Response of the International Center for Law & Economics to the AI Accountability Policy Request for Comment

National Telecommunications and Information Administration

Docket No. 230407-0093

June 12, 2023

Authored by:

Kristian Stout (Director of Innovation Policy, International Center for Law & Economics)

Brian Albrecht (Chief Economist, International Center for Law & Economics)

Mikolaj Barczentewicz (Senior Scholar, International Center for Law & Economics)

Eric Fruits (Senior Scholar, International Center for Law & Economics)

Geoffrey A. Manne (President Founder, International Center for Law & Economics)

Julian Morris (Senior Scholar, International Center for Law & Economics)

I. Introduction: How Do You Solve a Problem Like 'AI'?

On behalf of the International Center for Law & Economics (ICLE), we thank the National Telecommunications and Information Administration (NTIA) for the opportunity to respond to this *AI Accountability Policy Request for Comment* (RFC).

A significant challenge that emerges in discussions concerning accountability and regulation for artificial intelligence is the broad and often ambiguous definition of "AI" itself. This is demonstrated in the RFC's framing:

> This Request for Comment uses the terms AI, algorithmic, and automated decision systems without specifying any particular technical tool or process. It incorporates NIST's definition of an ''AI system,'' as ''an engineered or machine-based system that can, for a given set of objectives, generate outputs such as predictions, recommendations, or decisions influencing real or virtual environments.'' This Request's scope and use of the term ''AI'' also encompasses the broader set of technologies covered by the Blueprint: ''automated systems'' with ''the potential to meaningfully impact the American public's rights, opportunities, or access to critical resources or services.''[1]

As stated, the RFC's scope could be read to cover virtually all software.[2] But it is essential to acknowledge that, for the purposes of considering potential regulation, we lack a definition of AI that is either sufficiently broad as to cover all or even most areas of concern, and sufficiently focused as to be a useful lens for analysis. That is to say, what we think of as AI encompasses a significant diversity of discrete technologies that will be put to a huge number of potential uses.

One useful recent comparison is with the approach the Obama administration took in its deliberations over nanotechnology regulation in 2011.[3] Following years of consultation and debate, the administration opted for a parsimonious, context-specific approach precisely because

---

[1] *AI Accountability Policy Request for Comment*, Docket No. 230407-0093, 88 FR 22433, NATIONAL TELECOMMUNICATIONS AND INFORMATION ADMINISTRATION (Apr. 14, 2023) ("RFC").

[2] Indeed, this approach appears to be the default position of many policymakers around the world. *See, e.g.*, Mikolaj Barczentewicz, *EU's Compromise AI Legislation Remains Fundamentally Flawed*, TRUTH ON THE MARKET (Feb. 8, 2022), https://truthonthemarket.com/2022/02/08/eus-compromise-ai-legislation-remains-fundamentally-flawed; The fundamental flaw of this approach is that, while AI techniques use statistics, "statistics also includes areas of study which are not concerned with creating algorithms that can learn from data to make predictions or decisions. While many core concepts in machine learning have their roots in data science and statistics, some of its advanced analytical capabilities do not naturally overlap with these disciplines." *See*, *Explainable AI: The Basics*, THE ROYAL SOCIETY (2019) at 7 *available at* https://royalsociety.org/-/media/policy/projects/explainable-ai/AI-and-interpretability-policy-briefing.pdf ("Royal Society Briefing").

[3] John P. Holdren, Cass R. Sunstein, & Islam A. Siddiqui, *Memorandum for the Heads of Executive Departments and Agencies*, EXECUTIVE OFFICE OF THE WHITE HOUSE (Jun. 9, 2011), *available at* https://obamawhitehouse.archives.gov/sites/default/files/omb/inforeg/for-agencies/nanotechnology-regulation-and-oversight-principles.pdf.

"nanotechnology" is not really a single technology. In that proceeding, the administration ultimately recognized that it was not the general category of "nanotechnology" that was relevant, nor the fact that nanotechnologies are those that operate at very small scales, but rather the means by and degree to which certain tools grouped under the broad heading of "nanotechnology" could "alter the risks and benefits of a specific application."[4] This calls to mind Judge Frank Easterbrook's famous admonition that a "law of cyberspace" would be no more useful than a dedicated "law of the horse."[5] Indeed, we believe Easterbrook's observation applies equally to the creation of a circumscribed "law of AI."

While there is nothing inherently wrong with creating a broad regulatory framework to address a collection of loosely related subjects, there is a danger that the very breadth of such a framework might over time serve to foreclose more fruitful and well-fitted forms of regulation.

A second concern in the matter immediately at hand is, as mentioned above, the potential for AI regulation to be formulated so broadly as to encompass essentially all software. Whether by design or accident, this latter case runs a number of risks. First, since the scope of the regulation will potentially cover a much broader subject, the narrow discussion of "AI" will miss many important aspects of broader software regulation, and will, as a consequence, create an ill-fitted legal regime. Second, by sweeping in a far wider range of tools into such a regulation than the drafters publicly acknowledge, the democratic legitimacy of the process is undermined.

A. The Danger of Regulatory Overaggregation

The current hype surrounding AI has been driven by popular excitement, as well as incentives for media to capitalize on that excitement. While this is understandable, it arguably has led to oversimplification in public discussions about the underlying technologies. In reality, AI is an umbrella term that encompasses a diverse range of technologies, each with its own unique characteristics and applications.

For instance, relatively lower-level technologies like large language models (LLMs)[6] differ significantly from diffusion techniques.[7] At the level of applications, recommender systems can

---

[4] *Id.*

[5] Frank H. Easterbrook, *Cyberspace and the Law of the Horse*, 1996 U. CHI. L. FORUM 207 (1996).

[6] LLMs are a type of artificial-intelligence model designed to parse and generate human language at a highly sophisticated level. The deployment of LLMs has driven progress in fields such as conversational AI, automated content creation, and improved language understanding across a multitude of applications, even suggesting that these models might represent an initial step toward the achievement of artificial general intelligence (AGI). *See* Alejandro Peña *et al.*, *Leveraging Large Language Models for Topic Classification in the Domain of Public Affairs*, ARXIV (Jun. 5, 2023), https://arxiv.org/abs/2306.02864v1.

[7] Diffusion models are a type of generative AI built from a hierarchy of denoising autoencoders, which can achieve state-of-the-art results in such tasks as class-conditional image synthesis, super-resolution, inpainting, colorization, and stroke-based synthesis. Unlike other generative models, these likelihood-based models do not exhibit mode collapse and training instabilities. By leveraging parameter sharing, they can model extraordinarily complex distributions of natural images without necessitating billions of parameters, as in autoregressive models. *See* Robin Rombach *et al.*, High-Resolution *Image Synthesis with Latent Diffusion Models*, ARXIV (Dec. 20, 2021), https://arxiv.org/abs/2112.10752.

employ a wide variety of different machine-learning (or even more basic statistical) techniques.[8] All of these techniques collectively called "AI" also differ from the wide variety of algorithms employed by search engines, social media, consumer software, video games, streaming services, and so forth, although each also contains software "smarts," so to speak, that could theoretically be grouped under the large umbrella of "AI."

And none of the foregoing bear much resemblance at all to what the popular imagination conjures when we speak of AI—that is, artificial general intelligence (AGI), which some experts argue may not even be achievable.[9]

Attempting to create a single AI regulatory scheme commits what we refer to as "regulatory overaggregation"—sweeping together a disparate set of more-or-less related potential regulatory subjects under a single category in a manner that overfocuses on the abstract term and obscures differences among the subjects. The domains of "privacy rights" and "privacy regulation" are illustrative of the dangers inherent in this approach. There are, indeed, many potential harms (both online and offline) that implicate the concept of "privacy," but the differences among these recommend examining closely the various contexts that attend each.

Individuals often invoke their expectation of "privacy," for example, in contexts where they want to avoid the public revelation of personal or financial information. This sometimes manifests as the assertion of a right to control data as a form of quasi-property, or as a form of a right to anti-publicity (that is, a right not to be embarrassed publicly). Indeed, writing in 1890 with his law partner Samuel D. Warren, future Supreme Court Justice Louis Brandeis posited a "right to privacy" as akin to a property right.[10] Warren & Brandeis argued that privacy is not merely a matter of seclusion, but extends to the individual's control over their personal information.[11] This "right to be let alone" delineates a boundary against unwarranted intrusion, which can be seen as a form of intangible property right.[12]

This framing can be useful as an abstract description of a broad class of interests and concerns, but it fails to offer sufficient specificity to describe actionable areas of law. Brandeis & Warren were

---

[8] Recommender systems are advanced tools currently used across a wide array of applications, including web services, books, e-learning, tourism, movies, music, e-commerce, news, and television programs, where they provide personalized recommendations to users. Despite recent advancements, there is a pressing need for further improvements and research in order to offer more efficient recommendations that can be applied across a broader range of applications. *See* Deepjyoti Roy & Mala Dutta, *A Systematic Review and Research Perspective on Recommender Systems*, 9 J. BIG DATA 59 (2022), *available at* https://journalofbigdata.springeropen.com/counter/pdf/10.1186/s40537-022-00592-5.pdf.

[9] AGI refers to hypothetical future AI systems that possess the ability to understand or learn any intellectual task that a human being can do. While the realization of AGI remains uncertain, it is distinct from the more specialized AI systems currently in use. For a skeptical take on the possibility of AGI, *see* ROGER PENROSE, THE EMPEROR'S NEW MIND (Oxford Univ. Press 1989).

[10] Samuel D. Warren & Louis D. Brandeis, *The Right to Privacy*, 4 HARV. L. REV. 193 (1890).

[11] *Id.* at 200.

[12] *Id.* at 193.

concerned primarily with publicity;[13] that is, with a property right to control one's public identity as a public figure. This, in turn, implicates a wide range of concerns, from an individual's interest in commercialization of their public image to their options for mitigating defamation, as well as technologies that range from photography to website logging to GPS positioning.

But there are clearly other significant public concerns that fall broadly under the heading of "privacy" that cannot be adequately captured by the notion of controlling a property right "to be let alone." Consider, for example, the emerging issue of "revenge porn." It is certainly a privacy harm in the Brandeisian sense that it implicates the property right not to have one's private images distributed without consent. But that framing fails to capture the full extent of potential harms, such as emotional distress and reputational damage.[14] Similarly, cases in which an individual's cellphone location data are sold to bounty hunters are not primarily about whether a property right has been violated, as they raise broader issues concerning potential abuses of power, stalking, and even physical safety.[15]

These examples highlight some of the ways that, in failing to take account of the distinct facts and contexts that can attend privacy harms, an overaggregated "law of privacy" may tend to produce regulations insufficiently tailored to address those diverse harms.

By contrast, the domain of intellectual property (IP) may serve as an instructive counterpoint to the overaggregated nature of privacy regulation. IP encompasses a vast array of distinct legal constructs, including copyright, patents, trade secrets, trademarks, and moral rights, among others. But in the United States—and indeed, in most jurisdictions around the world—there is no overarching "law of intellectual property" that gathers all of these distinct concerns under a singular regulatory umbrella. Instead, legislation is specific to each area, resulting in copyright-specific acts, patent-specific acts, and so forth. This approach acknowledges that, within IP law, each IP construct invokes unique rights, harms, and remedies that warrant a tailored legislative focus.

The similarity of some of these areas does lend itself to conceptual borrowing, which has tended to enrich the legislative landscape. For example, U.S. copyright law has imported doctrines from patent law.[16] Despite such cross-pollination, copyright law and patent law remain distinct. In this way, intellectual property demonstrates the advantages of focusing on specific harms and remedies. This could serve as a valuable model for AI, where the harms and remedies are equally diverse and context dependent.

---

[13] *Id.* at 196-97.

[14] Notably, courts do try to place a value on emotional distress and related harms. But because these sorts of violations are deeply personal, attempts to quantify such harms in monetary terms are rarely satisfactory to the parties involved.

[15] Martin Giles, *Bounty Hunters Tracked People Secretly Using US Phone Giants' Location Data*, MIT TECH. REV. (Feb. 7, 2019), https://www.technologyreview.com/2019/02/07/137550/bounty-hunters-tracked-people-secretly-using-us-phone-giants-location-data.

[16] *See, e.g.*, *Sony Corp. of Am. v. Universal City Studios, Inc.*, 464 U.S. 417, 439 (1984) (The Supreme Court imported the doctrine of "substantial noninfringing uses" into copyright law from patent law).

If AI regulations are too broad, they may inadvertently encompass any algorithm used in commercially available software, effectively stifling innovation and hindering technological advancements. This is no less true of good-faith efforts to craft laws in any number of domains that nonetheless suffer from a host of unintended consequences.[17]

At the same time, for a regulatory regime covering such a broad array of varying technologies to be intelligible, it is likely inevitable that tradeoffs made to achieve administrative efficiency will cause at least some real harms to be missed. Indeed, NTIA acknowledges this in the RFC:

> Commentators have raised concerns about the validity of certain accountability measures. Some audits and assessments, for example, may be scoped too narrowly, creating a "false sense" of assurance. Given this risk, it is imperative that those performing AI accountability tasks are sufficiently qualified to provide credible evidence that systems are trustworthy.[18]

To avoid these unintended consequences, it is crucial to develop a more precise understanding of AI and its various subdomains, and to focus any regulatory efforts toward addressing specific harms that would not otherwise be captured by existing laws. The RFC declares that its aim is "to provide assurance—that AI systems are legal, effective, ethical, safe, and otherwise trustworthy."[19] As we discuss below, rather than promulgate a set of recommendations about the use of AI, NTIA should focus on cataloguing AI technologies and creating useful taxonomies that regulators and courts can use when they identify tangible harms.

II. AI Accountability and Cost-Benefit Analysis

The RFC states that:

> The most useful audits and assessments of these systems, therefore, should extend beyond the technical to broader questions about governance and purpose. These might include whether the people affected by AI systems are meaningfully consulted in their design and whether the choice to use the technology in the first place was well-considered.[20]

It is unlikely that consulting all of the people potentially affected by a set of technological tools could fruitfully contribute to the design of any regulatory system other than one that simply bans those tools.[21] Any intelligible accountability framework must be dedicated to evaluating the technology's

---

[17] A notable example is how the Patriot Act, written to combat terrorism, was ultimately used to take down a sitting governor in a prostitution scandal. *See* Noam Biale, *Eliot Spitzer: From Steamroller to Steamrolled*, ACLU, Oct. 29, 2007, https://www.aclu.org/news/national-security/eliot-spitzer-steamroller-steamrolled.

[18] RFC at 22437.

[19] *Id.* at 22433.

[20] *Id.* at 22436.

[21] Indeed, the RFC acknowledges that, even as some groups are developing techniques to evaluate AI systems for bias or disparate impact, "It should be recognized that for some features of trustworthy AI, consensus standards may be difficult or

real-world impacts, rather than positing thought experiments about speculative harms. Where tangible harms can be identified, such evaluations should encompass existing laws that focus on those harms and how various AI technologies might alter how existing law would apply. Only in cases where the impact of particular AI technologies represents a new kind of harm, or raises concerns that fall outside existing legal regimes, should new regulatory controls be contemplated.

AI technologies will have diverse applications and consequences, with the potential for both beneficial and harmful outcomes. Rather than focus on how to constrain either AI developers or the technology itself, the focus should be on how best to mitigate or eliminate any potential negative consequences to individuals or society.

NTIA asks:

> AI accountability measures have been proposed in connection with many different goals, including those listed below. To what extent are there tradeoffs among these goals?[22]

This question acknowledges that, fundamentally, AI accountability comes down to cost-benefit analysis. In conducting such analysis, we urge that the NTIA and any other agencies be sure to account not only for potential harms, but to take very seriously the massive benefits these technologies might provide.

A. The Law Should Identify and Address Tangible Harms, Incorporating Incremental Changes

To illustrate the challenges inherent to tailoring regulation of a new technology like AI to address the ways that it might generally create harm, it could be useful to analogize to a different existing technology: photography. If camera technology were brand new, we might imagine a vast array of harms that could arise from its use. But it should be obvious that creating an overarching accountability framework for all camera technology is absurd. Instead, laws of general applicability should address harmful uses of cameras, such as the invasion of privacy rights posed by surreptitious filming. Even where a camera is used in the commission of a crime—*e.g.*, surveilling a location in preparation to commit a burglary—it is not typically the technology itself that is the subject of legal concern; rather, it is the acts of surveillance and burglary.

Even where we can identify a tangible harm that a new technology facilitates, the analysis is not complete. Instead, we need to balance the likelihood of harmful uses of that technology with the likelihood of nonharmful (or beneficial) uses of that technology. Copyright law provides an apt example.

Sony,[23] often referred to as the "Betamax case," was a landmark U.S. Supreme Court case in 1984 that centered on Sony's Betamax VCR—the first consumer device that could record television shows

---

impossible to create." RFC at 22437. Arguably, this problem is inherent to constructing an overaggregated regulator, particularly one that will be asked to consulting a broad public on standards and rulemaking.

[22] *Id.* at 22439.

[23] *Sony Corp. of Am. v. Universal City Studios, Inc.*, 464 U.S. at 417.

for later viewing, a concept now referred to as time-shifting.[24] Plaintiffs alleged that, by manufacturing and selling the Betamax VCRs, Sony was secondarily liable for copyright infringement carried out by its customers when they recorded television shows.[25] In a 5-4 decision, the Supreme Court ruled in favor of Sony, holding that the use of the Betamax VCR to record television shows for later personal viewing constituted "fair use" under U.S. copyright law.[26]

Critical for our purposes here was that the Court found that Sony could not be held liable for contributory infringement because the Betamax VCR was capable of "substantial noninfringing uses."[27] This is to say that, faced with a new technology (recording relatively high-quality copies of television shows and movies at home), the Court recognized that, while the Betamax might facilitate some infringement, it would be inappropriate to apply a presumption against its use.

Sony and related holdings did not declare that using VCRs to infringe copyright was acceptable. Indeed, copyright enforcement for illegal reproduction has continued apace, even when using new technologies capable of noninfringing uses.[28] At the same time, the government did not create a new regulatory and licensing regime to govern the technology, despite the fact that it was a known vector for some illicit activity.

Note, the *Sony* case is also important for its fair-use analysis, and is widely cited for the proposition that so-called "time shifting" is permissible. That is not central to our point here, particularly as there is no analogue to fair use proposed in the AI context. But even here, it represents how the law adapts to develop doctrines that excuse conduct that would otherwise be a violation. In the case of copyright, unauthorized reproduction is infringement, period.[29] Fair use is raised as an affirmative defense[30] to excuse some unauthorized reproduction because courts have long recognized that, when viewed case-by-case, application of legal rules need to be tailored to make room for unexpected fact patterns where acts that would otherwise be considered violations yield some larger social benefit.

We are not suggesting the development of a fair-use doctrine for AI, but are instead insisting that AI accountability and regulation must be consistent with the case-by-case approach that has characterized the common law for centuries. Toward that end, it would be best for law relevant to AI to emerge through that same bottom-up, case-by-case process. To the extent that any new legislation is passed, it should be incremental and principles-based, thereby permitting the

---

[24] *Id.*

[25] *Id.*

[26] *Id.* at 456.

[27] *Id.*

[28] *See, e.g., Defendant Indicted for Camcording Films in Movie Theaters and for Distributing the Films on Computer Networks First Prosecution Under Newly-Enacted Family Entertainment Copyright Act*, U.S. DEPT OF JUSTICE (Aug. 4, 2005), *available at* https://www.justice.gov/archive/criminal/cybercrime/press-releases/2005/salisburyCharge.htm.

[29] 17 U.S.C. 106.

[30] *See* 17 U.S.C. 107; *Campbell v. Acuff-Rose Music, Inc.*, 510 U.S. 569, 590 (1994) ("Since fair use is an affirmative defense, its proponent would have difficulty carrying the burden of demonstrating fair use without favorable evidence about relevant markets.").

emergence of law that best fits particular circumstances and does not conflict with other principles of common law.

By contrast, there *are* instances where the law has recognized that certain technologies are more likely to be used for criminal purposes and should be strictly regulated. For example, many jurisdictions have made possession of certain kinds of weapons—e.g., nunchaku, shuriken "throwing stars," and switchblade knives—*per se* illegal, despite possible legal uses (such as martial-arts training).[31] Similarly, although there is a strong Second Amendment protection for firearms in the United States, it is illegal for a felon to possess a firearm.[32] The reason these prohibitions developed is because it was deemed that possession of these devices in most contexts had no other possible use than the violation of the law. But these sorts of technologies are the exception, not the rule. Many chemicals that can be easily used as poisons are nonetheless available as, *e.g.*, cleaning agents or fertilizers.

1. The EU AI Act: An overly broad attempt to regulate AI

Nonetheless, some advocate regulating AI by placing new technologies into various broad categories of risk, each with their own attendant rules. For example, as proposed by the European Commission, the EU's AI Act would regulate the use of AI systems that ostensibly pose risks to health, safety, and fundamental rights.[33] The proposal defines AI systems broadly to include essentially any software, and sorts them into three risk levels: unacceptable, high, and limited risk.[34] Unacceptable-risk systems are prohibited outright, while high-risk systems are subject to strict requirements, including mandatory conformity assessments.[35] Limited-risk systems face certain requirements related to adequate documentation and transparency.[36]

The AI Act defines AI so broadly that it would apply even to ordinary general-purpose software, as well as software that uses machine learning but does not pose significant risks.[37] The plain terms of the AI Act could be read to encompass common office applications, spam filters, and recommendation engines, thus potentially imposing considerable compliance burdens on businesses for their use of software that provides benefits dramatically greater than any expected costs.[38] A recently proposed amendment would "ban the use of facial recognition in public spaces, predictive

---

[31] *See, e.g.*, N.Y. Penal Law § 265.01; Wash. Rev. Code Ann. § 9.41.250; Mass. Gen. Laws Ann. ch. 269, § 10(b).

[32] *See, e.g.*, 18 U.S.C.A. § 922(g).

[33] Proposal for a Regulation Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act) and Amending Certain Union Legislative Acts, COM/2021/206 final. The latest proposed text of the AI Act is available at https://www.europarl.europa.eu/doceo/document/A-9-2023-0188_EN.html.

[34] *Id.* at amendment 36 recital 14.

[35] *Id.*

[36] *Id.*

[37] *See e.g.*, Mikolaj Barczentewicz, *supra* note 2.

[38] *Id.*

policing tools, and to impose transparency measures on generative AI applications OpenAI's ChatGPT."[39]

This approach constitutes a hodge-podge of top-down tech policing and one-off regulations. The AI Act starts with the presumption that regulators can design an abstract, high-level set of categories that capture the risk from "AI" and then proceeds to force arbitrary definitions of particular "AI" implementations into those categories. This approach may get some things right and some things wrong, but none of what good it does will be with principled consistency. For example, it might be the case that "predictive policing" is a problem that merits *per se* prohibition, but is it really an AI problem? What happens if the police get exceptionally good at using publicly available data and spreadsheets to approximate 80% of what they are able to do with AI? Or even just 50% efficacy? Is it the use of AI that is a harm, or is it the practice itself?

Similarly, a requirement that firms expose the sources on which they train their algorithms might be good in some contexts, but useless or harmful in others.[40] Certainly, it can make sense when thinking about current publicly available generative tools that create images and video, and have no ability to point to a license or permission for their training data. Such cases have a high likelihood of copyright infringement. But should *every* firm be expected to do this? Surely there will be many cases where firms use their own internal data, or data not subject to property-rights protection at all, but where exposing those sources reveals sensitive internal information, like know-how or other trade secrets. In those cases, a transparency obligation could have a chilling effect.

By contrast, it seems hard to believe that *every* use of public facial recognition should be banned. For instance, what if local authorities had limited access to facial recognition to find lost children or victims of trafficking?

More broadly, a strict transparency requirement could essentially make advanced machine-learning techniques illegal. By their nature, machine-learning systems and applications that employ LLMs make inferences and predictions that are, very often, not replicable.[41] That is, by their very nature they are not reviewable in a way that would be easily explained to a human in a transparency review. This means that strong transparency obligations could make it legally untenable to employ those techniques.

The broad risk-based approach taken by the AI Act faces difficult enforcement hurdles as well, as demonstrated by the EU's proposal to essentially ban the open-source community from providing access to generative models.[42] In other words, not only do the proposed amendments seek to prohibit

---

[39] Foo Yun Chee, Martin Coulter & Supantha Mukherjee, *EU Lawmakers' Committees Agree Tougher Draft AI Rules*, REUTERS (May 11, 2023), https://www.reuters.com/technology/eu-lawmakers-committees-agree-tougher-draft-ai-rules-2023-05-11.

[40] *See infra* at notes 71-77 and accompanying text.

[41] Explainable AI: The Basics, *supra* note 2 at 8.

[42] *See e.g.*, Delos Prime, *EU AI Act to Target US Open Source Software*, TECHNOMANCERS.AI (May 13, 2023), https://technomancers.ai/eu-ai-act-to-target-us-open-source-software.

large companies such as OpenAI, Google, Anthropic, Amazon, Microsoft, and IBM from offering API access to generative AI models, but they would also prohibit open-source developers and distributors such as GitHub from doing the same.[43] Moreover, the prohibitions have extraterritorial effects; for example, the EU might seek to impose large fines on U.S. companies for permitting access to their models *in the United States*, on grounds that those models could be imported into the EU by third parties.[44] These provisions reflect not only an attempt to control the distribution of AI technology but also the wider implications that such attempts would essentially require steering worldwide innovation down a narrow, heavily regulated path.

2. *Focus on the harm and the wrongdoers, not the innovators*

None of the foregoing is to suggest that it is impossible for AI to be misused. Where it is misused, there should be actionable legal consequences. For example, if a real-estate developer intentionally used AI tools to screen out individuals on the basis of protected characteristics from purchasing homes, that should be actionable. If a criminal found a novel way to use Chat GPT to commit fraud, that should be actionable. If generative AI is used to create "deep fakes" that further some criminal plot, that should be actionable. But in all those cases, it is not the AI itself that is the relevant unit of legal analysis, but the action of the criminal and the harm he causes.

To try to build a regulatory framework that makes it impossible for bad actors to misuse AI will be ultimately fruitless. Bad actors will always find ways to misuse tools, and heavy-handed regulatory requirements (or even strong suggestions of such) might chill the development of useful tools that could generate an enormous amount of social welfare.

B. Do Not Neglect the Benefits

A major complication in parsing the wisdom of potential AI regulation is that the technology remains largely in development. Indeed, this is the impetus for many of the calls to "do something" before it is "too late."[45] The fear that some express is that, unless a wise regulator intervenes in the development process, the technology will inevitably develop in ways that yield more harm than good.[46]

But trying to regulate AI in accordance with the precautionary principle would almost certainly stifle development and dampen the tremendous, but unknowable, good that would emerge as these

---

[43] *Id.*

[44] To be clear, it is not certain how such an extraterritorial effect will be obtained, and this is just a proposed amendment to the law. Likely, there will need to be some form of jurisdictional hook, *i.e.*, that this applies only to firms with an EU presence.

[45] Eliezer Yudkowsky, *Pausing AI Developments Isn't Enough. We Need to Shut it All Down*, TIME (Mar. 29, 2023), https://time.com/6266923/ai-eliezer-yudkowsky-open-letter-not-enough.

[46] *See, e.g.*, Kiran Stacey, *UK Should Play Leading Role on Global AI Guidelines, Sunak to Tell Biden*, THE GUARDIAN (May 31, 2023), https://www.theguardian.com/technology/2023/may/31/uk-should-play-leading-role-in-developing-ai-global-guidelines-sunak-to-tell-biden.

technologies mature and we find unique uses for them. Moreover, precautionary regulation, even in high-risk industries like nuclear power, can lead to net harms to social welfare.[47]

It is important here to distinguish two broad categories of concern about AI. First, there is the generalized concern about AGI, expressed as fear that we are inadvertently creating a super intelligence with the power to snuff out human life at its whim. We reject this fear as a legitimate basis for new regulatory frameworks, although we concede that it is theoretically possible that this presumption may need to be revisited as AI technologies progress. None of the technologies currently under consideration are anywhere *close* to AGI. They are essentially just advanced prediction engines, whether the predictions concern text or pixels.[48] It seems highly unlikely that we will accidentally stumble onto AGI by plugging a few thousand prediction engines into one another.

There are more realistic concerns that these very impressive technologies will be misused to further discrimination and crime, or will have such a disruptive impact on areas like employment that they will quickly generate tremendous harms. When contemplating harms that *could* occur, however, it is also necessary to recognize that many significant benefits could *also* be generated. Moreover, as with earlier technologies, economic disruptions will provide both challenges and opportunities. It is easy to see the immediate effect on the jobs of content writers, for instance, posed by ChatGPT, but less easy to measure the benefits that will be realized by firms that can deploy this technology to "in-source" tasks.

Firms often face what is called the "make-or-buy" decision. A firm that decides to purchase the services of an outside designer or copywriter has determined that doing so is more efficient than developing that talent in-house. But the fact that many firms employ a particular mix of outsourced and in-house talent to fulfill their business needs does not suggest a universally optimal solution to the make-or-buy problem. All we can do is describe how, under current conditions, firms solve this problem.

AI will surely augment the basis on which firms deal with the make-or-buy decision. Pre-AI, it might have made sense to outsource a good deal of work that was not core to a firm's mission. Post-AI, it might be the case that the firm can afford to hire additional workers who can utilize AI tools to more quickly and affordably manage the work that had been previously outsourced. Thus, the ability of AI tools to shift the make-or-buy decision, in itself, says nothing about the net welfare effects to society. Arguments could very well be made for either side. If history is any guide, however, it appears

---

[47] *See, e.g.*, Matthew J. Neidell, Shinsuke Uchida & Marcella Veronesi, T*he Unintended Effects from Halting Nuclear Power Production: Evidence from Fukushima Daiichi Accident*, NBER WORKING PAPER 26395 (2022), https://www.nber.org/papers/w26395 (Japan abandoning nuclear energy in the wake of the Fukushima disaster led to decreased energy consumption, which in turn led to increased mortality).

[48] *See, e.g.*, Will Knight, *Some Glimpse AGI in ChatGPT. Others Call It a Mirage*, WIRED (Apr. 10, 2023), https://www.wired.com/story/chatgpt-agi-intelligence ("GPT-4, like its predecessors, had been fed massive amounts of text and code and trained to use the statistical patterns in that corpus to predict the words that should be generated in reply to a piece of text input.")

likely that AI tools will allow firms to do more with less, while also enabling more individuals to start new businesses with less upfront expense.

Moreover, by freeing capital from easily automated tasks, existing firms and new entrepreneurs could better focus on their core business missions. Excess investments previously made in supporting, for example, the creation of marketing content could be repurposed into R&D-intensive work. Simplistic static analyses of the substitution power of AI tools will almost surely mislead us, and make us neglect the larger social welfare that could be gained from organizations improving their efficiency with AI tools.

Economists have consistently found that dynamic competition—characterized by firms vying to deliver novel and enhanced products and services to consumers—contributes significantly more to economic growth than static competition, where technology is held constant, and firms essentially compete solely on price. As Joseph Schumpeter noted:

> [I]t is not [price] competition which counts but the competition from the new commodity, the new technology, the new source of supply, the new type of organization.... This kind of competition is as much more effective than the other as a bombardment is in comparison with forcing a door, and so much more important that it becomes a matter of comparative indifference whether competition in the ordinary sense functions more or less promptly; the powerful lever that in the long run expands output and brings down prices is in any case made of other stuff.[49]

Technological advancements yield substantial welfare benefits for consumers, and there is a comprehensive body of scholarly work substantiating the contributions of technological innovation to economic growth and societal welfare. [50] There is also compelling evidence that technological progress engenders extensive spillovers not fully appropriated by the innovators.[51] Business-model innovations—such as advancements in organization, production, marketing, or distribution—can similarly result in extensive welfare gains.[52]

AI tools obviously are delivering a new kind of technological capability for firms and individuals. The disruptions they will bring will similarly spur business-model innovation as firms scramble to find innovative ways to capitalize on the technology. The potential economic dislocations can, in many cases, amount to reconstitution: a person who was a freelance content writer can be shifted to a different position that manages the output of generative AI and provides human edits to ensure

---

[49] JOSEPH A. SCHUMPETER, CAPITALISM, SOCIALISM AND DEMOCRACY 74 (1976).

[50] *See, e.g.*, Jerry Hausman, *Valuation of New Goods Under Perfect and Imperfect Competition*, in THE ECONOMICS OF NEW GOODS 209–67 (Bresnahan & Gordon eds., 1997).

[51] William D. Nordhaus, *Schumpeterian Profits in the American Economy: Theory and Measurement*, NBER Working Paper No. 10433 (Apr. 2004) at 1, http://www.nber.org/papers/w10433 ("We conclude that only a miniscule fraction of the social returns from technological advances over the 1948-2001 period was captured by producers, indicating that most of the benefits of technological change are passed on to consumers rather than captured by producers.").

[52] *See generally* OLIVER E. WILLIAMSON, MARKETS AND HIERARCHIES, ANALYSIS AND ANTITRUST IMPLICATIONS: A STUDY IN THE ECONOMICS OF INTERNAL ORGANIZATION (1975).

that content makes sense and is based in fact. In many other cases, the dislocations will likely lead to increased opportunities for workers of all sorts.

With this in mind, policymakers need to consider how to identify those laws and regulations that are most likely to foster this innovation, while also enabling courts and regulators to adequately deal with potential harms. Although it is difficult to prescribe particular policies to boost innovation, there is strong evidence about what sorts of policies should be avoided. Most importantly, regulation of AI should avoid inadvertently destroying those technologies.[53] As Adam Thierer has argued, "if public policy is guided at every turn by the fear of hypothetical worst-case scenarios and the precautionary mindset, then innovation becomes less likely."[54]

Thus, policymakers must be cautious to avoid unduly restricting the range of AI tools that compete for consumer acceptance. Key to fostering investment and innovation is not merely the endorsement of technological advancement, but advocacy for policies that empower innovators to execute and commercialize their technology.

By contrast, consider again the way that some EU lawmakers want to treat "high risk" algorithms under the AI Act. According to recently proposed amendments, if a "high risk" algorithm learns something beyond what its developers expect it to learn, the algorithm would need to undergo a conformity assessment.[55]

One of the prime strengths of AI tools is their capacity for unexpected discoveries, offering potential insights and solutions that might not have been anticipated by human developers. As the Royal Society has observed:

> Machine learning is a branch of AI that enables computer systems to perform specific tasks intelligently. Traditional approaches to programming rely on hardcoded rules, which set out how to solve a problem, step-by-step. In contrast, machine learning systems are set a task, and given a large amount of data to use as examples (and non-examples) of how this task can be achieved, or from which to detect patterns. The system then learns how best to achieve the desired output.[56]

By labeling unexpected behavior as inherently risky and necessitating regulatory review, we risk stifling this serendipitous aspect of AI technologies, potentially curtailing their capacity for innovation. It could contribute to a climate of regulatory caution that hampers swift progress in discovering the full potential and utility of AI tools.

---

[53] *See, e.g.,* NASSIM NICHOLAS TALEB, ANTIFRAGILE: THINGS THAT GAIN FROM DISORDER (2012) ("In action, [*via negativa*] is a recipe for what to avoid, what not to do.").

[54] ADAM THIERER, PERMISSIONLESS INNOVATION: THE CONTINUING CASE FOR COMPREHENSIVE TECHNOLOGICAL FREEDOM (2016).

[55] *See, e.g.,* Artificial Intelligence Act, *supra* note 33, at amendment 112 recital 66.

[56] EXPLAINABLE AI: THE BASICS, *supra* note 2 at 6.

C. AI Regulation Should Follow the Model of Common Law

In a recent hearing of the U.S. Senate Judiciary Committee, OpenAI CEO Sam Altman suggested that the United States needs a central "AI regulator."[57] As a general matter, we expect this would be unnecessarily duplicative. As we have repeatedly emphasized, the right approach to regulating AI is not the establishment of an overarching regulatory framework, but a careful examination of how AI technologies will variously interact with different parts of the existing legal system. We are not alone in this; former Special Assistant to the President for Technology and Competition Policy Tim Wu recently opined that federal agencies would be well-advised to rely on existing law and enhance that law where necessary in order to catch unexpected situations that may arise from the use of AI tools.[58]

As Judge Easterbrook famously wrote in the context of what was then called "cyberspace," we do not need a special law for AI any more than we need a "law of the horse."[59]

1. An AI regulator's potential effects on competition

More broadly, there are risks to competition that attend creating a centralized regulator for a new technology like AI. As an established player in the AI market, OpenAI might favor a strong central regulator because of the potential that such an agency could act in ways that hinder the viability of new entrants.[60] In short, an incumbent often can gain by raising its rivals' regulatory costs, or by manipulating the relationship between its industry's average and marginal costs. This dynamic can create strong strategic incentives for industry incumbents to promote regulation.

Economists and courts have long studied actions that generate or amplify market dominance by placing competitors at a disadvantage, especially by raising rivals' costs.[61] There exist numerous strategies to put competitors at a disadvantage or push them out of the market without needing to compete on price. While antitrust action focuses on private actors and their ability to raises rival's costs, it is well-accepted that "lobbying legislatures or regulatory agencies to create regulations that disadvantage rivals" has similar effects.[62]

---

[57] Cecilia Kang, *OpenAI's Sam Altman Urges A.I. Regulation in Senate Hearing*, NY TIMES (May 16, 2023), https://www.nytimes.com/2023/05/16/technology/openai-altman-artificial-intelligence-regulation.html; *see also* Mike Solana & Nick Russo, *Regulate Me, Daddy*, PIRATE WIRES (May 23, 2023), https://www.piratewires.com/p/regulate-me-daddy.

[58] Cristiano Lima, *Biden's Former Tech Adviser on What Washington is Missing about AI*, THE WASHINGTON POST (May 30, 2023), https://www.washingtonpost.com/politics/2023/05/30/biden-former-tech-adviser-what-washington-is-missing-about-ai.

[59] Frank H. Easterbrook, *supra* note 5.

[60] *See* Lima, *supra* note 58 ("I'm not in favor of an approach that would create heavy compliance costs for market entry and that would sort of regulate more abstract harms.")

[61] Steven C. Salop & David T. Scheffman, *Raising Rivals' Costs*, 73:2 AM. ECON. R. 267, 267–71 (1983), http://www.jstor.org/stable/1816853.

[62] Steven C. Salop & David T. Scheffman, *Cost-Raising Strategies*, 36:1 J. INDUS. ECON. 19 (1987), https://doi.org/10.2307/2098594.

Suppose a new regulation costs $1 million in annual compliance costs. Only companies that are sufficiently large and profitable will be able to cover those costs, which keeps out newcomers and smaller competitors. This effect of keeping out smaller competitors by raising their costs may more than offset the regulatory burden on the incumbent. New entrants typically produce on a smaller scale, and therefore find it more difficult to spread increased costs over a large number of units. This makes it harder for them to compete with established firms like OpenAI, which can absorb these costs more easily due to their larger scale of production.

This type of cost increase can often look benign. In the *United Mine Workers vs. Pennington*[63] case, a coal corporation was alleged to have conspired with the union representing its workforce to establish *higher* wage rates. How could higher wages be anticompetitive? This seemingly contradictory conclusion came from University of California at Berkeley economist Oliver Williamson, who interpreted the action to be an effort to maximize profits by raising entry barriers.[64] Using a model with a dominant incumbent and a fringe of other competitors, he demonstrated that wage-rate increases could lead to profit maximization if they escalated the fringe's costs more than they did the dominant firm's costs. Intuitively, while the dominant firm is dominant, the market price is determined by the marginal producers and the dominant company's price is determined by the prices of its competitors. If a regulation raises the competitors' per-unit costs by $2, the dominant company will be able to raise its price by as much as $2 per unit. Even if the regulation hurts the dominant firm, so long as its price increase exceeds its additional cost, the dominant firm can profit from the regulation.

As a result, while regulations might increase costs for OpenAI, they also serve to protect it from potential competition by raising the barriers to entry. In this sense, regulation can be seen as a strategic tool for incumbent firms to maintain or strengthen their market position. None of this analysis rests on OpenAI explicitly wanting to raise its rivals' costs. That is just the competitive implication of such regulations. Thus, while there may be many benign reasons for a firm like OpenAI to call for regulation in good faith, the ultimate lesson presented by the economics of regulation should counsel caution when imposing strong centralized regulations on a nascent industry.

2. *A central licensing regulator for AI would be a mistake*

NTIA asks:

> Are there ways in which accountability mechanisms are unlikely to further, and might even frustrate, the development of trustworthy AI? Are there accountability mechanisms that unduly impact AI innovation and the competitiveness of U.S. developers?[65]

---

[63] *United Mine Workers of Am. v. Pennington*, 381 U.S. 657, 661 (1965).

[64] Oliver E. Williamson, *Wage Rates as a Barrier to Entry: The Pennington Case in Perspective*, 82:1 Q. J. ECON. 85 (1968), https://doi.org/10.2307/1882246.

[65] RFC at 22439.

We are not alone in the belief that imposing a licensing regime would present just such a barrier to innovation.[66] In the recent Senate hearings, the idea of a central regulator was endorsed as means to create and administer a licensing regime.[67] Perhaps in some narrow applications of particular AI technologies, there could be specific contexts in which licensing is appropriate (*e.g.*, in providing military weapons), but broadly speaking, we believe this is inadvisable. Owing to the highly diverse nature of AI technologies, trying to license AI development is a fraught exercise, as NTIA itself acknowledges:

> A developer training an AI tool on a customer's data may not be able to tell how that data was collected or organized, making it difficult for the developer to assure the AI system. Alternatively, the customer may use the tool in ways the developer did not foresee or intend, creating risks for the developer wanting to manage downstream use of the tool. When responsibility along this chain of AI system development and deployment is fractured, auditors must decide whose data and which relevant models to analyze, whose decisions to examine, how nested actions fit together, and what is within the audit's frame.[68]

Rather than design a single regulation to cover AI, ostensibly administered through a single licensing regime, NTIA should acknowledge the broad set of industries currently seeking to employ a diverse range of AI products that differ in fundamental ways. The implications of AI deployment in health care, for instance, vastly differ from those in transportation. A centralized AI regulator might struggle to comprehend the nuances and intricacies of each distinct industry, thus potentially leading to ineffective or inappropriate licensing requirements.

Analogies have been drawn between AI and sectors like railroads and nuclear power, which have dedicated regulators.[69] These sectors, however, are more homogenous and discrete than the AI industry (if such an industry even exists, apart from the software industry more generally). AI is much closer to a general-purpose tool, like chemicals or combustion engines. We do not enact central regulators to license every aspect of the development and use of chemicals, but instead allow different agencies to treat their use differently as is appropriate for the context. For example, the Occupational Safety and Health Administration (OSHA) will regulate employee exposure to dangerous substances encountered in the workplace, while various consumer-protection boards will regulate the adulteration of goods.

---

[66] *See, e.g.*, Lima, *supra* note 58 ("Licensing regimes are the death of competition in most places they operate").

[67] Kang, *supra* note 57; *Oversight of A.I.: Rules for Artificial Intelligence: Hearing Before the Subcomm. on Privacy, Technology, and the Law of the S. Comm. on the Judiciary*, 118th Cong. (2023) (statement of Sam Altman, at 11), *available at* https://www.judiciary.senate.gov/download/2023-05-16-testimony-altman.

[68] RFC at 22437.

[69] *See*, e.g., *Transcript: Senate Judiciary Subcommittee Hearing on Oversight of AI*, TECH POLICY PRESS (May 16, 2023), https://techpolicy.press/transcript-senate-judiciary-subcommittee-hearing-on-oversight-of-ai ("So what I'm trying to do is make sure that you just can't go build a nuclear power plant. Hey Bob, what would you like to do today? Let's go build a nuclear power plant. You have a nuclear regulatory commission that governs how you build a plant and is licensed.")

The notion of licensing implies that companies would need to obtain permission prior to commercializing a particular piece of code. This could introduce undesirable latency into the process of bringing AI technologies to market (or, indeed, even of correcting errors in already-deployed products). Given the expansive potential to integrate AI technologies into diverse products and services, this delay could significantly impede technological progress and innovation. Given the strong global interest in the subject, such delays threaten to leave the United States behind its more energetic competitors in the race for AI innovation.

As in other consumer-protection regimes, a better approach would be to eschew licensing and instead create product-centric and harm-centric frameworks that other sectoral regulators or competition authorities could incorporate into their tailored rules for goods and services.

For instance, safety standards for medical devices should be upheld, irrespective of whether AI is involved. This product-centric regulatory approach would ensure that the desired outcomes of safety, quality, and effectiveness are achieved without stymieing innovation. With their deep industry knowledge and experience, sectoral regulators will generally be better positioned to address the unique challenges and considerations posed by AI technology deployed within their spheres of influence.

NTIA alludes to one of the risks of an overaggregated regulator when it notes that:

> For some trustworthy AI goals, it will be difficult to harmonize standards across jurisdictions or within a standard- setting body, particularly if the goal involves contested moral and ethical judgements. In some contexts, not deploying AI systems at all will be the means to achieve the stated goals.[70]

Indeed, the institutional incentives that drive bureaucratic decision making often converge on this solution of preventing unexpected behavior by regulated entities.[71] But at what cost? If a regulator is unable to imagine how to negotiate the complicated tradeoffs among interested parties across all AI-infused technologies, it will act to slow or prevent the technology from coming to market. This will make us all worse off, and will only strengthen the position of our competitors on the world stage.

D. The Impossibility of Explaining Complexity

NTIA notes that:

> According to NIST, "trustworthy AI" systems are, among other things, "valid and reliable, safe, secure and resilient, accountable and transparent, explainable and interpretable, privacy-enhanced, and fair with their harmful bias managed."[72]

---

[70] RFC at 22438.

[71] *See, e.g.*, Raymond J. March, *The FDA and the COVID-19: A Political Economy Perspective*, 87(4) S. ECON. J. 1210, 1213-16 (2021), https://www.ncbi.nlm.nih.gov/pmc/articles/PMC8012986 (discussing the political economy that drives incentives of bureaucratic agencies in the context of the FDA's drug-approval process).

[72] RFC at 22434.

And in the section titled "Accountability Inputs and Transparency," NTIA asks a series of questions designed to probe what can be considered a realistic transparency obligation for developers and deployers of AI systems. We urge NTIA to resist the idea that AI systems be "explainable," for the reasons set forth herein.

One of the significant challenges in AI accountability is making AI systems explainable to users. It is crucial to acknowledge that providing a clear explanation of how an AI model—such as an LLM or a diffusion model—arrives at a specific output is an inherently complex task, and may not be possible at all. As the UK Royal Society has noted in its paper on AI explainability:

> Much of the recent excitement about advances in AI has come as a result of advances in statistical techniques. These approaches – including machine learning – often leverage vast amounts of data and complex algorithms to identify patterns and make predictions. This complexity, coupled with the statistical nature of the relationships between inputs that the system constructs, renders them difficult to understand, even for expert users, including the system developers. [73]

These models are designed with intricate architectures and often rely on vast troves of data to arrive at outputs, which can make it nearly impossible to reverse-engineer the process. Due to these complexities, it may be unfeasible to make AI fully explainable to users. Moreover, users themselves often do not value explainability, and may be largely content with a "black box" system when it consistently provides accurate results.[74]

Instead, to the extent that regulators demand visibility into AIs, the focus should be on the transparency of the AI-development process, system inputs, and the general guidelines for AI that developers use in preparing their models. Ultimately, we suspect that, even here, such measures will do little to resolve the inherent complexity in understanding how AI tools produce their outputs.

In a more limited sense, we should consider the utility in transparency of AI-infused technology for most products and consumers. NTIA asks:

> Given the likely integration of generative AI tools such as large language models (e.g., ChatGPT) or other general-purpose AI or foundational models into downstream products, how can AI accountability mechanisms inform people about how such tools are operating and/or whether the tools comply with standards for trustworthy AI?[75]

As we note above, the proper level of analysis for AI technologies is the product into which they are incorporated. But even there, we need to ask whether it matters to an end user whether a product they are using relies on ChatGPT or a different algorithm for predictively generating text. If the product malfunctions, what matters is the malfunction and the accountability for the product. Most users do not really care whether a developer writes a program using C++ or Java, and neither should

---

[73] EXPLAINABLE AI: THE BASICS, *supra*, note 2 at 12.

[74] *Id.* at 20.

[75] *Id.* at 22439.

they explicitly care whether he incorporates a generative AI algorithm to predict text, or uses some other method of statistical analysis. The presence of an AI component becomes analytically necessary when diagnosing *how* something went wrong, but *ex ante*, it is likely irrelevant from a consumer's perspective.

Thus, it may be the case that a more fruitful avenue for NTIA to pursue would be to examine how a strict-liability or product-liability legal regime might be developed for AI. These sorts of legal frameworks put the onus on AI developers to ensure that their products behave appropriately. Such legal frameworks also provide consumers with reassurance that they have recourse if and when they are harmed by a product that contains AI technology. Indeed, it could very well be the case that overemphasizing "trust" in AI systems could end up misleading users in important contexts.[76] This would strengthen the case for a predictable liability regime.

1. The deepfakes problem demonstrates that we do not need a new body of law

The phenomenon of generating false depictions of individuals using advanced AI techniques—commonly called "deepfakes"—is undeniably concerning, particularly when it can be used to create detrimental false public statements,[77] facilitate fraud,[78] or create nonconsensual pornography.[79] But while deepfakes use modern technological tools, they are merely the most recent iteration of the age-old problem of forgery. Importantly, existing law already equips us with the tools needed to address the challenges posed by deepfakes, rendering many recent legislative proposals at the state level both unnecessary and potentially counterproductive. Consider one of the leading proposals offered by New York State.[80]

Existing laws in New York and at the federal level provide remedies for individuals aggrieved by deepfakes, and they do so within a legal system that has already worked to incorporate the context of these harms, as well as the restrictions of the First Amendment and related defenses. For example, defamation laws can be applied where a deepfake falsely suggests an individual has posed for an explicit photograph or video.[81] New York law also acknowledges the tort of intentional infliction of emotional distress, which likely could be applied to the unauthorized use of a person's likeness in

---

[76] EXPLAINABLE AI: THE BASICS, *supra* note 2 at 22. ("Not only is the link between explanations and trust complex, but trust in a system may not always be a desirable outcome. There is a risk that, if a system produces convincing but misleading explanations, users might develop a false sense of confidence or understanding, mistakenly believing it is trustworthy as a result.")

[77] Kate Conger, *Hackers' Fake Claims of Ukrainian Surrender Aren't Fooling Anyone. So What's Their Goal?*, NY TIMES (Apr. 5, 2022), https://www.nytimes.com/2022/04/05/us/politics/ukraine-russia-hackers.html.

[78] Pranshu Verma, *They Thought Loved Ones Were Calling for Help. It Was an AI Scam*, THE WASHINGTON POST (Mar. 5, 2023), https://www.washingtonpost.com/technology/2023/03/05/ai-voice-scam.

[79] *Video: Deepfake Porn Booms in the Age of A.I.*, NBC NEWS (Apr. 28, 2023), https://www.nbcnews.com/now/video/deepfake-porn-booms-in-the-age-of-a-i-171726917562.

[80] S5857B, NY State Senate (2018), https://www.nysenate.gov/legislation/bills/2017/s5857/amendment/b.

[81] *See, e.g.*, *Rejent v. Liberation Publications, Inc.*, 197 A.D.2d 240, 244–45 (1994); *see also*, *Leser v. Penido*, 62 A.D.3d 510, 510–11 (2009).

explicit content.[82] In addition, the tort of unjust enrichment can be brought to bear where appropriate, as can the Lanham Act §43(a), which prohibits false advertising and implied false endorsements.[83] Furthermore, victims may hold copyright in the photograph or video used in a deepfake, presenting grounds for an infringement action.[84]

Thus, while advanced deepfakes are new, the harms they can cause and the law's ability to address those harms is not novel. Legislation that attempts to carve out new categories of harms in these situations are, at best, reinventing the wheel and, at worst, risk creating confusing tensions in the existing legal system.

III. The Role of NTIA in AI Accountability

NTIA asks if "the lack of a federal law focused on AI systems [is] a barrier to effective AI accountability?"[85] In short, no, this is not a barrier, so long as the legal system is allowed to evolve to incorporate the novel challenges raised by AI technologies.

As noted in the previous section, there is a need to develop standards, both legal and technical. As we are in the early days of AI technology, the exact contours of the various legal changes that might be needed to incorporate AI tools into existing law remain unclear. At this point, we would urge NTIA—to the extent that it wants to pursue regulatory, licensing, transparency, and other similar obligations—to develop a series of workshops through which leading technology and legal experts could confer on developing a vision for how such legal changes would work in practice.

By gathering stakeholders and fostering an ongoing dialogue, NTIA can help to create a collaborative environment in which organizations can share knowledge, experiences, and innovations to address AI accountability and its associated challenges. By promoting industry collaboration, NTIA could also help build a foundation of trust and cooperation among organizations involved in AI development and deployment. This, in turn, will facilitate the establishment of standards and best practices that address specific concerns, while mitigating the risk of overregulation that could stifle innovation and progress. In this capacity, NTIA should focus on encouraging the development of context-specific best practices that prioritize the containment of identifiable harms. By fostering a collaborative atmosphere, the agency can support a dynamic and adaptive AI ecosystem that is capable of addressing evolving challenges while safeguarding the societal benefits of AI advancements.

In addressing AI accountability, it is essential for NTIA to adopt a harm-focused framework that targets the negative impacts of AI systems rather than the technology itself. This approach would recognize that AI technology can have diverse applications, with consequences that will depend on

---

[82] *See, e.g., Howell v. New York Post Co,*. 612 N.E.2d 699 (1993).

[83] *See, e.g., Mandarin Trading Ltd. v. Wildenstein*, 944 N.E.2d 1104 (2011); 15 U.S.C. §1125(a).

[84] 17 U.S.C. 106.

[85] RFC at 22440.

the context in which they are used. By prioritizing the mitigation of specific harms, NTIA can ensure that regulations are tailored to address real-world outcomes and provide a more targeted and effective regulatory response.

A harm-focused framework also acknowledges that different AI technologies pose differing levels of risk and potential for misuse. NTIA can play a proactive role in guiding the creation of policies that reflect these nuances, striking a balance between encouraging innovation and ensuring the responsible development and use of AI. By centering the discussion on actual harms and their causes, NTIA can foster meaningful dialogue among stakeholders and facilitate the development of industry best practices designed to minimize negative consequences.

Moreover, this approach ensures that AI accountability policies are consistent with existing laws and regulations, as it emphasizes the need to assess AI-related harms within the context of the broader legal landscape. By aligning AI accountability measures with other established regulatory frameworks, the NTIA can provide clear guidance to AI developers and users, while avoiding redundancy and conflicting regulations. Ultimately, a harm-focused framework allows the NTIA to better address the unique challenges posed by AI technology and foster an assurance ecosystem that prioritizes safety, ethics, and legal compliance without stifling innovation.

IV. Conclusion

Another risk of the current AI hysteria is that fatigue will set in, and the public will become numbed to potential harms. Overall, this may shrink the public's appetite for the kinds of legal changes that will be needed to address those actual harms that do emerge. News headlines that push doomsday rhetoric and a community of experts all too eager to respond to the market incentives for apocalyptic projections only exacerbate the risk of that outcome. A recent one-line letter, signed by AI scientists and other notable figures, highlights the problem:

> Mitigating the risk of extinction from AI should be a global priority alongside other societal-scale risks such as pandemics and nuclear war.[86]

Novel harms absolutely will emerge from products that employ AI, as has been the case for every new technology. The introduction of automobiles created new risks of harm from high-speed auto-related deaths, for example. But rhetoric about AI being an existential risk on the level of a pandemic or nuclear war is irresponsible.

Perhaps one of the most important positions NTIA can assume, therefore, is that of a calm, collected expert agency that helps restrain the worst impulses to regulate AI out of existence due to blind fear.

In essence, the key challenge confronting policymakers lies in navigating the dichotomy of mitigating actual risks presented by AI, while simultaneously safeguarding the substantial benefits it offers. It is undeniable that the evolution of AI will bring about disruption and may provide a conduit for

[86] *Statement on AI Risk*, CENTER FOR AI SAFETY, https://www.safe.ai/statement-on-ai-risk (last visited Jun. 7, 202).

malevolent actors, just as technologies like the printing press and the internet have done in the past. This does not, however, merit taking an overly cautious stance that would suppress the potential benefits of AI.

As we formulate policy, it is crucial to eschew dystopian science-fiction narratives and instead ground our approach in realistic scenarios. The proposition that computer systems, even those as advanced as AI tools, could spell the end of humanity lacks substantial grounding.

The current state of affairs represents a geo-economic competition to harness the benefits of AI in myriad domains. Contrary to fears that AI poses an existential risk, the real danger may well lie in attempts to overly regulate and stifle the technology's potential. The indiscriminate imposition of regulations could inadvertently thwart AI advancements, resulting in a loss of potential benefits that could be far more detrimental to social welfare.