1-1-2024

Lessons from GDPR for AI Policymaking

Christopher S. Yoo
University of Pennsylvania Carey Law School, csyoo@law.upenn.edu

Josephine Wolff
Fletcher School of Law and Diplomacy, Tufts University, josephine.wolff@tufts.edu

William Lehr
Massachusetts Institute of Technology, wlehr@mit.edu

Recommended Citation
27 Va. J. L. & Tech. 1 (2024)

VIRGINIA JOURNAL OF LAW & TECHNOLOGY

VOLUME 27 WINTER 2024 NUMBER 4

LESSONS FROM GDPR FOR AI POLICYMAKING

Josephine Wolff,* William Lehr,† and Christopher S. Yoo‡

ABSTRACT

The ChatGPT chatbot has not just caught the public imagination; it is also amplifying concern across industry, academia, and government policymakers interested in the regulation of Artificial Intelligence (AI) about how to understand the risks and threats associated with AI applications. Following the release of ChatGPT, some EU regulators proposed changes to the draft EU AI Act to classify AI systems like ChatGPT that generate complex texts without any human oversight as "high-risk" AI systems that would fall under the law's requirements. That classification was a controversial one, with other regulators arguing that technologies like ChatGPT, which merely generate text, are "not risky at all." This controversy risks disrupting coherent discussion and progress toward formulating sound AI regulations for Large Language Models (LLMs), AI, or Information and Communications Technologies (ICTs) more generally. It remains unclear where ChatGPT fits within AI and where AI fits within the larger context of digital policy and the regulation of ICTs despite nascent efforts by OECD.AI and the EU.

This paper aims to address two research questions around AI policy: (1) How are LLMs like ChatGPT shifting the policy discussions around AI regulations? (2) What lessons can regulators learn from the

* Associate Professor of Cybersecurity Policy, Fletcher School; Associate Professor of Computer Science; and Director of the Hitachi Center for Technology and International Affairs, Tufts University.

† Research Associate, Computer Science and Artificial Intelligence Laboratory (CSAIL), Massachusetts Institute of Technology.

‡ John H. Chestnut Professor of Law, Communication, and Computer & Information Science and Founding Director of the Center for Technology, Innovation & Competition, University of Pennsylvania.

EU's General Data Protection Regulation (GDPR) and other data protection policymaking efforts that can be applied to AI policymaking? The first part of the paper addresses the question of how ChatGPT and other LLMs have changed the policy discourse in the EU and other regions around regulating AI and what the broader implications of these shifts may be for AI regulation more widely. This section reviews the existing proposal for an EU AI Act and its accompanying classification of high-risk AI systems, considers the changes prompted by the release of ChatGPT and examines how LLMs appear to have altered policymakers' conceptions of the risks presented by AI. Finally, we present a framework for understanding how the security and safety risks posed by LLMs fit within the larger context of risks presented by AI and current efforts to formulate a regulatory framework for AI.

The second part of the paper considers the similarities and differences between the proposed EU AI Act and GDPR in terms of (1) organizations being regulated, or scope, (2) reliance on organizations' self-assessment of potential risks, or degree of self-regulation, (3) penalties, and (4) technical knowledge required for effective enforcement, or complexity. For each of these areas, we consider how regulators scoped or implemented GDPR to make it manageable, enforceable, meaningful, and consistent across a wide range of organizations handling many different kinds of data, as well as the extent to which they were successful in doing so. We then examine different ways in which those same approaches may or may not be applicable to the proposed EU AI Act and the ways in which AI may prove more difficult to regulate than issues of data protection and privacy covered by GDPR. We also look at the ways in which AI may make it more difficult to enforce and comply with GDPR since the continued evolution of AI technologies may create cybersecurity tools and threats that will impact the efficacy of GDPR and privacy policies. This section argues that the extent to which the proposed EU AI Act relies on self-regulation and the technical complexity of enforcement are likely to pose significant challenges to enforcement based on the implementation of the most technologically and self-regulation-focused elements of GDPR.

I. INTRODUCTION

IN February 2023, Brando Benifei and Dragoş Tudorache, the two members of the European Parliament who serve as co-rapporteurs for the proposed EU Artificial Intelligence Act, reportedly suggested a series of amendments to the list of "high-risk" AI applications that the Act would cover.[1] The initial draft released in 2021 had designated several AI systems as high risk, including those used for recruitment and hiring and those to determine eligibility for public assistance benefits and services, to provide law enforcement with help assessing the risk that someone might break the law, or to help courts research and interpret the law. The 2023 amendments suggested classifying several additional types of AI systems as high risk, including those "likely to influence democratic processes like elections," those that "may have serious effects on a child's personal development," and "generative AI systems such as ChatGPT."[2] It was a stark shift in tone from mentions of generative AI in the earlier 2021 proposal of the draft law, which had largely focused on deepfakes and specified, "For some specific AI systems, only minimum transparency obligations are proposed, in particular when chatbots or 'deep fakes' are used."[3] Tudorache himself had told Reuters in early 2023 that he believed generative AI was "not going to be covered" in the proposed EU AI Act in depth, saying, "That's another discussion I don't think we are going to deal with in this text."[4]

Contrary to Tudorache's prediction, in October 2023 the Spanish president of the EU Council of Ministers proposed adopting a tiered approach that

[1] Luca Bertuzzi, *AI Act: EU Parliament's Crunch Time on High-Risk Categorisation, Prohibited Practices*, Euractiv (Feb. 7, 2023), https://www.euractiv.com/section/artificial-intelligence/news/ai-act-eu-parliaments-crunch-time-on-high-risk-categorisation-prohibited-practices/.

[2] Ophélie Stockhem & Asha Allen, *CDT Europe's AI Bulletin: February 2023*, CTR. FOR DEMOCRACY & TECH. (Feb. 23, 2023), https://cdt.org/insights/cdt-europes-ai-bulletin-february-2023/.

[3] *Proposal for a Regulation of the European Parliament and of the Council Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act) and Amending Certain Union Legislative Acts*, COM (2021) 206 final (Apr. 21, 2021), https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex%3A52021PC0206 [hereinafter *Proposed EU AI Act*].

[4] Martin Coulter & Supantha Mukherjee, *Exclusive: Behind EU Lawmakers' Challenge to Rein in ChatGPT and Generative AI*, REUTERS (May 1, 2023), https://www.reuters.com/technology/behind-eu-lawmakers-challenge-rein-chatgpt-generative-ai-2023-04-28/.

would subject "high impact" foundation models and General Purpose AI (GPAI) systems to higher levels of regulation.[5] Despite concerns raised by France, Germany, and Italy that regulating foundation models would be inconsistent with the risk-based approach that focuses on the uses of AI rather than the technology itself,[6] a trilogue between the EU Commission, Council, and Parliament yielded a compromise agreement in December 2023 that subjected all GPAI with "systemic risk" foundation models to meet greater requirements.[7]

The speed with which the release of ChatGPT in November 2022 appeared to change European regulators' perceptions of the risks associated with generative AI and large language models (LLMs) and the need to regulate those systems hints at just how malleable and undecided regulators' understandings of the risks associated with different types of AI are. It also highlights the challenges associated with drafting a risk-based regulation for a technology whose risks are still relatively poorly understood. Over the course of less than a year, the release of ChatGPT led to significant changes in the draft EU AI Act that were specifically designed to respond to concerns about a type of technology that most European regulators had previously not viewed as high risk. It is not unusual or even necessarily surprising that a popular new technology should influence regulation or interest regulators in new risks with which they had not previously been concerned. However, the regulatory impact

[5] Luca Bertuzzi, *Spanish Presidency Pitches Obligations, for Foundation Models in EU's AI Law*, EURACTIV (Nov. 7, 2023), https://www.euractiv.com/section/artificial-intelligence/news/spanish-presidency-pitches-obligations-for-foundation-models-in-eus-ai-law/. The new proposal reportedly defines high-impact foundation models as "any foundation model trained with large amount of data and with advanced complexity, capabilities and performance well above the average for foundation models, which can disseminate systemic risks along the value chain, regardless there are integrated or not in a high-risk system." *Id. See also* Luca Bertuzzi, *AI Act: EU Countries Headed to Tiered Approach on Foundation Models amid Broader Compromise*, EURACTIV (Oct. 18, 2023), https://www.euractiv.com/section/artificial-intelligence/news/ai-act-eu-countries-headed-to-tiered-approach-on-foundation-models-amid-broader-compromise/
("Using the high-impact category as the basis of a two-tiered regulatory system replaced an earlier proposal that focused on "very capable foundation models.").

[6] Luca Bertuzzi, *France, Germany, Italy Push for 'Mandatory Self-Regulation' for Foundational Models in EU's AI Law*, EURACTIV (Nov. 21, 2023), https://www.euractiv.com/section/artificial-intelligence/news/france-germany-italy-push-for-mandatory-self-regulation-for-foundation-models-in-eus-ai-law/.

[7] European Parliament Press Release PR 15699, Artificial Intelligence Act: Deal on Comprehensive Rules for Trustworthy AI (Dec. 9, 2023), https://www.europarl.europa.eu/news/en/press-room/20231206IPR15699/artificial-intelligence-act-deal-on-comprehensive-rules-for-trustworthy-ai; *see also* European Commission Press Release IP/23/6473, Commission Welcomes Political Agreement on Artificial Intelligence Act (Dec. 9, 2023), https://ec.europa.eu/commission/presscorner/detail/en/ip_23_6473.

of ChatGPT in Europe is an instructive example of how policymakers around the world struggle with designing regulations intended to counter emerging technological risks related to AI systems that can be used for a variety of different purposes, some of which may be high risk and others of which may not.

This raises two questions: (1) how has the global concern over ChatGPT and LLMs shifted the discussion around Artificial Intelligence (AI) regulation, and (2) what might earlier experiences with GDPR teach us about crafting policies aimed at emerging technological risks? More broadly, this paper considers the question of how we should adapt our regulatory institutions in response to the transformative potential posed by the transition to a digital economy.

This is not a new question. Information technologies have already significantly transformed how many tasks are performed by partial or full-scale automation. The process of augmenting human physical and cognitive tasks with IT and machine assistance has been ongoing for a long time, but the release of ChatGPT in November 2022 confronted the world with significant progress in AI technologies. Much of the concern and fear is that this transformation might leave humans on the sidelines as super-capable, super-intelligent AIs take over the world.[8] Regardless of what one thinks of the likelihood of such an outcome, it is certainly timely that policymakers consider what role regulatory institutions should play in the digital future and with respect to AI. Unfortunately, the sudden amplification of sensationalist attention from all quarters prompted by the introduction of ChatGPT to the general public may not offer the right stimulus for sound policy progress.

[8] Two quotes from Weizenbaum and Bostrom are apt. *See* JOSEPH WEIZENBAUM, COMPUTER POWER AND HUMAN REASON: FROM JUDGMENT TO CALCULATION at X (1976) ("There are certain tasks which computers ought not be made to do, independent of whether computers can be made to do them."); NICK BOSTROM, SUPERINTELLIGENCE: PATHS, STRATEGIES, DANGERS 1 (2014) ("If some day we build machine brains that surpass human brains in general intelligence, then this new superintelligence could become very powerful. And, as the fate of the gorillas now depends more on us humans than on the gorillas themselves, so the fate of our species would depend on the actions of the machine superintelligence."). Joseph Weizenbaum, an MIT computer scientist and earlier pioneer in Natural Language Processing (NLP), developed ELIZA in the mid-1960s as an early example of a chatbot program and precursor to ChatGPT. After becoming alarmed when psychiatrists touted the potential for more advanced versions of ELIZA serving as automated therapists, Professor Weizenbaum emerged as one of the leading critics of AI. His concerns were not because of the limitations in what AI might be able to automate, but because of the social-economic harms that such automation might bring if unrestricted. The British philosopher Nicholas Bostrom opens his book about the potential and implications of an AI superintelligence by suggesting that AI may one day replace humans as the arbiters of humanity's future.

This essay is a collaborative effort between an economist, a lawyer, and a cybersecurity scholar who have been engaged in digital transformation-related research and policy analysis for several decades. Although we have each been engaged peripherally in matters related to AI, our principal research and professional activities have not been narrowly focused on AI technologies or their regulation. Herein, we take a broader view to offer our opinions and hypothesis relating to two themes: (1) the impact of ChatGPT on regulatory efforts related to AI in the EU,[9] and (2) the legacy of GDPR, the European Global Data Protection Regulation that was passed in 2016 and became effective in 2018.[10]

The motivation for this effort was prompted by the global reaction that followed the release of ChatGPT, including calls from many leading experts across academia, industry, and governments raising concerns about the risks posed by the new technology and LLMs more generally. We agree that AI poses a significant regulatory challenge that merits investigation. However, we see in the current furor a serious risk of harm to sound digital policymaking. It might induce hasty regulatory efforts, resulting in misdirected yet burdensome interventions that will raise costs and slow progress without effectively addressing the perceived risks. Alternatively, the furor may damage nascent sound policymaking efforts by stoking paranoia and the attention burnout that excessively overhyped risks or benefits can lead to. We do not doubt that AI comprises a host of important technologies that are likely to be beneficial and essential for our digital future. We also recognize that AI's use may further exacerbate global challenges, such as income disparities, cybercrime, and climate change. AI is a tool, and its welfare implications depend on how it is used. It is certainly possible that our collective failure to manage how AI is used will pose an existential threat to humanity. However, ChatGPT is hardly exemplary of that threat or indicative of world domination by a super-intelligent AI any time soon. Thus, it is worth distinguishing what ChatGPT does tell us about AI and its regulatory relevance.

It is also worth considering the legacy of GDPR, which, as a major piece of digital policymaking, is likely to influence the direction of AI policymaking. To extract those lessons, it is useful to compare and contrast the challenges that GDPR and AI seek to address. First, it will be important to recognize that AI's need for and use of data means that it will necessarily be heavily dependent

[9] *Proposed EU AI Act*, *supra* note 3; *see also* Cat Zakrzewski & Cristiano Lima, *Europe Moves Ahead on AI Regulation, Challenging Tech Giants' Power*, WASH. POST (June 14, 2023), https://www.washingtonpost.com/technology/2023/06/14/eu-parliament-approves-ai-act/.
[10] Regulation 2016/679, 2016 O.J. (L 119) (EU) [hereinafter *GDPR*].

on GDPR, which is a keystone of the European (and hence, global) data management regulatory framework. Second, the challenges addressed by AI regulation are much more complex, broader, and more uncertain than those tackled by GDPR, and efforts to advance AI regulation are likely to have significant implications for the future of GDPR. Moreover, it is worth considering the extent to which AI regulation is being influenced by individual technologies, with ChatGPT serving as a prime example of AI technology, but AI and the regulatory challenges it poses for societies and economies go well beyond ChatGPT. It is possible that in the near term, specialized rules for LLMs and other generative AI technologies may be adopted, but it is worth considering the broader implications of crafting general AI regulations specifically to address these technologies.

II. ChatGPT's Impact on the Proposed EU AI Act

On June 14, 2023, the European Parliament approved a draft version of the EU AI Act that included several amendments to the initial text proposed in 2021.[11] While the proposal to classify generative AI systems as high risk did not make it into the final approved amendments, several other changes seemingly designed to target issues raised by ChatGPT were approved. This section considers those changes and the impact ChatGPT had on the proposed EU AI Act overall. Negotiations over the final text of the AI Act are ongoing and, as discussed below, have proven to be quite controversial, so the language discussed here may not be the final version passed into law in the EU. However, as the first set of changes made following the release of ChatGPT, it offers a vivid picture of the early reactions to that technology among EU regulators.

A. Background on ChatGPT

Open.AI launched ChatGPT in November 2022.[12] ChatGPT is a Large Language Model (LLM), which is a class of AI technologies that are based on deep-learning technology that is a further development of Machine Learning

[11] Press Release, European Parliament, MEPs Ready to Negotiate First-Ever Rules for Safe and Transparent AI (June 14, 2023), https://www.europarl.europa.eu/news/en/press-room/20230609IPR96212/meps-ready-to-negotiate-first-ever-rules-for-safe-and-transparent-ai.

[12] *Introducing ChatGPT*, OpenAI (Nov. 30, 2022), https://openai.com/blog/chatgpt.

(ML) technology. It may be accessed on the web either as a free or fee-based service and generates text in response to user inputs.[13]

What makes ChatGPT so immediately impressive is that users with no experience with AI can often obtain quite coherent and informative responses to quite complex questions, such as a request for a simple explanation of quantum computing or an interpretation of a poem. ChatGPT can provide answers tailored to mimic the style of an author with an emotional tone. Users can also ask ChatGPT to regenerate a response and it will come forth with a new and often equally useful but different response, and because ChatGPT keeps track of earlier inputs,[14] ChatGPT can generate contextually relevant conversations.

Those conversations can highlight both the capabilities and limitations of ChatGPT. For example, a *New York Times* reporter exploring ChatGPT's capabilities reported his conversation that led to ChatGPT declaring its love for the reporter.[15] The fact that ChatGPT could appear to express emotions might seem to many like an eerily human behavior not generally thought feasible by a machine. Of course, the article, and others like it, highlight how easy it is to identify such behavior as a poor simulacrum of an actual human emotional interaction.

Additionally, the knowledge data set that was used to train the ChatGPT model was cut off as of September 2021, so it cannot give responses that depend on knowledge or events that occurred after that time, which highlights a fundamental limitation of ML models—they are only as good as their training data sets. If the world changes and renders their training data no longer relevant for the new context, then an ML can give predictions to which it might ascribe high precision even though those predictions might be completely wrong. Additionally, AI works on patterns of text observed in its training data set, but there is no guarantee that the responses it generates will be true. For example,

[13] See OPENAI PLATFORM, https://platform.openai.com/apps (last visited July 7, 2023), which provides access to three OpenAI apps: ChatGPT (language conversational interface), Dall-E (image generating interface), and API (for integrating OpenAI into other business software applications).

[14] It is also worth noting that ChatGPT uses its interactions with users to build its knowledge base and refine its answers. When ChatGPT generates or regenerates a response, it gives users an opportunity to comment on the quality of the response. This information can be collected from all users to allow the model to "learn"—in the sense that it will be able to provide answers that are perceived to be better by more users which may or may not be the relevant standard for objectively measuring whether the answers are, indeed, better.

[15] *See* Kevin Roose, *A Conversation with Bing's ChatBot Left Me Deeply Unsettled*, N.Y. TIMES (Feb. 16, 2023), https://www.nytimes.com/2023/02/16/technology/bing-chatbot-microsoft-chatgpt.html.

when asked to list papers written by a particular person, it may list papers that do not exist and were not written by the author even though they sound like papers that might have been written by the author. That is, ChatGPT is capable of creating new information based on the collection of information on which it has been trained in the past, including the history of the conversation with the user. Additionally, ChatGPT is programmed to follow "ethical guidelines," so it will not directly offer textual responses that could be used to defame, libel, or otherwise be deemed unlawful.

Artificial General Intelligence (AGI) is sometimes distinguished as being a distinct subset of AI technologies, most of which are focused on narrow, specific problems. Distinct categories of AI technologies include Decision Support Systems (e.g., expert systems that seek to mimic the decision-making capabilities of domain experts such as interpreting diagnostic data for cardiac illness; analyzing natural resource monitoring data, etc.,); Robotic Systems; Natural Language Processing (NLP); and Computer Vision (CV). More recently, significant progress has been made in the development of Machine Learning (ML) systems that use large data samples to develop predictive models in specific problem domains. ML systems have evolved from those based on structured data to ones capable of extracting insights from unstructured data, making use of neural net technologies and deep-learning algorithms.

The nuanced differences in the underlying AI technologies have relevance for their regulatory implications. The LLM that ran ChatGPT when it was released in November 2022 was GPT-3, which was the third generation of a software program that was first launched as GPT-1 in 2018.[16] That program was trained on an initial data set of web pages, books, and other textual material. Successive versions of GPT were trained on significantly larger and more diverse data sets spanning a wider breadth of textual source material. ChatGPT now runs on GPT-4, which was released in March 2023 and represents a

[16] GPT is an acronym coined by OpenAI which stands for Generative Pre-trained Transformer. The GPT works by feeding a massive amount of training data into a neural network that takes inputs and transforms them into outputs. The neural net is a computational algorithm that may be modeled as a set of connected nodes with different weights. The weights are adjusted so that the trained data set of sampled task inputs and associated outputs are matched. Once trained, the neural net can be used to predict the output from previously unseen, new input data. For example, given a large data set of insurance applications with a large collection of data about the characteristics of the applicants and the type of coverage they selected, as well as the results of those policies, a neural net can be trained to predict the likely default risk for a new candidate with characteristics of the sort included in the training data set. The input data could also be medical imaging data and doctor diagnoses of tumors. A neural net trained on such data can analyze new images and diagnose tumors, potentially with better accuracy and faster than human doctors.

significant expansion in the range of data and enhancements in the capabilities of the deep-learning ML system that is the engine that drives ChatGPT.[17] What this demonstrates is the rapid pace of innovation that has characterized the evolution of the LLM AI, illustrating what appears to be an exponential acceleration in the capabilities of the software relative to the much slower pace of improvements that characterized automated BOT conversation tools in the past (e.g., those associated with automated call-response systems and other computer-generated response systems). We say "appears" because the real improvements are harder to assess, but a major concern with AGI is that it has the capability of self-improvement.

B. The 2023 Amendments to the EU AI Act

One of the major changes made to the draft EU AI Act in June 2023 that appeared to be driven largely by the arrival of ChatGPT was the addition of language defining "foundation models" and "general purpose AI systems." The amendments adopted in 2023 define a "foundation model" as "an AI system model that is trained on broad data at scale, is designed for generality of output, and can be adapted to a wide range of distinctive tasks" and defines a "general purpose AI system" as "an AI system that can be used in and adapted to a wide range of applications for which it was not intentionally and specifically designed."[18] Both of these definitions seem aimed, at least to some extent, at trying to capture the general use nature of programs like ChatGPT, which can be applied in a wide variety of different contexts. This is of particular importance for a regulation like the proposed EU AI Act, in which all the risk tiers for AI systems are predicated on the application areas of those systems. A program like ChatGPT that can be applied in a variety of different contexts could be difficult to classify according to a specific application.

One possible option for dealing with this ambiguity about ChatGPT's risk level—one that was apparently considered by EU regulators—was explicitly classifying generative AI models as high-risk systems. Instead of doing that, however, the European Parliament ultimately chose to adopt a tiered approach to general purpose AI (GPAI) systems, requiring all non-open source GPAI

[17] *See* Fawad Ali, *GPT-1 to GPT-4: Each of OpenAI's GPT Models Explained and Compared*, MAKE USE OF (Apr. 11, 2023), https://www.makeuseof.com/gpt-models-explained-and-compared/.

[18] *Amendments Adopted by the European Parliament on 14 June 2023 on the Proposal for a Regulation of the European Parliament and of the Council on Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act) and Amending Certain Union Legislative Acts* amend. 169, COM (2021) 0206 (June 14, 2023), https://www.europarl.europa.eu/doceo/document/TA-9-2023-0236_EN.pdf. [hereinafter *June 2023 Amendments*].

systems to adhere to transparency requirements and copyright law while subjecting GPAI systems that pose "systemic risk" —defined as models trained with computing power above 10^{25} floating point operations (FLOPS) — to additional obligations.[19] These additional obligations include model evaluations, assessment and mitigation of systemic risks, adversarial testing, reports to the Commission of serious incidents, cybersecurity protection, and reports on energy efficiency.[20] Until the EU develops harmonized standards, GPAI models with systemic risk may comply with these requirements by adhering to codes of practice "developed by industry the scientific community, civil society, and other stakeholders."[21]

The 2023 text of the draft law also highlights the range of applications for these types of AI systems, with Amendment 99 noting that foundation models are "designed to optimize for generality and versatility of output" and are "often trained on a broad range of data sources and large amounts of data to accomplish a wide range of downstream tasks, including some for which they were not specifically developed and trained."[22] Moreover, the amended text states, "each foundation model can be reused in countless downstream AI or general purpose AI systems."[23] Clearly, one of the elements of ChatGPT that most confounded and complicated the regulatory issues for EU regulators was the idea that it could be used in so many different ways by so many different systems.

This range of applications meant that it was difficult for regulators to classify ChatGPT and other LLMs either as high risk or not: In some applications they might, indeed, be used for high-risk purposes, but in other contexts, they might be generating text for entirely benign purposes. The balance EU regulators struck in the amended draft bill approved in 2023 was essentially to make the owners and operators of these programs responsible for risk mitigation across all of the application areas for which their services were used unless they are willing to pass on all of their source code and information about how their model was trained. More specifically, Amendment 100 of the approved draft, states:

[19] FLOPS is a measure of computing power based on the number of calculations that a processor can perform in a second. For reference, media reports indicate that ChatGPT 3.5 was trained on a computer that could perform 10^{24} FLOPS. *See* Zosia Wainat et al., *EU to Put Extra Guardrails on AI Foundation Models like GPT-4*, SIFTED (Dec. 7, 2023), https://sifted.eu/articles/foundation-model-eu-ai-act.

[20] European Parliament Press Release PR 15699, *supra* note 7.

[21] European Commission Press Release IP/23/6473, *supra* note 7.

[22] *June 2023 Amendments*, *supra* note 18, amend. 99.

[23] *Id.*

> In the case of foundation models provided as a service such as through API [Application Programming Interface] access, the cooperation with downstream providers should extend throughout the time during which that service is provided and supported, in order to enable appropriate risk mitigation, unless the provider of the foundation model transfers the training model as well as extensive and appropriate information on the datasets and the development process of the system or restricts the service, such as the API access, in such a way that the downstream provider is able to fully comply with this Regulation without further support from the original provider of the foundation model.[24]

This approach essentially offers developers of LLMs and other "foundation models" the choice between giving users direct access not just to the outputs of their model but also to its algorithms and training processes or else bearing responsibility for "appropriate risk mitigation" for any "downstream provider" to which it grants access. The former approach is unlikely to be feasible for most companies since their entire business model often derives from having developed proprietary AI systems. Indeed, part of the point of granting access to those systems through APIs rather than selling the system code itself is often to avoid sharing the code base and inner workings of the AI system while still being able to sell its results. But the latter option—assuming responsibility for risk mitigation for all downstream providers—also places a potentially heavy burden on AI companies, requiring them to monitor everyone who uses their services and build in appropriate safeguards as needed, depending on the application areas those users are working in. Such an approach would make it difficult for an AI company to open access to their APIs freely to any interested user for fear that they might not be able to provide appropriate risk mitigation services to everyone who took advantage of their services. It is not entirely clear what it would mean for a company to try to take the third option presented by restricting their services and API access "in such a way that the downstream provider is able to fully comply with this Regulation without further support from the original provider of the foundation model," but it would presumably require users to have some access to the model's testing and auditing features in order to comply with the requirements of the proposed EU AI Act around risk assessment and mitigation.

Amendment 101 of the proposed 2023 EU AI Act text offers more specific clues as to what kinds of things the owners and operators of foundation models

[24] *Id.* amend. 100.

must do to comply with the law. In particular, it states that organizations developing foundation models "should assess and mitigate possible risks and harms through appropriate design, testing and analysis, should implement data governance measures, including assessment of biases, and should comply with technical design requirements to ensure appropriate levels of performance, predictability, interpretability, corrigibility, safety and cybersecurity and should comply with environmental standards."[25] In many ways, these obligations mirror the ones that the Act requires of the owners and operators of high-risk AI systems, including the specified obligations for foundation model operators to "prepare all necessary technical documentation for potential downstream providers to be able to comply with their obligations under this Regulation."[26] There is also an explicit transparency requirement for generative foundation models that they include some notification or label about "the fact the content is generated by an AI system, not by humans."[27] Documentation and transparency are central to the requirements the Act puts in place for designated high-risk AI systems, but the amendment goes to great lengths to distinguish between those requirements and the ones for foundation models, stating: "These specific requirements and obligations do not amount to considering foundation models as high risk AI systems, but should guarantee that the objectives of this Regulation to ensure a high level of protection of fundamental rights, health and safety, environment, democracy and rule of law are achieved."[28]

The tension in the EU's compromise over LLMs is evident here: the desire to apply many of the same safeguards to technologies like ChatGPT that are required of high-risk AI systems, while at the same time insisting that doing so does not constitute classifying those technologies as high risk. This tension arises from the generality of these models and the inability to clearly classify all of their uses and applications as either high risk or low risk. Helberger and Diakopoulos have argued that the necessity of classifying AI systems according to their level of risk makes the proposed EU AI Act's framework ill-suited to regulating generative AI systems like ChatGPT both because it is not feasible to sort such systems into high/no high-risk categories and because it is too difficult to predict their future risks. Writing before the June 2023 approval of the new draft law, they argued that under the 2021 draft, the regulations for

[25] *Id.* amend. 101.

[26] *Id.*

[27] *Id.*

[28] *Id.*

high-risk systems might "only take effect once the generative AI is being used in a high-risk area."[29] They continue:

> From the point of view of society and fundamental rights, this is too late. The whole point about generative AI as a general-purpose AI system is that because they can be used for so many different purposes, it is paramount to incentivise the providers of systems to think about the safety of these systems from the onset, starting with the difficult question of data quality.[30]

For this reason, they advocate creating a new category of general purpose AI systems in the Act and regulating those separately from the high-risk AI systems.

To a large extent, Helberger and Diakopoulos seem to get what they want from the 2023 revision of the proposed EU AI Act, in particular, their recommendation that general purpose AI systems should be "considered a general-risk category in their own right" rather than being categorized under the existing high-risk classification.[31] Somewhat perplexingly, this approach does not make clear how exactly the requirements that apply to foundation models differ from those that apply to high-risk applications beyond being somewhat vaguer and slightly more indirect because operators of foundation models are responsible for risk mitigation of their users' applications. Interestingly, the draft law's justification for treating foundation models in this special manner is couched not just in terms of the models' general applicability but also their "unpredictability." The approved draft states:

> Pre-trained models developed for a narrower, less general, more limited set of applications that cannot be adapted for a wide range of tasks such as simple multi-purpose AI systems should not be considered foundation models for the purposes of this Regulation, because of their greater interpretability which makes their behaviour less unpredictable.[32]

This notion that it is not just the generality of certain AI tools but also their lack of "interpretability" that makes them difficult to regulate in the same manner as more narrowly applied systems is notable because it seems to speak to the complexity and size of the data sets used to train these models rather than the broadness of their range of applications. The implication seems to be

[29] Natali Helberger & Nicholas Diakopoulos, *ChatGPT and the AI Act*, 12 INTERNET POL'Y REV. 1, 3 (2023).

[30] *Id.*

[31] *Id.*

[32] *June 2023 Amendments*, *supra* note 18, amend. 101.

that the large training data sets and high degree of complexity of these systems are somehow intrinsically linked to their generality and more significant than the complexity or size of more narrowly trained models.

There is some justification for this perspective. Indeed, whereas earlier ML models and narrow AI systems worked with structured datasets, the current frontier of deep-learning ML models that are core to AGI can work with unstructured data. Structured data is data that is collected and organized with a particular purpose in mind. For example, a structured image database might be a collection of pictures of trees with metadata labels that identify the tree and provide other salient information about the image. Then an ML can use such training data to develop its prediction capabilities, which, once trained, would allow the ML to identify a tree from an unlabeled tree picture.

Algorithms that can work with unstructured data could take in a wealth of image data and learn from those images how to identify trees without first being presented with the metadata that comes with structured data sets. This greatly expands the sorts of data a ML can make use of, but it can also make it much more difficult to determine how particular training data impacts the performance or predictions of the ML or to understand or explain the reasoning (in human intelligible form) that leads to the ML's forecast. Explaining the behavior of neural nets with multiple layers is difficult in all contexts, but explaining their behavior when applied to unstructured data is even more difficult.

Helberger and Diakopoulos have argued that the proposed EU AI Act is ill-suited to regulating generative AI models like ChatGPT because of those models' "dynamic context and scale of use."[33] They link the broad applicability of generative AI models to the scale of those models' training data, writing:

> Generative AI systems are not built for a specific context or conditions of use, and their openness and ease of control allow for unprecedented scale of use. The output of generative AI systems can be interpreted as media (text, audio, video) by people with ordinary communication skills, lowering, therefore, significantly the threshold of who can be a user. And they can be used for such a variety of reasons to some extent because of the sheer scale of extraction of data that went into their training.[34]

Still, in the context of the proposed EU AI Act's justification for singling out foundation models for special treatment, it is not immediately clear why

[33] Helberger & Diakopoulos, *supra* note 29, at 2.

[34] *Id.*

broadly applicable AI systems should always necessarily be trained on more data or be more complicated or less predictable than more narrowly targeted ones. In fact, this justification seems to suggest that in carving out special requirements for foundation models, the European Parliament has actually conflated two separate characteristics of ChatGPT—its broad applicability and its complexity—that need not always be linked. A simple AI model could apply in a range of different contexts, while an extremely complicated one trained on large amounts of data might be designed for only a very narrow application but still be difficult to interpret or predict.

EU regulators seem to have deliberately avoided splitting these two characteristics in the proposed EU AI Act text, where foundation models are defined as being both "trained on broad data at scale" and "designed for generality of output."[35] The definition of general purpose AI, on the other hand, makes no reference to training data or size and merely depends on a system being able to "be used in and adapted to a wide range of applications for which it was not intentionally and specifically designed."[36] This conflation undermines the logic of the special foundation model-specific provisions in the draft law, suggesting that they are necessary not because of how general they are but because of how complex they are when, in fact, it seems more likely that the reverse is true since the existing provisions for high-risk AI systems can apply regardless of system complexity and size. Surprisingly, however, the provisions in the proposed EU AI Act call out the operators of foundation models, not general purpose AI models, for special requirements and responsibilities, suggesting that regulators care more about the size and complexity of such systems than their broad range of applications. One possible explanation for this is provided in Amendment 102, which highlights concerns about the training data used for "generative AI systems" based on foundation models and the potential for such systems to violate copyright laws by exploiting publicly available, copyrighted text.

These concerns about protecting copyright and reining in training data may be one reason that regulators have sought to couch oversight of foundation models in terms of the size and complexity of their training data sets rather than just their broad applicability. In practice, however, it seems likely that these concerns will be largely sidestepped. Requiring models like ChatGPT that were built on copyrighted text to be abandoned or requiring their developers to incur the costs of acquiring consent from individuals seems largely unrealistic.

[35] *June 2023 Amendments*, *supra* note 18, art. 3.

[36] *Id.*

III. Reconciling General Purpose AI with GDPR

Because AI systems depend on the ability to access and use a variety of data for a range of purposes, the EU's General Data Protection Regulation (GDPR), in addition to the proposed EU AI Act, has significant implications for their development. Moreover, the shift in the EU AI Act's scope to encompass more general AI has the potential to raise even more significant GDPR-related concerns because so much of the focus of GDPR requirements is on establishing a purpose for the collection of data and notifying users of that purpose at the time of collection as well as minimizing collected data. This emphasis on identifying the purpose of collected data prior to its collection and data minimization raised concerns about GDPR's impacts on AI long before the release of the draft EU AI Act or ChatGPT.

For instance, writing in 2017, following the passage of GDPR, Zarsky wrote that the law was "incompatible with the data environment that the availability of Big Data generates."[37] In an analysis focused more specifically on the compatibility of GDPR and AI models, Kesa and Kerikmäe argued that "it is not possible to practice complete compliance with the requirement to ensure certain rights as they are guaranteed by GDPR, especially by data processors employing complex machine-learning systems that process vast amounts of data through complex multistep sets of instructions."[38] Other analyses, such as one carried out by the Scientific Foresight Unit of the European Parliamentary Research Service, have concluded that while GDPR may not be incompatible with the development of AI, it "does not provide sufficient guidance for controllers" and it may need to be "expanded and concretised" to offer clearer explanations of how its provisions apply to AI systems.[39] Yet another analysis posits that GDPR "may eventually help create the trust that is necessary for AI acceptance by consumers and governments."[40]

The recent responses of EU data protection authorities (DPAs) to ChatGPT eloquently illustrate the difficulties that general purpose AI faces in complying with GDPR. For example, on March 31, 2023, the Italian data

[37] Tal Zarsky, *Incompatible: The GDPR in the Age of Big Data,* 47 Seton Hall L. Rev. 995, 996 (2017).

[38] Aleksandr Kesa & Tanel Kerikmäe, *Artificial Intelligence and the GDPR: Inevitable Nemeses?*, 10 TalTech J. Eur. Stud. 67, 71 (2020).

[39] Giovanni Sartor, *The Impact of the General Data Protection Regulation (GDPR) on Artificial Intelligence*, Eur. Parliamentary Rsch. Serv. (June 2020), https://www.europarl.europa.eu/RegData/etudes/STUD/2020/641530/ EPRS_STU(2020)641530_EN.pdf.

[40] Kalliopi Spyridaki, *GDPR and AI: Friends, Foes or Something in Between?*, SAS Eur., https://www.sas.com/en_us/insights/articles/data-management/gdpr-and-ai--friends--foes-or-something-in-between-.html (last visited Nov. 22, 2023).

protection authority issued a temporary emergency decision banning ChatGPT based on four potential GDPR violations.[41] The DPAs from France, Spain, Ireland, and Germany have begun taking preliminary steps of their own, and the European Data Protection Bureau has launched a task force to foster cooperation among Member States in their potential responses to ChatGPT.[42] Although changes made by OpenAI led the Italian DPA to permit service to resume on April 28, those changes did not appear to address all of the concerns raised in the initial ban.[43]

The purpose of this section is not to reiterate all the ways in which GDPR's provisions may hinder (or support) the development of AI systems but rather to consider the specific ways in which general purpose AI—or the "foundation models" defined by the draft EU AI Act—may be especially difficult to reconcile with GDPR by virtue of its generality. Additionally, we aim to evaluate how the different structures and framings of GDPR and the proposed EU AI Act influence their respective enforceability and what lessons, if any, may be drawn from GDPR in designing the EU AI Act.

A. Background on GDPR

The General Data Protection Regulation (GDPR) is the EU law designed to provide a framework for protecting personal data online.[44] It was enacted in 2016 and became law in 2018. It is part of the EU privacy and human rights law. Work on drafting GDPR began in 2011, and when passed, GDPR superseded the EU's 1995 Data Privacy Protection Directive.[45] The basic goal of the framework remained the same, namely, to protect the privacy of individual data online. The 1995 Directive in turn had been crafted from the 1980 OECD Recommended Guidelines for Privacy Protection,[46] which were updated in 2013 and based on seven key principles: (1) Notice: individuals

[41] Matt Burgess, *ChatGPT Has a Big Privacy Problem*, WIRED (Apr. 4, 2023, 12:00 PM), https://www.wired.com/story/italy-ban-chatgpt-privacy-gdpr/.

[42] Toby Sterling, *European Privacy Watchdog Creates ChatGPT Task Force*, REUTERS (Apr. 13, 2023), https://www.reuters.com/technology/european-data-protection-board-discussing-ai-policy-thursday-meeting-2023-04-13/.

[43] Natasha Lomas, *ChatGPT Resumes Service in Italy after Adding Privacy Disclosures and Controls*, TECHCRUNCH (Apr. 28, 2023, 2:57 PM), https://techcrunch.com/2023/04/28/chatgpt-resumes-in-italy/.

[44] *See GDPR*, *supra* note 10.

[45] Council Directive 95/46, 1995 O.J. (L 281) _ (EC), https://eur-lex.europa.eu/legal-content/EN/LSU/?uri=celex: 31995L0046.

[46] Michael D. Birnhack, *The EU Data Protection Directive: An Engine of a Global Regime*, 24 COMPUT. L. & SEC. REV. 508 (2008).

should be notified when their personal data is collected; (2) Minimal & Consensual: collection should be with consent and limited; (3) Purpose: Data collection should be relevant, up to date, for specific purpose; (4) Shared: Not disclosed without individual consent unless for legal reasons; (5) Protected: security in place. (Privacy by design, by default); (6) Transparency: individuals have right to know what data is collected about them; and (7) Controllers of data accountable for complying with law.[47]

GDPR, however, represented a significant expansion in capabilities and requirements to provide a framework for stronger enforcement of online privacy. The intent of GDPR is to enhance individuals' control and rights over their personal data and to simplify the regulatory environment for international business by replacing what previously was a mishmash of different national and sector-specific privacy frameworks with a homogeneous framework applicable across all EU member states and data contexts where individually identifiable data may be used. It applied to all firms, and so its coverage and scope were broad.

GDPR expanded individual privacy rights by including the right to be forgotten, the right to correct on-line data (including deleting it from on-line files), portability rights (e.g., to enable end users to move their data to another business), and expanded transparency obligations to enable individuals to learn what their personal data is being used for, potentially by third parties. While it is conceptually clear how a firm might identify and share with an individual all the personal data associated with that individual so long as that data is structured (e.g., stored in database records that can be linked to the individual), it is unclear how these rules might be implemented for a business using unstructured data.

GDPR sought to enshrine privacy-by-design and privacy-by-default in business practices. It includes a number of special features, such as requiring Data Protection Impact Assessments (DPIA), which may be requested to determine whether data is being managed consistent with GDPR and to provide guidance to regulatory decision-making, including the adjudication of fines. DPAs were empowered in each EU member state to monitor and enforce compliance with GDPR. The potential fines that can be levied under GDPR are significant, running as high as the larger of 20 million Euros for a violation or up to 4% of the annual global revenue for a violator. As of June 2023, fines

[47] *See* Brian Daigle & Mahnaz Khan, *The EU General Data Protection Regulation: An Analysis of Enforcement Trends by EU Data Protection Authorities*, J. INT'L COM. & ECON. 1, 4 (2020). For a history of GDPR and enforcement actions undertaken under the Act, see *The History of the General Data Protection Regulation*, EUR. DATA PROT. SUPERVISOR (2018), https://edps.europa.eu/data-protection/data-protection/legislation/history-general-data-protectio n-regulation_en.

totaling over 4 billion Euros had been assessed against firms for violating GDPR, with 86% of those fines directed against Meta (Facebook, WhatsApp), Google, and Amazon.[48] Although these numbers are sizable, the compliance costs incurred globally—by EU businesses and businesses that share data with EU businesses, which includes most multinational enterprises—are likely several multiples higher. However, it is debatable how much of this additional spending may be sound cybersecurity investment as opposed to excess investment, incurred incrementally due to GDPR.

B. A Risk-Based Approach Versus a Rights-Based Approach

The proposed EU AI Act and GDPR center on very different approaches to technological regulation. While GDPR takes a rights-based approach to regulating data protection by articulating specific rights to which individuals are entitled regarding each transaction involving their data, the proposed EU AI Act frames its rules around a risk-based approach that imposes requirements on the owners and operators of AI systems according to the level of risk those systems pose.[49] Although the specific requirements that would apply to high-risk systems span forty-six articles,[50] the list of questions and answers that the European Commission issued when it proposed the Act summarized it as imposing five requirements: (1) quality data and data governance; (2) transparency for users; (3) human oversight; (4) accuracy, robustness, and cybersecurity; and (5) traceability and auditability.[51]

The regulatory regime imposed by GDPR poses significant and perhaps insurmountable problems for general purpose AI. To be successful, AI models require vast quantities of training data. Although OpenAI has not fully disclosed the sources of the data used to train GPT-4, a 2023 technical report reveals that it was trained on “publicly available data (such as internet data) and data licensed from third-party providers.”[52] This representation suggests that GPT-4 is in serious conflict with the rights-based approach of GDPR, which envisions data subjects possessing the right to control the processing of their personal data. Under GDPR, what matters is whether the data is *personal*, not whether it is *public* or *private*.

[48] *See GDPR Enforcement Tracker*, https://www.enforcementtracker.com/?insights (last visited June 28, 2023).

[49] *Proposed EU AI Act*, *supra* note 3, recital 14, art. 1(b).

[50] *Id.* arts. 6-51.

[51] European Commission, Artificial Intelligence – Questions & Answers (Dec. 14, 2023), https://ec.europa.eu/commission/presscorner/detail/en/QANDA_21_1683.

[52] OpenAI, *GPT-4 Technical Report* (Mar. 27, 2023), https://cdn.openai.com/papers/gpt-4.pdf.

Consider GDPR's requirement, laid out in Article 6, that any processing of personal data falls within one of six legal justifications.[53] Of these, three largely fall by the wayside with respect to training data: it is hard to characterize the use of data to train an AI model as necessary for compliance with one of the controller's legal obligations, necessary to protect a person's vital interests, or "necessary for the performance of a task carried out in the public interest or in the exercise official authority vested in the controller."[54]

A fourth justification authorizes processing that is "necessary for the legitimate interests pursued by the controller or a third party except where the interests are overridden by the interests or fundamental rights and freedom of the data subject."[55] Recital 47 indicates that such a legitimate interest exists "where there is a relevant and appropriate relationship between the data subject and the controller . . . such as where the data subject is a client or in the service of the controller" and where "a data subject can reasonably expect at the time and in the context of the collection of the personal data that processing for that purpose may take place."[56] Recital 69 emphasizes that data subjects whose data are processed under justification retain their right to object to the processing.[57] The lack of a client relationship between the AI provider and the data subjects whose information is contained in the training data, the lack of expectation in most cases that the data subject's personal data would be used as training data, and the difficulties in giving people control to data subjects whose data was scraped off of the public Internet essentially render this fourth justification inapplicable to general purpose AI.

The two remaining justifications require specific agreement of the data subject: consent or processing necessary for the performance of a contract, with GDPR further specifying that any consent by affirmative opt-in consent be in writing that grants specific consent to each form of processing. Those collecting personal data must also disclose certain information about the data and collect only the minimum amount necessary for the purposes of the processing.

These provisions reveal the deep tension between GDPR and general purpose AI. The rights-based approach reflected in GDPR, which requires an agreement or written consent from and specific disclosures to every data subject whose personal data was included in the data used to train ChatGPT, is inconsistent with the Internet scraping used to generate the large amounts of

[53] *See GDPR*, *supra* note 10, art. 6.
[54] *Id.*
[55] *Id.* art. 6(1)(f).
[56] *Id.* recital 47.
[57] *Id.* recital 69.

training data on which general purpose AI depends. The fact that general purpose AI expends considerable effort to extract features from unstructured data means that people whose personal data was contained in public websites contained in the training data would have to have agreed to permit that type of feature extraction in order for it to comply with the specificity required to constitute valid consent. Nor is it likely that developers of general purpose AI provide people whose data are included in their training data with the other rights mandated by GDPR, including access, rectification, erasure, restriction of processing, data portability, withdrawal of consent, and the right to object. Should a data subject exercise one of these rights, the controller would have to expunge the impact of that data from the algorithm or delete the algorithm altogether.

Moreover, the rights-based approach reflected in GDPR regulates data as an *input* into processing, whereas the risk-based approach reflected in the proposed EU AI Act focuses on the uses of the products of AI as *outputs*. Both of these framings make sense in the context of the motivations for these different regulations and the policy contexts in which they were developed, but neither is well-suited to the development of general purpose AI. For the proposed EU AI Act, the central problem posed by general purpose AI is that it is nearly impossible to assess the risks associated with it in any meaningful way. For GDPR, the key problem is that it is nearly impossible to perform any meaningful purpose limitation (or data minimization) for data used to train general AI systems. It is hard to imagine how AI system operators could feel any real confidence that they were in compliance with either law if they wanted to develop a general AI system using European personal data.

However, in terms of enforcement, GDPR appears to offer a much clearer path to imposing penalties on AI system operators than does the proposed EU AI Act. To show that an AI system operator is in violation of GDPR, a data protection authority need only demonstrate that personal information has been used to train an AI model without adequate lawful basis. It is much less clear what a regulator would have to demonstrate about a foundation model operator to show that they had failed to meet the bar of providing risk mitigation to their downstream providers. This reliance on self-regulation in the draft EU AI Act is perhaps the result of the complexity of the AI systems that the EU AI Act seeks to regulate—it may be that regulators have handed over responsibility for meeting their requirements to AI system owners and operators precisely because they fear they do not possess the requisite technological expertise to do so themselves. If this is in fact the case, then general AI systems and foundation models are likely to only exacerbate that problem.

This speaks to a broader weakness in the proposed EU AI Act, namely its reliance on AI companies to define their own risk mitigation, transparency, and explainability standards and documentation. This is true not just for foundation models but also of all the high-risk AI systems that the draft law proposes to regulate. All regulations, including GDPR, have some ambiguity and room for interpretation in their provisions, so it is perhaps unnecessary to assume that the proposed EU AI Act law will suffer from a lack of clarity any more than any other piece of legislation. But from an enforcement standpoint, it is perhaps notable that many of the largest fines that European regulators have imposed under GDPR have been for incidents that were already punishable under the previous Data Protection Directive—namely, data breaches, rather than privacy violations that only became violations under GDPR.[58] That suggests that regulators may be slower to enforce new provisions in the proposed EU AI Act than to rely on existing provisions, like those already enshrined in GDPR, with which they are already familiar.

IV. CONCLUSION

Regulators in the EU as well as other countries are still in the early stages of their efforts to frame policy framework for AI, but the impact of ChatGPT makes clear that these efforts can be easily and significantly influenced by the emergence of new, popular AI technologies. Interestingly, in the case of the proposed EU AI Act, this impact resulted in the creation of two new regulatory categories of AI systems (foundation models and general purpose AI systems) but did not appear to significantly alter the provisions applying to those categories. Put another way, the European Parliament chose not to classify foundation models like ChatGPT as high-risk AI systems but still imposed roughly the same set of requirements on the operators of those models as it does on the operators of high-risk AI systems. This suggests that while regulators may be struggling to come to terms with the broad range of risks presented by AI and the variety of AI systems and applications they want to address, they do not have a very diverse set of regulatory mechanisms or proposals to use to address those risks. This undermines, to some extent, the risk-based framework of the proposed EU AI Act, suggesting that rather than

[58] Josephine Wolff & Nicole Atallah, *Early GDPR Penalties: Analysis of Implementation and Fines Through May 2020*, 11 J. INFO. POL'Y 63, 94 (2021).

tailoring rules to different risks, regulators are instead merely designating which systems they believe are sufficiently impactful or important to merit regulation.

At the same time that ChatGPT has highlighted the ambiguity and gaps in the draft EU AI Act's designation of high-risk AI systems, it has also underlined the challenges for AI posed by GDPR, particularly that regulation's emphasis on purpose limitation and data minimization. While the broad applicability and large training data sets used by generative AI models make them harder to regulate under the EU AI Act, those same characteristics also make it, in some ways, easier for regulators to go after the owners and operators of those models using GDPR since by their very nature, such models do not rely on minimized or purpose-limited data sets. It is possible that regulators seeking to impose penalties on AI companies will therefore come to rely more on GDPR than the AI Act, ultimately rendering the latter more of a symbolic piece of regulation than a law that is actively enforced.