

**Statement of
Paul Ohm
Professor, Georgetown University Law Center**

**Before the
Subcommittee on Disaster Management,
District of Columbia, and Census
of the
Committee on Homeland Security and Governmental Affairs
U.S. Senate**

**Hearing on
Rogue AI: Securing the Homeland Against AI Agent Attacks**

September 30, 2026

Chairman Hawley, Ranking Member Kim, and Members of the Subcommittee, I appreciate the opportunity to be here with you today to discuss the role the law, liability, and our governance institutions should play in helping secure our critical infrastructure from attacks from AI agents.

I am a Professor of Law at the Georgetown University Law Center. My scholarship and teaching focus on artificial intelligence law, information privacy, computer crime, and technology and the law. For the past four years I have taught a class on Artificial Intelligence Law at Georgetown, using a soon-to-be-published law school casebook entitled ARTIFICIAL INTELLIGENCE LAW, which I have co-authored with Professors Margot Kaminski and Andrew Selbst. I serve as a faculty director for the [Institute for Technology Law and Policy](#), a faculty advisor of [The Privacy Center at Georgetown Law](#), and a member of the steering committee of the [Georgetown University Tech & Society Initiative](#).

My research and thoughts on the necessary and proper role of the legal system in helping secure our country's critical infrastructure from attacks by AI agents are informed by three positions I have held in federal and state government. I started my legal career twenty-five years ago as a federal prosecutor in the U.S. Department of Justice's Computer Crime and Intellectual Property Section, where I worked on cutting edge issues around the Computer Fraud and Abuse Act, Fourth Amendment, and Electronic Communications Privacy Act. Later, from 2012 to 2013, I served as a Senior Policy Advisor to the Federal Trade Commission, helping the FTC protect the privacy of consumers from online invasions. Most recently, I served as a Special Assistant Attorney General working for the Colorado Attorney General on the implementation of the Colorado Privacy Act, which not only protects privacy but also aims to assure the accountability and transparency of automated decisionmaking. I am speaking with you today solely in my personal and academic capacities.

The 2026 OpenAI hack of Hugging Face provides a concrete example of what many have long anticipated: the development and deployment of AI agents are exposing gaps in the legal frameworks we have developed to protect people and infrastructure from cyberattacks. Consider a cut-and-paste thought experiment. If you replace the words “AI agent” with “OpenAI employee” throughout the various technical reports that have been released,¹ there is little doubt that OpenAI and their employees would be liable to victims and guilty of committing federal crimes.² It is not so clear that these legal conclusions hold when machines are doing the hacking rather than humans, thus revealing worrisome gaps in our legal framework. We urgently need to change the law to close these gaps, but as we do so, it is important to intervene in a thoughtful and intentional way. My statement starts by listing five goals our reworked legal framework should achieve, surveys the legal tools we can already rely upon today without any reforms, proposes a few new laws at the federal and state levels, and explains why and to what extent the criminal law should play a role in our response. My overall message is that the mostly self-regulatory status quo is not nearly enough to protect us from the changing cyberthreat landscape and that there are better approaches we can and should adopt.

1 THE GOALS OF LEGAL RESPONSES TO AI’S CYBERSECURITY THREATS

The reports that OpenAI’s AI agents repeatedly breached the security of Hugging Face and other private and public entities should issue a clarion call for responses along multiple fronts. The time has long passed for relying on industry self-regulation, as the many company executives requesting government action appear to recognize. Let me begin by offering five goals for an effective legal response, positive changes our legal interventions should help bring about.

Change the incentives. Well-designed laws will give industry executives, line engineers, and other staff the incentive to invest more in AI safety and control. An economist might call this forcing companies to internalize the externalities; a products liability expert might frame this as a way to compel additional investment in product safety; a prosecutor might speak of creating specific and general deterrence. Whatever you call it, the idea is straightforward: well-designed legal frameworks can change the behavior of actors in this industry who otherwise seem

¹ OPENAI, OPENAI – HUGGING FACE INCIDENT: TECHNICAL REPORT, <https://openai.com/index/hugging-face-incident-and-the-road-ahead/> (Aug. 26, 2026) (available as link at URL); RYAN GREENBLATT, AJEYA COTRA & HJALMAR WIJK, METR, BRIEF INDEPENDENT INVESTIGATION OF AGENTS’ BEHAVIOR, REASONING AND COLLABORATION IN THE OPENAI / HUGGING FACE HACKING INCIDENT (Aug. 26, 2026), <https://metr.org/blog/2026-08-26-openai-hugging-face-incident-investigation/>; ALEX FORMAN ET AL., PARSE, REVEALING THE DETAILS OF HOW OPENAI AGENTS HACKED HUGGING FACE (Sept. 25, 2026), <https://swarmtraces.org/>.

² *E.g.*, 18 U.S.C. § 1030(a)(5)(A) (Computer Fraud and Abuse Act (CFAA) crime punishing “whoever . . . knowingly causes the transmission of a program, information, code, or command, and as a result of such conduct, intentionally causes damage without authorization, to a protected computer . . . shall be punished”); § 1030(g) (providing a cause of action for “any person who suffers damage and loss by reason of a violation” of the CFAA).

locked in a socially perilous race of careless competition. The right laws and enforcement strategies can spur everyone involved to take meaningful and sensible new steps to protect the public from avoidable and serious harm.

Increase transparency. OpenAI deserves some praise for speaking quickly, voluntarily, and candidly about what happened. While there are things I can fault about their public statements – some of the early announcements read more like marketing, and their laudably thorough technical report suggests a touch of pride in what they accomplished – we should hope that OpenAI and their competitors continue to speak openly and candidly about their failures. On the other hand, the kind of increased civil litigation I am calling for below might, if it comes to pass, dampen such candor going forward. We can adopt new legal mandates to ensure prompt, thorough, and actionable insights after incidents and near misses. These should require both company self-reporting and the involvement of truly independent auditors.

Establish baseline standards. To avoid a race to the bottom, the AI industry must have meaningful baseline standards about everything from safety to control to governance. Well-established and transparent standards will ensure industry actors adhere to at least minimum safety and control practices. They will also help consumers and citizens understand, scrutinize, and ultimately trust the safety of the products and services they buy. The law can both spur and shape the development of such standards: for example, tort liability regimes such as negligence and products liability help enumerate and refine standards of safety and care, one fact pattern at a time.

Broaden participation. The AI safety and control debates happen primarily inside small rooms in a few corporate headquarters among a limited number of people who draw on a very narrow set of expertise and experience. Legal frameworks can bring the outside in, inviting the participation of expert government workers, judges, juries, victims, senators and other federal and state legislators, policymakers, and ordinary people. They allow us to draw on different institutions with diverse toolkits and historical track records, from federal agencies to state attorneys general to civil society groups to unions to academics.

Provide relief to victims. A saving grace of the Hugging Face hack is, as far as we know, no victim suffered a significant loss of personal information or significant damage to critical infrastructure. Next time, we may not be so lucky: uncontrolled AI agents may damage critical infrastructure, steal sensitive and personal information, and even cause injuries or deaths. Filling gaps in our legal framework will give victims of future incidents the recourse they need to recover or be made whole.

2 TODAY’S LEGAL FRAMEWORK

Our present political climate is not well-suited for coming together to solve big challenges, yet in this case, the calls for urgent action are loud, justified, and increasing. We cannot and should not wait years for a fully-developed AI governance law to spring like Athena from the head of the Congress. Instead, we can and should

work piece by piece, setting up small foundation stones for later legal construction, including by enacting targeted pieces of federal legislation. Even before we do that, the current law already provides a few legal avenues to advance the five goals listed in Part 1. These create an admittedly incomplete framework that still forms a critically important core that Congress must actively preserve and extend.

The first piece is state tort law. Victims who suffer cyberattacks from uncontrolled AI agents can file lawsuits against those responsible for the development, deployment, monitoring, and control of those agents, alleging negligence and products liability, among other potential causes of action. Negligence cases will require judges and juries to decide whether defendant-companies have acted with reasonable care, and through the genius of the common law system, these decisions will gradually develop into a standard of care. Congress should reinforce and reaffirm the important roles of state tort law in setting the standard of care, strengthening incentives to control and monitor harmful AI agents, providing independent fact-finding with judicial oversight, and ensuring that if such harms come to pass, people will have recourse to pursue remedies.

For example, OpenAI's attacks on Hugging Face are alarming, and many would like to see someone held responsible, to send a message to the AI companies to improve their practices. It is possible that our centuries' old state tort law framework already holds the key. Having reviewed the multiple, detailed reports dissecting the causes and progression of the attacks, I cannot yet say whether OpenAI would be found liable for negligence if they were sued. The litigation would unearth additional facts about what truly happened to add to what the self-reporting investigations have already revealed. But based on what has been reported, multiple cybersecurity experts have publicly called the company's actions "negligent,"³ or otherwise suggested that OpenAI's security practices fell below industry standards,⁴ but those have been technical opinions, not sworn expert testimony, much less the basis for judicial rulings.

Congress should also reaffirm and take care to preserve state and federal unfair and deceptive acts and practices laws, such as section 5 of the Federal Trade Commission Act.⁵ In federal law, an unfair act is one "likely to cause substantial injury to consumers which is not reasonably avoidable by consumers themselves and

³ IANS, OpenAI's Postmortem Recasts the Hugging Face Breach as an Incident Response Failure, IANS News, (August 29, 2026) <https://www.iansresearch.com/resources/all-blogs/post/security-blog/2026/08/28/openai%27s-postmortem-recasts-the-hugging-face-breach-as-an-incident-response-failure> ("That OpenAI missed this, especially after identifying that Artifactory was a primary agentic target, is nothing short of negligence."); Derek B. Johnson, *The AI Hacking Apocalypse is not Inevitable*, CYBERSCOOP (Sept. 17, 2026), <https://cyberscoop.com/ai-agent-hacking-apocalypse-cybersecurity/> (quoting Juan Andres Guerrero-Saade criticizing the use of sandboxes in recent incidents as "neglect, negligence, people just mishandling things" without naming OpenAI specifically).

⁴ Marius Horatau, *The Hugging Face Incident Is Not an AI Story*, UPHACK (Aug. 31, 2026), <https://uphack.io/blog/post/the-hugging-face-incident-is-not-an-ai-story/> ("For any security engineer though, this apocalyptic story reads much more boring, as a major security operational failure rather than an AI capability result.").

⁵ 15 U.S.C. § 45.

not outweighed by countervailing benefits to consumers or to competition.”⁶ The FTC and state attorneys general may be able to prove that OpenAI’s attacks meet these requirements, and I imagine that some of them have begun to ask questions. Many state versions of this law allow injured people to sue, giving private actors another vehicle for seeking relief.

Finally, Congress should ensure that no matter how it approaches AI regulation in the years to come, it preserves other statutory private rights of action that supplement the tort system. If developers, deployers, and users launch AI agents that violate laws such as the Wiretap Act,⁷ Copyright Act,⁸ or Illinois Biometric Information Privacy Act,⁹ affected individuals can sue, and many of these laws allow for statutorily defined damages that set a higher floor on the potential monetary award.

3 NEW LAWS ARE NEEDED: NEW PRIVATE RIGHTS OF ACTION, STRICT LIABILITY, AND REGULATION

Although I urge Congress to preserve the critically important existing legal framework described above, additional measures are needed. Congress and state legislatures can bolster the current legal framework by enacting new statutes providing private rights of action tailored for the kinds of cyberattacks that AI agents may commit and establishing new systems of regulatory AI governance.

To start, Congress can enact a new private right of action for victims of cyberattacks initiated by AI agents. One possibility is to amend the Computer Fraud and Abuse Act (CFAA) to reach developers and deployers for the “damage” and “loss” (two already defined terms) caused by their agents.¹⁰ A place to start could be to amend 18 U.S.C. § 1030(a)(5)(A) which does not require any “access” to a computer but instead makes it illegal to “knowingly cause[] the transmission of a program, information, code, or command, and as a result of such conduct, intentionally cause[] damage without authorization to a protected computer.”¹¹ Simply striking the word “intentionally” from this provision might cover what happened in the OpenAI case, giving victims of the next similar incident a vehicle to sue for relief under section 1030(g).¹² I will say more about the CFAA in Part 4.

Congress might also consider creating a new law imposing *strict liability* for developers and deployers of AI agents that cause certain serious damage, such as physical injury, death, or damage to critical infrastructure. Under a regime of strict liability, victims of attacks would need to show that OpenAI’s agents caused their injuries, but they would not need to establish that OpenAI acted below a reasonable

⁶ § 45(n).

⁷ 18 U.S.C. § 2520(a).

⁸ 17 U.S.C. § 501(b).

⁹ 740 ILCS 14/20(a).

¹⁰ 18 U.S.C. § 1030(e)(8) (defining “damage”); § 1030(e)(11) (defining “loss”).

¹¹ § 1030(a)(5)(A).

¹² § 1030(g).

standard of care. Enacting this kind of law would echo instances in the past when strict liability has been applied to complex new technological innovations with immature industrial controls and the potential to cause significant damages, such as large reservoirs of water¹³ and early commercial aviation.¹⁴ In the case of aviation, strict liability gave the industry and legal system the time it needed to develop better controls and clearer best practices and standards, and many jurisdictions reverted to negligence standards once the risks of reservoir maintenance and flight were better understood and more controlled.¹⁵ We might walk a similar path with AI agents.

Congress should also consider new laws to reduce the likelihood that attacks by AI agents will invade the personal privacy of individuals, for example through the theft or disclosure of sensitive and intimate information about the users and employees of attacked services. If Congress enacts a meaningful comprehensive privacy law, such as it has considered in recent years,¹⁶ it would increase the pressure on companies to safeguard or minimize the collection of personal information, and it would give recourse to those who have their personal information stolen or misused.

The civil litigation system cannot alone bear the weight of safety and risk regulation in this emerging, complex space. We need to develop new forms of administrative oversight, and once again we can first turn to the states. Congress should avoid thwarting the regulatory innovation, agility, and proximity to the people of state legislatures – our nation’s “laboratories of democracy” – in this rapidly changing context.¹⁷ States are already testing new forms of regulation in AI and AI-adjacent areas, from data breach notification to the regulation of autonomous vehicles to automated decisionmaking. Of most direct relevance, California’s SB 53¹⁸ and New York’s RAISE Act¹⁹ both regulate so-called frontier model developers, and both provide incident reporting obligations. To give the laboratories of democracy the time

¹³ See *Rylands v Fletcher*, (1865-66) L.R. 1 Ex. 265; [1865-66] All E.R. Rep. 1; (1866) 4 Hurl. & C. 263 (Ex Chamber), affirmed by *Rylands v Fletcher*, (1868) L.R. 3 H.L. 330 (HL) (holding a landowner strictly liable for the escape of water stored in a reservoir); *Cahill v. Eastman*, 18 Minn. 324, 341–42 (1871) (adopting the *Rylands* rule where water from an underground tunnel escaped and damaged the plaintiff’s mill).

¹⁴ Uniform Aeronautics Act § 5 (Nat’l Conf. of Comm’rs on Unif. State Laws 1922) (imposing strict liability on aircraft owners for ground damage); Restatement (Second) of Torts § 520A (Am. L. Inst. 1977) (classifying strict liability for aircraft ground damage as an application of the “abnormally dangerous activity” doctrine).

¹⁵ See *Crosby v. Cox Aircraft Co.*, 746 P.2d 1198 (Wash. 1987) (summarizing the rise and fall of strict liability for aircraft ground damage and adopting negligence).

¹⁶ American Privacy Rights Act of 2024, H.R. 8818, 118th Cong. (2024); American Data Privacy and Protection Act, H.R. 8152, 117th Cong. (2022).

¹⁷ *New State Ice Co. v. Liebmann*, 285 U.S. 262, 311 (1932) (Brandeis, J., dissenting) (“It is one of the happy incidents of the federal system that a single courageous State may, if its citizens choose, serve as a laboratory; and try novel social and economic experiments without risk to the rest of the country.”).

¹⁸ Transparency in Frontier Artificial Intelligence Act, 2025 Cal. Stat. ch. 138 (S.B. 53) (codified at Cal. Bus. & Prof. Code § 22757.10 et seq.).

¹⁹ Responsible AI Safety and Education Act, ch. 699, 2025 N.Y. Laws, amended by S. 8828, 2026 Leg. (N.Y. 2026) (codified at N.Y. Gen. Bus. Law art. 44-B) (effective Jan. 1, 2027).

and freedom they need to do this critical work, Congress must not preempt AI regulation by the states.

Eventually, informed by the transparency-forcing and standards-developing features of a multi-layered legal approach, Congress will know enough to develop a robust overarching regulatory framework. Such a framework should have both broad powers and well-resourced, muscular enforcement capabilities.²⁰ There are useful precursors we can borrow, federal structures we created at earlier times in our history for regulating complex areas involving public health, safety, and welfare. Many of these offer ideas we might draw on for AI governance: the Food and Drug Administration’s premarket approval system for medical devices;²¹ the Environmental Protection Agency’s “cooperative federalism” framework of national standards implemented by the states under the Clean Air Act;²² the Nuclear Regulatory Commission’s resident inspector program, placing federal inspectors permanently on-site at every licensed nuclear power plant;²³ the National Transportation Safety Board’s authority to investigate transportation accidents and recommend safety improvements;²⁴ and the Cybersecurity and Infrastructure Security Agency’s forthcoming mandatory cyber incident reporting rules,²⁵ to name only a few.

4 THE ROLE FOR THE CRIMINAL LAW AND THE CFAA

A multi-layered civil and regulatory approach ought to form the core of our expanded legal framework for addressing, preventing, and mitigating cyberattacks from AI agents, but criminal law prohibitions must also play a role.

The criminal law can reach harms and risks that the tort system cannot, such as diffuse harms that affect many people and risks that have not yet ripened into harm, such as attempts, conspiracies, and endangerment. Criminal laws also serve an expressive function, signaling condemnation through prohibition, which is hard to ignore as a simple cost of doing business. They can be directed at responsible, culpable individuals, who can’t pass along their punishments to insurance or shareholders. Criminal laws also advance the interests of victims who lack the resources needed to file and win a complex lawsuit.

Criminal law also provides a useful set of transparency mechanisms such as search warrants, subpoenas, and testimonial immunity. It invites the direct

²⁰ My collaborative research proposing new regulatory approaches for governing the information economy can be found at the Redesigning the Governance Stack Project, Georgetown Institute for Technology, Law, and Policy, <https://www.law.georgetown.edu/tech-institute/programs-and-initiatives/redesigning-the-governance-stack-project/>.

²¹ 21 U.S.C. § 360e.

²² 42 U.S.C. §§ 7409-7410.

²³ 42 U.S.C. § 2201(o); 10 C.F.R. § 50.70.

²⁴ 49 U.S.C. §§ 1131.

²⁵ 6 U.S.C. §§ 681-681g; Cyber Incident Reporting for Critical Infrastructure Act of 2022 Reporting Requirements, 89 Fed. Reg. 23,644 (proposed Apr. 4, 2024).

involvement of subject matter experts in federal law enforcement, who are often on the frontline of coordinating our national cybersecurity response.

That said, the criminal law can be a blunt instrument. The threat of prosecution may overdeter, especially if laws are written vaguely or broadly. Criminal laws tend to turn on mental state – *mens rea* – requirements such as intent or knowledge, which can be difficult to apply when humans provide vague goals that are then converted into criminal action by AI agents.

What then, does the CFAA, the principal federal criminal law designed for facts like these, have to say? As currently written, and under the facts of the Hugging Face hack that have been publicly revealed, neither OpenAI nor any of its employees have likely committed a crime under the CFAA. The CFAA defines seven categories of crimes, several with multiple subcategories.²⁶ All of the relevant provisions require one to either “intentionally access” or “intentionally damage without authorization” a “protected computer.”²⁷ But to be clear, OpenAI admits in its technical report that it *intentionally* launched hundreds of agents to try to accomplish the tasks in the ExploitGym benchmark, “a cybersecurity evaluation that measures whether a model is capable of discovering and exploiting cybersecurity vulnerabilities.”²⁸ This is almost an admission of a violation of the CFAA crime defined in section 1030(a)(5)(A), which does not require “intentional access,”²⁹ but because OpenAI intentionally attacked only its own systems, it probably does not satisfy the element of “intentional damage *without authorization*.”³⁰

It is vital for Congress to understand that the CFAA is not the only criminal law with intent requirements that may soon be tested by AI agents. Many criminal laws separate culpable from innocent behavior based on the thoughts in the head of the human actor. Remove the human, his head, and his thoughts, and replace them with a mindless automation, and we may soon experience significant gaps in legal coverage, leading to the specter of uninvestigable, irremediable, unpunishable crime. Beyond computer hacking, AI agents are likely to take the steps that comprise the illegal conduct – in legal jargon, the *actus reus* – for market manipulation, price fixing, and wire fraud, through access to the Internet alone. AI agents with control over resources in the physical world may take the steps required for environmental crimes, export control violations, and controlled substance distribution. Agents will perform all of the acts and cause all of the consequences that these laws are designed to prevent, leaving behind the kinds of victims these laws were designed to protect, but with no criminal to charge. AI agents thus pose a significant “missing intent” problem throughout the criminal law.

²⁶ 18 U.S.C. § 1030(a)(1) – (a)(7).

²⁷ § 1030(a)(2)(C), (a)(3), (a)(4), (a)(5)(A), (a)(5)(B), (a)(5)(C).

²⁸ OPEN AI TECHNICAL REPORT, *supra* note 1, at 5.

²⁹ See Cal Newport, *Did OpenAI’s New Model “Go Rogue”?* (July 27, 2026), <https://calnewport.com/did-openais-new-model-go-rogue/> (“[C]ircumventing internet restrictions and hacking into servers are exactly the kinds of things these ExploitGym systems are designed to do. There was no ‘rogue’ agent or revelation of some surprising, devious new capability.”).

³⁰ § 1030(a)(5)(A) (emphasis added).

There is no simple, universal fix for this problem. It would be unwise to simply declare that all actions taken by an AI agent should be imputed as intended by the next deployer or developer up the chain. Instead of seeking a quick fix, we should look at the role that *mens rea* plays in each law and ask the following three questions: (1) Is the conduct described the kind that ought to be criminalized if committed by an AI agent? (2) If so, should we assign criminal culpability to a developer, deployer, user, or more than one of the above? (3) Can we set a *mens rea* requirement for that actor at a level other than intent?

On this last point, there are many criminal laws, including at the federal level, that punish certain kinds of conduct without intent. Some crimes punish negligence.³¹ Other crimes have been interpreted to punish reckless behavior.³² A few crimes affecting public welfare punish people without proof of any *mens rea*, in other words, as strict liability crimes.³³

But if Congress contemplates amending the CFAA, it should be mindful that this law has long been criticized. It has been the object of public protest,³⁴ academic criticism,³⁵ and judicial narrowing.³⁶ No doubt those who have opposed the CFAA in the past will object to proposals to use it to try to deal with threats from AI agents. These criticisms should be taken very seriously, but they may also be possible to accommodate. OpenAI did not simply “exceed authorized access” by violating Hugging Face’s terms of service, the part of the statute and the kind of fact pattern that has produced the most controversial applications. OpenAI’s AI agents engaged in good old fashioned computer hacking, the kind Congress had in mind when it first enacted the law in 1986. Four decades later, Congress might be able to extend the CFAA to cover developers and deployers who launch AI agents engaged in this kind of foreseeable conduct without reopening the kind of uses that cases like *Van Buren* have foreclosed.

³¹ E.g. § 1112 (involuntary manslaughter, interpreted by courts to require gross negligence or reckless disregard); § 1115 (seaman’s manslaughter, punishing mere negligence by a ship’s “captain, engineer, pilot, or other person employed on any steamboat or vessel” that leads to death of another); 33 U.S.C. § 1319(c)(1) (Clean Water Act, negligent violations punished as misdemeanors); 42 U.S.C. § 7413(c)(4) (Clean Air Act, negligent release of hazardous air pollutant that negligently places another in imminent danger of death or serious injury as misdemeanor).

³² See, e.g., 18 U.S.C. § 113(a)(6) (assault resulting in serious bodily injury); *United States v. Zunie*, 444 F.3d 1230, 1235 (10th Cir. 2006) (holding that recklessness is a sufficient *mens rea* for conviction under § 113(a)(6)).

³³ E.g., 21 U.S.C. 333(a)(1) (misdemeanor for certain violations of the Food, Drug, and Cosmetics Act).

³⁴ Elec. Frontier Found., Computer Fraud and Abuse Act Reform, <https://www.eff.org/issues/cfaa> (calling on people to email their representatives to urge CFAA reform).

³⁵ See, e.g., Laurent Sacharoff, *Criminal Trespass and Computer Crime*, 62 WM. & MARY L. REV. 571, 578 (2020) (criticizing the way the CFAA’s “breadth can lead to unjust applications”).

³⁶ See *Van Buren v. United States*, 593 U.S. 374, 390 (2021) (rejecting a broad reading of “exceeds authorized access” and adopting a “gates-up-or-down” approach to CFAA liability); *hiQ Labs, Inc. v. LinkedIn Corp.*, 31 F.4th 1180, 1197-98 (9th Cir. 2022) (holding, on remand in light of *Van Buren*, that the concept of “without authorization” does not apply to information available to the public on an otherwise open website).

5 CONCLUSION

If a primary goal of our tort and criminal law systems is to deter harmful behavior, we are failing to meet the mark when it comes to the threat of cyberattacks caused by AI agents. Fixing this will require more than the legal reforms I focus on above. We need to strengthen cybersecurity practices across all industry actors, large and small. We need to develop industry standards for AI safety and control, increase our use of red teams, after-action reports, and independent audits.

The legal system can and should also play a role. We should begin, piece-by-piece, to construct a framework of complementary and overlapping rules that hold companies to account, *ex post*, for the harmful acts of the AI agents they develop and deploy; we should set up state and federal regulatory structures that ensure *ex ante* that companies are meeting meaningful standards of safety and control; and we should provide a backstop of criminal punishment, tied to the most significant harms.

Most importantly, the distributed, multi-faceted, multi-jurisdictional approaches I have laid out invite all of us – as litigants, civil enforcers, agents, prosecutors, constituents, voters, legislators, policymakers, and civil society organizations – into active participation in the important deliberations we are having about what to do about technologies that are being invented, propagated, and marketed by very few of us. The people of this country are afraid and feel disempowered by this debate, and they are turning to their elected representatives to raise and address their concerns. For problems as complex and important as control over AI agents and protection of critical infrastructure, we need to invite many more voices and viewpoints to weigh in.