Endless Forms Most Similar: The Dearth of the Author in AI-Supported Art

MAX KREMINSKI, Santa Clara University, USA and Midjourney, USA

We characterize and discuss the "dearth of the author": a problem that emerges when AI-based creativity support tools (CSTs) allow their users to produce highly detailed output artifacts—such as lengthy written stories, or high-resolution pieces of visual art—from very small amounts of input (e.g., brief textual prompts). When small amounts of user input are extrapolated by a CST into a highly detailed output, the CST itself has to make many creative decisions that would otherwise fall to the user, disrupting the usual close relationship between input intentionality and output detail in the creative process and yielding output artifacts that contain much more of the tool's than the user's creative signature. Though users of such tools are often able to inject more intention into the output through the adoption of more sophisticated creative practices, these high-intentionality practices tend to be more effortful than the default modes of use of AI-based CSTs, and the use of such practices is not always immediately apparent in the output artifacts. We analyze the dearth of the author from an information-theoretic perspective, focusing on the informational ratio between intent and meaning in creative artifacts; examine implications of the dearth of the author, including homogenization of creative outputs and rhetorical confusion around the "soulfulness" or "soullessness" of art produced with AI assistance; and suggest strategies for mitigating the dearth of the author in future AI-based CSTs.

CCS Concepts: • **Applied computing** → **Arts and humanities**.

Additional Key Words and Phrases: creativity support tools, computational creativity, human-computer interaction, generative AI

1 INTRODUCTION

Research in the field of *creativity support tools* (CSTs) [Chung et al. 2021; Frich et al. 2019; Shneiderman 2007] aims to discover how computation can best be leveraged to support human creativity. Recent progress in generative artificial intelligence has led to the rapid introduction of a new wave of AI-based CSTs, including both traditional user-centered tools and mixed-initiative co-creative systems [Deterding et al. 2017; Liapis et al. 2016] that proactively intervene in the creative process. Existing AI-based CSTs target creative domains ranging from visual art [Chang et al. 2023] and writing [Lee et al. 2024] to music [Louie et al. 2022] and videogame design [Anjum et al. 2024], and tools targeting additional creative tasks and domains are now being introduced on a regular basis. Many of these tools are based on large pretrained generative models, particularly transformer language models [Vaswani et al. 2017] and diffusion image models [Ho et al. 2020]; the most widely deployed user interfaces for these models ask the user to supply a short textual prompt, which is then used to guide the generation of one or more new artifacts.

Although AI-based CSTs[1], especially those driven by prompting, have rapidly become very popular, some researchers and tool designers have begun to question some aspects of how these tools aim to support creativity. In particular, many researchers have expressed concerns [Buschek et al. 2021; Epstein et al. 2023; Hancock et al. 2020; Levent and Shroff 2023] about potential *homogenization effects* [Anderson et al. 2024] that these tools may exert on the creative output of their users, and a range of studies have found some degree of evidence for increased homogeneity of human creative output stemming from the use of AI-based CSTs in the creative process [Arnold et al. 2020; Doshi and Hauser 2023; Jakesch et al. 2023; Padmakumar and He 2023]. Given that

[1]Note that we do not aim to characterize generative AI models as CSTs in a universal sense, only when the outputs of these models are incorporated into human practices—such as art and design—that are traditionally deemed "creative". This characterization is broadly in line with the "normative ground" established by popular interfaces to these models [Li et al. 2023], which are often marketed as supporting user creativity: see, e.g., OpenAI's characterization of DALL·E 2 as "empower[ing] people to express themselves creatively" [OpenAI 2022].

Fig. 1. An illustration of how great an effect a more detailed prompt can have on the informational content of a large pretrained generative model's output. On the left are four example outputs from a diffusion image model prompted with the phrase "cat pirate", while on the right are three alternative, more precisely specified takes on the cat pirate concept. The underspecified cat pirates on the left reflect almost exclusively the generator's learned stereotypes of "a typical cat" and "a typical pirate", and little information can be gleaned from these images besides the concept of "a cat who is a pirate". The slightly more constrained cat pirates on the right are much more obviously evocative of distinctive settings and stories.

definitions of creativity often include *originality* [Torrance 1966] or *novelty* [Boden 2004] as one of creativity's key facets, it has become important to understand the mechanism by which AI-based CSTs can sometimes apparently undermine the originality of their users—and how these threats to originality can be mitigated or reversed.

In this article, we characterize and discuss the effects of AI-based CSTs on originality via an examination of the "dearth of the author": a problem that emerges when CSTs allow their users to produce highly detailed output artifacts—such as lengthy written stories, or high-resolution pieces of visual art—from very small amounts of input (e.g., brief textual prompts) [Kreminski 2024b]. When small amounts of user input are extrapolated by a CST into a highly detailed output, the CST itself has to make many creative decisions that would otherwise fall to the user, disrupting the usual close relationship between input intentionality and output detail in the creative process and yielding output artifacts that contain much more of the tool's than the user's creative signature. This phenomenon, we argue, is best understood via an information-theoretic perspective, which can also be used to relate the dearth of the author to well-known difficulties with classical generative methods. The information-theoretic perspective further allows us to suggest several distinct approaches to the mitigation or reversal of the dearth—one of which in particular ("lensing the imagination") stands out as especially promising for its simultaneous ability to increase the output diversity of AI-supported creative processes while also preserving or increasing user control.

It should be noted that scholarship in the humanities over the last several decades has done much to complicate and broaden traditional views of authorship. In particular, Barthes's original coinage of the "*death* of the author" [Barthes 1968] was meant to frame a rhetorical move away from the centering of authorial intent in the interpretation of texts. This decentering of authorship has aligned well with post-humanist attempts to characterize creativity as emergent or ecological [Harris and Holman Jones 2022], as well as practices in generative art [Galanter 2016; Ugander and Epstein 2024], computational fabrication [Albaugh et al. 2020; Twigg-Smith et al. 2021], and computational poetry [Kreminski 2024a] that lean into the aleatoric or beneficially unpredictable nature of creative machines. Nevertheless, the expressive or communicative dimension of creative work remains important to at least some creators and at least some audiences—and for these people, the dearth of the author represents a weakness of the typical outputs of many modern human-operated generative AI systems. The nature of this weakness is what we attempt to explain.

2 THE DEARTH OF THE AUTHOR

What's the point of creating art? Sometimes we create to learn, or create to think, or create for solitary enjoyment—but when the creative process results in an artifact that's meant to be shared with and interpreted by others, what is this artifact *for*? In this article, we take the view that in most cases, an artistic artifact is meant primarily to convey its author's *expressive intent*: to communicate information, invoke feelings, or otherwise express to the audience a set of ideas and sentiments that the creator intends to share. Indeed, this assumption forms the basis of the *expressive communication* framework [Louie et al. 2022] that has been used to evaluate AI-based CSTs in the past.

Although the word "intent" is sometimes used narrowly to mean an explicitly articulable intent that is fully known to the artist, originates fully within the artist, and fully precedes the creation of a particular artistic work, we use the term more broadly here: namely, to refer to the sometimes-ambiguous and gradually evolving cloud of mental and emotional factors that lead an artist to create a particular artifact and accept it as worthy of presentation to an audience. Among other things, this cloud of factors includes the vague and inarticulable early-stage creative impulse that Arieti refers to as the "endocept" [Arieti 1976]. In adopting this broad view of intent, we attempt to establish a significant point of commonality between several different fields of creative endeavor currently addressed by generative models (particularly creative writing and visual art, the two most directly impacted fields at the time of this writing) while also acknowledging that the human factors leading to expressive creation can be characterized very differently by different artists and theorists. Even a broad view of intent cannot cover every possible perspective on artistic motivation, and there is of course no guarantee that intent is successfully carried through to the audience—some have famously gone so far as to argue that authorial intent is ultimately irrelevant to the audience's artistic experience—but we nevertheless hold that an understanding of the communicative dimension of art is necessary for making sense of tool impacts on the socially embedded facets of the creative process.

The creative process can thus be viewed as a process of *decision-making* by the author. Consider writing: authors of fiction must decide on character names and personalities and appearances, on details of setting and background, on the balance between action and introspection, on moment-to-moment emotional tone; authors of argumentative essays must decide on the points they want to argue, on the evidence they want to use in support of each point, on the order in which to introduce this evidence; authors in general must decide on what words they want to use in what places to get their ideas across. A piece of writing that contains fifteen hundred words can be thought of as the result of *at least* fifteen hundred decisions: perhaps a bit fewer when some of these words are "forced choices" that could not be exchanged for another due to rules of grammar, but usually many more due to the need for decisions *beyond* simple word choices in composing any useful or interesting piece. In fact, the very choice of writing medium, form, or language can be taken as reflective of expressive intent, even as these choices condition what can be written and how. Much as musicians choose instruments in part for their characteristic "voice" [Tanaka 2006], writers have been known to work with typewriters over computers to avoid easy revision of their output [McDonell and Brinkley 2000] or write in a new language to distance themselves from expectations [Shapiro 2016] as a means of deliberately conditioning their own writing.

Historically, an author could not compose a substantive piece of writing without making the vast majority of these decisions directly.[2] Thus the whole of every piece of writing could be taken as reflecting its author's intent: the relationship between the length of a written piece and the number of decisions that its author had to make was relatively fixed, and an author could not lengthen a piece of writing without deciding what specific words to add.

But we suddenly find that this is no longer the case. In particular, when authors use AI-based CSTs to expand a small amount of input text (such as a one-sentence instruction) to a large amount of output text (such as a complete written story or essay), they *delegate* many of the creative decisions involved in producing the larger output to the CST—resulting in a piece of writing with an unusually low ratio of human decision-making to output length. In other words, the expressive intent of the author is underspecified relative to the amount of text that is generated, and the resulting piece of writing is unusually sparse in terms of expressive intent per word.

We refer to this unusual situation as the **dearth of the author**: the naïvely AI-augmented author is not absent or dead, but their intent is stretched so thinly over their writing that they may feel barely present. In lieu of authorial intent, creative decisions are made by the CST to which the author has delegated portions of the writing process; simple LLM-based CSTs like ChatGPT make these decisions by approximating highly probable choices that a certain set of raters might also score well [Ouyang et al. 2022], while other CSTs lean on knowledge encoded in handcrafted rulesets [Kreminski et al. 2022a], in large human-constructed databases of commonsense knowledge [Gero and Chilton 2019], or in more specialized corpuses of text [Mirowski et al. 2023]. Across the full range of AI-based CSTs for writing [Lee et al. 2024], the greater the discrepancy between the size of a minimum viable user input and the size of the output piece of writing, the more the naïve user's results tend to exhibit the dearth.

The same analysis applies just as well to other creative domains—including visual art, music, and a growing number of others—in which AI-based CSTs can be used to produce large outputs from comparably small user inputs. In visual art, for instance, the assumption of a one-to-one relationship between marks made on the canvas and authorial decisions—already disrupted in the 19th century and onward by the introduction of photography as an alternative to manual mark-making techniques of image production [Sontag 1977]—has been further unsettled by the widespread use of generative models to summon "paintings" that appear to consist of thousands of brushstrokes by means of prompts containing only a handful of words. Moreover, although AI-based CSTs based on large pretrained generative models have only emerged within the last few years, the perceived homogeneity of computer-generated expressive artifacts substantially predates the introduction of these models, suggesting that the dearth of the author may originate from something more fundamental than the particularities of recently introduced technologies. To fully characterize the dearth of the author, we argue that it is necessary to take an information-theoretic view of the computational generation of expressive artifacts in general.

[2] Though some words might be drawn verbatim from quotations, and in rare cases some might be chosen via mechanistic, aleatoric, or automatic writing processes, such as those adopted by the Oulipo movement [Montfort and Wardrip-Fruin 2003].

3 INFORMATION AND OATMEAL

The joys and pains of computationally supported creative expression have been discovered and rediscovered over the last several decades by many different communities of research and practice. One of these communities—the *procedural content generation* (PCG) research community [Shaker et al. 2016]—has a well-established name for the problem of perceived homogeneity in computer-generated artifacts: the "10,000 bowls of oatmeal" problem [Compton 2016], often referred to simply as the *oatmeal problem*. Suppose, the metaphor goes, we've created a computational generator that outputs bowls of oatmeal. By minutely varying the position and orientation of individual oats, we can trivially create an effectively infinite number of *mathematically* distinct bowls of oatmeal. But to a human observer, these bowls will not appear to be meaningfully, *perceptually* distinct: the human brain can quickly see through a vast number of minor, inconsequential variations to recognize and dismiss the generator's entire *expressive range* [Smith and Whitehead 2010] (the sum total of its possible outputs, both realized and unrealized) as "just a lot of oatmeal".

A central challenge for PCG research, then, is to somehow forestall, delay, or work around the *perceptual collapse* [Kreminski 2023a] that stems from the audience's recognition of a generator's fundamentally unchanging output patterns. In light of this challenge, PCG researchers and practitioners have proposed a wide variety of experience design solutions to the limited perceptual novelty of generated artifacts. Some researchers have called for a shift toward "orchard" or "forest" presentations of generator outputs; in the orchard mode, the goal of generation is not to generate individually distinctive-feeling artifacts, but to generate sets of artifacts that somehow aesthetically resonate with one another, while in the forest mode, the goal is to generate a gestalt background against which other sources of perceptual novelty (such as hand-crafted "landmarks") can be juxtaposed [Karth 2019]. Others have suggested presenting each audience member with only a single or very small number of generated artifacts, instead of allowing them to continue requesting more artifacts forever [Kreminski et al. 2018]. Still others have created experiences in which audience members interactively search through a space of possible generated artifacts in a *casually creative* manner [Compton and Mateas 2015; Petrovskaya et al. 2020], for instance by moving around within a rule-based generator's parameter space [Colton et al. 2018; Kreminski et al. 2020] or a data-driven generator's latent space [Epstein et al. 2020] to locate artifacts of interest and perhaps curate these artifacts for presentation to others somewhere down the line.

All of these strategies operate by changing the presentation—what computational creativity researchers would call the *framing* [Charnley et al. 2012; Cook et al. 2019]—of generated artifacts. But what interventions might be possible at the generator level to increase the distinctiveness of outputs? Suppose we attempt to expand the maximum possible output complexity $K^*(G)$ of a generator G, allowing the generator to produce more complex artifacts. Rabii and Cook [Rabii and Cook 2023] show that there exists a mathematical relationship between the size of the generator's possibility space, $log_2(\#\pi(G))$; the length of the generator's source code, $|G|$; and the Kolmogorov complexity of the most complex artifact that G is able to produce, $K^*(G)$. Kolmogorov complexity refers to the length of the shortest program that is capable of producing a particular output—with the intuition being that informationally simple outputs can be described very compactly, whereas informationally complex outputs must be described in correspondingly more complex ways to capture their intricacies in full. Especially complex outputs are often most compactly modeled by a program that simply states a verbatim string to print: a trivial, inflexible "generator" that can only produce one artifact in particular. Meanwhile, a compact generator that can produce a large number of possible outputs through mere random combination of values may nevertheless lack sufficient potential output complexity to yield artifacts that humans find perceptually interesting, especially after they have already begun to detect the similarities between a number of other outputs drawn from the same distribution.

This relationship helps to explain the origins of oatmeal, or perceptually similar artifacts in a generator's possibility space. A linear increase to $K^*(G)$ that is not accompanied by any increase to $|G|$ results in an exponential expansion of $\pi(G)$, essentially reducing the density of highly distinctive artifacts within the generator's possibility space—such that the majority of artifacts produced by G end up bearing substantial similarities to one another. Increasing $|G|$ (i.e., the size of the generator itself) tends to be expensive, because it involves the encoding of additional knowledge into the generator; but a generator whose maximum output complexity is scaled up *without* additional knowledge encoding taking place is doomed to decreased perceptual output diversity. This observation has been referred to as the "limited free breakfast" theorem [Kreminski 2023b]: "You can have free breakfast if you want, but it's going to be oatmeal."

4 THE DEARTH OF THE AUTHOR IS A DEARTH OF INFORMATION

Information-theoretically, the dearth of the author stems from essentially the same place as the oatmeal problem. Although Rabii and Cook focus predominantly on hand-programmed generators in their characterization of the oatmeal problem, the information-theoretic structure of their argument permits generalization to any generative system which functions by applying a finite range of possible transformations to a finite space of possible inputs—regardless of whether the logic by which these transformations are applied was hand-programmed (as in traditional procedural content generation) or learned from training data (as in procedural content generation via machine learning [Summerville et al. 2018]).

To characterize a modern pretrained generative model that accepts a text prompt as input and yields an output artifact, the formal definition of a generator G in Rabii and Cook's treatment of the oatmeal problem can be broken into two parts, both representable as bitstrings:

(1) The model m: a large pretrained generative model. This represents, essentially, a compressed corpus of patterns discovered in the sum of the model's training data.
(2) The input i: a user's prompt. This represents, essentially, the only variable portion of the logic that is used to sample from the model.

Under this bipartite division of the generator, the relationship described by Rabii and Cook between total generator size, possibility space size, and complexity of the generator's most intricate possible output artifact can again be found to obtain.[3] If the model m is treated as a black box that cannot be extended at sampling time (as is generally the case in present-day usage scenarios), the amount of information contained in the prompt i directly maps to the Kolmogorov complexity of the most complex artifact that the whole generative system G can produce. To increase $K^*(G)$, one must either expand the prompt, expand the model, or expand the logic used to post-process the model's output—in the latter case redefining the generative system more broadly to include not just the internal generative model but also any other computation performed after the model yields a result.

As a result, when multiple different users of the same large pretrained model supply the model with similar short prompts, the outputs that they receive are limited in how much they can possibly differ from one another. Only the information taken from the prompt itself and the information given by the random seed (when some degree of stochasticity is used in sampling) are permitted to

[3] Stochasticity in the sampling process—as is typically present in both transformer language models (modulated by the temperature parameter) and in text-to-image diffusion models—can be characterized by the addition of a few more bits of information to the input, representing the seed used to initialize random number generation. Because the amount of information drawn from this seed is comparable from one generation to the next with the same model and sampling procedure, this has negligible net influence on the overall analysis. Stochasticity driven by a seed is characterized the same way in Rabii and Cook's analysis.

vary from one run of the model to the next; any other information contained in the model output must be drawn from the fixed information contained in the model itself. Thus, the greater the ratio of output size to input size, the greater the expected perceptual similarity between different model outputs. When a short prompt like "write me a story about a cat pirate" is expanded to a story containing hundreds or thousands of words, the vast majority of words in the output story must be conditioned primarily on the information contained in the model (which is fixed and unchanging), rather than on the information contained in the input—which is obviously insufficient to precisely specify a *meaningful* cat pirate story, relative to the model's typical output, from an expressive perspective (Fig. 1). The end result is perceptual homogeneity [Begus 2023]: stories produced by the same language model from short prompts share similar structures, characters, tropes, and so on, because short inputs do not provide the model with sufficient extra conditioning information to yield creative decisions outside the defaults defined by model training data.

5 IMPLICATIONS OF THE DEARTH

Despite its simplicity, the dearth of the author helps to explain many of the disparate anxieties and difficulties around AI-based CSTs. We briefly discuss a few of its implications here.

Homogenization of creative output. Several recent studies of AI-based CSTs have found either direct [Anderson et al. 2024; Arnold et al. 2020; Doshi and Hauser 2023; Padmakumar and He 2023] or indirect [Bhat et al. 2023; Jakesch et al. 2023] evidence that these tools can exert a *homogenizing* effect on the creative outputs produced by different users: in other words, different users of the same CST may produce more similar outputs than they would without the CST. Homogenization effects can be explained by authors' delegation of creative decision-making to a tool that makes similar creative decisions in similar usage scenarios, with greater degrees of homogenization resulting from tools that make a greater proportion of creative decisions directly.

Limited feelings of ownership. Recent studies have also found that users of AI-based CSTs for writing tend to experience a limited sense of ownership of or responsibility for the outputs of their interaction with the CST [Draxler et al. 2024; Lee et al. 2022]. This is similarly explained by the delegation of creative decisions to the CST: delegating a greater proportion of creative decisions to the CST seems likely to result in a commensurately lower feeling of ownership toward the resulting text, as seen both in Draxler et al. [Draxler et al. 2024] and in the stronger sense of ownership reported by users of a narrower CST [Kreminski et al. 2022a].

Rhetorical confusion. The typical outputs of AI-supported creative processes are often referred to as "soulless". Simultaneously, users who build highly intentful creative processes around AI-based CSTs—sometimes considering and discarding many dozens of AI outputs before accepting one as complete [Chang et al. 2023] (Fig. 2)—are perplexed by these assertions. The dearth of the author helps explain both phenomena. The term "soullessness" reflects the sparseness of intent in the outputs that are easiest to create with many AI-based tools, and that therefore dominate most non-enthusiasts' impressions; meanwhile, the high-intent nature of some AI artists' processes may not be immediately apparent to onlookers, because AI tools permit the creation of comparably sized outputs from much smaller specifications of intent.

Greater impacts on inexperienced creators. Experienced writers (for instance) tend to have an ear for evocative language and a strong aversion to cliché, both of which are components of a sophisticated sense of *taste* built up over many years of paying close attention to language. These writers are readily able to identify problems in AI-generated text [Chakrabarty et al. 2023], and are often unwilling to simply accept creative decisions made by the machine when these decisions conflict with their own sensibilities [Mirowski et al. 2023]. Novice writers, on the other hand, may tend to treat the machine as a creative authority [Anderson et al. 2024] and delegate a greater

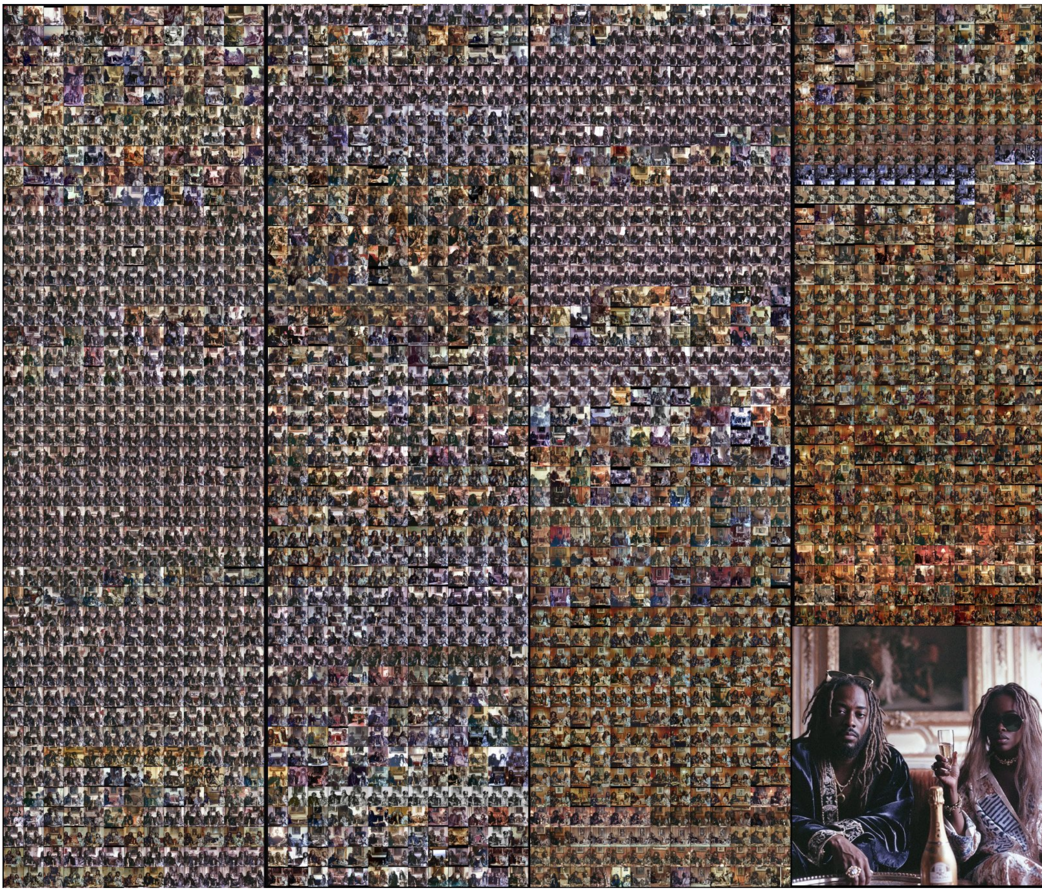

Fig. 2. An image quilt of the many generations and regenerations involved in producing a single final output image (shown enlarged at the bottom right) via an AI-based CST built around a diffusion image model. A large amount of user-controlled information has been brought into the creative process over the course of many rerolls and adjustments, but each image produced along the way is similar in *scale* to the final output, making it difficult to determine from output size alone how much more authorial intent went into the final image than into the first image generated. A better suite of tools for the specification of expressive intent could enable the generation of a similar final output image with fewer total regenerations, illustrating the importance of tool design to preserving detailed authorial intent in AI-based CSTs. [St. Pierre 2024]

proportion of their creative decisions—a form of *algorithmic loafing* [Inuwa-Dutse et al. 2023] that is likely to result in a stronger sense of authorial absence.

6 RESPONSES TO THE DEARTH

Fundamentally, the dearth of the author follows naturally from the fact that current generative models will tend to produce similar outputs in response to similar inputs, even when those inputs come from different users. Interventions in response to the dearth of the author, then, can come in three broad flavors: changing the generative models themselves; changing how we use their outputs; and changing the inputs that we feed into them.

6.1 Changing the Models

6.1.1 ...by training new models. Helena Sarin, a notable and frequently cited AI artist [Browne 2022; Hertzmann 2019], has been training generative models (particularly generative adversarial networks) on her own photography and hand-drawn artwork—explicitly as a means of avoiding the homogenous appearance of outputs from pretrained models—for nearly as long as modern generative AI architectures have existed [Bailey and Sarin 2018]. Given present tools, this approach to maintaining originality is relatively demanding: most artists do not possess the technical skills needed to train models from scratch on their own data, and even the assembly of a sufficiently large and diverse corpus of data to be useful in training a custom model may be both time-consuming and technically difficult, since even determining what data might make for a "good" model tends to require some degree of understanding of machine learning principles. Nevertheless, this approach can be quite powerful if pursued: an artist who follows it can essentially redefine their own creative process to include the process of training a characterful generative model, allowing expressive intent to re-enter the process "upstream" of the model usage phase.

As suggested in one discussion of Sarin's work [Bailey and Sarin 2018], another means by which artists can avoid tool-imposed homogeneity is to make use of relatively unknown generative models that impart a different grain onto their outputs, rather than adopting a model that is already widely used. However, this is easier said than done—in part because new models that are sufficiently powerful and general to produce a range of coherent outputs are not created often, and in part because the new architectures that power some novel models are often initially unapproachable to artists. It is not until new models are wrapped up in widely usable products that most artists tend to be able to access them, and as soon as a "fresh" model is productized, it is likely to be adopted by many users at once, resulting in a rapid exhaustion of its freshness.

Can the process of model training itself be made more approachable? Wekinator [Fiebrink et al. 2009] represents an early approach in this direction: by giving musicians tools that they can use to train machine learning models on bespoke input data without deep ML-specific technical skills, the Wekinator framework acts as a "meta-instrument" for the construction of new musical instruments that are driven by a very wide range of different inputs. However, a Wekinator-like toolkit to support the training of new *generative* models has not yet appeared as far as we are aware. In addition, it may not be possible to create such a toolkit for the types of large pretrained generative models that have received the most attention in recent years: the transformer and diffusion architectures on which these models are based require many orders of magnitude more data and computation to train than older architectures, placing the training of such models "from scratch" out of reach for the vast majority of potential users.

One alternative to from-scratch training involves the adaptation of large pretrained models to a more specific task using a much smaller amount of data and computation, using parameter-efficient fine-tuning methods [Mangrulkar et al. 2022] such as low-rank adaptation [Hu et al. 2021]. However, the outputs of fine-tuned models are still shaped substantially by the unadapted informational content of the original pretrained model on which the fine-tuned model is based, and the process of fine-tuning itself requires a greater degree of technical know-how and infrastructure than direct use of an approachable CST built around a fully pretrained model.

Broadly speaking, introducing new generative models and incorporating them into creative workflows can help to address the dearth of the author, especially if artists themselves are directly equipped to train new models on their own specialized datasets via approachable tools. However, the broad adoption of this approach faces many difficulties. Neural networks trained on similar datasets tend to converge [Li et al. 2015], so developing AI-based CSTs around a wider range of large pretrained generative models may not yield much improvement in outcome diversity if the underlying models are all trained on similar data. The increasingly widespread use of synthetic data to train new models results in additional convergence [Shumailov et al. 2023], further limiting the utility of simply switching which pretrained model is used as a foundation for CST design. Furthermore, training new models from scratch, especially using the currently most popular architectures, is substantially costly, requiring large amounts of infrastructure and specialized knowledge. For all of these reasons, the introduction of new models may remain of limited utility in mitigating the dearth of the author for the foreseeable future.

6.1.2 ...by sampling or fine-tuning for increased output diversity in general. If the creation of new models from scratch is generally too costly or difficult to pursue as a strategy for improving AI-based CSTs, can we simply adjust how we sample from existing pretrained generative models to improve the diversity of their outputs? A number of approaches in this family have shown promise, including diverse decoding [Ippolito et al. 2019; See et al. 2019], quality-diversity approaches [Bradley et al. 2023], and even finetuning existing models to force the generation of diffuse distributions [Zhang et al. 2024] in contexts where a greater degree of output unpredictability is desired. A finetuning approach in particular may essentially allow for the restoration of output diversity that seems to be lost in the process of finetuning a large language model to respond appropriately to user instructions [Padmakumar and He 2023].

Information-theoretically, however, these approaches seem to share a common limitation: they add diversity to model outputs by drawing on *random* components of model input (i.e., the seed used for random number generation) rather than on *user-supplied* components of model input (i.e., the prompt), and therefore may reduce user control over creative outcomes. Though diversity derived from randomness may prove useful for getting users out of creative ruts and providing inspiration, especially early in the creative process, it has yet to be evaluated whether randomness-derived approaches to improving model output diversity are beneficial to human creativity overall.

6.2 Changing the Use of Model Output

6.2.1 ...at the user level. When designing an AI-based CST that wraps a generative model with a user interface, it is easy to fall into the trap of assuming that the model's output must constitute the *final* output of the creative process: i.e., that model outputs may be *curated* by tool users, but not further modified or acted on by users before they are passed along to a user's audience directly. However, different AI-based CSTs for writing (for instance) vary widely in the nature of their output and their ability to make a large proportion of the decisions involved in writing on the user's behalf. When the outputs of a CST take the form of "sparks" [Gero et al. 2022], plot outlines [Kreminski et al. 2022a], questions [Kreminski and Chung 2024], or feedback [Kim et al. 2024; Stark et al. 2023] rather than output-ready prose, the user cannot as easily delegate creative decisions about the integration of these elements into a complete piece of writing—leaving open a space of *underdetermination* [Albaugh et al. 2020], or perhaps even *creative struggle* [Zhou and Sterman 2023], into which expressive intent must flow.

It has recently been suggested that the willingness of users to accept model outputs as authoritative or final is influenced both by user perception of AI systems as giving "correct answers" and by the extent to which model outputs resemble the target *form* of final output in a given creative context [Anderson et al. 2024]. Both of these factors suggest that, even when AI models are used to generate output in the target form of the overall creative process, it may be useful for CSTs wrapping these models to deliberately supply users with model outputs that contain obvious holes, errors, or points of incompleteness. These minor issues may increase the *inferential distance* between model outputs and pieces of a finished creative artifact, thereby adding friction to the direct or verbatim use of model outputs and leading users to interleave some of their own otherwise unexpressed creative intent into material output by the model during the process of repair.

6.2.2 ...at the tool level. Beyond requesting and framing model outputs for *users* to interpret or use as something other than final components of creative output, model outputs can also be used to trigger additional algorithmic processes, which are then used to supplement model outputs or adapt them into a different form prior to the output being displayed to tool users. For instance, a language model that has been trained to emit a control script for some sort of "agent" apparatus—one that permits the querying of external information sources or the execution of model-generated code—may be able to supplement its own outputs with information from these external sources, perhaps using a web API to supplement its own text suggestions with appropriately constrained random information when needed or a CST-provided API function to apply a particular filter to a generated image before displaying the image to the user. These tool-level adaptations of model output essentially extend the full generative pipeline of the CST beyond the model itself, and use the extra pipeline stages to inject additional information into the creative process; some of these adaptations also bleed into alterations of model inputs on subsequent invocations of the model, which are discussed further in the section below.

6.3 Changing Model Inputs

6.3.1 ...by injecting chaos from another source. If the outputs of an AI-based CST are intended to provide the tool's user with inspiration or raw material, it may be valuable for these outputs to be more diverse than model output conditioned on the user's input alone. Especially early on in the creative process when the user has likely not provided very much information that can be used to steer the model's generation, improved model output diversity may be most readily achievable through the supplementation of the user's input with information from an outside random or chaotic source: for instance by randomly selecting several keywords from a large vocabulary and supplying these to the model alongside the user's prompt.

As with boosting model output diversity via diverse decoding and related randomness-driven techniques that are applied *inside* the model's black box (Section 6.1.2), there are many potential pitfalls associated with chaos injection at the input level. Most importantly, the supplementation of user-controlled input with random input (i.e., input that is not under the user's control) changes the ratio of user-supplied to non-user-supplied information flowing through the whole creative process, decreasing the extent to which the user is in control of model outputs. If outside information is injected invisibly, the user may struggle to reproduce desirable results or avoid undesirable results on future runs of the model, because they are not directly privy to the input adjustments that are leading to these results; making the injection of outside information user-visible, however, will likely lead to increased user interface complexity and a steeper learning curve.

Additionally, if a traditional generative method (such as a generative grammar [Compton et al. 2015], planning domain [Cardona-Rivera et al. 2024], or answer set program [Smith and Mateas 2011]) is used to inject more structured random output into the model—either to give the user greater control over the random information injected, to ensure that random information is combined coherently, or both—all the difficulties associated with the selected generative method are now also brought into play at the tool design level. In particular, these methods tend to require substantial authoring effort [Jones 2023] and lack the easy adaptability to arbitrary creative domains that characterizes more recently introduced generative methods [Yao et al. 2019]. These difficulties have so far hindered the adoption of neurosymbolic approaches to user-controllable output diversification of large pretrained generative models.

6.3.2 ...by exploring possibility space. From the perspective of a user of a widely adopted AI-based CST, one approach to differentiating one's own output from the output achieved by other users of the same tool is to reframe the creative process as one of exploring the tool's input and output space in search of distinguishable and otherwise undiscovered oddities. This *possibility space exploration* approach is similar to the framing of creativity as curation adopted by many casual creator tools [Compton and Mateas 2015; Petrovskaya et al. 2020], particularly those that explicitly provide users with a means of rapidly adjusting parameter values in an explicitly defined parameter space [Colton et al. 2018; Kreminski et al. 2020] or a learned latent space [Epstein et al. 2020]. Searching for unusual (but short) inputs that led to unusual outputs represented an especially common approach for early "prompt artist" users of text-to-image models, who sometimes framed this process in terms of discovering and exploiting model glitches [Caramiaux and Fdili Alaoui 2022; Chang et al. 2023]; some of these artists tended to view the textual prompt as part of the output artifact (along with the image or images generated in response to this prompt) [Chang et al. 2023], while others hoarded unusual prompts to keep others from discovering and exploiting the remaining relatively unexplored corners of latent space.

Creative differentiation achieved through a possibility space exploration approach alone is often fragile. Changes to a particular AI-based CST's underlying generative model can wipe out glitches and artifacts that once led to reliably unusual outputs, while widespread discovery of a simple prompting trick can lead to outputs from the corresponding corner of latent space rapidly becoming much more prevalent (sometimes even to the point of becoming "overplayed"). In addition, an artistic practice that relies heavily on the discovery of new material characteristics of a particular generative model may not allow for the expression of a wide variety of different sentiments, ideas, or other aspects of communicative intent; the resulting artifacts may more closely resemble conceptual than traditionally expressive art [Karth and Compton 2023]. Ultimately, this approach to making art with generative models attempts to lean into the fact that the model outputs yielded by short prompts are necessarily conditioned primarily on the information contained within the model itself; thus, the results tend to be "about the model" more than they are about the user's creative goals independent of the model.

6.3.3 ...by inferring user intent automatically. Getting users of AI-based CSTs to specify their expressive intent more precisely is difficult, for a wide variety of reasons that we discuss further in the following section (Section 7). However, informationally diluting user-controlled inputs with non-user-controlled inputs (as proposed in several sections above) tends to reduce the extent to which model output is reflective of user intent, thereby reinforcing the dearth of the author. Can we instead make use of user modeling or personalization techniques to automatically infer a user's expressive intent from their actions, then append this inferred intent to model inputs as an additional facet of (implicitly) user-controlled input?

Although the automatic inference of expressive intent from user actions may be useful for steering model outputs, there is a need to evaluate whether the personalization paradox [Ontanon and Zhu 2021] is generally strong enough to substantially inhibit the application of these techniques in creative contexts. Particular caution should be taken around the potential of creating a self-reinforcing loop of inferred and intended user preference: if information that is inferred about the user's intent is then immediately used to steer the outputs of the generative model, it seems likely that early identifications of user preference may end up increasing user fixation [Alipour et al. 2018; Crilly 2019; Jansson and Smith 1991; Purcell and Gero 1996] on unexpected aspects of inferred intent, whether or not these aspects were actually present in the user's head as part of their expressive intent to begin with. Furthermore, especially early in the creative process, when user intent tends to be weakest and least well-defined, personalization attempts may be especially prone to displacing these initial sparks of intent, potentially replacing them with an inferred intent that a particular CST or user-modeling approach is especially prone to inferring—thereby once again reinforcing the dearth of the author. More research is needed to probe user susceptibility to these potential effects.

6.3.4 ...by lensing the imagination. Given the limitations of other approaches to mitigating the dearth of the author, one major approach remains to be discussed: getting the users of AI-based CSTs to increase the specificity of their input—without influencing their ideational process too strongly toward tool- or model-favored parts of idea-space along the way. This approach, which we refer to as "lensing the imagination", can perhaps best be understood metaphorically: if the ideal source of originality in the creative process is a human artist's explicit expressive intent, and this intent initially tends to arise within users in a vague and poorly defined form, the role that an AI-based CST must fulfill is that of a focusing and clarifying lens. An ideal CST in this framing is one that helps users to magnify and sharpen the initially ill-defined spark of intent without displacing it or snuffing it out.

Let's briefly return to the example of our cat pirates (Fig. 1). If the user of a basic AI-based CST for writing approaches the creative process with only the *explicit* knowledge that they want to tell a story about cat pirates, and they directly instruct the generative model within the CST to fulfill this briefly specified intent, they will immediately be confronted with the same kind of cat pirate story as every other user who makes a similar request: one in which the vast majority of the story's informational content reflects information drawn from the model itself, not from the user. If the user then accepts this story as complete, or substantially modifies their intent toward the details already contained in the first generated story before re-prompting the model for another story, the majority of whatever they really wanted to express will be displaced. But if the user is immediately confronted with a wide variety of different cat breeds they must choose between, or a question about what cat pirates mean to them thematically (violent oppressors of press-ganged mice? cute and cuddly residents of a sheltered tropical island? cunning raiders of heavily defended but resource-rich human installations?), further information may be drawn out of the user about their true artistic or expressive goals. Repetition of similar interactions, if they are carefully structured, may eventually draw enough intent out of the user that the final output artifacts of the creative process reflect predominantly information provided by the user and not the model—resulting in a neutralization, or even a reversal, of the dearth of the author.

7 FROM DEARTH TO ABUNDANCE

The dearth of the author arises when the ratio of authorial intent to output artifact size is small. AI-based CSTs that make creative decisions on the author's behalf tend to decrease this ratio. But AI-based CSTs can also *increase* this ratio by leading the author to make *a greater number of meaningful creative decisions per unit of output produced*, drawing out unexpressed elements of the author's intent and provoking them to refine this intent further. CSTs that do this well enough may even bring about an unexpected alternative condition: an **abundance of the author**, in which every word of a piece of writing (for instance) has been considered more carefully and from more different angles than the author could otherwise manage or afford.

One way to frame the task of lensing the imagination is through the recharacterization of AI-based CST design as designing for *intent elicitation* [Kreminski and Chung 2024]. Kreminski and Chung argue that user intent in creative processes is essentially always co-constructed with the materials of a particular "design situation" [Schön 1983], and that intent is therefore not just *underexpressed* by users early in the creative process (i.e., specified less completely in the user's input to the tool than in the user's head) but also *uncertain* (i.e., largely unspecified in the user's head, because they have not yet had a chance to discover completely what they want). Furthermore, because user actions are often constrained to some extent by fixation, attempts at inferring user intent from actions may often fail to capture important features of the true underlying intent. There are thus three key problems for an AI-based CST to solve: helping users *discover* what they want; helping users *express* what they want to the CST in a way that it can be used to steer generative model output; and helping users *avoid fixation* by showing them alternative ways of looking at the design situation they are currently in.

To facilitate user *expression* of intent, it may be necessary to move past the exclusive use of textual prompts—which have previously proven difficult for people to use effectively, especially when they are relatively inexperienced with the particular generative model they are trying to use [Zamfirescu-Pereira et al. 2023]. Prompting techniques that elicit improved performance on a particular task from one model may nevertheless worsen another model's performance on the same task [Mizrahi et al. 2023], and restricting users to textual prompting alone deprives them of more straightforward interaction techniques for many modes of content; parts of images, for instance, may be much more straightforwardly and precisely indicated via pointing and gesture than via textual description. Domain-specific alternatives to text-only prompting for model steering have been suggested by a number of researchers [Chung et al. 2022; Chung and Kreminski 2024; Lin et al. 2023], and many valuable new interaction techniques likely remain to be discovered in this area. In particular, easier-to-use control modes for specifying aspects of creative intent that are difficult to express through text may represent an especially fruitful focus for future work.

To facilitate user *discovery* of intent, meanwhile, prior work on a class of CSTs described as "reflective creators" [Kreminski and Mateas 2021] may be relevant. Broader design research on the key role of reflection in design processes [Schön 1983] and HCI research on designing systems to promote user reflection in general [Bentvelzen et al. 2022] may also apply. Design patterns for helping users to reflect on what they intend and why are directly applicable to the elicitation of explicitly stated expressive intents from users, and the deliberate introduction of friction into the creative process in order to provoke greater reflection is strongly compatible with recent calls to value the user's creative struggle [Zhou and Sterman 2023] in AI-based CST design.

7.1 Evaluating Progress Toward Abundance

How can we tell if our tools are making progress toward the goal of improving originality by drawing out distinctive user intent? Ever since the field was founded, the evaluation of CSTs has been considered one of the hardest problems in CST research and design [Hewett et al. 2005]—in part due to the diversity of creative domains, tasks, users, and goals that CSTs might aim to support, and in part due to fundamental difficulties with the definition and evaluation of *creativity* as a concept. Moreover, despite the introduction of a standardized survey instrument for CST evaluation [Cherry and Latulipe 2014], relatively little progress has been made toward a field-wide consensus on the best way to evaluate CSTs [Remy et al. 2020]. Nevertheless, we believe that some existing evaluation methods may prove especially well-suited to the task of evaluating tool effects on intentional originality.

To evaluate the raw distinctiveness of different users' tool-influenced creative products, several different forms of *homogenization analysis* [Anderson et al. 2024; Doshi and Hauser 2023; Padmakumar and He 2023] have been proposed. Generally speaking, homogenization analysis involves the production by tool users of many different artifacts, followed by the embedding of these artifacts in some sort of semantic embedding space (e.g., via sentence embeddings [Reimers and Gurevych 2019] for short textual expressions of ideas) and the quantitative comparison of semantic similarity between the resulting embedding vectors. Among other applications, homogenization analysis can be used to draw direct comparisons between two or more different CSTs in the same creative context; to quantify the degree to which homogenization effects arise from individual-level fixation [Alipour et al. 2018; Crilly 2019; Jansson and Smith 1991; Purcell and Gero 1996] as opposed to group-level similarity of AI suggestions; and to evaluate a single CST against multiple different creative contexts in order to determine whether it exhibits homogenization effects consistently or inconsistently across these different contexts.

A closely related family of evaluation approaches, which we term *overlap analysis* [Joshi and Vogel 2024; Padmakumar and He 2023; Roemmele and Gordon 2018], may be useful when the goal is to evaluate the extent to which authors are influenced by CST suggestions in the context of writing. These approaches examine the extent to which sequences of words from a tool's output are reproduced in the final output of a tool-influenced creative process, with higher degrees of overlap likely indicating a greater degree of tool influence on user creative choices. Note that different levels of overlap between intermediate tool output and final user output may be desirable in different contexts, and a middling level of tool influence on user output may arguably be best overall; middling levels of overlap would likely suggest that users consider tool suggestions useful, but are not inclined to let tool suggestions completely displace their own originality.

Because both homogenization analysis and overlap analysis can be conducted automatically to produce quantifiable results, they may be especially useful in conjunction with approaches that aim to visualize and characterize quantifiable dimensions of user movement through design spaces as they continue using a single tool [Alvarez et al. 2022; Kreminski et al. 2022b]. Visualizations of user artifacts in relation to broader *design spaces* defined by a particular generative method have been directly applied to the study of whether AI-based CSTs lead to homogeneity of output in the past [Kreminski et al. 2022b]; combining these successful past approaches with additional quantitative dimensions that aim to capture artifact originality may yield novel visual analytic techniques for CST evaluation.

Beyond directly measurable characteristics of the creative process, the user's self-reported sense of responsibility for creative outputs may function as a qualitative correlate of intent preservation in finished artifacts. One recent study [Joshi and Vogel 2024] finds that the use of longer prompts tends to increase users' feelings of responsibility for their own AI-mediated creative outputs, and hypothesizes that these increased feelings of responsibility may result from increased overlap between phrases written by users in prompts and phrases appearing in model output. Meanwhile, another recent study [Anderson et al. 2024] finds that users feel less responsible for their own creative outputs when using a tool that exerts a stronger homogenization effect on their creative process, again suggesting that subjective sense of responsibility may accurately reflect a CST's displacement of user intent. This possible correlation should be investigated further, but if it holds, it may serve as a relatively easy-to-evaluate means of cross-checking quantitative with qualitative data in attempts to determine whether a particular CST tends to produce a dearth of the author.

Finally, the *expressive communication* [Louie et al. 2022] framework for CST evaluation may be the gold standard for evaluating the preservation of user intent through the creative process. This framework involves first asking CST users to create an artifact that reflects a particular creative intent, then passing this artifact along to another set of human audience members and asking them to characterize the artifact. The extent to which the intent of the tool user is accurately received and reported by the audience members can then be evaluated to determine if a tool is effective at assisting its user in communicating feelings, ideas, or other aesthetic qualities to the audience. This framework may need to be adapted somewhat to account for the fact that an artist's creative intent is generally not fully formed at the beginning of the creative process (and in fact may be very ill-defined at first), but the broad concept of using a second pool of human evaluators to determine whether a tool is successful at supporting expressive communication still seems likely to be especially valuable for CST evaluation.

8 CONCLUSION

We have used an information-theoretic perspective to characterize a major recurring difficulty for AI-based CSTs: the *dearth of the author*, or relative sparseness of authorial intent in creative output, that may arise from their use. This characterization puts research on modern AI-based CSTs into direct conversation with parallel work on classical generative methods. An information-theoretic view of how generative models are able to produce large output artifacts from small user inputs enables us to conclude that the homogenization effects of AI-based CSTs can be addressed by making changes to generative models themselves, how we make use of their outputs, or what we provide them with as input. However, we also find that the only way to decrease homogenization while also preserving or increasing user *control* of creative outcomes involves the expansion of user input specifically—suggesting that the *elicitation of user intent* represents a central challenge for designers of AI-based CSTs.

This approach reflects a fundamental truth about the creative process: namely, that it is a process of making decisions. The ultimate fantasy of artificial intelligence is that it will simply "do what we mean" without further input—but if art is about *deciding* what we mean and why, there can ultimately exist no system, AI or otherwise, that is capable of simply extracting and executing a user's artistic intent without further involvement from the user. Instead, AI-based CSTs must take up the hard problem of helping users think through their intent—which is rarely as precisely defined at the start of the creative process as we initially believe it to be—and then gradually transform that intent into a finished artifact that accurately reflects the meaning the user is hoping to convey.

ACKNOWLEDGMENTS

This work was supported in part by Hackworth Grant GR102981, "Investigating Homogenization of Imagination by Generative AI Models", from the Markkula Center for Applied Ethics. Some of the ideas discussed here were initially formulated in response to discussions held at Stochastic Labs during the summer of 2022; thanks to all who participated in these discussions, and in particular to Joel Simon and Kyle Steinfeld. Thanks also to Isaac Karth, Phoebe J. Wang, and Barrett R. Anderson for conversations that helped shape this work.

REFERENCES

Lea Albaugh, Scott E Hudson, Lining Yao, and Laura Devendorf. 2020. Investigating underdetermination through interactive computational handweaving. In *Conference on Designing Interactive Systems*. 1033–1046.

Leyla Alipour, Mohsen Faizi, Asghar Mohammad Moradi, and Gholamreza Akrami. 2018. A review of design fixation: Research directions and key factors. *International Journal of Design Creativity and Innovation* 6, 1-2 (2018), 22–35.

Alberto Alvarez, Jose Font, and Julian Togelius. 2022. Toward Designer Modeling Through Design Style Clustering. *IEEE Transactions on Games* 14, 4 (2022), 676–686.

Barrett R Anderson, Jash Hemant Shah, and Max Kreminski. 2024. Homogenization Effects of Large Language Models on Human Creative Ideation. In *C&C '24: Proceedings of the 16th Conference on Creativity and Cognition*.

Asad Anjum, Yuting Li, Noelle Law, M Charity, and Julian Togelius. 2024. The Ink Splotch Effect: A Case Study on ChatGPT as a Co-Creative Game Designer. *arXiv preprint arXiv:2403.02454* (2024).

Silvano Arieti. 1976. *Creativity: The Magic Synthesis*. Basic Books.

Kenneth C Arnold, Krysta Chauncey, and Krzysztof Z Gajos. 2020. Predictive text encourages predictable writing. In *Proceedings of the 25th International Conference on Intelligent User Interfaces*. 128–138.

Jason Bailey and Helena Sarin. 2018. Helena Sarin: Why Bigger Isn't Always Better With GANs And AI Art. https://www.artnome.com/news/2018/11/14/helena-sarin-why-bigger-isnt-always-better-with-gans-and-ai-art. Accessed: 2024-04-25.

Roland Barthes. 1968. The death of the author. (1968).

Nina Begus. 2023. Experimental Narratives: A Comparison of Human Crowdsourced Storytelling and AI Storytelling. *arXiv preprint arXiv:2310.12902* (2023).

Marit Bentvelzen, Paweł W Woźniak, Pia SF Herbes, Evropi Stefanidi, and Jasmin Niess. 2022. Revisiting reflection in HCI: Four design resources for technologies that support reflection. *Proceedings of the ACM on Interactive, Mobile, Wearable* *and Ubiquitous Technologies* 6, 1 (2022).

Advait Bhat, Saaket Agashe, Parth Oberoi, Niharika Mohile, Ravi Jangir, and Anirudha Joshi. 2023. Interacting with Next-Phrase Suggestions: How Suggestion Systems Aid and Influence the Cognitive Processes of Writing. In *Proceedings of the 28th International Conference on Intelligent User Interfaces*. 436–452.

Margaret A Boden. 2004. *The Creative Mind: Myths and Mechanisms*. Routledge.

Herbie Bradley, Andrew Dai, Hannah Teufel, Jenny Zhang, Koen Oostermeijer, Marco Bellagente, Jeff Clune, Kenneth Stanley, Grégory Schott, and Joel Lehman. 2023. Quality-Diversity through AI Feedback. *arXiv preprint arXiv:2310.13032* (2023).

Kieran Browne. 2022. Who (or what) is an AI Artist? *Leonardo* 55, 2 (2022), 130–134.

Daniel Buschek, Lukas Mecke, Florian Lehmann, and Hai Dang. 2021. Nine Potential Pitfalls when Designing Human-AI Co-Creative Systems. In *Joint Proceedings of the ACM IUI 2021 Workshops*.

Baptiste Caramiaux and Sarah Fdili Alaoui. 2022. "Explorers of Unknown Planets": Practices and Politics of Artificial Intelligence in Visual Arts. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW2 (2022).

Rogelio E Cardona-Rivera, Arnav Jhala, Julie Porteous, and R Michael Young. 2024. The Story So Far on Narrative Planning. In *34th International Conference on Automated Planning and Scheduling*.

Tuhin Chakrabarty, Philippe Laban, Divyansh Agarwal, Smaranda Muresan, and Chien-Sheng Wu. 2023. Art or Artifice? Large Language Models and the False Promise of Creativity. *arXiv preprint arXiv:2309.14556* (2023).

Minsuk Chang, Stefania Druga, Alexander J Fiannaca, Pedro Vergani, Chinmay Kulkarni, Carrie J Cai, and Michael Terry. 2023. The Prompt Artists. In *Proceedings of the 15th Conference on Creativity and Cognition*. 75–87.

John William Charnley, Alison Pease, and Simon Colton. 2012. On the Notion of Framing in Computational Creativity. In *International Conference on Computational Creativity*. 77–81.

Erin Cherry and Celine Latulipe. 2014. Quantifying the creativity support of digital tools through the Creativity Support Index. *ACM Transactions on Computer-Human Interaction (TOCHI)* 21, 4 (2014).

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*. 1817–1833.

John Joon Young Chung, Wooseok Kim, Kang Min Yoo, Hwaran Lee, Eytan Adar, and Minsuk Chang. 2022. TaleBrush: Sketching stories with generative pretrained language models. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*.

John Joon Young Chung and Max Kreminski. 2024. Patchview: LLM-Powered Worldbuilding with Generative Dust and Magnet Visualization. *arXiv preprint arXiv:2408.04112* (2024).

Simon Colton, Mark Nelson, Edward Powley, Swen Gaudl, Rob Saunders, Blanca Perez Ferrer, Peter Ivey, and Michael Cook. 2018. A parameter-space design methodology for casual creators. In *International Conference on Computational Creativity*.

Kate Compton. 2016. So you want to build a generator... https://galaxykate0.tumblr.com/post/139774965871/so-you-want-to-build-a-generator. Accessed: 2024-04-25.

Kate Compton, Ben Kybartas, and Michael Mateas. 2015. Tracery: an author-focused generative text tool. In *Interactive Storytelling: 8th International Conference on Interactive Digital Storytelling, ICIDS 2015, Copenhagen, Denmark, November 30-December 4, 2015, Proceedings 8*. Springer, 154–161.

Kate Compton and Michael Mateas. 2015. Casual Creators. In *International Conference on Computational Creativity*. 228–235.

Michael Cook, Simon Colton, Alison Pease, and Maria Teresa Llano. 2019. Framing in Computational Creativity–A Survey and Taxonomy. In *International Conference on Computational Creativity*. 156–163.

Nathan Crilly. 2019. Creativity and fixation in the real world: A literature review of case study research. *Design Studies* 64 (2019), 154–168.

Sebastian Deterding, Jonathan Hook, Rebecca Fiebrink, Marco Gillies, Jeremy Gow, Memo Akten, Gillian Smith, Antonios Liapis, and Kate Compton. 2017. Mixed-initiative creative interfaces. In *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems*. 628–635.

Anil R Doshi and Oliver Hauser. 2023. Generative artificial intelligence enhances creativity. *Available at SSRN* (2023).

Fiona Draxler, Anna Werner, Florian Lehmann, Matthias Hoppe, Albrecht Schmidt, Daniel Buschek, and Robin Welsch. 2024. The AI Ghostwriter Effect: When Users Do Not Perceive Ownership of AI-Generated Text But Self-Declare as Authors. *ACM Transactions on Computer-Human Interaction* 31, 2 (2024).

Ziv Epstein, Océane Boulais, Skylar Gordon, and Matt Groh. 2020. Interpolating GANs to scaffold autotelic creativity. *arXiv preprint arXiv:2007.11119* (2020).

Ziv Epstein, Aaron Hertzmann, and Investigators of Human Creativity. 2023. Art and the science of generative AI. *Science* 380, 6650 (2023), 1110–1111.

Rebecca Fiebrink, Daniel Trueman, and Perry R Cook. 2009. A meta-instrument for interactive, on-the-fly machine learning. In *Proceedings of the International Conference on New Interfaces for Musical Expression*.

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*.

Philip Galanter. 2016. Generative art theory. In *A Companion to Digital Art*. Wiley Online Library, 146–180.

Katy Ilonka Gero and Lydia B Chilton. 2019. Metaphoria: An algorithmic companion for metaphor creation. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*.

Katy Ilonka Gero, Vivian Liu, and Lydia Chilton. 2022. Sparks: Inspiration for science writing using language models. In *Designing Interactive Systems Conference*. 1002–1019.

Jeffrey T Hancock, Mor Naaman, and Karen Levy. 2020. AI-mediated communication: Definition, research agenda, and ethical considerations. *Journal of Computer-Mediated Communication* 25, 1 (2020), 89–100.

Daniel Harris and Stacy Holman Jones. 2022. A manifesto for posthuman creativity studies. *Qualitative Inquiry* 28, 5 (2022), 522–530.

Aaron Hertzmann. 2019. Aesthetics of neural network art. *arXiv preprint arXiv:1903.05696* (2019).

Tom Hewett, Mary Czerwinski, Michael Terry, Jay Nunamaker, Linda Candy, Bill Kules, and Elisabeth Sylvan. 2005. Creativity support tool evaluation methods and metrics. In *Creativity Support Tools: A workshop sponsored by the National Science Foundation*. 10–24.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems* 33 (2020), 6840–6851.

Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. LoRA: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685* (2021).

Isa Inuwa-Dutse, Alice Toniolo, Adrian Weller, and Umang Bhatt. 2023. Algorithmic loafing and mitigation strategies in Human-AI teams. *Computers in Human Behavior: Artificial Humans* 1, 2 (2023), 100024.

Daphne Ippolito, Reno Kriz, Maria Kustikova, João Sedoc, and Chris Callison-Burch. 2019. Comparison of diverse decoding methods from conditional language models. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. 3752–3762.

Maurice Jakesch, Advait Bhat, Daniel Buschek, Lior Zalmanson, and Mor Naaman. 2023. Co-writing with opinionated language models affects users' views. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*.

David G Jansson and Steven M Smith. 1991. Design fixation. *Design Studies* 12, 1 (1991), 3–11.

Joey Donald Jones. 2023. Authorial Burden. In *The Authoring Problem: Challenges in Supporting Authoring for Interactive Digital Narratives*. Springer, 47–63.

Nikhita Joshi and Daniel Vogel. 2024. Writing with AI Lowers Psychological Ownership, but Longer Prompts Can Help. *arXiv preprint arXiv:2404.03108* (2024).

Isaac Karth. 2019. Preliminary poetics of procedural generation in games. *Transactions of the Digital Games Research Association* 4, 3 (2019).

Isaac Karth and Kate Compton. 2023. Conceptual Art Made Real: Why Procedural Content Generation is Impossible. In *Proceedings of the 18th International Conference on the Foundations of Digital Games*.

Jiho Kim, Ray C Flanagan, Noelle E Haviland, ZeAi Sun, Souad N Yakubu, Edom A Maru, and Kenneth C Arnold. 2024. Towards Full Authorship with AI: Supporting Revision with AI-Generated Views. *arXiv preprint arXiv:2403.01055* (2024).

Max Kreminski. 2023a. "Generator's Haunted"': A Brief, Spooky Account of Hauntological Effects in the Player Experience of Procedural Generation. In *Proceedings of the 18th International Conference on the Foundations of Digital Games*.

Max Kreminski. 2023b. one of my fav recent papers: "Why Oatmeal is Cheap: Kolmogorov Complexity and Procedural Generation" (by @pyrofoux and @mtrc). https://mastodon.social/@maxkreminski/110273233768792278. Retrieved: 2024-04-30.

Max Kreminski. 2024a. Computational Poetry is Lost Poetry. In *Proceedings of the Halfway to the Future Symposium*.

Max Kreminski. 2024b. The dearth of the author in AI-supported writing. In *Proceedings of the Third Workshop on Intelligent and Interactive Writing Assistants*. 48–50.

Max Kreminski and John Joon Young Chung. 2024. Intent Elicitation in Mixed-Initiative Co-Creativity. In *Joint Proceedings of the ACM IUI 2024 Workshops*.

Max Kreminski, Melanie Dickinson, Joseph Osborn, Adam Summerville, Michael Mateas, and Noah Wardrip-Fruin. 2020. Germinate: A mixed-initiative casual creator for rhetorical games. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, Vol. 16. 102–108.

Max Kreminski, Melanie Dickinson, Noah Wardrip-Fruin, and Michael Mateas. 2022a. Loose Ends: a mixed-initiative creative interface for playful storytelling. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, Vol. 18. 120–128.

Max Kreminski, Isaac Karth, Michael Mateas, and Noah Wardrip-Fruin. 2022b. Evaluating Mixed-Initiative Creative Interfaces via Expressive Range Coverage Analysis. In *IUI Workshops*. 34–45.

Max Kreminski and Michael Mateas. 2021. Reflective Creators. In *International Conference on Computational Creativity*. 309–318.

Max Kreminski, Noah Wardrip-Fruin, and N Wardrip. 2018. Gardening games: an alternative philosophy of PCG in games. In *Proceedings of the 13th International Conference on the Foundations of Digital Games*.

Mina Lee, Katy Ilonka Gero, John Joon Young Chung, Simon Buckingham Shum, Vipul Raheja, Hua Shen, Subhashini Venugopalan, Thiemo Wambsganss, David Zhou, Emad A. Alghamdi, Tal August, Avinash Bhat, Madiha Zahrah Choksi, Senjuti Dutta, Jin L.C. Guo, Md Naimul Hoque, Yewon Kim, Simon Knight, Seyed Parsa Neshaei, Antonette Shibani, Disha Shrivastava, Lila Shroff, Agnia Sergeyuk, Jessi Stark, Sarah Sterman, Sitong Wang, Antoine Bosselut, Daniel Buschek, Joseph Chee Chang, Sherol Chen, Max Kreminski, Joonsuk Park, Roy Pea, Eugenia Ha Rim Rho, Zejiang Shen, and Pao Siangliulue. 2024. A Design Space for Intelligent and Interactive Writing Assistants. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '24)*. Association for Computing Machinery, New York, NY, USA.

Mina Lee, Percy Liang, and Qian Yang. 2022. CoAuthor: Designing a human-AI collaborative writing dataset for exploring language model capabilities. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*.

Isabelle Levent and Lila Shroff. 2023. The Model is the Message. In *The Second Workshop on Intelligent and Interactive Writing Assistants*.

Jingyi Li, Eric Rawn, Jacob Ritchie, Jasper Tran O'Leary, and Sean Follmer. 2023. Beyond the Artifact: Power as a Lens for Creativity Support Tools. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*.

Yixuan Li, Jason Yosinski, Jeff Clune, Hod Lipson, and John Hopcroft. 2015. Convergent learning: Do different neural networks learn the same representations? *arXiv preprint arXiv:1511.07543* (2015).

Antonios Liapis, Georgios N Yannakakis, Constantine Alexopoulos, and Phil Lopes. 2016. Can Computers Foster Human Users' Creativity? Theory and Praxis of Mixed-Initiative Co-Creativity. *Digital Culture & Education* 8, 2 (2016), 136–153.

Zhiyu Lin, Upol Ehsan, Rohan Agarwal, Samihan Dani, Vidushi Vashishth, and Mark Riedl. 2023. Beyond Prompts: Exploring the Design Space of Mixed-Initiative Co-Creativity Systems. In *Proceedings of the Fourteenth International Conference on Computational Creativity*. 64–73.

Ryan Louie, Jesse Engel, and Cheng-Zhi Anna Huang. 2022. Expressive Communication: Evaluating Developments in Generative Models and Steering Interfaces for Music Creation. In *27th International Conference on Intelligent User Interfaces*. 405–417.

Sourab Mangrulkar, Sylvain Gugger, Lysandre Debut, Younes Belkada, Sayak Paul, and Benjamin Bossan. 2022. PEFT: State-of-the-art Parameter-Efficient Fine-Tuning methods. https://github.com/huggingface/peft.

Terry McDonell and Douglas Brinkley. 2000. Hunter S. Thompson, The Art of Journalism No. 1. https://www.theparisreview.org/interviews/619/the-art-of-journalism-no-1-hunter-s-thompson. *The Paris Review* 156 (2000).

Piotr Mirowski, Kory W Mathewson, Jaylen Pittman, and Richard Evans. 2023. Co-Writing Screenplays and Theatre Scripts with Language Models: Evaluation by Industry Professionals. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*.

Moran Mizrahi, Guy Kaplan, Dan Malkin, Rotem Dror, Dafna Shahaf, and Gabriel Stanovsky. 2023. State of What Art? A Call for Multi-Prompt LLM Evaluation. *arXiv preprint arXiv:2401.00595* (2023).

Nick Montfort and Noah Wardrip-Fruin. 2003. Six Selections by the Oulipo. In *The New Media Reader*, Noah Wardrip-Fruin and Nick Montfort (Eds.). MIT Press, 147–148.

Santiago Ontanon and Jichen Zhu. 2021. The personalization paradox: The conflict between accurate user models and personalized adaptive systems. In *26th International Conference on Intelligent User Interfaces-Companion*. 64–66.

OpenAI. 2022. DALL·E 2. https://openai.com/index/dall-e-2. Accessed: 2024-11-22.

Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F Christiano, Jan Leike, and Ryan Lowe. 2022. Training language models to follow instructions with human feedback. In *Advances in Neural Information Processing Systems*, Vol. 35. 27730–27744.

Vishakh Padmakumar and He He. 2023. Does Writing with Language Models Reduce Content Diversity? *arXiv preprint arXiv:2309.05196* (2023).

Elena Petrovskaya, Christoph Sebastian Deterding, and Simon Colton. 2020. Casual creators in the wild: A typology of commercial generative creativity support tools. In *ICCC'20: Eleventh International Conference on Computational Creativity*. Association for Computational Creativity (ACC).

A Terry Purcell and John S Gero. 1996. Design and other types of fixation. *Design Studies* 17, 4 (1996), 363–383.

Younès Rabii and Michael Cook. 2023. Why Oatmeal is Cheap: Kolmogorov Complexity and Procedural Generation. In *Proceedings of the 18th International Conference on the Foundations of Digital Games*.

Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics. https://arxiv.org/abs/1908.10084

Christian Remy, Lindsay MacDonald Vermeulen, Jonas Frich, Michael Mose Biskjaer, and Peter Dalsgaard. 2020. Evaluating creativity support tools in HCI research. In *Proceedings of the 2020 ACM Designing Interactive Systems Conference*. 457–476.

Melissa Roemmele and Andrew Gordon. 2018. Linguistic features of helpfulness in automated support for creative writing. In *Proceedings of the First Workshop on Storytelling (Story-NLP 2018)*. 14–19.

Donald Alan Schön. 1983. *The Reflective Practitioner: How Professionals Think in Action*. Basic Books.

Abigail See, Aneesh Pappu, Rohun Saxena, Akhila Yerukola, and Christopher D Manning. 2019. Do Massively Pretrained Language Models Make Better Storytellers? In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*. Association for Computational Linguistics, 843–861. https://doi.org/10.18653/v1/K19-1079

Noor Shaker, Julian Togelius, and Mark J Nelson. 2016. *Procedural Content Generation in Games*. Springer.

Ari Shapiro. 2016. Jhumpa Lahiri Finds Freedom In Italian Memoir: 'No One Expected Me To Do It'. https://www.npr.org/transcripts/466001114.

Ben Shneiderman. 2007. Creativity support tools: accelerating discovery and innovation. *Commun. ACM* 50, 12 (2007), 20–32.

Ilia Shumailov, Zakhar Shumaylov, Yiren Zhao, Yarin Gal, Nicolas Papernot, and Ross Anderson. 2023. The curse of recursion: Training on generated data makes models forget. *arXiv preprint arXiv:2305.17493* (2023).

Adam M Smith and Michael Mateas. 2011. Answer set programming for procedural content generation: A design space approach. *IEEE Transactions on Computational Intelligence and AI in Games* 3, 3 (2011), 187–200.

Gillian Smith and Jim Whitehead. 2010. Analyzing the expressive range of a level generator. In *Proceedings of the 2010 Workshop on Procedural Content Generation in Games*.

Susan Sontag. 1977. *On Photography*. Farrar, Straus & Giroux.

Nick St. Pierre. 2024. What gets generated versus what gets posted. Expecting perfect results every time you run a prompt is a fantasy. https://twitter.com/nickfloats/status/1744710730526421450. Accessed: 2024-04-30.

Jessi Stark, Anthony Tang, Young-Ho Kim, Joonsuk Park, and Daniel Wigdor. 2023. Can AI Support Fiction Writers Without Writing For Them? In *Proceedings of the Second Workshop on Intelligent and Interactive Writing Assistants*.

Adam Summerville, Sam Snodgrass, Matthew Guzdial, Christoffer Holmgård, Amy K Hoover, Aaron Isaksen, Andy Nealen, and Julian Togelius. 2018. Procedural content generation via machine learning (PCGML). *IEEE Transactions on Games* 10, 3 (2018), 257–270.

Atau Tanaka. 2006. Interaction, experience and the future of music. In *Consuming music together: Social and collaborative aspects of music consumption technologies*. Springer, 267–288.

E Paul Torrance. 1966. Torrance tests of creative thinking. *Educational and Psychological Measurement* (1966).

Hannah Twigg-Smith, Jasper Tran O'Leary, and Nadya Peek. 2021. Tools, Tricks, and Hacks: Exploring Novel Digital Fabrication Workflows on #PlotterTwitter. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*.

Johan Ugander and Ziv Epstein. 2024. The art of randomness: Sampling and chance in the age of algorithmic reproduction. *Harvard Data Science Review* 6, 4 (2024).

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in Neural Information Processing Systems* 30 (2017).

Lili Yao, Nanyun Peng, Ralph Weischedel, Kevin Knight, Dongyan Zhao, and Rui Yan. 2019. Plan-and-write: Towards better automatic storytelling. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 7378–7385.

JD Zamfirescu-Pereira, Richmond Y Wong, Bjoern Hartmann, and Qian Yang. 2023. Why Johnny can't prompt: how non-AI experts try (and fail) to design LLM prompts. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*.

Yiming Zhang, Avi Schwarzschild, Nicholas Carlini, Zico Kolter, and Daphne Ippolito. 2024. Forcing Diffuse Distributions out of Language Models. *arXiv preprint arXiv:2404.10859* (2024).

David Zhou and Sarah Sterman. 2023. Creative Struggle: Arguing for the Value of Difficulty in Supporting Ownership and Self-Expression in Creative Writing. In *The Second Workshop on Intelligent and Interactive Writing Assistants (In2Writing)*.