

When happy accidents spark creativity: Bringing collaborative speculation to life with generative AI

Ziv Epstein^{1*}, Hope Schroeder^{1*}, Dava Newman^{1,2}

¹MIT Media Lab

²Human Systems Laboratory, Department of Aeronautics and Astronautics, MIT
{ziv, hopes, dnewman}@mit.edu

*These authors contributed equally

Abstract

Generative AI techniques like those that synthesize images from text (text-to-image models) offer new possibilities for creatively imagining new ideas. We investigate the capabilities of these models to help communities engage in conversations about their collective future. In particular, we design and deploy a facilitated experience where participants collaboratively speculate on utopias they want to see, and then produce AI-generated imagery from those speculations. In a series of in-depth user interviews, we invite participants to reflect on the generated images and refine their visions for the future. We synthesize findings with a bespoke community zine on the experience. We observe that participants often generated ideas for implementing their vision and drew new lateral considerations as a result of viewing the generated images. Critically, we find that the unexpected difference between the participant’s imagined output and the generated image is what facilitated new insight for the participant. We hope our experimental model for co-creation, computational creativity, and community reflection inspires the use of generative models to help communities and organizations envision better futures.

Introduction

New methods in generative machine learning, such as GANs, have created an explosion of opportunities and possibilities for computational creativity. One possibility GANs afford is the empowerment of people in casual creation (Compton and Mateas, 2015; Epstein et al., 2020a; Berns and Colton, 2020). For such casual creation, text-to-image models built on technologies like CLIP (Radford et al., 2021) have enormous promise by providing people an intuitive and user-friendly possibility space to query GAN-generated images via prompts (see Colton et al. (2021) for an overview). With such a technology, what are the possibilities for human creativity, and how might these applications impact communities? One opportunity for these text-to-image models is to aid in *imaginative idea visualization* (Colton et al., 2021) as a way to bootstrap the creative process. An issue of particular relevance to this goal is the inherent variety/fidelity trade-off of generative models: as the outputs of models become more realistic, they become less diverse (Ramesh et al., 2022). Yet it remains unclear how this trade-

off impacts downstream tasks like imaginative idea visualization. On one hand, high-fidelity outputs might provide helpful details for how a given idea might be actually implemented. On the other, “happy accidents” from diverse but low-fidelity outputs might help evolve an idea via lateral thinking.

To explore how these two possibilities trade off, we design, deploy, and evaluate a computational creativity system for imagining and visualizing new ideas. Participants collaboratively speculated on utopian ideas for the future. These speculations were then fed as text prompts into a generative AI model to visually manifest them. We conducted a series of user interviews to learn about the experience from participants, and surface key themes. In this paper, we present a field report of the experience and use the system to trace broader questions about the social and collaborative aspects of creativity, such as when generative visual imagery can inspire ideas about the future, and how the variety/fidelity trade-off in generative models might impact creativity.

This work builds on the tradition of speculative design and design fiction, which use design to imagine and prototype alternative worlds (Dunne and Raby, 2013). The social nature of the collaborative speculation allows individuals to build these alternative worlds together. Furthermore, it explores how organizations or communities could use speculative design and design fiction to chart a course for the future.

There are several key contributions in this field report. First, we introduce a novel approach to prompt engineering, which leverages collaborative speculation from a facilitated co-design experience, rather than a single individual sourcing the prompts. Second, we highlight the importance of high-variance, low-fidelity images in inducing creative insights. Finally, we highlight the potential for computational creativity to aid in community dialogue and the collective elicitation of an organization’s values. We explore this possibility by creating a zine to document and synthesize the findings, which we then present back to the community.

Methods

On October 8th, 2021, we organized a facilitated co-design experience at a solarpunk¹-themed event at the MIT Media

¹Solarpunk is an artform and aesthetic that imagines near-distant futures where humans have become climate change-

Lab, which had over 400 RSVPs. Participants in the exercise paired up and responded to the following prompt: “How will we re-imagine the following categories in utopia? Team up with someone, and each pick one of the following sectors: (see Table 1). Brainstorm how the future could be better at their intersection! Pick up a leaf and write your vision down on it. Add any visual representation you want. Tape the leaves to a stalk, add a flower, and put them in the solarpunk garden together!”

Money	Medicine	Cities	Transportation
Space	Agriculture	Music	Environment
Economy	Relationships	Family	Healthcare
Arts	Civil rights	Fashion	Infrastructure
Trade	Social justice	Oceans	Natural land & Wildlife
Education	Government	Energy	Community

Table 1: The 20 sectors used for prompt generation. Pairs of participants picked two to blend.

After co-creating their visions, participants placed their flowers with the vision written on it in the solarpunk garden (Figure 1). A total of 32 visions of the future were co-created by over 64 participants.



Figure 1: Solarpunk garden where participants placed their written visions. The garden was conceptualized as a “seed vault,” a place for hopeful seeds for the future to be stored and preserved.

Next, we took the 32 visions, and ran them through VQGAN+CLIP, a common model of text-to-image synthesis. This produced visual representations of the participants’ visions, and served as the output of the speculation experience. We then used these images for evaluation of the paradigm and incorporated them into a zine for additional community impact.

Evaluation via user interviews To evaluate the experience of facilitated speculation augmented with generative AI, we conducted a series of ten 15-20 minute semi-structured interviews from March 30th to April 11th, 2022 about 10 unique visions both in person and by video chat.

First, to measure relatedness between prompt and image, we tested if participants could recognize the image generated by their prompt. For each participant, we paired the image that corresponded to their prompt with three additional images from other prompts and presented the set of four to the participant in random order. We invited participants to

resilient and learned to live in harmony with nature.

identify the image generated from the prompt they wrote in the original activity. Regardless of their answer, we then revealed the correct image. We then asked questions related to their interpretation of the image, starting by asking them to describe what they saw. Based on directions from this initial description, we then asked follow-up questions about if and how the image differed from their expectations. Finally, we asked if the participant had ideas for a follow-up prompt that better reflected their vision. After interviews were complete, we coded them using standard qualitative coding practices for themes that emerged across conversations (Saldaña, 2013).

Results

While all loosely related to the solarpunk theme, the written prompts from participants in the original activity had a remarkable amount of diversity. Some participants generated futuristic ideas, while others called for a return to traditional wisdom or practices. Four exemplary images, with their corresponding prompts are shown in Figure 2.

User Interviews We conducted 10 interviews with individuals who participated in the original activity. We focused our interviews to participants who remembered the conversation they had from several months prior and still recalled the vision that they wrote down.

We coded our 10 participant interviews for major themes distilled through repeated interviewing. In the process of synthesizing recurring themes across conversations, we noticed that many participants gleaned new insight as a result of viewing the image. We recorded when participants mentioned new information gleaned about their original idea after viewing the generated image. A majority of participants reported gaining new ideas as a result of viewing the image generated by their prompt, and this new information fell into two main categories.

The first was the emergence of new, unexpected ideas for implementing the vision they had written (40% of interviews, $n = 4$). For example, the creator of the prompt “Biophilic vertical gardens lining neighborhood roads, creating function and beautiful public spaces” already had a “pretty concrete” image of what their vision could look like given that biophilic community gardens already exist. However, they had imagined “pumps and tubes” as a visual feature given the complex engineering required to create such gardens. When viewing the generated image, the creator noted that the wall looked like it was made of a “natural, rocky substrate” rather than one with exposed hydroponic engineering (see Figure 2, left). This led the participant to note that perhaps the wall of the garden could in fact be beautiful and natural-looking as well as functional, and that this aesthetic would be an improvement over exposed pipes.

The second main area of insight participants reported in interviews related to generating unexpected connections between a vision and lateral concepts (40%, $n = 4$). For example, the prompt “Public spaces: solidarity-building. The intersection of oceans and relationships. Publicly accessible oceanic vistas” generated an image of humans on a beach (see Figure 2, center left). Something like sand is present,



Figure 2: Four images from the following prompts: Biophilic vertical gardens lining neighborhood roads, creating function and beautiful public spaces (left), Public spaces: solidarity-building. The intersection of oceans and relationships. Publicly accessible oceanic vistas (center left), Holistic traditional medicine as an art form (center right), Dye the ocean purple to prevent global warming (right)

but the lack of an actual ocean in the image was surprising to the prompt’s creator. However, the repeated visual motif of people in relationship yielded a new perspective for the creator, who then reflected on the importance of relationship for organizing in public space. This was not a framing he had not been considering as central to this prompt before being presented with the image. The lack of ocean in an image generated from a prompt with two references to the ocean could be considered low-fidelity and ultimately undesirable behavior from VQGAN+CLIP. Yet the unexpectedness of the image led the prompt’s creator in a new direction that was ultimately valuable for ideation. As a result, we consider this a happy accident in the context of this exercise in social dreaming.

In another example, the participant who wrote “Holistic traditional medicine as an art form” reported noticing the centrality of hands as a visual motif in the image (see Figure 2, center right). Her interpretation was that the hands made the image “focused on the making process” and that the image “emphasizes labor.” She reported that the human labor aspect behind traditional medicine was not a major association she had with the topic before viewing the image.

Most (80%, $n = 8$) interviewees mentioned at least one idea for a follow-up prompt to refine the image or clarify their vision. For example, the biophilic garden image featured a visual element on the bottom left that looked like a road. The prompt’s creator noticed the curb and said they might use language like “neighborhood path” instead of “road” to generate an image with explicitly pedestrian streets in the future.

Despite the fact that it was easy for participants to identify the image created from their prompt (90%, $n = 9$) out of a set of random images, most participants interpreted the generated images as being either partly or substantially different from what they imagined prior to seeing the image (70%, $n = 7$). Three interviewees that did not gain any new insight from viewing the images were concerned that the image was too abstract to be useful. On the other hand, two other participants commented that they specifically appreciated the “whimsical” image of a technical concept.

“When the place inspires the zine inspires the place”²

In the wake of the solarpunk event, we created a zine outlining the project and showcasing all the co-written visions and corresponding GAN images to aid in community reflection (see Figure 3). We printed over 500 copies and distributed them to community stakeholders. As community members, we wanted to give these images back to the community, hoping they would act as a mirror for additional conversation and “close the loop” of the reflective process.

The event, co-creation experience, and zine all came at a time when our organization was emerging from the throes of the pandemic and actively setting its strategy and vision for the next decade and beyond. At such a critical inflection point, the combination of these three elements had several key cultural impacts. First, there was the physical aspect of people coming together to imagine and plant the seeds for the future after a long period of pandemic-induced isolation. Second, in an organization with varied interests and priorities, the process of brainstorming through images helped visualize and synthesize collective threads across the organization. Finally, cementing the images in familiar, tangible form of printed media provided a snapshot for both disseminating the outputs and orienting discussion around the organization’s future (Pérez y Pérez and Ackerman, 2020). The zine has since been used to communicate interests and priorities to both visitors and external stakeholders.

Discussion

Initially, we predicted a main contribution of this work would be designing a tool to concretize visions that teams co-created. We found that occasionally the image did meaningfully crystallize a vision for the future in a literal way, and many participants noted the image gave them ideas for ways to implement their vision. However, we often found that it was the unexpected differences between the prompt and the generated image’s interpretation of it that yielded new insight for and excitement from participants. This suggests that it was the high-variance, low-fidelity behavior of

²Quote by Sidebody <https://www.instagram.com/p/CbtJnW1loui/>



Figure 3: Cover and four select pages from the Seedvault zine. Written in the style of speculative fiction, the zine takes place in the year 2121, recounting an event that took place 100 years ago and how the previous 100 years unfolded from there.

VQGAN+CLIP that yielded these unexpected and whimsical differences, which in turn induced novel and creative insights. This finding has important implications for the design of image synthesis systems for computational creativity. As new systems become increasingly realistic (such as new models like DALL-E 2 (Ramesh et al., 2022)), there is a danger that this increased fidelity will come at the cost of these unexpected quirks that we found actually sparked positive lateral thinking in our participants.

By offering a low-friction and provocative new way to integrate visual storytelling with community engagement, text-to-image models with collaboratively sourced prompts have an opportunity for cultural and organizational impact. This approach could be used in contexts as diverse as crafting organization mission statements, co-designing community interventions, and facilitating a mediation process.

There are, however, dangers to such an approach. For one, AI-generated images are bounded by training data, which inherits historical biases and cultural practices (Ganivals Blog, 2020; Crawford and Paglen, 2019). Therefore uncritically relying on model outputs as an “oracle” may entrench users in existing inequities and stereotypes, rather than freeing them to envision radical new possibilities. Relatedly, if such an approach becomes commonplace, there is the risk that such visualization strategies could be a crutch if used too much, with users becoming overly reliant on a machine’s vision of the future rather than their own. Finally, there is the risk that anthropomorphizing the AI can undermine human credit and responsibility (Epstein et al., 2020b).

It is also important to reflect on the notion of “utopias,” which we used to frame this collaborative speculation. Literally meaning “no place” in the original Greek, utopias represent idealized non-existent or impossible societies. Rarely in utopian thinking are questions like “A utopia for whom?” or “A utopia at what cost?” considered, which in turn leads to imagined futures that are culturally homogeneous and can perpetuate existing inequities. Furthermore, the very concept of utopias imagines a future society different from and evolved upon the current one, a notion rooted in endless growth, a potentially colonial and capitalistic value (Morri-

son, 2017). Rather than envisioning radical new alternatives to be manifested, many indigenous cultures instead hold a worldview that humans are a part of a complex web of ecological relationships that can exist in a perpetual steady-state equilibrium (Kimmerer, 2013; Sepie, 2017). In the prompts we received, we saw varied interpretations of and priorities for utopias: some participants longed for a return to local economies or traditional medicine, while others dreamt of futuristic space suits and wind turbines. Finally, it is important to note that utopias are but one frame for imaginative idea visualization. While we used the idea of utopia to foster brainstorming about the future, we believe the paradigm could instead focus inward for any community or space looking to collaboratively speculate about its future.

Our work has several limitations which open up exciting possibilities for future work. For one, we focused on one particular community – the MIT Media Lab, a technology-focused research institution embedded in a university. Future work could explore how this approach works for other communities with distinct cultures and practices, and how the effectiveness of the approach varies across community contexts. In addition, we relied on analog media throughout the process: from handwritten prompts by participants and a paper zine distributed in person to us showing participants generated images asynchronously in the evaluation phase. Future work could explore digital versions of these methods, such as an online interface for collaborative prompt generation, an online gallery for the prompts and their corresponding images, or real-time dialogue with a generative model. We also did not include in-depth explanation for how the generative model worked. Future work should consider more explicitly users’ understanding of the involved technology, as well as consider other types of models, such as physically-informed models of climate futures (Lütjens et al., 2021b,a) or those that integrate across larger, more complex systems (Lavin et al., 2021). Finally, there is the possibility of using this method in diverse settings. We hope that this approach could be applied to other community contexts, whether it be designs for a local community garden or bold new tactics to fight the global climate emergency.

Acknowledgements

We thank Eyal Perry, Micah Epstein, Amy Smith, Simon Colton, Rubez Chong, Océane Boulais, Maggie Hughes, Valdemar Danry and Tobin South for helpful feedback, comments and support. We thank all participants for their time.

Author Contributions

ZE, HS, and DN conceptualized the project. HS conducted the user interviews. ZE and HS wrote the paper, with input from DN.

References

- Berns, S., and Colton, S. 2020. Bridging generative deep learning and computational creativity. In *ICCC*, 406–409.
- Colton, S.; Smith, A.; Berns, S.; Murdock, R.; and Cook, M. 2021. Generative search engines: Initial experiments. In *ICCC*.
- Compton, K., and Mateas, M. 2015. Casual creators. In *ICCC*, 228–235.
- Crawford, K., and Paglen, T. 2019. Excavating ai: The politics of images in machine learning training sets. *AI and Society*.
- Dunne, A., and Raby, F. 2013. *Speculative everything: design, fiction, and social dreaming*. MIT press.
- Epstein, Z.; Boulais, O.; Gordon, S.; and Groh, M. 2020a. Interpolating gans to scaffold autotelic creativity. *International Conference on Computational Creativity Causal Creator Workshop*.
- Epstein, Z.; Levine, S.; Rand, D. G.; and Rahwan, I. 2020b. Who gets credit for ai-generated art? *Iscience* 23(9):101515.
- Ganimals Blog. 2020. Beware the training data: The baracuda effect. ganimals.media.mit.edu.
- Kimmerer, R. 2013. *Braiding sweetgrass: Indigenous wisdom, scientific knowledge and the teachings of plants*. Milkweed editions.
- Lavin, A.; Gilligan-Lee, C. M.; Visnjic, A.; Ganju, S.; Newman, D.; Ganguly, S.; Lange, D.; Baydin, A. G.; Sharma, A.; Gibson, A.; et al. 2021. Technology readiness levels for machine learning systems. *arXiv preprint arXiv:2101.03989*.
- Lütjens, B.; Crawford, C. H.; Veillette, M.; and Newman, D. 2021a. Pce-pinns: Physics-informed neural networks for uncertainty propagation in ocean modeling. *arXiv preprint arXiv:2105.02939*.
- Lütjens, B.; Leshchinskiy, B.; Requena-Mesa, C.; Chishtie, F.; Díaz-Rodríguez, N.; Boulais, O.; Sankaranarayanan, A.; Piña, A.; Gal, Y.; Raïssi, C.; et al. 2021b. Physically-consistent generative adversarial networks for coastal flood visualization. *arXiv preprint arXiv:2104.04785*.
- Morrison, M. I. 2017. *Decolonizing Utopia: Indigenous Knowledge and Dystopian Speculative Fiction*. University of California, Riverside.
- Pérez y Pérez, R., and Ackerman, M. 2020. Towards a methodology for field work in computational creativity. *New Generation Computing* 38(4):713–737.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, 8748–8763. PMLR.
- Ramesh, A.; Dhariwal, P.; Nichol, A.; Chu, C.; and Chen, M. 2022. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*.
- Saldaña, J. 2013. *The coding manual for qualitative researchers*. SAGE.
- Sepie, A. J. 2017. More than stories, more than myths: Animal/human/nature (s) in traditional ecological worldviews. *Humanities* 6(4):78.