

Generative artificial intelligence, human creativity, and art

Eric Zhou ^{a,*} and Dokyun Lee ^{a,b,*}

^aDepartment of Information Systems, Boston University Questrom School of Business, Boston, MA 02215, USA

^bComputing & Data Sciences, Boston University, Boston, MA 02215, USA

*To whom correspondence should be addressed: Email: ebzhou@bu.edu (E.Z.); Email: dokyun@bu.edu (D.L.)

Edited By: Matthew Harding

Abstract

Recent artificial intelligence (AI) tools have demonstrated the ability to produce outputs traditionally considered creative. One such system is text-to-image generative AI (e.g. Midjourney, Stable Diffusion, DALL-E), which automates humans' artistic execution to generate digital artworks. Utilizing a dataset of over 4 million artworks from more than 50,000 unique users, our research shows that over time, text-to-image AI significantly enhances human creative productivity by 25% and increases the value as measured by the likelihood of receiving a favorite per view by 50%. While peak artwork Content Novelty, defined as focal subject matter and relations, increases over time, average Content Novelty declines, suggesting an expanding but inefficient idea space. Additionally, there is a consistent reduction in both peak and average Visual Novelty, captured by pixel-level stylistic elements. Importantly, AI-assisted artists who can successfully explore more novel ideas, regardless of their prior originality, may produce artworks that their peers evaluate more favorably. Lastly, AI adoption decreased value capture (favorites earned) concentration among adopters. The results suggest that ideation and filtering are likely necessary skills in the text-to-image process, thus giving rise to “generative synesthesia”—the harmonious blending of human exploration and AI exploitation to discover new creative workflows.

Keywords: generative AI, human–AI collaboration, creative workflow, impact of AI, art

Significance Statement

We investigate the implications of incorporating text-to-image generative artificial intelligence (AI) into the human creative workflow. We find that generative AI significantly boosts artists' productivity and leads to more favorable evaluations from their peers. While average novelty in artwork content and visual elements declines, peak Content Novelty increases, indicating a propensity for idea exploration. The artists who successfully explore novel ideas and filter model outputs for coherence benefit the most from AI tools, underscoring the pivotal role of human ideation and artistic filtering in determining an artist's success with generative AI tools.

Introduction

Recently, artificial intelligence (AI) has exhibited that it can feasibly produce outputs that society traditionally would judge as creative. Specifically, generative algorithms have been leveraged to automatically generate creative artifacts like music (1), digital artworks (2, 3), and stories (4). Such generative models allow humans to directly engage in the creative process through text-to-image systems (e.g. Midjourney, Stable Diffusion, DALL-E) based on the latent diffusion model (5) or by participating in an open dialog with transformer-based language models (e.g. ChatGPT, Bard, Claude). Generative AI is projected to become more potent to automate even more creative tasks traditionally reserved for humans and generate significant economic value in the years to come (6).

Many such generative algorithms were released in the past year, and their diffusion into creative domains has concerned

many artistic communities which perceive generative AI as a threat to substitute the natural human ability to be creative. Text-to-image generative AI has emerged as a candidate system that automates elements of humans' creative process in producing high-quality digital artworks. Remarkably, an artwork created by Midjourney bested human artists in an art competition,^a while another artist refused to accept the top prize in a photo competition after winning, citing ethical concerns.^b Artists have filed lawsuits against the founding companies of some of the most prominent text-to-image generators, arguing that generative AI steals from the works upon which the models are trained and infringes on the copyrights of artists.^c This has ignited a broader debate regarding the originality of AI-generated content and the extent to which it may replace human creativity, a faculty that many consider unique to humans. While generative AI has demonstrated the capability to automatically create new digital

Competing Interest: The authors declare no competing interest.

Received: January 4, 2024. **Accepted:** January 23, 2024

© The Author(s) 2024. Published by Oxford University Press on behalf of National Academy of Sciences. This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivs licence (<https://creativecommons.org/licenses/by-nc-nd/4.0/>), which permits non-commercial reproduction and distribution of the work, in any medium, provided the original work is not altered or transformed in any way, and that the work is properly cited. For commercial re-use, please contact reprints@oup.com for reprints and translation rights for reprints. All other permissions can be obtained through our RightsLink service via the Permissions link on the article page on our site—for further information please contact journals.permissions@oup.com.

artifacts, there remains a significant knowledge gap regarding its impact on productivity in artistic endeavors which lack well-defined objectives, and the long-run implications on human creativity more broadly. In particular, if humans increasingly rely on generative AI for content creation, creative fields may become saturated with generic content, potentially stifling exploration of new creative frontiers. Given that generative algorithms will remain a mainstay in creative domains as it continues to mature, it is critical to understand how generative AI is affecting creative production, the evaluation of creative artifacts, and human creativity more broadly. To this end, our research questions are 3-fold:

1. How does the adoption of generative AI affect humans' creative production?
2. Is generative AI enabling humans to produce more creative content?
3. When and for whom does the adoption of generative AI lead to more creative and valuable artifacts?

Our analyses of over 53,000 artists and 5,800 known AI adopters on one of the largest art-sharing platforms reveal that creative productivity and artwork value, measured as favorites per view, significantly increased with the adoption of text-to-image systems.

We then focus our analysis on creative novelty. A simplified view of human creative novelty with respect to art can be summarized via two main channels through which humans can inject creativity into an artifact: *Contents* and *Visuals*. These concepts are rooted in the classical philosophy of symbolism in art which suggests that the contents of an artwork is related to the meaning or subject matter, whereas visuals are simply the physical elements used to convey the content (7). In our setting, *Contents* concern the focal object(s) and relations depicted in an artifact, whereas *Visuals* consider the pixel-level stylistic elements of an artifact. Thus, *Content* and *Visual Novelty* are measured as the pairwise cosine distance between artifacts in the feature space (see Materials and methods for details on feature extraction and how novelty is measured).

Our analyses reveal that over time, adopters' artworks exhibit decreasing novelty, both in terms of *Concepts* and *Visual* features. However, maximum *Content Novelty* increases, suggesting an expanding yet inefficient idea space. At the individual level, artists who harness generative AI while successfully exploring more innovative ideas, irrespective of their prior originality, may earn more favorable evaluations from their peers. In addition, the adoption of generative AI leads to a less concentrated distribution of favorites earned among adopters.

Results

We present results from three analyses. Using an event study difference-in-differences approach (8), we first estimate the causal impact of adopting generative AI on creative productivity, artwork value measured as favorites per view, and artifact novelty with respect to *Content* and *Visual* features. Then, using a two-way fixed effects model, we offer correlational evidence regarding how humans' originality prior to adopting generative AI may influence postadoption gains in artwork value when artists successfully explore the creative space. Lastly, we show how adoption of generative AI may lead to a more dispersed distribution of favorites across users on the platform.

Creative productivity

We define creative productivity as the log of the number of artifacts that a user posts in a month. Figure 1a reveals that upon adoption, artists experience a 50% increase in productivity on average, which then doubles in the subsequent month. For the average user, this translates to approximately 7 additional artifacts published in the adoption month and 15 artifacts in the following month. Beyond the adoption month, user productivity gradually stabilizes to a level that still exceeds preadoption volume. By automating the execution stage of the creative process, adopters can experience prolonged productivity gains compared to their nonadopter counterparts.

Creative value

If users are becoming more productive, what of the quality of the artifacts they are producing? We next examine how adopters' artifacts are evaluated by their peers over time. In the literature, creative Value is intended to measure some aspect of utility, performance and/or attractiveness of an artifact, subject to temporal and cultural contexts (9). Given this subjectivity, we measure Value as the number of favorites an artwork receives per view after 2 weeks, reflecting its overall performance and contextual relevance within the community. This metric also hints at the artwork's broader popularity within the cultural climate, suggesting a looser definition of Value based on cultural trends. Throughout the paper, the term "Value" will refer to these two notions.

Figure 1b reveals an initial nonsignificant upward trend in the Value of artworks produced by AI adopters. But after 3 months, AI adopters consistently produce artworks judged significantly more valuable than those of nonadopters. This translates to a 50% increase in artwork favorability by the sixth month, jumping from the preadoption average of 2% to a steady 3% rate of earning a favorite per view.

Content Novelty

Figure 1c shows that average *Content Novelty* decreases over time among adopters, meaning that the focal objects and themes within new artworks produced by AI adopters are becoming progressively more alike over time when compared to control units. Intuitively, this is equivalent to adopters' ideas becoming more similar over time. In practice, many publicly available fine-tuned checkpoints and adapters are refined to enable text-to-image models to produce specific contents with consistency. Figure 1d, however, reveals that maximum *Content Novelty* is increasing and marginally statistically significantly within the first several months after adoption. This suggests two possibilities: either a subset of adopters are exploring new ideas at the creative frontier or the adopter population as a whole is driving the exploration and expansion of the universe of artifacts.

Visual Novelty

The result shown in Fig. 1e highlights that average *Visual Novelty* is decreasing over time among adopters when compared to nonadopters. The same result holds for the maximum *Visual Novelty* seen in Fig. 1f. This suggests that adopters may be gravitating toward a preferred visual style, with relatively minor deviations from it. This tendency could be influenced by the nature of text-to-image workflows, where prompt engineering tends to follow a formulaic approach to generate consistent, high-quality images with a specific style. As is the case with contents, publicly available fine-tuned checkpoints and adapters for these models

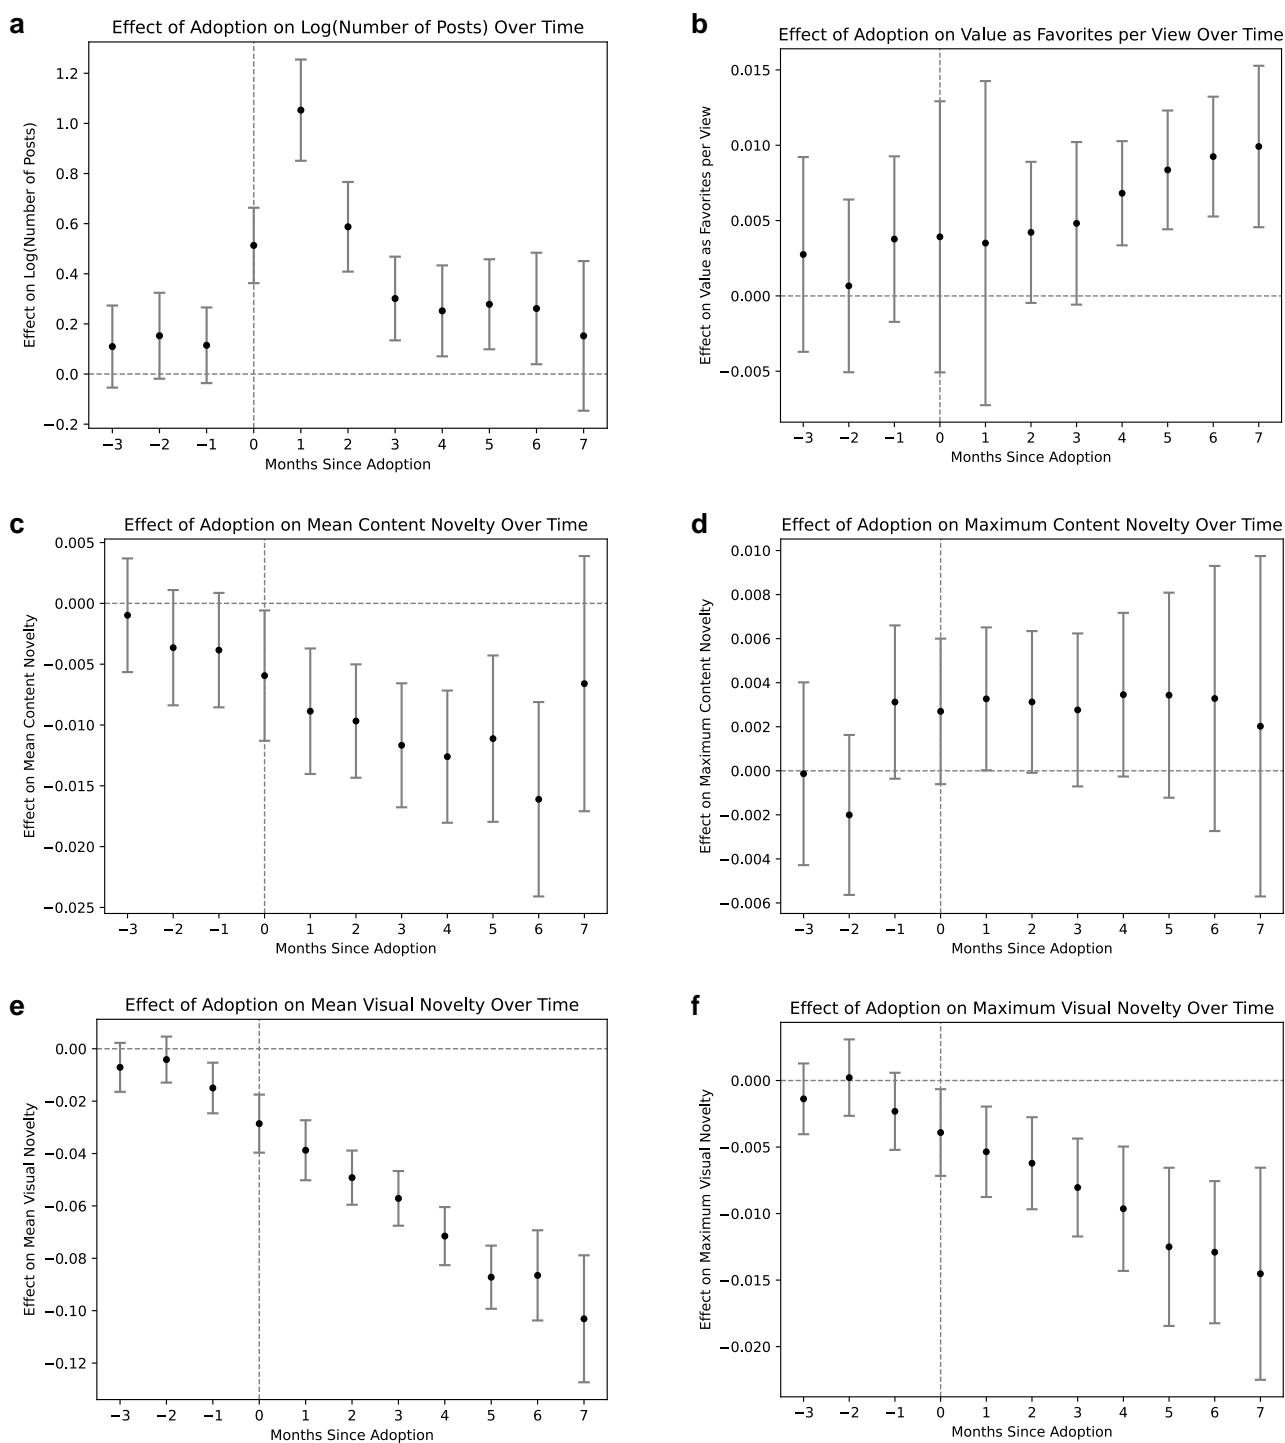


Fig. 1. Causal effect of adopting generative AI on a) creative productivity as the log of monthly posts; b) creative value as number of favorites per view; c) mean Content Novelty; d) maximum Content Novelty; e) mean Visual Novelty; f) maximum Visual Novelty. The error bars represent 95% CI.

may be designed to capture specific visual elements from which users can sample from to maintain a particular and consistent visual style. In effect, AI may be pushing artists toward visual homogeneity.

Role of human creativity in AI-assisted value capture

Although aggregate trends suggest novelty of ideas and aesthetic features is sharply declining over time with generative AI, are

there individual-level differences that enable certain artists to successfully produce more creative artworks? Specifically, how does humans' baseline novelty, in the absence of AI tools, correlate with their ability to successfully explore novel ideas with generative AI to produce valuable artifacts? To delve into this heterogeneity, we categorize each user into quartiles based on their average Content and Visual Novelty without AI assistance to capture each users' baseline novelty. We then employ a two-way fixed effects model to examine the interaction between adoption, pretreatment novelty quartiles, and posttreatment

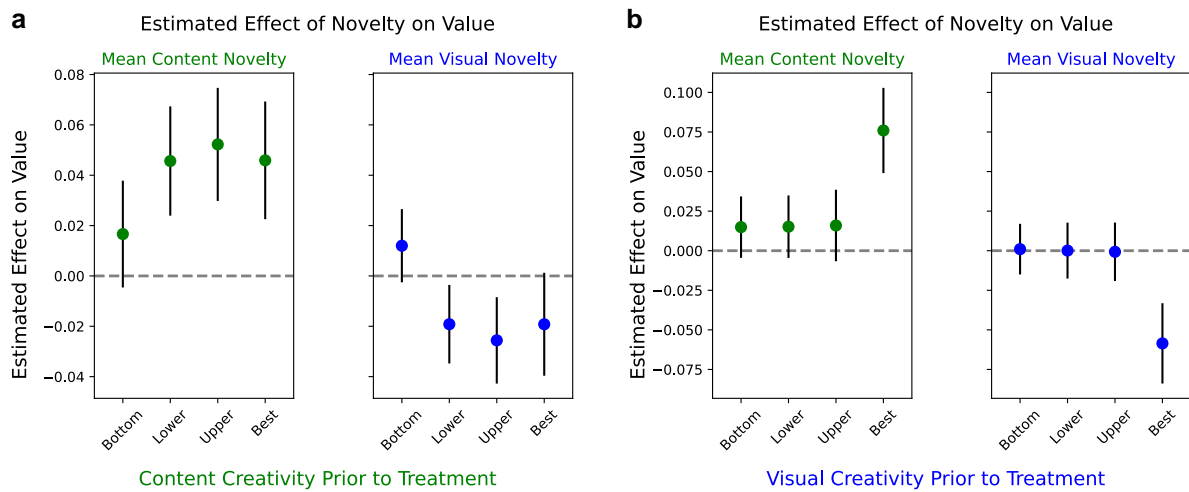


Fig. 2. Estimated effect of increases in mean Content and Visual Novelty on Value post-adoption based on a) average Content Novelty quartiles prior to treatment; b) average Visual Novelty quartiles prior to treatment. Each point shows the estimated effect of postadoption novelty increases given creativity levels prior to treatment on Value. The error bars represent 95% CI.

adjustments in novelty. Each point in Fig. 2a and b represents the estimated impact of increasing mean Content (left) or Visual (right) Novelty on Value based on artists' prior novelty denoted along the horizontal axis. Intuitively, these estimates quantify the degree to which artists can successfully navigate the creative space based on prior originality in both ideation and visuals to earn more favorable evaluations from peers. Refer to [SI Appendix, Section 2B](#) for estimation details.

Figure 2a presents correlational evidence that users, regardless of their proficiency in generating novel ideas, might be able to realize significant gains in Value if they can successfully produce more novel content with generative AI. The lowest quartile of content creators may also experience marginally significant gains. However, those same users who benefit from expressing more novel ideas may also face penalties for producing more divergent visuals.

Next, Fig. 2b suggests that users who were proficient in creating exceedingly novel visual features before adopting generative AI may garner the most Value gains from successfully introducing more novel ideas. While marginally significant, less proficient users can also experience weak Value gains. In general, more novel ideas are linked to improved Value capture. Conversely, users capable of producing the most novel visual features may face penalties for pushing the boundaries of pixel-level aesthetics with generative AI. This finding might be attributed to the contextual nature of Value, implying an "acceptable range" of novelty. Artists already skilled at producing highly novel pixel-level features may exceed the limit of what can be considered coherent.

Despite penalties for pushing visual boundaries, the gains from exploring creative ideas with AI outweigh the losses from visual divergence. Unique concepts take priority over novel aesthetics, as shown by the larger Value gains for artists who were already adept at Visual Novelty before using AI. This suggests users who naturally lean toward visual exploration may benefit more from generative AI tools to explore the idea space.

Lastly, we estimate Generalized Random Forests (10) configured to optimize the splitting criteria that maximize heterogeneity in Value gains among adopters for each postadoption period. With each trained model, we extract feature importance weights quantified by the SHAP (SHapley Additive exPlanations)

method (11). This method utilizes ideas from cooperative game theory to approximate the predictive signal of covariates, accounting for linear and nonlinear interactions through the Markov chain Monte Carlo method. Intuitively, a feature of greater importance indicates potentially greater impacts on treatment effect heterogeneity among adopters.

Figure 3 offers correlational evidence that Content Novelty significantly increases model performance within several months of adoption, whereas Visual Novelty remains marginally impactful until the last observation period. This suggests that Content Novelty plays a more significant role in predicting posttreatment variations in Value gains compared to Visual Novelty. In summary, these findings illustrate that content is king in the text-to-image creative paradigm.

Platform-level value capture

One question remains: do individual-level differences within adopters result in greater concentrations of value among fewer users at the platform-level? Specifically, are more favorites being captured by fewer users, or is generative AI promoting less concentrated value capture? To address these questions, we calculate the Gini coefficients with respect to favorites received of never-treated units, not-yet-treated units, and treated units and conduct permutation tests with 10,000 iterations to evaluate if adoption of generative AI may lead to a less concentrated distribution of favorites among users. The Gini coefficient is a common measure of aggregate inequality where a coefficient of 0 indicates that all users make up an equal proportion of favorites earned, and a coefficient of 1 indicates that a single user captures all favorites. Thus, higher values of the Gini coefficient indicate a greater concentration of favorites captured by fewer users. Figure 4 depicts the differences in cumulative distributions as well as Gini coefficients of both control groups and the treated group with respect to a state of perfect equality.

First, observe that platform-level favorites are predominantly captured by a small portion of users, reflecting an aggregate concentration of favorites. Second, this concentration is more pronounced among not-yet-treated units than among never-treated units. Third, despite the presence of aggregate concentration,

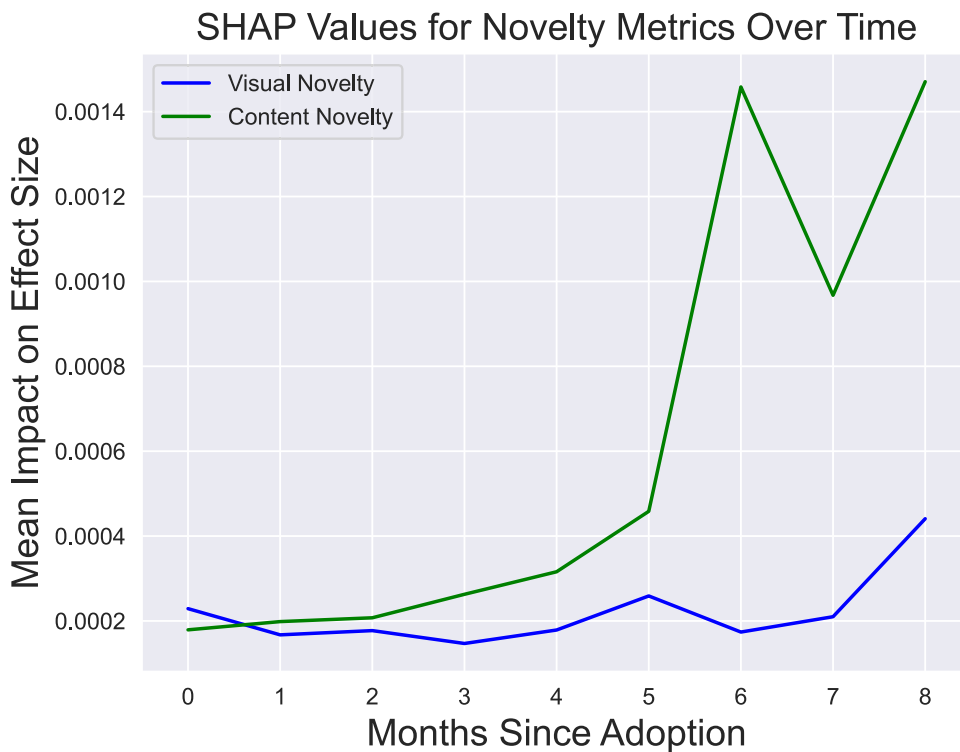


Fig. 3. SHAP values measuring importance of mean Content and Visual Novelty on Value gains.

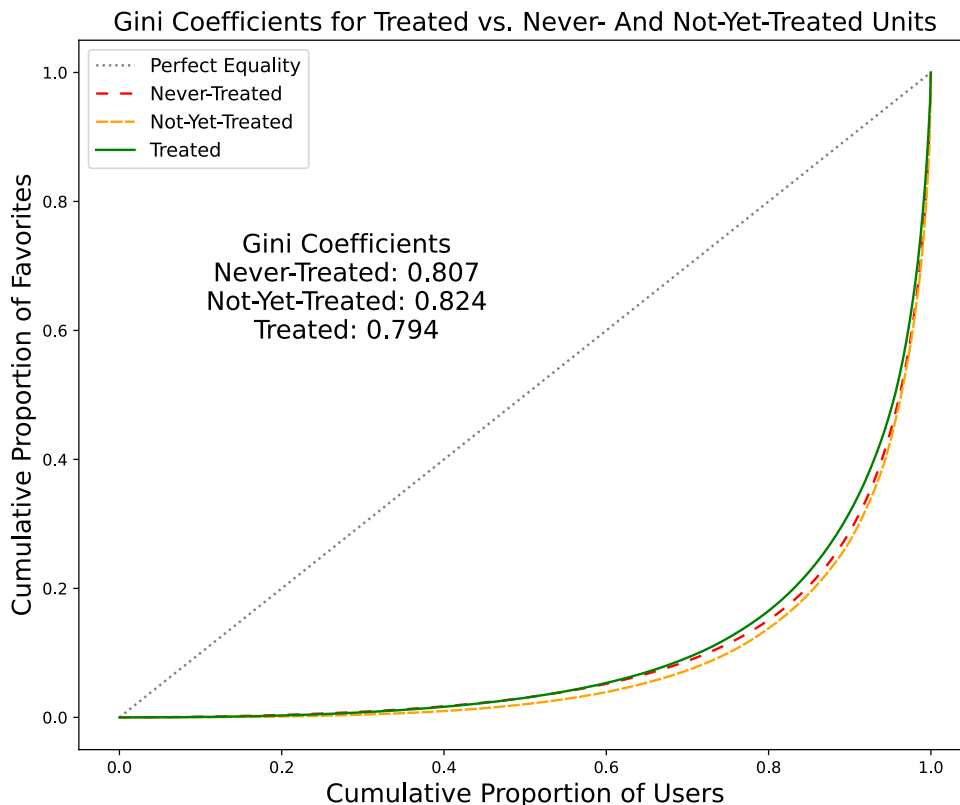


Fig. 4. Gini coefficients of treated units vs. never-treated and not-yet-treated units.

favorites captured among AI adopters are more evenly distributed compared to both never-treated and not-yet-treated control units. The results from the permutation tests in Table 1, where column

D shows the difference between the treated coefficient and the control group coefficients, show that the differences in coefficients are statistically significant between never-treated and

Table 1. Permutation tests for statistical significance.

	Coefficient	D	P-value
Never-treated	0.807	−0.0128	0.0673
Not-yet-treated	0.824	−0.0298	0.0026
Treated	0.794		

The column *D* denotes the difference in Gini coefficients relative to the treated population.

not-yet-treated groups vs. the treated group. This suggests that generative AI may lead to a broader allocation of favorites earned (value capture from peer feedback), particularly among control units who eventually become adopters.

Robustness checks and sensitivity analyses

To reinforce the validity of our causal estimates, we employ the generalized synthetic control method (12) (GSCM). GSCM allows us to relax the parallel trends assumption by creating synthetic control units that closely match the pretreatment characteristics of the treated units while also accounting for unobservable factors that may influence treatment outcome. In addition, we conduct permutation tests to evaluate the robustness of our estimates to potential measurement errors in treatment time identification and control group contamination. Our results remain consistent even when utilizing GSCM and in the presence of substantial measurement error.

Because adopting generative AI is subject to selection issues, one emergent concern is the case where an artist who experiences renewed interest in creating artworks, and thus is more “inspired,” is also more likely to experiment with text-to-image AI tools and explore the creative space as they ramp up production. In this way, unobservable characteristics like a renewed interest in creating art or “spark of inspiration” might correlate with adoption of AI tools while driving the main effects rather than AI tools themselves. Thus, we also provide evidence that unobservable characteristics that may correlate with users’ productivity or “interest” shocks and selection into treatment are not driving the estimated effects by performing a series of falsification tests. For a comprehensive overview of all robustness checks and sensitivity analyses, please refer to [SI Appendix, Section 3](#).

Discussion

The rapid adoption of generative AI technologies poses exceptional benefits as well as risks. Current research demonstrates that humans, when assisted by generative AI, can significantly increase productivity in coding (13), ideation (14), and written assignments (15) while raising concerns regarding potential disinformation (16) and stagnation of knowledge creation (17). Our research is focused on how generative AI is impacting and potentially coevolving with human creative workflows. In our setting, human creativity is embodied through prompts themselves, whereas in written assignments, generative AI is primarily used to source ideas that are subsequently evaluated by humans, representing a different paradigm shift in the creative process.

Within the first few months post-adoption, text-to-image generative AI can help individuals produce nearly double the volume of creative artifacts that are also evaluated 50% more favorably by their peers over time. Moreover, we observe that peak Content Novelty increases over time, while average Content and Visual

Novelty diminish. This implies that the universe of creative possibilities is expanding but with some inefficiencies.

Our results hint that the widespread adoption of generative AI technologies in creative fields could lead to a long-run equilibrium where in aggregate, many artifacts converge to the same types of content or visual features. Creative domains may be inundated with generic content as exploration of the creative space diminishes. Without establishing new frontiers for creative exploration, AI systems trained on outdated knowledge banks run the risk of perpetuating the generation of generic content at a mass scale in a self-reinforcing cycle (17). Before we reach that point, technology firms and policy makers pioneering the future of generative AI must be sensitive to the potential consequences of such technologies in creative fields and society more broadly.

Encouragingly, humans assisted by generative AI who can successfully explore more novel ideas may be able to push the creative frontier, produce meaningful content, and be evaluated favorably by their peers. With respect to traditional theories of creativity, one particularly useful framework for understanding these results is the theory of blind variation and selective retention (BVSR) which posits that creativity is a process of generating new ideas (variation) and consequently selecting the most promising ones (retention) (18). The blindness feature suggests that variation is not guided by any specific goal but can also involve evaluating outputs against selection criteria in a genetic algorithm framework (19).

Because we do not directly observe users’ process, this discussion is speculative but suggestive that a text-to-image creative workflow models after a BVSR genetic process. First, humans manipulate and mutate known creative elements in the form of prompt engineering which requires that the human deconstruct an idea into atomic components, primarily in the form of distinct words and phrases, to compose abstract ideas or meanings. Then, visual realization of an idea is automated by the algorithm, allowing humans to rapidly sample ideas from their creative space and simply evaluate the output against selection criteria. The selection criteria varies based on humans’ ability to make sense of model outputs, and curate those that most align with individual or peer preferences, thus having direct implications on their evaluation by peers. Satisfactory outputs contribute to the genetic evolution of future ideas, prompts, and image refinements.

Although we can only observe the published artworks, it is plausible that many more unobserved iterations of ideation, prompt engineering, filtering, and refinement have occurred. This is especially likely given the documented increase in creative productivity. Thus, it is possible that individuals with less refined artistic filters are also less discerning when filtering artworks for quality which could lead to a flood of less refined content on platforms. In contrast, artists who prioritize coherence and quality may only publish artworks that are likely to be evaluated favorably.

The results suggest some evidence in this direction, indicating that humans who excel at producing novel ideas before adopting generative AI are evaluated most favorably after adoption if they successfully explore the idea space, implying that ability to manipulate novel concepts and curate artworks based on coherence are relevant skills when using text-to-image AI. This aligns with prior research which suggest that creative individuals are particularly adept at discerning which ideas are most meaningful (20), reflecting a refined sensitivity to the artistic coherence of artifacts (21). Furthermore, all artists, regardless of their ability to produce novel visual features without generative AI, appear to be evaluated more favorably if they can capably explore more novel

ideas. This finding hints at the importance of humans' baseline ideation and filtering abilities as focal expressions of creativity in a text-to-image paradigm. Finally, generative AI appears to promote a more even distribution of platform-level favorites among adopters, signaling a potential step toward an increasingly democratized, inclusive creative domain for artists empowered by AI tools.

In summary, our findings emphasize that humans' ideation proficiency and a refined artistic filter rather than pure mechanical skill may become the focal skills required in a future of human-AI cocreative process as generative AI becomes more mainstream in creative endeavors. This phenomenon in which AI-assisted artistic creation is driven by ideas and filtering is what we term "generative synesthesia"—the harmonization of human exploration and AI exploitation to discover new creative workflows. This paradigm shift may provide avenues for creatives to focus on what ideas they are representing rather than how they represent it, opening new opportunities for creative exploration. While concerns about automation loom, society must consider a future where generative AI is not the source of human stagnation, but rather of symphonic collaboration and human enrichment.

Materials and methods

Identifying AI adopters

Platform-level policy commonly suggests that users disclose their use of AI assistance in the form of tags associated with their artworks. Thus, we employ a rule-based classification scheme. As a first-pass, any artwork published before the original DALL-E in January 2021 is automatically labeled as non-AI generated. Then, for all artworks published after January 2021, we examine postlevel title and tags provided by the publishing user. We use simple keyword matching (AI-generated, Stable Diffusion, Midjourney, DALL-E, etc.) for each post to identify for which artworks a user employs AI tools. As a second-pass, we track artworks posts published in AI art communities which may not include explicit tags denoting AI assistance. We compile all of these artworks and simply label them as AI-generated. Finally, we assign adoption timing based on the first-known AI-generated post for each use (SI Appendix, Fig. S2).

Measuring creative novelty

To measure the two types of novelty, we borrow the idea of conceptual spaces which can be understood as geometric representations of entities which capture particular attributes of the artifacts along various dimensions (9, 22). This definition naturally aligns with the concept of embeddings, like *word2vec* (23), which capture the relative features of objects in a vector space. This concept can be applied to text passages and images such that measuring the distance between these vector representations captures whether an artifact deviates or converges with a reference object in the space.

Using embeddings, we apply the following algorithm: take all artifacts published before 2022 April 1, as the baseline set of artworks. We use this cutoff because nearly all adoption occurs after May 2022, so all artifacts in future periods are compared to non-AI-generated works in the baseline period, and it provides an adequate number of pretreatment and posttreatment observations (on average 3 and 7, respectively) for the majority of our causal sample. Then, take all artifacts published in the following month and measure the pairwise cosine distance between those artifacts and the baseline set, recovering the mean, minimum,

and maximum distances for each artifact. This month's artifacts are then added to the baseline set such that all future artworks are compared to all prior artworks, effectively capturing the time-varying nature of novelty. Continue for all remaining months. We apply this approach to all adopters' artworks and a random sample of 10,000 control users due to computational feasibility.

Content feature extraction

To describe the focal objects and object relationships in an artifact, we utilize state-of-the-art multimodal model BLIP-2 (24) which takes as input an image and produces a text description of the content. A key feature of this approach is the availability of controlled text generation hyperparameters that allow us to generate more stable descriptions that are systematically similar in structure, having been trained on 129M images and human-annotated data. BLIP-2 can maintain consistent focus and regularity while avoiding the noise added by cross-individual differences.

Given the generated descriptions, we then utilize a pretrained text embedding model based on BERT (25), which has demonstrated state-of-the-art performance on semantic similarity benchmarks while also being highly efficient, to compute high-dimensional vector representations for each description. Then, we apply the algorithm described above to measure Content Novelty.

Visual feature extraction

To capture visual features of each artifact at the pixel level, we use a more flexible approach via the self-supervised visual representation learning algorithm DINOv2 (26) which overcomes the limitations of standard image-text pretraining approaches where visual features may not be explicitly described in text. Because we are dealing with creative concepts, this approach is particularly suitable to robustly identify object parts in an image and extract low-level pixel features of images while still exhibiting excellent generalization performance. We compute vector representations of each image such that we can apply the algorithm described above to obtain measures of Visual Novelty.

Notes

^a An AI-generated picture won an art prize. Artists are not happy.

^b Artist wins photography contest after submitting AI-generated image, then forfeits prize.

^c The current legal cases against generative AI are just the beginning.

Acknowledgments

The authors acknowledge the valuable contributions from their Business Insights through Text Lab (BITLAB) research assistants Animikh Aich, Aditya Bala, Amrutha Karthikeyan, Audrey Mao, and Esha Vaishnav in helping to prepare the data for analysis. Furthermore, the authors are grateful for Stefano Puntoni, Alex Burnap, Mi Zhou, Gregory Sun, our audiences at the Wharton Business & Generative AI Workshop (23/9/8), INFORMS Workshop on Data Science (23/10/14), INFORMS Annual Meeting (23/10/15) and seminar participants at the University of Wisconsin-Milwaukee (23/9/22), University of Texas-Dallas (23/10/6), and MIT Initiative on the Digital Economy (23/11/29) for their insightful comments and feedback.

Supplementary Material

Supplementary material is available at PNAS Nexus online.

Funding

The authors declare no funding.

Author Contributions

D.L. and E.Z. designed the research and wrote the paper. E.Z. analyzed data and performed research with guidance from D.L.

Preprints

A preprint of this article is available at SSRN.

Data Availability

Replication archive with code is available at Open Science Framework at <https://osf.io/jfzyp/>. Data have been anonymized for the privacy of the users.

References

- Dong H-W, Hsiao W-Y, Yang L-C, Yang Y-H. 2018. MuseGAN: multi-track sequential generative adversarial networks for symbolic music generation and accompaniment. In: Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence and Thirtieth Innovative Applications of Artificial Intelligence Conference and Eighth AAAI Symposium on Educational Advances in Artificial Intelligence. p. 34–41.
- Tan WR, Chan CS, Aguirre H, Tanaka K. 2017. ArtGAN: artwork synthesis with conditional categorical GANs. In: 2017 IEEE International Conference on Image Processing (ICIP). IEEE. p. 3760–3764.
- Elgammal A, Liu B, Elhoseiny M, Mazzone M. 2017. CAN: creative adversarial networks, generating “art” by learning about styles and deviating from style norms, arXiv, arXiv:1706.07068, preprint: not peer reviewed.
- Brown TB, et al. 2020. Language models are few-shot learners. *Adv Neural Inf Process Syst.* 33:1877–1901.
- Rombach R, Blattmann A, Lorenz D, Esser P, Ommer B. 2022. High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition. p. 10684–10695.
- Huang S, Grady P, GPT-3. 2022. Generative AI: a creative new world. Sequoia Capital US/Europe. <https://www.sequoiacap.com/article/generative-ai-a-creative-new-world/>
- Wollheim R. 1970. Nelson Goodman’s languages of art. *J Philos.* 67(16):531–539.
- Callaway B, Sant’Anna PHC. 2021. Difference-in-differences with multiple time periods. *J Econom.* 225(2):200–230.
- Boden MA. 1998. Creativity and artificial intelligence. *Artif Intell.* 103(1–2):347–356.
- Athey S, Tibshirani J, Wager S. 2019. Generalized random forests. *Ann Statist.* 47(2):1148–1178.
- Lundberg S, Lee S-I. 2017. A unified approach to interpreting model predictions. In: Proceedings of the 31st International Conference on Neural Information Processing Systems. p. 4768–4777.
- Xu Y. 2017. Generalized synthetic control method: causal inference with interactive fixed effects models. *Polit Anal.* 25(1):57–76.
- Peng S, Kalliamvakou E, Cihon P, Demirel M. 2023. The impact of AI on developer productivity: evidence from github copilot, arXiv, arXiv:2302.06590, preprint: not peer reviewed.
- Noy S, Zhang W. 2023. Experimental evidence on the productivity effects of generative artificial intelligence. *Science.* 381(6654): 187–192.
- Dell’Acqua F, et al. 2023. Navigating the jagged technological frontier: field experimental evidence of the effects of AI on knowledge worker productivity and quality. Harvard Business School Technology & Operations Mgt. Unit Working Paper, (24-013).
- Spitale G, Biller-Andorno N, Germani F. 2023. AI model GPT-3 (dis)informs us better than humans. *Sci Adv.* 9(26):eadh1850.
- Burtch G, Lee D, Chen Z. 2023. The consequences of generative AI for UGC and online community engagement. Available at SSRN 4521754.
- Campbell DT. 1960. Blind variation and selective retentions in creative thought as in other knowledge processes. *Psychol Rev.* 67(6):380–400.
- Simonton DK. 1999. Creativity as blind variation and selective retention: is the creative process Darwinian? *Psychol Inq.* 10(4):309–328.
- Silvia PJ. 2008. Discernment and creativity: how well can people identify their most creative ideas? *Psychol Aesthet Creat Arts.* 2(3): 139–146.
- Ivcevic Z, Mayer JD. 2009. Mapping dimensions of creativity in the life-space. *Creat Res J.* 21(2–3):152–165.
- McGregor S, Wiggins G, Purver M. 2014. Computational creativity: a philosophical approach, and an approach to philosophy. In: International Conference on Innovative Computing and Cloud Computing. p. 254–262.
- Mikolov T, Chen K, Corrado G, Dean J. 2013. Efficient estimation of word representations in vector space, arXiv, arXiv:1301.3781, preprint: not peer reviewed.
- Li J, Li D, Savarese S, Hoi S. 2023. Bootstrapping language-image pre-training with frozen image encoders and large language models, arXiv, arXiv:2301.12597, preprint: not peer reviewed.
- Reimers N, Gurevych I. 2019. Sentence-bert: sentence embeddings using siamese bert-networks, arXiv, arXiv:1908.10084, preprint: not peer reviewed.
- Oquab M, et al. 2023. DINOv2: learning robust visual features without supervision, arXiv, arXiv:2304.07193, preprint: not peer reviewed.