

The role of interface design on prompt-mediated creativity in Generative AI

Maddalena Torricelli
City, University of London
London, United Kingdom

Andrea Baronchelli
andrea.baronchelli.1@city.ac.uk
City, University of London
The Alan Turing Institute
London, United Kingdom

Mauro Martino
IBM Research
Cambridge, MA, United States

Luca Maria Aiello
luai@itu.dk
IT University of Copenhagen
Pioneer Centre for AI
Copenhagen, Denmark

ABSTRACT

Generative AI for the creation of images is becoming a staple in the toolkit of digital artists and visual designers. The interaction with these systems is mediated by *prompting*, a process in which users write a short text to describe the desired image's content and style. The study of prompts offers an unprecedented opportunity to gain insight into the process of human creativity. Yet, our understanding of how people use them remains limited. We analyze more than 145,000 prompts from the logs of two Generative AI platforms (Stable Diffusion and Pick-a-Pic) to shed light on how people *explore* new concepts over time, and how their exploration might be influenced by different design choices in human-computer interfaces to Generative AI. We find that users exhibit a tendency towards exploration of new topics over exploitation of concepts visited previously. However, a comparative analysis of the two platforms, which differ both in scope and functionalities, reveals some stark differences. Features diverting user focus from prompting and providing instead shortcuts for quickly generating image variants are associated with a considerable reduction in both exploration of novel concepts and detail in the submitted prompts. These results carry direct implications for the design of human interfaces to Generative AI and raise new questions regarding how the process of prompting should be aided in ways that best support creativity.

CCS CONCEPTS

• **Human-centered computing** → **Empirical studies in HCI**.

KEYWORDS

Prompting, explore-exploit, creativity, Stable Diffusion, Pick-a-Pic

ACM Reference Format:

Maddalena Torricelli, Mauro Martino, Andrea Baronchelli, and Luca Maria Aiello. 2024. The role of interface design on prompt-mediated creativity

in Generative AI. In *ACM Web Science Conference (WEBSCI '24)*, May 21–24, 2024, Stuttgart, Germany. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3614419.3644000>

1 INTRODUCTION

Recent years have witnessed the rapid advancement of Generative AI [9], a branch of artificial intelligence dedicated to developing computer models capable of creating original media such as text, images, speech, or music. Due to the impressive quality of their outputs, these models have drawn considerable public attention, which has been further sustained by the proliferation of increasingly accessible models [5] and openly available applications for generating a variety of synthetic content [3]. The rapid democratization of Generative AI [7] has arguably originated the largest shift in the participatory paradigm of the Web since the Web 2.0 revolution [1], with the crowd uploading massive amounts of AI-generated content online. This trend appears to be particularly pronounced in the realm of visual art, with an estimated 15 billions of AI-generated images uploaded to the Web as of August 2023 [19]. The collective process of mass creation of artistic content through artificial intelligence represents an unprecedented cultural phenomenon whose complex facets have attracted the interest of researchers from multiple disciplines [4, 11, 12].

One crucial aspect that characterizes the innovative creative process enabled by Generative AI tools lies in their user interface. The generation of visual artifacts is achieved through *prompting*, a form of human-computer interaction that uses short descriptive texts as instructions for the AI models to create a desired image [13]. The ways in which people compose their prompts offers two significant insights into the behavioral aspects of AI-assisted creation of art. First, prompts embody the human creative intent, thereby serving as a valuable resource for understanding how people navigate the realm of ideas. Second, prompts are shaped around the distinctive capabilities and constraints inherent to the models' different interfaces, which presents an opportunity to quantify the extent to which design decisions regarding how Generative AI interacts with humans influence the creative process. Therefore, understanding the prompting practices that emerge during interactions with different AI tools is essential to uncover the behavioral mechanisms that define the cultural process of artistic production in the AI era and that are not yet fully understood.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

WEBSCI '24, May 21–24, 2024, Stuttgart, Germany

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0334-8/24/05

<https://doi.org/10.1145/3614419.3644000>

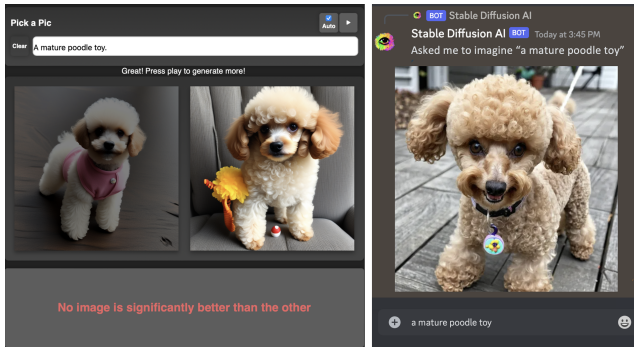


Figure 1: User interface of Pick-a-Pic (left) and Stable Diffusion (right).

In this context, we analyze textual prompts used for image generation to address two key questions:

Q1) To what extent people *explore* different concepts in the process of creating generative art?

Q2) How do different designs of Generative AI tools either foster or hinder this exploration?

By answering these questions, we aim to shed light on the process of human creativity itself, and offer insights on how technology can be designed to support it.

To answer these questions, we analyze more than 145,000 prompts from two text-to-image generator platforms which differ significantly in their structure, design, and objectives: DiffusionDB [21] and Pick-a-Pic [10]. We examine the sequences of prompts submitted by each user in terms of *exploration* and *exploitation*, two aspects of a decision making pattern that refers to how individuals balance between exploring new available options and exploiting familiar choices to optimize a given outcome [2, 6, 18]. In the context of this study, users can either engage in exploration by experimenting with unfamiliar prompts, or they can lean towards exploitation by composing prompts similar to those they have previously submitted, indicating a preference for specific image styles or content. We quantify these patterns by measuring the evolution of textual similarity between prompts and visual similarity between images within user sessions. In our analysis, we adopt the working hypothesis that exploration is loosely yet positively associated with creativity: given a fixed number of interactions with the platform, users who submit more diverse prompts create on average a wider range of artistic content.

Our findings reveal that users have a propensity towards exploration over exploitation. However, the presence of features to generate new image variants with a simple click, without revising previously submitted prompts, is associated both with a reduced degree of exploration and shorter prompts. Our findings raise new questions about how to best aid user interactions with Generative AI tools without constraining the creative process that arises from those interactions.

2 DATASET

Figure 1 shows the interfaces of both platforms we collected data from. To interact with the Stable Diffusion model, users submit

prompts via dedicated channels on the official Stable Diffusion Discord server and receive automated responses containing the corresponding generated image. DiffusionDB [21] is a large-scale repository of images created by the Stable Diffusion model, paired with the prompts that users submitted to generate them. This data collection was obtained by scraping the Stable Diffusion Discord channel at the beginning of its activity, from 1st to 31st August 2022. The collection contains 10,177 unique users and 1.5 million prompts, each prompt paired with the image it generated.

Pick-a-Pic is a web platform that allows users to submit prompts for image generation and nudges them to provide annotations of image quality. Upon receiving a prompt, the platform generates a pair of image variants and invites the user to select their preferred option. The non-preferred variant is replaced with a new one derived from the same prompt, and the new pair can be rated again. At any time, users can decide to cast a vote with a simple click or to generate a different pair of images by writing a new prompt. In the dataset, a click on the interface counts as a submission of the same prompt it was used in the previous interaction. The platform creators logged all user interactions in January 2023 (the month in which the platform launched) and compiled them into a dataset comprising 992 users and 77,929 prompts.

It is important to stress that the data from both platforms was collected right after the platforms were launched, and a time in which very few open tools for AI image generation were accessible to the general public. It is therefore reasonable to assume that all users were roughly equally familiar with this type of platforms during the data collection period.

All images are free from illegal, hateful, or NSFW content, and are published under CC0 license that waives all copyright and allows uses for any purpose. The information from the official Stable Diffusion Discord server is public. The Pick-a-Pic dataset was obtained previous explicit consent from the platform users.

For each dataset, we focus on the essential data to study temporal patterns of human-AI interaction: the *userID*, the textual *prompt*, the *timestamp* of prompt submission, and the set of *images* returned. We call *prompt sequence* a user’s temporally-sorted list of prompts.

The platforms exhibit different levels of popularity. To mitigate potential activity biases when comparing user sessions across both platforms, we harmonize their characteristics by selecting a subset of users from DiffusionDB that align more closely with the activity of Pick-a-Pic users. To achieve that, we match each of the 992 Pick-a-Pic users with a randomly chosen DiffusionDB user who had submitted a number of prompts within 50% of those submitted by the corresponding Pick-a-Pic user. To ensure the robustness of our results to this step of random selection, we repeat this sampling process ten times with replacement, and report all our results as the average over these ten randomized selections. On average, the DiffusionDB samples contain 70,081 prompts.

3 METHODS

Image and prompt similarity. For each user, we calculate the similarity between all pairs of images they generated. To achieve that, we map each image into a numerical embedding space using DinoV2 [15], a family of foundation visual models designed to extract




Figure 2: To identify topical transitions in prompting, we calculate a similarity matrix between pairs of prompts sorted by their submission time (a), binarize the matrix (b), and identify blocks of highly-similar, consecutive prompts around the matrix diagonal (c).

image features that are suitable for a number of image processing tasks, including calculating visual similarity. DinoV2 models achieve state-of-the-art performance across most common benchmarks, outperforming other openly available weakly-supervised models, such as OpenCLIP [8]. Specifically, we used the lvd142m model, that produces embeddings with 769 dimensions. Given two DinoV2 vectors $\mathbf{v}_{\text{image}_i}$ and $\mathbf{v}_{\text{image}_j}$, we calculate the cosine similarity between them:

$$\text{sim}(\text{image}_i, \text{image}_j) = \frac{\mathbf{v}_{\text{image}_i} \cdot \mathbf{v}_{\text{image}_j}}{\|\mathbf{v}_{\text{image}_i}\| \|\mathbf{v}_{\text{image}_j}\|}. \quad (1)$$

We also calculate the similarity between all pairs of prompts submitted by any given uses. We do that using the Jaccard index, namely the intersection over the union of the set of tokens composing prompts i and j :

$$\text{sim}(\text{prompt}_i, \text{prompt}_j) = \frac{|\text{tokens}_i \cap \text{tokens}_j|}{|\text{tokens}_i \cup \text{tokens}_j|}. \quad (2)$$

Topical variation. The degree of exploration or exploitation that users exhibit during their interaction with the AI can be quantified by examining the pairwise similarity between their prompts or images. In particular, the similarity between consecutive prompts in a user’s sequence being low indicates a *topical variation* in that user’s focus. This is illustrated in Figure 2(a), which presents the Jaccard similarity matrix of the prompts of an example user. Here, rows and columns represent the user’s prompt sequence, and the cells encode the Jaccard similarity between all prompt pairs in the sequence. High similarity values along the diagonal suggest an exploitation process, where the user submits a series of consecutive prompts that are semantically similar. Conversely, low similarity scores around the diagonal indicate exploration, where consecutive prompts are markedly different from each other.

To measure the frequency of topical variations in a user’s prompts, we first binarize the similarity matrix. This is achieved by assigning a value of 1 to elements above the matrix average and 0 to the rest, as shown in Figure 2(b). We then count the number of 1-valued cell *blocks* adjacent to the diagonal (Figure 2(c)). A boundary between blocks is identified when the binarized similarity between two consecutive prompts is zero ($\text{sim}_{\text{bin}}(\text{prompt}_i, \text{prompt}_{i+1}) = 0$). Finally, we quantify the number of topical variations as the number of blocks minus one ($n_{\text{variation}} = n_{\text{blocks}} - 1$). This number ranges from one (when all prompts are similar) to the number of prompts minus one (when each prompt is dissimilar from all others). From

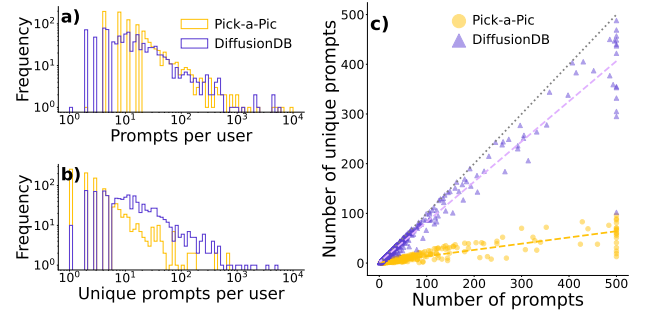


Figure 3: Distribution of number of prompts (a) and unique prompts (b) per user. Total number of unique prompts submitted after n interactions (c).




Figure 4: Probability distributions of image similarity for: (a) all pairs of different prompts; (b) all pairs of identical prompts within the same prompt sequence.

this count, we calculate the probability of a topical variation for each user:

$$P(\text{variation}) = \frac{n_{\text{variation}}}{n_{\text{prompt}} - 1} = \frac{d_{\text{blocks}} - 1}{n_{\text{prompt}} - 1}. \quad (3)$$

4 RESULTS

4.1 Image and prompt similarity

The distribution of the number of prompts per user is broad, with the majority of users submitting a limited number of prompts (*median* = 12 in Pick-a-Pic, *median* = 11 in DiffusionDB), while few power users submit thousands during the observed period (Figure 3a). The data preprocessing we conducted to align the two datasets results in a comparable distribution of prompts per user across both platforms (two-sample Kolmogorov Smirnov test $p = 0.0$). However, the platforms’ distinct interaction designs lead to significant differences in the types and range of prompts users submit and, consequently, the variety of images produced. Specifically, every time a user rates an image in Pick-a-Pic, the system generates a new image variant for the same prompt, while Stable Diffusion tends to produce identical or very similar variants for identical prompts (Figure 4).

The average number of *unique* prompts submitted by users is significantly lower in Pick-a-Pic than in Stable Diffusion (*avg.* = 8.7 vs. *avg.* = 58.0), Figure 3b). Given the option to generate new images with a simple click without revising their prompt, Pick-a-Pic users typically explore fewer prompts than Stable Diffusion users after the same number of rounds, and this gap widens with the number of




Figure 5: Distribution of similarity of consecutive prompts (a) and consecutive images (b) in a user sequence.

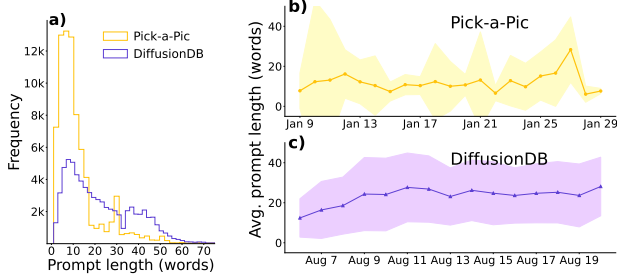


Figure 6: Distribution of number of words per prompt (a) and evolution of average prompt length over time (b,c).

interactions with the AI (Figure 3c). These distinct practices directly influence the similarity between consecutive prompts and images. In Pick-a-Pic, the distribution of similarity between consecutive prompts is relatively even, with the exception of notable peaks at the extreme values of zero and one, indicating users either sticking to the same prompt when generating the next image or switching to a completely non-overlapping prompt (Figure 5a). Conversely, the same similarity distribution in DiffusionDB shows a slight U-shape and a lower frequency of pairs with similarity equal to 1, suggesting that repetitions of the same prompts are less common, and there is a tendency towards more gradual prompt modifications. These two different approaches to prompting directly impact the variety of images generated. The distributions of visual similarity for consecutive images in a users' sequence are bimodal for both platforms, with two peaks around 0 and 0.7 indicating processes of exploring new visual content and exploiting variations of content generated with the previous prompt in the sequence (Figure 5b). However, DiffusionDB exhibits a more pronounced exploration peak, while Pick-a-Pic shows a more prominent exploitation peak.

These results indicate that an AI system capable of generating a diverse set of images from the same prompt may discourage users from trying new prompts. More surprisingly, Pick-a-pic prompts are not only repeated more frequently, but they are also significantly shorter (Figure 6a). This suggests that interfaces facilitating easy ways to generate images without the need of prompting may divert user attention from the process of prompt composition, potentially resulting in missed opportunities for users to learn how to craft their prompts to enhance the quality of the desired output. This hypothesis is supported by the trend of the average length of prompts over time (Figure 6b): the average prompt length remains roughly constant in Pick-a-pic, while it increases steadily in Stable

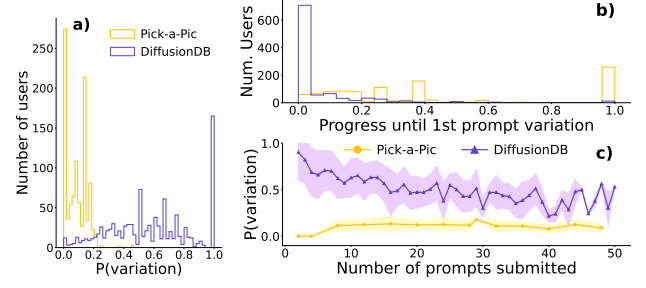


Figure 7: Topical variation in prompting. (a) Distribution of the user-specific probability of changing topic when submitting a new prompt. (b) Number of users who experienced their first topical variation at a given point of their prompt sequence, normalized by the length of the sequence. (c) Average probability of topical variation for users who submitted a fixed number of prompts; standard deviation is shown.

Diffusion, a possible indication that Stable Diffusion users learn to refine and improve their prompts over time.

4.2 Exploration vs. exploitation in prompting

The diversity of user-submitted prompts varies significantly across platforms, leading to two distinct experiences in terms of the breadth of visual content that users encounter during their interaction with the AI. To quantify these differences within the context of an exploration-exploitation process, we aggregate similar consecutive prompts for each user and calculate the probability $P(\text{variation})$ of exploring an entirely new topic with their next prompt, as defined in Formula 3. The probability distributions reveal a stark contrast between the two platforms (Figure 7a). Users of Stable Diffusion exhibit a wide range of behaviors, but the majority tend towards exploration, with 63% changing topics more often than not ($P(\text{variation}) \geq 0.5$) and 16% changing topic at every prompt ($P(\text{variation}) = 1$). Conversely, almost all Pick-a-Pic users change topic less than once every five prompts, and 26% of them never do. Moreover, most Stable Diffusion users explore new topics immediately after their first prompt, whereas Pick-a-Pic users typically attempt topical variations much later in their sequence of interactions with the AI (Figure 7b). The proportion of exploratory behavior over time changes significantly in Stable Diffusion, decreasing from approximately 0.9 at the second prompt to around 0.5 after 15 prompts (Figure 7c). This trend aligns with a scenario where users gain confidence in trying prompt variations within the same topic as they gain experience with the AI interaction. In contrast, the probability of a topic transition remains stable over time in Pick-a-Pic, suggesting that its human-computer interaction paradigm does not encourage users to alter their approach to prompt variation, at least within the timeframe of our dataset.

The Pick-a-Pic feature that allows for generating new images without the need for prompt revision appears to slow down exploration compared to the Stable Diffusion interface that solely relies on prompting to interact with the AI interaction. This raises the question of how many prompts it takes for the average Pick-a-Pic users to explore a topical space as diverse as that typically navigated

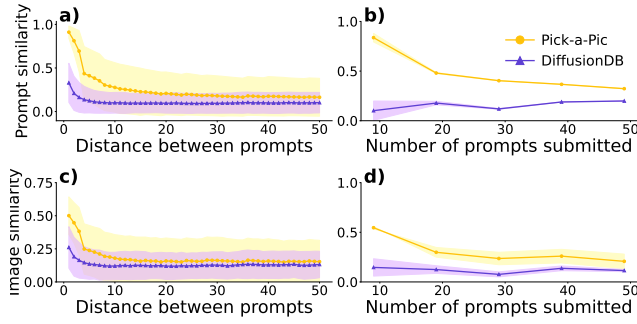


Figure 8: Similarity of prompts and images over time. Average and standard deviation of: (a) the similarity between pairs of prompts submitted by the same user at distance d in their sequence of prompts; (b) similarity between all pairs of prompts of users who submitted a fixed number of prompts. Panels (c) and (d) show the same trends for image similarity.

by Stable Diffusion users. To address this, we measure the similarity of prompt pairs at a distance d within a user’s prompt sequence to gauge the rate of topical drift (Figure 8a). Three trends become apparent. First, as anticipated by our prior analysis, the starting offsets of the similarity curves differ noticeably between the two platforms, with the similarity at $d = 1$ in Pick-a-Pic being approximately three times greater than that in Stable Diffusion. Second, similarity in both platforms decreases rapidly with d . In Pick-a-Pic, the steepest rate of decay occurs at $d = 4$, aligning with the observation that users typically do not transition to new topics until after five image generations. In Stable Diffusion, the average similarity stabilizes around 0.1 after just a few iterations. Last, the similarity curve of Pick-a-Pic converges with that of Stable Diffusion at slow pace. At $d = 4$, the standard deviation intervals of the two curves begin to overlap, and only at $d = 8$ the average Pick-a-Pic similarity becomes comparable with the average in Stable Diffusion at $d = 1$. This suggests that a Pick-a-Pic user requires approximately 4 to 8 iterations to experience a topical shift equivalent to that typically experienced by Stable Diffusion users in a single iteration. In aggregate, the average diversity of images explored by Pick-a-Pic users approaches, but never quite equals, that of Stable Diffusion users (Figure 8b). Similar patterns emerge when considering image similarity instead of prompt similarity (Figure 8c,d). In conclusion, while users across both platforms demonstrate a similar propensity for exploration, mechanisms that divert attention from prompt editing significantly impede this process.

5 CONCLUSIONS

Generative AI solutions are quickly becoming a staple in the digital toolbox for computer-aided art production. Ensuring that these tools function as expected requires aligning their design with their intended use. We examined the prompt logs of two synthetic image generation platforms through the lens of the exploration-exploitation trade-off, and obtained two findings that have direct implications for the design of Generative AI interfaces.

A1) When their interaction with the AI is focused exclusively on prompting, users tend to explore a diverse range of prompts rather

than reusing previously seen concepts, often shifting topics entirely from one prompt to the next.

A2) When Generative AI tools provide functionalities for generating image variations without composing new prompts, the exploration process significantly slows down, prompts become less complex, and the breadth of content to which users are exposed is reduced.

These findings suggest that while Generative AI systems hold a great potential for inspiring diversity in creative thinking, this potential may be curtailed by the introduction of supportive functionalities that limit prompt modification.

Our study is constrained by its focus on two platforms and the inability to control for external factors such as differing user attitudes towards prompting, art, or technology. These factors and other confounders that are not accounted for in this study might contribute to the observed differences between the two platforms. To fully account for these factors and confirm the generality of our findings, it is necessary to extend our study to a broader range of platforms. More in general, our study opens the question of how to strike a good balance between building accessible interfaces that alleviate the burden of composing complex prompts and creating the best conditions for users to express and enhance their creativity while interacting with Generative AI tools.

6 RELATED WORK

Our work is inspired by previous research in this domain and builds upon it by providing a quantitative study of a pattern of human-computer interaction that has not been quantified so far. In the domain of image generation, prompting has been primarily explored from the user experience perspective through small-scale qualitative studies. Prior research has employed surveys and focus groups to delve into the motivations behind the use of AI art, indicating that a significant proportion of users are motivated purely by recreational use [16]. A few interview-based studies have provided insights into the strategies people employ to approach prompting, examining aspects such as how people evaluate and enhance the quality of their prompts [14] the cognitive processes involved in composing prompts [17], and the strategies for exploring prompt variations [23]. Only a few large-scale quantitative studies have been conducted to date. For example, Xie et al. compared a vast corpus of image prompts with traditional query logs from Web search engines, discovering that prompts are typically longer and semantically broader than traditional queries [22]. Vodrahalli et al. measured the diversity of chains of prompt reformulations from a dataset of 50k human-AI interactions, observing that diversity does not decrease as users refine their prompts [20].

REFERENCES

- [1] Paul Anderson et al. 2007. *What is Web 2.0?: ideas, technologies and implications for education*. Vol. 1. JISC Bristol.
- [2] Andrea Baronchelli and Filippo Radicchi. 2013. Lévy flights in human behavior and cognition. *Chaos, Solitons & Fractals* 56 (2013), 101–105.
- [3] Erik Brynjolfsson, Danielle Li, and Lindsey R Raymond. 2023. *Generative AI at work*. Technical Report. National Bureau of Economic Research.
- [4] David De Cremer, Nicola Morini Bianzino, and Ben Falk. 2023. How generative AI could disrupt creative work. *Harvard Business Review*. (2023). <https://t.ly/IEtLO>
- [5] Roberto Gozalo-Brizuela and Eduardo C Garrido-Merchan. 2023. ChatGPT is not all you need. A State of the Art Review of large Generative AI models. *arXiv:2301.04655* (2023).

- [6] Quanqin He, Hao Liu, Guiyan Ding, and Liangping Tu. 2023. A modified Lévy flight distribution for solving high-dimensional numerical optimization problems. *Mathematics and Computers in Simulation* 204 (2023), 376–400.
- [7] Krystal Hu. 2023. ChatGPT sets record for fastest-growing user base. <https://www.reuters.com/technology/chatgpt-sets-record-fastest-growing-user-base-analyst-note-2023-02-01/>
- [8] Gabriel Ilharco, Mitchell Wortsman, Ross Wightman, Cade Gordon, Nicholas Carlini, Rohan Taori, Achal Dave, Vaishaal Shankar, Hongseok Namkoong, John Miller, et al. 2021. OpenCLIP. *Zenodo* 4 (2021), 5.
- [9] Mladan Jovanovic and Mark Campbell. 2022. Generative artificial intelligence: Trends and prospects. *Computer* 55, 10 (2022), 107–112.
- [10] Yuval Kirstain, Adam Polyak, Uriel Singer, Shahbuland Matiana, Joe Penna, and Omer Levy. 2023. Pick-a-pic: An open dataset of user preferences for text-to-image generation. *arXiv:2305.01569* (2023).
- [11] Alexis Newton and Kaustubh Dhole. 2023. Is AI Art Another Industrial Revolution in the Making? *arXiv:2301.05133* (2023).
- [12] Jonas Oppenlaender. 2022. The creativity of text-to-image generation. In *Proceedings of the 25th International Academic Mindtrek Conference*. 192–202.
- [13] Jonas Oppenlaender. 2022. Prompt engineering for text-based generative art. *arXiv:2204.13988* (2022).
- [14] Jonas Oppenlaender, Rhema Linder, and Johanna Silvennoinen. 2023. Prompting ai art: An investigation into the creative skill of prompt engineering. *arXiv:2303.13534* (2023).
- [15] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. 2023. Dinov2: Learning robust visual features without supervision. *arXiv:2304.07193* (2023).
- [16] Téó Sanchez. 2023. Examining the Text-to-Image Community of Practice: Why and How do People Prompt Generative AIs?. In *Proc. of the 15th Conference on Creativity and Cognition*. 43–61.
- [17] Johannes Schleith, Milda Norkute, Mary Mikhail, and Daniella Tsar. 2022. Cognitive Strategy Prompts: Creativity Triggers for Human Centered AI Opportunity Detection. In *Proc. of the 14th Conference on Creativity and Cognition*. 29–37.
- [18] Wataru Toyokawa, Hye-rin Kim, and Tatsuya Kameda. 2014. Human collective intelligence under dual exploration-exploitation dilemmas. *PloS one* 9, 4 (2014).
- [19] Alina Valyaeva. 2023. AI Has Already Created As Many Images As Photographers Have Taken in 150 Years. <https://journal.everyapixel.com/ai-image-statistics>
- [20] Kailas Vodrahalli and James Zou. 2023. ArtWhisperer: A Dataset for Characterizing Human-AI Interactions in Artistic Creations. *arXiv:2306.08141* (2023).
- [21] Zijie J Wang, Evan Montoya, David Munechika, Haoyang Yang, Benjamin Hoover, and Duen Horng Chau. 2022. Diffusiondb: A large-scale prompt gallery dataset for text-to-image generative models. *arXiv:2210.14896* (2022).
- [22] Yutong Xie, Zhaoying Pan, Jinge Ma, Luo Jie, and Qiaozhu Mei. 2023. A prompt log analysis of text-to-image generation systems. In *Proceedings of the ACM Web Conference 2023*. 3892–3902.
- [23] JD Zamfirescu-Pereira, Richmond Y Wong, Bjoern Hartmann, and Qian Yang. 2023. Why Johnny can't prompt: how non-AI experts try (and fail) to design LLM prompts. In *Proc. of CHI Conference on Human Factors in Computing Systems*. 1–21.