAD funding to parliamentary procedures to litigation risks facing states. We briefed Congressional staff and our memos were widely circulated and reported on in the [media](#). Our Director of US Policy, Mackenzie Arnold, and the team's Senior Research Fellow, Charlie Bullock, also published nuanced [analysis](#) ahead of the Senate vote.

Uncovering the legal implications of the Executive Order on AI preemption

In November, we published [commentary](#) on the leaked Executive Order on AI preemption, expressing doubts around the AI Litigation Task Force's legal grounds and raising concerns around the Executive Order's provision to restrict or withhold federal grants. Our Senior Research Fellow, Charlie Bullock, was subsequently invited to discuss the potentially chilling effects on state legislation of this Executive Order with [The Verge](#). Bullock also published a [piece in Lawfare](#), examining how the federal government could approach AI preemption, and arguing for an iterative and gradual approach that involves legislative preemption of narrow categories of state law in exchange for implementation of substantive federal AI policy.

Helping major news outlets understand and report on the legal implications of AI policy

Meanwhile, our Director of US Policy, Mackenzie Arnold, spoke to CNN, the [Financial Times](#), [Politico](#), [NBC](#), and other outlets about the legal implications of the Executive Order and the legal authority of the Trump administration to enact such an order. Alongside our features in [The Verge](#), these media collaborations serve to improve the public discourse around AI policy and support LawAI's mission to promote security, welfare, and the rule of law.



Featured publications

Building AI Surge Capacity: Mobilizing Technical Talent Into Government for AI-Related National Security Crises →

Cara Labrador, Shefali Agrawal, Alex Jumper, Joe O'Brien, Mackenzie Arnold, Zoe Williams, Nandidi Shiralkar

AI Federalism: The Right Way to Do Preemption →

Charlie Bullock

The AI Moratorium-the Blackburn Amendment and New Requirements for “Generally Applicable” Laws →

Mackenzie Arnold, Charlie Bullock

The Role of Compute Thresholds for AI Governance →

Matteo Pistillo, Suzanne Van Arsdale, Lennart Heim, Christoph Winter

EU Law

On 1 August 2024, the landmark EU AI Act Regulation 2024/1689 (“AI Act”) came into force across all 27 EU Member States, creating the world’s most comprehensive legal framework to regulate artificial intelligence.



Emily Gillett
Interim Director of EU Law

The significance of the AI Act is heightened by the stalled development of some comparative initiatives, like SB-1047, and the ongoing regulatory uncertainty in many other countries. As the first of its kind, the AI Act can promote safe, human-centric and trustworthy AI. It also stands the chance of global standard-setting, echoing the “Brussels effect” seen in other regulatory domains, such as data protection where the General Data Protection Regulation is – for better or for worse – a global benchmark. Therefore, high-quality legal research into the AI Act has real potential for impact: robust insights may not only support an effective EU regulatory approach, but also trigger regulatory diffusion and shape the trajectory of global AI governance.

The AI Act is a complex and ambitious legal instrument: it comprises 180 Recitals, 113 Articles, and 13 Annexes, which impose horizontal obligations across the AI value chain i.e., for providers, deployers, importers, and distributors of AI models and systems with a nexus to the EU internal market. Many of these obligations become applicable in the near future.

Our current research focuses on general-purpose AI models (“GPAI models”) (Chapter V AI Act) and related enforcement issues (Chapters IX and XII AI Act) given the harm potential, technological novelty, and consequential regulatory uncertainties. The chapters of the AI Act we focus on primarily concern systemic risks posed by advanced AI models as well as the responsibilities and procedures for assessing and mitigating those risks. Furthermore, the AI Act provides for the issuance of secondary instruments – such as codes of practice, common rules, guidelines, and harmonised standards – which raise important legal questions given that they aim to contribute to the objectives and proper application of the Regulation.

The team’s work has contributed significantly to AI regulatory implementation in the European Union and will continue to be influential as the enforcement powers concerning GPAI under the EU AI Act become applicable in the latter part of 2026.

2025 Activities

Advising the European AI Office on the EU GPAI Code of Practice

Our Director, Christoph Winter, performed a unique legal advisory role for the European AI Office concerning the drafting of the Safety and Security Chapter of the EU GPAI Code of Practice (“CoP”) that was published in July 2025. The CoP is a voluntary guidance framework designed to help developers of general-purpose AI models to comply with the EU AI Act. While voluntary, its significance is not to be understated. It has been signed by a wide range of leading AI providers operating in the EU (including OpenAI, Anthropic, and Google) and may have implications beyond the EU if other jurisdictions address requests to harmonize AI regulation.

Researching and writing a commentary on the EU GPAI Law

In August 2024, the EU AI Act came into force across all EU Member States. With a view to providing up-to-date developments and analyses of the EU AI Act, to inform both governance bodies and legal scholars, our team has begun writing a specialist legal commentary. Rigorous legal research into the provisions of the EU AI Act relevant to GPAI, and their effective enforcement, was therefore the focus of our EU Law team in the latter half of 2025. This has been a significant undertaking, involving a dozen independent contributors and a team of in-house editors with expertise in EU law. In collaboration with the University of Cambridge, we began publishing the *Cambridge Commentary on EU General-Purpose AI Law* in early 2026.



Making recommendations to the European Parliament on AI agents and democracy

In July, our Director and Founder, Christoph Winter, was invited to speak before the European Parliament’s Special Committee on the European Democracy Shield with participation of IMCO and LIBE Committee members. Professor Winter was asked to present on AI governance, regulation, and democratic safeguards. He **spoke** about the challenges that AI agents may present and how democracies could approach these challenges. In doing so, he made two substantive recommendations. First, introduce Law-Following AI:

“Law-following AI is democratically compelling for three reasons: First, it is democratically legitimate. Laws represent our collective will, refined through democratic deliberation, rather than unilaterally determined corporate values. Second, it enables democratic adaptability. Laws can be changed through democratic processes, and AI agents designed to follow law can automatically adjust their behaviour. Third, it offers a democratic shield — because without these constraints, we risk creating AI agents that blindly follow orders, and history has shown where blind obedience leads.”

Second, strengthen the European AI Office:

“The AI Office currently lacks the institutional capacity to develop these crucial capabilities over time [...] To illustrate what the current resource gap looks like: Recent reports suggest Meta offered one AI researcher a salary package of €190 million. The AI Office — tasked with overseeing the entire industry — operates on less.”

Preparing the next generation of AI law leaders

We recognize that the scale of AI’s legal challenges far exceeds what any single organization can address. That is why we have made identifying, training, and connecting exceptional legal minds a core part of our mission. To that end, we are pleased that two former LawAI Research Scholars began working at the European AI Office in 2025, and two more have accepted offers to join in 2026. A former Summer Research Fellow will also begin working at the European AI Office in 2026. These appointments consolidate, and further illustrate the success of, our work on this aspect of our mission, whilst other alumni from our programs have taken positions at leading AI companies, academia, and think tanks across the world.

Legal Frontiers

The legal and policy issues raised by advanced AI are profound. LawAI therefore finds it fruitful to approach them from multiple angles.



Cullen O'Keefe
Director of Research

The US and EU teams tend to focus largely on law and policy questions tethered to existing legal and policy debates, such as the proper application of existing laws or live policy debates. The Legal Frontiers team takes a different, complementary approach: We focus on *articulating, incubating, and scaling* new legal and policy paradigms that tackle the largest legal and policy challenges we expect to confront as AI progress continues. We aim to combine creative policy design with analytical rigor to produce policy proposals that are:

- **Anticipatory**, in that they respond to a reasonable forecast of the legal and policy challenges that further advances in AI will produce;
- **Actionable**, in that we can make progress within these workstreams even under significant uncertainty;
- **Accommodating** to a wide variety of worldviews and technological trajectories, given the shared challenges that AI will create and the uncertainties we have about likely developments; and
- **Ambitious**, in that they both significantly reduce some of the largest risks from AI while also enabling society to reap its benefits.

Our past work on the **Chips for Peace framework** for international AI governance is a classic example. This year, we focused primarily on **Law-Following AI**. We expect to work on scaling the Law-Following AI paradigm for the foreseeable future, while also making a concerted effort to develop **new workstreams** that are similarly ambitious.

We have also assembled a talented roster of team members and external contributors who share our passion for producing clear legal analysis that helps humanity navigate the coming AI revolution. Even as our team continues to grow, however, we envision that our impact will come mostly not from our own work, but from a broader constellation of scholars, analysts, and practitioners who help refine our ideas and put them into practice. We therefore expect to continue a strong emphasis on fieldbuilding and convening work that amplifies our impact.

2025 Activities

Designing AI agents to obey human laws

Increasingly agentic AI systems are the next great frontier in general-purpose AI. While our US Law and Policy team examines when and how humans should be made liable for the behavior of their AI agents, our Legal Frontiers team investigates a promising complementary **approach**: directly requiring AI systems to obey some set of laws. We have begun to **catalyze a field** of technical, legal, and policy research to develop the idea of law-following AI (“LFAI”) more fully, including through our own subsequent research and our Workshop on Law-Following AI in early August. We look forward to hosting the second annual Workshop on Law-Following AI in June 2026.



Designing AI agents to follow international treaties

In December, we followed up our pioneering introduction to LFAI with an in-depth **research article** on treaty-following AI (“TFAI”). We examined whether TFAI agents are technically feasible, and considered how such agents could function as a novel commitment mechanism for states, and as a self-execution mechanism for a new generation of treaties. Our paper also explored the legal, technical, and political hurdles to overcome, the ways treaty-following AI agents would function under existing legal frameworks, practical implementation pathways, and potential applications.

Anticipating potential concentration of presidential power

Advances in AI could eliminate traditional human constraints on executive power, and enable an unprecedented concentration of presidential authority. We term this hypothesized AI-enabled executive the “Unitary Artificial Executive.” In November, our Visiting Senior Fellow Prof Alan Rozenshtein delivered **remarks** on this topic at the University of Toledo Law School’s Stranahan National Issues Forum, which were subsequently published in *Lawfare*. Prof Rozenshtein argued that the Unitary Artificial Executive could become both constitutionally and technologically feasible, and encouraged scholars and courts to become sensitive to this trajectory.

Unbundling AI openness

The debate over AI openness – whether to make components of an artificial intelligence system available for public inspection and modification – forces policymakers to balance innovation, democratized access, safety and national security. In August, our affiliated researchers **introduced** a novel taxonomy of “differential openness,” unbundling AI into its constituent components and illustrating how each one has its own spectrum of openness. They operationalized these insights to provide policymakers with a playbook for how law can be precisely calibrated to achieve optimal configurations of component openness. The research article will be published in the 2026 *Wisconsin Law Review*.

Incubating new law and policy paradigms

As AI technology advances, we will likely need new legal and policy paradigms to help address the unique challenges that accompany technological progress. We aim to incubate ideas that are simultaneously anticipatory, actionable, accommodating, and ambitious. In 2025, we managed two workstreams – AI Agents and the Rule of Law, and International Regulatory Institutions – but our goal is to grow new paradigms until they can be sustained by a larger research community, both internally and externally. To that end, we have identified further **future frontiers** that we might work on in the future, including:

- Regulating Government-Developed Frontier AI
- Accelerating Technologies that Defend against Risks from AI
- Regulating Internal Deployment
- Fostering Legal Resilience by Rapidly Patching Legal Loopholes
- Responsibly Advancing AI-Enabled Governance
- Responsibly Automating Legal Processes
- Accelerating Legal Technologies that Empower Citizens
- Approval Regulation in a Decentralized World

We don't want people to wait for us to start working on these questions: they are already ripe for scholarly attention.

Featured publications

Law-Following AI: Designing AI Agents to Obey Human Laws →

Cullen O'Keefe, Ketan Ramakrishnan, Janna Tay, Christoph Winter

Treaty-Following AI →

Matthijs Maas, Tobi Olasunkanmi

Architectures of Global AI Governance →

Matthijs Maas

The Unitary Artificial Executive →

Alan Rozenshtein

Instrument Choice in AI Governance: Liability as the Indispensable Core →

Gabriel Weil

Unbundling AI Openness →

Parth Nobel, Alan Rozenshtein, Chinmayi Sharma

Balancing Safety and Privacy: Regulatory Models for AI Misuse →

JP de Mello Barreto

New AI Transparency Rules Have a Trade Secrets Problem →

Julius Hattingh

Fellowships

LawAI has made remarkable strides in establishing and advancing the field of AI, law and policy. Nevertheless, there is only so much we can achieve on our own, and an ecosystem reliant on a single source of expertise and advice is not robust to the quickening influx of new challenges and questions posed by regulatory and legislative efforts around the globe.

Part of our responsibility and our contribution to the broader ecosystem is to nurture the next generation of legal scholars and experts in AI. To that end, twice a year we offer Seasonal Research Fellowships to talented and ambitious law students, scholars, and professionals who want to work on projects at the leading edge of AI, law, and policy. This includes our flagship Summer Research Fellowship, as well as our inaugural Winter Research Fellowship.

Summer Research Fellowship

Our Summer Research Fellowship is a paid, highly-selective program, and one of the most competitive entry points into careers in AI, law, and policy. Designed for graduate law students, professionals, and academics, the fellowship offers participants the opportunity to work at the leading edge of legal research.



Each of the program's three tracks (US, EU Law, and Legal Frontiers) combines tailored research mentorship, career planning, and opportunities for direct engagement with policymakers, government officials, and private-sector leaders. Fellows are also invited to participate in an in-person week hosted in either Washington, DC, or Cambridge, UK, featuring expert-led sessions. Speakers in 2025 included Helen Toner, Interim Director at the Center for Security and Emerging Technology (CSET) and former OpenAI board member, and Max Katz, Policy Advisor to the US Senate.

Demand for this program continues to grow sharply. In 2025, we received 4,971 applications, a 6.8x increase over the previous year. Building on this momentum, we recently launched our inaugural [Winter Research Fellowship](#).

Fellowship alumni have gone on to pursue law and policy roles at the European AI Office, UK AI Safety Institute, leading AI labs, academia, and think tanks.

Testimonials from our Summer Research Fellowship

“It felt like being thrown into the deep end,
in the best way possible.”

— Summer Research Fellow (US), 2025

“During this fellowship, I gained a much
better understanding of the legal issues
raised in the context of frontier AI, particularly
in the AI Act. Thanks to the thorough review
process in the commentary, I have also
improved my reasoning and writing skills.”

— Summer Research Fellow (EU Law), 2025

“My time at LawAI has significantly shifted
my perspective about what it means to do
high-impact research. I'm quite confident
this has changed my career trajectory.”

— Julius Hattingh, Summer Research Fellow (Legal Frontiers), 2025

“To be clear about how good I thought this
in-person week was: it was one of the best
weeks of my life.”

— Summer Research Fellow (US), 2025, on the in-person week

Featured Alumni



Alexandra Jumper
Research Scholar, LawAI

Education: Harvard Law School (JD, magna cum laude); Harvard College (BA, magna cum laude)

Fellowship track: US

Before joining LawAI, Alex served as an Attorney Adviser in the Office of the Legal Adviser at the US Department of State, practiced at Boies Schiller Flexner LLP, and clerked for Judge Cornelia T.L. Pillard on the DC Circuit and Judge Jia M. Cobb on the US District Court for DC. At Harvard Law School, she served on the *Harvard Law Review*, was a teaching assistant for courses on constitutional law and civil procedure, and worked with the International Human Rights Clinic. Before law school, she was a Fulbright English Teaching Assistant in Brazil.



Madalina Nicolai
Research Scholar, LawAI

Education: University of Cambridge (LLM); University of Groningen (LLB)

Fellowship track: EU Law

Madalina is a researcher focused on the regulation of general-purpose AI models and systemic risk under the EU AI Act. Before joining LawAI, she was a Winter Fellow at an AI think tank, where she worked on best practices for reporting model evaluations in system cards. Madalina previously worked in Brussels, first at a public affairs consultancy focusing on EU digital and competition policy, and later in the public policy team of a semiconductor manufacturing company. She also spent two years in brand protection and intellectual property enforcement.



Henry Thompson
Assistant Professor of Economics
at the University of Mississippi

Education: George Mason University (PhD, MA); Clemson University (BA)

Fellowship track: Legal Frontiers

Henry studies the political economy and law and economics of AI and existential risk. While at LawAI, Henry worked on the economics of rogue AI agents. His project argued that misaligned AI agents need not imply misaligned behavior. With a simple model, he shows that ordinary constraints such as budgets and government sanctions can induce misaligned agents to act normally despite their alien objectives. His results suggest that technical and institutional alignment are substitutes.

Events

We host a number of annual workshops and forums on AI, law, and policy that cannot be found elsewhere. Our events aim to equip current and future legal scholars with the knowledge, tools, and networks needed to ensure that advances in artificial intelligence are beneficial and safe for everyone.

In doing so, we are making an effort to create a more robust ecosystem of scholars and practitioners who can pursue impactful careers both inside and outside of LawAI, as well as forge new bonds and strengthen existing networks between the key players in this space.

In 2025, we held three flagship events: the Summer Institute on Law and AI, the Workshop on Law-Following AI, and the Cambridge Forum on Law and AI.

Summer Institute on Law and AI

The Summer Institute on Law and AI hosts law students, professionals, and academics eager to explore pressing issues at the intersection of AI, law, and policy. This five-day gathering in Airlie, Virginia, featured presentations, fireside chats and Q&As with researchers, practitioners, and policymakers. Speakers included a former member of the National Security Council who worked on export controls, leaders of AI policy think tanks, and world-leading researchers on AI law and governance, including Turing Award winner Yoshua Bengio. Through expert talks and small discussion groups, participants examined fundamental challenges in applying legal analysis to the rapidly evolving AI landscape, as well as applied questions around AI liability, insurance law, export controls, and more.



“Thank you for planning and executing an exciting weekend in Airlie! You really curated a great group of folks, and I enjoyed many robust conversations on AI (and non-AI topics too!)”

– Michael Goodyear, Associate
Professor at New York Law School

“I really enjoyed the conference and felt like I got a lot out of it, both through forming valuable connections and through getting a better sense of the things people are working on and excited about.”

– Participant of the Summer
Institute on Law and AI, 2025

Workshop on Law-Following AI

The Workshop on Law-Following AI aims to catalyze research into the design of agentic AI systems that follow human laws. This two-day workshop at Christ's College, Cambridge, featured presentations and discussions from scholars in the fields of law and artificial intelligence, with a particular focus on crucial questions in legal scholarship that must be resolved in order to design AI agents that follow human laws.



“This was such a fantastic workshop. The best part was the curation of ideas, topics, and speakers. It was all top-notch content: I found myself strongly agreeing with speakers, strongly disagreeing with speakers, and also coming away with more formed, organized thoughts about how to do both.”

— Participant of the Workshop on Law-Following AI, 2025

“The workshop was the summer’s professional highlight. LFAI has been such a generative topic. I’m walking away with paper ideas, collaborations, and future research agendas.”

— Mihailis Diamantis, Ben V. Willie
Professor of Excellence at University of Iowa College of Law

“It was a fantastic event: the group of people, the content of presentations, the atmosphere and surroundings. The mix of long and short panels — as well as the hands-on activities — was refreshing!”

— Participant of the Workshop on Law-Following AI, 2025

Cambridge Forum on Law and AI

The Cambridge Forum on Law and AI aims to equip participants with the knowledge, tools, and networks they need to address challenges in the EU's AI laws. This four-day gathering of law students, professionals, and academics featured presentations, fireside chats, and Q&As, complemented by opportunities for smaller group meetings and one-on-one conversations between participants. Participants had the opportunity to give short lightning talks to present their working papers, projects, and ideas to spur discussion and get feedback.



“The agenda and organization of the different sessions allowed for deeper exchanges and many networking opportunities.”

– Participant of the Cambridge Forum on Law and AI, 2025

“This is how meaningful policy events should look.”

– Participant of the Cambridge Forum on Law and AI, 2025

“It was extremely well organised. The programme was focused on 2-3 topical issues in AI governance, bringing together leading experts in this area. The relatively short presentations worked well to keep the event dynamic. I also appreciated the diversity of backgrounds and affiliations of the speakers and participants.”

– Participant of the Cambridge Forum on Law and AI, 2025

Finances

LawAI relies on the philanthropic support of a diverse group of organizations and individuals who invest in our research, programs, and operations each year.

Over the past year, we benefited from substantial support from foundations and individual donors, alongside many smaller contributions from individual supporters. We remain grateful for the generous support and trust of our donors who enable us to advance our mission to promote security, welfare, and the rule of law in the age of AI.

In 2025, our expenses totalled \$5.1M. Approximately 71% of this covered staff costs, including compensation, taxes, and employee benefits for our researchers, as well as our communications and operations teams. The remaining 29% supported operational and program-related expenses, including events, facilities, contracted services, compensation for seasonal research fellows, and other organizational costs.

LawAI welcomes philanthropic contributions to ensure that we continue to advance our mission while remaining independent of commercial or political influences. You can **donate to LawAI** via every.org. To learn more about our work, funding needs, and specific opportunities, please contact our Director, Christoph Winter.

Team

Our full-time legal experts are based in locations where they have access to policymakers. Our two primary locations are Washington, DC, giving us close proximity to policymakers on the Hill, and Cambridge, UK, where we collaborate with top academics while remaining close to government decisionmakers in the EU and UK.

Board of Directors

Eric Martínez
Fellow and Instructor in Law,
University of Chicago Law School

Prof Robert Trager
Director of the Oxford Martin AIGI,
University of Oxford

Jonas Freund
Senior Research Fellow,
GovAI

Prof Christoph Winter
Assistant Professor of Law and AI,
University of Cambridge

Leadership

Prof Christoph Winter
Director

Cullen O’Keefe
Director of Research

Mackenzie Arnold
Director of US Policy

Research & Policy

Cullen O’Keefe
Director of Research

Mackenzie Arnold
Director of US Policy

Emily Gillett
Interim Director of EU Law

Koji Flynn-Do
Chief of Staff, US Policy

Dr Matthijs Maas
Senior Research Fellow

Charlie Bullock
Senior Research Fellow

Janna Tay
Research Manager

Prof Gabriel Weil
Non-Resident Senior Fellow

Prof Alan Rozenshtein
Visiting Senior Fellow

Prof Peter Salib
Visiting Senior Fellow

Maarten Herbosch
Senior Research Scholar

Bahrad Sokhansanj
Senior Research Scholar

Mark Reddish
Senior Research Scholar

Benjamin Brady
Senior Research Scholar

Thomas Lilamand
Senior Research Scholar

Alex Jumper
Research Scholar

Madalina Nicolai
Research Scholar

Zlatko Grigorov
Research Scholar

Hannes Bastians
Research Scholar

Dev Basumallik
Research Scholar

Stephan Llerena
Research Scholar

Communications

Cody Fenwick
Senior Communications Manager

Andrew Leeke
Editor

Programs & Operations

Kristina Fort
Research Manager

Malou Estier
Program Associate

Jackson Gardner
Senior Operations Manager (US)

Simon Jarvis
Senior Operations Manager (EU)

Margisel Adams
Senior Operations Manager

Sophia Rubio
Operations Manager, Programs

Executive

Prof Christoph Winter
Director

Chloe Clayton-Green
Executive Associate

Research Affiliates

Prof Jeff Sebo
Associate Professor of Environmental
Studies, New York University

Prof John Bliss
Assistant Professor, University
of Denver Sturm College of Law

Prof Yonathan A. Arbel
Professor of Law, University
of Alabama School of Law

Dr José Jaime Villalobos
Multilateral Governance Lead,
Future of Life Institute

Cecil Abungu
PhD student,
University of Cambridge

Prof Kevin Frazier
AI Innovation & Law Fellow,
University of Texas School of Law

Jonas Freund
Senior Research Fellow,
GovAI

Prof Jack Thorlin
Assistant Professor, University
of Arkansas School of Law

Prof Chinmayi Sharma
Associate Professor, Fordham
University School of Law

Prof Jack Boeglin
Sharswood Fellow, University of
Pennsylvania Carey Law School

Prof Noam Kolt
Assistant Professor, Hebrew University
Faculty of Law and School of Computer
Science & Engineering

Dr Anat Lior
Assistant Professor, Drexel University's
Thomas R. Kline School of Law

“What drew me to LawAI was the mission, but what keeps me here are the people. I work alongside individuals who are not only exceptionally gifted and kind, but also driven by a genuine desire to get things right and to make a difference, along with a degree of nuance and epistemic modesty that I've come to cherish. The support and freedom I have received at LawAI have allowed me to contribute across our research efforts and field-building work in ways that feel truly meaningful, and have allowed me to produce research in an environment that for me combines the best virtues of the academic and policy worlds. I've also been impressed by the ability of the organization to scale ambitiously while remaining focused, mission-aligned and caring.”

— Matthijs Maas, Senior Research Fellow

“LawAI has been an excellent place for me to move from being a private practice lawyer into working in research and making an impact in an area I care deeply about. It has been a privilege to learn from some of the best and brightest at the forefront of AI and legal research — they have guided and pushed me in ways I wouldn't have received anywhere else. The Legal Frontiers team has a unique vision and orientation that I am really excited to help realise.”

— Janna Tay, Research Manager

“Working at LawAI has been a genuine privilege. The team brings together people who are exceptionally bright and mission-driven, while also being thoughtful, kind, and generous with their time. It's a rare environment where ambition and integrity are aligned, and where I've been able to take meaningful ownership of our core programs and events as we scale.”

— Sophia Rubio, Operations Manager, Programs

Looking Ahead

As we step further into the AI era, we're motivated to deepen our research on law and AI, and expand our collaborations across academia, government, and industry.

On the research front, our US Law and Policy team will continue to develop and promote adaptive governance: research and consulting work that deliberately preserves flexibility across technical, institutional, and policy choices. We published an essay introducing adaptive governance (termed “radical optionality”) in early 2026. We are also hiring a Research Manager to support our growing number of full-time researchers in the US.

Meanwhile, our EU Law team will continue to focus on general-purpose AI models and related enforcement issues. In particular, we are focusing on systemic risks posed by advanced AI models as well as the responsibilities and procedures for assessing and mitigating those risks as mandated by the EU AI Act. In collaboration with the University of Cambridge, we will continue to publish our in-depth commentary on the EU AI Act over the course of 2026.

Finally, our Legal Frontiers team will continue exporting the results of its new Automated Governance workstream, beginning with its paper on [automated compliance](#). We intend to hire a Director of Automated Governance in 2026, to investigate how governments and governed parties can leverage frontier AI technologies to automate core governance tasks—safeguarding security, welfare, and the rule of law in the face of rapid AI progress.

Meanwhile, in early 2026 our Programs team commenced its first Winter Research Fellowship program, offering law students, professionals, and academics the opportunity to work at the leading edge of AI, law, and policy. Our annual flagship Summer Research Fellowship program will run for 10 weeks in the summer of 2026. We remain excited about the calibre of our applicant pool as well as the trajectories of our alumni, who have gone on to pursue law and policy roles at the European AI Office, UK AI Safety Institute, leading AI labs, academia, and think tanks. We intend to hire a Director of Programs in 2026 to lead our growing portfolio of fellowships and events.

We will also host three flagship events in 2026: our [Workshop on Law-Following AI](#), the [Summer Institute on Law and AI](#), and the [Cambridge Forum on Law and AI](#). We will also launch our inaugural [Summer School on Law-Following AI](#), aimed at gathering early-stage academics to be introduced to, and learn about, law-following AI and its emerging research agenda.

On the operations side, LawAI is expanding its international presence in three ways. First, we have acquired additional office space at our workspace in Cambridge, UK, to support more staff, research affiliates, and visiting scholars. Second, we are establishing entities in the UK and EU. Third, we are hiring a Director of Talent to source and evaluate promising candidates from across the world.

We remain grateful for the generous support and trust of our donors and staff who enable these developments, allowing us to promote security, welfare, and the rule of law at scale.



INSTITUTE FOR LAW & AI

6 Liberty Square, #6079
Boston, MA, 02109, United States

20 Station Road,
Cambridge, CB1 2JD,
United Kingdom

law-ai.org

hello@law-ai.org