FEATURE

Governing Artificial Intelligence: Designing Professional Structures for the Predictive Age

Wendy H. Wong and David A. Lake

It seems like ages ago that OpenAI CEO Sam Altman embarked on a worldwide tour to ask governments to govern artificial intelligence, though it was only in 2024.[1] As the political winds have shifted, world leaders are largely adopting permissive stances toward AI (even those in Europe, since Mario Draghi's 2024 report that emphasized the need for European competitiveness in AI).[2] The latest AI Action Summit in Paris pushed for AI that takes "into account international frameworks for all"[3] but, ironically, *without* an emphasis on safety. The administration of U.S. president Donald Trump has rolled back Biden-era regulations and announced its own AI Action Plan that emphasizes innovation, building AI infrastructure, and leading in diplomacy and security, with scant attention to potential harms.[4] A recent executive order signed by the president overrides state efforts to regulate AI.[5] After a meeting of the UN General Assembly concerning AI held in September 2025, Michael Kratsios, director of the Office of Science and Technology Policy, stated that "we totally reject all efforts by international bodies to assert centralized control and global governance of AI."[6] With the limited exception of the EU Artificial Intelligence Act (EUAIA) and a bill

Wendy H. Wong, University of British Columbia, British Columbia, Canada (w.wong@ubc.ca)

David A. Lake, University of California San Diego, San Diego, California, United States (dlake@ucsd.edu)

Acknowledgments: We thank Terry Chapman, Clara H. Suong, Graeme Auld, and anonymous reviewers for helpful comments and Elisa Ehrenburg, Karyn Kao, Michael Monclou, Aleksandra Ristova-Sanyal, and Tony Wu for research assistance.

Ethics & International Affairs, 39, no. 4 (2025), pp. 312–335.
© The Author(s), 2026. Published by Cambridge University Press on behalf of Carnegie Council for Ethics in International Affairs. This is an Open Access article, distributed under the terms of the Creative Commons Attribution-NoDerivatives licence (http://creativecommons.org/licenses/by-nd/4.0), which permits re-use, distribution, and reproduction in any medium, provided that no alterations are made and the original article is properly cited.
doi:10.1017/S0892679426100343

https://doi.org/10.1017/S0892679426100343 Published online by Cambridge University Press

recently passed by the state of California, which we discuss below, little regulation appears to be in the offing.

Government regulation and Silicon Valley have rarely gone hand in hand, at least since the counterculture years.[7] Since the 1960s, companies have rejected government interference even as they sought government contracts. Even as some larger companies (so-called Big Tech) have called for the regulation of AI, smaller companies are more skeptical, worrying they do not have the resources to adapt to more onerous rules.[8] Even some Big Tech firms are now on record opposing most forms of regulation as an impediment to innovation.

Despite its idealized promise of increasing efficiency and perhaps accuracy in making decisions now made by humans, AI also has the potential to impose vast and unanticipated harms on society. Governing the sector seems like the obvious thing to do, yet efforts have been stymied, as the available options seem too flawed. One cognitive block lies in thinking of regulation in either-or terms. If we reject government intervention as too cumbersome to properly regulate a fast-moving sector, the solution is industry self-regulation. If we think industry is unable to tame its profit ambitions to protect socially beneficial outcomes, public regulation seems to be the answer.[9] Yet, there are other ways to think about the choices we have for regulation. To this end, we propose a "hybrid model of public-private regulation" for AI: the creation and enforcement of professional standards by industry, backed by government-enforced licensure.

This article is both normative and analytical. We advance an argument that a hybrid model of regulation holds promise in terms of articulating ethical principles and credentialing professionals in the AI industry, or "AI developers." At the same time, these principles need enforceability, with noncompliance having concrete consequences. Thus, AI developers ought to be governed under a regulatory regime with strong ethical standards similar to those already developed in other professions, such as in medicine, law, or psychology. These professions create disciplinary standards, ensure that individuals bearing the title of "physician," "attorney," or "psychologist" have met common educational requirements, maintain professional knowledge and standards, and honor ethical boundaries in order to continue being certified as licensed practitioners. These licenses are regularly renewed and may come with professional development or other requirements. These certifications allow those with such designations to demand higher salaries and influence how future generations of professionals are trained through standardized training and exams. They also create a focal point of responsibility.

https://doi.org/10.1017/S0892679426100343 Published online by Cambridge University Press

Doctors are sued for malpractice, lawyers can be disbarred for unethical behavior, and so on. Such models of professional regulation create ex ante accountability.

We also provide an analytical basis for our hybrid regulation proposal. AI has successfully disrupted human societies, changing how we understand ourselves and one another,[10] leaving many governments at a loss for the right mix of policies to encourage technological innovation while curbing the worst of AI's social and political effects. The divisions between tech optimists and tech pessimists, and Big Tech and the rest, signal that AI developers cannot be left to their own devices to self-regulate. We also know that public regulation is unlikely to prove effective on its own. Posed against a lightning-fast industry such as AI, governments can be ponderous. Current efforts to regulate AI, such as the EUAIA, do not inspire much optimism despite heavy boosterism from policymakers. The EUAIA (and other public governance instruments) took years to develop, making many of its strictures outdated on delivery. Besides the time lag, the EUAIA also relies on self-reporting by companies and a citizen-complaint mechanism that requires adjudicatory capacity that is still developing. The EUAIA focuses mostly on registration of AI models with a risk-ranking system that applies *only* to private firms with harms identified only ex-post. It is a step forward in highlighting ethical concerns, but an inadequate destination for AI regulation.

We first briefly discuss the promise and perils of AI. Second, we examine ongoing efforts at industry self-regulation and public regulation. The third section outlines a proposal for professional regulation of AI and why we view it as a plausible alternative to the status quo. Drawing on an analogy to the medical field, we introduce the history and reasoning behind medical licensing as practiced in the United States by the American Medical Association (AMA), and map some of these logics onto regulatory issues related to AI. We conclude with a discussion of possible objections to reiterate why we think professional regulation is necessary.

The Promise and Perils of Artificial Intelligence

Artificial intelligence is made up of three components: algorithms (instructions for the computer); data (vast amounts of binary or digital data); and "compute" (computing resources, often about processing power and capacity and data storage). Modern forms of AI that make headlines, such as generative AI (gen AI) models like ChatGPT, Midjourney, and Sora, demand massive amounts of all three components. That is, the AI systems that elicit the most public excitement demand sophisticated and intricate algorithms to direct them, unfathomable amounts of

https://doi.org/10.1017/S0892679426100343 Published online by Cambridge University Press

data to train those algorithms, and equally stunning computing power.[11] The vast scale of contemporary AI systems means that the industry is highly concentrated, with only a handful of firms developing foundational models, such as OpenAI (ChatGPT), Anthropic (Claude), and DeepSeek.[12] Using the EUAIA differentiation between types of companies, firms that produce foundational models are classed as "providers." Beyond providers, there is a large and rapidly growing niche of firms that use or operate AI systems. This includes firms selling AI services to other firms, firms using AI for their own purposes, and nonprofits that use AI. Firms in this tier are referred to as "deployers" in the EUAIA. Together, we refer to AI providers and deployers as the "AI complex." Finally, an even larger number of consumers use AI to write essays, find travel deals, or query the Internet.

The promise of AI is that it will improve everything: the economy, our relationships, warfare, and our personal lives.[13] It is for this reason that investors and states are pouring billions of dollars into AI research and development. Although it has the potential to displace workers in many previously "unautomated" and "white collar" occupations, it is an open question as to whether the net effect of AI on overall employment will be positive or negative,[14] a question we do not take up here. Exogenous shocks to social and political systems, such as AI, can take a long time to filter out.

The perils of AI are already being identified and sometimes felt, adding urgency to regulation. Problems include: the large carbon footprint of cutting-edge AI systems and the expanded datafication of our everyday lives,[15] which affect how human rights values such as equality, dignity, community, and autonomy are realized;[16] more generally reproduce distortions in the digital economy;[17] and concentrate political power in a handful of companies, which can then shape future legislation and regulations, further undermining principles and practices of democracy.[18]

Beyond these macrolevel effects, AI systems can have pernicious effects at the individual level. Most importantly, AI systems are inherently discriminatory in at least two ways. First, they are *by design* discriminatory. AI systems take massive pools of data, and categorize and distinguish between types of things in the data. Second, because AI systems are fed data about human behaviors and choices, they are informed by human bigotry. Humans are racist, sexist, ableist, and ageist. It should be no surprise to anyone, then, that AI systems replicate human biases that are present in the data, by being racist, sexist, ableist, and ageist. Much has been written about these issues elsewhere, so we will not attempt to recreate that work here.[19] While individual effects are important, it is essential to note that some of the

https://doi.org/10.1017/S0892679426100343 Published online by Cambridge University Press

harms expected to come from the widespread diffusion of AI are likely to be large-scale "negative externalities"; that is, the aggregate harms will affect society beyond the individual users. Were harms simply at the individual level, people could opt out by refusing to use AI in their daily lives or using only AI systems that reflect their ethical values—as we see individuals do with ethical consumption.[20] Not only are people's data scooped up in the training sets for large language models, or LLMs, but the use of these data in public decision-making about, say, policing or the allocation of social benefits is beyond the control of any individual person.

Finally, AI also challenges human dignity and, especially, creativity, one of the hallmarks of being human. Lee Saedol, at one point the top Go player in the world, who was beaten by an AI player, remarked recently, "People used to be in awe of creativity, originality and innovation . . . . But since A.I. came, a lot of that has disappeared."[21] AI systems take products of human creativity—whether written, drawn, or vocalized—and turn them into digital data. These data form the training data for AI systems, which do not recognize these creations as anything "artistic" but rather as information to analyze. In short, such systems can distort the very human sense of aesthetic appreciation, even as they easily render outcomes in the style of famous poets or artists.[22]

Some of these perils are about technological change that has led to widespread social and political change (for example, in terms of datafication and corporate power). Some of these perils are consequences of the maturing of the AI technologies themselves—unrealized economic promises, ecological effects, and winner-take-all dynamics. And some are perils that the makers, boosters, and users of these systems did not anticipate would become vexing societal problems in the way that they have (such as discrimination and threats to creativity and meaning). All of these perils are magnified by the lack of transparency and accountability in AI systems. It is an unfortunate characteristic of deep learning and gen AI systems that how they reach their predictions is largely opaque, though there are some methods for deducing the importance of certain weights in AI models.[23] With sometimes billions of parameters and weights that evolve through the training process, even their designers cannot understand or explain how the models arrive at their predictions. This implies that decisions made by AI systems are more difficult to contest; that is, a person harmed by a decision to, say, deny them social benefits may struggle to challenge the decision or hold a decision-maker accountable for harm or liable for any damages.[24] This is because existing laws are premised on a "mechanical" understanding in which failure or blame can be identified at some

https://doi.org/10.1017/S0892679426100343 Published online by Cambridge University Press

step in the process, or the guarantee that there is a "human in the loop" who somehow ensures an AI system works properly. A victim of an AI-based decision might appeal to a human to override the AI recommendation, but many systems preclude humans from explaining *why* decisions were made with specificity.[25] This will severely diminish the ability of individuals to challenge authorities and violations of their rights. Any ethical governance of AI systems, then, should tackle these problems.

Ethical Standards for AI

A well-designed AI ethics system will preserve the promises while minimizing the perils of such technologies. In their comprehensive analysis of AI ethics guidelines, Anna Jobin, Marcello Ienca, and Effy Vayena identify and catalog ethical principles that have emerged from various professional groups and other nonstate actors, noting coalescence around common themes.[26] They find that "transparency, justice and fairness, non-maleficence, responsibility, privacy, beneficence, freedom and autonomy, trust, dignity, sustainability, and solidarity" were the general principles that could be derived from the eighty-four statements they evaluated, though no document had all eleven principles.[27] Only five (transparency, justice and fairness, nonmaleficence, responsibility, and privacy) emerged with global consensus, though unsurprisingly there were marked differences in the scope and interpretation of these ideas across documents. The perils we identify above demonstrate that it is exceedingly challenging to figure out what ideals (such as nonmaleficence) mean in the context of technologies deployed widely in real-life markets and products. Furthermore, the AI complex, and especially cutting-edge gen AI models, is largely built on the backs of proprietary algorithms that run on datasets of unknown content. Transparency, and the accompanying value of responsibility, are also not just about algorithms—addressing these kinds of concerns requires targeting not just the AI systems themselves but also the assumptions built into those systems by their developers and as a result the overall economic context in which these technologies are created. Finally, the principle of privacy, though touted as a value by many technologists, is not the only human right threatened by AI.[28]

Still, AI regulations and ethics do not work well if the target of those efforts is left ambiguous. We argue here that, as implied in the EUAIA, the target should be the AI complex, but instead of fully top-down regulation it should be in the form of a

https://doi.org/10.1017/S0892679426100343 Published online by Cambridge University Press

license for AI developers that is articulated by those in the AI profession, mandates ethical processes and outcomes, and is enforced at the level of firms and universities, and ultimately, the state. What we propose is a hybrid structure, similar to those found in other professions, that mixes expertise with external enforcement. We argue that AI insiders should embrace such a model not only to create trust in their systems by a skeptical public but also to exert a measure of control over their industry, increase professionalism, and even possibly raise wages for licensed developers and the firms that employ them. Further to the point, most "professions" eventually develop licensing as a way to police the boundaries of which professionals count among them, and as a way to distribute the benefits of attaining extra accreditation.

Alternative Approaches to Regulation

At present, regulation of the AI complex is a governance anarchy in which firms are largely free to do whatever they want. With the exception of the EUAIA, discussed below, there is little in the way of limits on how firms develop and use AI. In the United States, the White House Office of Science and Technology Policy published the *Blueprint for an AI Bill of Rights* in October 2022, meant to guide businesses in implementing ethical AI systems, and President Biden issued an executive order titled "Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence" (EO 14110), both now rescinded by President Trump. Instead, the Trump administration issued its own executive order titled "Removing Barriers to American Leadership in Artificial Intelligence" (EO 14179) that removes regulatory barriers, slim as they already were. The Council of Europe launched its own Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law in 2024, which obligates states to defend human rights in the life cycle of AI. Given the paucity of restrictions on firms creating and using AI, it is no wonder they have largely done what they want.

Industry Self-Regulation

There is active interest within the AI complex in some form of self-regulation. Advocates argue that self-regulation would enable the technical experts who know best to set industry standards, allow for flexibility in handling fast-moving technology, and build trust through transparency. Nonetheless, others in the industry

https://doi.org/10.1017/S0892679426100343 Published online by Cambridge University Press

are opponents of self-regulation—or regulation in any form—and argue that any limits will inhibit innovation and delay the benefits that AI will bring.

Previously confined to industry insiders, this disagreement broke out into the open with a March 2023 open letter signed by over thirty thousand people advocating an AI "pause."[29] Things became more dramatic in November 2023 when Sam Altman, previously on the "go slow" side, was dismissed by the board of OpenAI because some members no longer trusted him to build AI in a way that would benefit humanity. There is also disagreement around how the code for AI systems should be distributed: open source, as in the case of Meta's Llama or DeepSeek, or held proprietarily, as in the case of OpenAI's ChatGPT or Anthropic's Claude.

Critics from outside the industry are also concerned that self-regulation led by for-profit firms may not give sufficient attention to public interest and ethical concerns. Likewise, self-regulation may lock in the favored positions of early innovators, concentrating political and economic power in those firms.[30] This risks stifling innovation not as a result of regulation per se, but by suppressing healthy competition within the complex.

Given a lack of trust by the public in AI and especially in self-regulation, a degree of public regulation is likely inevitable, even if erratic or different across jurisdictions. Nonetheless, we argue that some means of allowing expert involvement in setting standards and facilitating innovation should be built into the regulatory mechanism. Our proposed model of decentralized hybrid regulation seeks to do just that.

Public Regulation

The most developed form of public regulation of AI to date is the EUAIA, adopted in March 2024, although it will not come fully into force until 2027.[31] Relatedly, the EU's 2018 General Data Protection Regulation has also been hailed as a way to protect digital rights, but largely protects data rather than its uses in AI.[32] The EUAIA does two main things. First, it prohibits firms from deploying what it defines as unacceptably risky forms of AI. Second, it requires firms to register their high-risk AI applications in a central, EU- and member state–maintained repository.

The purpose of the EUAIA is to ensure that AI systems that operate within the EU are "safe, transparent, traceable, non-discriminatory, and environmentally friendly."[33] Three principles underlie the EUAIA, all relating to the fundamental

https://doi.org/10.1017/S0892679426100343 Published online by Cambridge University Press

principle of transparency: (1) explainability, such that the provider can explain in a comprehensible fashion how the AI system produced its output; (2) data transparency, including access to the original training data; and (3) user awareness; that is, individuals must be notified that they are interacting with an AI system.

Central to the EUAIA is its identification of risk levels. "Unacceptable risk" includes systems that engage in cognitive behavioral manipulation, social scoring, and biometric identification (including facial recognition). Such systems are banned within the EU (as of February 2025) for private actors, although governments are exempt from this restriction. Notably, the EUAIA allows law enforcement to use "real time" biometric identification in serious cases and remote biometric identification with delay to prosecute serious crimes with court approval. "High-risk products" include private-sector applications in the management and operation of critical infrastructure, education and vocational training, employment, worker management, access to self-employment, and access to essential public and private services. These applications must be assessed for transparency before being put on the market and over their life cycle must be registered in a publicly accessible EU database. Limited-risk applications must be transparent in both machine- and human-readable form but are otherwise without restrictions.[34] Providers and deployers self-designate their products to these categories, subject to a conformity assessment by the European Commission's European AI Office in conjunction with designated national supervisory authorities (starting thirty-six months after the EUAIA comes into force).[35] Nonconformance or failure to register are subject to significant fines. Although all European states were required to create national competent authorities by August 2025, only four states (Ireland, Spain, Malta, and Italy) have publicly announced such bodies.[36] Overall, the EUAIA does little to restrict the development of AI or its applications, relying on transparency before and after release to ensure that human rights are protected.

The EUAIA incorporates many of the ethical principles articulated in prior documents. Unfortunately, it has two broad limitations. First, like any public regulation, decision-making processes within the regulatory system are slow and cumbersome, unlikely to keep up with a fast-moving industry that still embodies the ethos of the early motto of Facebook (now Meta) to "move fast and break things." Second, it largely relies on ex-post monitoring by third parties. Beyond conformance assessments that focus on designating risk levels, the EUAIA relies on third parties at the state level, private entities, and individuals to identify violations after release. Though the EUAIA grants EU and state-level authorities rights to

https://doi.org/10.1017/S0892679426100343 Published online by Cambridge University Press

demand information, there are no formal or scheduled auditing procedures. The act largely relies on "downstream" users to initiate complaints with designated authorities, and fines can be assessed for violations of EUAIA provisions.

In addition to this European initiative, California recently passed the Transparency in Frontier Artificial Intelligence Act, which went into effect January 1, 2026. It is one of the strongest sets of regulations in the United States, despite being a diluted version of a previous bill that was vetoed by California's governor in 2024.[37] The act requires advanced AI companies with revenues greater than $500 million to report safety protocols used in building their technologies and mandates companies to report the greatest risks posed by their systems, subject to civil penalties for noncompliance. The bill also strengthens whistleblower protections for employees who warn about potential dangers. It was passed in September 2025 despite heated opposition from the tech industry. Though it emphasizes transparency, the California regulation, like the EUAIA, still relies on ex-post reporting of violations and harm.[38] This initiative is unlikely to take effect given President Trump's recent executive order blocking state-level AI regulations.[39]

Both the EUAIA and the California regulation are significant steps in the right direction. The transparency and civil penalties in these regulations can form part of the "public" side of our hybrid public-private model. Our proposed licensing model for AI developers, however, will require companies to certify ex ante that their systems meet designed ethical standards, provide common grounds by which to evaluate the qualifications of AI developers, and result in other professional benefits that we document next.

Professional Regulation: A Proposal

Given the shortcomings in both self- and public regulation, we propose a hybrid and decentralized regulatory regime, analogous to those found in other professions. In our view, public law is necessary to (1) set the broad ethical principles that guide professional standards, most likely requiring specification in statutes; (2) mandate that AI systems released to the public be certified as meeting professional standards by a licensed AI developer; (3) require licensing of AI developers; and (4) permit civil suits under which victims can receive compensation for violations of professional standards. In turn, involving professional organizations is necessary to set codes of conduct or professional standards based on the ethical principles defined in public law, and amend and adjust these codes in practice as

https://doi.org/10.1017/S0892679426100343 Published online by Cambridge University Press

the technology changes. These codes of professional conduct will serve as the standard for certification against which professional practitioners will be held in cases of professional discipline and civil suits. A mixed public-private body should write and administer licensing exams and accredit educational programs, with both including a strong ethical component. This same mixed public-private body should hear complaints of ethical violations from individuals and, if warranted, discipline AI developers with fines and other measures, up to revoking a license to operate.

Being a "professional" implies a set of core ethics. Professionals, of course, are good at their respective craft, whether it is soccer, medicine, or cleaning. They are experts: they know the ins and outs of what they do and have the training to view things differently from nonprofessionals. They are also expected to maintain a sense of ethical conduct that differentiates them from nonprofessionals *and* gives them a common framework that defines those who know best about the profession. As Leonard Seabrooke and Duncan Wigan argue, "Expertise is not simply a claim to superior knowledge about how things work, but also a claim about how things should be, which relies on moral authority."[40] Thus, when we propose professional regulation, with it must come not only expert knowledge and procedures on how to maintain and gain knowledge on a subject but also a key set of ethical principles that define the profession. Professionals know *X* subject matter and follow *Y* ethical principles; nonprofessionals do not. Professionals have insider (albeit, of course, biased) knowledge on how best to regulate their craft and thus should be part of the regulatory process.

Reasoning by Analogy: Modern Medicine

Medicine is a reasonable (but imperfect) analogy for AI professionals. Not only is medicine a technical field but ethics play a central role in establishing principles of care. Harms are possible at the individual level, but these same harms can affect public health as well. Doctors must prove both competence at a technical level and an understanding of ethical principles through various licensing exams developed by professional bodies. They are subject to state-enforced disciplinary procedures and the possibility of private suits for medical malpractice. Although the types of harms in medicine and AI are different, the process of regulating the practice of medicine is similar to what we propose here.[41]

Medicine in the twenty-first century has strict content expectations and ethical standards for practitioners. The reasoning for this is sound. Doctors can mistreat

https://doi.org/10.1017/S0892679426100343 Published online by Cambridge University Press

patients, causing further medical problems, violations of patient dignity, or even death. The modern setting and enforcement of ethical standards in the practice of medicine reflects a hybrid set of professional standards that have elements of both state and private associational enforcement. Private professional associations such as the American Medical Association develop codes of conduct that are then incorporated into professional education, licensing exams, and finally licensing standards. Although the privately formulated standards are not themselves law, they become incorporated into the profession at all levels and are enforced by public and nongovernmental entities, having the practical effect of law.[42] Equally, they form standards against which practice is judged in civil suits filed by victims.

Professional organizations and their ethical standards emerged to "professionalize" medicine, distinguishing between "quacks" and "qualified" doctors.[43] Prior to the creation of professional organizations, nearly anyone could hang out a shingle and advertise themselves as a doctor. To develop public trust in the profession, better-trained and more competent practitioners wanted some means to distinguish themselves to the public, and thus organized to create professional bodies that, in turn, set standards of training, competence, and ethics enforced by state-licensing boards. In many cases, licensure also served to limit competition in the profession, especially from an unlicensed fringe, leading to higher incomes than would otherwise have been the case.[44] Over time, standards within the profession increased further to distinguish specialists from generalists, and so on, creating a set of ever increasing professional standards.[45] We use examples from the United States to illustrate our logic, because the U.S. model has been a standard bearer for the profession.[46]

The AMA's Code of Medical Ethics is a cornerstone of the practice of medicine in the United States. Medical ethics have long been recognized and categorized into four main principles: beneficence (what physicians "should do"), nonmaleficence ("do no harm"), autonomy (patient's rights to decide care), and justice (equal treatment for all). The AMA is a private organization and the code itself does not have the force of law. Rather, the code guides law, practice, and education at multiple levels.

The AMA was founded in 1847 but for decades it struggled to establish its legitimacy against older competing medical systems. At the time of its creation, it was not clear that the AMA would succeed as a principal professional organization against its competitors, especially homeopathy.[47] Homeopathy is based on the philosophy and teachings of Samuel Hahnemann, a German physician, whose

https://doi.org/10.1017/S0892679426100343 Published online by Cambridge University Press

career spanned the 1770s–1840s. Dissatisfied with the state of the art in medical treatments that mainstream "allopathic" doctors used in his day—bloodletting, administering large doses of mercury, and purging—Hahnemann founded homeopathy as an alternative. Allopathic, or orthodox medicine, or what is sometimes called "biomedicine," became the foundation for the AMA.[48] In its heyday, homeopathy proved compelling enough that a professional association, the American Institute of Homeopathy, was founded in 1844, predating the AMA by several years. The allopath-homeopath distinction endures to this day. Hahnemann's techniques of creating remedies through "proving" (testing the negative effects of various substances, or "causes," on healthy individuals, and then using an extreme dilution of the causes as remedy) have ultimately not stood up to scientific testing, but homeopathic practices continue.[49] In spite of accusations of quackery dating back to at least the 1850s[50] and outright characterizations of Hahnemann as insane,[51] a détente seemed to be reached at the turn of the twentieth century when the AMA invited homeopaths to join its ranks. In the end, the AMA "won" the professional competition, defining "medicine" as what is practiced by those certified as MDs, doctors of osteopathic medicine, and select other designations (but *not* homeopaths).[52]

Even after the establishment of the AMA, American medical training was spotty and inconsistent, at least by our modern standards. The AMA was initially unsuccessful in lobbying for standardized medical training, even after establishing the Council on Medical Education in 1904 with the goal of revamping U.S. medical training.[53] The majority of medical schools were not attached to universities, existing instead as for-profit storefronts with little or even no applied training, involving varying entry and completion standards, and consisting of full apprenticeships that did not require classroom learning. It was not until after *Medical Education in the United States and Canada*, or the "Flexner Report," was published in 1910, commissioned by the Carnegie Foundation, that the American medical system became systematized. Abraham Flexner, the author of the report, visited all 155 U.S. and Canadian medical schools, finding only thirty-one schools up to par and recommending the rest be shut down.[54] It is in this context that many "alternative" medical approaches—including homeopathy—got pushed aside in medical education and certification. In the same decade as the report came out, state licensing boards began pushing more rigorous standards of admission and curriculum on medical schools.[55] It should be noted that the dramatic closure and restructuring of American medical schools was underwritten by major philanthropic organizations,

https://doi.org/10.1017/S0892679426100343 Published online by Cambridge University Press

most significantly the Carnegie and Rockefeller foundations.[56] In this way, allopathic medicine, embodied in the AMA, eventually emerged as the primary association for medicine by demeaning homeopathy as "quackery"[57] and continually upgrading its own standards of professionalism.[58] The key point, however, is that physicians professionalized to distinguish themselves from other practitioners and, once begun, continued to upgrade standards to stay ahead of competing professional organizations.

Individuals today are beneficiaries of the tangle of standards creation, turf wars, and the generously funded restructuring of doctors as professionals. The overall structure of the medical system, including standards of physician competence and behavior, is now set by state-level medical practice acts, typically specified in business and professions codes that govern licensed practitioners; codes protecting the rights of patients and clients; and health and safety codes. Under medical practice acts, physicians are required to obtain a license to practice medicine within each state, typically issued by a publicly appointed state medical board. To obtain a license, the physician must pass the United States Medical Licensing Examination (USMLE) during their medical education and residency training. The USMLE is administered by the Federation of State Medical Boards (FSMB) and the National Board of Medical Examiners, a private nonprofit organization. The test itself is designed by the faculties of medical schools, members of the state medical boards, and practicing physicians. Questions of medical ethics, based on the AMA's Code of Medical Ethics, are integral to the exams at both the educational and residency levels.

To obtain a license, physicians must also graduate from an accredited medical education program. Many of the accreditation organizations are nonstate bodies. Medical schools are accredited by the Liaison Committee on Medical Education (LCME), sponsored by the Association of American Medical Colleges and the AMA. The LCME requires all medical programs to offer a curriculum in medical ethics. In addition, physicians must fulfill a residency program accredited by the nonprofit Accreditation Council for Graduate Medical Education, which also mandates a curriculum in medical ethics. The ethics curriculum at both the MD and residency levels is based on the AMA Code.

Medical ethics are enforced by the state medical boards and civil suits for malpractice. After receiving a complaint from a patient, a state medical board investigates and, if the physician is found to be in violation of the standard of professionalism, the board may impose a variety of punishments, from a fine to

https://doi.org/10.1017/S0892679426100343 Published online by Cambridge University Press

revocation of the practitioner's license. The FSMB maintains a public database of licensee information and disciplinary actions. Physicians are also liable in civil courts for medical malpractice, which includes violations of ethical standards. The question in civil suits is whether the physician followed standard practices as defined by the profession in general and interpreted in specific situations by other medical professionals who may testify as "expert witnesses." In short, the professionally defined standards serve as the baseline against which questions of appropriate practice are assessed in specific situations.

Professionalizing AI

In the medical analogy, practice and ethical standards defined by the profession are reinforced by and inform public law, creating a hybrid system of rules and enforcement procedures. We propose a similar system for AI, combining public law and professional standards and the setting of standards.

In a fast-moving field such as AI, public law cannot set specific rules to regulate practice, as demonstrated by the limits of the EUAIA. Legislators are not experts in the relevant fields, nor should they be. Lawmaking is cumbersome and slow. Law is also a blunt instrument that is not effective at regulating conditions in which judgement in real time by practitioners is essential. For these reasons, public law is best suited to define the ethical principles of regulation, articulate minimal standards for practitioners in a given industry, and create the framework for determining competence and enforcement of professionally determined standards. What constitutes "good work" or "appropriate work" in a field should otherwise be left to those in the profession.

A key component of hybrid regulation, as in medicine, is a license to practice as a professional. Licensure is also central to regulating AI. We propose that all AI systems should be certified as compliant with legally defined principles and professional standards by a licensed AI developer before being released to the public. This should apply to all providers and deployers operating within a specific jurisdiction as well as foreign firms selling AI systems within that territory. Certification in any case might be conducted in an "audit" model, wherein any firm offering an AI system to the public must submit its final product to a licensed professional within that jurisdiction, or "internally" where the firm hires its own AI developer to certify the system.[59] Certification will ensure that, say, the data used to train a foundational model is not copyrighted or used without permission from the

https://doi.org/10.1017/S0892679426100343 Published online by Cambridge University Press

originator, is anonymized to protect privacy, and so on, comparing the system to the ethical principles outlined in the professional standards. It could also ensure that the output of any AI system is unbiased and transparent as defined by law and the profession. Licensing can lead to best practices if the contents of certifications are made accessible to professionals and to the public, helping spread transparency norms. As is the case in any doctor's office, not all health workers need to be licensed doctors, but in terms of AI systems, any system must be certified by a licensed AI developer who then bears ultimate responsibility for its compliance with standards set by the profession. The standards—including ethical standards—are likely best set by professional bodies, but licenses and the certification of AI systems must be required by public law to be enforceable. In this scheme, individuals have the obligation to comply not only from an ethical standpoint but also because of legal constraints.

In tandem with licensing, public law must set procedures for disciplining licensed professionals. One component of any enforcement mechanism is placing restrictions on the licenses of AI developers who violate professional standards, including limiting the kinds of systems a licensee might certify, suspending a license for some period of time, or revoking the license completely. A second component is creating procedures for civil suits and penalties for harms experienced by victims of violations of professional standards—much like malpractice suits for doctors. The issue in such suits is whether the licensee violated professional standards, which would be contested in trials with other licensed practitioners testifying about what would be expected of a professional in similar cases. Again, the analogy holds for AI providers and deployers as to whether they engaged in due diligence according to professional standards before certifying an AI system.

The framework set by public law, in turn, must be filled by standards set by practitioners that are consistent with the goals of regulation as defined by public law. As in medicine, professionals in AI should, as a body, define what competence means in its fullest sense, including appropriate ethics. To set standards will require the creation of one or more parallel professional associations with the infrastructure necessary to deliberate and articulate professional standards across AI systems. Associations may specialize across different segments of the AI complex, with perhaps one association and set of standards for providers working on foundational models and a different association and set of standards for the deployers of these models.[60] As the AI complex grows and the technology becomes increasingly

https://doi.org/10.1017/S0892679426100343 Published online by Cambridge University Press

sophisticated and differentiated, such specialization is expected. Professional associations should be charged with continually monitoring developments in the field and updating standards accordingly. This means that the current demands on compute, data, and algorithms may change as AI technologies mature. Given that innovation creates a moving target for datafication, one of the perils identified above, one could expect that if governments move to regulate datafication practices, a professional association could set some boundaries or help encourage different kinds of AI developments, such as less data-demanding LLMs.[61]

Rigorous licensing exams that embody these professional standards will generate incentives for educational programs to incorporate training in ethics into their curricula. This is already happening voluntarily at some universities, but it is far from expected for graduates of computer science programs and other fields to have ethical training.[62] Ethics training can address problems of discrimination and bias, another peril identified above. Even though AI systems are inherently discriminatory, budding AI developers can be taught that the assumptions they build into such systems, whether on the data or algorithm side, can cascade through a system in many unanticipated ways. Educational programs that want their students to pass the licensing exam will naturally want to prepare them for the content and questions on the exam. Equally, as the professional standards get elaborated, professional associations may accredit educational programs, as in medicine, that incorporate the relevant standards into their coursework. In this way, ethical standards will diffuse across the profession.

Discussion

Our proposed governance structure for AI is an exercise in mechanism design, providing an architecture for professionalization. As a proposal, we cannot provide evidence in advance that it will "work." Rather, we can only reason by analogy with an understanding of the purposes of regulation and the incentives of the relevant actors within the AI complex. Central to our proposal is the requirement that AI systems be certified as meeting professionally agreed-upon standards before being released to the public, and that this certification must be performed by a licensed professional who is then held accountable for the system's performance. Unlike either open-source code, advocated by some in the AI complex, or registration systems, as in the EUAIA, certification is performed before release and does not rely upon sleuths to discover problems after they may have already imposed harm.

https://doi.org/10.1017/S0892679426100343 Published online by Cambridge University Press

The possibility of civil suits, however, also allows for ex-post monitoring and penalties when professional standards are violated.

We anticipate some pushback to our proposal. One objection might be that relying on AI developers to be integrally involved in setting professional standards may lead to lowest common denominator regulations. Yet, given the technical knowledge required and a fast-moving field, expertise held by professionals is likely to be required for creating any effective rules. At the same time, it is appropriate that professionals from the nonprofit sector (including universities) be fully engaged in the process of determining what is appropriate and expected behavior by AI developers to safeguard against weak standards advocated by for-profit firms.[63] It is also true that standards may begin with relatively general and perhaps "loose" expectations, but as professionals seek to distinguish themselves from nonprofessionals and aim to develop public trust in their systems the standards are likely to rise over time. From a compliance and legitimacy view, it is better to start with those who do the work than to presume they cannot possibly contribute to governing the industry. As the rewards of licensure become clear, the costs associated with becoming certified should become appealing for those who want the added benefits (and responsibility) of additional designations. The hybridity of the system puts states in the role of ensuring that standards do not result in a race to the bottom.

A second possible objection to our proposal is that national firms in regulated markets will be at a competitive disadvantage relative to foreign firms in unregulated markets. In this view, since we cannot move directly to a global regulatory standard and mechanism, it is unfair to U.S. (or other national) firms to impose a national regulatory standard on them alone. There are, at least, four possible responses to this objection.

First, decentralized, hybrid regulation does not automatically stifle innovation. In medicine, a highly regulated field, innovation is robust and new surgical techniques, drugs, and treatment practices diffuse rapidly via incentives to provide the best-possible care as defined by frequently updated professional standards. Although any form of regulation may inhibit some innovation at the margin, professionally set standards will also safeguard society against ideas and practices that might harm the public (for example, the idea that injecting bleach is a prophylactic or cure for COVID-19). Innovation and public safety are not at odds with one another. If anything, establishing baselines helps set common ground for competition, rather than a contest for who can be the most ruthless, amoral, or cunning.

https://doi.org/10.1017/S0892679426100343 Published online by Cambridge University Press

Second, regulation can increase public trust in products and thereby increase "uptake" or consumer demand, which also serves to stimulate innovation. Even as the tech industry and enthusiasts pour billions of dollars into AI, the broader public remains skeptical and concerned about the perils discussed earlier. Setting professional standards reassures the public that innovations are safe and can be trusted. At the very least, it makes clear who is declaring the standards, based on clear reasoning. By analogy again, standardizing the medical profession was key to building public trust in doctors as authorities in health. AI is in a similar situation today wherein certifying that systems must meet professional standards is necessary to building public trust and acceptance. For this reason, it is in the interest of the AI complex to support the setting of rigorous professional standards—even if they apply only in some countries.

Third, and most importantly, although in its initial stages professional standards in AI may be limited to national standards, there is a strong potential for a "race to the top" in standards setting. Here, the analogy with medicine breaks down because of the localization of standards.[64] Foreign medical licenses are rarely accepted in any country, and in the United States even state licenses are rarely accepted across state lines—even as the standards themselves may differ only slightly. As a technology that can be easily transported over the Internet, AI (to date) lacks that same geographic specificity. For some, this counsels against any attempt at national regulation. Nonetheless, national regulation in large markets like the United States or EU can rapidly become an international standard. Here, market power matters. If foreign firms wishing to sell systems in the United States must have their products certified by a U.S. licensed AI developer as meeting the standards set by U.S. professionals, they are more likely to produce systems for all markets that also meet these standards. A better analogy here is to California emissions standards for automobiles, which apply only in the state but are typically met for cars sold elsewhere; it is not cost effective to produce two different types of autos for two different markets, so the "tighter" regulations come to prevail.[65] Likewise, it is expensive to train foundational models, likely making it inefficient to develop a U.S. model that operates somewhat differently from, say, an Indian model. If adopted in major markets, certification requirements create incentives for national standards to converge internationally.[66] Were the United States and the EU to adopt similar standards for AI, the pull of the two markets combined might be inexorable. Once diffused sufficiently, broadly accepted standards for AI might be codified and reinforced by an international organization, analogous to the World Health Organization for medicine.

https://doi.org/10.1017/S0892679426100343 Published online by Cambridge University Press

Finally, it is important to acknowledge that professional standards do not always travel neatly between countries. Drugs approved for use in one country are not automatically approved for use elsewhere. Attitudes toward risk and, therefore, food and safety standards differ across countries as well. A second and perhaps more apt comparison is with human rights. Though human rights are universal in principle, they are geographically particular in practice.[67] That China does not respect freedom of the press, for instance, does not mean that the United States should abandon its principle of freedom of the press. Similarly, countries that do not respect freedom of the press do not automatically affect journalism in countries that do. The same goes for prohibitions on importing goods produced with forced labor. We can insist that AI firms operating in the U.S., wherever they may originate, respect ethical principles as they are defined and practiced in the United States. If other countries allow AI firms to violate ethical standards held by liberal democracies, for instance, this does not mean that these firms should be allowed to violate those rights when working in the United States, Canada, or Europe.

We have already seen how Big Tech adjusts its practices for certain markets or refuses to participate in them (for example, Apple's relationship to China vs. Google's exit from the market in 2010). Firms in many industries could increase their profits if unregulated or allowed to operate everywhere under the rules of the least-regulated country. But some countries choose to protect certain rights of citizens, even at the cost of lower profits for firms. Models do travel, as we saw with the spread of U.S. governance of medicine. In the tech space, some have argued that the "Brussels Effect" spreads European standards beyond the continent.[68] AI should be no different.

In this article, we draw explicitly upon medicine for existing examples of hybrid governance. More broadly, licensing and professional standards are used to regulate lawyers, mental health professionals, engineers, architects, building contractors, plumbers, and aestheticians. Although we associate the tech industry with a hacker ethos, this should not be a barrier to regulation. We see no reason why professionalization will not work to robustly ensure competence and ethical practices in AI as it does in these other fields. Professionalizing AI as an industry simultaneously includes the most relevant stakeholders—AI developers—in creating and maintaining governance. It also allows for third-party involvement in the creation, implementation, and enforcement of professional standards backed by public law that requires certification and provides a framework for accountability.

https://doi.org/10.1017/S0892679426100343 Published online by Cambridge University Press

Notes

[1] Rishi Iyengar, "OpenAI's CEO Goes on a Diplomatic Charm Offensive," *Foreign Policy*, June 20, 2023, foreignpolicy.com/2023/06/20/openai-ceo-diplomacy-artificial-intelligence/.

[2] Mario Draghi, *The Future of European Competitiveness*, Part A, "A Competitiveness Strategy for Europe" (Luxembourg: Publications Office of the European Union, 2024), commission.europa.eu/topics/competitiveness/draghi-report_en

[3] Dan Milmo and Eleni Courea, "US and UK Refuse to Sign Paris Summit Declaration on 'Inclusive' AI," *Guardian*, February 11, 2025, www.theguardian.com/technology/2025/feb/11/us-uk-paris-ai-summit-artificial-intelligence-declaration.

[4] "White House Unveils America's AI Action Plan," White House, President Donald J. Trump, July 23, 2025, www.whitehouse.gov/articles/2025/07/white-house-unveils-americas-ai-action-plan/.

[5] Donald J. Trump, "Ensuring a National Policy Framework for Artificial Intelligence" (executive order, December 11, 2025), www.whitehouse.gov/presidential-actions/2025/12/eliminating-state-law-obstruction-of-national-artificial-intelligence-policy/.

[6] Michael Kratsios, quoted in Jared Perlo, "U.S. Rejects International AI Oversight at U.N. General Assembly," NBC News, September 27, 2025, www.nbcnews.com/tech/tech-news/us-rejects-international-ai-oversight-un-general-assembly-rcna233478.

[7] See Margaret O'Mara, *The Code: Silicon Valley and the Remaking of America* (New York: Penguin Press, 2019); and Fred Turner, *From Counterculture to Cyberculture: Stewart Brand, the Whole Earth Network, and the Rise of Digital Utopianism* (Chicago: University of Chicago Press, 2006).

[8] Gerrit De Vynck, "Big Tech Wants AI Regulation. The Rest of Silicon Valley Is Skeptical," *Washington Post*, November 9, 2023, www.washingtonpost.com/technology/2023/11/09/ai-regulation-silicon-valley-skeptics/.

[9] Alyssa Wong, "Regulatory Gaps and Democratic Oversight: On AI and Self-Regulation," Schwartz Reisman Institute for Technology and Society, University of Toronto, September 21, 2023, srinstitute.utoronto.ca/news/tech-self-regulation-democratic-oversight.

[10] John Cheney-Lippold, *We Are Data: Algorithms and the Making of Our Digital Selves* (New York: NYU Press, 2017); and E. Stefan Kehlenbach, "The Subatomic Person: A New Ontology of Big Data," *Theory & Event* 25, no. 4 (October 2022), pp. 851–72.

[11] Wendy H. Wong, "Don't Forget That There Are People in the Data: LLMs in the Context of Human Rights: A Response to Samuel R. Bowman," *Critical AI* 2, no. 2 (October 2024).

[12] Amba Kak, Sarah Myers West, and Meredith Whittaker, "Make No Mistake—AI Is Owned by Big Tech," Opinion, *MIT Technology Review*, December 5, 2023, www.technologyreview.com/2023/12/05/1084393/make-no-mistake-ai-is-owned-by-big-tech/; and Jai Vipra and Sarah Myers West, *Computational Power and AI* (AI Now Institute, September 27, 2023), ainowinstitute.org/publication/policy/compute-and-ai#h-what-is-compute-and-why-does-it-matter.

[13] Wim Naudé, "AI's Current Hype and Hysteria Could Set the Technology Back by Decades," Conversation, July 24, 2019, theconversation.com/ais-current-hype-and-hysteria-could-set-the-technology-back-by-decades-120514.

[14] Steve Lohr, "How One Tech Skeptic Decided A.I. Might Benefit the Middle Class," *New York Times*, April 1, 2024, www.nytimes.com/2024/04/01/business/ai-tech-economy.html.

[15] See Nick Couldry and Ulises A. Mejias, *The Costs of Connection: How Data Is Colonizing Human Life and Appropriating It for Capitalism* (Stanford, Calif.: Stanford University Press, 2019); Alexandra Sasha Luccioni and Alex Hernandez-Garcia, "Counting Carbon: A Survey of Factors Influencing the Emissions of Machine Learning," arXiv, Cornell University, 2023, arxiv.org/abs/2302.08476; and Sasha Luccioni, Yacine Jernite, and Emma Strubell, "Power Hungry Processing: Watts Driving the Cost of AI Deployment?," in *FAccT '24: Proceedings of the ACM Conference on Fairness, Accountability, and Transparency* (New York: Association for Computing Machinery, June 5, 2024), pp. 85–99; Ulises A. Mejias and Nick Couldry, "Datafication," *Internet Policy Review* 8, no. 4 (November 2019), pp. 1–10; and Jose van Dijck, "Datafication, Dataism and Dataveillance: Big Data between Scientific Paradigm and Ideology," *Surveillance & Society* 12, no. 2 (May 9, 2014), pp. 197–208.

[16] Wendy H. Wong, *We, the Data: Human Rights in the Digital Age* (Cambridge, Mass.: MIT Press, 2023).

[17] Nikos Smyrnaios, *Internet Oligopoly: The Corporate Takeover of Our Digital World* (Bingley, U.K.: Emerald, 2018).

[18] Wendy H. Wong, Jamie Duncan, and David A. Lake, "Why Data about People Are So Hard to Govern," *Regulation & Governance* 19, no. 1 (January 2025), pp. 236–52.

[19] See, for example, Ruha Benjamin, *Race after Technology: Abolitionist Tools for the New Jim Code* (Cambridge, U.K.: Polity, 2019); Joy Buolamwini, *Unmasking AI: My Mission to Protect What Is Human in a World of Machines* (New York: Random House, 2023); Virginia Eubanks, *Automating Inequality:*

https://doi.org/10.1017/S0892679426100343 Published online by Cambridge University Press

How High-Tech Tools Profile, Police, and Punish the Poor (New York: St. Martin's, 2018); Safiya Umoja Noble, *Algorithms of Oppression: How Search Engines Reinforce Racism* (New York: NYU Press, 2018); and Cathy O'Neil, *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy* (Portland: Broadway Books, 2017).

[20] See Tim Bartley, "Certifying Forests and Factories: States, Social Movements, and the Rise of Private Regulation in the Apparel and Forest Products Fields," *Politics & Society* 31, no. 3 (September 2003), pp. 433–64; and Tim Bartley, *Rules without Rights: Land, Labor, and Private Authority in the Global Economy* (New York: Oxford University Press, 2018).

[21] Lee Saedol, quoted in Daisuke Wakabayashi and Jin Yu Young, "Defeated by A.I., a Legend in the Board Game Go Warns: Get Ready for What's Next," *New York Times*, updated July 11, 2024, www.nytimes.com/2024/07/10/world/asia/lee-saedol-go-ai.html.

[22] Rebecca Marrone, David Cropley, and Kelsey Medeiros, "How Does Narrow AI Impact Human Creativity?," *Creativity Research Journal* (2024), pp. 1–11.

[23] Adityanarayanan Radhakrishnan, Daniel Beaglehole, Parthe Pandit, and Mikhail Belkin, "Mechanism for Feature Learning in Neural Networks and Backpropagation-Free Machine Learning Models," *Science* 383, no. 6690 (March 7, 2024), pp. 1461–67.

[24] Frank Pasquale, *The Black Box Society: The Secret Algorithms That Control Money and Information* (Cambridge, Mass.: Harvard University Press, 2015).

[25] Bowen Long, Enjie Liu, Renxi Qiu, and Yanqing Duan, "Explainable AI: The Latest Advancements and New Trends," arXiv, Cornell University, 2025, arxiv.org/abs/2505.07005.

[26] Anna Jobin, Marcello Ienca, and Effy Vayena, "The Global Landscape of AI Ethics Guidelines," *Nature Machine Intelligence* 1, no. 9 (September 2, 2019), pp. 389–99.

[27] Ibid., p. 391.

[28] Filippo A. Raso, Hannah Hilligoss, Vivek Krishnamurthy, Christopher Bavitz, and Levin Kim, "Artificial Intelligence & Human Rights: Opportunities & Risks" (Research Publication No. 2018-6, Berkman Klein Center for Internet & Society, Harvard University, September 15, 2018); and Office of the United Nations High Commissioner for Human Rights (OHCHR), "UN B-Tech Generative AI Human Rights Summit: Advancing Rights-Based Governance and Business Practice: Summary Note" (Geneva, November 30, 2023); and Wong, *We, the Data*.

[29] "Pause Giant AI Experiments: An Open Letter," Future of Life Institute, March 22, 2023, futureoflife.org/open-letter/pause-giant-ai-experiments/.

[30] Wong et al., "Why Data about People Are So Hard to Govern."

[31] The OECD maintains a comprehensive database of national-level AI governance initiatives. It covers sixty-nine countries, territories, and the EU, and a thousand different policy instruments. Furthermore, the UN High Commissioner for Human Rights has recently engaged with business and other nonstate actors on human rights concerns brought on by gen AI (OHCHR, "UN B-Tech Generative AI Human Rights Summit") as part of a broader sweep of moving away from a state-centric focus for human rights. Within the United States, thirty-eight states have passed over a hundred regulations, a concern even for AI companies that support some form of regulation but that complain that they cannot comply with so many different and shifting standards.

[32] Anu Bradford, *Digital Empires: The Global Battle to Regulate Technology* (Oxford: Oxford University Press, 2023); and Chris Jay Hoofnagle, Bart van der Sloot, and Frederik Zuiderveen Borgesius, "The European Union General Data Protection Regulation: What It Is and What It Means," *Information & Communications Technology Law* 28, no. 1 (March 2019), pp. 65–98.

[33] "EU AI Act: First Regulation on Artificial Intelligence," European Parliament, February 19, 2025, www.europarl.europa.eu/topics/en/article/20230601STO93804/eu-ai-act-first-regulation-on-artificial-intelligence. See also David Nolan, Hajira Maryam, and Michael Kleinman, "The Urgent But Difficult Task of Regulating Artificial Intelligence," Amnesty International, January 16, 2024, www.amnesty.org/en/latest/campaigns/2024/01/the-urgent-but-difficult-task-of-regulating-artificial-intelligence/.

[34] Alex Engler, "The EU AI Act Will Have Global Impact, But a Limited Brussels Effect," Brookings, June 8, 2022, www.brookings.edu/articles/the-eu-ai-act-will-have-global-impact-but-a-limited-brussels-effect/.

[35] "Artificial Intelligence (AI) Act: Council Gives Final Green Light to the First Worldwide Rules on AI," European Council and Council of the European Union, May 21, 2024, www.consilium.europa.eu/en/press/press-releases/2024/05/21/artificial-intelligence-ai-act-council-gives-final-green-light-to-the-first-worldwide-rules-on-ai/.

[36] Cristina Villasante, "The EU AI Act Is Ready. And the Regulators? Not So Much," ECIJA, August 11, 2025, www.ecija.com/en/news-and-insights/la-ley-de-ia-de-la-ue-esta-lista-los-reguladores-no-tanto/.

[37] "Governor Newsom Signs SB 53, Advancing California's World-Leading Artificial Intelligence Industry," Governor Gavin Newsom, September 29, 2025, www.gov.ca.gov/2025/09/29/governor-newsom-signs-sb-53-advancing-californias-world-leading-artificial-intelligence-industry/.

https://doi.org/10.1017/S0892679426100343 Published online by Cambridge University Press

[38] Hayden Field, "SB 53, the Landmark AI Transparency Bill, Is Now Law in California," Verge, September 29, 2025, www.theverge.com/ai-artificial-intelligence/787918/sb-53-the-landmark-ai-transparency-bill-is-now-law-in-california; and Cecilia Kang, "California Governor Signs Sweeping A.I. Law," *New York Times*, updated September 30, 2025, www.nytimes.com/2025/09/29/technology/california-ai-safety-law.html?smid=nytcore-ios-share&referringSource=articleShare.

[39] Jonathan J. Cooper, "Trump Signs Executive Order to Block State AI Regulations," Associated Press, updated December 11, 2025, apnews.com/article/trump-ai-regulation-executive-order-state-laws-9cb4dd1bc249e40426ob3dc233217388.

[40] Leonard Seabrooke and Duncan Wigan, "Powering Ideas through Expertise: Professionals in Global Tax Battles," in "Ideas, Political Power, and Public Policy," special issue, *Journal of European Public Policy* 23, no. 3 (2016), pp. 357–74, at p. 360.

[41] Other possible analogies include law, social work, and engineering, which also require licensing. Other analogies are less appropriate, with public and private regulation focusing on technical, not ethical standards. The Institute of Electrical and Electronics Engineers, for instance, is not involved in licensing practicing engineers, but is an academic organization for promoting research. Underwriters Laboratories and the International Organization for Standardization are professional bodies that establish standards with the aim of improving consumer goods and manufacturing practices. Thus, UL's approval ensures that your toaster will not explode and ISO standards ensure that companies have standard risk management plans, not whether your toaster uses too much power or whether computers ought to process data about people.

[42] Andrew H. Beck, "The Flexner Report and the Standardization of American Medical Education," *JAMA* 291, no. 17 (2004), pp. 2139–40, at p. 2140.

[43] Carol Rice Andrews, "Standards of Conduct for Lawyers: An 800-Year Evolution Annual Survey of Texas Law," *SMU Law Review* 57 (2004), pp. 1385–1458.

[44] Mike Mitka, "The Flexner Report at the Century Mark," *JAMA* 303, no. 15 (2010), pp. 1465–66.

[45] Molly Cooke, David M. Irby, William Sullivan, and Kenneth M. Ludmerer, "American Medical Education 100 Years after the Flexner Report," *New England Journal of Medicine* 355, no. 13 (September 28, 2006), pp. 1339–44; and Atul Gawande, *Being Mortal: Medicine and What Matters in the End* (New York: Henry Holt, 2014).

[46] Thomas P. Duffy, "The Flexner Report—100 Years Later," *Yale Journal of Biology and Medicine* 84, no. 3 (2011), pp. 269–76.

[47] A similar history can be told about chiropractic and the American Chiropractic Association.

[48] Stephen Barrett, "Homeopathy: Is It Medicine?," in "Medical Controversies," *Skeptical Inquirer* 12, no. 1 (Fall 1987), skepticalinquirer.org/1987/10/homeopathy-is-it-medicine/; Edzard Ernst and Ted J. Kaptchuk, "Homeopathy Revisited," *JAMA Internal Medicine* 156, no. 19 (October 28, 1996), pp. 2162–64; Wayne B. Jonas, Ted J. Kaptchuk, and Klaus Linde, "A Critical Overview of Homeopathy," *Annals of Internal Medicine* 138, no. 5 (March 2003), pp. 393–99; and Irvine Loudon, "A Brief History of Homeopathy," *Journal of the Royal Society of Medicine* 99, no. 12 (December 2006), pp. 607–10.

[49] Cemre Cukaci, Michael Freissmuth, Christopher Mann, Joshua Marti, and Veronika Sperl, "Against All Odds—the Persistent Popularity of Homeopathy," *Wiener Klinische Wochenschrift* 132, no. 9 (May 2020), pp. 232–42.

[50] See, for example, J. G. Affleck, "Annual Reports," *Belmont Medical Society*, 1855.

[51] W. W. Parker, "Was Hahnemann Insane?," *Journal of the American Medical Association* 26, no. 1 (January 4, 1896), pp. 7–9.

[52] "AMA Membership Eligibility," American Medical Association, www.ama-assn.org/member-benefits/member-eligibility-dues/ama-membership-eligibility.

[53] Beck, "Flexner Report and the Standardization of American Medical Education."

[54] Mitka, "Flexner Report at the Century Mark."

[55] Beck, "Flexner Report and the Standardization of American Medical Education."

[56] Duffy, "Flexner Report," p. 273.

[57] See, for example, Affleck, "Annual Reports"; and Parker, "Was Hahnemann Insane?"

[58] Cukaci et al., "Against All Odds."

[59] Graeme Auld, Ashley Casovan, Amanda Clarke, and Benjamin Faveri, "Governing AI through Ethical Standards: Learning from the Experiences of Other Private Governance Initiatives," in "The Governance of Artificial Intelligence," special issue, *Journal of European Public Policy* 22, no. 11 (2022), pp. 1822–44.

[60] An analogy here is the difference between certified financial analysts, who actually make and direct investments, and certified financial planners, who advise clients on investment strategies.

[61] See, for example, Brenden M. Lake and Marco Baroni, "Human-Like Systematic Generalization through a Meta-Learning Neural Network," *Nature* 623, no. 7985 (November 2023), pp. 115–21.

https://doi.org/10.1017/S0892679426100343 Published online by Cambridge University Press

[62] Casey Fiesler, "Tech Ethics Curricula: A Collection of Syllabi," Medium, July 5, 2018, cfiesler.medium.com/tech-ethics-curricula-a-collection-of-syllabi-3eedfb76be18.
[63] Wong, *We, the Data*, pp. 58–60.
[64] Localization was not inevitable but a function of path-dependent developments in regulatory structures.
[65] David Vogel, *California Greenin': How the Golden State Became an Environmental Leader* (Princeton, N.J.: Princeton University Press, 2018).
[66] On diffusion more generally, see Beth A. Simmons, Frank Dobbin, and Geoffrey Garrett, "Introduction: The International Diffusion of Liberalism," *International Organization* 60, no. 4 (October 2006), pp. 781–810. For a model of the "upgrading" of standards, see Susan D. Hyde, *The Pseudo-Democrat's Dilemma: Why Election Monitoring Became an International Norm* (Ithaca, N.Y.: Cornell University Press, 2011).
[67] Beth A. Simmons, *Mobilizing for Human Rights: International Law in Domestic Politics* (Cambridge, U.K.: Cambridge University Press, 2009).
[68] Anu Bradford, *The Brussels Effect: How the European Union Rules the World*, eBook (New York: Oxford University Press, 2019).

Abstract: The consensus on the need to regulate artificial intelligence is clear, but the how remains elusive. Private regulation, as proposed by the tech industry itself, and state regulation, as embodied in the recent EU Artificial Intelligence Act, are two common forms of governance. We advance a third option that has received very little attention to date: professional regulation. Professional regulation is modeled after hybrid public-private regulatory structures found in medicine, such as those put forth by the American Medical Association. Such governance schemes develop both technical and ethical standards, shaping professional training, continuing knowledge, and conduct. We contend that it is the most practical means of ensuring the development of human-centered AI in an era of rapid technological change and intensely opposing views of what regulation ought to do. This article places the responsibility of acting ethically on the group that knows the technology best and can anticipate its effects: AI developers. But unlike other voluntary standards, professional regulation articulates and enforces standards to certify individuals. Professional licensing is an alternative that provides public protections based on privately developed standards that ensure the safety of AI prior to their release.

Keywords: AI, public law, private regulation, certification, ethics, professional standards, professions, AI governance

https://doi.org/10.1017/S0892679426100343 Published online by Cambridge University Press