


# Homogenization Effects of Large Language Models on Human Creative Ideation

Barrett R. Anderson  
Independent Researcher  
Santa Cruz, California, USA  
barrettrees@gmail.com

Jash Hemant Shah  
Santa Clara University  
Santa Clara, California, USA  
jshah5@scu.edu

Max Kreminski  
Santa Clara University  
Santa Clara, California, USA  
mkreminski@scu.edu



**Figure 1: Homogenization analysis involves semantic similarity comparisons between artifacts produced by users of creativity support tools (CSTs). We apply homogenization analysis to two different CSTs for divergent ideation, and find that users of the Oblique Strategies deck (on the left) and ChatGPT (on the right) each produce similarly homogenous sets of ideas as *individuals*—but collectively, users of ChatGPT produce a more homogenous set of ideas at the *group* level (as shown by the higher degree of overlap between the sets of ideas produced by each user).**

## ABSTRACT

Large language models (LLMs) are now being used in a wide variety of contexts, including as creativity support tools (CSTs) intended to help their users come up with new ideas. But do LLMs actually support user creativity? We hypothesized that the use of an LLM as a CST might make the LLM’s users feel more creative, and even broaden the range of ideas suggested by each individual user, but also homogenize the ideas suggested by different users. We conducted a 36-participant comparative user study and found, in accordance with the homogenization hypothesis, that different users tended to produce less semantically distinct ideas with ChatGPT than with an alternative CST. Additionally, ChatGPT users generated a greater number of more detailed ideas, but felt less responsible for the ideas they generated. We discuss potential implications of these findings for users, designers, and developers of LLM-based CSTs.

## CCS CONCEPTS

• **Human-centered computing** → Empirical studies in HCI;  
• **Applied computing** → Arts and humanities; • **Computing methodologies** → Natural language processing.

## KEYWORDS

creativity support tools, divergent ideation, large language models, user study

## ACM Reference Format:

Barrett R. Anderson, Jash Hemant Shah, and Max Kreminski. 2024. Homogenization Effects of Large Language Models on Human Creative Ideation. In *Creativity and Cognition (C&C '24)*, June 23–26, 2024, Chicago, IL, USA. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3635636.3656204>

## 1 INTRODUCTION

Success in creative contexts (such as creative writing and product design) often hinges on the ability to come up with ideas that are—in line with prominent creativity researcher Margaret Boden’s definition of creativity—simultaneously “new, surprising, and valuable” to some extent [9]. Though research in *creativity support tools* (CSTs) has long aimed to produce software systems that can support parts of human creative processes [19, 31, 78], it was largely not until the recent wave of developments in large language models (LLMs) that software tools capable of directly generating potentially creative ideas in arbitrary domains began to seem feasible. Especially since the release of ChatGPT [67] in November 2022, large numbers of people have begun to use LLM-based systems as CSTs, consulting LLMs for ideas in creative contexts as wide-ranging as recipe



This work is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike International 4.0 License.

C&C '24, June 23–26, 2024, Chicago, IL, USA  
© 2024 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-0485-7/24/06  
<https://doi.org/10.1145/3635636.3656204>

creation [24], tabletop roleplaying game scenario design [64], and marketing slogan generation [61].

The adoption of LLMs in creative contexts has raised questions about the extent to which these models can assist in the production of genuinely creative outputs. In particular, some researchers have expressed concern that the widespread use of a small number of highly centralized, data-driven AI systems (such as ChatGPT) may lead to decreased diversity in the outputs of creative processes that incorporate these tools [13, 29, 36, 55]. These concerns resonate with earlier work that has proposed diversity of output as a potential evaluation criterion for AI-based CSTs in general [49] and LLM-based CSTs in particular [14]. Despite these earlier discussions, however, it is only very recently that researchers have begun to directly study the question of whether the use of AI-based CSTs leads to homogenization of human creative output [4, 25, 41, 49, 69].

To advance understanding of homogenization effects in human-AI co-creative contexts, we conducted a 36-participant comparative user study of ChatGPT and an alternative, non-AI CST. Participants completed four divergent ideation tasks based on a subset of the Torrance Tests of Creative Thinking (TTCT) [85]; each participant completed half of these tasks with ChatGPT and half with the non-AI CST. Participants produced 1271 ideas in total. Based on the resulting data, we investigate four research questions:

- RQ1 With which CST do participants produce more semantically similar ideas at the *group* level? (A: **ChatGPT**)
- RQ2 With which CST do participants produce more semantically similar ideas at the *individual* level? (A: **no difference**)
- RQ3 Do ChatGPT users feel more or less *responsible* for the ideas they produce? (A: **less responsible**)
- RQ4 Beyond originality, do ChatGPT and non-ChatGPT users differ in terms of *other facets* of creativity, such as fluency, flexibility, and elaboration? (A: **ChatGPT** → **higher fluency, flexibility, and elaboration**)

In the remainder of this paper, we first summarize prior work on the creative homogenization effects of AI-based CSTs, reasons we might expect LLM-based CSTs to result in creative homogenization, and approaches to the evaluation of CSTs in general. We then describe our experimental procedure and results. Finally, we discuss potential implications of our findings for users, designers, and developers of LLM-based CSTs. Uniquely among recent studies, we compare the homogenization effects of LLMs to those of a potential alternative CST; tease apart individual user-level from group-level homogenization effects; and extend the study of LLM-driven homogenization effects outside the domain of writing.

Collectively, our results suggest that LLM-driven homogenization stems from the LLM providing different users with similar ideas, rather than by increasing individual user-level fixation; that low inferential distance between LLM outputs and apparently finished creative products may contribute to homogenization; and that users may be able to resist homogenization effects if they are given a sense of what the model tends to suggest in similar contexts. These conclusions imply potential mitigations for homogenization effects at the CST design level and clear research directions for follow-up work.

## 2 BACKGROUND AND RELATED WORK

### 2.1 Prior Studies of Homogenization

To date, there have been relatively few direct studies of creative homogenization resulting from the use of AI-based CSTs. Prior to the widespread adoption of LLMs, Arnold et al. found that “predictive text encourages predictable writing” [4] in the context of single-word suggestions given by smartphone keyboards. Subsequently, Kreminski et al. evaluated whether an AI-based poetry composition tool caused users to produce more or less similar poems over time as they continued to use the tool [49].

Two recent user studies of LLMs provide direct evidence of a homogenization effect in LLM-supported writing. Padmakumar and He evaluate the baseline GPT-3 model [12] versus the instruction-finetuned variant InstructGPT [68] in the context of a short-form argumentative essay writing task and find a homogenization effect from LLM assistance at both the lexical and content levels, but only for the instruction-tuned LLM [69]. Meanwhile, Doshi and Hauser evaluate GPT-4 in the context of short-form fictional narrative writing and observe a similar homogenization effect [25].

While not addressing homogenization effects directly, several other recent studies examine how LLMs change the writing of humans who use them for writing support [53]. Lee et al. [54] find that LLM support tends to increase the diversity of a user’s vocabulary but may reduce their feelings of ownership for the text they produce. Jakesch et al. [41] find that an opinionated LLM-based CST influences the opinions expressed by its users in argumentative writing. Similarly, Bhat et al. [8] find that LLM-supplied next-phrase suggestions may alter the form and content of a human user’s writing even when the user dislikes these suggestions. Potentially explaining these results, Roemmele [74] proposes and provides evidence for an “inspiration through observation” model of how human writers are influenced by LLMs: observation of LLM-generated text shapes the ideas expressed by the LLM’s user even when the user does not incorporate LLM-generated text directly into their writing.

Still other studies investigate the homogeneity of LLM outputs in creative contexts without examining their effects on human-in-the-loop creative processes (where a human user is more actively involved in curating and refining LLM output). Begus [7] compares human-written to GPT-generated short stories, observing that GPT-generated stories are less diverse than human-written stories in structure but exhibit greater gender and sexual diversity in character description. In a similar vein, Chakrabarty et al. [15] compare LLM-generated short stories to stories by expert human writers and find that the human-written stories substantially outperform the LLM-generated stories on all dimensions of creativity measured by the TTCT, including the originality dimension.

Altogether, we are aware of only two previous direct studies of LLM-driven creative homogenization. Both of these studies only compare LLM-assisted to tool-unassisted users, making it difficult to determine how the homogenization effects of LLMs compare to those of potential alternative CSTs. Additionally, both studies focus specifically on writing tasks, and consequently both studies solicited only one creative product (i.e., written essay or short story) from each participant; thus, it is unclear whether homogenization effects come into play predominantly at the individual user or group level. Existing evidence seems to support the existence of homogenization

effects overall, but more research is needed to gauge the severity of these effects; to confirm their existence outside of writing tasks specifically; and to clarify how they emerge.

## 2.2 Reasons to Expect Homogenization

Concerns about LLM-driven homogenization of creative outputs [29, 55] have so far mostly been articulated in terms of *algorithmic monoculture* [44]. The use of a single consistent AI system to perform tasks that used to be performed by many very different humans or systems can lead to increased homogenization of outcomes [10]; similar effects might be expected on creative processes if many different people all begin using ChatGPT (or another singular LLM-based system) as a CST.

Beyond monoculture, we might also expect LLM-based CSTs to result in homogenization for several other reasons. Creative design processes are subject to measurable *fixation* effects [21, 42, 71], through which the final proposed solutions to a design problem are inappropriately constrained by features of earlier solution candidates [2]; LLMs may induce fixation by presenting users with complete-seeming ideas early in the ideation process, thereby reducing variation in later ideas. Past work in HCI has proposed that underdetermination may be valuable in human-machine co-creation [1]; by producing text that looks “finished”, LLM-based CSTs may foreclose the space of desirable ambiguity [32] that results in legitimate creativity. If people trust LLMs due to *trustworthiness cues* [57] such as an authoritative-sounding writing style, or due to a general trust in machines [84], they may treat LLM-suggested ideas as good or valuable by default, causing them to take up these ideas without closely examining them. Further, groups tend to be less creative than individuals when majoritarian convergence processes are used, and minority dissent within group creative processes tends to lead to greater creativity [30]; an LLM trained to reproduce statistically likely results may seem to speak with the authority of the majority and therefore reduce creativity. If LLMs indeed exert homogenization effects on creative processes, any of these proposed mechanisms might be at work.

## 2.3 Evaluating Creativity and CSTs

Evaluation of CSTs has remained an open problem since essentially the beginning of CST research [38], due in part to the ambiguous and multifaceted nature of “creativity” as a phenomenon, in part to the wide range of (sometimes contradictory) user needs associated with different creative contexts, and in part to the lack of a clear consensus around what aspects of CSTs should be evaluated [73]. Broadly speaking, approaches to the evaluation of CSTs can be divided into two categories: those that primarily evaluate aspects of the creative *process* when the CST is used, and those that primarily evaluate the creative *products* that emerge from this process.

On the process side, CSTs are most frequently evaluated by means of subjective self-reports of experience from tool users. The Creativity Support Index (CSI) [17] is a widely used and psychometrically validated survey instrument that attempts to standardize some aspects of this experience reporting process across different CSTs. Other (often bespoke) survey instruments are also deployed in CST evaluation, either as a supplement or an alternative to the CSI (e.g., [18, 35, 47, 48, 88]). Process is sometimes also evaluated

via observation of user *actions* during the creative process (e.g., [3, 22, 47, 49, 88]). Evaluations of LLM-based CSTs have largely followed this pattern to date: most such CSTs are evaluated primarily through subjective experience reports and secondarily through observation of user actions (e.g., [14, 81–83, 86]).

On the product side, CSTs can also be evaluated by examining the quantity, quality, or other characteristics of the artifacts that their users produce. The Torrance Tests of Creative Thinking (TTCT) [85] evaluate the creativity of test-takers according to four facets of creative output: *fluency*, or sheer quantity of artifacts created; *flexibility*, or quantity of distinct categories of artifacts created; *originality*, or dissimilarity of created artifacts to others’ creations; and *elaboration*, or level of detail in created artifacts. These same criteria can also be applied to the evaluation of CSTs by comparing a novel tool’s users to users of another baseline tool along these lines. Notably, the TTCT does not include artifact *quality* as an evaluation criterion; where studies of CSTs have attempted to evaluate output quality, human raters have usually been employed to judge the results (e.g., [23, 33]). This pattern holds for studies of LLM-based CSTs as well (e.g., [20, 25, 59]).

Our comparison of ChatGPT to a non-AI CST makes use of both process and product data, with a particular focus on the assessment of homogenization effects via examination of creative products: the ideas that study participants produce. Homogenization effects are most closely linked to the *originality* dimension of the TTCT; like other parallel studies of homogenization effects [25, 69], we investigate originality primarily by means of semantic similarity, using a well-performing sentence embedding model [72] whose direct predecessors [70] have been found to agree well with human judgments of originality in creativity research [6, 26]. We also use product data to assess CST effects on fluency, flexibility, and elaboration (the other three facets of creativity that the TTCT attempts to gauge) via simple idea count, manual categorization of ideas, and stoplisted word count [27] respectively. Self-reported user experience data (collected via both the CSI and bespoke survey items) is used to assess CST effects on participant feelings of responsibility and other user experience qualities.

## 3 METHODS

We conducted a within-subjects experiment to evaluate the effects of using two different CSTs for idea generation: ChatGPT and the Oblique Strategies (OS) deck.

### 3.1 Participants

**3.1.1 Recruitment.** Participants were recruited from academic mailing lists, forums, and solicitations posted by the experimenters on social media. Our study protocol and recruitment materials were approved by the Santa Clara University IRB. Participation was incentivized with a \$17.50 gift card for a one hour session. Participants were required to have a stable Internet connection, a device capable of screen-sharing, and access to a quiet place for the duration of the session in order to participate in the study. Of the 36 participants originally recruited, three were excluded from all analyses for failure to follow direction (e.g., not using ChatGPT when prompted to do so).

**3.1.2 Demographics.** Our sample included 33 participants, ranging in age from 22 to 44 ( $M=28.36$ ,  $SD=6.84$ ), and including 63.63% ( $n=21$ ) men and 33.36% ( $n=12$ ) women. We had 39.39% Black or African American ( $n=13$ ), 36.36% Asian ( $n=12$ ), and 24.24% ( $n=8$ ) White participants. Participant occupations included 36.36% students ( $n=12$ ), 30.30% creative professionals (e.g. game designer, writers,  $n=10$ ), and 33.33% other professionals (e.g. registered nurse, social worker, customer service,  $n=11$ ). Educational experience included 15.15% high school graduates ( $n=5$ ), 15.15% participants with some college ( $n=5$ ), 48.48% college graduates ( $n=16$ ), and 21.21% participants with a Master’s degree or higher ( $n=7$ ). Experience with text-based generative AI (e.g. ChatGPT) was varied, with 72.72% of participants reporting daily ( $n=16$ ) or weekly ( $n=8$ ) usage of LLM tools, and 27.27% of participants reporting that they had used them once a month or less ( $n=3$ ), or that they had never used them before the study ( $n=6$ ). Experience with generative art AI tools (e.g. Midjourney, Stable Diffusion) was less common: 63.63% of participants reported that they had either never used such tools ( $n=13$ ) or used them once a month or less ( $n=9$ ), and 33.33% of participants reported that they used them about once a week ( $n=5$ ) or daily ( $n=6$ ).

## 3.2 Materials

**3.2.1 ChatGPT.** ChatGPT is a popular LLM-based tool trained to respond to text instructions [67]. Participants in this study used the versions of ChatGPT 3.5 released on May 3rd 2023 ( $n=6$ , 16.6%) and on August 3rd, 2023 (30, 83.3%).

**3.2.2 Oblique Strategies Deck.** The Oblique Strategies (OS) deck, originally created by the artists Brian Eno and Peter Schmidt [28], consists of a collection of cards with prompts designed to support creative work. Example prompts include “Turn it upside down”, “Don’t avoid what is easy”, “Destroy the most important thing”, and “How would someone else do it?” We directed participants to a web app version of the deck [76] as an alternative CST to ChatGPT in our control condition.

**3.2.3 Creative Ideation Prompts.** We provided creative ideation prompts for two types of divergent thinking tasks: *Product Improvement* (PI) and *Improbable Consequences* (IC). Prompts included:

- How could you make a stuffed toy animal more fun to play with? (PI\_A)
- How could you make a jigsaw picture puzzle more interesting and engaging? (PI\_B)
- Suppose that a great fog has fallen over the earth and all we can see of people is their feet. What would happen? (IC\_A)
- Suppose that gravity suddenly became incredibly weak, and objects could float away easily. What would happen? (IC\_B)

For each prompt, participants were instructed to generate as many ideas as possible and to try to come up with ideas that no one else would think of.

**3.2.4 Creativity Support Index.** The Creativity Support Index (CSI) is a survey instrument for assessing the ability of a CST to assist a user engaged in creative work [17]. Its design was inspired by the NASA TLX, a questionnaire designed to evaluate mental task load [37]. We administered the CSI to capture participant experiences with each CST after they used that CST to complete a creative ideation task.

## 3.3 Procedure

All experimental sessions were remote-moderated over videoconferencing software. In each session, participants were asked to generate ideas in response to several ideation prompts, first while using one of two CSTs (ChatGPT or OS) and then while using the other tool. Participants were instructed to generate as many ideas as they could, and to try to come up with ideas that no one else would think of.

During the session, participants were encouraged to think-aloud, to the degree that they felt doing so would not interfere with their performance. The time was held constant at 8 minutes per ideation prompt, and participants responded to two prompts with each support tool. With each tool, the first prompt asked participants to come up with ideas for improving an existing product (Product Improvement). The second prompt asked them to consider an impossible situation and imagine as many possible consequences as they could think of (Improbable Consequences). The order of CSTs and prompts was randomized per participant and balanced across the entire experiment.


After using each tool the participants responded to the Creativity Support Index questionnaire, and indicated the degree to which they felt personally responsible for their output, or that they felt their output came from the tool that they used. Each session concluded with an open-ended discussion of each participant’s experience with both ChatGPT and OS.

## 4 RESULTS

We observed three categories of data about working with each CST: creative process, creative outcomes, and retrospective reflections on the experience of working with the tool. Creative process data included video-recordings of moderated sessions conducted with a talk-aloud protocol, which resulted in observations about prompting styles and iteration, and how participants used ChatGPT output. Creative outcomes included the list of ideas generated with each CST. Retrospective reflections included Creativity Support Index (CSI) ratings, ratings of how personally responsible participants felt for the ideas they generated (vs. crediting the CST), and a brief open-ended discussions about the experience of using each CST conducted at the end of each experimental session. The results reported below are organized by research question, addressing homogenization (at group and individual levels), participants’ sense of responsibility for the ideas they produced, and other facets of creativity. We also provide a summary of observations from participant interviews, and some observations of their creative process. Data from three participants who did not follow directions (e.g., not using ChatGPT when instructed to do so) were excluded from our analysis.

### 4.1 Homogenization

**4.1.1 Evaluating Homogenization via Semantic Similarity.** To evaluate creative outcomes we followed two approaches to quantifying participant-generated ideas, complementing a traditional human rater-based approach with semantic similarity assessment via sentence embeddings [72]. Sentence embeddings allow us to quantify homogenization in the form of semantic similarity, comparing the



**Figure 2: Participant responses were more homogenous at the group level (i.e., more semantically similar to the average embedding of *all* participant ideas) when using ChatGPT.<sup>1</sup>**

cosine similarity of each idea a participant generated to the average embedding of all participant ideas, to evaluate how distinct an individual participant’s ideas are from the group. By comparing to an average embedding for the individual, rather than the group, we are also able to evaluate the diversity of each participant’s ideas.

Our semantic similarity-based approach to homogenization analysis closely follows recent psychological studies of creativity—e.g., [6]. However, the specific sentence embeddings that we used for our homogenization analysis (though high-performing in general) have not previously been validated for creativity assessment. As a result, we performed a small experiment to validate the agreement of these embeddings with human judgments of semantic similarity on our dataset. This experiment is discussed in Appendix A.


**4.1.2 Group-Level Homogenization.** When participants used ChatGPT, the ideas they produced were less divergent from the average embedding of all ideas generated for that task, ( $M = .24$ ,  $SD = .07$ ), compared to the ideas that they produced when using OS, ( $M = .28$ ,  $SD = .08$ ),  $t(32) = 2.154$ ,  $p = 0.038$ ,  $d = .47$ , 95% CI [.00,.07]. See Figure 2. At the group level, ideas produced with the help of ChatGPT were more homogenized.

**4.1.3 Individual-Level Homogenization.** When participants used ChatGPT, the ideas they produced were not observably more divergent from the average embedding of all of the other ideas that they themselves generated for the same ideation prompt, ( $M = .65$ ,  $SD = .07$ ), compared to semantic dissimilarity for ideas that they produced when using OS, ( $M = .66$ ,  $SD = .08$ ),  $t(32) = .944$ ,  $p = 0.352$ ,  $d = .12$ , 95% CI [-.04,.01]. See Figure 3. At the individual level, we did not observe a difference in homogenization for ideas generated with the help of ChatGPT.


## 4.2 Sense of Responsibility

Participants assigned less responsibility to themselves (and more to the tool) for ideas generated while using ChatGPT ( $M = 48.17\%$ ,  $SD = 26.22\%$ ), compared to ideas generated while using OS ( $M = 63.63\%$ ,

<sup>1</sup>All error bars are 95% Confidence Intervals, with Cousineau-Morey (2008) corrections for within-subjects data [62].



**Figure 3: Participant responses were not observably more or less homogenous at the individual level (i.e., more semantically similar to the average embedding of *this participant’s* ideas) when using ChatGPT.**




**Figure 4: Participants assigned less creative responsibility to themselves, and more to the CST, when working with ChatGPT.**

$SD = 17.36\%$ ),  $t(32) = 3.21$ ,  $p = 0.003$ ,  $d = .67$ , 95% CI [-24.60%, -5.51%]. See Figure 4.


## 4.3 Other Facets of Creativity

In addition to the TTCT dimension of originality, quantified in our homogenization analysis above, we also studied the three other TTCT dimensions of creativity quantitatively: fluency by simple idea count, flexibility by human coding of ideas into categories, and elaboration via stoplisted word count [27]. Further, we assessed originality from another angle via several different means of gauging idea uniqueness. We did not evaluate idea quality, which is subjective; may vary substantially depending on the assumed use case for the idea; and is both time-consuming and costly to evaluate via the most widely used method (annotation by human raters).

**4.3.1 Fluency: Simple Idea Count.** Participants generated about one additional idea—approximately a 15% increase—when using ChatGPT ( $M = 8.39$ ,  $SD = 3.39$ ), compared to the number of ideas



**Figure 5: Participants generated more ideas with ChatGPT than with OS.**



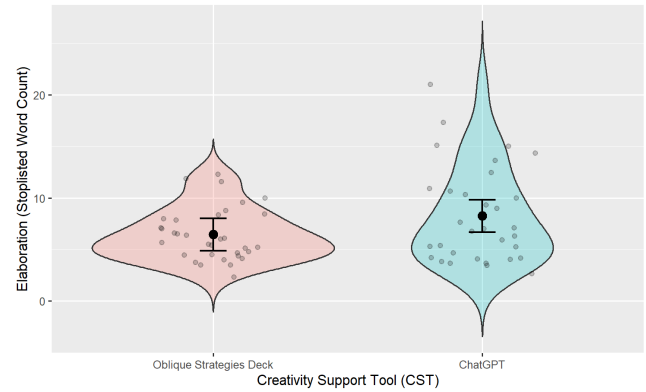
**Figure 6: Participants generated ideas across more categories with ChatGPT than with OS.**

they generated when using OS ( $M = 7.32$ ,  $SD = 3.22$ ),  $t(32) = 2.10$ ,  $p = 0.044$ ,  $d = .32$ , 95% CI [.03,2.11]. See Figure 5.

**4.3.2 Flexibility: Idea Categories.** Following an iterative grounded theory based approach [16], we reviewed the ideas generated for each ideation prompt and observed the categories of ideas that emerged, ultimately generating 181 distinct idea categories from 1271 individual responses. Each participant’s responses were then tagged with all relevant idea categories. This process resulted in a count of idea categories hit by each participant, including a sub-count of each participant’s idea categories hit that were unique within our sample. Coding was conducted while blind to the CST condition (i.e., coders did not know if an idea was generated with or without LLM support).

Participants generated ideas that hit about 27% more categories when using ChatGPT ( $M = 8.58$ ,  $SD = 2.90$ ), compared to the number of categories covered when using OS ( $M = 6.77$ ,  $SD = 3.69$ ),  $t(32) = 3.50$ ,  $p = 0.001$ ,  $d = .54$ , 95% CI [.75,2.86]. See Figure 6.

**4.3.3 Elaboration: Stoplised Word Count.** Among several potential computational means of gauging idea elaboration (which is roughly synonymous with complexity or level of detail), stoplised



**Figure 7: Participants generated more elaborated ideas with ChatGPT than with OS.**

word count (i.e., counting the words used to express the idea while excluding a “stolist” of common low-information filler words) demonstrates the strongest correlation with human judgments of elaboration [27]. Participant ideas generated while using ChatGPT had a greater stoplised word count ( $M = 8.25$ ,  $SD = 4.61$ ) than ideas generated while using OS, ( $M = 6.46$ ,  $SD = 2.55$ ),  $t(32) = 2.32$ ,  $p = 0.0237$ ,  $d = .48$ , 95% CI [.22,3.37]. See Figure 7.

**4.3.4 Originality: Unique Ideas.** We did not observe a difference in the number of unique ideas (i.e., ideas that no other participant also generated)<sup>2</sup> between participants using ChatGPT ( $M = .74$ ,  $SD = .69$ ) and OS, ( $M = .97$ ,  $SD = 1.00$ ),  $t(32) = 1.65$ ,  $p = 0.108$ ,  $d = .26$ , 95% CI [-.05,.50].

Because uniqueness is sensitive to sample size [79], we also attempted two common alternative approaches to assessing originality: counting idea categories that less than 5% of the participants produced as unique [58], and weighting each idea category by the frequency of its occurrence in the entire sample for this study to produce a weighted flexibility score [34, 75].

However, neither of these approaches changed our findings. Applying a 5% uniqueness threshold to idea categories, we again did not observe a difference in the number of unique responses (with a 5% threshold) produced by participants using ChatGPT ( $M = .87$ ,  $SD = .88$ ) or OS ( $M = .90$ ,  $SD = 1.07$ ),  $t(32) = .134$ ,  $p = 0.894$ ,  $d = .02$ , 95% CI [-.34,.39]. Furthermore, as measured by our weighted flexibility score, participants using ChatGPT did not generate more unique ideas ( $M = 6.34$ ,  $SD = 3.31$ ) compared to participants using OS, ( $M = 5.50$ ,  $SD = 3.01$ ),  $t(32) = 1.424$ ,  $p = 0.164$ ,  $d = .26$ , 95% CI [-.35,2.02].

## 4.4 Retrospective Reflections

**4.4.1 Interview.** Participants were asked to discuss their own experiences using both tools, and several themes emerged from those discussions. The single most common theme was that ChatGPT was easier to use but less rewarding (27.78%,  $n = 10$ ). The second most common theme was a positive sentiment regarding the speed

<sup>2</sup>Creativity research sometimes also deals with *historical* uniqueness, analogous to Boden’s “H-creativity” [9]. However, we neither expected nor saw any historically unique ideas submitted by participants in our study.

Theme	Example Responses	n	%
Effort/Reward Tradeoff (Oblique Strategies)	<i>Oblique Strategies got me thinking more creatively, but I got more responses with ChatGPT.</i>	10	27.78%
OS was more challenging to use, but also more rewarding.	<i>It was harder to use Oblique Strategies, but it was more fun and it got me to more interesting places.</i>		
Speed/Accuracy (ChatGPT)	<i>ChatGPT gave me the right answers.</i>	9	25.00%
ChatGPT was fast, and its responses were accurate.	<i>Using ChatGPT is a very nice experience... It's very fast and accurate.</i>		
Low Engagement (ChatGPT)	<i>ChatGPT allowed me to turn my brain off. It did more of the heavy lifting.</i>	8	30.56%
Using ChatGPT was less engaging.	<i>ChatGPT reduced the confidence I had to come up with creative things on my own.</i>		
Low Task Relevance (Oblique Strategies)	<i>I didn't really understand Oblique Strategies. It didn't relate to most of the questions.</i>	7	19.44%
Responses from OS were less task-relevant.	<i>The cards were inspirational, but most of them were just random thoughts.</i>		
Repetitive Responses (ChatGPT)	<i>ChatGPT is a more research-based tool. ChatGPT is a bit repetitive, but it has a lot of data.</i>	3	8.33%
ChatGPT responses were repetitive.	<i>When I asked for more [ChatGPT] repeated half... When I want more, I want different more.</i>		
High Engagement (Oblique Strategies)	<i>I got into a flow with Oblique Strategies.</i>	3	8.33%
OS was more engaging.	<i>[Oblique Strategies cards] were more interesting than ChatGPT.</i>		
Premature Closure (ChatGPT)	<i>ChatGPT feels like it can go really specific really quickly. Almost more than you need.</i>	2	5.56%
The ChatGPT responses became too specific too quickly.	<i>With ChatGPT, I felt like it was more guided and way more specific.</i>		
Self-Doubt (ChatGPT)	<i>I felt like ChatGPT was a tool I could use in a deeper way.</i>	2	5.56%
Self-deprecation regarding technical ability.	<i>I didn't really feel like I knew good questions to ask ChatGPT.</i>		

**Table 1: Reflections on experiences with idea generation using both CSTs (ChatGPT and OS).**

and accuracy of the LLM-provided responses (25.00%,  $n = 9$ ). Participants also reported finding ChatGPT to be less engaging (22.22%,  $n = 8$ ), that the LLM responses were too repetitive (8.33%,  $n = 3$ ) and that the LLM responses became too specific too quickly (5.56%,  $n = 2$ ). Participants also made self-deprecating remarks regarding their familiarity and technical ability with ChatGPT (5.56%,  $n = 2$ ), but made no such remarks regarding OS.

**4.4.2 Creativity Support Index.** We did not observe any differences in Creativity Support Index (CSI) ratings for the LLM-based tool ( $M=78.03\%$ ,  $SD=18.82\%$ ) and for OS ( $M=73.98\%$ ,  $SD=15.35\%$ ),  $t(35) = 1.028$ ,  $p = 0.312$ ,  $d = .24$ , 95% CI [-3.94%, 12.02%]. We also observed no differences for any of the CSI sub-scales (*Exploration, Engagement, Effort/Reward Tradeoff, Tool Transparency, Expressiveness*). We did not collect responses for the *Collaboration* subscale, which is irrelevant and often omitted in exclusively single-user contexts like that of our study (e.g., [5, 87]).

## 4.5 Creative Process

**4.5.1 Initial Ideation.** When using ChatGPT, about two-thirds (63.89%,  $n=23$ ) of our participants began by providing a number of their own ideas ( $M=5.65$ ,  $SD=3.90$ ) before interacting with the tool at all. This was similar to the interaction pattern we observed with OS, and contrary to our expectation that participants might rely on the CST to overcome a blank page [50, 52].

**4.5.2 Prompting Styles and Iteration.** When interacting with ChatGPT, the majority of participants (86.11%,  $n = 31$ ) directly copy/pasted the creative ideation prompt for their current task. Other prompting strategies included asking for relevant scientific or factual information (e.g. "What are the most popular dolls?", "How common is face blindness?", etc.), asking for responses from a specific perspective (e.g. project manager, scientist, athlete, etc.), or adding creative constraints (e.g. make this toy for a dog, respond in the form of a story, etc.). Most participants (72.22%,  $n = 26$ ) also iterated on their prompt, either by adding more context to their initial prompt, probing with more specific questions or variations, or by asking ChatGPT to regenerate answers or provide additional responses.

**4.5.3 ChatGPT Output Usage.** When using ideas from ChatGPT, a slight plurality of participants copy and pasted from the output

directly (41.67%,  $n = 15$ ). The majority of these participants copied selectively: only 1 copied the entirety of the ChatGPT output without any attempt at curation or editing. A minority of participants entirely avoided direct copying, providing ideas either paraphrased from or inspired by the ChatGPT output (16.67%,  $n = 6$ ). A substantial fraction of participants combined both approaches, copying a few ideas directly, modifying others, and later adding more of their own ideas with no obvious connection to the ChatGPT output (38.89%,  $n = 14$ ).

**4.5.4 Process Impact on Outcome.** We observed a significant positive relationship between the number of prompts a participant entered into ChatGPT and the number of ideas that they generated,  $r(32)=.43$ ,  $p = .008$ , and between the number of prompts entered and weighted uniqueness scores,  $r(32)=.46$ ,  $p = .004$ . We did not observe any relationship between the number of LLM prompts entered and average homogeneity of ideas generated,  $r(32)=.23$ ,  $p = .185$ .

We did not observe any difference in the number of ideas generated from participants who started by entering their own ideas first ( $M = 8.38$ ,  $SD = 3.39$ ), compared to those who went directly to ChatGPT, ( $M = 8.71$ ,  $SD = 3.58$ ),  $t(32) = .276$ ,  $p = .784$ , 95% CI [-2.11, 2.77]. We also observed no difference in the weighted number of unique ideas (Own Ideas First:  $M = 7.21$ ,  $SD = 3.12$ , LLM First:  $M = 7.46$ ,  $SD = 3.12$ ),  $t(32) = .222$ ,  $p = .826$ , 95% CI [-1.98, 2.47], or in the homogeneity of their ideas, between participants who took these different approaches, (Own Ideas First:  $M = .99$ ,  $SD = .03$ , LLM First:  $M = .99$ ,  $SD = .03$ ),  $t(32) = .208$ ,  $p = .836$ , 95% CI [-.01, .02].

We did not observe any difference in the number of ideas generated from participants who used any creative prompting strategy (e.g. role-based, creative constraints, etc.) ( $M = 9.18$ ,  $SD = 3.77$ ), compared to those who directly copied the creative ideation prompt, ( $M = 8.01$ ,  $SD = 3.13$ ),  $t(32) = 1.016$ ,  $p = 0.317$ , 95% CI [-3.51, 1.17]. We also observed no difference in the weighted number of unique ideas, (Unusual Prompt:  $M = 7.98$ ,  $SD = 3.12$ , Copied Prompt:  $M = 6.82$ ,  $SD = 3.12$ ),  $p = .278$ , 95% CI [-3.29, .98], or in the homogeneity of their ideas between participants who took these different approaches, (Unusual Prompt:  $M = .99$ ,  $SD = .03$ , Copied Prompt:  $M = .99$ ,  $SD = .03$ ),  $t(32) = .091$ ,  $p = 0.927$ , 95% CI [-.02, .02].

## 5 DISCUSSION

Overall, we found that—in line with to our expectation of LLM-induced homogenization—ideas generated with assistance from ChatGPT were significantly less semantically diverse at the group level than ideas generated with assistance from the non-AI-based CST. When supported by ChatGPT, participants produced a larger number of ideas, but the set of ideas generated by each individual participant was similarly diverse in the ChatGPT and non-ChatGPT conditions, suggesting that the increase in quantity of ideas generated did not result in a commensurate increase in diversity as might be expected if diversity is a linear function of quantity [80].

We interpret our results as supporting three key takeaways. First, we find that homogenization stems not from individual-level increases in fixation when working with the LLM, but from group-level suggestion of similar ideas to different users by the LLM. Second, we suggest that the homogenization effect of LLMs on creative ideation is attributable in part to the low inferential distance between LLM outputs and apparently complete ideas. Third, we infer that LLM users might be able to resist the homogenization effect if they are given information about what kinds of stereotyped outputs the LLM tends to produce in a particular creative context, though it may still be difficult for them to adapt their prompting strategies to elicit more diverse responses from the LLM. We discuss these and several other takeaways below.

### 5.1 Creative Fixation

Creative fixation [2, 21, 42, 71] occurs when people engaged in a creative activity become incapable of seeing past an inappropriately narrow assumption that they have made about the conceptual space in which they are working. If LLMs induced creative fixation (for instance by presenting users with ideas to which they cannot readily envision compelling alternatives), we would expect to see individual-level homogenization brought about by fixation: decreased diversity within the set of ideas proposed by each individual user, due to users becoming fixated on specific subsets of the space of all possible ideas. We did not witness this in practice, suggesting that LLMs do not increase homogenization by causing or worsening creative fixation.

This remains the case despite the fact that some participants in our study viewed ChatGPT as a “very fast and accurate” authoritative source that “gave [...] the right answers”, a view which might be expected to provoke user fixation on ideas suggested by the LLM. Even among participants who viewed ChatGPT as an authority, they still generally suggested some ideas unlike those generated by the LLM, perhaps recognizing (as some users explicitly noted) that the LLM tended to “go really specific really quickly” or become “repetitive” in its outputs. Altogether, we conclude that users viewing ChatGPT as authoritative may have led them to *accept* LLM-suggested ideas as valid, but it did not seem to *constrain* the process of ideation by displacing ideas that were somehow incompatible with those suggested by ChatGPT.

### 5.2 Inferential Distance

In planning this study, we considered it a possibility that the Oblique Strategies (OS) deck might cause greater homogenization than ChatGPT due to the fixed nature of the OS cards. The deck consists of

a relatively small set of fixed text strings, and if different users drew the same card in the context of the same ideation prompt, we expected that the similar stimulus might push them to think in similar directions. In practice, however, this did not turn out to be the case. We believe that this might be due to higher *inferential distance* [43] between the text on the OS cards and the expected form of the task responses: because the user has to do more work to interpret the text of a particular OS card in relation to a particular ideation prompt (compared to text generated by ChatGPT in response to that prompt), the OS deck creates more room within the overall ideation process for the user’s psychological uniqueness to influence the eventual output ideas.

Subjective sense of responsibility data further substantiate this interpretation. Compared to Oblique Strategies users, ChatGPT users reported feeling less responsible for the ideas they produced, and several mentioned their own felt sense of non-responsibility for ideas produced with ChatGPT in post-test interview responses. In the words of one participant, “this was a creativity task, and it felt like I should have some stake in the answer”. For at least some users, the apparent completeness of ChatGPT responses resulted in users feeling disengaged from the process of idea generation; according to another participant, “ChatGPT allowed me to turn my brain off” and did “the heavy lifting” during ideation. However, we note that even disengaged users typically did not simply copy-paste ChatGPT output unexamined into their list of ideas: we only observed one participant in the entire study directly copy-paste the entirety of a ChatGPT response without any further curation or modification, suggesting that even low-engagement users may be experiencing something closer to “inspiration through observation” [74] than to total subsumption of the ideation process by the CST.

From one perspective, low inferential distance between CST outputs and finished-looking artifacts can be viewed as enabling *algorithmic loafing* [39]: simply accepting and passing along the decisions made by the algorithm, regardless of whether these decisions agree with the decisions the user might have made without the algorithm’s input. Consequently, one way to mitigate the homogenization effect of LLM-based CSTs may be to design the CST to output deliberately oblique or gnomic responses, analogous to those printed on Oblique Strategies cards: outputs that can provoke user ideation in potentially unexpected directions, but cannot be employed as straightforward substitutes for user ideas (instead requiring a degree of user interpretation before they can be used). In this way, future LLM-based CSTs may be able to reintroduce some of the underdetermination [1] that has proven valuable in other computationally engaged creative processes.

### 5.3 User Adaptation to Homogenization

The absence of individual-level homogenization in our study suggests that participants were able to gauge the diversity of their own responses and continue producing ideas until they felt the set of ideas they have produced is sufficiently diverse. This may be because participants can see all of their own ideas, so they can actively modulate the diversity of the set of ideas that they have generated—for instance by generating new ideas (either on their own, by prompting ChatGPT and curating the output, or both) until they are satisfied with the level of diversity in the resulting idea set



overall. However, participants cannot easily gauge the similarity or dissimilarity of their own responses to other people’s responses, so they cannot actively intervene to modulate diversity at the group level. The ability of participants to observe and work around homogenization at the individual level is backed up to some extent by the fact that a few ChatGPT users explicitly noted the repetitiveness of the model’s responses. “When I want more”, one user asserts, “I want *different* more.”

The apparent success of individual-level user adaptation to homogenization suggests that one LLM-based CST design strategy for mitigating homogenization may involve detecting and alerting the user to clichés within model output, as has been suggested before [51]. By making it apparent to users when a particular piece of model output is semantically similar to outputs that other users have previously encountered, it may be possible to give users more information about the typicality of model output, and thereby to help them notice and resist model-induced homogenization at the group level as well.

More generally, users may also be able to adapt to homogenization in the real world via adoption of and horizontal movement between a variety of different CSTs, each of which imposes a different “normative ground” on users and their creative inputs [56]. To mitigate homogenization, CST designers might thus seek to create a plurality of different tools between which users can move, rather than a small number of monolithic, one-size-fits-all tools.

#### 5.4 Mitigating Homogeneity of LLM Outputs

Because the homogenization effect of ChatGPT seems to stem largely from the LLM generating similar responses to different users’ queries in the same creative context, one obvious approach to reducing homogenization involves getting LLMs (or LLM-based CSTs) to output more diverse responses. Several strategies for achieving LLM output diversity might be implementable in the future.

At the CST user level, it may seem tempting to suggest that users of current LLMs can resist homogenization by implementing more sophisticated prompting strategies, such as deliberately designing prompts that incorporate some aspect of the user’s unique perspective (to differentiate one’s own prompt from the prompts one expects others to use in similar contexts). However, our analysis of prompting process (Section 4.5.4) found that—although different ChatGPT users in our study did employ different prompting strategies—there were no reliable differences in homogenization to be found between users of different approaches to prompting. Additionally, LLM prompt design is difficult for novice users [90], and changes in prompt structure can dramatically alter both the relative and absolute performance of different LLMs on a single task [60], suggesting that what makes for a good LLM prompt is likely opaque to users in general. Consequently, user adoption of more sophisticated prompting strategies may not yield reliable improvements in terms of group-level homogenization without additional interventions at the level of CST design or model development.

At the CST design level, injecting randomness into LLM prompts (for instance, by drawing stimuli at random from a large pool of potential stimuli—such as the Oblique Strategies deck—and then instructing the LLM to consider these random stimuli in responding to the user’s input) might serve as a stopgap approach to increasing

output diversity when different users are expected to input very similar prompts. However, due again to the complexity of prompting, thorough testing would need to be undertaken of any such prompt-level randomness injection strategy to ensure that it does in fact lead to increased output diversity.

Meanwhile, at the model level, strategies such as quality-diversity optimization (e.g., QDAIF [11]) and diverse decoding methods [40, 77] may be used to improve the diversity of LLM responses. Compared to prompt-level randomness injection, these algorithmic strategies attempt to achieve output diversity more directly, and can thus be expected to yield more reliably diverse results in a wide variety of different creative contexts.

#### 5.5 Implications for LLM-Based CST Design

Many of our findings follow naturally from the fact that an LLM will tend to produce similar outputs in response to similar inputs, even when those inputs come from different users. This property of LLMs is often desirable, especially when the goal is to build systems that respond reliably to queries in a particular domain. However, the fact that LLM outputs are strongly reflective of direct user inputs raises a key question for designers of LLM-based CSTs: when users are in the early stages of a creative process and their creative intent is still relatively undefined, how can an LLM-based tool avoid guiding them toward well-trodden ground?

A vaguely defined creative intent, expressed as (for instance) a short text string containing only a handful of words, will often tend—for level-of-detail reasons alone—to be similar to vaguely defined intents expressed by other users. For an LLM to directly transform a brief expression of intent into a substantially larger artifact (e.g., a passage of text), it will necessarily have to make a large number of creative decisions (e.g., individual word choices) on the user’s behalf, and these decisions will be made by reference to the relative likelihood of different alternatives. Brief statements of creative intent therefore yield “samey” artifacts, which exhibit a phenomenon characterized elsewhere as the *dearth of the author* [45]. In other words, to get highly original outputs from an LLM, the user must supply the LLM with an input that contains enough detail to strongly differentiate it from other users’ inputs; the LLM itself functions most directly as a means of refracting user-supplied inputs through the lens of statistical patterns learned from its training data.

In light of this limitation, we believe that it may be necessary to move toward a view of creativity support as *intent elicitation* [46]: progressively drawing a more and more detailed, idiosyncratic specification of creative intent out of the CST user, with the goal of getting users to think *more* rather than *less* about each creative decision that they make along the way. By scaffolding the multi-turn discovery and expression of creative intent, we hope that LLM-based CSTs can eventually operate as something like a focusing lens for the imagination: a means of gradually clarifying and magnifying the spark of originality that first emerges in the user’s head, without substituting the LLM’s implicitly statistical and majoritarian process of creative decision-making for the user’s own. Whether this design vision can be effectively realized remains to be seen, but we believe it to be an especially worthwhile goal for the next generation of AI-based CSTs to pursue.

## 6 LIMITATIONS

Our study was conducted in a lab context and consequently may not reflect some aspects of how people use LLMs for creativity support “in the wild”. In order to keep study sessions to a manageable length, participants were only allowed a fixed and relatively short amount of time to respond to each ideation prompt; additionally, the ideation prompts in our study were supplied by us, rather than being formulated independently by CST users. Most organic divergent ideation scenarios are not subject to these constraints.

Because our study made use of a general-purpose LLM-based chatbot (ChatGPT) rather than an LLM-based CST specifically designed to support divergent ideation, it should be noted that alternative deployments of LLMs may be able to avoid homogenizing user ideas. Padmakumar and He have found evidence that feedback-tuned LLMs may induce creative homogenization where base (non-feedback-tuned) LLMs do not [69], and task-specific CST design has previously shown promise for reducing conformity effects in crowdsourced divergent ideation [89]. These findings suggest that interventions at both the AI and UI layers may help to mitigate homogenization effects in future LLM-based CSTs.

## 7 CONCLUSION

We have presented evidence that LLM-based CSTs exert a stronger homogenization effect on human-in-the-loop divergent ideation processes than at least some plausible alternative CSTs. We have further clarified that the homogenization effect is group-level rather than individual-level, and that it may be partly attributable to low inferential distance between LLM outputs and apparently finished ideas. Coupled with evidence that ChatGPT users exhibit greater fluency, flexibility, and elaboration than users of an alternative CST, these results suggest that current general-purpose instruction-tuned LLMs (such as ChatGPT) are capable of functioning as useful CSTs by enabling the rapid enumeration of relatively obvious possibilities that users might otherwise fail, or take longer, to consider. However, these systems are not currently well-suited to helping users develop truly original ideas.

We believe that the style of homogenization analysis employed here is suitable for wider adoption as a technique for the evaluation of CSTs—including, but not limited to, AI-based CSTs—in the future. By providing both evidence for the existence of and proposed mechanisms of action for creative homogenization effects of AI-based CSTs, we also hope to stimulate future work on the adaptation of AI-based CSTs to mitigate these effects. Data-driven AI technologies clearly have a role to play in creative ideation processes, but interventions at both the CST design and model development levels may be required to realize the full potential of these technologies for creativity support.

## ACKNOWLEDGMENTS

This work was supported by Hackworth Grant GR102981, “Investigating Homogenization of Imagination by Generative AI Models”, from the Markkula Center for Applied Ethics.

## REFERENCES

- [1] 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.

- [2] 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.
- [3] Alberto Alvarez, Jose Font, and Julian Togelius. 2022. Toward Designer Modeling Through Design Style Clustering. *IEEE Transactions on Games* 14, 4 (2022), 676–686.
- [4] 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.
- [5] Ilhan Aslan, Katharina Weitz, Ruben Schlagowski, Simon Flutura, Susana Garcia Valesco, Marius Pfeil, and Elisabeth André. 2019. Creativity support and multimodal pen-based interaction. In *2019 International Conference on Multimodal Interaction*. 135–144.
- [6] Roger E Beaty and Dan R Johnson. 2021. Automating creativity assessment with SemDis: An open platform for computing semantic distance. *Behavior Research Methods* 53, 2 (2021), 757–780.
- [7] Nina Begus. 2023. Experimental Narratives: A Comparison of Human Crowdsourced Storytelling and AI Storytelling. *arXiv preprint arXiv:2310.12902* (2023).
- [8] 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.
- [9] Margaret A Boden. 2004. *The Creative Mind: Myths and Mechanisms*. Routledge.
- [10] Rishi Bommasani, Kathleen A Creel, Ananya Kumar, Dan Jurafsky, and Percy S Liang. 2022. Picking on the Same Person: Does Algorithmic Monoculture lead to Outcome Homogenization? In *Advances in Neural Information Processing Systems*, Vol. 35. 3663–3678.
- [11] 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).
- [12] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language Models are Few-Shot Learners. In *Advances in Neural Information Processing Systems*, Vol. 33. 1877–1901.
- [13] 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*.
- [14] Alex Calderwood, Vivian Qiu, Katy Ilonka Gero, and Lydia B Chilton. 2020. How Novelists Use Generative Language Models: An Exploratory User Study. In *Workshop on Human-AI Co-Creation with Generative Models (HAI-GEN 2020)*.
- [15] 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).
- [16] Kathy Charmaz. 2006. *Constructing Grounded Theory: A Practical Guide through Qualitative Analysis*. Sage.
- [17] 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).
- [18] John Joon Young Chung and Eytan Adar. 2023. PromptPaint: Steering Text-to-Image Generation Through Paint Medium-like Interactions. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*.
- [19] John Joon Young Chung, Shiqing He, and Eytan Adar. 2021. The intersection of users, roles, interactions, and technologies in creativity support tools. In *Designing Interactive Systems Conference 2021*. 1817–1833.
- [20] 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*.
- [21] Nathan Crilly. 2019. Creativity and fixation in the real world: A literature review of case study research. *Design Studies* 64 (2019), 154–168.
- [22] Nicholas Davis, Chih-Pin Hsiao, Kunwar Yashraj Singh, Brenda Lin, and Brian Magerko. 2017. Creative sense-making: Quantifying interaction dynamics in co-creation. In *Proceedings of the 2017 ACM SIGCHI Conference on Creativity and Cognition*. 356–366.
- [23] Nicholas Davis, Alexander Zook, Brian O’Neill, Brandon Headrick, Mark Riedl, Ashton Grosz, and Michael Nitsche. 2013. Creativity support for novice digital filmmaking. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 651–660.
- [24] Alexandra Domrongchai. 2023. Can ChatGPT Make Me a Better Cook? <https://www.foodandwine.com/can-chatgpt-make-meal-planning-easier-7368840>. Accessed: 2023-11-11.

- [25] Anil R Doshi and Oliver Hauser. 2023. Generative artificial intelligence enhances creativity. *Available at SSRN* (2023).
- [26] Denis Dumas, Peter Organisciak, and Michael Doherty. 2021. Measuring divergent thinking originality with human raters and text-mining models: A psychometric comparison of methods. *Psychology of Aesthetics, Creativity, and the Arts* 15, 4 (2021), 645.
- [27] Denis Dumas, Peter Organisciak, Shannon Maio, and Michael Doherty. 2021. Four text-mining methods for measuring elaboration. *The Journal of Creative Behavior* 55, 2 (2021), 517–531.
- [28] Brian Eno and Peter Schmidt. 1975. Oblique Strategies, hand produced deck of cards. <http://www.rtqe.net/ObliqueStrategies/>
- [29] Ziv Epstein, Aaron Hertzmann, and Investigators of Human Creativity. 2023. Art and the science of generative AI. *Science* 380, 6650 (2023), 1110–1111.
- [30] Umer Farooq, John M Carroll, and Craig H Canoe. 2008. Designing for creativity in computer-supported cooperative work. *International Journal of e-Collaboration (IJEC)* 4, 4 (2008), 51–75.
- [31] 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*.
- [32] William W Gaver, Jacob Beaver, and Steve Benford. 2003. Ambiguity as a resource for design. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 233–240.
- [33] 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*.
- [34] Agnes Guenther, Boris Eisenbart, and Andy Dong. 2017. Creativity as a way to innovate successfully. In *DS 87-8 Proceedings of the 21st International Conference on Engineering Design (ICED 17) Vol 8: Human Behaviour in Design, Vancouver, Canada, 21-25.08.2017*, 389–398.
- [35] Matthew Guzdial, Nicholas Liao, Jonathan Chen, Shao-Yu Chen, Shukan Shah, Vishwa Shah, Joshua Reno, Gillian Smith, and Mark O Riedl. 2019. Friend, collaborator, student, manager: How design of an AI-driven game level editor affects creators. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*.
- [36] 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.
- [37] Sandra G Hart and Lowell E Staveland. 1988. Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research. In *Advances in Psychology*. Vol. 52. Elsevier, 139–183.
- [38] 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.
- [39] 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.
- [40] 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.
- [41] 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*.
- [42] David G Jansson and Steven M Smith. 1991. Design fixation. *Design Studies* 12, 1 (1991), 3–11.
- [43] David Kirsh. 2011. Creative cognition in choreography. In *Proceedings of the 2nd International Conference on Computational Creativity*.
- [44] Jon Kleinberg and Manish Raghavan. 2021. Algorithmic monoculture and social welfare. *Proceedings of the National Academy of Sciences* 118, 22 (2021), e2018340118.
- [45] Max Kreminski. 2024. The Dearth of the Author in AI-Supported Writing. *arXiv preprint arXiv:2404.10289* (2024).
- [46] Max Kreminski and John Joon Young Chung. 2024. Intent Elicitation in Mixed-Initiative Co-Creativity. In *Joint Proceedings of the ACM IUI Workshops*.
- [47] 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.
- [48] Max Kreminski, Melanie Dickinson, Noah Wardrip-Fruin, and Michael Mateas. 2022. 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.
- [49] Max Kreminski, Isaac Karth, Michael Mateas, and Noah Wardrip-Fruin. 2022. Evaluating mixed-initiative creative interfaces via expressive range coverage analysis. In *IUI Workshops*. 34–45.
- [50] Max Kreminski and Chris Martens. 2022. Unmet creativity support needs in computationally supported creative writing. In *Proceedings of the First Workshop on Intelligent and Interactive Writing Assistants (In2Writing 2022)*. 74–82.
- [51] Max Kreminski and Michael Mateas. 2021. Reflective creators. In *International Conference on Computational Creativity*. 309–318.
- [52] Max Kreminski and Noah Wardrip-Fruin. 2019. Generative games as storytelling partners. In *Proceedings of the 14th International Conference on the Foundations of Digital Games*.
- [53] Mina Lee, Katy Ilonka Gero, John Joon Young Chung, Simon Buckingham Shum, Vipul Raheja, Hua Shen, Subhashini Venugopalan, Thimo Wambsgans, 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.
- [54] 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*.
- [55] Isabelle Levent and Lila Shroff. 2023. The Model is the Message. In *The Second Workshop on Intelligent and Interactive Writing Assistants*.
- [56] 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*.
- [57] Q Vera Liao and S Shyam Sundar. 2022. Designing for responsible trust in AI systems: A communication perspective. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*. 1257–1268.
- [58] Roberta M Milgram and Norman A Milgram. 1976. Creative thinking and creative performance in Israeli students. *Journal of educational psychology* 68, 3 (1976), 255.
- [59] 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*.
- [60] 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).
- [61] Ethan Mollick. 2023. On-boarding your AI Intern. <https://www.oneusefulthing.org/p/on-boarding-your-ai-intern>. Accessed: 2023-11-11.
- [62] Richard D Morey. 2008. Confidence intervals from normalized data: A correction to Cousineau (2005). *Tutorials in Quantitative Methods for Psychology* 4, 2 (2008), 61–64.
- [63] Niklas Muennighoff, Nouamane Tazi, Loic Magne, and Nils Reimers. 2023. MTEB: Massive Text Embedding Benchmark. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*. Association for Computational Linguistics, 2014–2037. <https://doi.org/10.18653/v1/2023.eacl-main.148>
- [64] Matthew J Neagley. 2023. Campaign Prep with ChatGPT. <https://gnomestew.com/campaign-prep-with-chatgpt/>. Accessed: 2023-11-11.
- [65] Arvind Neelakantan, Tao Xu, Raul Puri, Alec Radford, Jesse Michael Han, Jerry Tworek, Qiming Yuan, Nikolas Tezak, Jong Wook Kim, Chris Hallacy, Johannes Heidecke, Pranav Shyam, Boris Power, Tyna Eloundou Nekoul, Girish Sastry, Gretchen Krueger, David Schnurr, Felipe Petroski Such, Kenny Hsu, Madeleine Thompson, Tabarak Khan, Toki Sherbakov, Joanne Jang, Peter Welinder, and Lilian Weng. 2022. Text and Code Embeddings by Contrastive Pre-Training. *arXiv preprint arXiv:2201.10005* (2022).
- [66] Dmitry Nikolaev and Sebastian Padó. 2023. Representation biases in sentence transformers. *arXiv preprint arXiv:2301.13039* (2023).
- [67] OpenAI. 2023. ChatGPT: A Large-Scale Conversational Language Model. <https://www.openai.com/research/chatgpt>. Accessed: August 3rd, 2023.
- [68] 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.
- [69] Vishakh Padmakumar and He He. 2023. Does Writing with Language Models Reduce Content Diversity? *arXiv preprint arXiv:2309.05196* (2023).
- [70] Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. GloVe: Global Vectors for Word Representation. In *Empirical Methods in Natural Language Processing (EMNLP)*. 1532–1543. <http://www.aclweb.org/anthology/D14-1162>
- [71] A Terry Purcell and John S Gero. 1996. Design and other types of fixation. *Design Studies* 17, 4 (1996), 363–383.

- [72] 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>
- [73] 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.
- [74] Melissa Roemmele. 2021. Inspiration through observation: Demonstrating the influence of automatically generated text on creative writing. In *Proceedings of the 12th International Conference on Computational Creativity (ICCC '21)*.
- [75] Mark A Runco, Shawn M Okuda, and Becky J Thurston. 1987. The psychometric properties of four systems for scoring divergent thinking tests. *Journal of Psychoeducational Assessment* 5, 2 (1987), 149–156.
- [76] Matt Ruten. 2012. Oblique Strategies App. <https://obliquestrategies.ca>
- [77] 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>
- [78] Ben Shneiderman. 2007. Creativity support tools: accelerating discovery and innovation. *Commun. ACM* 50, 12 (2007), 20–32.
- [79] Paul J Silvia, Beate P Winterstein, John T Willse, Christopher M Barona, Joshua T Cram, Karl I Hess, Jenna L Martinez, and Crystal A Richard. 2008. Assessing creativity with divergent thinking tasks: exploring the reliability and validity of new subjective scoring methods. *Psychology of Aesthetics, Creativity, and the Arts* 2, 2 (2008), 68.
- [80] Dean Keith Simonton. 2010. Creative thought as blind-variation and selective-retention: Combinatorial models of exceptional creativity. *Physics of Life Reviews* 7, 2 (2010), 156–179.
- [81] Nikhil Singh, Guillermo Bernal, Daria Savchenko, and Elena L Glassman. 2023. Where to hide a stolen elephant: Leaps in creative writing with multimodal machine intelligence. *ACM Transactions on Computer-Human Interaction* 30, 5 (2023).
- [82] Sangho Suh, Meng Chen, Bryan Min, Toby Jia-Jun Li, and Haijun Xia. 2023. Structured Generation and Exploration of Design Space with Large Language Models for Human-AI Co-Creation. *arXiv preprint arXiv:2310.12953* (2023).
- [83] Yuqian Sun, Xingyu Li, Jun Peng, and Ze Gao. 2023. Inspire Creativity with ORIBA: Transform Artists’ Original Characters into Chatbots through Large Language Model. In *Adjunct Proceedings of the 2023 ACM International Joint Conference on Pervasive and Ubiquitous Computing & the 2023 ACM International Symposium on Wearable Computing (UbiComp/ISWC '23 Adjunct)*. 78–82. <https://doi.org/10.1145/3594739.3610695>
- [84] S Shyam Sundar and Jinyoung Kim. 2019. Machine heuristic: When we trust computers more than humans with our personal information. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*.
- [85] E Paul Torrance. 1966. Torrance tests of creative thinking. *Educational and Psychological Measurement* (1966).
- [86] Qian Wan, Siying Hu, Yu Zhang, Piaohong Wang, Bo Wen, and Zhicong Lu. 2023. “It Felt Like Having a Second Mind”: Investigating Human-AI Co-creativity in Prewriting with Large Language Models. *arXiv preprint arXiv:2307.10811* (2023).
- [87] Qian Wan and Zhicong Lu. 2023. GANCollage: A GAN-Driven Digital Mood Board to Facilitate Ideation in Creativity Support. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference*. 136–146.
- [88] Chuan Yan, John Joon Young Chung, Yoon Kiheon, Yotam Gingold, Eytan Adar, and Sungsoo Ray Hong. 2022. FlatMagic: Improving flat colorization through AI-driven design for digital comic professionals. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*.
- [89] Lixiu Yu, Aniket Kittur, and Robert E Kraut. 2014. Distributed analogical idea generation: inventing with crowds. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 1245–1254.
- [90] 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*.

## A VALIDATING SENTENCE EMBEDDINGS FOR HOMOGENIZATION ANALYSIS

Our primary homogenization analysis uses a transformer-based sentence embedding model—all-MiniLM-L6-v2, one of the standard general-purpose sentence embedding models provided by the Python SentenceTransformers library [72]—to evaluate the semantic similarity between participant ideas expressed as short strings

of text. Our methodology here is similar to that employed in several other recent psychological studies of creativity. Broadly speaking, semantic similarity approaches to creativity research involve the use of some algorithm to produce a numeric score representing the similarity of a pair of creative artifacts; the originality of multiple different artifacts can then be determined and compared relative to a fixed reference point. The SemDis platform [6], a key example of this approach, automates large-scale determination of semantic similarity scores between ideas (expressed as short strings of text) and the creative ideation prompt in response to which these ideas were generated; similarity scores are determined by means of cosine similarity between aggregated word embeddings [70], a metric which has been found to agree well with human judgments of semantic similarity [6, 26].

Aggregated word embeddings are generally outperformed on semantic similarity tasks by more recent transformer-based sentence embeddings, which (unlike aggregated word embeddings) are able to take sentence structure into account. However, the rapid pace of progress in machine learning research means that transformer-based sentence embeddings have not yet been validated against human judgments of semantic similarity in the context of creativity research specifically. Therefore, in order to validate our use of all-MiniLM-L6-v2, we conducted a small experiment to determine whether this model agrees strongly with human judgments of semantic similarity on our participant ideas dataset.

Our experiment took the human-constructed idea categories produced by our flexibility analysis (Section 4.3.2) as a source of ground truth for semantic similarity judgments and evaluated several candidate embedding models in terms of their agreement with the human coders’ manual classification of ideas. First, for each category of ideas in our dataset, we produced a category embedding by averaging together the individual embeddings of the ideas belonging to this category. We then iterated over each idea in the dataset, sorted the category embeddings by their cosine similarity to the embedding of the idea being categorized, and assigned the idea to the  $n$  categories represented by the  $n$  most similar category embeddings (where  $n$  = the number of categories human coders assigned to this idea). To avoid producing artificially high similarity scores between ideas and their actual human-assigned categories across the board, we also excluded each idea’s own embedding from the average category embedding when testing similarity to the idea’s actual human-assigned categories.

We then compared the model-assigned categories for each idea to the actual categories human coders assigned to this idea, and noted the percentage of overlap between these category sets. Finally, we repeated this process for several different embedding models—as well as a pessimistic baseline “model” that assigned each idea to  $n$  categories at random—and computed the human-agreement percentage of each model on participant ideas generated in response to each of our four creativity tasks.

Results are reported in Table 2. Notably, our chosen sentence embedding model (all-MiniLM-L6-v2) agrees with human idea categorizations more than half the time across all four creativity tasks; it therefore substantially outperforms both GloVe 840B (an aggregate word embedding model previously assessed as state-of-the-art for creativity research [6, 26]) and the random baseline (which GloVe itself beats by more than an order of magnitude). It

Model	IC_A	IC_B	PI_A	PI_B	Average
all-MiniLM-L6-v2	60.94%	58.61%	51.48%	64.63%	58.92%
all-mpnet-base-v2	58.91%	57.08%	51.30%	59.78%	56.77%
GloVe 840B	46.59%	47.25%	41.16%	40.02%	43.76%
Random	4.31%	2.37%	4.14%	2.94%	3.44%

**Table 2: Percentage agreement of several different embedding models with human idea categorization judgments. Columns IC\_A, IC\_B, PI\_A, and PI\_B report performance on ideas generated in response to specific ideation tasks (IC = “Improbable Consequences”, PI = “Product Improvement”).**

also consistently outperforms all-mpnet-base-v2, the theoretically best overall general-purpose SentenceTransformers model, by a small margin.

We did not directly evaluate OpenAI’s sentence embedding models [65], despite the use of an unspecified OpenAI embedding model in another study of creative homogenization [25], because OpenAI models (unlike the SentenceTransformers models) are proprietary: they cost money to use, cannot be run locally, and are not open source. This limits the replicability of research results derived from

OpenAI models and therefore the suitability of these models for research pipelines. However, due to the prior use of one of these models in a similar research context, a comparison of some sort is nevertheless merited. For this, we consult the Massive Text Embedding Benchmark [63], which shows that the most powerful OpenAI embeddings model (text-embedding-ada-002) is not clearly much better or worse than a variety of SentenceTransformers models (including all-MiniLM-L6-v2) across a range of tasks.

Sentence embedding models remain imperfect arbiters of semantic similarity. In our categorization experiment, even the best-performing embeddings model achieved only 59% agreement with human coders on average. There also exists some evidence that cosine similarity between sentence embeddings is more strongly influenced by overlap in the set of nouns than by other similarities [66], suggesting that these models do not take all of the nuances of sentence meaning into account when computing similarity scores. However, the agreement between models like all-MiniLM-L6-v2 and human judgments of semantic similarity strike us as high enough to justify the use of these sentence embedding models for homogenization analysis in creativity research, in particular for the increased scale of analysis that these models enable.