# As Generative Models Improve, People Adapt Their Prompts*

Eaman Jahani
University of Maryland

Benjamin S. Manning
MIT

Joe Zhang
Stanford University

Hong-Yi TuYe
MIT

Mohammed Alsobay
MIT

Christos Nicolaides
University of Cyprus

Siddharth Suri†
Microsoft Research

David Holtz†
University of California, Berkeley

July 22, 2024

## Abstract

In an online experiment with N = 1,891 participants, we collected and analyzed over 18,000 prompts to explore how the importance of prompting will change as the capabilities of generative AI models continue to improve. Each participant in our experiment was randomly and blindly assigned to use one of three text-to-image diffusion models: DALL-E 2, its more advanced successor DALL-E 3, or a version of DALL-E 3 with automatic prompt revision. Participants were then asked to write prompts to reproduce a target image as closely as possible in 10 consecutive tries. We find that task performance was higher for participants using DALL-E 3 than for those using DALL-E 2. This performance gap corresponds to a noticeable difference in the similarity of participants' images to their target images, and was caused in equal measure by: (1) the increased technical capabilities of DALL-E 3, and (2) endogenous changes in participants' prompting in response to these increased capabilities. More specifically, despite being blind to the model they were assigned, participants assigned to DALL-E 3 wrote longer prompts that were more semantically similar to each other and contained a greater number of descriptive words. Furthermore, while participants assigned to DALL-E 3 with prompt revision still outperformed those assigned to DALL-E 2, automatic prompt revision reduced the benefits of using DALL-E 3 by 58%. Taken together, our results suggest that as models continue to progress, people will continue to adapt their prompts to take advantage of new models' capabilities.

The release of ChatGPT in November 2022 has fueled a surge of interest in generative AI and large language models across industry, the public sector, and academia, with early research suggesting that generative AI may be a general purpose technology (Eloundou et al., 2024). Indeed, generative AI models have already been found to outperform humans on text annotation tasks (Gilardi et al., 2023) and improve the productivity of workers across various professional settings (Brynjolfsson et al., 2023; Dell'Acqua et al., 2023; Noy and Zhang, 2023). However, the extent to which generative AI can provide value is contingent on the quality of the textual instructions, or "prompts," that users write to a model. As generative AI has proliferated, the practice of creating and iterating on such prompts, often using a variety of specialized techniques (e.g., chain-of-thought and few-shot prompting) (Schulhoff et al., 2024), has come to be known as prompt engineering. Some firms now even offer six-figure salaries for dedicated prompt engineers (Quilty-Harper, 2023), and an expanding academic literature focuses exclusively on prompting-related topics (Oppenlaender, 2023; Schulhoff et al., 2024; Don-Yehiya et al., 2023a; Xie et al., 2023). Despite the growing interest in prompting, some argue that as models become more capable at transforming human language into output, the need to craft prompts in specific, and sometimes esoteric, ways will diminish (Acar, 2023; Mollick, 2023). If this comes to pass, the benefits of generative AI will be accessible to a much larger population.

In this paper, we bring experimental evidence to bear on this debate by presenting results from a pre-registered online lab experiment with $N = 1,891$ participants recruited on Prolific.[1] In the experiment, participants were randomly and blindly assigned to complete the same task using one of three generative AI models with varying capabilities: (1) DALL-E 2, (2) DALL-E 3, and (3) a version of DALL-E 3 with automatic LLM-based prompt revisions (henceforth referred to as "DALL-E 3 with revision"). These prompt revisions were not made visible to participants during the experiment.[2] The task required each participant to make at least 10 attempts at recreating a "target image" as closely as possible by prompting their assigned model. Each participant's target image was randomly selected from a curated set of 15 images; this set of images was constructed to span realistic use cases of AI-generated images. Throughout the task, participants could view all of their past prompts and generated images alongside the target image. Participants were paid $4 USD to complete the task, and to incentivize performance, they received $8 USD if they were in the top 20% of participants, as measured by the similarity of their best attempt to the target image.

[1]This paper presents a subset of pre-registered analyses, with additional results and deviations from pre-registration detailed in the **SI**.

[2]Automatic prompt revision is the default behavior of the DALL-E 3 API endpoint. We followed the DALL-E 3 API documentation and attempted to disable prompt revision in the second, "DALL-E 3" treatment arm by prepending a system prompt to participants' prompts. While this drastically reduced the number and extent of LLM-based prompt revisions, it did not completely eliminate them. Therefore, the "DALL-E 3" treatment arm still includes some automatic prompt revisions. For brevity and simplicity, we refer to this treatment arm as "DALL-E 3" throughout the paper. Full details on the system prompt's efficacy are provided in the **SI**.
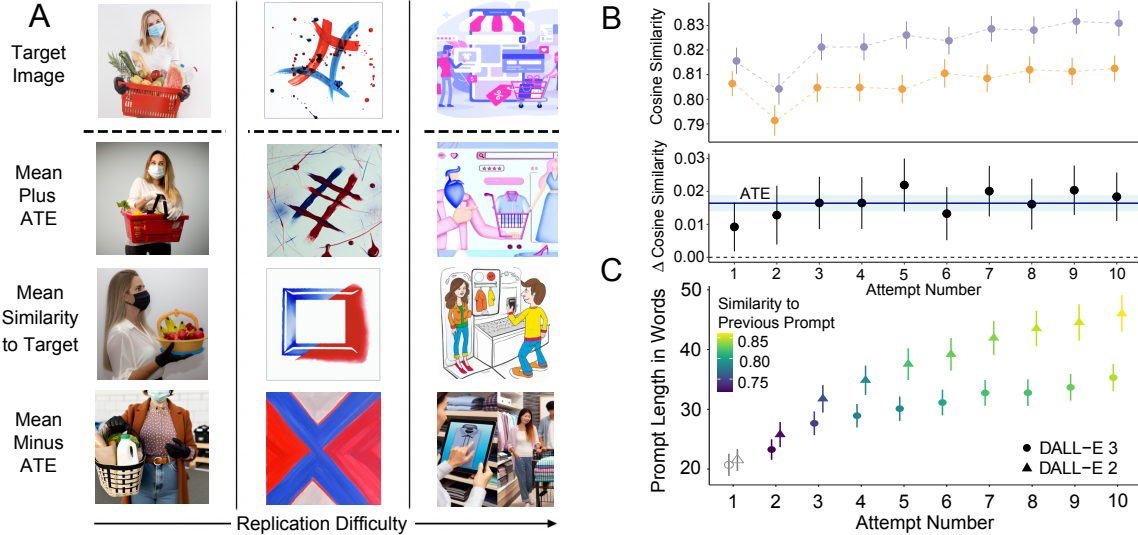
Figure 1: Panel (A): The top row shows three example target images, ordered by the difficulty participants had in replicating them. The middle row below the dashed line shows images representing the mean similarity to each target image based on all relevant attempts in either the DALL-E 2 or DALL-E 3 treatment arms. The images in the row above (below) show images that are the ATE more (less) similar than the mean image to their relevant target image. Panel (B): The top pane shows the average CLIP embedding cosine similarity of participant-generated images to their target image by model per replication attempt. The bottom pane shows the difference between averages in the top pane; i.e., the per-attempt ATE, with the dark blue line corresponding to the overall ATE ($\Delta CoSim = 0.0164$) and blue shading depicting the 95% confidence interval. In the **SI**, we show that the results in Panel (B) still hold when standardizing our data within-image. Panel (C): This plot shows the average prompt length in words by model averaged across participants for each of their ten attempts. The color scale corresponds to the average cosine similarity of each prompt's text embedding vector to that of the previous attempt's prompt. The 1st attempt does not have an average cosine similarity since there is no prior prompt to compare it to. All error bars depict 95% confidence intervals.

# Results

The majority of our analyses focus on the DALL-E 2 and DALL-E 3 treatment arms, as they are more directly comparable. We first estimate the effect of access to a more advanced generative model on task performance by evaluating how similar each generated image is to the relevant target image.[3] To construct a quantitative measure of image similarity, we generated CLIP embedding vectors (Radford et al., 2021) for each participant-generated image and the relevant target image and then calculated the cosine similarity ($CoSim$) between those vectors (see **SI** for further details on how we quantify image similarity and account for the stochasticity of model output). Unsurprisingly, we find that providing access to a more advanced generative model improved task performance. Participants using DALL-E 3 (the superior model) produced images that were, on average, $z = 0.20$

---

[3] All main text analyses use "replay" data from Figure 2 for consistency and to avoid model drift issues. All results hold with the "original" data. See the **SI** for details.

standard deviations closer to the target image ($\Delta CoSim = 0.0164$, $p < 10^{-7}$) than those produced by participants using DALL-E 2 (the inferior model). Figure 1A presents three examples from our data to illustrate the qualitative magnitude of this average treatment effect (ATE), which corresponds to a considerable increase in similarity to the target image. We also find that the treatment effect increased slightly as participants made successive attempts to replicate the target image (See Figure 1B; $\beta = 0.0008$, $p = 0.042$). In other words, the performance gap widened between DALL-E 2 and DALL-E 3 as participants gained more experience with the respective models, potentially due to changes in prompting.

To investigate how participants' prompts changed in response to the superior model's capabilities, we conducted an exploratory analysis of the contents of participant prompts and how those prompts evolved as participants iterated. We find that participants assigned the superior model wrote prompts that were 24% longer ($\Delta Words = 6.6$, $p < 10^{-4}$), and that the difference in prompt length increased as participants made successive replication attempts (Figure 1C; $\beta = 1.17$ extra words per attempt, $p < 10^{-7}$). Despite differing in length, prompts written for both models contained similar proportions of nouns and adjectives (48% superior vs. 49% inferior; $p = 0.148$). This suggests that prompts submitted to DALL-E 3 were longer because they conveyed additional descriptive information, not because they were padded with superfluous words. We also find that participants blindly assigned to DALL-E 3 wrote prompts that were more similar to one another, both sequentially from prompt-to-prompt (as shown by the color scale in Figure 1C; $\beta = 0.0166$, $p = 0.02$) and in the aggregate ($\beta = 0.0169$, $p = 0.008$) (see the **SI** for the details of our prompt-related textual analyses). These findings indicate that exposure to different models did in fact induce differences in prompting, and that these differences grew over successive replication attempts.

To understand the extent to which the ATE we observe is attributable to these behavioral differences, as opposed to DALL-E 3's more advanced capabilities, we conducted another exploratory analysis in which we regenerated images with (or "replayed") all participant prompts on both DALL-E 2 and DALL-E 3, regardless of the model the participant used during the experiment. This allows us to decompose the ATE into two components: the "model effect," i.e., the average increase in similarity to the target when replaying the prompts written by inferior model participants on the superior model, and the "prompting effect," i.e., the difference between the average similarity to the target of all prompts generated by superior model participants and the average similarity to the target of all prompts generated by inferior model participants, both replayed on the superior model. Figure 2 shows that the model effect (red) accounts for 53% of the ATE ($\Delta CoSim = 0.00864$, $p < 10^{-8}$), whereas the prompting effect (blue) accounts for 47% of the ATE ($\Delta CoSim = 0.00769$, $p = 0.0066$). Thus, nearly half of the ATE is driven by endogenous changes in prompting strategy that emerged after exposure to DALL-E 3. The rightmost point in Figure 2 shows that these endogenous changes in prompting took advantage of the unique capabilities of the superior model; when replaying both sets of prompts on the inferior model, the difference in performance is not statistically significant and is close to zero ($\Delta CoSim = 0.0018$,
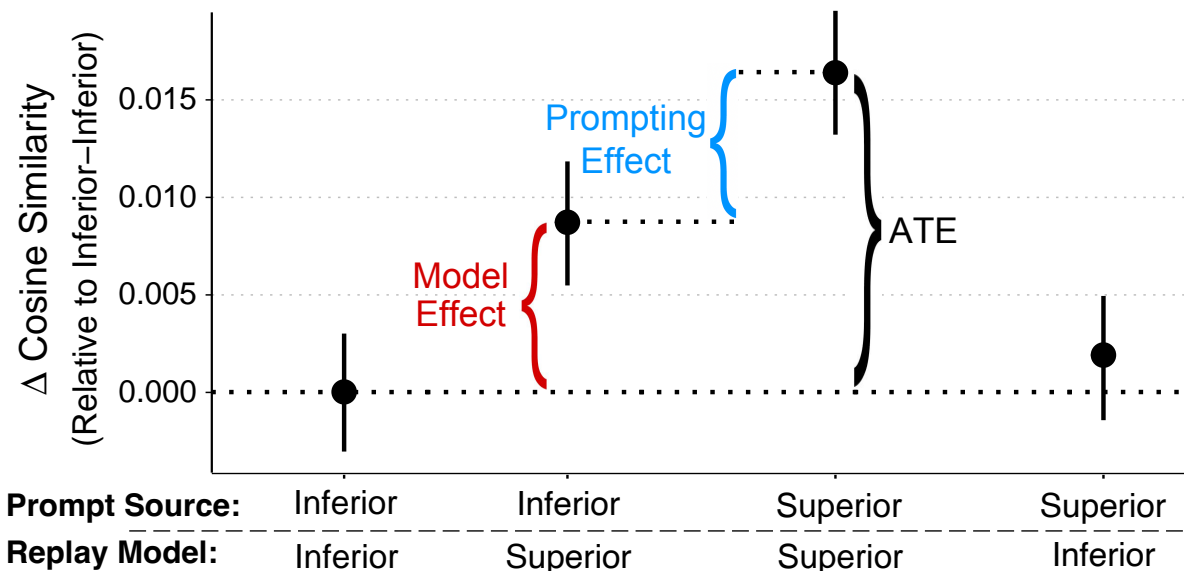
Figure 2: The decomposition of the ATE (black) into model (red) and prompting (blue) effects. The x-axis indicates the model for which the prompts were written (above the dashed line) and the model used for replay (below the dashed line). Effects are shown relative to a baseline of prompts written for and replayed on the inferior model. Error bars show 95% confidence intervals based on bootstrapped standard errors clustered by participant.

$p = 0.52$). In other words, DALL-E 3 was more capable than DALL-E 2 at effectively rendering the additional information contained in the longer and more descriptive prompts (Betker et al., 2023), and participants adapted their prompting behavior accordingly.

One might think that the need to adapt one's prompts to the capabilities of different generative models is obviated by LLM-based, automated prompt revision (Li et al., 2024; Betker et al., 2023). We use data from our third treatment arm—DALL-E 3 with revision—to test this hypothesis. We find that although participants assigned to DALL-E 3 with revisions performed better than those assigned to DALL-E 2 ($\Delta CoSim = 0.0068$; $p = 0.013$), prompt revision reduced the benefit of using DALL-E 3 by nearly 58% (95% CI: [22%, 94%]). Importantly, the positive impact of access to DALL-E 3 with revision was less than the positive impact of directly passing prompts written for DALL-E 2 to DALL-E 3 (i.e., the model effect shown in Figure 2). In other words, prompt revision actually caused participants to perform *worse* in our context. This finding suggests that as currently implemented, AI-assisted prompt revisions are not a panacea, and can actually inhibit people's capacity to leverage a model's capabilities when misaligned with an end user's goals.

## Discussion

In summary, we find that the progression of generative AI models from DALL-E 2 to DALL-E 3 caused performance to increase, not only because of DALL-E 3's superior capabilities (the model effect), but also because of changes in prompting behavior that took advantage of those capabilities

(the prompting effect). More specifically, participants using DALL-E 3 wrote longer prompts that contained more descriptive information, and DALL-E 3 was able to render this additional information more effectively than DALL-E 2. These changes in prompting emerged despite the fact that participants were unaware of which model they were assigned. Although automated LLM-revision features are designed to and have been shown to help people generate higher-quality content (Betker et al., 2023; Li et al., 2024), we also find that LLM-based prompt revision failed to further improve performance; instead, automatic prompt revision reduced the benefits of using DALL-E 3. Overall, we find that shifting to a superior model, whether or not accompanied by LLM-based revisions, did not eliminate the need for prompting.

The primary limitation of our work is that we only studied the transition from DALL-E 2 to DALL-E 3. We leave the study of future transitions and different types of generative AI models to future work. However, if the trends we report generalize and continue into the future, one could imagine a lock-step dynamic where as models continually improve, people continually respond by adapting their prompts to take advantage of the newest model's capabilities. Such a pattern suggests that as generative AI models advance, newer models are unlikely to render prompting obsolete. Rather, prompting will be the mechanism by which people unlock new models' capabilities. Thus, we anticipate that prompting will remain relevant in the future, and that it will be an important determinant of who is best positioned to use and directly benefit from generative AI moving forward.

# References

Acar, O. (2023). Ai prompt engineering isn't the future. *HARVARD BUSINESS REVIEW*.

Adams, G. S., Converse, B. A., Hales, A. H., and Klotz, L. E. (2021). People systematically overlook subtractive changes. *Nature*, 592(7853):258–261.

Betker, J., Goh, G., Jing, L., Brooks, T., Wang, J., Li, L., Ouyang, L., Zhuang, J., Lee, J., Guo, Y., et al. (2023). Improving image generation with better captions. *Computer Science. https://cdn. openai. com/papers/dall-e-3. pdf*, 2(3):8.

Bosker, H. R. (2021). Using fuzzy string matching for automated assessment of listener transcripts in speech intelligibility studies. *Behavior Research Methods*, pages 1–9.

Brynjolfsson, E., Li, D., and Raymond, L. R. (2023). Generative ai at work. Technical report, National Bureau of Economic Research.

Bubeck, S., Chandrasekaran, V., Eldan, R., Gehrke, J., Horvitz, E., Kamar, E., Lee, P., Lee, Y. T., Li, Y., Lundberg, S., Nori, H., Palangi, H., Ribeiro, M. T., and Zhang, Y. (2023). Sparks of artificial general intelligence: Early experiments with gpt-4.

Cui, K. Z., Demirer, M., Jaffe, S., Musolff, L., Peng, S., and Salz, T. (2024). The Productivity Effects of Generative AI: Evidence from a Field Experiment with GitHub Copilot. *An MIT Exploration of Generative AI.*

Dell'Acqua, F., McFowland, E., Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Krayer, L., Candelon, F., and Lakhani, K. R. (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).

Don-Yehiya, S., Choshen, L., and Abend, O. (2023a). Human learning by model feedback: The dynamics of iterative prompting with midjourney.

Don-Yehiya, S., Choshen, L., and Abend, O. (2023b). Human learning by model feedback: The dynamics of iterative prompting with midjourney. In Bouamor, H., Pino, J., and Bali, K., editors, *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 4146–4161, Singapore. Association for Computational Linguistics.

Eloundou, T., Manning, S., Mishkin, P., and Rock, D. (2024). Gpts are gpts: Labor market impact potential of llms. *Science*, 384(6702):1306–1308.

Fu, S., Tamir, N., Sundaram, S., Chai, L., Zhang, R., Dekel, T., and Isola, P. (2023). Dreamsim: Learning new dimensions of human visual similarity using synthetic data. *arXiv preprint arXiv:2306.09344.*

Gilardi, F., Alizadeh, M., and Kubli, M. (2023). Chatgpt outperforms crowd-workers for text-annotation tasks. *arXiv preprint arXiv:2303.15056.*

Harwell, D. (2023). Tech's hottest new job: Ai whisperer. no coding required. `https://www.washingtonpost.com/technology/2023/02/25/prompt-engineers-techs-next-big-job/`. [Accessed 13-05-2024].

Kenter, T. and De Rijke, M. (2015). Short text similarity with word embeddings. In *Proceedings of the 24th ACM international on conference on information and knowledge management*, pages 1411–1420.

Li, C., Zhang, M., Mei, Q., Kong, W., and Bendersky, M. (2024). Learning to rewrite prompts for personalized text generation. In *Proceedings of the ACM on Web Conference 2024*, WWW '24.

Mollick, E. (2023). Working with ai: Two paths to prompting. *One Useful Thing.*

Neelakantan, A., Xu, T., Puri, R., Radford, A., Han, J. M., Tworek, J., Yuan, Q., Tezak, N., Kim, J. W., Hallacy, C., Heidecke, J., Shyam, P., Power, B., Nekoul, T. E., Sastry, G., Krueger, G., Schnurr, D., Such, F. P., Hsu, K., Thompson, M., Khan, T., Sherbakov, T., Jang, J., Welinder, P., and Weng, L. (2022). Text and code embeddings by contrastive pre-training.

Noy, S. and Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654):187–192.

Oppenlaender, J. (2023). A taxonomy of prompt modifiers for text-to-image generation. *Behaviour & Information Technology*, 0(0):1–14.

Quilty-Harper, C. (2023). Ai's hottest job: Prompt engineer. `https://www.bloomberg.com/news/videos/2023-07-05/ai-s-hottest-job-prompt-engineer`. [Accessed 26-10-2023].

Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al. (2021). Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR.

Sävje, F., Higgins, M. J., and Sekhon, J. S. (2021). Generalized full matching. *Political Analysis*, 29(4):423–447.

Schulhoff, S., Ilie, M., Balepur, N., Kahadze, K., Liu, A., Si, C., Li, Y., Gupta, A., Han, H., Schulhoff, S., Dulepet, P. S., Vidyadhara, S., Ki, D., Agrawal, S., Pham, C., Kroiz, G., Li, F., Tao, H., Srivastava, A., Costa, H. D., Gupta, S., Rogers, M. L., Goncearenco, I., Sarli, G., Galynker, I., Peskoff, D., Carpuat, M., White, J., Anadkat, S., Hoyle, A., and Resnik, P. (2024). The prompt report: A systematic survey of prompting techniques.

Singla, N. and Garg, D. (2012a). String matching algorithms and their applicability in various applications.

Singla, N. and Garg, D. (2012b). String matching algorithms and their applicability in various applications. *International journal of soft computing and engineering*, 1(6):218–222.

Torricelli, M., Martino, M., Baronchelli, A., and Aiello, L. M. (2023). The role of interface design on prompt-mediated creativity in generative ai. *arXiv preprint arXiv:2312.00233*.

van der Maaten, L. and Hinton, G. (2008a). Visualizing data using t-sne. *Journal of Machine Learning Research*, 9(86):2579–2605.

van der Maaten, L. and Hinton, G. E. (2008b). Visualizing data using t-sne. *Journal of Machine Learning Research*, 9:2579–2605.

Wei, J., Wang, X., Schuurmans, D., Bosma, M., Xia, F., Chi, E., Le, Q. V., Zhou, D., et al. (2022). Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35:24824–24837.

Xie, Y., Pan, Z., Ma, J., Jie, L., and Mei, Q. (2023). A prompt log analysis of text-to-image generation systems. In *Proceedings of the ACM Web Conference 2023*, WWW '23, page 3892–3902.

# Supplementary Information (SI)

## A    Experiment Design

### A.1    Task Design

Participants were asked to reproduce a single target image as closely as possible using a text-to-image generative AI model (i.e., DALL-E 2, DALL-E 3 without prompt revision, or DALL-E 3 with prompt revision, all developed by OpenAI). They did so by successively submitting prompts. In response to each submitted prompt, the model would generate an image, which was then displayed to the participant next to their assigned target image. Participants were instructed to make at least 10 attempts at trying to recreate the target image within a 25-minute window, with no upper limit on their number of attempts.

All interactions between participants and the generative AI models occurred on a custom-built online interface designed to resemble OpenAI's ChatGPT interface but with some adjustments related to our task (e.g., displaying the target image and the total number of attempts so far to the user). On the right-hand side of the interface, participants were shown the target image they were randomly assigned to recreate. On the left-hand side, participants were shown their previously submitted prompts as well as the resulting generated images. We placed the text box where participants were able to write and submit their prompts at the bottom of the interface. Prompts were limited to a maximum of 1,000 characters. Participants were informed that their interactions with their assigned model would be memory-less, i.e., the model retained no memory of previous prompts and only used the current prompt to generate each image. Before the task, participants were provided with written and video instructions on how to interact with our experiment interface. Our task did not assume nor require prior experience with *any* generative AI tools.

After the task, we surveyed participants' opinions and preferences regarding generative AI tools. We also inquired about their self-assessed occupational skills and how often they 1) engaged in creative writing, 2) wrote specific instructions, and 3) engaged in any sort of computer programming. Finally, we collected socio-demographic data, such as age, gender, and occupation.

### A.2    Randomization

We randomized participants across two dimensions: the target image and the text-to-image generative AI model that participants had access to. We randomized participants across both dimensions simultaneously using complete randomization, generating 45 possible target image-model cells. We conducted a balance check after the conclusion of the experiment with a $\chi^2$ test for homogeneity of proportions across all cells. With $\chi^2 = 0.374$, $df = 28$, the resulting p-value equals to 1 and thus we cannot reject the null hypothesis that the proportions are equal across all 45 groups:

$$H_0 : p_1 = p_2 = \cdots = p_{45}$$

Participants were unaware of this randomization.

### A.2.1 Generative Models

We randomly assigned participants to 1 of 3 generative models:

1. DALL-E 2, which is referred to at points in the main text as the inferior model.

2. DALL-E 3 (Verbatim), which is referred to at points in the main text as the superior model.

3. DALL-E 3 (Revised), which is referred to at points in the main text as DALL-E 3 with Revision

Both the "verbatim" and "revised" versions of the DALL-E 3 treatment utilize the same underlying image-generating model; the distinction lies in the pre-processing applied before submitting user prompts to OpenAI's image-generating API. OpenAI's DALL-E 3 system, by design, employs a GPT-4 model to rewrite user prompts, adding more detail before processing the modified prompt using the DALL-E image-generating model. During our experiment, it was not possible to explicitly disable this prompt rewriting feature of the DALL-E 3 system. To manage this behavior, we defined two treatments utilizing the DALL-E 3 model.

In the DALL-E 3 (Revised) treatment arm, we submit the participant's prompt directly to OpenAI's API and do not interfere with the default prompt rewriting process. In the DALL-E 3 (Verbatim) treatment arm, we prepend a string instructing the GPT-4 model to not modify the participant's prompt before passing it forward to DALL-E 3. This string is never visible to participants and was modeled after a prefix specifically suggested in OpenAI's online documentation for the DALL-E 3 endpoint.[4] We modified the recommended prefix slightly to account for the fact that we did not expect our participants to always submit "extremely simple prompts." The string we prepended to prompts is found below:

> "I NEED to test how the tool works with my prompt as it is written. DO NOT add any detail; just use it AS IS:"

Prepending this string to participants' prompt did reduce the rate at which OpenAI's endpoint modified prompts, but compliance was not perfect. Thus, we view the "verbatim" treatment arm as more of an intent-to-treat intervention. The GPT model still modified 59% of participant prompts. The average token sort ratio (TSR) between the original prompt and the modified prompt was 77 for the DALL-E 3 (Verbatim) arm, compared to an average token sort ratio of 44 across the entire DALL-E 3 (Revised) treatment arm (a TSR of 100 denotes an exact string match). Conditional on any modification (any observations with TSR < 100), the average TSR between the original prompt and the modified prompt was 61 for the DALL-E 3 (Verbatim) arm, compared to an average TSR of 44 across the entire DALL-E 3 (Revised) treatment arm.

---

[4]See here and here for online documentation.

### A.2.2 Target Images

We randomly assigned participants to 1 of 15 target images. The set of target images consisted of 5 images each from 3 different broad categories: business and marketing, graphic design, and photography. We chose these to represent the use cases suggested by the prompt categories on https://promptbase.com/, a leading marketplace for image generation prompts. The images vary in color, style, content, and complexity within and across categories. These images can be found online linked to the pre-registration document: https://osf.io/ejbtp. As we discuss in Section F, performance, and variability of performance varied substantially across images. In other words, some images were much easier than others to replicate with the generative models, which we view as additional evidence that the set of 15 images was reasonably diverse.

## A.3 Subjects

Our Prolific-recruited US sample (N = 2,059) was limited to fluent English speakers, and we prevented participants from completing the task more than once. We also prevented users from completing the task on mobile devices or tablets. Data was collected between December 12, 2023 and December 19, 2023. Participants were guaranteed a payment of $4 USD for completing the task and could earn an additional $8 USD (a 200% bonus) if they ranked in the top 20% of participants in DreamSim of their image most similar to the target (construction of DreamSim is described in section B.5.1). The median time to complete our entire task, including a demographic survey, was 22 minutes. Given that 20% of subjects received a bonus, the average compensation for participants in our study was $5.60 USD per person, or about $15 USD per hour. We explained the payment and incentive scheme to participants in full multiple times during the onboarding phase of the experiment, and asked participants to confirm their understanding before they were allowed to complete the task. The onboarding process also included multiple attention checks; participants who failed the first check were immediately disqualified. For subsequent checks, participants were required to retry until they demonstrated understanding.

## A.4 Model Endpoints

We used the following model endpoints and parameters to generate images from prompts:

1. **OpenAI API:** We used the image generation endpoint of the official OpenAI Node.js library to generate images for user prompts during the experiment. For all treatment arms, we set the image size parameter to be 1024 x 1024 pixels. For the DALL-E 3 (Revised) and DALL-E 3 (Verbatim) treatment arms, we set the quality parameter to standard and the style parameter to natural.

2. **Azure OpenAI Service:** We used the image generation endpoints in the Python implementation of Azure OpenAI Service to generate all replay images based off user prompts collected

during the experiment. For prompts replayed through the DALL-E 2 treatment arm, we deployed a set of DALL-E 2 models on Azure OpenAI Service and set the API version for each to the `2023-06-01-preview` version. For DALL-E 2, we created replay images in batches of 5. For prompts replayed through the DALL-E 3 (Revised) and DALL-E 3 (Verbatim) treatment arms, we deployed a set of DALL-E 3 models on Azure OpenAI Service and set the API version for each to the `2023-12-01-preview` version. The parameter values for image size, and quality and style for the DALL-E 3 treatment arms, were set to the same values as in the experiment.

# B  Measurement and Variables

## B.1  Survey Data

For each participant, we used a Qualtrics survey to collect demographic information, information on the participant's skills that may be relevant to generative AI use, and information on the participant's attitude towards generative AI. This data includes:

- **Demographics:** Ethnicity, Gender, Age, Highest level of education attained (some high school, high school, some college, associate's degree, bachelor's degree, master's degree, doctoral degree, professional degree, other), Years of work experience, Annual Income (0-$25k, $25.001k-$50k, $50.001k-75k, $75.001k-$100k, $100.001k-$150k, $150k+), and elicitation of sets of O*NET job skills that participants used in their occupation (reading comprehension, active listening, writing, speaking, critical thinking, social perceptiveness, coordination, instructing, programming, judgment and decision making, systems evaluations, science, active learning, learning strategies, monitoring, complex problem analysis, technology design, troubleshooting, quality control analysis, systems analysis).

- **Opinions and Skills:** Computer programming proficiency and usage frequency (self-reported), Structured and creative writing proficiency and usage frequency (self-reported), Generative AI tool proficiency and usage frequency (self-reported), Attitudes towards net social impact of Generative AI (self-reported), Advice for (hypothetical) future participants on how to perform well on the task.

## B.2  Prompt Data

For each prompt, we record the text of the participant's prompt, the order in which it was submitted, the timestamp of submission, and for the DALL-E 3 treatment arms, the revised prompt returned by the model.

## B.3  Image Data

For each prompt, the following images were collected:

1. **The participant-facing images (OpenAI API endpoint):** The image shown to the participant during the experiment, generated by the model they were assigned to using the prompt they submitted. These images were generated from December 12-19, 2023.

2. **Post-hoc resampled images (Azure OpenAI endpoint):** For any given prompt, the output of the text-to-image model is stochastic. To better approximate the expected image from a given prompt, we generated 20 additional images for each prompt after the experiment concluded. We provide full details on this procedure in Section C. These images were generated from December 26, 2023 - January 27, 2024. These images are not used for analyses presented in the main text, but were used for other pre-registered analyses. These additional analyses are discussed in Section F.

3. **Post-hoc replayed images (Azure OpenAI endpoint):** To decompose our overall effects into model and prompting effects, we generated "counterfactual images" for each prompt written under the DALL-E 2 and DALL-E 3 (Verbatim) treatments. In other words, we submitted all prompts written under both the DALL-E 2 and DALL-E 3 (Verbatim) treatments to both the DALL-E 2 and DALL-E 3 (Verbatim) endpoints. Similarly to the resampling procedure outlined above, we generated 10 images per prompt per model: we generated a single replay for each prompt-model pair from March 16-18, 2024, and then, to increase power, generated the replications for these replay images from June 14-27, 2024. This replay process produced a total of 20 images per prompt—10 under the original model, 10 under the counterfactual model. We re-submitted prompts to their original model to account for potential model drift, as this exploratory analysis was conducted multiple months after our initial data collection. For consistency, this replay data is used throughout the main text of our paper.

## B.4  Sample Construction

The sample we analyze in the main text of our paper is constructed using the process described as follows.

- The initial "raw" dataset collected during the experiment is comprised of 24,672 rows of raw prompt data (one prompt per row) generated by 2,059 participants.

- We first removed rows with blank prompt entries, invalid prolific IDs, and unsuccessful attempts (logging errors). These exclusion criteria were pre-registered. This left us with 2,029 participants and 24,123 prompts.

- We next removed participants from our sample if they failed to submit at least 10 prompts or if they did not make a "good-faith" effort to complete the task as requested. Both of these exclusion criteria were pre-registered. A good-faith effort was not made if a participant submitted the same prompt at least five times in a row at any point during the task. These exclusion criteria were also explained to participants, who were told that payment was contingent on submitting at least 10 successful prompts and a "good-faith effort." To avoid reward

hacking, we did not specify the "no more than 5 repeated prompts" criterion for "good-faith effort." This left us with 1,899 participants.

- Although participants were allowed to submit as many prompts as they desired in the 25-minute time span, we limited all analyses to each participant's first 10 prompts—the minimum required to receive payment for the task. This exclusion criteria was not pre-registered, and is noted in the list of deviations from pre-registration in Section F.F.2. We restrict our analysis dataset in this way because participants who chose to submit more than 10 prompts may have been systematically different than those who did not. Excluding any prompt beyond the $10^{\text{th}}$ attempt allows us to alleviate selection bias concerns. This left us with 18,990 prompt observations from 1,899 participants.

- We next removed participants who failed to complete the Qualtrics survey. This exclusion criteria was pre-registered. This left us with 1,893 participants and 18,930 prompts.

- We also removed prompts from our dataset according to a number of post-hoc, non-pre-registered exclusion criteria to ensure data quality and avoid selection bias. If a prompt had any of the following flags, it was removed from the sample:

  - Prompts sometimes trigger errors in OpenAI's safety system because they contain language that might be deemed unsafe under openAI's policies. The specific language that triggers these errors is constantly changing and not available publicly. If a prompt triggered a safety error during the replication or replay process, we re-submitted the prompt up to 50 times or until the 10 original arm replications/replay samples had been collected. We removed prompts if they failed to generate 10 replications on the original model or 10 replay samples under the counterfactual model during the replication/replay process. This affected 305 prompts between the DALL-E 2 and DALL-E 3 (Verbatim) treatment arms. It did not affect any DALL-E 3 (Revised) prompts, as we did not conduct replay analysis with the prompts from this treatment arm.

  - Due to rare latency issues, some prompts were assigned duplicate attempt numbers by the mongoDB database that we used to collect our data. This data collection error led to issues in the data analysis process. Thus, we excluded prompts with duplicate attempt numbers. This affected 34 prompts across all three treatment arms, and 20 prompts between the DALL-E 2 and DALL-E 3 (Verbatim) treatment arms, approximately 0.1% of the original data.

  - While analyzing our data, we found that our sample contained a number of prompts that were not "good-faith efforts" to complete our task, despite the use of the exclusion criteria described above. Thus, we used the following process to identify and remove additional "low-effort" prompts. First, we generated embeddings for each prompt using OpenAI's `text-embedding-3-small` model. We then calculated the mean embedding for each target image. Next, we calculated the Euclidean distance between each prompt's

14

embedding vector and the mean embedding vector for prompts corresponding to the focal prompt's assigned target image. Finally, we removed the 2.5% of prompts that were most distant from the mean image-level prompt embedding vector. We observed that the prompts that were dropped were often off-topic and likely were not good-faith efforts to complete our task. This led to the removal of 481 prompts across all three treatment arms, and 338 prompts between the DALL-E 2 and DALL-E 3 (Verbatim) treatment arms. Our results are robust to the inclusion of these filtered prompts.

- Our final sample included 1,891 participants and 18,152 prompts.

## B.5  Dependent Variables

### B.5.1  Image Similarity

We pre-registered two quantitative measures of image similarity: the cosine similarity of CLIP embedding vectors and a recently developed measure called 'DreamSim' (Fu et al., 2023). In the main text, we present analyses using CLIP embedding cosine similarity, since it is likely more familiar to readers. Our results are qualitatively and quantitatively similar using DreamSim instead.

- **CLIP Embedding Cosine Similarity:** To calculate CLIP embedding cosine similarity, we first generated CLIP embedding vectors (Radford et al., 2021) from Hugging Face (**?**) for each participant-generated image and for each target image. Unlike traditional image embeddings that only encode visual features, CLIP embeddings also capture semantic relationships between images and descriptive text. We then calculated the cosine similarity between each participant-generated image's CLIP embedding and the relevant target image's CLIP embedding.

- **DreamSim:** DreamSim is an image similarity measure proposed recently by (Fu et al., 2023). The authors claim that relative to something like CLIP embedding cosine similarity, DreamSim measures image similarity in a way that more effectively captures human visual perceptions of similarity. Because the original DreamSim metric outputs a distance measure, we invert this score $\tilde{D} = 1 - (\text{original DreamSim})$ to recast it as a similarity score. After doing so, both the inverted DreamSim and CLIP embedding cosine similarity are closer to 1 when two images are more similar and closer to 0 when two images are more dissimilar.

We find that these two measures of image similarity are highly correlated in our sample ($\rho_{pearson} = 0.763$, 95% CI: [0.755 0.770]), and our main results are robust to the use of either measure. We present the results obtained when conducting our main text analyses using DreamSim in Section E.E.1.

### B.5.2  Prompt Length

We measure the lengths of prompts written by participants in our sample, both in terms of the number of *words* in a given prompt and in terms of the number of *characters* in a given prompt.

In our main text analysis, we present results only in terms of the number of words, since the two outcomes are highly correlated ($\rho_{pearson} = 0.9952$, 95% CI: [0.99507, 0.99543]).

### B.5.3 Embedding-based Prompt Similarity

We calculate two measures of embedding-based prompt similarity: successive similarity and aggregate similarity. Both measures use the vector embedding representation of each prompt in our sample, which we obtained using OpenAI's `text-embedding-3-small` model (Neelakantan et al., 2022). The two similarity measures are defined as follows:

- **Successive similarity:** The successive similarity ($ss$) is a measure of the similarity of a participant's prompt to their immediately preceding prompt. We define the successive similarity of a prompt $p_{i,n}$ written by user $i$ to the their immediately preceding prompt $p_{i,n-1}$ as:

$$ss_{i,n,n-1} = \frac{\mathbf{E}(\mathbf{p_{i,n}}) \cdot \mathbf{E}(\mathbf{p_{i,n-1}})}{||\mathbf{E}(\mathbf{p_{i,n}})|| \, ||\mathbf{E}(\mathbf{p_{i,n-1}})||}, \tag{1}$$

  where $\mathbf{E}(\mathbf{p_{i,n}})$ is the vector embedding representation of participant $i$'s $n^{\text{th}}$ prompt, $p_{i,n}$. This measure starts with participant $i$'s 2$^{\text{nd}}$ attempt, as the calculation requires a previous attempt.

- **Aggregate similarity:** The aggregate similarity ($as$) is a measure of how dispersed each user's prompts are around their "average prompt" (calculated by taking the element-wise average of all prompt embeddings produced by the user). We define the aggregate similarity for the 10 prompts written by a given user as:

$$as_i = \frac{1}{10} \sum_{n=1}^{10} ||\mathbf{E}(\mathbf{p_{i,n}}) - \overline{\mathbf{E}(\mathbf{p_{i,n}})}||_2^2, \tag{2}$$

  where $\mathbf{E}(\mathbf{p_{i,n}})$ is again the vector embedding representation of participant $i$'s $n^{\text{th}}$ prompt, $p_{i,n}$, and $\overline{\mathbf{E}(\mathbf{p_{i,n}})}$ is the element-wise mean of all 10 of participant $i$'s prompts.

### B.5.4 Successive Prompt Token Sort Ratio

Starting with each participant's second prompt, we also calculated the token sort ratio (TSR) of each prompt $p_{i,n}$ to the immediately preceding prompt $p_{i,n-1}$. TSR is a fuzzy string-matching technique (Singla and Garg, 2012b) that provides a continuous measure of how similar two strings are. We refer the reader to (Bosker, 2021) for a more in-depth description of how TSR is calculated.

### B.5.5 Successive Prompt 'Contains Previous Prompt' Dummy

Starting with each participant's second prompt, we record whether each prompt $p_{i,n}$ contains the immediately preceding prompt $p_{i,n-1}$ as an exact substring.

### B.5.6 Inversions of Successive Measures

In some of our pre-registered analyses, we use inverted versions of the measures in Sections B.E.3, B.E.4, and B.E.5 such that they are forward-looking instead of backward-looking. This means that instead of comparing the similarity of a prompt to the prompt immediately *preceding* it, we compare the similarity of a prompt to the prompt immediately *succeeding* it. The succeeding measurements are identical to all their preceding counterparts in the above equations, except that $n - 1$ is replace with $n + 1$, and measurement begins with each participant's 1$^{\text{st}}$ attempt and ends with a participant's 9$^{\text{th}}$ attempt. To refer to these measurements, we simply called them the "inverted" version of the above equations, i.e., inverted successive similarity, inverted aggregate similarity, inverted successive prompt token sort ratio, and inverted successive prompt 'contains previous prompt' dummy.

### B.5.7 Prompt Composition

We use the `spaCy v3.7.4` Python package's `en_core_web_sm` model to tag the parts of speech (POS) in each prompt. SpaCy's models utilize the "universal POS tags" from the Universal Dependencies framework for grammar annotation **?**. These tags encompass parts of speech such as adjectives, adverbs, nouns, and verbs. The model tags each word in a prompt according to this framework, after which we count the total number of words corresponding to each part of speech for each prompt.

### B.5.8 Strategic Shifts

In addition to calculating the successive and aggregate similarity of prompts written by particular users, we also attempt to identify particular moments when participants shift their approach to prompting. In order to do so, we adapt a method proposed in (Torricelli et al., 2023) (because they are conducting research in a different context, (Torricelli et al., 2023) refer to these shifts as "topical transitions" as opposed to "strategic shifts"). To identify these strategic shifts, we first calculate the mean cosine similarity ($MCS$) for the embedding vectors of every possible pair of prompts submitted in response to a given target image, $t$:

$$MCS_t = \frac{2}{P_t(P_t - 1)} \sum_{a=1}^{P_t} \sum_{b=a+1}^{P_t-1} \text{CosineSim}(\mathbf{E}(\mathbf{p_{a,t}}), \mathbf{E}(\mathbf{p_{b,t}})). \tag{3}$$

where $P_t$ is the total number of prompts submitted in response to a given target image, and $a$ and $b$ are indices representing individual prompts for that target.

We then label any given prompt as a strategic shift ($SS$) if the cosine similarity of its embedding vector with that of the previous prompt is lower than this target-image-level mean:

$$SS(\mathbf{p}_{i,n,t}) = \begin{cases} 1 & \text{if CosineSim}(\mathbf{E}(\mathbf{p_{i,n,t}}), \mathbf{E}(\mathbf{p_{i-1,n,t}})) < MCS_t \\ 0 & \text{otherwise} \end{cases} \tag{4}$$

It is worth noting that Torricelli et al. (2023) uses the participant-level mean, as opposed to the task-level mean, as the cutoff for a topical shift. We instead use the task-level mean because in our setting, it did not seem appropriate that half of each participant's submitted prompts would be strategic shifts.

## C   Methods

### C.1   Stratification

The results shown in the main text and in the SI are mostly stratified by reference image and iteration. In some analyses, we have stratified only on the reference image (e.g., for analyses presented at the iteration level). The exact stratification for each finding is indicated in section D. To stratify our results, we take a weighted average across $j = 1, ..., J$ cells defined by our stratification variables.

$$\overline{Y}_{strat} = \sum_{j=1}^{J} \frac{N_j}{N} \bar{Y}_j$$

To calculate the variance (and standard error) of this sample mean we apply the following:

$$\widehat{\text{Var}}(\overline{Y}_{strat}) = \widehat{\text{Var}}\left(\sum_{j=1}^{J} \frac{N_j}{N} \bar{Y}_j\right) = \sum_{j=1}^{J} \left(\frac{N_j}{N}\right)^2 \frac{s_j^2}{N_j}$$

where:

- $N_j$ is the population size of stratum $j$.

- $N$ is the total population size across all strata.

- $\bar{Y}_j$ is the sample mean for stratum $j$.

- $J$ is the total number of strata.

- $s_j$ is the sample standard deviation of stratum $j$. Therefore, $s_j^2$ is sample variance stratum $j$.

### C.2   Z-Scoring

We find statistically significant evidence for differences in the variability of performance across the 15 target images used in our experiments, which we discuss in Section F. In the main text, we also showed that performance increases across the attempts. To test whether our results are in some way due to this image-level or attempt-level variations, we replicate all analyses using the

within-image-attempt Z-score of CLIP-cosine similarity of each image produced by participants in our experiment. Formally, this is:

$$Z(\text{CosineSim}_{i,n,t}) = \frac{\text{CosineSim}_{i,n,t} - \text{Mean}_{n,t}\left(\text{CosineSim}_{i,n,t}\right)}{\text{SD}_{n,t}\left(\text{CosineSim}_{i,n,t}\right)}, \tag{5}$$

where $\text{CosineSim}_{i,n,t}$ is the cosine similarity of user $i$'s image in attempt $n$ to target image $t$. The mean and standard deviation are computed per each image-attempt, but over the DALL-E 2 and DALL-E 3 treatment arms. We also use the rescaling above to test the robustness of our DreamSim-based analyses.

## C.3   Accounting for Model Stochasticity

The output that generative AI models return in response to a given prompt is stochastic. The strength of this stochasticity is controlled by a model parameter referred to as the temperature, which could not be edited using the DALL-E 3 API at the time of our experiment. To account for this model stochasticity, we generated 10 images for each prompt submitted by participants for all arms. We were then able to calculate the similarity between each replication and its corresponding target image, and calculate an "expected" CLIP cosine similarity and DreamSim score for each prompt by averaging over these samples.

We generated these additional samples for both the original prompts on their assigned treatment arms, as well as replaying on the counterfactual arms, as introduced in Figure 2 in the main text. Importantly, OpenAI updated its content filters between our initial experiment and image resampling. As a result, some prompts that originally produced images either generated no images or fewer images than requested during our regeneration attempts. This affected 1.8% (371 out of 18,990 prompts) of the data in our sample under the under the "replaying" procedure (Section B.B.3.3).

# D   Main Text Analyses

## D.1   Task Performance and ATEs

The top pane of Figure 1B compares the average performance across models and attempt numbers (also referred to as iterations). It shows the average cosine similarity score stratified by the reference image. A notable feature in this figure is the performance dip during the second recreation attempt across both treatment arms. This is likely due to participants' initial misunderstanding of the model's "memoryless" nature. Participants failed to recognize that context from previous prompts was not carried over to new iterations. We observed numerous prompts in the second iteration across users that explicitly referenced the first prompt, a behavior that rarely occurred in subsequent attempts. However, from the third prompt onward, participants appeared to grasp the independence of each attempt, as evidenced by a marked decrease in cross-prompt references and a corresponding rebound in performance.

Next, the bottom pane of Figure 1B shows the average treatment effect (ATE) per iteration, which is the difference between the stratified averages of DALL-E 3 and DALL-E 2 in the top pane.[5] To test the widening impact of using DALL-E 3 on performance relative to DALL-E 2, we also run the following fixed effects linear model with participant-level ($i$) clustered standard errors:

$$Y_{i,n,t} = \beta_0 + \beta_1 \text{iteration} + \beta_2 \mathbb{I}[\text{dalleVersion} = 3]_i + \beta_3 \text{iteration} \times \mathbb{I}[\text{dalleVersion} = 3]_i + \gamma_t + \epsilon_{i,n,t} \quad (6)$$

The coefficient estimates generated by this model are:

- $\hat{\beta}_1 = 0.0015$, $\hat{SE}(\beta_1) = 0.0003$, $p = 9.678 \times 10^{-7}$

- $\hat{\beta}_2 = 0.0126$, $\hat{SE}(\beta_2) = 0.003$, $p = 1.7 \times 10^{-4}$

- $\hat{\beta}_3 = 0.00086$, $\hat{SE}(\beta_3) = 0.0004$, $p = 0.0418$

The overall ATEs that we report between different pairs of treatment arms (DALL-E 2, DALL-E 3, and DALL-E 3 with revisions) in the main text are estimated from a two-way fixed effect (iteration and target image) model per each pair. Standard errors are cluster robust at the participant level.

## D.2   Prompt Characteristics

Figure 1C compares the prompt length and prompt similarity of the two models. To generate these results, we first remove any prompt that does not resemble a good-faith attempt according to the sample construction procedure detailed in Section B.B.4. The prompt length is the average number of words per model and iteration stratified by the reference image. The prompt similarity is the average cosine similarity between all consecutive pairs of user prompts, which are both determined to be valid attempts, stratified by the reference image (see Section B.B.5.3 for details on similarity calculations). The color scale in figure 1C shows the stratified average similarity to the previous prompt across all users per each model. We find that superior model users write prompts that, on average, have $\beta = 0.0166$ higher in cosine similarity to their previous prompts using cluster robust standard errors at the participant level ($p = 0.02$).

Comparing the aggregate similarity of all attempts made by a given participant, we also find the prompts from the superior model participants were, on average, more similar than the prompts of the inferior model participants. For this analysis, we use the dispersion around the centroid in the prompt embedding space, explained in Section B.B.5.3, as the dependent variable. When we average across all participants by model, we find that the average distance of prompts written by superior model participants to their centroid is $\beta = -0.0169$ smaller than inferior users ($p = 0.008$). Standard errors are cluster-robust at the participant level.

---

[5]In Section D, when we refer to "DALL-E 3", we mean "DALL-E 3 (Verbatim)" unless otherwise specified.

## D.3 ATE Decomposition

Figure 2 in the main text decomposes the ATE into the model and prompting effects. This decomposition is conceptually similar to a simple mediation analysis, with an important difference being that we can observe counterfactual outcomes (e.g., prompting the superior model as if it is the inferior model). This is not typically the case in mediation analysis, and makes causal identification rely on fewer assumptions. To obtain counterfactual outcomes, we fed or "replayed" the participant prompts when interacting with one model (e.g., inferior) on another model (e.g., superior). The notation *(prompt, model)* specifies which treatment arm the prompts were written under and which model was used in the replay. For example, (2,3) indicates replaying prompts written under DALL-E 2 on DALL-E 3. To be clear, (2,2) and (3,3) correspond to the original observed treatment arms, while (2,3) and (3,2) are the counterfactual outcomes of interest.[6]

The left-most point in Figure 2 corresponds to the average CLIP cosine similarity to the target image of (2,2). To make the interpretation of the results clearer, we have subtracted this quantity from all average quality scores and added a dashed line throughout. The second point from the left corresponds to average similarity to the target of (2,3), the third point from the left to (3,3), and the rightmost points to (3,2). All average similarity scores are stratified by iteration and reference image, and the standard errors are bootstrapped and cluster-robust at the participant level. The model effect, as shown by the red braces in Figure 2, corresponds to the average increase in quality of (2,3) relative to (2,2). In the terminology of mediation analysis, the model effect would be referred to as the direct effect. The prompting effect, as shown by the blue braces in Figure 2, corresponds to the average increase in quality of (2,3) relative to (3,3). In the terminology of mediation analysis, the prompting effect would be referred to as the indirect effect. We can also test the difference in average quality between (3,3) and (3,2), as well as the difference between (3,2) and (2,2). Both of these differences are visible in Figure 2; the second is small and not statistically significant.

The standard errors in Figure 2 correspond to the uncertainty around the estimated average score for each of the four replay conditions. These uncertainty estimates are insufficient for exact inference on the direct, indirect, and treatment effects. The statistics and significance values report in the main text, which correspond to such effects (i.e., the difference between average estimates in two conditions) are obtained using a two-way (iteration and target image) fixed effect model with the effect type as the main independent variable:

$$Y_{i,n,t} = \beta_0 + \beta_1 \text{effect} + \alpha_n + \gamma_t + \epsilon_{i,n,t} \tag{7}$$

where $\beta_1$ is the coefficient on the effect type in question (i.e., model or prompting). To estimate the different effects, we simply use the above model and filter the data as appropriate. For example, to estimate the ATE, the data contains all (2,2) and (3,3) scores, and in this case effect=1 for observations in (3,3) group. Similarly, to estimate the direct or model effect, the data contains all

---

[6]To avoid problems with model drift, we regenerated images for all four possible combinations at the same time and used these images for all analyses in the main text.

(2,2) and (2,3) scores and effect=1 for observations in (2,3) group. Finally, to estimate the indirect or prompting effect, the data contains all (2,3) and (3,3) scores and effect=1 for observations in (3,3) group. The standard errors for each estimated model are cluster robust at the participant level, and p-values are adjusted accordingly.

# E    Robustness Checks

## E.1    DreamSim-based Analysis

As discussed in Section B, we repeat all main-text analyses with DreamSim. The results of these analyses are reported below.

### E.1.1    Overall ATEs

In terms of DreamSim, participants using DALL-E 3 (the superior model) produced images that were, on average, $z = 0.261$ standard deviations (95% CI = [0.180, 0.342]) closer to the target image ($\Delta DreamSim = 0.0304$, $p < 10^{-9}$) than those produced by participants using DALL-E 2 (the inferior model).

### E.1.2    Figure 1

We reran the regression in Section D.D.1 with $Y_{i,n,t}$ representing the DreamSim outcome instead of cosine similarity:

$$Y_{i,n,t} = \beta_0 + \beta_1 \text{iteration} + \beta_2 \mathbb{I}[\text{dalleVersion} = 3]_i + \beta_3 \text{iteration} \times \mathbb{I}[\text{dalleVersion} = 3]_i + \gamma_t + \epsilon_{i,n,t} \quad (8)$$

The coefficient estimates generated by this analysis are:

- $\hat{\beta}_1 = 0.0039$, $\hat{SE}(\beta_1) = 0.00048$, $p = 4.3 \times 10^{-16}$

- $\hat{\beta}_2 = 0.0179$, $\hat{SE}(\beta_2) = 0.0060$, $p = 0.0029$

- $\hat{\beta}_3 = 0.0023$, $\hat{SE}(\beta_3) = 0.00071$, $p = 0.001$

### E.1.3    Figure 2

Decomposing the ATE as measured in terms of DreamSim, we find similar results to those in the main text. The model effect accounts for 55% of the ATE ($\Delta DreamSim = 0.0168$, $p < 10^{-10}$), whereas the prompting effect accounts for 45% of the ATE ($\Delta CoSim = 0.01364$, $p = 0.0053$).

## E.2 Z-score-based Analysis

As we discuss in Section B, we repeat all main-text analyses with the within-image Z score of CLIP-cosine similarity to better account for variation between images. The results hold across the board and sometimes, the differences between the superior and inferior model are even starker than when using cosine similarity.

### E.2.1 Overall ATEs

As mentioned in the main text, participants using DALL-E 3 (the superior model) produced images that were, on average, $z = 0.20$ standard deviations (obtained from ATE in terms of Z-Scored Cosine Sim $= 0.20$, 95% CI $= [0.126, 0.284]$) closer to the target image ($\Delta CoSim = 0.0164$, $p < 10^{-7}$) than those produced by participants using DALL-E 2 (the inferior model). Standard errors are cluster robust by participants.

### E.2.2 Figure 1

On average, participants using the superior model produced images that were $z = 0.20$ standard deviations closer (the ATE) to the target image than those using the inferior model. Like in the main text with CLIP cosine similarity, this treatment effect increased as participants made successive attempts to replicate the target image.

$$\text{ZScore}_{i,n} = \beta_0 + \beta_1 \text{iteration} + \beta_2 \mathbb{I}[\text{dalleVersion} = 3]_i + \beta_3 \text{iteration} \times \mathbb{I}[\text{dalleVersion} = 3]_i + \epsilon_{i,n} \quad (9)$$

- $\hat{\beta}_1 = 0.02012$, $\hat{SE}(\beta_1) = 0.0041$, $p = 1.19 \times 10^{-6}$

- $\hat{\beta}_2 = 0.1439$, $\hat{SE}(\beta_2) = 0.0466$, $p = 0.002$

- $\hat{\beta}_3 = 0.01253$, $\hat{SE}(\beta_3) = 0.0058$, $p = 0.03$

### E.2.3 Figure 2

When we decompose the ATE into the model effect ($z = 0.0895$; $p = 1.02 \times 10^{-5}$) and prompting effect ($z = 0.1139$; $p = 0.0044$), they account for 44% and 56% of the treatment effect, respectively. And when we replay the inferior model prompts on the superior model, the difference in similarity to the target from these prompts played on the inferior is not statistically significant and close together ($z = -0.03$; $p = 0.43$). In short, Figure 2 in the main text is quantitatively and qualitatively unchanged when using the within-image Z score of the cosine similarity.

## F  Pre-registration

Prior to our experiment, we pre-registered a number of hypotheses and a pre-analysis plan. This pre-registration is deposited at OSF at the following URL: https://osf.io/ejbtp. The main text of

this paper contains a subset of our pre-registered analysis, as well as complementary exploratory analyses that are also mentioned in our pre-registration. We chose to present this subset of our pre-registered analyses because we believe this subset constitutes an important and timely set of results best-suited to a short-form paper.

Below, we provide a high-level description of all of our pre-pregistered analyses, along with a description of the results of those analyses. In our pre-registration, we declared the intent to conduct each of our analyses using eight possible outcome variables, all of which are different transformations of the same underlying data: CLIP embedding cosine similarity and DreamSim, both of which rescaled in all of four ways:

- **No rescaling**: The outcome variable is used as-is.

- **Z-score rescaling**: We rescale the outcome variable into a Z-score according to the procedure describe in these supplementary materials.

- **Percentile rank rescaling**: We rescale the outcome variable into a percentile rank. This is done by calculating the percentile rank of a given prompt relative to all other prompts submitted for the relevant target image.

- **GPT-4V baseline rescaling**: We rescale the outcome variable by measuring its distance to the outcome variable obtained using a prompt generated by GPT-4V in response to the target image.

We intend to complete each pre-registered analysis with each of these eight possible outcome variables. Although we have not yet done so, we have no reason to believe that our results will not be robust to these different transformations of our data. However, we did not want to delay the publication of our preprint until all of these analyses were complete. Thus, we are posting preprint with our pre-registered analyses complete for only a subset of these pre-registered outcomes. As we complete more of these analyses, we will update this supplement. We also intend to post full results, data, and replication code online shortly.

## F.1  Hypotheses and Results

Below, we list the hypotheses precisely as they are written in our pre-registration document. For each hypothesis, we also describe our results.

### H1

**There are differences in prompt engineering ability (as measured through metrics such as average expected prompt quality, initial expected prompt quality, and max expected prompt quality) across demographic attributes and other observables, such as educational background and occupational skills.**

**Analysis approach:** For this hypothesis, we conducted multiple ANOVA tests, one for each of the demographic variables, against the relevant outcome variable and treatment variable (the model) as the covariate. As a robustness check, we repeated the same procedure using the Kruskal-Wallis U test. To adjust for multiple testing, we used the Benjamini-Hochberg adjustment with a false discovery rate of 0.05.

**Results:** The results of our analysis are as follows:

- **CLIP embedding cosine similarity**

    **No rescaling:** Among all the demographic variables, the following variables have a significant association with the similarity to the target image when testing with ANOVA: computer programming frequency, self-reported programming ability, outlook towards generative AI, age, gender, generative AI use, education, writing frequency, and self-reported occupational skills of critical thinking and social perceptiveness.

    When we fit a linear model to evaluate directionality, we find that those who use critical thinking and social perceptiveness as job skills are, on average, better at the task. Conversely, we find that older people, men, those with a more positive outlook regarding generative AI, and those who programmed more frequently performed worse on our experiment task.

    When conducting the non-parametric Kruskal-Wallis test, we found fewer significant variables. In this case, the statistically significant relationships with performance are: self-reported programming frequency, outlook towards generative AI, self-reported programming skill, gender, age, and the critical thinking occupational skill.

    **Z-score rescaling:** *Forthcoming*

    **Percentile rank rescaling:** *Forthcoming*

    **GPT-4V baseline rescaling:** *Forthcoming*

- **DreamSim**

    **No rescaling:** *Forthcoming*

    **Z-score rescaling:** *Forthcoming*

    **Percentile rank rescaling:** *Forthcoming*

    **GPT-4V baseline rescaling:** *Forthcoming*

**H2**

**There are observable differences in the prompting techniques of successful prompt engineers and unsuccessful prompt engineers. Such prompting techniques might include the use of longer prompts, the use of structured prompting techniques, and/or specific patterns in the way that the participant iterates on their prompts over time.**

**Analysis approach:** To investigate this hypothesis, we pre-registered evaluating how exploration/exploitation in the prompting space is related to performance. Hence, most of the pre-registered independent variables measured the level of similarity (i.e., exploitation) across a user's prompts. We refer to participants being more "exploitative" if they wrote prompts that are more similar to their previous prompts and more "exploratory" if they wrote prompts that deviated more from their previous prompts. As described below, we operationalized "similarity" in a variety of ways.

1. **Positively associated with exploitation**: average token sort ratio compared to the previous prompt, average cosine similarity between the embeddings of a given prompt with the previous prompt, fraction of times a prompt contains the previous prompt as an exact substring

2. **Negatively association with exploitation**: the variance of the prompt embedding, the number of topical transitions

We also pre-registered measuring each participant's average prompt length. These variables were measured per each user and across their first 10 attempts. Their association with performance, measured in terms of each of our eight outcome variables, was estimated in a linear model with DALL-E version fixed effects.

We also conducted a similar analysis at the iteration level to answer the following question: how is exploration/exploitation associated with subsequent performance in the next attempt, and is the strength of this relationship mediated by the quality of the previous attempt? To answer this question, first, we divided user-iteration observations into 6 equal-sized brackets by performance in the previous iteration. The bracketing allows us to explore heterogeneity in prompting behavior by the quality of previous attempts. We then estimated the effect of textual similarity to the previous prompt on the quality of the next attempt within each bracket. Our estimates adjusted for other covariates by matching user-iteration observations on the target image, DALL-E model, iteration and the exact quality of the previous attempt (Sävje et al., 2021).

**Results:** The results of our analysis are as follows:

- **CLIP embedding cosine similarity**

    **No rescaling:** We find strong and statistically significant evidence of an association between performance and our prompting variables, even after adjusting for Benjamini-Hochberg multiple testing. Token sort ratio, cosine similarity with the previous prompt, and frequency of including the previous prompt were positively associated with performance, while embedding variance and the number of topical transitions negatively correlated with performance. Taken together, these findings suggest that more successful users engaged in more exploitation and wrote prompts that were similar to one another. We also found that longer prompts were associated with higher performance, as shown in the main text.

We also find that when the previous performance was poor, higher exploration (or lower cosine similarity with the previous prompt) was associated with improved performance, although extremely high levels of exploration did not improve performance. In contrast, in cases where previous performance was high, higher exploitation monotonically increased performance. Using token sort ratio with the previous prompt as the measure of exploration also generated similar results. However, when using binary measures of exploration, i.e. topical transition or containing the previous prompt, we found that exploitation is associated with higher performance in the next iteration regardless of the bracket. In particular, we did not find the non-linear pattern we observed with continuous measures of exploration in the low-performing group, likely due to the binary nature of this measure that fails to distinguish between low, medium, and high levels of exploration/exploitation.

> **Z-score rescaling:** *Forthcoming*
>
> **Percentile rank rescaling:** *Forthcoming*
>
> **GPT-4V baseline rescaling:** *Forthcoming*

- **DreamSim**

  > **No rescaling:***Forthcoming*
  >
  > **Z-score rescaling:** *Forthcoming*
  >
  > **Percentile rank rescaling:** *Forthcoming*
  >
  > **GPT-4V baseline rescaling:** *Forthcoming*

**H3**

**There are differences in prompt engineering techniques (as measured through metrics such as prompt length and iteration-to-iteration token sort ratio) across demographic attributes and other observables, such as educational background and occupational skills.**

**Analysis approach:** To test this hypothesis, we estimated a linear model per each demographic trait as the independent variable, treatment arm fixed effects, and the prompting behaviors outlined below and described in-depth in Section B.B.5 as the outcome variables of interest. To account for multiple testing, we adjusted the p-values for all these models according to Benjamini-Hochberg procedure.

**Results:** The results of our analysis are as follows. As a reference, the definitions of these dependent variables are provided in Section B.B.5.

- **Prompt embedding variance:** We did not find any significant differences across various demographic traits.

- **Strategic Shifts:** We observed statistically significant differences based on age and reported programming frequency. Older participants and those with high programming frequency demonstrated increased topical transitions across their prompts.

- **Successive Prompt Token Sort Ratio:** We found significant differences by age, programming frequency, and writing frequency. Older users, those with high programming frequency and writing frequency, both precise instructions and imagery, write prompts that are less similar to each other on average.

- **Successive similarity:** We found significant differences by age, gender, generative AI outlook, programming skill/frequency and writing frequency. On average, older users, males, those who reported frequently computer programming, those who reported being strong computer programmers, those with a positive outlook on Generative AI, those who frequently use generative AI, and those who write a lot, those who self-describe as being good at writing instructions and writing with a lot of clear imagery write prompts that are less similar to each other (as measured by cosine similarity) on average.

- **Successive Prompt 'Contains Previous Prompt' Dummy:** We found significant differences by age, gender, outlook towards generative AI, and writing frequency. Older users and those who write frequently are less likely to write prompts that exactly contain the previous. In contrast, males and those with a positive outlook on generative AI are more likely to keep using their previous prompt.

### H4

**Insofar as the output returned by a generative AI model in response to a prompt is stochastic, the subsequent prompting strategies and prompting outcomes of participants that get lower-than-expected, higher-than-expected, or approximately expected outputs in response to their first prompt are different.**

**Analysis approach:** To test this hypothesis, we first calculated the z-score of each participant's realized image observed during our experiment relative to the sampling distribution that we approximated for the prompt that generated that image according to the procedure outlined in Section B.B.3 (see Section C for details on calculating the z-score). This z-score quantifies the random variation in observed image quality relative to the prompt's true underlying quality. Images with higher z-scores represent instances where the realized image quality exceeded expectations based on the prompt, while lower z-scores indicate instances where the image quality was randomly lower than expected for a given prompt. By examining the relationship between these z-scores and subsequent performance and prompting behavior, we can assess the causal impact of this stochasticity on user behavior.

To estimate the relationship between this stochasticity on subsequent prompting behavior and performance, we transformed the z-score into a trichotomous variable where any z-score less than

-0.45 as "lower-than-expected," between -0.45 and 0.45 as "expected" and greater than 0.45 as "higher-than-expected." We then perform two-sample t-tests with comparing the three groups' relative performance. We also estimate a linear model as a robustness check, regressing participant performance on the trichotomous z-score variable with treatment arm fixed effects. We also repeat this analysis, treating the z-score as a continuous variable without the trichotomous transformation. Finally, we also perform this analysis at the user level with the trichotomous variable, where only the z-score of the observed image of the first prompt is measured, and we test if that affects the average performance of all subsequent user attempts.

**Results:** The results of our analysis are as follows, with each outcome variable clearly labeled.

- **CLIP embedding cosine similarity**

  **No rescaling:** We find statistically significant evidence that increases in the z-score of the image realized in a previous prompt caused an increase in cosine similarity of the next prompt when conducting two-sample t-tests stratified by treatment arm. When we test differences in the trichotomous variable, the difference between the top bracket and the other two brackets is significant, but we do not find significant differences between the bottom and middle brackets. However, we do not find a statistically significant relationship between the observed z-score of the prompt and performance in the next iteration when we estimate a linear model. This may be due to the non-linearity of the relationship between these variables. Additionally, when we estimate the effect at the user-level, we do not find statistically significant effects of the z-score of the first prompt on average subsequent performance.

  **Z-score rescaling:** *Forthcoming*

  **Percentile rank rescaling:** *Forthcoming*

  **GPT-4V baseline rescaling:** *Forthcoming*

- **DreamSim**

  **No rescaling:** *Forthcoming*

  **Z-score rescaling:** *Forthcoming*

  **Percentile rank rescaling:** *Forthcoming*

  **GPT-4V baseline rescaling:** *Forthcoming*

- **Prompting Behaviors**

  **Prompt Length:** *Forthcoming*

  **Successive Similarity:** We find statistically significant evidence that increases in the z-score of the image realized in a previous prompt caused increases in successive similarity. When we compare the lowest z-score bracket (worse than expected images) to the highest z-score bracket (better than expected images), we observe a statistically significant difference

as well. However, we do not observe statistically significant differences between either the top or bottom bracket with the middle bracket. Additionally, when we estimate the effect at the user-level, we do not find statistically significant effects of the z-score of the first prompt on average successive similarity.

**Successive Prompt Token Sort Ratio:** We find statistically significant evidence that increases in the z-score of the image realized in a previous prompt caused increases in the Token Sort Ratio between successive prompts. When we compare the lowest z-score bracket to the highest z-score bracket, we observe a statistically significant difference. However, we do not observe statistically significant differences between either the top or bottom bracket with the middle bracket. Additionally, when we estimate the effect at the user-level, we do not find statistically significant effects of the z-score of the first prompt on average successive prompt token sort ratio.

**Successive Prompt 'Contains Previous Prompt' Dummy:** We find statistically significant evidence that increases in the z-score of the image realized in a previous prompt caused increases in the likelihood that the subsequent prompt contains the previous prompt. When we compare the lowest z-score bracket to the highest z-score bracket, we observe a statistically significant difference. However, we do not observe statistically significant differences between either the top or bottom bracket with the middle bracket. Additionally, when we estimate the effect at the user-level, we do not find statistically significant effects of the z-score of the first prompt on the average probability of a prompt containing the preceding prompt.

## H5

**Average prompt engineering ability (as measured through metrics such as average expected prompt quality, initial expected prompt quality, and max expected prompt quality) and prompting strategies will depend on the capacity of the model that participants are interacting with.**

**Analysis approach:** To test this hypothesis, we conducted three two-sample t-tests that compare the prompting performance and prompting behavior of the three possible pairs of treatment assignments. All outcome variables are specified in the results below. For robustness, we repeat these analyses with ANOVA to test whether any of the three treatment arms has an effect on the same variables. To account for multiple testing, we adjusted the p-values for tests with the Benjamini-Hochberg procedure.

**Results:** The results of our analysis are as follows:

- **CLIP embedding cosine similarity**

    **No rescaling:** *Forthcoming*

**Z-score rescaling:** We find statistically significant evidence for differences in performance across treatment arms by comparing participants' first attempt at recreating the target images, their average across all attempts, and their best attempts. Those using DALLE-E 3 (verbatim) performed better than both those using DALL-E 3 (Revised) and DALL-E 2, although there was no statistically significant difference between the latter two treatment arms. When we compare participants' average performance, we find comparatively identical resultsin terms of statistical significance. Finally, when we compare participants' best attempts pairwise across treatment arms, we find statistically significant differences between all three treatment arms: those using DALLE-E 3 (verbatim) were, on average, better than those using DALL-E 3 (Revised), who were, on average, better than those using DALL-E 2. Logically, those using DALL-E 3 (Verbatim) also performed better than those using DALL-E 2. When we test the relationship between the three treatment arms and these three outcome variables (first, average, and best performance) with ANOVA, all are statistically significant.

**Percentile rank rescaling:** *Forthcoming*

**GPT-4V baseline rescaling:** *Forthcoming*

- **DreamSim**

   **No rescaling:** *Forthcoming*

   **Z-score rescaling:** *Forthcoming*

   **Percentile rank rescaling:** *Forthcoming*

   **GPT-4V baseline rescaling:** *Forthcoming*

- **Prompting Behaviors**

   **Mean prompt Length:** We found statistically significant differences in prompt length between treatment arms. Participants using DALL-E 2 used significantly shorter prompts compared to both DALL-E 3 (Revised) and DALL-E 3 (Verbatim) groups. There was no significant difference in prompt length between the two DALL-E 3 groups. ANOVA results confirmed a significant effect of DALL-E version on prompt length.

   **Aggregate Similarity:** We found no statistically significant differences between any of the treatment arms. This suggests that the overall variability in prompts was similar across all three DALL-E versions. ANOVA results confirmed no significant effect of DALL-E version on this measure.

   **Successive Similarity:** Analysis of the average cosine similarity between successive prompts revealed no statistically significant differences between the treatment arms. ANOVA results confirmed no significant effect of DALL-E version on successive prompt similarity.

   **Successive Prompt Token Sort Ratio:** We found no statistically significant differences between the treatment arms. ANOVA results confirmed no significant effect of DALL-E version on this measure.

**Successive Prompt 'Contains Previous Prompt' Dummy:** Examining the probability of a current prompt being a superset of the previous prompt showed no statistically significant differences between any of the treatment arms. ANOVA results confirmed no significant effect of DALL-E version on this measure.

### H6

**Variability in participants' ability to prompt engineer effectively and prompting strategies will depend on the capacity of the model that participants are interacting with.**

**Analysis approach:** To test this hypothesis, we conducted two analyses. First, would conducted F-tests comparing the variance of participant performance and prompting behaviors between all 3 pairs of models. Second, we estimated the quantile treatment effects (QTEs) between all 3 pairs models on participant performance and prompting behaviors. We also pre-registered our intent to visually inspect whether the QTEs we observe are consistent with dispersion/"inequality" being reduced or increased when participants use different models (e.g., positive effects for low quantiles and negative/null effects for high quantiles would be consistent with inequality reduction).

**Results:**

- **CLIP embedding cosine similarity**

    **No rescaling:** *Forthcoming*

    **Z-score rescaling:** *Forthcoming*

    **Percentile rank rescaling:** *Forthcoming*

    **GPT-4V baseline rescaling:** *Forthcoming*

- **DreamSim**

    **No rescaling:** *Forthcoming*

    **Z-score rescaling:** We observed statistically significant reductions in the variance ratio of prompt-length, participant-level prompt embedding variance, and the percentage of prompts containing the previous prompt when comparing DALL-E 2 to DALL-E 3. We also found statistically significant reductions in the variance ratio of prompt length and participant-level prompt embedding variance when comparing DALL-E 2 and DALL-E 3 (Revised). Our QTE results comparing performance under DALL-E 2 and DALL-E 3 appear in the main text. The results we have seen when comparing DALL-E 3 (Revised) to DALL-E 2 are qualitatively similar, although the QTEs are attenuated due to the prompt revision. Almost all QTEs when comparing DALL-E 3 and DALL-E 3 (Revised) are positive.

    **Percentile rank rescaling:** *Forthcoming*

    **GPT-4V baseline rescaling:** *Forthcoming*

**H7**

**As participants repeatedly try to complete a task with a given model, the quality of their attempts will increase, and the extent to which the quality increases varies as a function of model capacity.**

**Analysis approach:** In order to test for differences in prompting performance across participant iterations, we pre-registered and conducted stratified two-sample tests with the 3 treatment arms as our strata. We run this test on each participant's performance improvement, defined as the performance gap between the initial iteration and the best-performing iteration.

- **CLIP embedding cosine similarity**

  **No rescaling:** We do not find evidence for statistically significant relationships between DALL-E 2 versus DALL-E 3 (Verbatim), nor between DALL-E 3 (Verbatim) and DALL-E 3 (Revised). However, we do find a statistically significant difference between DALL-E 2 and DALL-E 3 (Revised), whereby DALL-E 2 users showed greater improvement in cosine similarity compared to DALL-E 3 (Revised) users between their first and best attempts. The latter relationship holds even after adjusting for Benjamini-Hochberg multiple testing.

    **Z-score rescaling:** *Forthcoming*

    **Percentile rank rescaling:** *Forthcoming*

    **GPT-4V baseline rescaling:** *Forthcoming*

- **DreamSim**

    **No rescaling:** *Forthcoming*

    **Z-score rescaling:** *Forthcoming*

    **Percentile rank rescaling:** *Forthcoming*

    **GPT-4V baseline rescaling:** *Forthcoming*

**H8**

**The extent to which participants can recreate images using models such as DALL-E 2/3 will vary across images.**

**Analysis approach:** To evaluate this hypothesis, we pre-registered and performed two analyses. First, we compared the extent to which different images can be replicated using GPT-4V—a multimodal generative AI model that can take in both text and images and output text. For each of the 15 target images, we prompted GPT-4V to "Write a DALL-E {2, 3} prompt to recreate this image verbatim as closely and as detailed as possible." We do this once for DALL-E 2 and once for DALL-E 3, generating two "AI prompts" per target image. We then provided these AI

prompts to DALL-E 2, DALL-E 3 (Verbatim), and DALL-E 3 (Revised), respectively, 20 times each, generating 60 replicated images per target image (the DALL-E 3 AI prompt was sent to both DALL-E 3 (Verbatim) and (revised)). For each image, we measure the cosine similarity of the CLIP embedding vectors to those of the relevant target image. We then average these 60 similarity measures by target image, resulting in a mean similarity score that represents GPT-4V's ability to generate prompts that recreate each target image. Second, we simply measured the similarity of every participant-generated image to that of the relevant target image, and averaged over these participant-generated similarities.

**Results:**

- **CLIP embedding cosine similarity**: Below is the ranking of GPT-4V's ability to generate prompts that recreate each target image, as measured through CLIP embedding cosine similarity:

  1. Business Image #3
  2. Business Image #5
  3. Business Image #1
  4. Photography Image #1
  5. Design Image #2
  6. Business Image #2
  7. Photography Image #5
  8. Design Image #4
  9. Business Image #4
  10. Design Image #1
  11. Photography Image #2
  12. Design Image #5
  13. Design Image #3
  14. Photography Image #3
  15. Photography Image #4

  The maximum average cosine similarity of the CLIP embeddings in the above list (the easiest image for GPT-4V to replicate) was $CoSim = 0.944$ for Business Image #3, and the lowest score (the hardest image for GPT-4V to replicate was $CoSim = 0.734$ for Photography Image #4.

  Below is the ranking of participants' ability to generate prompts that recreate each target image, as measured through CLIP embedding cosine similarity:

  1. Business Image #3
  2. Business Image #5
  3. Business Image #1
  4. Photography Image #5
  5. Design Image #2
  6. Design Image #4
  7. Business Image #2
  8. Design Image #5
  9. Design Image #1
  10. Photography Image #1
  11. Design Image #3
  12. Photography Image #2
  13. Photography Image #3
  14. Photography Image #4
  15. Business Image #4

  The maximum average cosine similarity of the CLIP embeddings in the above list (the easiest image for participants to replicate) was $CoSim = 0.892$ for Business Image #3, and the lowest score (the hardest image for participants to replicate was $CoSim = 0.669$ for Business Image #4.

- **DreamSim**: Below is the ranking of GPT-4V's ability to generate prompts that recreate each target image, as measured through 1-DreamSim:

1. Design Image #3
2. Business Image #2
3. Business Image #4
4. Design Image #2
5. Business Image #1
6. Business Image #5
7. Design Image #4
8. Business Image #3
9. Photography Image #1
10. Photography Image #5
11. Photography Image #2
12. Design Image #1
13. Design Image #5
14. Photography Image #3
15. Photography Image #4

The maximum 1-DreamSim score in the above list (the easiest image for GPT-4V to replicate) was $\tilde{D} = 0.75$ for Design Image #3, and the lowest score (the hardest image for GPT-4V to replicate was $\tilde{D} = 0.40$ for Photography Image #4.

Below is the ranking of participants' ability to generate prompts that recreate each target image, as measured through 1-DreamSim:

1. Business Image #3
2. Business Image #2
3. Business Image #5
4. Business Image #1
5. Design Image #2
6. Design Image #3
7. Photography Image #2
8. Photography Image #5
9. Design Image #5
10. Design Image #4
11. Business Image #4
12. Photography Image #1
13. Photography Image #4
14. Design Image #1
15. Photography Image #3

The maximum 1-DreamSim score in the above list (the easiest image for participants to replicate) was $\tilde{D} = 0.575$ for Design Image #3, and the lowest score (the hardest image for participants to replicate was $\tilde{D} = 0.356$ for Photography Image #4.

## H-Exploratory

In our pre-registration, we also mention a number of exploratory analyses that are not described with the same level of detail. Multiple of these exploratory analyses appear in our main text. These pre-registered exploratory analyses are copied verbatim below:

"We plan to investigate whether differences in prompt engineering ability across demographic and other observed variables will vary depending on the complexity of the task, e.g., the difficulty of the image participants are being asked to replicate. We anticipate power for this analysis will be very low, so we chose to label it as an exploratory analysis rather than a pre-registered hypothesis.

We anticipate that we may conduct additional analysis of the prompts submitted by participants (and how these prompts evolve over the course of a session). Furthermore, we might explore the tips that participants provide after completing the task on how to prompt engineer effectively.

We also may take original and revised prompts submitted to DALL-E 3 treatment arms and submit them to DALL-E 2 (and vice versa) to see how participants would have counterfactually performed under different treatment assignments than the one to which they were assigned."

## F.2 Deviations From Pre-registration

Here, we also report any deviations from our pre-registered analysis. By and large, these deviations occurred because certain aspects of our pre-registration were not appropriate from a statistical analysis perspective, or were infeasible.

- In our pre-registration, we had anticipated using t-tests and Mann-Whitney U tests for many hypotheses. However, this turned out to be impossible for many variables since most of the demographic traits have multiple categories. Thus, we applied ANOVA and Kruskal-Wallis tests instead.

- In our pre-registration, we had forgotten to include Education and GenAI outlook in our list of demographics with respect to which we would measure task performance heterogeneity. We decided to include these important variables in our analysis, despite the accidental omission.

- In addition to our pre-registered prompt exclusion criteria, we removed additional prompts that did not appear to be "good-faith efforts" to complete our task based on the text of those prompts (see Section B.4 for more details). We find that our results are the robustness to the inclusion of prompts that were removed by this procedure.

- Because they were easier to implement, we ran t-tests rather than z-tests to conduct our tests of H5. Because the two tests are asymptotically equivalent, we do not believe this will make a difference in our analysis.

- When conducting our analyses to test H4, we observed that the distribution of the z-scores (our independent variable) did not conform to a normal distribution, displaying extremely large (7.1) or small (-21.9) values. To prevent our findings from being disproportionately influenced by these outliers, particularly in linear models, we excluded observations with absolute z-scores greater than 3, which constituted 2.5% of the user-attempt observations. For robustness, we also conducted the stratification analysis including these outliers and found similar results, with the exception that the Average Treatment Effect (ATE) for performance between the top and middle brackets was only marginally significant.