2-2023

Regulating Machine Learning: The Challenge of Heterogeneity

Cary Coglianese
University of Pennsylvania Carey Law School

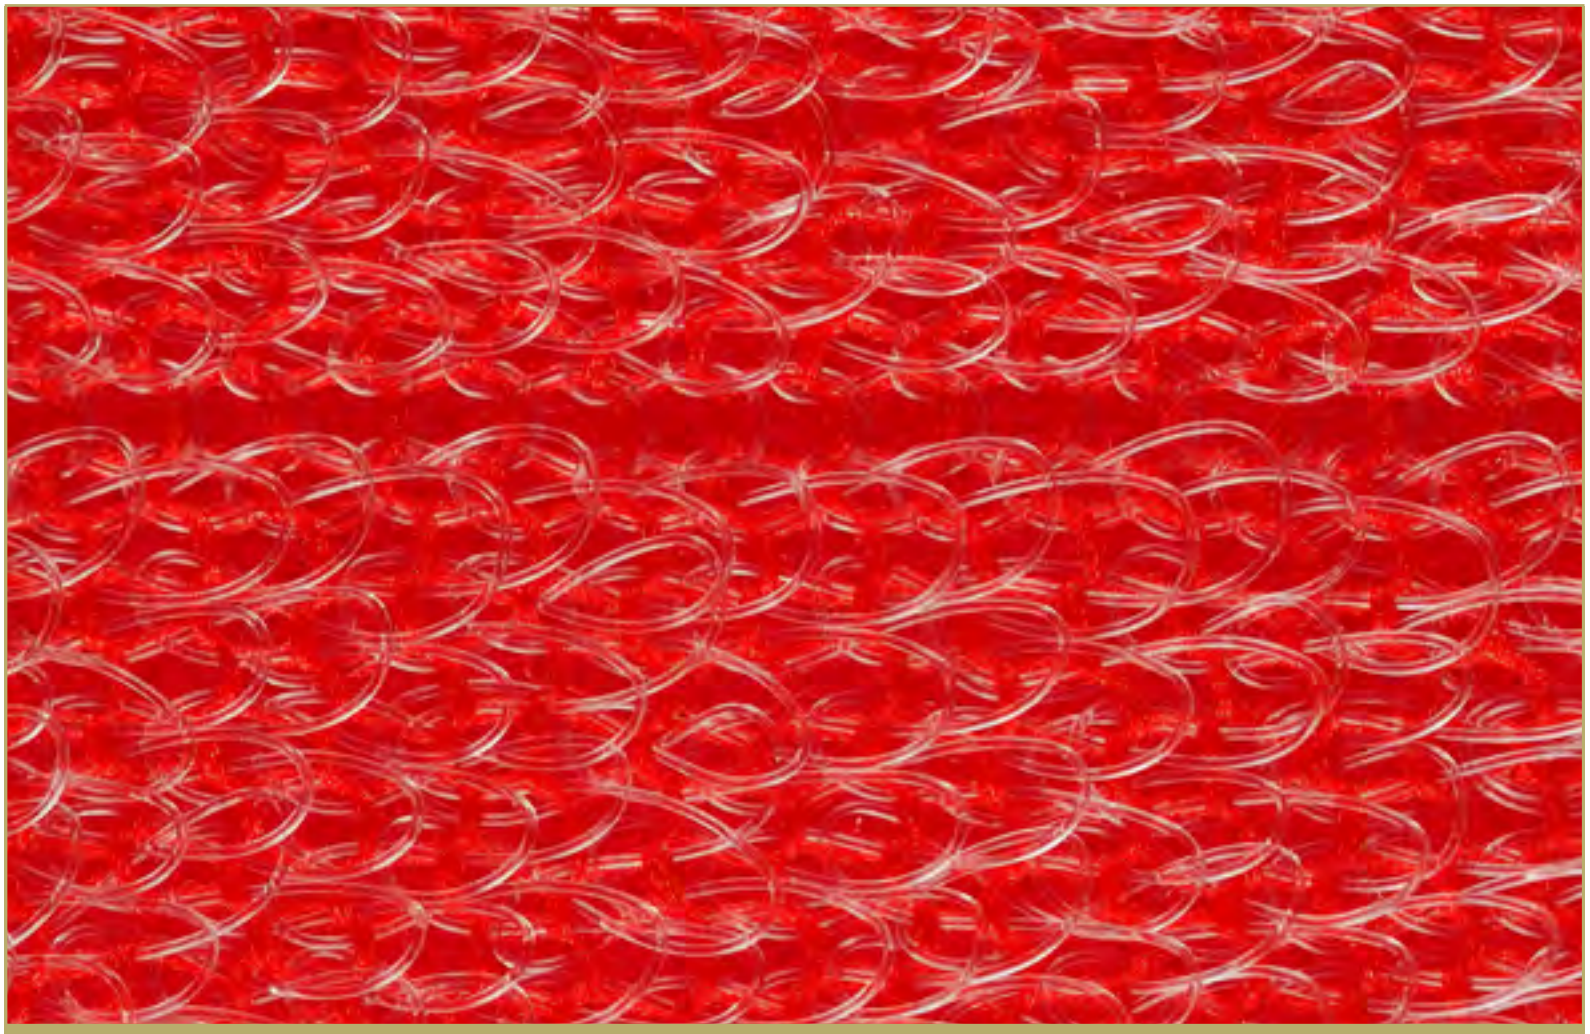

REGULATING MACHINE LEARNING: THE CHALLENGE OF HETEROGENEITY

BY

CARY COGLIANESE

Edward B. Shils Professor of Law and Professor of Political Science, and Director, Penn Program on Regulation, University of Pennsylvania.

Machine-learning algorithms increasingly drive technological advances that deliver valuable improvements for society and the economy. But these algorithms also raise important concerns. The way machine-learning algorithms work autonomously to find patterns in large datasets has given rise to fears of a world that will ultimately cede critical aspects of human control to the dictates of artificial intelligence. These fears seem only exacerbated by the intrinsic opacity surrounding how machine-learning algorithms achieve their results. To a greater degree than with other statistical tools, the outcomes generated by machine learning cannot be easily interpreted and explained, which can make it hard for the public to trust the fairness of products or processes powered by these algorithms.

For these reasons, the autonomous and opaque qualities of machine-learning algorithms make these digital tools both distinctive and a matter of public concern. But when it comes to *regulating* machine learning, a different quality of these algorithms matters most of all: their heterogeneity. The Merriam-Webster Dictionary defines "heterogeneity" as "the quality or state of consisting of dissimilar or diverse elements." Machine learning algorithms' heterogeneity will make all the difference in deciding when to regulate them, who should regulate them, and how to design regulations imposed on their development and use.

01
MACHINE LEARNING'S HETEROGENEITY

One of the most important sources of machine learning's heterogeneity derives from the highly diverse uses to which it is put. These uses could hardly vary more widely. Consider just a small sample of ways that different entities use machine-learning algorithms:

· Social media platforms use them to select and highlight content for users;
· Hospital radiology departments use them to detect cancer in patients;
· Credit card companies use them to identify potential fraudulent charges;
· Commercial airlines use them to operate aircraft with auto-piloting systems;
· Online retailers use them to make product recommendations to visitors to their websites; and
· Political campaigns use them in deciding where and how advertise.

Even within the same organizations, different machine-learning algorithms can perform different functions. An automobile manufacturer, for example, might use one type of machine-learning algorithm to automate certain on-road operations of their vehicles, while using other machine-learning algorithms as part of its manufacturing processes or for managing its supply chain and inventory.

In addition to their varied uses, machine-learning algorithms can themselves take many different forms and possess diverse qualities. These algorithms are often grouped into several main categories: supervised learning, unsupervised learning, semi-supervised learning, and reinforcement learning. Within each category, the range of algorithms and their forms can be highly diverse. Naïve Bayesian models, decision trees, random forests, and neural networks are just a few types of supervised learning models.[1] Even within any single type, finer points about how each model generated by an algorithm is structured, not to mention differences in the data used to train it, can lead each application of machine learning almost to fall within a category of its own.

Despite the wide variation in algorithms, it also remains that the same machine-learning model can be put to different uses within a single organization. For example, Meta—the corporation that owns Facebook and Instagram—has noted that, even though its "image classification models are all designed to predict what's in a given image, they may be used differently in an integrity system that flags harmful content versus a recommender system used to show people posts they might be interested in."[2]

Added to the extreme variation in uses and designs of algorithms is the fact that, for many uses, multiple different algorithms are used in combination with each other to support automated systems. What may at times be referred to as "an" algorithm is often actually a suite or family of algorithms, integrated into an automated system or process in a manner designed to perform a specified function. Furthermore, these algorithms and their combinations are updated and changed over time, as new or refined algorithms are shown to do better. Today's ChatGPT, for example, runs on models that are markedly different than earlier language models, and it will only be updated, enhanced, and modified repeatedly in the years to come.

Finally, these changes in machine-learning models come on top of the fact that when the data processed by a learning

1 Differences of expert opinion even exist over what counts as machine learning, with some data scientists treating forms of what others see as standard regression analysis as a type of machine learning.

2 MetaAI, *System Cards, A New Resource for Understanding How AI Systems Work* (Feb. 23, 2022), https://ai.facebook.com/blog/system-cards-a-new-resource-for-understanding-how-ai-systems-work/.

algorithm changes, then so too can its performance. This means that, for some algorithms, their performance can be constantly evolving as they encounter and process new data.[3]

In short, machine-learning algorithms place the definition of heterogeneity on steroids. These algorithms vary widely across different types and different uses at any given time — and they are highly dynamic, with their performance evolving over time. All this heterogeneity holds crucial implications for whether and how machine-learning algorithms should be regulated.

02
DECIDING TO REGULATE MACHINE LEARNING

The first question to ask, of course, is whether machine learning needs to be regulated at all.[4] Regulation is a tool designed to respond to and help solve social and economic problems. But by themselves, machine-learning algorithms are just mathematical constructs and create no social or economic problems.[5] If they were used only for intellectual pleasure—say, as a hobby pursued by a mathematically inclined subset of the population — then there would surely be no need to consider regulating them. Regulating ma-chine learning becomes a topic of conversation only when it is used in ways that have tangible effects on people.

If machine learning is to be a candidate for regulation, then, it is because of *the uses* for which it gets employed. This is not unlike other physical machines. When other machines have had consequential effects on the public, they have come to be regulated. The National Highway Traffic Safety Administration ("NHTSA"), for example, long ago starting imposing regulatory standards on different parts of an automobile not because of something intrinsic about the parts themselves, but rather because of how they are used in vehicles and how those uses affect the safety of the vehicle. Machine-learning algorithms are much the same. They are or will become objects of regulation because of the systems in which they are situated and how they ultimately affect system outcomes in ways that touch people's lives and livelihoods.

Because machine-learning algorithms can be used in so many different ways, this means that the regulatory problems they can create will vary quite widely as well. Looking across a host of different uses of machine learning, it is possible to say that the potential problems cover the gamut of classic market failures that justify regulation. Machine-learning algorithms used as part of automated pricing systems by online retailers, for example, may contribute to anti-competitive behavior in the marketplace.[6] Machine-learning algorithms used in medical treatments and consumer products can contribute to the kind of information asymmetries that typically justify consumer protection regulation.[7] And any pedestrian put at an increased risk from a self-driving car should easily be able to see another obvious market failure—an externality—created by vehicles that operate autonomously using sensors and machine-learning algorithms.

Regulation is often justified by more than just these classic market failures. It can also be used, for example, as a tool for preventing injustices and protecting civil rights, such as when regulations aim to combat employment discrimination.[8] Grounds exist for regulating machine learning on this basis as well. When society's prevailing biases have been reflected in the design of machine-learning algorithms or in the data on which they are trained, these algorithms can end up reinforcing, if not even exacerbating, existing in-

3 See, e.g., Jessa Boubker, *When Medical Devices Have a Mind of Their Own: The Challenges of Regulating Artificial Intelligence*, 47 Am. J.L. & Med. 427, 434 (2021) (indicating that, if an algorithm is continuously learning, it "will not always be able to predict how a software is going to react in real-time based on new data").

4 In posing the question in terms of whether to "regulate machine learning," I mean to distinguish it from the question of whether to impose antitrust regulation on the structural or other business decisions of firms that rely heavily on machine learning—namely, the so-called big tech firms. Deciding to impose regulatory scrutiny on mergers and acquisitions in the big tech space is not what I mean here by regulating machine learning. Only if machine-learning tools are themselves directly used to impede competition or concentrate market power would antitrust law become relevant for regulating machine learning in the sense I mean here.

5 This is putting to the side, of course, the fact that processing data using machine-learning algorithms can result in externalities from the production of energy needed to power the necessary computer hardware.

6 Cary Coglianese & Alicia Lai, *Antitrust by Algorithm,* Stan. Computational Antitrust, Vol. 2, no. 1, 2022, at 4.

7 *Cf. id.* at 18 (describing the difficulty in supporting algorithmic forecasts with intuitive explanations, which may run in some tension with consumer protection principles favoring disclosure and transparency).

8 See, e.g., Olatunde C.A. Johnson, Beyond the Private Attorney General: Equality Directives in American Law, 87 N.Y.U. L. Rev. 1339 (2012) (providing an overview of civil rights regulation in the United States).

justices.[9] Machine learning used as part of an employer's hiring process, for example, can thus create the problems that antidiscrimination regulation has been established to solve.[10]

Privacy is another civil rights concern that is often raised in the context of calls for regulation of machine learning. One worry centers on protecting the private information contained in the extensive data on which these algorithms draw — as well as ensuring individual notice of or consent to the use of such information. Still another concern arises from the ability of machine-learning algorithms to make accurate inferences about certain private characteristics that are not contained in the data themselves. Yet another concern centers on how machine-learning algorithms can make possible the use of facial recognition and other tools that can track individuals' whereabouts and contribute to fears of a "surveillance state."[11]

And then there are a host of other public policy concerns surrounding machine-learning algorithms that lie at the heart of many conversations about regulating artificial intelligence.[12] The availability of ChatGPT, for example, has raised new questions about what artificial intelligence means for education.[13] Social media platforms use machine-learning algorithms to push content to users in ways that accentuate conflict, keep users distracted, or make them crave more time on their smart phones.[14] Digital tools driven by machine-learning algorithms can also generate new artwork from existing works, raising questions about ownership rights and rules about appropriation.[15] These tools can be used perniciously too, such as by facilitating new opportunities for fraud through deep fakes.[16] Pernicious actors can also use artificial intelligence to propagate cyberattacks that threaten both digital and physical assets.[17]

As should be evident, the heterogeneous uses for machine-learning algorithms lead to a variety of regulatory concerns. It is surely axiomatic to observe that when the types of reg-ulatory problems vary, regulation itself must vary as well to fit the nature of the problem. At the very least, regulation must be designed in a way that accommodates variation in uses and either targets diverse problems or provides appro-priate incentives for regulated entities to find and address those problems.[18]

03

WHO SHOULD REGULATE MACHINE LEARNING?

Before turning to how regulation might be designed to accommodate machine learning's heterogeneity, a prior question arises about what type of institution should regulate machine learning, whenever that regulation is justified.

With respect to other technologies and their regulatory problems, the need for regulation to be adapted to fit different circumstances has led governments to establish different regulatory bodies, each targeting a circumscribed range

9 See, e.g., Dorothy Roberts, *Digitizing the Carceral State,* 132 Harv. L. Rev. 1695, 1698 (2019) (reviewing Virginia Eubanks, Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor (2018)); Sandra G. Mayson*, Bias in, Bias Out*, 128 Yale L.J. 2218 (2019).

10 Jeffrey Dastin, *Amazon Scraps Secret Ai Recruiting Tool That Showed Bias Against Women*, Reuters (Oct. 10, 2018, 7:04 PM), https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G.

11 A number of jurisdictions have prohibited law enforcement agencies from using facial recognition tools. *See* Cary Coglianese & Kat Hefter, *From Negative to Positive Algorithm Rights*, 30 Wm. & Mary Bill Rts J. 883, 886 n.15 (2022).

12 *Id*. at 886-893.

13 Kalley Huang, *Alarmed by A.I. Chatbots, Universities Start Revamping How They Teach*, N.Y. Times (Jan. 16, 2023), https://www.nytimes.com/2023/01/16/technology/chatgpt-artificial-intelligence-universities.html.

14 Barbara Ortutay & David Klepper, *Facebook Whistleblower Testifies: Five Highlights*, Assoc. Press (Oct. 5, 2021), https://apnews.com/article/facebook-frances-haugen-congress-testimony-af86188337d25b179153b973754b71a4. *See generally* Tim Wu, The Attention Merchants: The Epic Scramble to Get Inside Our Heads (2016).

15 Elizabeth Penava, *AI Art Is in Legal Greyscale*, Regul. Rev. (Jan. 24, 2023), https://www.theregreview.org/2023/01/24/penava-ai-art-is-in-legal-greyscale/.

16 Todd C. Helmus, RAND Corp., Artificial Intelligence, Deepfakes, and Disinformation: A Primer (2022).

17 Blessing Guembe, Ambrose Azeta, Sanjay Misra, Victor Chukwudi Osamor, Luis Fernandez-Sanz & Vera Pospelova, *The Emerging Threat of AI-Driven Cyber Attacks: A Review*, 36 Applied A.I. 1 (2022).

18 For a related discussion, see Cary Coglianese, *Regulating New Tech: Problems, Pathways, and People*, TechReg Chron., Dec. 2021, at 65-73.

of problems. The problems created by anticompetitive behavior, after all, are different than those created by industrial pollution, which are in turn different than the problems of unsafe and ineffective consumer products. As a result, antitrust regulatory institutions exist to target anticompetitive behavior; environmental regulatory bodies specialize in reducing pollution; and drug and consumer safety regulators aim to protect consumers from unsafe products. A single firm will need to comply with the regulations of several distinct regulators with respect to different facets of its operations and market behavior.

These different, specialized regulatory bodies have the advantage over a general legislature in that they can draw upon the specialized knowledge needed to address the different types of problems, their origins in different industries, and their effects on different subsets of the population. This is not to say that, even within their specializations, regulators do not confront heterogeneity. On the contrary, antitrust regulators are usually tasked with looking across all sectors of the economy for different ways businesses might engage in anticompetitive behavior. Environmental regulators are commonly tasked with regulating a variety of types of pollution, such as to the air, water, and land, and from a myriad of different businesses, large and small. Even regulatory bodies with relatively narrow targets — such as the U.S. Nuclear Regulatory Commission, which targets a single industry for the important but still circumscribed problem of nuclear safety[19] — will face some degree of heterogeneity in the different sources of risks and different scenarios that must be accounted for if regulation is to be effective. Nevertheless, because of the value of specialized expertise, nuclear regulators exist to look at nuclear safety and are not responsible for, say, ensuring the safety and soundness of banks. This is why, as a prescriptive matter, environmental regulators do not also seek to combat anticompetitive market conduct, and antitrust regulators are not responsible for addressing pollution problems.

It may be tempting to conclude that machine-learning algorithms are like nuclear power plants and that they need their own regulator. Recently, U.S. Representative Ted Lieu, for example, has argued that "[w]hat we need is a dedicated agency to regulate A.I."[20] Certainly, machine-learning algorithms do require specialized skills to understand how they work and how they can go awry. Regulating machine-learning algorithms' impact on any segment of society or the economy will require sophisticated knowledge about artificial intelligence. But because the regulatory problems that machine-learning algorithms are associated with can be so varied—and often so closely connected to long-standing regulatory problems that already have dedicated regulatory institutions—it is unrealistic to expect that any single regulator could ever sufficiently regulate all the problematic aspects of machine learning. Regulating algorithmic stock market trading will necessarily require great expertise about financial markets. A similar need for substantive expertise will apply when regulating the effects of machine-learning algorithms on the safety of medical devices, the operation of automobiles, and the pricing behavior of firms. No dedicated AI regulatory agency could possibly possess all of the additional related technical knowledge and capacity needed to regulate algorithms' many uses.

"It may be tempting to conclude that machine-learning algorithms are like nuclear power plants and that they need their own regulator

Given the many ways that machine-learning algorithms are intertwined with different problems, many of which are already addressed by existing regulatory bodies, it is not surprising that these existing regulators have so far taken the lead in responding to potential problems related to machine learning. Within the Department of Transportation, for example, NHTSA has issued regulatory guidance for automobile manufacturers on safety assessments for autonomous vehicle technology.[21] It ordered these manufacturers to file reports on crashes involving their autonomous vehicles.[22] NHTSA also recently prodded Tesla to recall more than 350,000 of its vehicles over safety concerns related to its driver assistance software.[23]

19 *About NRC*, U.S. Nuclear Regul. Comm'n, https://www.nrc.gov/about-nrc.html (last visited Feb. 4, 2023).

20 Ted Lieu, *I'm a Congressman Who Codes. A.I. Freaks Me Out*., N.Y. Times (Jan. 23, 2023), https://www.nytimes.com/2023/01/23/opinion/ted-lieu-ai-chatgpt-congress.html.

21 U.S. Dep't Transp. Nat'l Highway Traffic Safety Admin., *Federal Automated Vehicles Policy* (Sept. 2016), https://www.nhtsa.gov/sites/nhtsa.gov/files/documents/av_policy_guidance_pdf.pdf.

22 First Amended Standing General Order, U.S. Dep't Transp. Nat'l Highway Traffic Safety Admin., Incident Reporting for Automated Driving Systems (ADS) and Level 2 Advanced Driver Assistance Systems (ADAS), Order No. 2021-01 (August 2021), https://www.nhtsa.gov/sites/nhtsa.gov/files/2021-08/First_Amended_SGO_2021_01_Final.pdf.

23 Neal E. Boudette, *Tesla to Recall 362,000 Cars With Its "Full Self Driving" System*, N.Y. Times (Feb. 16, 2023), https://www.nytimes.com/2023/02/16/business/tesla-recall-full-self-driving.html.

Separately, the U.S. Food and Drug Administration (FDA) has developed an action plan for addressing the use of machine learning in medical devices, announcing it will treat them under a separate category for innovative devices.[24] In 2020, FDA approved the first AI-based cardiac ultrasound software under this alternative track.[25]

As existing regulatory bodies go forward to address AI-related problems within their domains, they will certainly need to develop further their data science expertise. It is not inconceivable that they could benefit from a centralized expert body that can provide guidance and support. Already, the National Institute of Standards and Technology (NIST) within the U.S. Department of Commerce has issued a generalized risk management framework for artificial intelligence that could be of value if customized to fit the needs of other more specialized regulatory settings.[26] NIST's framework joins other similar documents issued by other federal entities — such as the U.S. Government Accountability Office,[27] the White House Office of Science and Technology,[28] and the Administrative Conference of the United States[29] — that articulate general principles to follow when using machine-learning tools. The federal government has also established an AI Center of Excellence within the General Services Administration.[30]

Nevertheless, as helpful as these general, cross-cutting initiatives may be, existing regulators still need to build up their own capacity to understand and regulate AI tools, given how intertwined they can be with so many longstanding regulatory problems. Admittedly, even with sufficient capacity within existing agencies, some kinds of new problems will fall through the cracks. Ill effects from social media platforms' use of algorithms, for example, have so far have elided serious governmental oversight. Nevertheless, rather than hoping that a new omnibus AI regulatory body can swoop in to save the day by regulating all uses of machine learning, policymakers would do well to look instead to empower existing centers of regulatory expertise. Where gaps or overlaps exist in current regulatory authority, policymakers can then work to fill those gaps or work out any conflicting jurisdictions. Gaps could be filled either by creating new regulatory bodies focused on unattended problems or by assigning those new problems to existing regulators with relevant expertise.

04
HOW TO REGULATE MACHINE LEARNING

No matter which institutions take responsibility for regulating machine learning, they will still confront heterogeneity. Even within a specified industry and even with respect to some identical uses of machine learning, heterogeneity will remain because both the algorithms themselves and the data they use vary so widely. Moreover, the algorithms and the automated systems of which they are a part are changing over time. As a result, even within specialized domains, regulators will need to pursue measures that take into account the varied and dynamic nature of these algorithms.

For this reason, it is impossible to specify a tidy, one-size-fits-all formula for how regulators should approach their task of regulating machine learning. But at a broad brush, it is possible to say that regulators will need to approach their work with agility, flexibility, and vigilance.

24 U.S. Food & Drug Admin., *Artificial Intelligence and Machine Learning (AI/ML) Software as a Medical Device Action Plan* (Sept. 22, 2021), https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-and-machine-learning-aiml-enabled-medi-cal-devices; U.S. Food & Drug Admin., *Clinical Decision Support Software Guidance for Industry and Food and Drug Administration Staff* (Sept. 28, 2022), https://www.fda.gov/media/109618/download.

25 Press Release, U.S. Food & Drug Admin., *FDA Authorizes Marketing of First Cardiac Ultrasound Software That Uses Artificial Intelligence to Guide User* (Feb. 7, 2020), https://www.fda.gov/news-events/press-announcements/fda-authorizes-marketing-first-cardiac-ultrasound-software-uses-artificial-intelligence-guide-user.

26 Nat'l Inst. of Standards & Tech. (NIST), Artificial Intelligence Risk Management Framework (Jan. 2023), https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf.

27 U.S. Gov't Accountability Off., GAO-21-519SP, Artificial Intelligence: An Accountability Framework for Federal Agencies and Other Entities (June 2021), https://www.gao.gov/assets/gao-21-519sp.pdf.

28 White House Off. Of Sci. & Tech. Pol'y, Blueprint For An Ai Bill Of Rights: Making Automated Systems Work For The American People, https://www.whitehouse.gov/ostp/ai-bill-of-rights.

29 Admin. Conf. of the U.S., Administrative Conference Statement #20: Agency Use of Artificial Intelligence, 86 Fed. Reg. 6616, 6616 n.1 (Jan. 22, 2021).

30 Gen. Servs. Admin., Accelerate Adoption of Artificial Intelligence to Discover Insights at Machine Speed, https://coe.gsa.gov/docs/2020/AIServiceCatalog.pdf.

1. Regulate with agility. Regulators will need to be active and adaptive. Regulation of machine learning cannot be approached as a matter of finding the "right" rule and then moving on simply to enforcing that rule. Instead, regulators need to think of their work as incremental and constantly provisional. When the world that regulators seek to regulate keeps changing, the last thing regulators can do is remain static.

To regulate machine learning with agility, regulators need to build up their capacity to keep pace with changes in industry.[31] This requires building up a regulator's internal technological infrastructure and human capital with expertise in data sciences. It also means finding ways to engage with and gather information from industry.[32] Industry, after all, will be best-positioned to know the most about their algorithms and how they are used. Regulators cannot avoid active engagement with industry if they are to adopt smart approaches to regulation.

> ***No matter which institutions take responsibility for regulating machine learning, they will still confront heterogeneity.***

Of course, in seeking to engage with industry, regulators should never lose sight of their distinctive role as protectors of public value. To be sure, the public does gain from technological innovation in the private sector and regulation that unduly impedes innovation should be avoided. But regulators also should avoid embracing a perspective that values innovation for its own sake. They should not take their eyes off of the risks and other regulatory problems that innovations might bring.[33] Private firms will see some of these problems too, but if regulation is needed, that is because the firms lack the socially optimal incentives to ferret out and redress these problems, especially when the solutions are costly.

2. Deploy flexible rules. Machine learning's heterogeneity will make flexible rules strong candidates for adoption. A one-size-fits-all "prescriptive" or "specification" standard will not make sense, as that would necessitate the regulator telling firms exactly how to design, train, and use their algorithms. Regulators will almost surely never have sufficient capacity to regulate with such specificity.

An obvious alternative would be for the regulator to adopt performance standards that specify outcomes to be achieved (or avoided) but then give regulated firms the flexibility to decide how to proceed as long as they meet (or avoid) the outcome in the regulatory standard.[34] As appealing as performance standards may be, they necessitate that the regulator will be able to specify the desired outcome in a clear, monitorable fashion—and then have the capacity to do the actual monitoring.[35] Sometimes that might be the case, such as when machine learning is embedded in a larger system that can be observed independently and subjected to sufficient testing and monitoring. But in many cases it will be unlikely that regulators can develop sufficiently clear, monitorable performance tests for algorithms themselves.

When standard-setting organizations around the world have adopted voluntary performance guidelines for algorithms, they have tended to do so by articulating general performance *principles* calling for algorithms to yield outcomes that are "fair," "safe," "explainable," and so forth.[36] Although these principles-based approaches may be helpful in offering general guidance to industry, they are far from operational. It remains to be seen whether and how regulators could articulate with greater precision outcome values

31 Cary Coglianese, *Optimizing Regulation for an Optimizing Economy*, 4 U. Pa. J.L. & Pub. Affs. 1, 2 (2018).

32 Cary Coglianese, Richard Zeckhauser & Edward Parson, *Seeking Truth for Power: Informational Strategy and Regulatory Policy Making*, 89 Minn. L. Rev. 277, 278-79 (2004).

33 Cary Coglianese, *Regulatory Vigilance in a Changing World*, Regul. Rev. (Feb. 25, 2019), https://www.theregreview.org/2019/02/25/coglianese-innovation-regulatory-vigilance/.

34 Cary Coglianese, Jennifer Nash & Todd Olmstead, *Performance-Based Regulation: Prospects and Limitations in Health, Safety, and Environmental Regulation*, 55 Admin. L. Rev. 705 (2003).

35 Cary Coglianese, *The Limits of Performance-Based Regulation*, 50 U. Mich. J.L. Reform 525 (2017).

36 Gary E. Marchant, Lucille Tournas & Carlos Ignacio Gutierrez, *Governing Emerging Technologies Through Soft Law: Lessons for Artificial Intelligence*, 61 Jurimetrics J. 1, 5-6 (Fall 2020).

such as fairness and explainability.[37] Even with safety, one must surely ask: Exactly how safe is safe enough? Absent an ability to specify outcome values in measurable and monitorable terms, it is hard to see how regulators could rely on a performance-based approach to the regulation of machine learning.

In situations where neither a one-size-fits-all prescriptive rule nor a performance-based rule seem likely to work, regulators have turned to an alternative regulatory strategy called *management-based* regulation.[38] Under a management-based approach, the regulator requires the firm to engage in systemic managerial activities that seek to identify problems and then create internal responses to correct them. This approach has been widely applied to address other regulatory problems where heterogeneity dominates, such as food safety and chemical facility security. In these situations, the sources of the underlying regulatory problem are highly diverse and dynamic. The management-based approach typically calls for a regulated entity to develop a management plan, monitor for potential risks, produce internal procedures and trainings to address those risks, and maintain documentation on the operation of the firm's management system. Sometimes these regulations also require firms to subject their management systems to third-party auditing and certification.

In situations where neither a one-size-fits-all prescriptive rule nor a performance-based rule seem likely to work, regulators have turned to an alternative regulatory strategy called management-based regulation

Management-based regulation will be an obvious option to consider for machine learning. This regulatory option does not demand that the regulator have the same level of knowledge as regulated firms themselves, nor does it require that the regulator be able to specify and measure all the relevant outcomes. It also gives firms considerable flexibility and thereby accommodates heterogeneity across firms and over time.

Unsurprisingly, many emerging soft law standards for machine learning are taking a management-based approach. The voluntary framework that NIST recently issued to improve the trustworthiness of machine-learning applications, for example, bears all the hallmarks of a management-based approach. Specifically, it calls for firms to develop "structures, systems, processes, and teams" for "[a]nticipating, assessing, and otherwise addressing potential sources of negative risks" and to put in place "rigorous software testing and performance assessment methodologies," "[s]ystematic documentation practices," and "plans for prioritizing risk and regular monitoring and improvement."[39]

Although the NIST framework is not mandatory, similar approaches are starting to emerge in regulations or proposed regulations in various parts of the world. Canada, for example, has imposed a requirement that its own federal government agencies conduct algorithmic impact assessments, quality assurance auditing, and various documentation measures before launching algorithmic systems that substitute for human decision-makers.[40] A proposed European Union regulation would impose similar impact assessment and auditing requirements on both public and private sector machine-learning systems.[41] These auditing and impact assessment requirements are management-based. They do not impose any specific prescriptions for the design and use of algorithms nor what outcomes they achieve — but they do direct firms to undertake a series of risk management steps.

In other contexts, management-based regulations have sometimes required firms to disclose publicly their plans and audit results. Mandatory disclosure is another likely option for the future regulation of machine-learning algorithms. Already, big-tech firms are starting to develop their own semi-standardized means of disclosing information

37 For a discussion of principles-based regulation in other contexts, see Julia Black, *Forms and Paradoxes of Principles-Based Regulation*, 3 Cap. Markets L.J. 425 (2008); Cristie L. Ford, *New Governance, Compliance, and Principles-Based Securities Regulation*, 45 Am. Bus. L.J. 1 (2008). For treatment in the context of artificial intelligence, see Julia Black & Andrew Murray, *Regulating AI and Machine Learning: Setting the Regulatory Agenda*, 10 Eur. J. L. & Tech. 1 (2019).

38 Cary Coglianese & David Lazer, *Management-Based Regulation: Prescribing Private Management to Achieve Public Goals*, 37 L. & Soc. Rev. 691 (2003); Cary Coglianese, *Management-Based Regulation: Implications for Public Policy*, *in* Risk and Regulatory Policy: Improving the Governance of Risk (Gregory Bounds & Nikolai Malyshev, eds., 2010); Cary Coglianese & Shana Starobin, *Management-Based Regulation*, *in* Policy Instruments in Environmental Law 292 (Kenneth R. Richards & Josephine van Zeben, eds., 2020).

39 NIST, *supra* note 26.

40 Government of Canada, Directive on Automated Decision-Making (2021), https://www.tbs-sct.canada.ca/pol/doc-eng.aspx?id=32592.

41 European Commission, Proposal for a Regulation Laying Down Harmonised Rules on Artificial Intelligence (2021), https://digital-strategy.ec.europa.eu/en/library/proposal-regulation-laying-down-harmonised-rules-artificial-intelligence.

about their uses of machine learning as well as the basic properties of the algorithms and the data on which they are trained and deployed.[42] These voluntary disclosure efforts — what are currently known as "model cards" — could provide a template in the future for mandatory disclosure of information about machine-learning algorithms.[43] For the same reasons that performance-based standards are unlikely to prove viable as a regulatory strategy, it is unlikely that any disclosure regulation could demand a unified outcome metric to be applied to all algorithms and all use cases.[44] But any firm that has an internal management process supportive of the responsible use of artificial intelligence will necessarily generate some common types of information that could be disclosed.[45] The disclosure of information from firms' management of their algorithms would go some distance toward addressing concerns about machine learning's opacity as well as providing consumers and the public better assurance that firms are testing, validating, and deploying machine learning in a responsible manner.[46]

3. Remain vigilant. Research in other regulatory domains shows that management-based regulation can lead firms to reduce risks.[47] But as much as management-based regulation has been demonstrated to work in other contexts and is conceptually well-suited for regulating machine learning, it is hardly a panacea. The evidence for the long-term efficacy of this strategy remains less clear and worries exist that managerial rigor and steadfastness by firms can atrophy over time. The possibility exists that, even if firms subjected to AI impact assessment and auditing requirements take their required risk management responsibilities seriously at first, these management-based requirements can become rote paperwork exercises over time.[48] It is crucial that regulators build the capacity to assess the quality of firms' management efforts and that regulators sustain rigor in their oversight of their management-based regulatory regime.

Vigilance is also needed simply because of the rapid pace of change. Machine learning's future is a dynamic one and regulators need to equip themselves to make smart decisions in a changing environment. This means regulators must remain engaged with the industry they are overseeing and continue learning constantly. Regulators will make mistakes—they always have. But the key will be to try to minimize the consequences of those mistakes and, most of all, to learn from failures. Responsible regulation, like the responsible use of AI, requires vision, attentiveness, and the capacity to learn and adapt. If regulation of machine learning is to succeed, it must be viewed as an ongoing pursuit of continuous improvement.

05

REGULATING MACHINE LEARNING WITH MACHINE LEARNING?

A final aspect of the regulation of machine learning should not be overlooked: using machine learning to regulate machine

42 Vasi Philomin & Peter Hallinan, *Introducing AWS AI Service Cards: A New Resource to Enhance Transparency and Advance Responsible AI* (Nov. 30, 2022), https://aws.amazon.com/blogs/machine-learning/introducing-aws-ai-service-cards-a-new-resource-to-enhance-transparency-and-advance-responsible-ai/; *The Value of a Shared Understanding of AI Models*, Google Cloud, https://modelcards.withgoogle.com/about (last visited Feb. 16, 2023); Meta AI, System Cards, A New Resource for Understanding How AI Systems Work (Feb. 23, 2022), https://ai.facebook.com/blog/system-cards-a-new-resource-for-understanding-how-ai-systems-work/.

43 Margaret Mitchell et al., *Model Cards for Model Reporting* 221 (Jan. 14, 2019) (paper prepared for FAT* '19: Proceedings of the Conference on Fairness, Accountability, and Transparency), https://dl.acm.org/doi/10.1145/3287560.3287596 ("Model cards provide a way to inform users about what machine learning systems can and cannot do, the types of errors they make, and additional steps that could create more fair and inclusive outcomes with the technology.").

44 *See supra* notes 34-35 and accompanying text. Model cards, on the other hand, "are designed to be flexible in both scope and specificity in order to accommodate the wide variety of machine learning model types and potential use cases." *Id*. at 228.

45 *Cf*. *Service Cards and ML Governance with Michael Kearns* (January 2, 2023), https://twimlai.com/podcast/twimlai/service-cards-and-ml-governance/ (discussing the quantitative technical assessments and extensive internal reviews that underlie AWS service cards and noting that "a lot of work went into these cards").

46 *Cf*. Cary Coglianese & David Lehr, *Transparency and Algorithmic Governance*, 71 Admin. L. Rev. 1, 49-55 (2019) (discussing emerging technical advances that can enhance machine learning's transparency).

47 See Lori S. Bennear, *Are Management-based Regulations Effective? Evidence from State Pollution Prevention Programs,* 26 J. Pol'y Analysis & Mgmt. 327 (2007); Travis Minor & Matt Parrett, *The Economic Impact of the Food and Drug Administration's Final Juice HACCP Rule*, 68 Food Pol'y 206 (2017).

48 See, e.g., Cary Coglianese & Jennifer Nash, *Compliance Management Systems: Do They Make a Difference?*, *in* Cambridge Handbook of Compliance 571 (D. Daniel Sokol & Benjamin van Rooij, eds., 2021); Garry C. Gray & Susan S. Silbey, *Governing Inside the Organization: Interpreting Regulation and Compliance*, 96 Amer. J. Soc. 120 (2014).

learning. Algorithms, after all, are not merely tools for private sector firms seeking to innovate and enhance value. Regulators can also look to machine-learning algorithms as tools for improving their own performance.[49] At present, some regulators use them to identify firms that are likely in violation of applicable rules. Rather than sending out auditors or inspectors at random, and thereby using limited oversight resources to monitor firms that will be in compliance, regulators can vastly improve the detection of violators by using machine learning to decide how to target their limited resources.[50]

This same approach could be used by regulators when al-locating limited resources to oversee firms' compliance with machine-learning regulation. With so many different uses for machine learning, and the prospect of vast numbers of firms using this digital technology, regulators will have to be smart about how to allocate their oversight resources. This may include using natural language processing algorithms to identify firms with inadequate risk management plans. It may include using algorithms to select firms for regulatory audit-ing that are most likely to be treating required management-based planning in a pro forma fashion. The kind of vigilance that regulators will need to maintain will require that regulators themselves use the most sophisticated tools in their arsenals.

This same approach could be used by regulators when allocating limited resources to oversee firms' compliance with machine-learning regulation

The time may also come when regulators develop automated regulatory tools that match the speed and heterogeneity of private sector machine learning with the speed and heterogeneity of regulatory machine learning. When businesses rely on machine-learning tools to make subtle but anticompetitive pricing decisions in real time, for example, antitrust regulators might do well to use machine-learning tools to detect these collusive pricing patterns at the same speed.[51] When high-speed algorithms facilitate ever-so-slight but profitable forms of stock market manipulation, securities regulators would likely do well to use similarly sophisticated algorithms to discover that manipulation.[52] Over time, regulators' own algorithms might even be used as part of larger automated systems that can detect and algorithmically punish at the same time.

Perhaps the idea of regulatory robots seems a bit fanciful, but it is already becoming a reality, even if in seemingly banal ways. Automated regulatory systems already are already being used in one of the most familiar venues of daily life: the roadway. Several cities around the United States have installed automated rule-makers and rule-enforcers on their streets and highways to optimize traffic flow.[53] These digital traffic light systems rely on sensors and machine-learning algorithms to determine when signals turn red and green. Other jurisdictions have installed automated systems on highways that can detect vehicles traveling at excessive speeds and then send tickets to the vehicles' owners.[54]

It is not hard to imagine a future in which machine-learning systems that operate self-driving cars are integrated into automated systems of traffic control and management, making the regulation of the nation's roadways run entirely on machine learning. Nor is it difficult to envision a world in which many other activities and business practices are reg-ulated by automated systems driven by machine-learning algorithms.[55]

Admittedly, the regulatory tasks involved in detecting vehicle speed and changing traffic lights may seem simple compared with the tasks regulators face in overseeing all the myriad uses of machine learning. And technology will not erase the regulatory challenges created by machine learning's heterogeneity. But the existence of even crude automated regulatory systems today on the nation's roadways offers a vision of a future in which at least some private sector uses of machine-learning algorithms will be overseen by

49 Cary Coglianese & David Lehr, *Regulating by Robot: Administrative Decision-Making in the Machine Learning Era*, 105 Geo. L. J. 1147 (2017); Cary Coglianese, *Algorithmic Regulation: Machine Learning as Governance Tool*, *in* The Algorithmic Society: Power, Knowledge and Technology in the Age of Algorithms 35 (Marc Schuilenburg & Rik Peeters, eds., 2021).

50 Cary Coglianese & Alicia Lai, *Algorithm vs. Algorithm*, 72 Duke L.J. 1281, 1311 (2021).

51 Coglianese & Lai, *supra* note 6.

52 Coglianese & Lehr, *supra* note 49.

53 Cary Coglianese & Lavi M. Ben Dor, *AI in Adjudication and Administration*, 86 Brook. L. Rev. 791, 824-25 (2021).

54 Coglianese & Hefter, *supra* note 11.

55 Coglianese & Lehr, *supra* note 49; Cary Coglianese & Alicia Lai, *Assessing Automated Administration*, *in* Oxford Handbook of AI Governance (Justin Bullock et al., eds., forthcoming), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4084844.

regulatory systems driven themselves by machine-learning algorithms.[56]

06
MEETING THE CHALLENGE OF HETEROGENEITY

Regulating machine-learning algorithms sensibly will not be easy. Their complexity, self-learning autonomy, and opacity create reasons for, as well as challenges to, sound regulation. But it is machine learning's *heterogeneity* that poses regulators' greatest challenge of all. These algorithms' varied forms, multiple uses, and dynamic properties make most conventional regulatory strategies obsolete. The tradition of a regulatory body that establishes and then enforces rigid, general commands will not fit well in a world of rapidly evolving, highly varied digital tools.

Regulating machine learning well must draw upon the expertise of multiple regulatory institutions that can target machine learning's multiple uses. These specialized regulators will need to deploy flexible regulatory instruments, such as management-based regulation, and use smart oversight strategies, such as by using algorithmic tools for prioritizing resources.

In the end, effective governance in a world driven by heterogeneous algorithmic machines will depend on sophisticated decision-making and top-level performance by human institutions tasked with regulatory oversight. Regulating machine learning well will demand the utmost levels of vigilance and excellence by regulatory officials as they practice their craft.[57]

Regulating machine learning well must draw upon the expertise of multiple regulatory institutions that can target machine learning's multiple uses

56 Cary Coglianese, *Moving Toward Personalized Law*, U. Chi. L. Rev. Online (2022), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4051776.

57 Achieving Regulatory Excellence (Cary Coglianese, ed., 2017); Malcolm K. Sparrow, The Regulatory Craft: Controlling Risks, Solving Problems & Managing Compliance (2000).

CPI SUBSCRIPTIONS

CPI reaches more than **35,000 readers** in over **150 countries** every day. Our online library houses over **23,000 papers**, articles and interviews.

Visit **competitionpolicyinternational.com** today to see our available plans and join CPI's global community of antitrust experts.

CPI COMPETITION POLICY® INTERNATIONAL