

The Prompt Artists

Minsuk Chang (<https://orcid.org/0000-0002-9441-3337>), Research, Google, United States, minsukchang@google.com (<mailto:minsukchang@google.com>)

Stefania Druga (<https://orcid.org/0000-0002-5475-8437>), Google Labs, Google, United States and Information School, University of Washington, United States, druga@google.com (<mailto:druga@google.com>)

Alexander J. Fiannaca (<https://orcid.org/0000-0001-8981-1114>), Research, Google, United States, afiannaca@google.com (<mailto:afiannaca@google.com>)

Pedro Vergani (<https://orcid.org/0009-0000-0305-319X>), Google, United Kingdom, pvergani@google.com (<mailto:pvergani@google.com>)

Chinmay Kulkarni (<https://orcid.org/0000-0003-4922-6601>), Research, Google, United States and Computer Science Department, Emory University, United States, ckulkarni@google.com (<mailto:ckulkarni@google.com>)

Carrie J Cai (<https://orcid.org/0000-0001-9421-7128>), Research, Google, United States, cjcai@google.com (<mailto:cjcai@google.com>)

Michael Terry (<https://orcid.org/0000-0003-1941-939X>), Research, Google, United States, michaelterry@google.com (<mailto:michaelterry@google.com>)

DOI: <https://doi.org/10.1145/3591196.3593515> (<https://doi.org/10.1145/3591196.3593515>)

C&C '23: [Creativity and Cognition](https://doi.org/10.1145/3591196) (<https://doi.org/10.1145/3591196>), Virtual Event, USA, June 2023

This paper examines the art practices, artwork, and motivations of prolific users of the latest generation of text-to-image models. Through interviews, observations, and a user survey, we present a sampling of the artistic styles and describe the developed community of practice around generative AI. We find that: 1) artists hold the text prompt *and* the resulting image can be considered collectively as a form of artistic expression (*prompts as art*), and 2) *prompt templates* (prompts with “slots” for others to fill in with their own words) are developed to create *generative art styles*. We discover that the value placed by this community on unique outputs leads to artists seeking specialized vocabulary to produce distinctive art pieces (e.g., by reading architectural blogs to find phrases to describe images). We also find that some artists use “glitches” in the model that can be turned into artistic styles of their own right. From these findings, we outline specific implications for design regarding future prompting and image editing options.

CCS Concepts: • Human-centered computing → Empirical studies in HCI;

Keywords: AI art, Artists using AI, Text-to-Image models

ACM Reference Format:

Minsuk Chang, Stefania Druga, Alexander J. Fiannaca, Pedro Vergani, Chinmay Kulkarni, Carrie J Cai, and Michael Terry. 2023. The Prompt Artists. In *Creativity and Cognition (C&C '23), June 19–21, 2023, Virtual Event, USA*. ACM, New York, NY, USA 13 Pages. <https://doi.org/10.1145/3591196.3593515> (<https://doi.org/10.1145/3591196.3593515>)



Figure 1: Prompt artists develop descriptive text-based prompts that are rendered by text-to-image models. Highly skilled prompt artists will develop 1) distinct visual concepts and styles (S1), 2) prompts that can also serve as titles of the art piece (“prompts as art”, P1), and 3) “prompt templates” (P2), which encapsulate specific visual concepts to be customized by others. Artists strive to discover unique *natural language* that produces unique *visual outputs* (G1), and/or model “glitches” (G2) that can be elevated to artistic styles in their own right. Finally, some prompt artists validate the novelty of their work by conducting an image search for similar images (C1).

1 INTRODUCTION

Advances in text-to-image(TTI) models have led to significant improvements in the quality of computer-generated, synthetic images [20, 31]. A new generation of text-to-image models enable the creation of high-fidelity images via descriptive text prompts by leveraging advances in large language models [3, 49, 56, 57, 62]. With broadening access to these models, communities of practice have emerged, enabling people to share designs, prompts, and example images. For instance, there are now tools to help people write prompts ¹, and even marketplaces for successful prompts ².

Prior work has examined the phenomenon of computer-generated art in a variety of contexts [6, 13, 41]. For example, as an art historian analyzing the AI-assisted art movement, Mazzone et al. [41] describe how the artist's role has adapted to include pre-curation, tweaking, and post-curation. More recently, Hertzmann argues that text-to-image (TTI) models like DALL-E do not themselves create art, but that the artists and technologists who apply them as tools are the ones creating art [30]. With the emergence of this new class of models, which are capable of producing extremely high quality images from textual descriptions (e.g., [2, 4, 49]), we are motivated to understand how this new technology is being adopted and used by creators.

In this research, we provide a snapshot of a vibrant community of art practice that has arisen around text-to-image models, sharing insights into the ingenuity and creativity of the users of these models.³ Within a US-based technology company that has produced its own TTI models, we sent a survey to the TTI models users to gain a basic understanding of how and why they are used. We also interviewed and observed 11 prominent users of these models who are using the models as an art medium, recruiting from survey respondents and by directly asking prominent users. They have each generated thousands of images with both the internal and other publicly available models, and are actively sharing their creations in multiple communities. In studying the artistically-driven members of this local community, we sought to understand their practices, their artifacts, and their motivations for engaging extensively with the models. For the purposes of this research, we scope our inquiry to studying interfaces that only accept text as input (recognizing that a wide variety of model capabilities and interfaces are available, including those that enable more fine-grained editing of images). We restrict our scope to text-only interfaces because these were the first interfaces available for models such as DALL-E, and these text-based interfaces have seen considerable use by the public and our internal users.

Our study reveals that users of these models have developed a range of artistic styles, including origami figurines, fashion (e.g., dresses) made out of materials like bricks, and reality “mash-ups” that create hybrids of animals or of fruits and vegetables (see Figures 2b, 3a, 3b, 4b). However, we also found that the artistic outputs of this community of users are not limited to the images themselves. For example, the *prompt itself* is an important output, and a piece of the art: a parsimonious, descriptive prompt accompanying the image is seen as a virtuous goal beyond just the image, as it simultaneously acts as a “title” for, and description of, the art piece (see Figure 1, P1). Similarly, a *prompt template*—a text prompt with one or more empty “slots” for others to fill in—is considered an art piece all on its own (see Figure 1, P2). Among other characteristics, a well-designed prompt template has the property of encapsulating an artistic vision that can nonetheless be customized by future users of that prompt template.

Our results also reveal the lengths some users go to when searching for unique, distinctive outputs. In particular, some creators turn to thesauri or online, domain-specific blogs (e.g., architectural blogs) in search of vocabulary that elevates the model output beyond the ordinary. This focus on vocabulary suggests that capable TTI model artists may also benefit from being highly skilled with natural language. Another creator explicitly seeks unique model outputs, but through identification of “glitches” that can be elevated to styles all on their own. For example, this latter artist found the model did not render reflections in mirrors perfectly, and explored this concept through a number of pieces (see Figure 1, G2).

Finally, we find that the artists interviewed place a premium on originality, with some turning to image search to validate that their outputs are, in fact, unique.

In sum, this paper presents results from a survey and interview study of heavy users of TTI models, making the following contributions:

- **Sample styles:** We provide a sampling of artistic styles developed by study participants to contextualize the types of outputs being produced by new text-to-image models.
- **Prompt as art:** We find that the prompt itself is often considered a part of the artistic output (in addition to the actual image), with artists pursuing a goal of creating parsimonious, descriptive input prompts.
- **Prompt templates as art:** We discover that artists also produce *prompt templates* to encapsulate a unique visual concept that others can customize.
- **Natural language mastery for visual language artistry:** We describe how TTI artists seek unique natural language in an attempt to elevate their pieces beyond the norm.
- **Glitches as art:** We show how some artists look for “glitches” that can be reliably transformed into new styles.
- **Validating originality:** We describe artists’ concerns in validating their outputs as original, and how they currently validate through image search.

Together, these findings suggest new directions for interactive interfaces and aids for prompt-centric uses of TTI models: 1) Methods and tools to help users locate novel language and capabilities of the model, 2) aiding users in validating the originality of their outputs, and 3) reifying the notion of a prompt template into a standalone

computational artifact that supports richer interaction by users of the template. Importantly, while our results derive from a study of internal TTI models, the implications for design are generally applicable to use of any TTI model (e.g., the notion of a prompt template is useful for any TTI model, as it captures a particular artistic vision in a portable, yet customizable. form).

In the rest of the paper, we review related work, describe our study method, present results from the survey and interview study, and conclude with a discussion that draws implications for design from the study data.

2 RELATED WORK

Advances in deep learning have led to the development of generative machine learning models capable of producing both images that are highly realistic and images that are highly creative in both existing and novel artistic styles [12]. For example, Generative Adversarial Networks (GANs) [24] learn to generate images and simultaneously distinguish real and fake images to generate highly realistic images. Many variations of the GAN architecture have been investigated (e.g., [9, 38, 64]). A particularly relevant variation of this architecture is the Creative Adversarial Network (CAN) from Elgammal et al. [21], designed to generate images with novel artistic styles. This artwork was subsequently featured in multiple exhibitions [1] where human observers could not distinguish the CAN-generated art from human-authored artwork. Outside of GANs, Gatys et al. [23] introduced a method to apply learned artistic styles to random images (a technique now known as Neural Style Transfer [36]). Additionally, work originally designed to make convolutional neural networks more explainable, now referred to as Deep Dreams, became popular for generating art [43] due to its ability to generate psychedelic versions of images [45]. While these techniques allowed for generating images with creative and novel artistic styles [21], none provided significant affordances to end-users for controlling what was generated outside of the training data scope.

Mansimov et al. [40] addressed this issue by showing that a generative model could produce novel images from natural text when conditioned on image captions. As text-to-image models rely on language modeling techniques, recent advances in the scaling of large language models [19, 53] have enabled the development of correspondingly large text-to-image models with impressive results. The most recent of these models include: DALL·E [50, 55] and DALL·E 2 [49, 54], Stable Diffusion [5, 56], Midjourney[3], Parti [4, 62], and Imagen [2, 57].

Fueled by the latest advances in text-to-image models, current image generation applications are becoming mainstream. With this broader adoption comes the question of how these new models’ capabilities impact art practices, which we examine in this paper.

2.1 AI in Creativity Support Tools & Human-AI Co-Creation

AI tools have played a prominent role in creativity support tool (CST) research [22, 34]. AI-based CST systems have been produced to support artistic generation in domains such as fashion and product design [35, 52, 58], music creation [33, 39, 42], drawing [17, 18, 37, 48], visual design and story-boarding [59, 63], and storytelling [32, 51]. Most of these tools support the artistic implementation process as either production aids (i.e., tools that perform most of the work of generating the art, e.g., generative text-to-image models) or as execution aids (e.g., AI-powered brush tools in drawing applications) [16].

Hwang et al. [34] further characterize how these tools apply AI models in the creative process as falling into four general categories: *editors* (facilitate execution of processes), *transformers* (aid in changing existing content), *blenders* (combine 2 or more content sources), and *generators* (produce novel content). In this framing, the large-scale TTI models that this work is focused on fall into the *generators* category.

Additionally, research in the related field of Human-AI Co-Creation (a sub-field of Mixed-Initiative Co-Creativity [61]) is highly relevant. In the *Library of Mixed-Initiative Creative Interfaces* [60], Spoto et al. proposed a framework to understand mixed-initiative co-creation as a process involving seven potential actions: *ideate*, *constrain*, *produce*, *suggest*, *select*, *assess*, and *adapt*. Muller et al. [46] extended this framework to generative AI applications, while Grabe et al. further simplified this extension and characterized four primary interaction patterns concerning GAN applications: curating, exploring, evolving, and conditioning [25]. In our work, we observe similar themes, especially the notion of users of TTI models feeling like they are using the models to explore, and acting as curators of their outputs.

Work in this field has also identified core challenges in creating human-AI co-creation systems. For example, Chung et al. [15] identified the limited ability to control the output of generative AI models and proposed using gestural input to constrain/guide the model output. Likewise, Buschek et al. [11] identified a set of nine challenges system designers could encounter when developing human-AI co-creative systems. In the context of TTI models, they identified challenges of *invisible AI boundaries* (“A (generative) AI component imposes unknown restrictions on creativity and exploration”) and *conflicts of territory* (“AI overwrites what the user has manually created/edited”) as particularly salient. Participants in our study encountered similar challenges of identifying what the models are capable (and not capable) of, and in building upon prior inputs to the system.

Building on this prior CST research, we examine the emergent practices and goals that have evolved in concert with the latest generation of TTI models.

2.2 Generative AI as an Artistic Medium

Given the rapid advancement of AI models for generating images with ever more creative and novel styles, art historians and technologists have been actively discussing how to conceptualize AI-assisted art in relation to other artforms. Experts in these communities have disagreed as to whether generative models should be considered artists in-and-of themselves [41] (as with the famous sale of *Portrait of Edmond Belamy* auctioned by Christie's in 2018 [14]) or whether they should be considered as merely a tool employed by artists. Hertzmann [29, 30] argues that generative AI models are similar to the camera as it relates to the art of photography: a tool the enables the art. Hertzmann further theorizes that “art is an interaction between social agents,” and generative AI models are therefore best considered an agent in this interaction. Agüera y Arcas [6] provides a similar argument with an in-depth discussion of the similarity between the emerging field of AI-generated art and photography, particularly surrounding the historical reaction of painters to the introduction of the camera. Grba [27, 28] provides a framework to critically evaluate art created with a generative AI model. In this framework, he echoes arguments above, and critiques what he refers to as “the ever-receding artist”: the repeated occurrence of technologists referring to models as artists, thereby minimizing the contribution of the human who employed the model to create art. Finally, Browne [10] explored what it means to be an “AI artist” and proposed the framing of an AI artist as a *bricoleur* (building upon [26]), saying, “*Bricolage* is common to generative art, where ideas are developed through playful experimentation with existing tools and techniques.”

In our paper, we present an analysis of interviews with artists who employ a large TTI model for the generation of their art, highlighting the work of three of these artists and provide context around their motivations and goals. In our results, we find thematic alignment with perspectives advanced by Hertmann, Agüera y Arcas, and Grba above: While the models can produce surprisingly high quality output, dedicated users of these models employ the models as tools to explore specific themes and concepts. Accordingly, they have intentionally developed processes (e.g., locating domain-specific terminology) to improve their ability to achieve their individual goals.

2.3 Diffusion and Auto-Regressive Models

Diffusion models (e.g., Imagen [2] or DALL-E [49]) are trained by gradually adding noise to an image, until all of the image is completely noise. The model then learns to reverse the noising process to generate the original image. In this way, a diffusion model learns to synthesize an image from noisy images, and is capable of generating images from arbitrary “noise.” For text-to-image models, the models are also trained (conditioned) on text inputs [47], allowing the model to produce an image from a noisy image input and text input, where the resulting image bears a resemblance to the input text.

Autoregressive models, such as Parti [4], treat text-to-image generation as a sequence-to-sequence problem, akin to machine translation or other language modeling tasks. In the case of TTI models, the “translation” is from text to image (i.e., text tokens to image tokens).

In our study, participants used both types of models.

3 STUDY DESIGN

To understand current practices, motivations, and goals when using modern text-to-image models, we sent a survey to internal users of two internal TTI models (Imagen [2] and Parti [4]) to collect basic information about their use of these models (e.g., time spent using the models, motivations, desired capabilities, and prompting strategies). We also interviewed and observed 11 power users of TTI models (8 identifying as Male and 3 identifying as Female) in a 50-minute study to uncover their motivations and practices. The latter participants were prolific users of one or more TTI models.

3.1 Participants

For the survey, participants were recruited from an internal chat channel dedicated to TTI models (where the internal chat channel has thousands of members) and internal TTI model mailing lists (with hundreds of members). From the survey respondents, we identified interview candidates who had reported having created artwork over more than ten sessions and having spent more than five hours in the previous week using a TTI model. To create a pool of participants, we recruited eight of these latter respondents, and further recruited three prominent artists in the internal artist community to participate. These three artists are quite visible in the internal artist community, and have shared their unique artwork collections within that community. Participants were also actively engaged in external communities, sharing knowledge, expertise, and artwork. Participants were given a 60 USD gift card for participation.

3.2 Interview structure

The study consisted of four parts intended to understand participants’ practices. Each participant was first asked to create an image of their choice to allow the researcher to observe their natural practices. In the second part of the study, participants were asked to reflect on 1) an artwork they were proud of, 2) a piece they found most successful, and 3) the piece that was least successful. In the third part of the interview, participants were asked to reflect on someone else's work by examining only the prompt, and specifically asked to either improve or change the prompt in their style. In the last part of the interview, participants were asked to discuss envisioned uses for these text-to-image models.

Interviews were conducted remotely. We recorded the shared screen and automatically generated transcripts for the interviews. We constrained our focus to interfaces that only use text prompts as input to the models. In addition, some participants voluntarily shared their collection of generated artwork after the interview.

3.3 Qualitative Data Analysis

For the qualitative analyses, the authors analyzed the video transcriptions and also noted comments on participants non-verbal interactions. The final corpus of automatically generated transcripts was 164 pages (60614 words). The first two authors each reviewed the transcripts data independently, looking for ways of explaining the artistic practices [44]. In this process, the authors separately analyzed each transcript to extracted salient themes, and independently generated hypotheses and points of discussions [8]. Using these data, all authors participated in two rounds of interpretation sessions to arrive at the primary themes reported in this paper, and resolve any discrepancies and disagreements. During the interpretation sessions, authors also analyzed the prompts and the images created by the study participants to identify unique artistic styles and practices. These sessions were inspired by existing analysis practices from qualitative media analysis [7].

4 SURVEY RESULTS

We received 161 responses to the survey. Of these responses, 160 answered the question, “At present, why are you using the [TTI] models?” Of the responses to this question, 79 (49%) indicated they use the models to create art, 33 (21%) reported using it as part of their creative work pipeline, and 126 (79%) indicated their use was curiosity-driven (not work-related).

In the survey, we also asked participants to estimate the length of time they work with a TTI model when they use one (“When you interact with a model, how much time do you typically spend interacting with it?”). We received 157 responses to this question, with 20% of respondents indicating that they use a model for one or more hours at a time when they use it (11% reporting using it for 1-2 hours at a time, 9% using it for 2 or more hours at a time), 53% indicating they use a model for 10 minutes to an hour, and 27% reporting use for less than 10 minutes.

When asked about observed strengths and weaknesses of the various models they've interacted with, survey responses indicated a number of desired capabilities, such as the ability to render text in images, the ability to have more control over spatial arrangements, and the ability for models to handle complex prompts. We also asked participants for desired capabilities when interacting with the model. Seventy six percent of the respondents requested a “how to build a prompt” guide, and 75% desired the ability to fork and remix images, especially for spatial refinement. Seventy six percent of the responses also indicated they would like features like bookmarking, and the ability to directly share outputs to internal chat groups or social media. As we will see in the artist spotlights below, there is a clear social component to working with these TTI models for our study participants.

Finally, 63% of the responses desired greater control over the model, such as the ability to assign specific values to each of the prompt words.

When asked to provide prompting techniques they have learned, common themes for the strategies included 1) producing specific art styles and eras, such as “impressionist style”, 2) use of keywords that describe camera lenses and aperture (e.g., “DSLR photo”, “3D render”, “24 mm, f8, ISO1000”), and 3) domain-specific terms (e.g., “Line Art”, “black and white”).

5 PROMPT ARTISTS: STYLES, MOTIVATIONS, PRACTICES

In this section, we provide a sampling of the vibrant internal TTI artistic community by spotlighting the work of three highly active creators (P1, P2, P3) 4. For these three artists, we describe and present examples of the styles they have developed and summarize their artistic motivations and goals. We then provide a summary of motivations, styles, and practices observed across the 11 interview participants. In the section that follows, we describe salient high-level, emergent themes arising from the interviews and observations. We want to credit the artists whose artwork are featured, and they had expressed the desire to be associated with their artwork. We use their full name in places where the artwork appears.

5.1 Shai Noy: The Explorer (P1)

Artist 1 is a software engineer with no training in the visual arts or design. However, they have produced thousands of images with TTI models. The styles developed by this artist include “super macro photography” images (i.e., extremely close-up views of objects, Figure 2a), and fashion (dresses, suits) made out of unusual materials, such as wood, grass, brick, or ice (Figure 2b).

Elements of both discovery and community were emphasized as rewarding for this artist, such as being the first to explore particular concepts and the ability to share discoveries: “Everything is more fun when you can share it” and, “Art doesn't live in a vacuum, nobody starts from scratch, everything is based on something else. I am proud of being able to recognize the potential” (P1).



Figure 2: Selections from Artist 1 (The Explorer)

5.2 Irina Blok: The Art Director (P2)

Artist 2 is a designer who has done visual work for many years using stock images and applications like Photoshop. This artist's styles include origami dancers (Figure 3a) and reality “mash-ups,” such as sliced produce that contains different textures internally (e.g., a sliced head of cabbage that reveals a cross-section of an orange, Figure 3b).


Driving these styles is a passion for developing “impossible objects, something we haven't seen before, mashups” (P2). They also seek to create images that differ significantly from the images it was trained on: “The further away the generated images are from what it was trained on, the more the satisfaction” and, “they're [the resulting images] also defying the rules of its training set, like defying [...] gravity. [...] A creative prompt breaks that” (P2).

Importantly, this artist's output also includes *prompt templates* that describe a particular, parameterized image or visual concept, such as this prompt for generating a house: *Audacious and whimsical fantasy house shaped like <object> with windows and doors, <location>*. These prompt templates are carefully crafted to reliably produce pleasing results for others. We describe this concept more fully in a later section.

Notably, when working with the text-to-image models, P2 considers the model to be akin to an artist itself, with themselves the art director. In this relationship, P2 acknowledges a certain lack of control: “[I] don't have full control, and there's beauty in this” (P2). At the same time, they also consider the model to be a tool: “[The model] is a brush [...] you just learn to speak its language” (P2). With these dual views (model as an artist, model as a tool), they note that “the hardest part is [the] conceptual aspect, being skillful with [the] prompt” and that “it's a thought exercise, it's not a visual exercise” (P2). It's really about how to make people think like an artist” (P2). In this latter sentiment, this participant was specifically speaking to the need to think like an artist in formulating a prompt, as opposed to formulating a prompt as if one was talking to a machine.



(a) Origami dancers



(b) Reality mash-ups

Figure 3: Selections from Artist 2 (The Art Director)

5.3 Dan Smith: The Social Commentator (P3)

Artist 3 comes from a background in visual media, and is partially driven by the desire to deliver a message about the climate crisis, in order to facilitate change and awareness: “[I’m] not just having fun... but [actually] making something that has some power” (P3). In working towards this goal, they want to make “something that you look at, and it makes you feel something” while also making “something that people would want to look at” (P3). P3’s image styles align with these overall goals: They have a number of pieces that put nature in “situations where you wouldn’t see it” (Figure 4a) and have created numerous “hybrid animals” (Figure 4b). This particular style—hybrid animals—is also inline with a motivation to create images that “would be hard to [...] visualize or create, if you were [...] a really skilled Photoshop artist” (see Figure 4).

P3 also heavily considers the quality of the image when assessing its outputs: it must have near expert-level composition, photorealism, and/or artistry. When describing the artwork they created and how they created them, they noted, “I think the ones that I pick out are [...] standouts for various reasons, and just [...] like what I said, [...] composition photorealism, artistic quality” (P3). Consistent with their emphasis on overall image quality, they have found ways to address undesirable outputs. For example, in generating images that include animals, they found they needed to adopt a specific strategy to create aesthetically pleasing images: “I would do ‘Tall Grass’ a lot because early on I discovered that limbs and fingers and paws can get a little wonky” (P3). The “tall grass” addition was their creative strategy to hide feet or paws and suggests an understanding of model limitations, but also a sense of how to cope with these limitations.



(a) Climate change



(b) Hybrid animals

Figure 4: Selections from Artist 3 (The Social Commentator)

5.4 Summarizing Motivations, Styles, and Practices

One of the primary motivations for interacting with the models was that participants found them fun—the models’ output quality enabled people to feel creative, and they were generally interested in interacting with this new class of model. Some people also noted that the model enabled them to engage in their domain interests in a new way. For example, one participant said that the models allowed them to explore their interest in Swiss trains with the model, while another found it compelling to try to create new forms of currency (e.g., new types of coins). One participant used it to create the equivalent of clip art for presentations: “I use [TTI] model when I need some [...] clipart to use in my presentation, and [TTI] model would be amazing for that...” (P5).

In comparing the artists spotlighted above to the other participants, one notable difference in practices is that the spotlighted artists placed a particular emphasis on exploring specific themes in depth (e.g., origami figurines). The other participants did not pursue concepts with the same rigor or depth.

6 THEMES: ART BEYOND IMAGES, DISCOVERING UNIQUE POINTS, AND VALIDATING ORIGINALITY

Across our interviews and observations, a number of themes emerged. First, the notion of what constituted the final artistic output was not always just an image: Some creators consider the *prompt itself* as part of the art. Similarly, a *prompt template* can be considered an artistic output. Second, there was a clear desire for creators to discover new styles possible with the models. This focus on originality extended to one creator going so far as to conduct an image search on successful outputs to validate their originality. We unpack these themes further in this section: 1) Prompt as art, 2) Prompt template as art, 3) Discovering new capabilities of the model, and 4) Validating originality.

6.1 Prompt as Art

For some participants, the prompt itself was part of the overall art, and thus worthy of attention. For these participants, it was important to “[create] aesthetically pleasing images” and “[develop] art concepts” that *were inherently tied to prompts*. We detail these motivations and behaviors below.

While generating aesthetically pleasing, “glitch-free” images was a common goal of the creators, other goals were also present in their practices. For some participants, the prompt itself was part of the overall art, and thus worthy of attention 1) on its own, and 2) as it relates to the image: “It’s part of the aesthetic” (P3), where the prompt is “like a title of the piece, but you don’t get to choose it independently” (P3). Hinted at in this last comment is the notion that while a prompt could also serve as a title of the art piece, there is a clear dependency on, and perfunctory role for, that prompt as well: The prompt serves as the source material for the model generating the resultant image. Given this dependency, finding a prompt that produces the desired result *and* that can serve as a title for the piece can be challenging, but rewarding when it happens.



(a) “a photo of an owl turning into a croissant. f2.2”



(b) “a photo of a bananasaurus rex. a bananasaurus rex has the legs and arms of a banana and the body and head of a tyrannosaurus rex. sigma 85mm f2.2. studio lighting.”

Figure 5: Artist Ian Fischer (P6) considers (a) as a successful art work given the brevity of the prompt, while (b) a failed attempt because the prompt is too long to get the image they want

This same artist (P6) further described the prompt’s role as “communicating the image, and the idea of the image, and how I got it all at the same time” (P6). This quote sheds light on how art with TTI models can, in some sense, be considered multimodal (text and image) for both the artist and viewer: The prompt and the final output combine into a single, mutually-reinforcing, art piece. Seen in this light, one can consider the *prompt as art* itself—a well-crafted prompt that creates a compelling image but also accompanies that image, saying something about the image. See Figure 5 a for an example of “prompt as art,” as well a prompt that wasn’t able to achieve this same level of aesthetic (Figure 5 b).

Diving deeper into this theme, we observed that the artists believe the “concept” is a critical component of an artwork. In an internal blog article, P2 describes, “There’s a common misconception art is largely about drawing and painting skills. Art is not only about how something looks, it’s about what it says, it tells a story, and has a concept. Art can surprise, provoke, teach, delight and inspire. Art is not just about drawing, art is a way of thinking.” P1 also suggested “the unit of shareable artwork is not necessarily a specific image but maybe it’s the whole exploration of the concept of those images.”

In pursuing “prompt as art,” we observed participants impose different constraints when creating their prompts, with some wanting it as descriptive as possible, while others attempting to make it as simple as possible. One participant even discovered delight in their accidental discovery that their “random string” produced beautiful output, and named that prompt as their proudest achievement. In this latter case, the pride comes from the joint

pairing of the random string and the beautiful output—without the context of the random string, the image has less value, as the random string reveals an unexpected feature of the model.

6.2 Prompt Templates As Art



Figure 6: Example of the prompt “Audacious and whimsical fantasy house shaped like <object> with windows and doors, <location>”. All images above use “strawberry” as the shape, with the location varied (“countryside”, “Paris”, “Tokyo”).

In a spirit similar to “prompt as art,” five artists sought to produce *prompt templates*. A prompt template is an image description with “slots” for someone else to fill in. For example, we previously noted this prompt template created by P2: *Audacious and whimsical fantasy house shaped like <object> with windows and doors, <location>* (see Figure 6 for example outputs of this prompt template).

These prompt templates leverage the capabilities of the models as well as the lightweight, accessible features of prompts. More specifically, when an artist identifies a compelling composition, they can create a text prompt that allows others to create a similar composition, but with their own unique customization to it. The ease with which the templates can be shared also introduces social motivations for producing and distributing templates (e.g., to participate and contribute to a larger community of practice). We elaborate on these points below.

Prompt templates have a number of key features:

- They are tightly coupled to and represent a particular artistic concept, vision, or composition (such as a “whimsical fantasy house”). These features make prompt templates conceptually richer than words or phrases used to specify stylistic characteristics of the image (e.g., “35mm” or “watercolors”, which may be used to produce a particular effect, but don't specify a larger composition).
- Others can make use of prompt templates by filling in the blanks. The richness of the models means that the templates guide the overall generated image, but users’ unique input can yield a diverse variety of outputs.
- There is the intent for the prompt template to provide consistently high quality and delightful output when used by others.

Unpacking these concepts in the context of the example prompt above, the skeleton of this prompt embodies a particular (visual) concept: A fantasy home in a given location. Someone making use of this prompt can customize it through two key variables: A shape for the house and a location for the house. While these are seemingly simple variables to customize, the template gives the user great flexibility in terms of the final outputs produced by the model: Any number of shapes can be provided in any number of locations (in fact, the user is free to substitute any text in the slots they wish).

Simultaneously, the prompt cues the model to the *types* of output to produce, as well as details to guide the generation. The phrase “audacious and whimsical fantasy house” defines desired attributes of the house, while the specification of “with windows and doors” provides additional details that should be included in the generated house. These details in the prompt help increase the reliability of the model output and reduce the likelihood that the downstream user of the prompt template needs to experiment further with the overall prompt structure and content to obtain a good output.

A noteworthy characteristic of these templates and the text-to-image models is the rich interplay that results between the original template and the phrases the user substitutes in the template. For example, a particular choice of location can profoundly influence the resulting house generated by the model—the model does not simply generate the same type of house and situate it in a different location. Instead, the choice of location can directly interact with the other parts of the prompt. For example, creating a strawberry house in the “countryside”, “Paris”, or “Tokyo” yields qualitatively different outputs for the house, with the house style meshing more naturally in the chosen location (e.g., having styles of windows more typical for a European building for the house in Paris, and doors more typical of Japan for Tokyo—see Figure 6). These interactions between the template and the users’ choices enable a diverse variety of outputs that allow a user to explore a wide range of ideas.

To produce prompt templates, one participant described a process whereby they would 1) input a prompt, 2) identify high-quality outcomes, and 3) rewrite the prompt to try to produce that same outcome again. This process could involve several iterations until they get a reliable prompt. Once a prompt is working, participants would sometimes remove content to make it more succinct. For example, they would first emphasize a specific characteristic (like “high resolution”) by applying many descriptors representing a similar effect (such as “high res”, “DSLR”, “crystal clear”, “photo realistic”), then start removing the repeated qualifiers in the prompt until the prompt could reliably generate similar outputs to the original, verbose version.

As with the notion of a “prompt as art,” the creation of these prompt templates can also be considered an artistic outcome in and of itself: The prompt author must first develop a compelling concept, then ensure it has enough capacity to enable people to produce their own unique creations within the frame of the concept. When done well, a prompt template has the quality of attracting users (because of the compelling output produced by the prompt) *and* continually delighting users with the outputs produced using their unique input.

Elaborating further on this notion of “prompt template as art,” a prompt template represents a particular *artistic vision* , with the prompt text capturing the overall design, composition, and aesthetic intentions of the author. Notably, because the artistic vision is captured in natural language text, a prompt template can be used *across* TTI models: Users can customize the template as desired, then tweak the prompt to produce the desired output using whatever model they have at their disposal.

6.3 Discovering Unique Points in the Model's Latent Landscape

A common goal of many artists was to discover new capabilities of the model that others had not yet found. As one participant put it, they tried to “break” the model, pushing it to its limits. As with prompt templates, these newly discovered capabilities are often shared so others can apply the concept in their own images. For example, an artist may identify the ability of the model to produce physical sculptures out of unusual materials or to create origami-like figurines (as P2 did). Once a concept and ability have been identified, others can build on the core idea and create their own images or permutations. In particular, participants felt proud when they “made the model do X (new concept)”, especially when they discovered a model capability for the first time in the community in which they heavily engage, as the discovery could influence the discourse in the community. We expand on these points below.

new words to steer the model in new directions. For example, one artist mentioned referencing architectural blogs to learn that domain's vocabulary so it could be applied to their prompts (P2). Another artist mentioned turning to a thesaurus to enrich their vocabulary for the prompt. As one concrete example, P1 (the “Explorer”) likes to employ “unusual words” to steer the model, such as “intricate” instead of “detailed”, explaining: “If you choose common words, then you get a bit of an uninspiring result quite often. But if you use something a bit more unusual, then you really narrow down [...] the set [of images], and you're going to get into things that use this less common word” (P1).

What's noteworthy here is that in the pursuit of novel imagery, creators would carefully research and choose *words* to produce specific *visual* outcomes: *What you say and how you say it is critical to producing high-quality imagery with text-to-image models*, requiring TTI users to *enrich their natural language vocabulary* in order to develop and skillfully execute a unique *visual language*.

Driving these practices of discovering new capabilities was a clear desire to push the model away from “average” outputs and get it into more unique spaces. In this sense, the artists are navigating and charting the vast latent space of the model and sharing back the most interesting places discovered. When a new, unique output space was discovered, artists expressed a certain satisfaction in having discovered that space: “ A good rule of thumb is to be more descriptive than not [...] if you see that something is lacking, then you try to add more descriptors that will encourage the model, in this direction [...] sometimes you just have to reword a sentence or move something from one sentence to another [...] it's mostly like binary search” (P1); “ I think your choice of vocabulary is very interesting because you want a large tree. But then, instead of saying a large tree you say mature tree [...] Where do these [...] choices come from? They just come from trial and error” (P2).

In seeking unique spaces, one participant (P11) expressly hunted for interesting imperfections; instead of *avoiding* glitches, they *sought* glitches. For example, this participant found one of their prompts created imperfections in its output for “a hybrid of a clock and a snail on an infinite mirror. Steampunk. DSLR photo. astrophotography” (Figure 7). Here, P11 explores the imperfections of the infinite mirror: it's “technically a failure but still amazing” (P11). P11 also discovered that the model sometimes has trouble generating the backs of things (like cats) and developed that behavior into its own art style (Figure 7).




Figure 7: Artwork by Paul Emmerich(P11). Examples of “breaking the model”: This artist seeks out and elevates “glitches” in the model's output to create new styles. The artist takes advantage of the model's imperfect reflections (on the left) while making creative use of its tendency to transform the backs of cats into cat faces (see the bottom part of the cat's back, which looks like a cat's face).

6.4 Validating Originality

One question that often arises on the topic of human-in-the-loop, AI-generated art is one of creativity and originality—how much can be attributed to the AI versus the person [30]. Our participants also struggled with this question, with some seeking to ensure their outputs were novel. Participants wanted to validate the originality of the artistic concept, but could only check the originality of the artifact. For example, one participant described a practice of using Google's image search on compelling outputs to ensure originality: After producing a creative output, they use image search to search for that image (or similar images) to ensure it is, in fact, original. This participant mentioned they do this in part because they are sensitized to the fact that machine learning models can sometimes memorize portions of the training data. Given this, they want to ensure their output is *not* in the training data.

7 DISCUSSION

In this research, we have provided a snapshot of the emerging artistic scene enabled by the latest generation of text-to-image models. Our interviews and observations document the types of imagery artists are developing with these new models, as well as an enhanced understanding of the types of results creators seek beyond the images themselves (e.g., prompts and prompt templates as important artifacts in their own right).

In this section, we discuss some implications deriving from artists’ goals of seeking novelty, validating originality, and producing reusable prompt templates. In reviewing design implications for these models, we note that there are countless ways the input and editing interfaces could be improved for interacting with these models (for example, one can find many proposals and working demos online). In this discussion, we restrict ourselves to interfaces that only take text as input, with no post-processing of the image (such techniques that allow in-filling of regions by the model). We limit our discussion to this input modality in part because a number of the artists interviewed were seemingly attracted to the simplicity (and challenge) that this single input modality provides.

7.1 Aiding Novelty

In our interviews, the use of thesauri and domain-specific blogs (e.g., architectural blogs) illustrate the desire of artists to identify unique terms that help them produce results that rise above the average. Embedding these search capabilities directly into the tooling could be useful (e.g., quick access to a thesaurus or an embedded search engine). Pushing this idea further, there may also be opportunities to make use of large language models (LLMs) and/or the TTI model's training data to help surface salient terms for a given topic. For example, an LLM may be able to generate terms-of-art for architecture using a prompt like this: *Here are some terms specific to describing the architectural design of a house: 1).* In our own test with an LLM 5, this prompt yielded terms like *entryway, foyer, entry hall*, and *elevation*. To use the training data itself to identify new spaces, it may be possible to collect terms associated with a particular topic, such as “house,” then identify other terms that are more frequently found near that topic in the training data compared to the rest of the training data (e.g., using a method such as term frequency-inverse document frequency (TF-IDF)).

To help users understand how unique their word choices are, one could also visualize the input prompt with respect to how common each individual input token is. For templates, one could also show what the common terms would be for filling in the blanks to help people move beyond those common terms to find more distinct, unique inputs. One possible outcome in helping people find more novel inputs is that their inputs may lie outside of the training distribution, leading to unpredictable results. Providing feedback through mechanisms like visualizations (e.g., showing frequency in the training data) could help users better understand these types of issues should they arise.

7.2 Validating Originality

As mentioned, there is a desire to validate the originality of the image produced by the model. Streamlining the process of doing an image search with an online search engine is one obvious way to address this issue. However, as is the case in aiding novelty (described above), there may be opportunities to take advantage of the training data itself. More specifically, in addition to external search, one could also search for the closest images in the training data to the image produced.

While the above mechanisms would help validate the outputs generated, there may also be opportunities to help validate the originality of the *input*. For example, one might be able to search the training data to identify which parts of a prompt exist within the training data.

7.3 Materializing Prompt Templates

Prompt templates provide a way for an artist to derive a new style, then share it with others so they can produce their own unique images. One could imagine embracing this practice and transforming a prompt template into its own first-class interface.

For example, one could imagine allowing users to provide multiple inputs for each slot of a prompt template, then generating the cross-product of all the inputs. One could also make use of the fact that the prompt text for the templates is tokenized into vectors. Specifically, by supplying two different inputs for the same slot, embedding vectors for each input could be obtained and then automatically interpolated between the embeddings to produce a spectrum of outputs. For example, for the shape of a house (in the previous house prompt template), the user could provide two inputs: strawberry and apple. The system could then produce the embedding vectors for those inputs, then interpolate between those embeddings to create a series of images that morph from a strawberry to an apple. However, one thing to keep in mind with these interpolations is they are occurring in the text space rather than image space—the model will be interpolating not between shapes (per se) but between the *linguistic concepts* of strawberries and apples (which may still produce an interesting morphing between these conceptual entities).

7.4 Prompts and TTI Models as an Art Medium

One the primary ways the three spotlighted artists (A1, A2, and A3) distinguished themselves from others was their perception of TTI models as an art medium, with a clear focus on exploring the capabilities and limits of medium itself, rather than only on individual outcomes. In this spirit, they embraced the limitation of not being able to edit the images directly, accepting the text-only input as a defining feature and characteristic of the medium. For example, A2 expressed “there's beauty in it” when describing the prompting interaction with the models. In contrast, other participants focused more on the outcome, and expressed desire for additional features, such as direct editing of the generated images. Embracing the TTI models as-is further reinforces the idea that text prompts are part of the artwork, rather than simply a means to an end.

This observation suggests that the design implications for TTI models can be considered from multiple perspectives: prompt-only artists, and creative professionals using the models to achieve specific goals. Creative professionals with a specific design goal may require and request specific features that offer more fine-grained control, perhaps with tools to help understand model behavior. For example, features related to directly editing the generated images were among the most frequently requested features in the survey. Given that non-design experts can also learn to use TTI models to quickly demonstrate and visualize artifacts, we hypothesize that this may facilitate more active and iterative communication between design professionals and their clients, with clients more directly engaged in the creative process (as opposed to the more traditional pipeline-like design workflow). If this proves to be true, feedback and collaboration features could become more important.

7.5 Limitations

Our participant pool was drawn among employees of one large US-based corporation, and does not cover other possible ways that culture, community, and collaboration might shape use of TTI models (e.g. on social media). Also, since our analysis was episodic rather than longitudinal, we do not document how artistic prompting strategies may evolve within individuals. For the interactions we *could* observe, observing participants’ interactions with the TTI model does not definitively indicate their conceptions of how the model works or how best to prompt it. We also acknowledge that different models behave in different ways due to their structure, training data, and the design of the interface. In that regard, we also acknowledge some of the findings might be model dependent, and are specific to the models used in the study. Similarities and differences in art forms and art practices might be observable with different models and communities. Another limitation is that while our participants are actively involved in multiple communities, we did not ask participants about their experiences with other models or other communities in depth.

8 CONCLUSION

In this paper, we have described a unique moment in time: Recent text-to-image models have given rise to an exceptionally vibrant community of practice, complete with new ideas about what constitute notable outcomes (e.g., new styles, prompts as art, and prompt templates as art). As the larger artistic and creator community adopts these new models and forms of art, there are clear ways in which the tools can improve to better support desired practices including: 1) helping to discover and create novel outputs, 2) providing methods to validate novelty, and 3) elevating the notion of a prompt template into a standalone, first class, interactive object. In considering design implications for these TTI models, our results also suggest the value in distinguishing between prompt artists—those users who embrace the constraint of creating images using only an input prompt—and practitioners, who may desire more fine-grained input and editing controls in comparison.

ACKNOWLEDGMENTS

We gratefully appreciate the valuable input from the artists who have participated in our study. Also we thank Dan Liebling, Meredith Ringel Morris, and Google PAIR team for their helpful comments.

REFERENCES

[1] [n. d.]. AICAN. <https://www.aican.io/> (<https://www.aican.io/>). (Accessed on 08/31/2022).

Navigate to ▾

[2] [n. d.]. Imagen: Text-to-Image Diffusion Models. <https://imagen.research.google/> (<https://imagen.research.google/>). (Accessed on 08/31/2022).

Navigate to ▾

[3] [n. d.]. Midjourney. <https://www.midjourney.com/home/> (<https://www.midjourney.com/home/>). (Accessed on 09/01/2022).

Navigate to ▾

[4] [n. d.]. Parti: Pathways Autoregressive Text-to-Image Model. <https://parti.research.google/> (<https://parti.research.google/>). (Accessed on 08/31/2022).

Navigate to ▾

[5] [n. d.]. Stable Diffusion launch announcement — Stability.Ai. <https://stability.ai/blog/stable-diffusion-announcement> (<https://stability.ai/blog/stable-diffusion-announcement>). (Accessed on 08/31/2022).

Navigate to ▾

[6] Blaise Agüera y Arcas. 2017. Art in the Age of Machine Intelligence. *Arts* 6, 4 (2017).
<https://doi.org/10.3390/arts6040018> (<https://doi.org/10.3390/arts6040018>)

Navigate to ▾

[7] David L Altheide and Christopher J Schneider. 2012. *Qualitative media analysis*. Vol. 38. Sage publications.

Navigate to ▾

[8] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative research in psychology* 3, 2 (2006), 77–101.

Navigate to ▾

[9] Andrew Brock, Jeff Donahue, and Karen Simonyan. 2018. Large Scale GAN Training for High Fidelity Natural Image Synthesis. *CoRR* abs/1809.11096 (2018). arXiv:[1809.11096](https://arxiv.org/abs/1809.11096) (<https://arXiv.org/abs/1809.11096>) <http://arxiv.org/abs/1809.11096> (<http://arxiv.org/abs/1809.11096>)

Navigate to ▾

[10] Kieran Browne. 2022. Who (or What) Is an AI Artist?*Leonardo* 55, 2 (04 2022), 130–134.
https://doi.org/10.1162/leon_a_02092 (https://doi.org/10.1162/leon_a_02092)
arXiv:https://direct.mit.edu/leon/article-pdf/55/2/130/2004755/leon_a_02092.pdf (https://direct.mit.edu/leon/article-pdf/55/2/130/2004755/leon_a_02092.pdf)

Navigate to ▾

[11] Daniel Buschek, Lukas Mecke, Florian Lehmann, and Hai Dang. 2021. Nine Potential Pitfalls when Designing Human-AI Co-Creative Systems. *Workshops at the International Conference on Intelligent User Interfaces (IUI)* (2021).

Navigate to ▾

[12] Eva Cetinic and James She. 2022. Understanding and Creating Art with AI: Review and Outlook. *ACM Trans. Multimedia Comput. Commun. Appl.* 18, 2, Article 66 (Feb 2022), 22 pages.
<https://doi.org/10.1145/3475799> (<https://doi.org/10.1145/3475799>)

Navigate to ▾

[13] Eugene Ch'ng. 2019. Art by Computing Machinery: Is Machine Art Acceptable in the Artworld?*ACM Trans. Multimedia Comput. Commun. Appl.* 15, 2s, Article 59 (Jul 2019), 17 pages.
<https://doi.org/10.1145/3326338> (<https://doi.org/10.1145/3326338>)

Navigate to ▾

[14] Christie's. 2018. The first piece of AI-generated art to come to auction | Christie's.
<https://www.christies.com/features/a-collaboration-between-two-artists-one-human-one-a-machine-9332-1.aspx> (<https://www.christies.com/features/a-collaboration-between-two-artists-one-human-one-a-machine-9332-1.aspx>). (Accessed on 09/06/2022).

Navigate to ▾

[15] John Joon Young Chung, Minsuk Chang, and Eytan Adar. 2021. Gestural Inputs as Control Interaction for Generative Human-AI Co-Creation. *Workshops at the International Conference on Intelligent User Interfaces (IUI)* (2021).

Navigate to ▾

[16] John Joon Young Chung, Shiqing He, and Eytan Adar. 2021. The Intersection of Users, Roles, Interactions, and Technologies in Creativity Support Tools. In *Designing Interactive Systems Conference 2021 (Virtual Event, USA)* (DIS '21). Association for Computing Machinery, New York, NY, USA, 1817–1833.
<https://doi.org/10.1145/3461778.3462050> (<https://doi.org/10.1145/3461778.3462050>)

Navigate to ▾

[17] Nicholas Davis, Chih-PIn Hsiao, Kunwar Yashraj Singh, Lisa Li, Sanat Moningi, and Brian Magerko. 2015. Drawing Apprentice: An Enactive Co-Creative Agent for Artistic Collaboration. In *Proceedings of the 2015 ACM SIGCHI Conference on Creativity and Cognition (Glasgow, United Kingdom) (C&C '15)*. Association for Computing Machinery, New York, NY, USA, 185–186. <https://doi.org/10.1145/2757226.2764555>. (<https://doi.org/10.1145/2757226.2764555>)

Navigate to ▾

[18] Nicholas Davis, Chih-PIn Hsiao, Kunwar Yashraj Singh, Lisa Li, and Brian Magerko. 2016. Empirically Studying Participatory Sense-Making in Abstract Drawing with a Co-Creative Cognitive Agent. In *Proceedings of the 21st International Conference on Intelligent User Interfaces (Sonoma, California, USA) (IUI '16)*. Association for Computing Machinery, New York, NY, USA, 196–207.
<https://doi.org/10.1145/2856767.2856795> (<https://doi.org/10.1145/2856767.2856795>)

Navigate to ▾

[19] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. Association for Computational Linguistics, Minneapolis, Minnesota, 4171–4186. <https://doi.org/10.18653/v1/N19-1423> (<https://doi.org/10.18653/v1/N19-1423>)

Navigate to ▾

[20] Prafulla Dhariwal and Alexander Nichol. 2021. Diffusion Models Beat GANs on Image Synthesis. In *Advances in Neural Information Processing Systems*, M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan (Eds.). Vol. 34. Curran Associates, Inc., 8780–8794.
<https://proceedings.neurips.cc/paper/2021/file/49ad23d1ec9fa4bd8d77d02681df5cfa-Paper.pdf> (<https://proceedings.neurips.cc/paper/2021/file/49ad23d1ec9fa4bd8d77d02681df5cfa-Paper.pdf>)

Navigate to ▾

[21] Ahmed Elgammal, Bingchen Liu, Mohamed Elhoseiny, and Marian Mazzone. 2017. CAN: Creative Adversarial Networks, Generating "Art" by Learning About Styles and Deviating from Style Norms.
<https://doi.org/10.48550/ARXIV.1706.07068> (<https://doi.org/10.48550/ARXIV.1706.07068>)

Navigate to ▾

[22] Jonas Frich, Lindsay MacDonald Vermeulen, Christian Remy, Michael Mose Biskjaer, and Peter Dalsgaard. 2019. Mapping the Landscape of Creativity Support Tools in HCI. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (Glasgow, Scotland Uk) (CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–18. <https://doi.org/10.1145/3290605.3300619> (<https://doi.org/10.1145/3290605.3300619>)

Navigate to ▾

[23] Leon A. Gatys, Alexander S. Ecker, and Matthias Bethge. 2015. A Neural Algorithm of Artistic Style.
<https://doi.org/10.48550/ARXIV.1508.06576> (<https://doi.org/10.48550/ARXIV.1508.06576>)

Navigate to ▾

[24] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2020. Generative Adversarial Networks. *Commun. ACM* 63, 11 (Oct 2020), 139–144. <https://doi.org/10.1145/3422622> (<https://doi.org/10.1145/3422622>)

Navigate to ▾

[25] Imke Grabe, Miguel González-Duque, Sebastian Risi, and Jichen Zhu. 2022. Towards a Framework for Human-AI Interaction Patterns in Co-Creative GAN Applications. *Workshops at the International Conference on Intelligent User Interfaces (IUI)* (2022).

Navigate to ▾

[26] Dejan Grba. 2019. Forensics of a molten crystal: challenges of archiving and representing contemporary generative art. *ISSUE Annual Art Journal: Erase* 8, 3-15 (2019), 5.

Navigate to ▾

[27] Dejan Grba. 2021. Brittle Opacity: Ambiguities of the Creative AI. In *Proceedings of the xCoAx, 9th Conference on Computation, Communication, Aesthetics & X Proceedings, xCoAx, Graz, Austria*. 12–16.

Navigate to ▾

[28] Dejan Grba. 2022. Deep Else: A Critical Framework for AI Art. *Digital* 2, 1 (2022), 1–32.
<https://doi.org/10.3390/digital2010001> (<https://doi.org/10.3390/digital2010001>)

Navigate to ▾

[29] Aaron Hertzmann. 2018. Can Computers Create Art?*Arts* 7, 2 (2018).
<https://doi.org/10.3390/arts7020018> (<https://doi.org/10.3390/arts7020018>)

Navigate to ▾

[30] Aaron Hertzmann. 2020. Computers Do Not Make Art, People Do. *Commun. ACM* 63, 5 (Apr 2020), 45–48. <https://doi.org/10.1145/3347092> (<https://doi.org/10.1145/3347092>)

Navigate to ▾

[31] Jonathan Ho, Chitwan Saharia, William Chan, David J Fleet, Mohammad Norouzi, and Tim Salimans. 2022. Cascaded Diffusion Models for High Fidelity Image Generation.*J. Mach. Learn. Res.* 23 (2022), 47–1.

Navigate to ▾

[32] Rania Hodhod and Brian Magerko. 2016. Closing the Cognitive Gap between Humans and Interactive Narrative Agents Using Shared Mental Models. In *Proceedings of the 21st International Conference on Intelligent User Interfaces (Sonoma, California, USA) (IUI '16)*. Association for Computing Machinery, New York, NY, USA, 135–146. <https://doi.org/10.1145/2856767.2856774>. (<https://doi.org/10.1145/2856767.2856774>)

Navigate to ▾

[33] Cheng-Zhi Anna Huang, Hendrik Vincent Koops, Ed Newton-Rex, Monica Dinculescu, and Carrie J. Cai. 2020. AI Song Contest: Human-AI Co-Creation in Songwriting. (2020).
<https://doi.org/10.48550/ARXIV.2010.05388> (<https://doi.org/10.48550/ARXIV.2010.05388>)

Navigate to ▾

[34] Angel Hsing-Chi Hwang. 2022. Too Late to Be Creative? AI-Empowered Tools in Creative Processes. In *Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems (New Orleans, LA, USA) (CHI EA '22)*. Association for Computing Machinery, New York, NY, USA, Article 38, 9 pages.
<https://doi.org/10.1145/3491101.3503549> (<https://doi.org/10.1145/3491101.3503549>)

Navigate to ▾

[35] Youngseung Jeon, Seungwan Jin, Patrick C. Shih, and Kyungsik Han. 2021. FashionQ: An AI-Driven Creativity Support Tool for Facilitating Ideation in Fashion Design. In Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (Yokohama, Japan) (CHI ’21). Association for Computing Machinery, New York, NY, USA, Article 576, 18 pages. <https://doi.org/10.1145/3411764.3445093> (<https://doi.org/10.1145/3411764.3445093>)
[Navigate to ▼]

[36] Yongcheng Jing, Yezhou Yang, Zunlei Feng, Jingwen Ye, Yizhou Yu, and Mingli Song. 2020. Neural Style Transfer: A Review. *IEEE Transactions on Visualization and Computer Graphics* 26, 11 (2020), 3365–3385. <https://doi.org/10.1109/TVCG.2019.2921336> (<https://doi.org/10.1109/TVCG.2019.2921336>)
[Navigate to ▼]

[37] Pegah Karimi, Nicholas Davis, Mary Lou Maher, Kazjon Grace, and Lina Lee. 2019. Relating Cognitive Models of Design Creativity to the Similarity of Sketches Generated by an AI Partner. In Proceedings of the 2019 on Creativity and Cognition (San Diego, CA, USA) (C&C ’19). Association for Computing Machinery, New York, NY, USA, 259–270. <https://doi.org/10.1145/3325480.3325488> (<https://doi.org/10.1145/3325480.3325488>)
[Navigate to ▼]

[38] Tero Karras, Samuli Laine, and Timo Aila. 2019. A Style-Based Generator Architecture for Generative Adversarial Networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
[Navigate to ▼]

[39] Ryan Louie, Andy Coenen, Cheng Zhi Huang, Michael Terry, and Carrie J. Cai. 2020. Novice-AI Music Co-Creation via AI-Steering Tools for Deep Generative Models. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (Honolulu, HI, USA) (CHI ’20). Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3313831.3376739> (<https://doi.org/10.1145/3313831.3376739>)
[Navigate to ▼]

[40] Elman Mansimov, Emilio Parisotto, Jimmy Lei Ba, and Ruslan Salakhutdinov. 2015. Generating Images from Captions with Attention. <https://doi.org/10.48550/ARXIV.1511.02793> (<https://doi.org/10.48550/ARXIV.1511.02793>)
[Navigate to ▼]

[41] Marian Mazzone and Ahmed Elgammal. 2019. Art, Creativity, and the Potential of Artificial Intelligence. *Arts* 8, 1 (2019). <https://doi.org/10.3390/arts8010026> (<https://doi.org/10.3390/arts8010026>)
[Navigate to ▼]

[42] Jon McCormack, Toby Gifford, Patrick Hutchings, Maria Teresa Llano Rodriguez, Matthew Yee-King, and Mark d’Inverno. 2019. In a Silent Way: Communication Between AI and Improvising Musicians Beyond Sound. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (Glasgow, Scotland Uk) (CHI ’19). Association for Computing Machinery, New York, NY, USA, 1–11. <https://doi.org/10.1145/3290605.3300268> (<https://doi.org/10.1145/3290605.3300268>)
[Navigate to ▼]

[43] Matt McFarland. 2016. Google’s psychedelic ‘paint brush’ raises the oldest question in art - The Washington Post. <https://www.washingtonpost.com/news/innovations/wp/2016/03/10/googles-psychedelic-paint-brush-raises-the-oldest-question-in-art/> (<https://www.washingtonpost.com/news/innovations/wp/2016/03/10/googles-psychedelic-paint-brush-raises-the-oldest-question-in-art/>). (Accessed on 09/14/2022).
[Navigate to ▼]

[44] Matthew B Miles and A Michael Huberman. 1984. Drawing valid meaning from qualitative data: Toward a shared craft. *Educational researcher* 13, 5 (1984), 20–30.
[Navigate to ▼]

[45] Alexander Mordvintsev, Christopher Olah, and Mike Tyka. 2015. Inceptionism: Going Deeper into Neural Networks. <https://research.googleblog.com/2015/06/inceptionism-going-deeper-into-neural.html> (<https://research.googleblog.com/2015/06/inceptionism-going-deeper-into-neural.html>)
[Navigate to ▼]

[46] Michael Muller, Justin D Weisz, and Werner Geyer. 2020. Mixed Initiative Generative AI Interfaces: An Analytic Framework for Generative AI Applications. In *Proceedings of the Workshop The Future of Co-Creative Systems-A Workshop on Human-Computer Co-Creativity of the 11th International Conference on Computational Creativity (ICCC 2020)*.
[Navigate to ▼]

[47] Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. 2021. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. *arXiv preprint arXiv:2112.10741* (<https://arXiv.org/abs/2112.10741>) (2021).
[Navigate to ▼]

[48] Changhoon Oh, Jungwoo Song, Jinhan Choi, Seonghyeon Kim, Sungwoo Lee, and Bongwon Suh. 2018. I Lead, You Help but Only with Enough Details: Understanding User Experience of Co-Creation with Artificial Intelligence. In Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (Montreal QC, Canada) (CHI ’18). Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3173574.3174223> (<https://doi.org/10.1145/3173574.3174223>)
[Navigate to ▼]

[49] OpenAI. [n. d.]. DALL·E 2. <https://openai.com/dall-e-2/> (<https://openai.com/dall-e-2/>). (Accessed on 08/31/2022).
[Navigate to ▼]

[50] OpenAI. [n. d.]. DALL·E: Creating Images from Text. <https://openai.com/blog/dall-e/> (<https://openai.com/blog/dall-e/>). (Accessed on 08/31/2022).
[Navigate to ▼]

[51] Allison Perrone and Justin Edwards. 2019. Chatbots as Unwitting Actors. In Proceedings of the 1st International Conference on Conversational User Interfaces (Dublin, Ireland) (CUI ’19). Association for Computing Machinery, New York, NY, USA, Article 2, 2 pages. <https://doi.org/10.1145/3342775.3342799> (<https://doi.org/10.1145/3342775.3342799>)
[Navigate to ▼]

[52] Brian Quanz, Wei Sun, Ajay Deshpande, Dhruv Shah, and Jae-eun Park. 2020. Machine learning based co-creative design framework. *arXiv preprint arXiv:2001.08791* (<https://arXiv.org/abs/2001.08791>) (2020).
[Navigate to ▼]

[53] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, Peter J Liu, *et al.* 2020. Exploring the limits of transfer learning with a unified text-to-text transformer.*J. Mach. Learn. Res.* 21, 140 (2020), 1–67.
[Navigate to ▼]

[54] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. 2022. Hierarchical Text-Conditional Image Generation with CLIP Latents. <https://doi.org/10.48550/ARXIV.2204.06125> (<https://doi.org/10.48550/ARXIV.2204.06125>)
[Navigate to ▼]

[55] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. 2021. Zero-Shot Text-to-Image Generation. In Proceedings of the 38th International Conference on Machine Learning(Proceedings of Machine Learning Research, Vol. 139), Marina Meila and Tong Zhang (Eds.). PMLR, 8821–8831. <https://proceedings.mlr.press/v139/ramesh21a.html> (<https://proceedings.mlr.press/v139/ramesh21a.html>)
[Navigate to ▼]

[56] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-Resolution Image Synthesis With Latent Diffusion Models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 10684–10695.
[Navigate to ▼]

[57] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S. Sara Mahdavi, Rapha Gontijo Lopes, Tim Salimans, Jonathan Ho, David J Fleet, and Mohammad Norouzi. 2022. Photorealistic Text-to-Image Diffusion Models with Deep Language Understanding. <https://doi.org/10.48550/ARXIV.2205.11487> (<https://doi.org/10.48550/ARXIV.2205.11487>)
[Navigate to ▼]

[58] Othman Sbai, Mohamed Elhoseiny, Antoine Bordes, Yann LeCun, and Camille Couprie. 2018. DesIGN: Design Inspiration from Generative Networks. In *Proceedings of the European Conference on Computer Vision (ECCV) Workshops*.
[Navigate to ▼]

[59] Yang Shi, Nan Cao, Xiaojuan Ma, Siji Chen, and Pei Liu. 2020. EmoG: Supporting the Sketching of Emotional Expressions for Storyboarding. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (Honolulu, HI, USA) (CHI ’20). Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/3313831.3376520> (<https://doi.org/10.1145/3313831.3376520>)
[Navigate to ▼]

[60] Angie Spoto and Natalia Oleynik. [n. d.]. Library of Mixed-Initiative Creative Interfaces. <http://mici.codingconduct.cc/aboutmicis/> (<http://mici.codingconduct.cc/aboutmicis/>). (Accessed on 08/31/2022).
[Navigate to ▼]

[61] Georgios N Yannakakis, Antonios Liapis, and Constantine Alexopoulos. 2014. Mixed-initiative co-creativity. (2014).
[Navigate to ▼]

[62] Jiahui Yu, Yuanzhong Xu, Jing Yu Koh, Thang Luong, Gunjan Baid, Zirui Wang, Vijay Vasudevan, Alexander Ku, Yinfei Yang, Burcu Karagol Ayan, Ben Hutchinson, Wei Han, Zarana Parekh, Xin Li, Han Zhang, Jason Baldridge, and Yonghui Wu. 2022. Scaling Autoregressive Models for Content-Rich Text-to-Image Generation. <https://doi.org/10.48550/ARXIV.2206.10789> (<https://doi.org/10.48550/ARXIV.2206.10789>)
[Navigate to ▼]

[63] Nanxuan Zhao, Nam Wook Kim, Laura Mariah Herman, Hanspeter Pfister, Rynson W.H. Lau, Jose Echevarria, and Zoya Bylinskii. 2020. ICONATE: Automatic Compound Icon Generation and Ideation. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (Honolulu, HI, USA) (CHI ’20). Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3313831.3376618> (<https://doi.org/10.1145/3313831.3376618>)
[Navigate to ▼]

[64] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. 2017. Unpaired Image-To-Image Translation Using Cycle-Consistent Adversarial Networks. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*.
[Navigate to ▼]

FOOTNOTE

¹ <https://promptomania.com/prompt-builder/> (<https://promptomania.com/prompt-builder/>)

² <https://promptbase.com/shop> (<https://promptbase.com/shop>)

³We thank our anonymous reviewers for the specific phrase recognizing our contribution.

⁴In the text, we use the terms “artists,” “creators,” and “study participants” interchangeably.

⁵Specific model anonymized for review



This work is licensed under a [Creative Commons Attribution International 4.0 License \(https://creativecommons.org/licenses/by/4.0/\)](https://creativecommons.org/licenses/by/4.0/).

C&C '23, June 19–21, 2023, Virtual Event, USA

© 2023 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0180-1/23/06.

DOI: <https://doi.org/10.1145/3591196.3593515> (~~<https://doi.org/10.1145/3591196.3593515>~~)