
Community Norms as Self-Regulation of Generative AI in Creative Industries

Peter Hall^{*1} Betty Li Hou^{*1} Falaah Arif Khan^{*1}

1. Introduction

History has seen many instances of technologies that provided major enhancements to industry work through automation. These developments inevitably came with regulatory concern, such as, what should effective regulation look like to ensure that the technology is used productively and beneficially, and that the societal harms introduced by it are mitigated? The recent meteoric rise of generative AI (Gen AI) has raised concerns on how creative industries will be affected, as these models offer an automated method of generating artistic works with unprecedented speed and ease, due to being scalable, inexpensive, and widely accessible. The question of effective regulation arises again, as policy-makers seek to balance the interests of working creatives and the conditions conducive to the rapid advancement of this exciting new technology.

In the years following the release of ChatGPT, it has become the prevailing public sentiment that the AI of today is a novel, unprecedented technology, but in this work we ask: in what way exactly? We analyze historical cases of automation, including the **printing press**, which automated transcription; the **camera**, which automated portrait painting; **sampling and digital audio technology**, which automated the playing of instruments and arrangement of different musical elements; **search engines**, which automated information retrieval; **photo-editing software**, which automated the manipulation and enhancement of visual imagery; and **3-D printing**, which automated small-scale fabrication, and construct a taxonomy of dimensions along which generative AI is similar (Section 3) and different (Section 4) to prior forms of automation in creative industries.

Further, we study the regulatory approaches that were employed in each of these past cases, and evaluate what the regulatory approach for Gen AI should be, in light of the similarities and differences from our taxonomy. Our analysis reveals community norms to be an existing, powerful form of regulation in creative industries, which has so far re-

^{*}Equal contribution, listed alphabetically ¹New York University. Correspondence to: Falaah Arif Khan <fa2161@nyu.edu>.

Proceedings of the 41st International Conference on Machine Learning, Vienna, Austria. PMLR 235, 2024. Copyright 2024 by the author(s).

ceived less regulatory capture than the push for new legislation such as the EU AI Act and the AI Bill of Rights. While community norms are less formally defined (in that they are unwritten) and are highly community-specific, for example conventions around the use of other artists' work are vastly different in music (sampling), literature (fan fiction) and film-making (referencing iconic scenes from other films), they have shaped the adoption of new tools in the creative workflow, and generative AI, despite all its bells and whistles, is just another tool at the dispense of the artist, albeit with the potential to transform the creative process. Hence, while we wait for new legislation to be drafted, debated, passed and enforced, we should empower creatives to lead the way forward with community-driven self-regulation.

2. What counts as “automation”?

The impact of automation on labor has been of sustained interest to the academic community, mainly economists (Acemoglu & Restrepo, 2019a; Acemoglu et al., 2022; Acemoglu & Restrepo, 2022; Autor & Salomons, 2019; Acemoglu & Johnson, 2023; Walkowiak & MacDonald, 2023). These works define ‘automation’ quite broadly, including mechanization of agriculture by harvesters and ploughs, to industrial robots and automated machinery, to more recently, automated software that replaces white-collar workers. While our work is situated within this larger discussion of the impact of automation on labor, we limit our discussion to historic cases of automation in **creative industries**, such as the digitalization of books, audio and video (Potts, 2011).

3. What’s the same?

Previous cases of automation in creative domains such as the printing press, camera and photo-editing software enabled the development of new creative works and effectively changed the labor landscape of creative industries such as visual art, music, and journalism — much like generative AI is premised to (Bordàs Vives, 2023; Walkowiak & Potts, 2024; Walkowiak & MacDonald, 2023).

Many of these cases were met with regulatory and legal concern, such as, the copyright and intellectual property rights concerns raised in *Burrow-Giles Lithographic Co. v. Sarony*, *Grand Upright Music, Ltd. v. Warner Bros. Records Inc.*

and *Kelly v. Arriba Soft Corp.* (Bur; Gra; Kel). Despite this, technological advancements have in many cases become widely-used methods and tools for creative processes, some even giving rise to entirely new artistic fields such as photography and graphic design.

Regulatory approaches to automation in these cases have endorsed the new form of creation, attempting to incorporate them fairly into existing structures and avoid stifling innovation, for example: *Burrow-Giles Lithographic Co. v. Saroni* and *Kelly v. Arriba Soft Corp.* both ultimately endorsed the use of the automation with respect to copyright law. Policymakers have little to gain by limiting automation, as going too far can risk staying economically competitive on the world stage. Instead, governmental initiatives can attempt to *promote* automation by way of promoting innovation. The patent system, for example, aims to promote developments of new technologies and designs, many of which are tools of automation.

4. What's different?

Gen AI is a novel form of automation in both the way it is constructed and the way it is used. We discuss these aspects below.

HOW IS THE AUTOMATION CONSTRUCTED?

Derivative nature. Previous forms of automation have been intrinsically referential, i.e. merely using or pointing to existing information. A traditional search engine, for instance, links to different websites and resources. Cameras capture a scene as it exists, with metadata providing specifics as to how and of what a photo is taken. On the other hand, generative AI models must draw from vast datasets by design, and in this way the process of automating (as well as the resulting system) is inherently derivative of its inputs—it creates new outputs by transforming or processing existing information.

While other forms of automation discussed have already been able to take in vast quantities of data, e.g., a search engine, this is at runtime of the automated process (when it is being executed). At construction (when it is being built), the amount of data needed is relatively small. In other words, pre-existing automated processes that handle large amounts of data have been dynamic, handling the data as the process is in action but requiring relatively little to set up. By contrast, Gen AI models require vast amounts of input for training, often including many terabytes of internet data. The increased computational and data requirements for building Gen AI models presents a fundamental shift from how previous forms of automated systems were designed and operate.

Opacity. Gen AI is built through a complex supply

chain (Lee et al., 2024) involving several different stakeholders. The developers of commercial Gen AI systems have been reluctant to disclose details such as the sources or even the size of the training dataset (Carlini et al., 2023). While this complexity and opacity is not unique to Gen AI systems (and was a challenge with search engines and social media platforms as well), the derivative nature of Gen AI coupled with this opacity compounds the regulatory challenges. For example, groups like The New York Times (NYT) and Getty Images (Get) have alleged that derivations were made from their intellectual property without proper compensation, but without sufficient disclosure it is difficult to rigorously validate such claims.

WHAT IS THE NATURE OF THE AUTOMATION?

Stochasticity. Traditional automated systems or tools have been deterministic in the following sense: Given a set of user-defined inputs (such as aperture, shutter speed and size, for instance), the machine produces a single output (e.g., a photograph). Gen AI, on the other hand, utilizes randomness to produce novel outputs, such as giving the user multiple choices (Adobe's generative fill, for example) or new responses on repeated queries (most chatbots). This presents a novel relationship between the automation and the users of it.

General-purpose. Traditionally, automated systems have been designed for a single, specific task, be it image capture, information retrieval, manufacturing tasks, or otherwise. Gen AI, on the other hand, is envisioned to be a general-purpose, adaptable form of automation with a variety of potential uses ranging from article summary (Liu & Lapata, 2019) to concept art (Cetinic & She, 2022), newswriting (Broussard et al., 2019), tutoring (Limo et al., 2023), voice acting (Niroula, 2023), general assistant work (Brown et al., 2020) and beyond. Labor effects of automated are popularly studied using a task-based approach (Walkowiak & Potts, 2024; Acemoglu & Restrepo, 2019a;b), and the fact that generative AI is (or is premised to be) proficient at several tasks, complicates the analysis of impending labor effects, and the necessary regulatory mechanisms to safeguard workers in creative domains.

HOW IS THE AUTOMATION PERCEIVED?

Accessibility. Many past forms of automation at their time of introduction have been limited in their immediate range of accessibility. This has historically provided a natural gatekeeping effect that in theory gives some lag-time to affected industries (see, e.g., (Savini & Savini, 2015) with respect to 3-D printing). This has allowed creatives to adapt and integrate these tools into their workflows, without replacing them outright. Gen AI, on the other hand, is touted

for its ease of use for all levels of expertise¹. The startup cost to using Gen AI for creative uses is small, sometimes as simple as constructing a good prompt (not considering the initial effort to train these models). By contrast, other forms of automation for commercial purposes in the past have demanded a higher investment, such as in the case of warehouse automation.

Anthropomorphism. Gen AI is commonly described in agentic language and anthropomorphized (Bender et al., 2021). This may suggest expectations of the system’s capabilities as an agent rather than a tool, and in turn affect perceptions of its artistic contributions relative to the human using it.

5. What’s at stake?

It is widely believed that Gen AI has the potential to impact all aspects of society. Working creatives have already expressed concerns of the impact of the technology on their livelihoods, notably actors (Drescher & Crabtree-Ireland, 2023), writers (Limbong, 2024), journalists (Darcy, 2024), and visual artists (Roose, 2022).

Affected groups may have to deal with a greater intensity and quantity of competition from automation than previous forms (Acemoglu et al., 2022). Broader labor effects in creative industries include economic losses for artists (Jiang et al., 2023; Epstein et al., 2023) and changes in the nature and quality of jobs (of the kind that occurred with warehouse automation (Weidinger et al., 2021)). There have additionally been claims made of a cannibalization of the creative market with Gen AI tools replacing the very artists whose intellectual property was trained on (Weidinger et al., 2021; Acemoglu et al., 2022).

6. Regulating Generative AI

Regulation of automation in creative industries has historically centered innovation over protection for artists. From our analysis in Section 4, Gen AI presents novel regulatory challenges given its derivative nature, reliance on user generated data, and scale of impact. It is thereby insufficient to follow the historical trend of taking a hands-off approach to regulation, as is starting to be emphasized in legal and policy discourse (Sobel, 2017). Attempts at regulating AI have been largely swayed by powerful technology corporations, who assert either that self-regulation is sufficient or that their expertise is essential for crafting new regulation to prevent unintended societal harms (Cusumano et al., 2021). We argue that, as a result, current discussions around regulation of Gen AI fail to center the perspectives of the most important stakeholders: working creatives. Specifically, we

suggest existing community norms in creative fields should serve as the basis for self-regulation of Gen AI.

Community norms are a compelling lens from which to view the regulatory landscape of Gen AI because they are already a strong mediating force in creative industries (Du et al., 2021). We use the term ‘community norms’ here in a broad sense to refer to the social standards and expectations in various artistic communities surrounding creative processes and labor, attribution, discourse, and other aspects of the social identity of individual artists (Malle, 2023). Note that norms are highly community-specific, for example conventions around the use of other artists’ work are vastly different in music (sampling), literature (fan fiction) and film-making (referencing iconic scenes from other films).

The integration of Gen AI into artistic processes is not inevitable, and its acceptability and adoption are reliant on community perception. Development and deployment must center the opinions of artists and consensus of artistic communities on symbiotic uses of Gen AI in the creative process. Mirowski et al. (Mirowski et al., 2024) and Shumakova, Lloyd, and Titova (Shumakova et al., 2023), for example, interviewed creative professionals on their perceptions and preferred uses of AI, revealing varying normative positions regarding use of systems trained on other artists’ intellectual property.

In conclusion, Gen AI is a novel form of automation with the potential to instigate dramatic changes within creative industries. It requires novel regulatory mechanisms, including a robust self-regulatory approach shaped by community norms, in order to safeguard artistic expression and the livelihoods of creative professionals (Shumakova et al., 2023; Jiang et al., 2023).

References

- Burrow-Giles Lithographic Company v. Saronoy, 111 U.S. 53 (1884). URL <https://supreme.justia.com/cases/federal/us/111/53/>.
- Getty Images, Inc. v. Stability AI, Inc., No. 1:23-cv-00135 (Del. 2023). URL <https://admin.bakerlaw.com/wp-content/uploads/2023/09/ECF-1-Complaint-1.pdf>.
- Grand Upright Music Ltd. v. Warner Brothers Records, Inc. et al., 780 F. Supp. 182 (S.D. N.Y. 1991). URL <https://law.justia.com/cases/federal/district-courts/FSupp/780/182/1445286/>.
- Kelly v. Arriba Soft Corporation, 336 F.3d 811 (9th Cir. 2002). URL <https://law.justia.com/cases/federal/appellate-courts/F3/336/811/468703/>.

¹OpenAI’s spring 2024 update “Introducing GPT-4o.”

- The New York Times Company v. Microsoft Corporation et al., No. 1:23-cv-11195 (S.D. N.Y. 2023). URL https://nytco-assets.nytimes.com/2023/12/NYT_Complaint_Dec2023.pdf.
- Acemoglu, D. and Johnson, S. *Power and Progress: Our Thousand-Year Struggle Over Technology and Prosperity*. PublicAffairs, New York, 2023.
- Acemoglu, D. and Restrepo, P. Artificial intelligence, automation, and work. In Agrawal, A., Gans, J., and Goldfarb, A. (eds.), *The Economics of Artificial Intelligence: An Agenda*, chapter 2, pp. 197–236. University of Chicago Press, 2019a. ISBN 978-0-226-61333-8. URL <http://www.nber.org/books/agra-1>.
- Acemoglu, D. and Restrepo, P. Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2):3–30, Spring 2019b.
- Acemoglu, D. and Restrepo, P. Tasks, automation, and the rise in U.S. wage inequality. *Econometrica*, 90(5): 1973–2016, September 2022.
- Acemoglu, D., Autor, D., Hazell, J., and Restrepo, P. Artificial intelligence and jobs — evidence from online vacancies. *Journal of Labor Economics*, 40 (1):293–340, 2022. doi: 10.1086/718327. URL <https://economics.mit.edu/sites/default/files/publications/AI%20and%20Jobs%20-%20Evidence%20from%20Online%20Vacancies.pdf>.
- Autor, D. and Salomons, A. New frontiers: The evolving content and geography of new work in the 20th century. Working Paper, 2019. URL <https://app.scholarsite.io/david-autor/articles/new-frontiers-the-evolving-content-and-geography-of-new-work-in-the-20th-century>.
- Bender, E. M., Gebru, T., McMillan-Major, A., and Shmitchell, S. On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT '21)*, pp. 610–623, New York, NY, USA, Virtual Event, Canada 2021. Association for Computing Machinery. doi: 10.1145/3442188.3445922. URL <https://doi.org/10.1145/3442188.3445922>.
- Bordàs Vives, A. Artificial intelligence and the creative industries. 2023.
- Broussard, M., Diakopoulos, N., Guzman, A. L., Abebe, R., Dupagne, M., and Chuan, C.-H. Artificial intelligence and journalism. *Journalism & mass communication quarterly*, 96(3):673–695, 2019.
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33: 1877–1901, 2020.
- Carlini, N., Hayes, J., Nasr, M., Jagielski, M., Sehwal, V., Tramèr, F., Balle, B., Ippolito, D., and Wallace, E. Extracting training data from diffusion models. 2023. doi: 10.48550/ARXIV.2301.13188. URL <https://doi.org/10.48550/ARXIV.2301.13188>.
- Cetinic, E. and She, J. Understanding and creating art with ai: Review and outlook. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)*, 18(2):1–22, 2022.
- Cusumano, M. A., Gawer, A., and Yoffie, D. Can self-regulation save digital platforms? *Industrial and Corporate Change*, July 14 2021. doi: 10.2139/ssrn.3900137. URL <https://ssrn.com/abstract=3900137>.
- Darcy, O. News publishers sound alarm on google’s new ai-infused search, warn of ‘catastrophic’ impacts, May 2024. URL <https://www.cnn.com/2024/05/15/media/google-gemini-ai-search-news-outlet-impact/index.html>.
- Drescher, F. and Crabtree-Ireland, D. Member message: Strike authorization vote. SAG-AFTRA Correspondence, May 2023. URL <https://www.sagaftra.org/member-message-strike-authorization-vote>.
- Du, X., Zhang, H., Zhang, S., Zhang, A., and Chen, B. Creativity and leadership in the creative industry: A study from the perspective of social norms. *Frontiers in Psychology*, 12, 2021. doi: 10.3389/fpsyg.2021.651817. URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC8057347/pdf/fpsyg-12-651817.pdf>.
- Epstein, Z., Hertzmann, A., Herman, L., Mahari, R., Frank, M. R., Groh, M., Schroeder, H., Smith, A., Akten, M., Fjeld, J., et al. Art and the science of generative ai: A deeper dive. *arXiv preprint arXiv:2306.04141*, 2023.
- Jiang, H. H., Brown, L., Cheng, J., Khan, M., Gupta, A., Workman, D., Hanna, A., Flowers, J., and Gebru, T. Ai art and its impact on artists. In *AAAI/ACM Conference on AI, Ethics, and Society (AIES '23)*, pp. 12, Montréal, QC, Canada, August 2023. ACM. doi: 10.1145/3600211.3604681. URL <https://doi.org/10.1145/3600211.3604681>.
- Lee, K., Cooper, A. F., and Grimmelmann, J. Talkin’ ’bout ai generation: Copyright and the generative-ai supply

- chain. *Forthcoming, Journal of the Copyright Society*, July 27 2024. doi: 10.2139/ssrn.4523551. URL <https://ssrn.com/abstract=4523551>.
- Limbong, A. Authors push back on the growing number of ai ‘scam’ books on amazon, March 2024. URL <https://www.npr.org/2024/03/13/1237888126/growing-number-ai-scam-books-amazon>.
- Limo, F. A. F., Tiza, D. R. H., Roque, M. M., Herrera, E. E., Murillo, J. P. M., Huallpa, J. J., Flores, V. A. A., Castillo, A. G. R., Peña, P. F. P., Carranza, C. P. M., et al. Personalized tutoring: Chatgpt as a virtual tutor for personalized learning experiences. *Przestrzeń Społeczna (Social Space)*, 23(1):293–312, 2023.
- Liu, Y. and Lapata, M. Text summarization with pretrained encoders. *arXiv preprint arXiv:1908.08345*, 2019.
- Malle, B. F. What are norms, and how is norm compliance regulated? In Berg, M. K. and Chang, E. C. (eds.), *Motivation and Morality: A Multidisciplinary Approach*, pp. 45–75. American Psychological Association, 2023. doi: 10.1037/0000342-003. URL <https://doi.org/10.1037/0000342-003>.
- Mirowski, P. W., Love, J., Mathewson, K. W., and Mohamed, S. A robot walks into a bar: Can language models serve ascreativity support tools for comedy? an evaluation of llms’ humour alignment with comedians, 2024.
- Niroula, B. The impact of ai on the future of voice acting. 2023.
- Potts, J. *Creative Industries and Economic Evolution*. Number 12977 in Books. Edward Elgar Publishing, 2011. URL <https://ideas.repec.org/b/elg/eebook/12977.html>.
- Roose, K. Ai-generated art won a prize. artists aren’t happy., September 2022. URL <https://www.nytimes.com/2022/09/02/technology/ai-artificial-intelligence-artists.html>.
- Savini, A. and Savini, G. G. A short history of 3d printing, a technological revolution just started. *HISTELCON*, August 2015. doi: 10.1109/HISTELCON.2015.7307314.
- Shumakova, N., Lloyd, J., and Titova, E. Towards legal regulations of generative ai in the creative industry. *Journal of Digital Technologies and Law*, 1:880–908, 12 2023. doi: 10.21202/jdtl.2023.38.
- Sobel, B. L. Artificial intelligence’s fair use crisis. *Colum. JL & Arts*, 41:45, 2017.
- Walkowiak, E. and MacDonald, T. Generative ai and the workforce: What are the risks? *Available at SSRN*, 2023.
- Walkowiak, E. and Potts, J. Generative ai, work and risks in cultural and creative industries. May 16 2024. doi: 10.2139/ssrn.4830265. URL <https://ssrn.com/abstract=4830265>.
- Weidinger, L., Mellor, J., Rauh, M., Griffin, C., Uesato, J., Huang, P.-S., Cheng, M., Glaese, M., Balle, B., Kasirzadeh, A., et al. Ethical and social risks of harm from language models. *arXiv preprint arXiv:2112.04359*, 2021.