



One Does Not Simply Meme Alone: Evaluating Co-Creativity Between LLMs and Humans in the Generation of Humor

Zhikun Wu
KTH Royal Institute of Technology
Stockholm, Sweden
zhikun@kth.se

Thomas Weber
LMU Munich
Munich, Germany
thomas.weber@ifi.lmu.de

Florian Müller
TU Darmstadt
Darmstadt, Germany
florian.mueller@tu-darmstadt.de

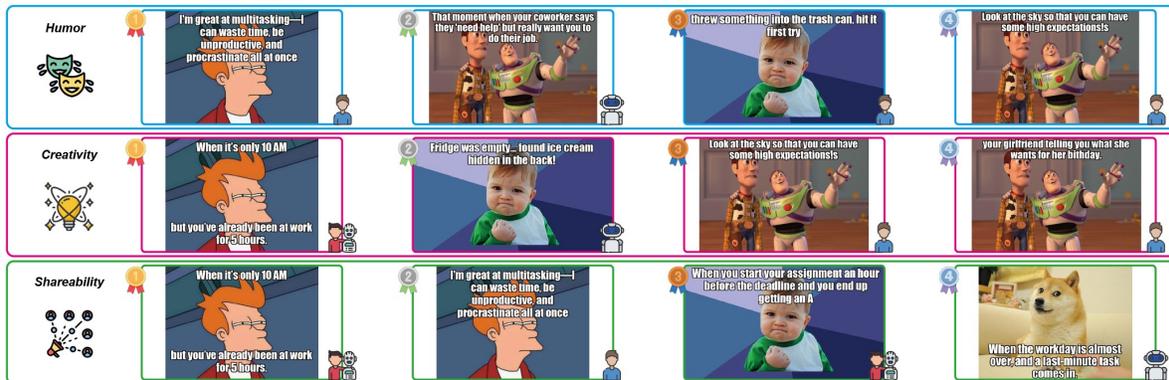


Figure 1: Top 4 Memes Generated by AI, Humans, and Human-AI Collaboration Across Humor, Creativity, and Shareability Metric.

Abstract

Collaboration has been shown to enhance creativity, leading to more innovative and effective outcomes. While previous research has explored the abilities of Large Language Models (LLMs) to serve as co-creative partners in tasks like writing poetry or creating narratives, the collaborative potential of LLMs in humor-rich and culturally nuanced domains remains an open question. To address this gap, we conducted a user study to explore the potential of LLMs in co-creating memes—a humor-driven and culturally specific form of creative expression. We conducted a user study with three groups of 50 participants each: a human-only group creating memes without AI assistance, a human-AI collaboration group interacting with a state-of-the-art LLM model, and an AI-only group where the LLM autonomously generated memes. We assessed the quality of the generated memes through crowdsourcing, with each meme rated on creativity, humor, and shareability. Our results showed that LLM assistance increased the number of ideas generated and reduced the effort participants felt. However, it did not improve the quality of the memes when humans were collaborated with LLM. Interestingly, memes created entirely by AI performed better than both human-only and human-AI collaborative memes in all areas on average. However, when looking at the top-performing memes, human-created ones were better in humor, while human-AI collaborations stood out in creativity and shareability. These

findings highlight the complexities of human-AI collaboration in creative tasks. While AI can boost productivity and create content that appeals to a broad audience, human creativity remains crucial for content that connects on a deeper level.

CCS Concepts

• Human-centered computing → Empirical studies in HCI.

Keywords

Human-AI Collaboration, LLM, Co-Creativity, Memes, Humor

ACM Reference Format:

Zhikun Wu, Thomas Weber, and Florian Müller. 2025. One Does Not Simply Meme Alone: Evaluating Co-Creativity Between LLMs and Humans in the Generation of Humor. In *30th International Conference on Intelligent User Interfaces (IUI '25)*, March 24–27, 2025, Cagliari, Italy. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3708359.3712094>

1 Introduction

From co-authoring articles to discussing vacation plans or pair programming, collaborating with other people is a core part of our daily lives in many ways. While we could perform all these activities on our own, the diverse perspectives and problem-solving strategies [22], increased motivation in groups [22] and continuous feedback [53] help to support creative processes [2]. Prior work showed that such collaborative work increases the performance of project teams [12], improves quality [54], and, in particular, enhances creativity [38] in numerous domains. With the rise of Large Language Models (LLMs), there is a growing trend of replacing human collaboration partners with such systems in typically creative activities in areas such as art [30, 59], music [16, 17, 21], and



This work is licensed under a Creative Commons Attribution 4.0 International License. *IUI '25, Cagliari, Italy*
© 2025 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-1306-4/25/03
<https://doi.org/10.1145/3708359.3712094>

literature [28, 46, 51, 57]. In these areas, LLMs have demonstrated remarkable capabilities in generating content, often matching or even surpassing human performance in tasks like divergent thinking, which involves producing numerous and varied ideas [29].

However, much of the prior work on the creative aspect of LLMs focuses on their outputs as standalone creations, often neglecting their potential as true co-creative partners. In these studies, the researchers presented LLMs with various tasks that typically require creativity and evaluated the result produced by the LLM [23, 25, 26, 29]. These assessments focus on attributes like originality and fluency, where LLMs demonstrate strong performance on metrics such as the Torrance Test of Creative Thinking [25, 29]. While these works offer important insights into the possibilities of such models for creative tasks, such approaches fail to capture the iterative nature of human co-creativity. Unlike simple task delegation, co-creativity involves iterative co-creation, where humans and AI systems actively refine ideas through dialog and feedback loops, aligning with established frameworks of collaborative creativity [7, 38]. Recently, research has started to investigate this co-creative process in the joint creative work of humans and LLMs in a number of domains [27, 28, 35, 49] and found that LLMs can contribute novel suggestions, enrich human ideas and enhance the joint creative output [4, 36].

Additionally, one important aspect has not yet been investigated in the area of co-creative collaboration with LLMs: humor. Humor is an interesting area because it is one of the most sophisticated and complex forms of human creativity. Humor strengthens social bonds, addresses difficult topics, and offers new perspectives on everyday situations [3]. It relies on surprise, contrast, cultural context, and emotional resonance [45]. What we think is funny depends on our personal cultural, linguistic and political background [15]. With the rise of online social networks and increasing globalization, parts of this context such as pop culture are aligning across borders, resulting in humor that incorporates elements that are personal and local as well as elements that are globally valid [44]. A well-known example of such humor, that incorporates global and local contexts, are internet memes [1], often in the form of captioned images. Such memes, as a cultural phenomenon, have emerged as a universal language of the internet and are used to express emotions, convey messages or appropriation and recontextualization familiar elements [55]. This also made them relevant for research as a means of evaluating creative humor [48]. However, while prior studies have explored the autonomous generation of memes by LLMs [50], there exists no prior work in examining how human-AI collaboration with a multimodal LLM affects the creativity, humor, and shareability of memes.

In this paper, we add to the body of work on human-AI co-creativity by exploring the potential of LLMs as co-creative partners for generating humor. While ‘collaboration’ in HCI traditionally entails shared goals and mutual interdependence between multiple human collaborators [7, 38], we follow prior work that defines co-creativity in terms of iterative, dialog-based refinement of ideas [24, 59]. In such collaborative systems, the AI can serve as a valuable, but not necessarily equal, partner in the creative process. Consequently, we investigate how people interact with a “humor assistant” in the creation of internet memes and how the availability of such an assistant affects productivity. Additionally,

we evaluate how memes generated with such an assistant compare to memes that were created purely by a human and purely by AI in terms of their humor and shareability. For this, we conducted two user studies: in the first, we asked participants to generate ideas for memes, either collaboratively, using an LLM, or without any assistance, and rate the experience. Our study showed that participants who worked with the LLM assistant generated more ideas during the meme creation process than those who worked independently while perceiving the process as less arduous. This first study also yielded 335 images from the human-only group and 307 from the collaborative group which serve as the basis for the subsequent evaluation. For it, we sampled 150 images from each group, as well as 150 fully AI generated images, and asked a second group of participants to rate them in terms of how funny they considered them, how creative, and how likely it would be that they would share them. This evaluation showed that memes generated entirely by AI surpassed both human and human-AI collaborations in all three dimensions.

These findings indicate that while LLM assistance can significantly boost productivity and reduce perceived effort in creative tasks, it does not necessarily enhance the quality of creative output. This suggests that AI models, by drawing from vast datasets, are quite adept at producing content that appeals to a wide audience. However, many of the top-rated memes were created with human involvement, suggesting that AI models primarily produce solid but average quality, while human input can help to iterate and curate on it to lift it to a higher level of quality.

These results emphasize how AI can be a tool to quickly and easily produce large volumes of ideas that can often already meet a broad, average appeal. However, it also demonstrates a need for better methods, tools and processes for integrating AI into an iterative creative process where the AI may produce quantity while the humans acts as a curator that selectively pushes towards the AI towards better results. Designing smarter and more human-centered AI systems that can simplify the challenging steps of the creative process while enriching and amplifying humans’ unique creative abilities will continue to be an important challenge in the future.

2 Related Work

Our research is informed by prior work on human co-creativity, the complexity of humor, human-AI collaboration in creativity, LLMs in creative content generation and the evaluation of creative outputs.

2.1 Human Co-Creativity and the Complexity of Humor

Collaborative creativity between humans has long been recognized as a powerful means of enhancing the creative process [7]. Co-creativity allows individuals to combine diverse perspectives, skills, and ideas, leading to more innovative and high-quality outcomes [2, 38]. In group settings, social interactions can stimulate creativity by fostering motivation, providing immediate feedback, and encouraging risk-taking [22, 53].

Humor, as a sophisticated and complex form of human creativity, plays a significant role in social bonding and communication [39]. Creating humor is particularly challenging because it relies on timing, cultural context, shared knowledge, and the ability to subvert

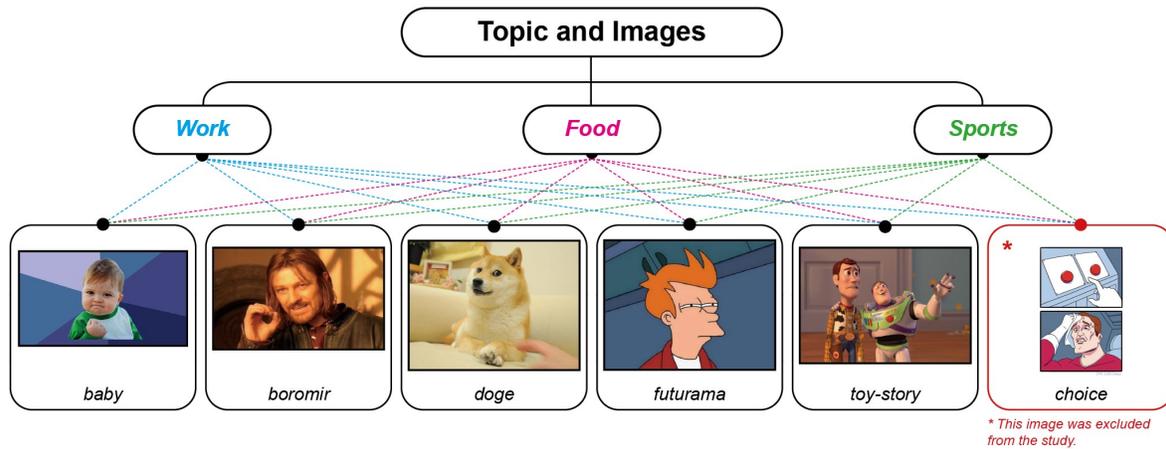


Figure 2: Mapping of Meme Templates to Topics (Work, Food, Sports) in the Study.

expectations [52]. What individuals find humorous is deeply influenced by their personal experiences, cultural backgrounds, and social environments [15]. The complexity of humor makes it a rich area for exploring the dynamics of co-creativity, as collaborators must navigate these nuances to produce content that resonates with others [32].

2.2 Human-AI Collaboration in Creativity

Human-AI collaboration involves humans working alongside AI systems to co-create content by leveraging the strengths of both [47, 56]. In creative fields, this collaboration can enhance human creativity by providing novel ideas and alternative perspectives [24]. For instance, in collaborative writing, AI tools generate alternative text suggestions, acting as a "second mind" to stimulate divergent thinking [49]. In design, systems like the Creative Sketching Partner offer visual stimuli to help designers overcome fixation and explore new directions [31].

In the realm of meme creation, however, there is limited research on human-AI collaborative processes. While LLMs can assist in generating meme content, the interplay between human creativity and AI-generated suggestions remains underexplored. Studies in other creative domains suggest that human-AI collaboration can lead to more creative outputs than either humans or AI alone [28]. Yet, challenges persist, such as managing the AI's lack of contextual sensitivity and ensuring that the collaboration enhances rather than hinders the creative process [43].

Moreover, ethical considerations arise in human-AI co-creation, including issues of authorship and bias introduced during AI training [11, 43]. Understanding how humans interact with AI systems in creative tasks is crucial for designing tools that effectively support and enhance human creativity.

2.3 LLMs in Creative Meme Generation

Large Language Models (LLMs) have demonstrated capabilities in generating human-like text, enabling applications in various creative domains [41]. Recent studies have explored the use of LLMs

in autonomous creative content generation, including narratives [46, 51, 57], humor [58], and particularly memes [50].

In the context of meme generation, MemeCraft [50] utilizes LLMs to produce stance-driven memes with minimal human intervention, showcasing the models' ability to create contextually rich, multimodal content. However, while LLMs can generate humorous and contextually appropriate memes, they often face challenges in capturing nuanced cultural references and emotional subtleties inherent in human creativity [20, 33]. Studies have indicated that LLMs may produce homogenized content, lacking diversity and originality compared to human-generated content [6, 18].

Despite these limitations, LLMs have been found to outperform humans in certain divergent thinking tasks, exhibiting remarkable originality and fluency [25, 29]. In humor generation, some LLMs produce jokes rated comparably to human-written humor [23, 26], although they may still fall short in capturing the depth of human humor in various contexts [20]. These findings highlight both the potential and the challenges of using LLMs for autonomous creative tasks such as meme generation.

2.4 Evaluation Metrics for Creative Outputs

Evaluating creative outputs, such as memes, involves assessing aspects like creativity, humor, and shareability [48, 55]. Memes are unique cultural artifacts that blend visual and textual elements to convey messages resonating with diverse audiences [10]. The shareability of a meme reflects its potential to be widely circulated, influenced by factors like humor, relatability, and relevance to current cultural topics [34, 40].

Humor is a key driver of engagement and virality in memes, often relying on incongruity and the juxtaposition of unexpected elements [37, 45]. Memes that effectively utilize humor can facilitate social bonding and amplify sociopolitical discourse [3]. Creativity, encompassing originality and novelty, is critical in making memes stand out in the vast online content landscape [14].

Prior studies have employed both qualitative and quantitative methods to evaluate memes. Qualitative analyses involve content

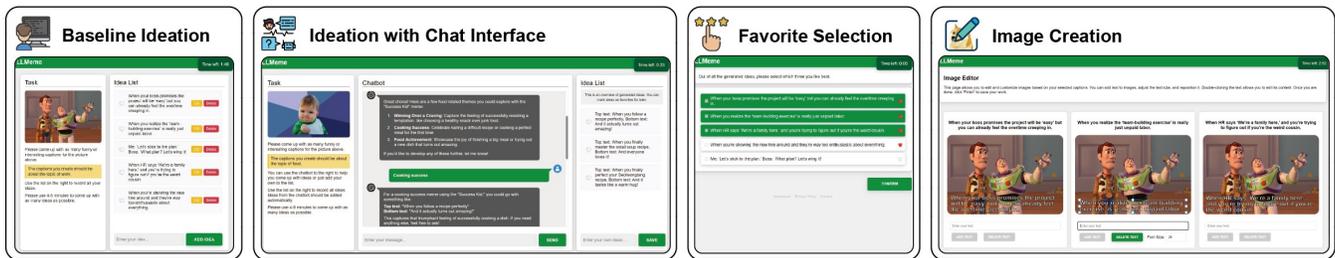


Figure 3: User Interface Overview: Baseline Ideation, Ideation with Chat Interface, Favorite Selection, and Final Image Creation.

analysis of themes and cultural relevance [19, 45], while quantitative approaches use machine learning models and sentiment analysis to predict meme virality based on textual and visual attributes [13, 34]. However, evaluating creative outputs poses challenges due to the subjective nature of creativity and humor, and the dynamic, context-dependent nature of memes [8].

3 Methodology

To explore the impact of human collaboration with LLMs on creative meme generation, we conducted a between-subject user study with three experimental groups. The following section presents the methodology of the user study.

3.1 Task

The participants’ task in the study was to generate captions for memes. More specifically, the task consisted of three steps:

Ideation In the first step, we displayed one of six background images of popular memes (Figure 2) to the participants and asked them to come up with as many captions as they could within five minutes. We asked participants to focus their ideas on one of three topics: *work*, *food*, and *sports*. The goal was to keep the ideas relatively constrained, for comparability but also not to overwhelm users with having to come up with arbitrary ideas. The interface (Figure 3) displayed a blank meme template as well as the instructions to the user. Users could then enter any ideas for image captions and they were then displayed in a list next to the instructions. Once the user had created ideas, they were also able to mark them as favorite, edit or remove them. For users in the treatment condition that had access to a LLM, this part of the interface also featured a chat interface where they could prompt the LLM (Figure 3). Any responses by the LLM were additionally processed to automatically determine whether the response contained any ideas. If this was the case, they were automatically extracted and added to the idea list.

Favorite Selection Once participants had completed the ideation step, they moved on to an overview of all the ideas they had come up with (Figure 3). From this full list, they had to select their top three ideas. These three ideas were then used in the last step.

Image Creation In the last step, we asked our participants to add their ideas as captions to the meme template. The meme editor allowed users to add text to the image in arbitrary

chunks (Figure 3). Each chunk could then be positioned and resized, edited or removed.

Each experimental group used different methods for creating memes, comparing the effects of creativity driven solely by humans, human-AI collaboration, and entirely AI-driven creation.

The first group (baseline) participants independently generated ideas and created memes without external assistance by an AI tool or otherwise. The second group also involved human participants who had to come up with memes but had access to a conversational interface. Through it, they were able to prompt an LLM to support them with generating ideas. In the third group, ideas were generated fully autonomously by the LLM.

Following the main study, we evaluated the memes generated by the three groups in terms of their funniness, shareability, and creativity using a second online survey.

3.2 Procedure

For generating memes, after recording for their informed consent, we asked participants to spend at least four and at most five minutes on coming up with captions using our UI (Figure 4). Following this ideation phase, they selected their three favorite ideas and could then edit the image to add their idea as captions. After generating and downloading their creation, they moved on to the next image with a next topic. Each participant had to produce memes for three different combinations of images and topics. The permutation of image and topic was selected randomly but in a way that each participant used each image and each topic at most once. After generating ideas for three different topics and images, participants completed a survey recording feedback on their experience. The overall process of ideation, selecting favorites, editing the images, and completing the survey was scheduled to take no more than 40 minutes. For their work, participants received compensation equivalent to 15 USD.

Following this first phase, we continued with rating the generated ideas (Figure 5). Since each participant had to select three favorite ideas for each of the three images they captioned, we gathered 882 ideas marked as favorites across both study conditions with human involvement. Due to technical problems, only 415 for the baseline conditions and all 441 for the collaborative condition were usable.

For each idea, we re-generated the captioned image to ensure consistent placement of the text. We then curated the images, excluding all those where the participants clearly entered a caption not matching the task or where the length of the caption obscured

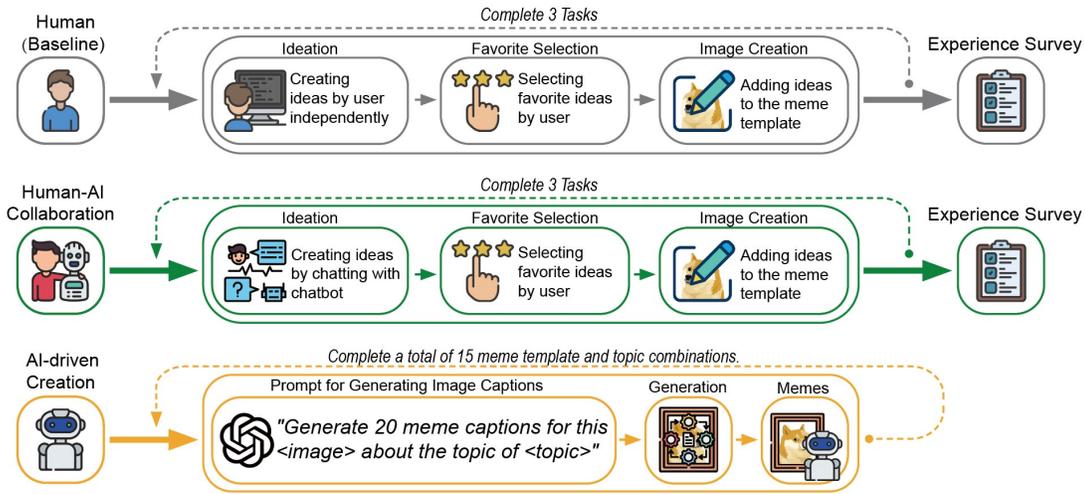


Figure 4: Meme Generation Workflow: Human (Baseline), Human-AI Collaboration, and AI-Driven Creation.

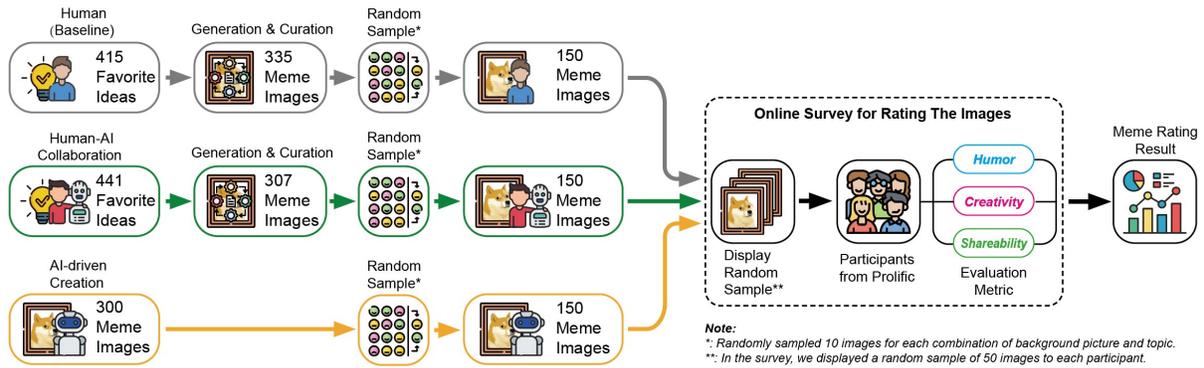


Figure 5: Meme Evaluation Workflow: This diagram illustrates the evaluation process of memes created by humans, human-AI collaboration, and AI-driven approaches.

the majority of the image. Since for one of the images, we had to exclude more than two thirds of the images, we decided to fully exclude it from the study. This left us with 335 images from the baseline and 307 images from the collaborative condition.

For rating the quality of these images, we then randomly sampled 10 images for each combination of background picture and topic, which, at five remaining images and three topics, left us with 150 images from the baseline and 150 images from the collaborative condition.

We then leveraged LLM to create fully AI generated captions for the third study condition. To this end, we prompted the model to generate captions for each combination of image and topic, giving us additional 150 images.

For assessing the subjective quality of these images, we asked a second group of participants to complete an online survey for rating the images. In the survey, we displayed a random sample of 50 images to each participant. For each image, the participated provided feedback along three dimensions: humor, creativity and shareability. These categories were selected based on prior work.

We estimated that each rating would take 10–15 seconds, so participants should complete the task in about 10-15 minutes. For their participation, they received compensation equivalent to 10 USD.

3.3 Prompting

For conducting the study, we used LLM in two functions: first, as part of the UI where participants could generate ideas with the assistance of a conversational UI. In this interface, participants were free to enter any prompt into the system. However, we set a system prompt to constrain the functionality and output of the system. This system prompt set the context for the LLM, including the fact that the goal of the system was to help users in creating meme ideas, the tone of the interaction to be helpful and polite, and it constrained the system to produce at most three ideas with a single response. Additionally, we always sent the current image to the LLM before any user prompt. The full prompts are available in the supplementary material.

Secondly, we used the LLM to generate image captions for generating the memes for the pure AI condition. For this, we again sent

the image first and then instructed the model to “generate 20 meme captions for this <image> about the topic of <topic>”, where <topic> was one of the three topics and <image> was a brief description of the image of no more than 10 words. A full list of the generated captions is also part of the supplementary material.

3.4 Apparatus

The user interface for the study was implemented using React while any data collection and the interaction with the OpenAI API for GPT-4o was performed by a NodeJS server. All processing of the prompts, randomization of tasks, etc. was performed on the server to ensure the integrity of the data.

Both parts of the study were conducted fully online using our implementation of the meme-creation interface for the first part of the study and a commercial survey platforms for any subsequent surveys.

3.5 Data Collection

While participants created memes, we recorded all ideas they came up with, both text and images, as well as a full log of their interaction with the LLM and its responses on the server. For recording the subjective perception, we used a commercial survey platform. The survey included questions for the participants to self-assess their creativity, as well as the NASA-TLX and general questions about the interface and the ideation process. We used the same platform for rating the generated ideas as well. Demographic data was provided via Prolific, which we used for participant recruiting.

3.6 Participants

For this first part of the study, we recruited 124 participants using the online platform Prolific. 26 participants were excluded due to not completing the task. The number of participants was determined after an initial power analysis for the study design. Given how the success of humor can be highly dependent on language skill, we selected only participants with good English skills. Additionally, we required participants to have used a LLM interface before at least once, to ensure they would be familiar with the concepts and interactions. This resulted in a diverse participant sample from 30 different countries. Of the participant, 63 indicated to identify as male and 35 as female. The average age was 28.8 years (sd: 8.7).

For the second phase of the study, we similarly recruited a second set of N=100 participants with the same prerequisites for language skills but knowledge of LLMs was not a requirement. 98 of these completed the task, rating at least 50 images. Participants in this group were equally split between identifying as male and female with an average age of 32.6 (sd: 11.1) and originating in 29 different countries.

4 Results

The following section will describe the quantitative findings and their statistical analysis.

4.1 Meme Creation

4.1.1 Idea Generation. During ideation, participants created an average of 6.1 ideas (sd: 3.2) with one participant managing to come with a total of 21 ideas for one of the images. As seen in Figure 7,

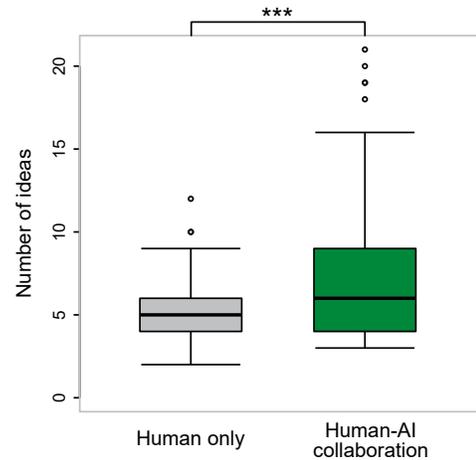


Figure 6: Participants using the LLM were able to produce significantly more ideas than participants who had no external support, according to the Mann-Whitney-U test (*: $p < 0.001$)**

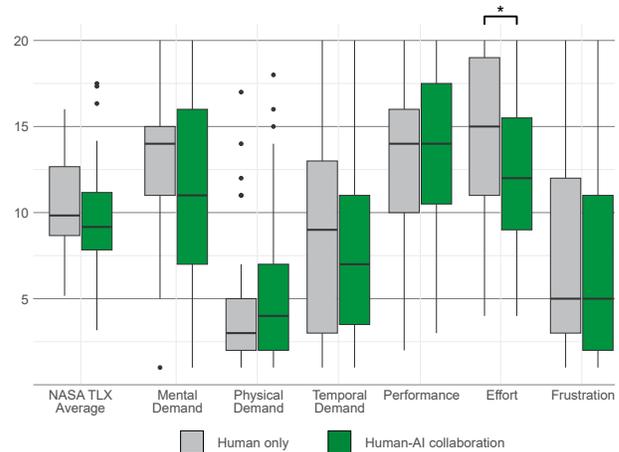


Figure 7: While there were no significant differences in overall workload, the “Effort” subscale of the NASA TLX was significantly different according to the Mann-Whitney-U test (*: $p < 0.05$)

participants that were able to use the LLM created noticeably more ideas than the participants in the baseline group. To get further insights how the presence of the chat affected the ideation process, we conducted statistical hypothesis tests on the number of ideas per participant. Following a Shapiro-Wilk test to determine non-normality for both the absolute number of ideas ($W = 0.811, p < 0.001$) and the average number of ideas per participant ($W = 0.820, p < 0.001$), we used the Mann-Whitney-U test. This test indicated significant differences for the absolute count ($W = 12652, p < 0.001$) and the average number of ideas (1519.5, $p < 0.001$).

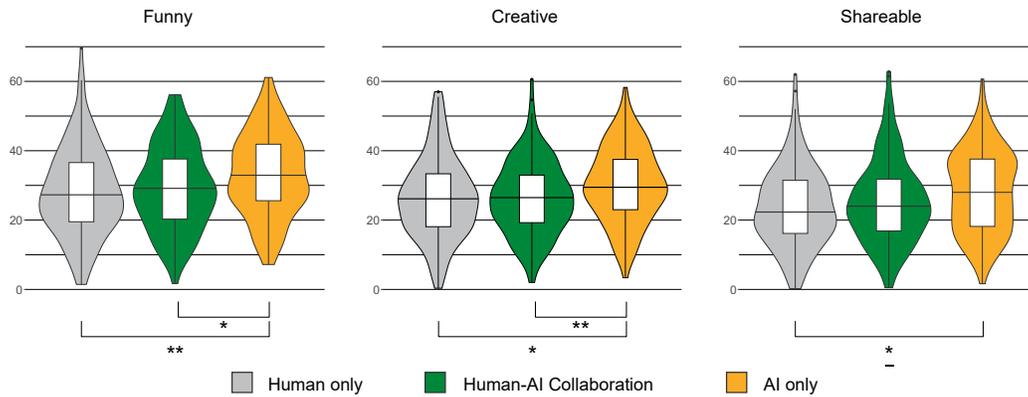


Figure 8: Pairwise comparison of how participants rated the memes with respect to the three scales “funny”, “creative”, and “shareable”. (*: $p < 0.05$, **: $p < 0.01$, *: $p < 0.001$, pairwise t-test/Mann-Whitney-U test, Bonferroni adjusted)**

4.1.2 Workload. While there are significant differences for the number of created ideas, there is no evidence that this also affected the workload that was required to achieve this result. Statistical analysis of the Raw TLX showed no significant differences (Shapiro-Wilk test: $W = 0.980$, $p = 0.1632$, t-test: $t = -0.955$, $df = 88.811$, $p = 0.342$). We found the same to be true for each of the six TLX subscales, except for the question “How hard did you have to work to accomplish your level of performance?”, where participants using the LLM entered significantly lower values (Shapiro-Wilk test: $W = 0.934$, $p < 0.001$, Mann-Whitney-U test: $W = 755$, $p = 0.027$).

4.1.3 General Feedback. Similarly, of the general questions about the user’s experience while creating the memes and the ideation process, we received responses that indicated no significant differences except for two questions.

For the first of these was the question whether participants felt that they created a lot of ideas. Results here match the actual idea count, with the LLM-supported users also subjectively noting that they created significantly more ideas (Shapiro-Wilk test: $W = 0.909$, $p < 0.001$, Mann-Whitney-U test: $W = 1308.5$, $p = 0.043$), although the difference is less stark than with the actual number of ideas.

For the question on perceived ownership “The generated captions are my ideas”, participants that did not use the LLM perceived a higher degree of ownership for the generated ideas (Shapiro-Wilk test: $W = 0.0766$, $p < 0.001$, Mann-Whitney-U test: $W = 562$, $p < 0.001$). However, even when using the chat, participants still generally felt ownership for the ideas.

4.2 Meme Rating

In the second phase of our experiments, we had a group of people rate the memes according to three criteria: how funny they thought they were, how creative they considered them, and how likely it was that they would share them. Along with the memes from the two conditions before, we had an additional condition with memes created exclusively using AI with no human input.

Considering the fact that the data from the “funny” and “creative” scale were likely normally distributed (Shapiro-Wilk test: $W = 0.994$, $p = 0.062$ and $W = 0.995$, $p = 0.155$ respectively), we used the ANOVA and pairwise t-tests, Bonferroni adjusted, to compare

these two. The third scale, “shareability” was likely not normally distributed (Shapiro-Wilk test: $W = 0.988$, $p < 0.001$), we used the Kruskal-Wallis test instead as well as pairwise Mann-Whitney-U tests, also Bonferroni adjusted.

According to these tests, each condition showed significant differences, as shown in Table 1. The pairwise comparison highlighted that it were consistently the memes generated by the LLM alone that were rated more positive than those where created with human involvement. The only exception to this was “shareability” where the comparison between the cooperative and pure AI creation was not significant. Memes created by humans with the help of the LLM were not rated significantly different than those from the baseline, i.e. without AI, for any of the three dimensions.

To ensure any unintended side-effects by of the image or topic selection, we performed the same statistical analysis to determine whether the image or the topic had any notable influence on the rating. This showed that the ratings across the images are relatively consistent, with only one pairing of images showing significantly different ratings for the question how funny the memes were perceived. The topic, on the other hand, seems to have had an impact on the rating of the memes, since work related memes were consistently rated significantly more funny, creative, and shareable.

We therefore analyzed the three scales from the survey again for each topic individually (see Table 1), which demonstrated that the previously found significant differences do not come evenly from across all data but seem to stem primarily from the memes about the topic of “work”.

5 Discussion

The results of our study provide some early insights into how the availability of LLM support influences people’s creative process. Further, we investigated how the output of human-LLM co-creation is viewed, compare to purely LLM-generated content. In the following, we discuss our results with regard to the research questions.

Table 1: Results of the statistical analysis of the meme ratings using ANOVA and Kruska-Wallis test (underlined).

	All topics									
	Shapiro-Wilk test		ANOVA / Kruskal-Wallis test							
	W	p	df	F/ χ^2	p					
Funny	0.994	0.062	2	6.971	0.001	**				
Creative	0.995	0.155	2	5.793	0.003	**				
Shareable	0.988	0.001	2	<u>11.761</u>	0.003	**				

	Memes about “work”			Memes about “sports”			Memes about “food”									
	Shapiro-Wilk		ANOVA/KW	Shapiro-Wilk		ANOVA/KW	Shapiro-Wilk		ANOVA/KW							
	W	p	df	F/ χ^2	p	W	p	df	F/ χ^2	p						
Funny	0.984	0.075	2	7.1	0.001	**	0.992	0.592	2	2.036	0.134	0.977	0.014	2	<u>3.455</u>	0.177
Creative	0.981	0.034	2	<u>11.22</u>	0.004	**	0.991	0.435	2	2.095	0.127	0.985	0.095	2	0.608	0.546
Shareable	0.981	0.033	2	<u>8.470</u>	0.014	*	0.983	0.063	2	1.712	0.184	0.977	0.013	2	<u>4.701</u>	0.095

5.1 LLM support increases content output without increasing effort but might diminish the feeling of ownership

In our study, participants who worked with the LLM assistant came up with significantly more ideas when creating memes compared to those participants working alone. Interestingly, even though they generated more ideas, they did not feel like the task was more demanding. The NASA-TLX results showed that the overall workload was not much different between the two groups, but participants in the LLM-assisted group did report that they had to put in less effort.

These results suggest that using the LLM can make the creative process of crafting humorous content more efficient by helping people generate more ideas without feeling overwhelmed. The AI assistant seemed to help them explore more options without putting in extra effort. This is consistent with previous studies, which suggest that AI tools can support users in generating more creative ideas by reducing obstacles associated with brainstorming and creative development[24, 31].

Additionally, participants who used the LLM reported a slightly reduced sense of ownership over their creations, indicating that AI assistance might affect the user’s connection to their work. However, participants generally still felt that they owned the idea. We attribute this finding to the fact that, in our online testing environment, AI assistance mainly contributed during the idea generation stage, whereas in the stages of idea screening and creating the meme image, no AI assistance or suggestions were provided. The final decision was always left up to the participants. Since feeling a sense of ownership and personal investment plays a significant role in creative motivation and satisfaction[5], it is essential to think about how we can balance the involvement of AI in the creative process.

5.2 The increased productivity of human-AI teams does not lead to better results - just to more results

The participants in our study developed more ideas in collaboration with the AI than when they worked alone. However, this did not translate into higher quality in the memes selected by our participants - which always happened without LLM support - in terms of the metrics we collected. This raises questions about the link between quantity and quality in human-AI collaboration. Although coming up with more ideas could increase the chances of producing something high-quality, our study did not find a significant difference.

In relation to our findings, several prior studies suggest similar outcomes in the context of human-AI co-creativity. For example, Wan et al. [49] found that participants who used an LLM during prewriting activities produced more creative ideas. However, these ideas were not significantly better in quality compared to those created by participants working alone. Similarly, Rezwana and Maher[43]pointed out ethical and practical challenges in human-AI co-creativity. They noted that while AI can help generate more content, it does not always improve the depth or quality of the work because it struggles with understanding the context and subtleties of creativity. These studies align with our results, showing that while human-AI collaboration can increase productivity and the number of ideas, it does not always lead to higher-quality output.

From this, we conclude that the use of AI support in the context of the metrics investigated in this study does not lead to better results in terms of humor. Conversely, we can also assume that users with AI support achieve a consistent result faster and with more variations, without this representing an additional mental burden for the users.

5.3 LLMs appeal to a broad taste in humor, but humans can be wittier still

Over all assessed metrics, we found that memes created solely by AI performed better than memes created solely by humans or in collaboration between humans and AI. We attribute this initially

quite surprising result to the following: The LLM used was trained on large data sets with many cultural references and different types of humor. During the training process, the LLM was most likely to learn the types of humor that it saw most frequently during the training process, i.e. those that resonated best with the crowd. As a result, such LLMs are good at creating content that appeals to a wide audience. The AI picks up on general trends in its training data, which helps it to produce content that most people find appealing. On the other hand, human-created content tends to draw from personal experiences and specific cultural backgrounds[9]. This broad appeal of AI-generated memes, however, may also be influenced by the composition of the evaluators. The people in the online study, as well as those who evaluated the memes, were random users from a crowdsourcing platform. While this approach brings in a variety of perspectives, it might not capture the subtle differences in humor appreciation across specific demographic groups. Humor is personal and influenced by factors like life experiences, cultural background, and social norms. Therefore, understanding these subtleties would likely require a more targeted evaluation across specific audience segments, which could provide deeper insights into how different groups perceive humor. Even when people work with AI, they often rely on their own experiences when choosing ideas, which creates a similar challenge. It is hard to compete with the AI's ability to cater to popular tastes.

This interpretation is further supported by the analysis of the top-performing memes (Figure 1). The funniest memes were mainly created by humans, while those rated highest for creativity and shareability were the result of human-AI collaborations. Among the top 4 humor memes, human creators claimed most of the top positions. In terms of creativity, humans still took half of the rankings, with the rest coming from human-AI teams. For shareability, human-AI collaborations made up half of the highest-ranked memes.

These results highlight that while AI is effective at creating broadly appealing content on average, the individual human touch resonates most deeply on certain dimensions. The top-ranked human creators likely brought in personal experiences, cultural nuances, or innovative ideas that AI, limited to patterns from existing data, cannot fully replicate. At the same time, human-AI collaboration showed real potential, particularly in creativity and shareability, suggesting that AI can offer new ideas or perspectives that, combined with human creativity, lead to content that's both original and widely appealing.

6 Limitations and Future Work

We are convinced that the presented user study provides valuable insights into the creativity process when generating humorous content together with LLMs and the perceived funniness of AI-generated content compared to human (co-) authored content. However, the design and results of our study imply a number of limitations and directions for future work, which we discuss below.

6.1 Short-term vs. long-term interaction with the AI

Our study only investigated short-term interactions with the system in a single session. We did not explore how prolonged use of

such an AI support system might affect creativity, satisfaction, or the development of new or improved skills. This short-term focus limits our understanding of how users' creative strategies evolve or how they rely on AI over time, which could lead to a decline in quality of the output. Future studies could therefore give participants the opportunity to use AI tools over a longer period of time. By following changes in creative strategies over time, reliance on AI and sense of ownership of generated content, this could provide a better understanding the long-term effects of AI collaboration in creative tasks.

6.2 Limited collaboration between humans and the LLM

We found that many participants did not fully utilize the potential of the LLM. Some users limited themselves to fulfilling the minimum requirements without revising their ideas in collaboration with the system. Less than half of the participants interacted with the LLM multiple times, and only six participants had more than eight interactions, possibly affecting the quality of their results. A possible reason for this might be our implementation of the study interface. Although the chatbox interface used during the idea generation phase was familiar and easy to use, it lacked structure to guide the creative process. This open-ended approach resulted in varied interactions depending on participants' individual backgrounds and ideation strategies. Further reasons could include the setting of unambitious goals for the participants, or even an inherent problem with crowdsourcing platforms for conducting such studies [42]. Future work could integrate more structured prompts or more collaborative tools into AI systems to encourage deeper engagement and iterative idea development.

6.3 Cultural and Social Influence on Humor and Creativity

What we find funny is heavily influenced by personal backgrounds such as social and cultural factors. In our study, the participants who contributed content and those who rated content came from a wide range of different cultural and social backgrounds. This approach is likely to have led to a wide range of interpretations and preferences for what is considered humorous or creative, thus impacting the outcomes. However, we did not systematically collect detailed demographic information (e.g., ethnicity or language proficiency), limiting our ability to fully explore how varying cultural backgrounds shape humor perception and meme creation. Future work should gather richer demographic and qualitative data—for instance, through open-ended survey questions or interviews—to capture the nuances in how different cultural, linguistic, or social backgrounds perceive and produce humor. Such qualitative insights could also reveal how participants interpret creativity in the context of meme-making, providing a more holistic picture of why certain ideas resonate (or fail to resonate) with different audiences.

7 Conclusion

In this paper, we examined the role of LLMs as co-creators in generating humorous content, focusing on the creation of internet memes. Our findings demonstrated that participants who collaborated with an LLM assistant produced a significantly higher number of ideas

without reporting an increase in perceived workload, suggesting improvements in both productivity and efficiency. However, this increase in ideas did not consistently lead to higher-quality content when humans were involved. Memes created through human-AI collaboration were rated about the same as those made by humans alone in terms of humor, creativity, and shareability. Interestingly, memes generated entirely by AI scored better, on average, than both human-only and human-AI collaborative memes. But when we looked at the top-performing memes, human-created content was strongest in humor, while human-AI collaborations excelled in creativity and shareability.

These findings show that human-AI collaboration in creative tasks is complex. While AI can increase productivity and produce content that appeals to a wide audience, human creativity is still key for creating content that connects more deeply in certain areas. Participants working with the LLM reported feeling less ownership over their work, suggesting that integrating AI into the creative process needs to be done carefully to keep users connected to their creations. Also, the short-term nature of our study, the limited use of the AI's full potential due to the open-ended interface, and the similar backgrounds of our participants suggest that more research is needed to understand the long-term effects of AI assistance on creativity and collaboration.

Looking forward, future studies should explore how long-term use of AI affects creative strategies, satisfaction, and skill development. AI interfaces could be improved by providing more structured guidance and encouraging deeper engagement, helping users make better use of AI's capabilities while maintaining ownership over their work. By addressing these challenges, we can develop smarter, more human-centered AI systems that not only boost productivity but also enhance human creativity.

Acknowledgments

This work was supported by HumanE-AI-Net under Grant Agreement ID: 952026.

References

- [1] 2016. Laughing across Borders. *The European Journal of Humour Research* 4, 4 (2016), 26–49.
- [2] Jock Abra. 1994. Collaboration in Creative Work: An Initiative for Investigation. *Creativity Research Journal* 7, 1 (1994), 1–20. doi:10.1080/10400419409534505
- [3] Preetham Gopalakrishna Adiga and Padmakumar K. 2024. Humor, Politics and the Global North: A Systematic Literature Review of Internet Memes. *Internet reference services quarterly* (04 2024), 1–19. doi:10.1080/10875301.2024.2335920
- [4] Ahmad Ali and Mina Eshaq. 2023. Using Artificial Intelligence for enhancing Human Creativity. *Journal of Art, Design and Music* 2 (06 2023). doi:10.55554/2785-9649.1017
- [5] Teresa M. Amabile. 1983. The social psychology of creativity: A componential conceptualization. *Journal of Personality and Social Psychology* 45 (1983), 357–376. doi:10.1037/0022-3514.45.2.357
- [6] Barrett R Anderson, Jash Hemant Shah, and Max Kreminski. 2024. Homogenization Effects of Large Language Models on Human Creative Ideation. *Creativity and Cognition* (06 2024). doi:10.1145/3635636.3656204
- [7] Cecilia R. Aragon and Alison Williams. 2011. Collaborative creativity: a complex systems model with distributed affect. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Vancouver, BC, Canada) (CHI '11). Association for Computing Machinery, New York, NY, USA, 1875–1884. doi:10.1145/1978942.1979214
- [8] Kate Barnes, Péter Juhász, Marcell Nagy, and Roland Molontay. 2024. Topicality boosts popularity: a comparative analysis of NYT articles and Reddit memes. *Social Network Analysis and Mining* 14 (06 2024). doi:10.1007/s13278-024-01272-3
- [9] Lucas Bellaïche, Rohin Shahi, Martin Harry Turpin, Anya Ragnhildstveit, Shawn Sprockett, Nathaniel Barr, Alexander Christensen, and Paul Seli. 2023. Humans versus AI: Whether and why we prefer human-created compared to AI-created artwork. *Cognitive Research: Principles and Implications* 8 (07 2023). doi:10.1186/s41235-023-00499-6
- [10] Susan Blackmore. 2000. The Power of Memes. *Scientific American* 283 (10 2000), 64–73. doi:10.1038/scientificamerican1000-64
- [11] Daniel Buschek, Lukas Mecke, Florian Lehmann, and Hai Dang. 2024. Nine Potential Pitfalls when Designing Human-AI Co-Creative Systems. <https://arxiv.org/abs/2104.00358v1>
- [12] François Chiocchio, Simon Grenier, Thomas A. O'Neill, Karine Savaria, and J. Douglas Willms. 2012. The Effects of Collaboration on Performance: A Multilevel Validation in Project Teams. *International Journal of Project Organisation and Management* 4, 1 (Jan. 2012), 1–37. doi:10.1504/IJPO.2012.045362
- [13] Victor W Chu, Raymond K Wong, Fang Chen, and Chi-Hung Chi. 2017. Prediction-as-a-Service for Meme Popularity. (06 2017). doi:10.1109/scc.2017.56
- [14] Michele Coscia. 2014. Average is Boring: How Similarity Kills a Meme's Success. *Scientific Reports* 4 (09 2014). doi:10.1038/srep06477
- [15] Christie Davies. 1990. *Ethnic Humor around the World: A Comparative Analysis*. Indiana University Press, Bloomington, IN, US, x, 404 pages.
- [16] Shuangrui Ding, Zihan Liu, Xiaoyi Dong, Pan Zhang, Rui Qian, Conghui He, Dahua Lin, and Jiaqi Wang. 2024. SongComposer: A Large Language Model for Lyric and Melody Composition in Song Generation. *arXiv (Cornell University)* (02 2024). doi:10.48550/arxiv.2402.17645
- [17] SeungHeon Doh, Keunwoo Choi, Jongpil Lee, and Juhan Nam. 2023. LP-MusicCaps: LLM-Based Pseudo Music Captioning. *arXiv (Cornell University)* (01 2023). doi:10.48550/arxiv.2307.16372
- [18] Anil R Doshi and Oliver P Hauser. 2024. Generative AI enhances individual creativity but reduces the collective diversity of novel content. *Science advances* 10 (07 2024). doi:10.1126/sciadv.adn5290
- [19] Yuhao Du, Muhammad Aamir Masood, and Kenneth Joseph. 2020. Understanding Visual Memes: An Empirical Analysis of Text Superimposed on Memes Shared on Twitter. *Proceedings of the International AAAI Conference on Web and Social Media* 14 (05 2020), 153–164. doi:10.1609/icwsm.v14i1.7287
- [20] Philippe Esling and Ninon Devis. 2020. Creativity in the era of artificial intelligence. doi:10.48550/arXiv.2008.05959
- [21] Josh Gardner, Simon Durand, Daniel Stoller, and Rachel M Bittner. 2023. LLaRC: A Multimodal Foundation Model for Music. *arXiv (Cornell University)* (01 2023). doi:10.48550/arxiv.2310.07160
- [22] Russell G. Geen. 1994. Social Motivation. In *Companion Encyclopedia of Psychology*. Routledge.
- [23] Drew Gorenz and Norbert Schwarz. 2024. How funny is ChatGPT? A comparison of human- and A.I.-produced jokes. *PLoS one* 19 (07 2024), e0305364–e0305364. doi:10.1371/journal.pone.0305364
- [24] Matthew Guzdial and Mark O Riedl. 2019. An Interaction Framework for Studying Co-Creative AI. *arXiv (Cornell University)* (03 2019). doi:10.48550/arxiv.1903.09709
- [25] Erik E. Guzik, Christian Byrge, and Christian Gilde. 2023. The originality of machines: AI takes the Torrance Test. *Journal of Creativity* 33 (12 2023), 100065. doi:10.1016/j.jyoc.2023.100065
- [26] Carlos Gómez-Rodríguez and Paul Williams. 2023. A Confederacy of Models: A Comprehensive Evaluation of LLMs on Creative Writing. doi:10.48550/arXiv.2310.08433
- [27] Jessica He, Stephanie Houde, Gabriel E Gonzalez, Silva Moran, Steven I Ross, Michael Muller, and Justin D Weisz. 2024. AI and the Future of Collaborative Work: Group Ideation with an LLM in a Virtual Canvas. (06 2024). doi:10.1145/3663384.3663398
- [28] Jimpei Hitsuwari, Yoshiyuki Ueda, Woojin Yun, and Michio Nomura. 2022. Does human-AI collaboration lead to more creative art? Aesthetic evaluation of human-made and AI-generated haiku poetry. *Computers in Human Behavior* 139 (10 2022), 107502. doi:10.1016/j.chb.2022.107502
- [29] Kent F Hubert, Kim N Awa, and Darya L Zabelina. 2024. The current state of artificial intelligence generative language models is more creative than humans on divergent thinking tasks. *Scientific Reports* 14 (02 2024). doi:10.1038/s41598-024-53303-w
- [30] Jianan Jiang, Di Wu, Hanhui Deng, Yidan Long, Wenyi Tang, Xiang Li, Can Liu, Zhanpeng Jin, Wenlei Zhang, and Tangquan Qi. 2024. HAIGEN: Towards Human-AI Collaboration for Facilitating Creativity and Style Generation in Fashion Design. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 8 (08 2024), 1–27. doi:10.1145/3678518
- [31] Pegah Karimi, Jeba Rezwana, Safat Siddiqui, Mary Lou Maher, and Nasrin Dehbozorgi. 2020. Creative sketching partner. *Proceedings of the 25th International Conference on Intelligent User Interfaces* (03 2020). doi:10.1145/3377325.3377522
- [32] Aaron Kozbelt. 2019. Chapter 10 - Evolutionary Explanations for Humor and Creativity. In *Creativity and Humor*, Sarah R. Luria, John Baer, and James C. Kaufman (Eds.). Academic Press, 205–230. doi:10.1016/B978-0-12-813802-1.00010-7
- [33] Hye-Kyung Lee. 2022. Rethinking creativity: creative industries, AI and everyday creativity. *Media, Culture & Society* 44 (03 2022), 016344372210770. doi:10.1177/01634437221077009
- [34] Chen Ling, Ihab AbuHilal, Jeremy Blackburn, Emiliano De Cristofaro, Savvas Zannettou, and Gianluca Stringhini. 2021. Dissecting the Meme Magic: Understanding Indicators of Virality in Image Memes. *Proceedings of the ACM on*

- Human-Computer Interaction* 5 (04 2021), 1–24. doi:10.1145/3449155
- [35] Yiren Liu, Si Chen, Haocong Cheng, Mengxia Yu, Xiao Ran, Andrew Mo, Yiliu Tang, and Yun Huang. 2024. How AI Processing Delays Foster Creativity: Exploring Research Question Co-Creation with an LLM-based Agent. (05 2024). doi:10.1145/3613904.3642698
- [36] Li-Chun Lu, Shou-Jen Chen, Tsung-Min Pai, Chan-Hung Yu, Hung-yi Lee, and Shao-Hua Sun. 2024. LLM Discussion: Enhancing the Creativity of Large Language Models via Discussion Framework and Role-Play. doi:10.48550/arXiv.2405.06373
- [37] Suresh Malodia, Amandeep Dhir, Anil Bilgihan, Pranao Sinha, and Tanishka Tikoo. 2022. Meme marketing: How Can Marketers Drive Better Engagement Using Viral memes? *Psychology & Marketing* 39 (06 2022). doi:10.1002/mar.21702
- [38] Lena Mamykina, Linda Candy, and Ernest Edmonds. 2002. Collaborative Creativity. *Commun. ACM* 45, 10 (Oct. 2002), 96–99. doi:10.1145/570907.570940
- [39] ROD A. MARTIN. 2007. CHAPTER 5 - The Social Psychology of Humor. In *The Psychology of Humor*, ROD A. MARTIN (Ed.). Academic Press, Burlington, 113–152. doi:10.1016/B978-012372564-6/50024-1
- [40] Maria D. Molina. 2020. What Makes an Internet Meme a Meme? Six Essential Characteristics. *Handbook of Visual Communication* (04 2020), 380–394. doi:10.4324/9780429491115-35
- [41] OpenAI. 2023. GPT-4 Technical Report. *arXiv (Cornell University)* (03 2023). doi:10.48550/arxiv.2303.08774
- [42] Jonas Oppenlaender, Kristy Milland, Aku Visuri, Panos Ipeirotis, and Simo Hosio. 2020. Creativity on Paid Crowdsourcing Platforms. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (*CHI '20*). Association for Computing Machinery, New York, NY, USA, 1–14. doi:10.1145/3313831.3376677
- [43] Jeba Rezwana and Mary Lou Maher. 2023. User Perspectives on Ethical Challenges in Human-AI Co-Creativity: A Design Fiction Study. *Creativity and Cognition* (06 2023). doi:10.1145/3591196.3593364
- [44] Limor Shifman, Hadar Levy, and Mike Thelwall. 2014. Internet Jokes: The Secret Agents of Globalization?*. *Journal of Computer-Mediated Communication* 19, 4 (July 2014), 727–743. doi:10.1111/jcc4.12082
- [45] Kohtaro Tanaka, Hiroaki Yamane, Yusuke Mori, Yusuke Mukuta, and Tatsuya Harada. 2022. Learning to Evaluate Humor in Memes Based on the Incongruity Theory. *ACL Anthology* (10 2022), 81–93. <https://aclanthology.org/2022.cai-1.9>
- [46] Saranya Venkatraman, Nafis Irtiza Tripto, and Dongwon Lee. 2024. CollabStory: Multi-LLM Collaborative Story Generation and Authorship Analysis. *arXiv (Cornell University)* (06 2024). doi:10.48550/arxiv.2406.12665
- [47] Florent Vinchon, Todd Lubart, Sabrina Bartolotta, Valentin Gironnay, Marion Botella, Samira Bourgeois-Bougrine, Jean-Marie Burkhardt, Nathalie Bonnardel, Giovanni Emanuele Corazza, Vlad Petre Glaveanu, Michael Hanchett Hanson, Zorana Ivcevic, Maciej Karwowski, James C Kaufman, Takeshi Okada, Roni Reiter-Palmon, and Andrea Gaggioli. 2023. Artificial Intelligence & Creativity: A Manifesto for Collaboration. *Artificial Intelligence & Creativity: A Manifesto for Collaboration* 57 (06 2023). doi:10.1002/jocb.597
- [48] Camilla Vásquez and Erhan Aslan. 2021. “Cats be outside, how about meow”: Multimodal humor and creativity in an internet meme. *Journal of Pragmatics* 171 (01 2021), 101–117. doi:10.1016/j.pragma.2020.10.006
- [49] Qian Wan, Siying Hu, Yu Zhang, Piaohong Wang, Bo Wen, and Zhicong Lu. 2024. “It Felt Like Having a Second Mind”: Investigating Human-AI Co-creativity in Prewriting with Large Language Models. *Proceedings of the ACM on human-computer interaction* 8 (04 2024), 1–26. doi:10.1145/3637361
- [50] Han Wang and Roy Ka-Wei Lee. 2024. MemeCraft: Contextual and Stance-Driven Multimodal Meme Generation. *arXiv (Cornell University)* (05 2024). doi:10.1145/3589334.3648151
- [51] Tiannan Wang, Jiamin Chen, Qingrui Jia, Shuai Wang, Ruoyu Fang, Huilin Wang, Zhaowei Gao, Chunzhao Xie, Chuou Xu, Jihong Dai, Yibin Liu, Jialong Wu, Shengwei Ding, Long Li, Zhiwei Huang, Xinle Deng, Teng Yu, Gangan Ma, Han Xiao, Zixin Chen, Danjun Xiang, Yunxia Wang, Yuanyuan Zhu, Yi Xiao, Jing Wang, Yiru Wang, Siran Ding, Jiayang Huang, Jiayi Xu, Yilihamu Tayier, Zhenyu Hu, Yuan Gao, Chengfeng Zheng, Yueshu Ye, Yihang Li, Lei Wan, Xinyue Jiang, Yujie Wang, Siyu Cheng, Zhule Song, Xiangru Tang, Xiaohua Xu, Ningyu Zhang, Huajun Chen, Yuchen Eleanor Jiang, and Wangchunshu Zhou. 2024. Weaver: Foundation Models for Creative Writing. <http://arxiv.org/abs/2401.17268>
- [52] Caleb Warren, Adam Barsky, and A. Peter McGraw. 2021. What Makes Things Funny? An Integrative Review of the Antecedents of Laughter and Amusement. *Personality and Social Psychology Review* 25, 1 (2021), 41–65. doi:10.1177/1088868320961909 arXiv:<https://doi.org/10.1177/1088868320961909> PMID: 33342368.
- [53] Gillian Wigglesworth and Neomy Storch. 2012. What Role for Collaboration in Writing and Writing Feedback. *Journal of Second Language Writing* 21, 4 (Dec. 2012), 364–374. doi:10.1016/j.jslw.2012.09.005
- [54] Laurie Ann Williams. 2000. *The Collaborative Software Process*. Ph. D. Dissertation. The University of Utah.
- [55] Willmore and Hocking. 2017. Internet Meme Creativity as Everyday Conversation. *Journal of Asia-Pacific Pop Culture* 2 (2017), 140. doi:10.5325/jasiapacpopcult.2.2.0140
- [56] Zhuohao Wu, Danwen Ji, Kaiwen Yu, Xianxu Zeng, Dingming Wu, and Mohammad Shidujaman. 2021. AI Creativity and the Human-AI Co-creation Model. *Human-Computer Interaction. Theory, Methods and Tools* 12762 (2021), 171–190. doi:10.1007/978-3-030-78462-1_13
- [57] Ann Yuan, Andy Coenen, Emily Reif, and Daphne Ippolito. 2022. Wordcraft: Story Writing With Large Language Models. *27th International Conference on Intelligent User Interfaces* (03 2022). doi:10.1145/3490099.3511105
- [58] Shanshan Zhong, Zhongzhan Huang, Shanghua Gao, Wushao Wen, Liang Lin, Marinka Zitnik, and Pan Zhou. 2024. Let’s Think Outside the Box: Exploring Leap-of-Thought in Large Language Models with Creative Humor Generation. *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* 35 (06 2024), 13246–13257. doi:10.1109/cvpr52733.2024.01258
- [59] Jichen Zhu, Antonios Liapis, Sebastian Risi, Rafael Bidarra, and G. Michael Youngblood. 2018. Explainable AI for Designers: A Human-Centered Perspective on Mixed-Initiative Co-Creation. 8 pages. doi:10.1109/CIG.2018.8490433

Received 10 October 2024; accepted 12 December 2024