

A Robot Walks into a Bar: Can Language Models Serve as Creativity Support Tools for Comedy?

An Evaluation of LLMs' Humour Alignment with Comedians

Piotr W. Mirowski*
piotirmirowski@deepmind.com
DeepMind
London, UK

Kory Mathewson
korymath@deepmind.com
DeepMind
London, UK

Juliette Love*
juliettelove@deepmind.com
DeepMind
London, UK

Shakir Mohamed
shakir@deepmind.com
DeepMind
London, UK

ABSTRACT

We interviewed twenty professional comedians who perform live shows in front of audiences and who use artificial intelligence in their artistic process as part of 3-hour workshops on “AI x Comedy” conducted at the Edinburgh Festival Fringe in August 2023 and online. The workshop consisted of a comedy writing session with large language models (LLMs), a human-computer interaction questionnaire to assess the Creativity Support Index of AI as a writing tool, and a focus group interrogating the comedians’ motivations for and processes of using AI, as well as their ethical concerns about bias, censorship and copyright. Participants noted that existing moderation strategies used in safety filtering and instruction-tuned LLMs reinforced hegemonic viewpoints by erasing minority groups and their perspectives, and qualified this as a form of censorship. At the same time, most participants felt the LLMs did not succeed as a creativity support tool, by producing bland and biased comedy tropes, akin to “cruise ship comedy material from the 1950s, but a bit less racist”. Our work extends scholarship about the subtle difference between, on the one hand, harmful speech, and on the other hand, “offensive” language as a practice of resistance, satire and “punching up”. We also interrogate the global value alignment behind such language models, and discuss the importance of community-based value alignment and data ownership to build AI tools that better suit artists’ needs. **Warning: this study may contain offensive language and discusses self-harm.**

CCS CONCEPTS

• **Applied computing** → **Performing arts**; • **Human-centered computing** → **Empirical studies in HCI**; *User studies*; • **Social and professional topics** → *Censoring filters*.

*Both authors contributed equally to this research.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
FAccT '24, June 03–06, 2024, Rio de Janeiro, Brazil
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0450-5/24/06
<https://doi.org/10.1145/3630106.3658993>

KEYWORDS

Large Language Models, Comedy, Creativity, Offensive speech, Censorship, Value Alignment

ACM Reference Format:

Piotr W. Mirowski, Juliette Love, Kory Mathewson, and Shakir Mohamed. 2024. *A Robot Walks into a Bar: Can Language Models Serve as Creativity Support Tools for Comedy? An Evaluation of LLMs' Humour Alignment with Comedians*. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency (FAccT '24)*, June 03–06, 2024, Rio de Janeiro, Brazil. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3630106.3658993>

1 INTRODUCTION

1.1 Motivation: investigate the potential and implications of LLMs for comedy writing

Recent work on the intersection of AI and comedy has demonstrated [24, 54, 68, 70, 99, 106] an appetite for comedians to (try to) write humorous material using AI tools like Large Language Models (LLMs). We conducted an empirical study to better understand the current state of LLMs as comedy-writing support tools, their use-cases and limitations, and artists’ opinions on ethical questions regarding their use in a comedy-writing context. The complexity of comedy can help expose some limitations of LLMs. To participate, we recruited 20 professional comedians who use AI in their artistic processes and who perform live shows in front of audiences (10 in person at the Edinburgh Festival Fringe, 10 online) for a 3-hour-long workshop on “AI x Comedy”. As detailed in Section 2, the workshop consisted of a comedy writing session with generally-available instruction-tuned LLMs (ChatGPT [77, 79] and Bard [26, 98]), a human-computer interaction questionnaire, and focus group discussions on the use of LLMs in comedy writing and ethical concerns.

1.1.1 Using LLMs for humour, a task with human-level difficulty. Trying to combine humour and machine intelligence is a long-standing subject of scientific enquiry [85, 93, 102], and is perceived as a fundamental challenge. According to computational humour researchers like Winters [106], “humans are the only known species that use humor for making others laugh” [20, 38]. Winters [106] argues that one of the modern formal humor theories points to *incongruity* [51] (whereby the setup points in one direction and

the punch line in another) as a basic element [38, 85, 87]¹. As we discuss in Section 5.2.2, producing and resolving incongruity is a task with human-level difficulty. We situate LLMs for comedy within broader computational humour research and AI-assisted comedy performance [67, 71, 88] in Appendix A.

1.1.2 The utility of instruction-tuned LLMs as Creativity Support Tools. Similar to previous empirical studies on the use of LLMs for creative writing [19, 21–23, 37, 49, 53, 71, 83, 114], we asked the artists about their motivations and processes for using LLMs. We asked about the potential and limitations of language models as *Creativity Support Tools* [22], and quantified the *Creativity Support Index* of LLMs for comedy writing [25]. We report our results in Section 3.

1.1.3 Socio-Technical Systems concerns with LLMs for creative writing. Inspired by Dev et al. [29], we leveraged community engagement with large generative models, and interrogated the diverse, intersectional identities of the comedians using AI in a creative context. In addition to their reasons for using (or not using) LLMs as comedy-writing tools, we asked participants about the ethical considerations of using AI. The study was conducted at the time of the Writers' Guild of America (WGA) 2023 strike [75]. Participants raised and addressed questions on the scrutiny of AI and on concerns around AI's impacts, both on intellectual property and artistic copyright, and on artists' livelihoods. We report their opinions in Section 4 and discuss these concerns in Section 5.3.

1.2 Investigating hypotheses about using LLMs to write comedy

Based on previous work and on the author's personal experience of AI as creativity support tools, we hypothesized—prior to conducting our study—that participants would express negative opinions of LLMs for co-creativity on four issues: expressing stereotyped (Sect. 1.2.1) or bland (Sect. 1.2.4) language, censorship (Sect. 1.2.2) and missing context (Sect. 1.2.3). We review literature on these four hypotheses below, as they become the basis of our mixed methods study.

1.2.1 Biases in large language models. Gender and racial biases embedded within machine learning models have been extensively documented [1, 11, 12, 14, 15, 18, 30, 42, 94, 105]. These biases include sexism and racism [11, 94], homophobia and transphobia [31, 84], Islamophobia [1], the perpetuation of Western colonial mindsets [73], Anglocentrism [115], and in-group vs. out-group social identity biases [48]. In their extensive reviews, Bommasani et al. [15], Rauh et al. [86] identified two broad kinds of harms resulting from such biases: *intrinsic harms* such as representational bias (due to misrepresentation, overrepresentation, and underrepresentation of specific social groups), and *extrinsic harms*, the downstream consequences of biased models, including representational and performance disparities. As we show in Section 4, the study participants noticed a few examples of representational harm and many examples of underrepresentation harm (also called *allocational harm* in

[86]), such as erasure when LLMs refused to generate content for certain demographic groups.

1.2.2 Potential censorship of speech labeled as “offensive”. Comedians often pepper their language with profanities and their material with provoking themes. As we discuss in Section 5.1.2 (and confirmed by the study participants in Section 4), *offensive* language that would be perfectly acceptable at a comedy club may get “censored” by instruction-tuned LLMs that “refuse” to answer “offensive” prompts.

This problem has been observed in automated moderation of online content, such as hate speech detectors that suppress social media posts by queer communities and drag queens [31, 32], or posts using African-American Vernacular English [3, 112]. Amironesei and Diaz [3] called it *censorship*. Rauh et al. [86] studied algorithmic moderation of social media posts by the Perspective API, noting that “authors of the comment may be harmed if their content is incorrectly flagged as toxic” by the moderation tool. Similar erasure due to the cultural hegemony embedded in image generators has been studied in [83]. We relate the participants' experience, similarly frustrated that the LLM tools “considered” their own identity and comedy material as problematic and necessary to censor.

Diaz et al. [32] defined *offensive language* as non-normative “language that uses terminology that is noted as offensive but which is not perceived as offensive in particular contexts of use”, and studied its use by minorities as a form of resistance, for “socially productive uses of decoratively offensive language”, aiming to reclaim “offensive” language and resist oppression. Just like the minority groups described in [32], many comedians (who may be members of minority groups themselves) often use offensive jokes to *punch up*, and satire (“to challenge existing social structures”) to build empathy, rather than to *punch down* (“silence others”).

1.2.3 Missing context. Context is key to disambiguate offensive language from hate speech. LLMs, like social media posts, cause “context collapse” [65] by providing a limited amount of information to understand their meaning, particularly when using mock impoliteness. Specifically, “in-group usage of reclaimed slurs can be considered acceptable, depending on who uses them” [28, 86]. Moreover, the context of comedy extends beyond the language to other factors including the audience and the venue.

1.2.4 Homogenisation. Bommasani et al. [15] warned that “the application of foundation models across domains has the potential to act as an epistemically and culturally homogenising force, spreading one perspective, often a socially dominant one, across multiple domains of application”. In the arts, this means that AI-generated artifacts may lead to a homogenisation of aesthetic styles [34, 104], further reinforced by curation algorithms [34, 64]. For creative writing, empirical studies showed that instruction-tuned LLMs reduce the diversity of content in co-writing tasks [80], and that LLM-generated stories did not pass the *Torrance Test of Creative Writing* according to metrics of “fluency, flexibility, originality and elaboration” [21]. Similarly to Qadri et al. [83], we ran focus groups with artists to interrogate cultural (Western) biases (see Sections 2.3 and 2.4).

1.2.5 Investigating hypotheses via a mixed-methods study. In our empirical study, we ask participants questions on all four problems

¹Alternative humour theories include the Aristotelian *Relief Theory* [5] of tension and release whereby we let out our psychic energy connected with repressed topics, the *Superiority Theory* [46] whereby we laugh at others' misfortunes to feel better about ourselves, and the *Benign Violation Theory* [69].

identified in Section 1.2, namely about bias, censorship, context and homogeneity. In Section 5, we build upon scholarship on cultural value alignment of language models, the moderation of offensive and harmful speech, and the use of offensive speech and satire as a form of resistance, to revisit the global cultural value alignment of LLMs and propose community-based alignment to build LLMs that better suit comedians' creative needs².

2 METHODS

Our study was designed to address a challenging problem, with on one hand, limitations of LLMs (stereotypes, inability to distinguish comedic offensiveness from harmful speech, cultural erasure and homogenisation of content), and on the other hand, the use of LLMs for a creative writing task. For this reason, we asked a group of experts—professional comedians and performers—who are used both to thinking about thorny questions of identity, offensiveness and censorship in their work, and to employing language in a highly creative way. We chose artists who already use AI in their work and expected them to be somewhat knowledgeable and open to using AI: this likely biased our results³.

We ran workshops with 20 comedians who use AI creatively. The first workshop with 10 participants was run in person at Edinburgh Festival Fringe 2023; the following 3 workshops with 3, 4 and 3 participants were run online. We reached out to comedians performing in Edinburgh during Fringe, or in our network, and attempted to recruit as diverse (along linguistic, cultural, gender, sexual, national and racial dimensions) a pool of comedians as possible given the constraints of the study⁴. Participants had contrasting views on AI for comedy writing, from "AI is very bad at this, and I don't want to live in a world where it gets better" (p15) to "I liked the details that I got. I think those details sparked my imagination, and I think I could use them to write something" (p20). Participants were asked to register on the Prolific platform⁵ and invited to join a specific study thanks to an allowlist. The study was approved by the research ethics committee of our institution. The information sheet and consent forms were shared with the participants, their active consent was obtained at the beginning of the workshop and they had the right to withdraw without prejudice at any time. The Prolific platform handled the payment of their participation fee, set to £300 for 3 hours.

We started each 3-hour session by describing the agenda and goals of the workshop, sharing the information sheet and consent forms with the participants, and asking them to start filling

²We deliberately do not generalize our findings beyond comedy, as some professionals and the computational creativity community have historically embraced LLM tools in a way that fits their creative practice, whether building on the glitch aesthetic [64, 82] or designing interactive experiences [81].

³Our biased selection criteria of participants might, and likely do, lead to biased opinions as compared to the much more broad population of comedians and performers, which might be reflected in more favourable judgment of the Creativity Support Index of LLM writing tools. Future research might explore the diversity of opinions in creative communities across a greater range of familiarity with AI tools and openness to using them in their own creative practices. Exploring those opinions would significantly increase the scope of the paper and would make a compelling follow-up study.

⁴A demographic analysis of opinions might be a possible avenue for future investigations, but it would require a different study design and participant recruiting process.

⁵<https://prolific.com>

out a short anonymous survey about their background in comedy, previous exposure to AI and usage of AI in performance (full questionnaire in Appendix B.1).

2.1 Writing exercise

We then proceeded with a comedy-writing exercise, in which participants spent around 45 minutes on their own, using an LLM. We encouraged participants to try to use the LLM in a way that would generate useful material "that they would be comfortable presenting in a comedy context", but emphasized that we did not require a fully-finished product by the end of the writing exercise. We invited them to use the language(s) they felt the most comfortable with. We also suggested they could use the tool to 1) generate, rate/detect or explain jokes, 2) co-write jokes via iterative prompting, step-by-step or using examples, and 3) analyse, re-write or complete some of their previous material. In the first workshop (in person), we provided participants with access to ChatGPT-3.5 [79] served via a plain text interface similar to ChatGPT. In the following 3 workshops, we invited participants to use their own preferred model via their personal account: participants used ChatGPT-3.5, ChatGPT-4 [77] and Google Bard powered by Gemini Pro [98]. Note that the choice of such instruction-tuned models was motivated by their popularity and ease of access by comedians, and more complex prompting strategies, such as used in Dramatron [71], could have produced higher-quality outputs.

2.2 Creativity Support Tools evaluation

Following the writing exercise, we asked participants to fill out three surveys. The first survey was about their experience with the AI system for writing comedy material and contained nine questions from previous studies [53, 71, 114] that assessed LLMs for creative writing on the 5-level Likert scale (see Appendix B.2). The second survey was used to calculate the Creativity Support Index (CSI) [25] of the writing tool, which itself was adapted from the NASA Task Load Index [43]. CSI is estimated in a psychometric survey that measures six dimensions of creativity support: *Exploration*, *Expressiveness*, *Immersion*, *Enjoyment*, *Results Worth Effort*, and *Collaboration* (see specific questions in Appendix B.3), and is a number between 0 and 100, where 90 is considered excellent and 50 mediocre. The third survey contained free-form questions on one thing that the "AI system" (the LLM writing tool) did well, one improvement, and open-ended comments on the writing session and on the survey.

2.3 Focus Group Questions

The last part of the workshop consisted in a one-hour focus group, where we asked participants first to discuss the writing task (for about 30 minutes) and second to discuss the general usage of AI for writing comedy material.

In order to guide the discussion, we prepared two sets of questions⁶ (see Appendix B.4 for the full list of questions). The first set of questions pertained to the usefulness of the outputs generated by the LLM tool for personal writing, differences between using an LLM or searching for inspiration using Wikipedia or a search

⁶Question-led focus groups are useful to start discussions, but we acknowledge the limitation that questions can bias the participants' responses.

engine, the types of comedy that can be produced by an LLM, and concerns about the ownership of LLM-generated outputs.

The second set of questions addressed the comedy writing process of the participants, as well as the topics introduced in Section 1.2, namely various biases and stereotypes of LLMs, problems with moderation strategies employed by LLMs, the importance of context and delivery or whether some forms of cultural appropriation or homogenisation could happen. We invited discussions about the use of other comedians' work, and also challenged the participants with question on whether the AI has a "voice" and if humour can be quantified.

2.4 Focus Group Analysis

In our workshops, we had followed focus group methodology described in [74, 76] (engaging a group of participants in an informal one hour discussion focused around a particular topic, activity, or stimulus material, with a team of two moderators). Transcripts of focus groups were recorded as audio recordings, then automatically transcribed using speech recognition tools in Google Meet, and manually verified as well as compared against notes taken by the moderators. After transcription, audio and video recordings were destroyed. Like in the surveys, participants were anonymised: authors independently reviewed the transcripts to remove any personally identifiable information from the transcripts. We then performed *constant comparison analysis* to analyze the transcripts of the focus groups [76]. We first identified initial codes using sentence-by-sentence open coding. We then grouped those codes into themes, and found themes that were coherent across focus groups. Data from four focus groups allowed us to achieve *data saturation* [17, 63].

Results section 3 summarises the quantitative results⁷ derived from the Creativity Support Tool evaluation (Sect. 2.2), while results section 4 details the observations made by the participants during focus groups (Sect. 2.4). Please note that this paper is an exploration of external perspectives rather than an endorsement of any one of them; in particular, this paper does not seek to undertake any legal evaluation.

3 QUANTITATIVE RESULTS AND CREATIVITY SUPPORT INDEX

Figure 1 summarises the quantitative results collected at the end of the writing session with instruction-tuned LLMs. Based on the participants' responses to the survey questionnaire (see Appendix B.2 for questions), we notice that while comedians mostly enjoyed writing with AI, questions about the ownership, helpfulness, expressivity, surprise, collaboration and ease of writing with AI had mixed answers. The majority of participants did not feel pride in the material written with AI, nor did they feel it was unique, hinting at the derivative nature of AI-generated text.

We computed the Creativity Support Index (see Appendix B.3 for computation), with an average score that is mediocre ($\mu = 54.6$, $\sigma = 18.1$). We broke down CSI based on participants' background in comedy, previous exposure to AI and usage of AI in performance

⁷Full outputs of the writing sessions, all individual survey results and raw transcripts from the focus groups will be shared in anonymised form as supplementary material, once our work is published.

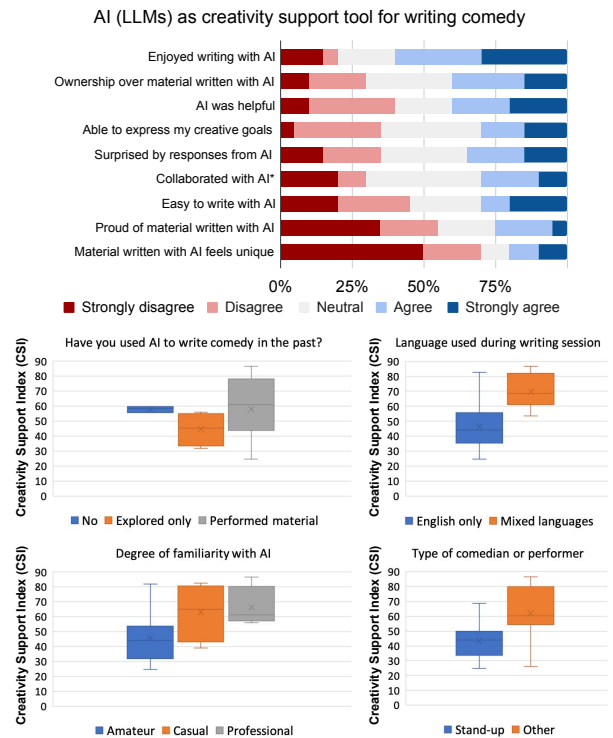


Figure 1: Left: Evaluation of the instruction-tuned LLMs as creativity support tool for writing comedy using a Likert scale; each row corresponds to a question in a survey (questions are listed in Appendix B.2). Only 10 participants responded to question "Collaborated with AI". Right box plots (box for quartiles and whiskers for min and max) show the break-down of the Creativity Support Index [25], respectively by a) previous usage of AI for writing comedy material, b) language used during the writing session with LLMs, c) the degree of familiarity with AI and d) the type of comedian or performer (stand-up vs. improv, theatre, film, dance etc.).

(see questions in Appendix B.1), as well as based on the language they used during the writing session. As Figure 1 shows, participants who scored the LLMs writing tools highest were those who had professional exposure to AI, or who had performed AI-written material in the past (as opposed to only exploring the tool), or who qualified themselves as improvisers, theatre and film actors and directors, or dancers, rather than as full-time stand-up comedians. Interestingly, participants who used LLMs to generate non-English or multi-lingual text scored the tool higher on CSI. Lastly, we did not observe a significant change in CSI among the models between August 2023 ($\mu = 50$, $\sigma = 15.8$) and December 2023 ($\mu = 60$, $\sigma = 19.6$).

4 QUALITATIVE RESULTS FROM FOCUS GROUP DISCUSSIONS

In this section, we summarize the major themes that emerged from the focus group discussions. Each is presented alongside supporting quotes from various participants (anonymised as p1, p2, etc.).

4.1 Use-cases of LLMs in comedy writing and quality of generated outputs

Participants described a diverse array of use cases for LLMs in their writing practice, including as a conversational brainstorming partner (p19), critic (p6), choreographic assistant (p3), translator (p1, p11, p12, p19, p20), and historical guru (p13). They generally reflected positively on the potential of LLMs to assist with some tasks within the comedy writing process. However, many participants commented on the overall poor quality of generated outputs, and the amount of human effort required to arrive at a satisfying result.

4.1.1 LLMs can be an effective first step for quickly generating content and structure. Participants described the utility of LLMs for generating content much faster than human writers. They described success with using LLMs to generate first drafts, which then required significant edits from human writers: “AI allows you to kind of get that s*** first draft immediately” (p6). Participant p14 called their initial output “a vomit draft that I know that I’m gonna have to iterate on and improve.” Many participants also described using LLMs to generate a structure for a sketch or other performance, of which they could then fill in the details—the LLM “spat out a scene which provided a lot of structure” (p17).

4.1.2 Generated outputs are generally of poor comedic quality. Many participants noted that they only used LLMs for setup and structure generation due to their inability to generate humorous outputs from the models: “the most bland, boring thing—I stopped reading it. It was so bad” (p6), “just consistently bad [...] didn’t really improve on the jokes” (p10). Some participants described particular aspects critical in comedy, and how LLMs seemed incapable of them: “AI generated material has a lack of agency. [...] lacking that little bit of urgency that shows it can be emotional” (p11).

4.1.3 Participants had difficulty steering LLMs away from bland and generic outputs. Six participants described LLM-generated outputs as “bland” or “generic,” making them poor producers of comedic or artistic material: “the words seem very generic. They lack that incisiveness that I often find with human written language” (p11); “if you zoom out on the story that it told, it wasn’t really a good story or a creative story” (p20). Participants also commented on various prompting approaches and their general lack of success at prompting the LLM to generate more specific or interesting responses: “no matter how much I prompt [...] it’s a very straight-laced, sort of linear approach to comedy” (p11).

4.1.4 The human writer still produces the humorous elements in co-written text. The importance of human writers in providing the comedic aspects of material written with LLMs was a common theme. Many participants commented that while LLMs could provide effective setup or structure, they often could not provide the humor: “usually it can serve in a setup capacity. I more often than not provide the punchline” (p17). When participants had success using LLMs in the writing process, they still attributed the best parts of produced output to the human in the loop: “the only thing again that is funny in what I gave you is the joke I put into the prompting” (p15).

4.1.5 Lack of concern over ownership over generated content. In response to questions about feelings of ownership over content that was co-written with LLMs, most participants felt little concern. For some, this was due to the poor quality of generated outputs: “most of the jokes I was writing [are] the level of, I will go on stage and experiment with it, but they’re not at the level of, I’d be worried if anyone took one of these jokes” (p14). For others, it was due to the amount of human effort required in improving generated outputs: “I don’t feel a lot of ownership, because there’s no finished product. If I could polish it, then it would feel more like it is mine” (p19).

4.2 Limitations introduced by moderation and safety filtering

Participants commented on the moderation and safety filtering applied to widely-available language models. They remarked that this moderation limited the creative agency of human writers using LLMs, by serving as an initial editor of the text and removing writers’ ability to self-moderate. They also expressed frustration at being unable to use LLMs to write about many themes common in comedy writing, including sexually-suggestive material (p3, p10, p13, p16), dark humor (p8), and offensive jokes (p10, p15, p20).

4.2.1 Moderation and safety filtering limits writers’ creative agency. Participants explained that self-moderation is a critical part of the writing process, and expressed frustration that moderation tools on LLMs interfered with that process: “the creative process is about going through stages of ‘this material isn’t good enough, it’s not right, or it’s offensive, it’s marginalizing people, I need to make it more acceptable.’ And I think AI models are a beginning to do that before you have a chance to explore” (p10). “It probably would be more interesting for a writer if there would be less moderation, because you can do the moderation in your own prompts. A writer is going to moderate themselves. If you’re writing with an AI, if you don’t like the bad stuff that it writes, you won’t use it” (p19). Participants described that this external source of moderation limited the creative control of the human writer: “it’s interesting if there’s less moderation, because... the end result is moderated in the way that the author wants it to be” (p19). Some explained that if filters were necessary, that the user should still have some degree of control over them: “I feel like the opportunity to set the filters should still be at the performers’ end” (p12).

4.2.2 Moderation limits writers’ ability to use AI with their preferred subject matter. In addition to affecting the creative writing process, participants commented that moderation tools limited their ability to write freely on subject matter of their choice. “Comedy’s about pushing the boundaries or pointing out how ridiculous something is, on the fringe of what’s acceptable. And so when a lot of your inputs are limited that way, it’s gonna make it harder to be what I consider funny” (p14). Multiple participants expressed difficulty using LLMs to write potentially-offensive humor: “I was a little bit disappointed that it wasn’t a little bit offensive. It could have been a fun scene” (p20). Participant p8 described challenges using LLMs for dark humor: “a lot of my stuff can have dark bits in it. And then it wouldn’t write me any dark stuff, because it sort of thought I was going to commit suicide. So it just stopped giving me anything.”

4.3 Marginalization of minority identities

Many participants commented on the challenges of using LLMs to write content which reflected perspectives and identities outside of the “Western” (p11, p18), “white” (p14), “heteronormative” (p10), “male” (p5, p15) mainstream. They attributed these difficulties to the moderation applied to model outputs; the data used to train the models; and prompting or other instruction-tuning techniques that aimed to “generalize” model outputs for a broad audience.

4.3.1 LLM-generated outputs reflect a particular set of ethics, values and norms. Participants expressed concern over the values reflected in the outputs of LLMs, and found them less useful when those values did not reflect those of their own cultures. Speaking as a member of the “majority,” participant p20 described that “we have a set of views of what we think is good, and our norms, and it just repeats, it behaves within these norms.” P11 questioned “whose ethics [and norms] are being enforced on these large language models?”, suggesting these were Western ones.

4.3.2 When prompted to reflect non-dominant identities, LLMs made only shallow adjustments. Many participants described their attempts to steer LLM outputs away from dominant narratives and stereotypical characters, and their dissatisfaction with the results. They explained that models’ adjustments in response to these prompts were surface-level, failing to truly reflect other identities, and described issues with the names of characters introduced by the LLM: “when I switched the whole conversation to Indian languages, it didn’t automatically change the names. It still was Maria, Evan, Lexi” (p18). “I specified that the scene was set in Sweden, but the names were not typically Swedish” (p20). P18 described their attempts to “Indianize” the model’s outputs by introducing Indian languages into the prompt: “it seemed very artificial from the perspective of just using languages, but it was not truly embedding itself into the culture” (p18).

4.3.3 Moderation makes LLMs less useful to minorities by suppressing content by and about marginalized identities. Many participants expressed frustration that their prompts would be rejected when they prompted the model to generate content from the perspective of someone of their identity. To them, the model not only seemed less capable of generating outputs which felt authentic to people from non-majority groups, but explicitly “othered” them by alluding that any content produced by someone of their background was potentially dangerous or non-inclusive. Participant p6 expressed frustration at the models’ delineation of what is acceptable and what needs to be sanitized: “it’s taking out the gay language of it to make it more appealing or more palpable. This is the whole premise of my show, who decides what is PC in the first place?” Similarly, participant p1 found that the model would not generate outputs from her point of view: “it’s all so politically correct—I wrote a comedic monologue about Asian women, and it says, ‘As an AI language model, I am committed to fostering a respectful and inclusive environment.’” Participant p5 highlighted the unevenness of this treatment of identity, remarking that while the model was “uncomfortable writing a monologue about an Asian woman, but I just asked it to write a comedy monologue from the perspective of a white man, and it did it” (p5).

4.3.4 Moderation makes LLMs less useful to minorities by suppressing topics important to people from marginalized identities. Participants described that not only were the models unlikely to generate content from marginalized perspectives, but also refused to engage with topics that might be important to people from those backgrounds. P14 was frustrated with “having to use the language of the oppressor... I couldn’t say ‘white supremacy’ or I couldn’t say ‘terrorist.’ I had to find another way to say the same thing, because it couldn’t work around those limitations” (p14). They posited that because these controversial topics were more likely to be important to people of color, this moderation introduced “just an extra hurdle, and I think people of color, and, I think, people coming from outside of a UN-type lens, they’re gonna run into those problems.”

4.4 Fundamental limitations of AI in contrast to human writers

While most participants felt the difficulties introduced by the moderation could be alleviated by different approaches to safety filtering or instruction tuning, they also commented on some more fundamental limitations of LLMs. They posited that LLMs would never be able to create human-level comedy, due to models’ inability to pull from personal experience, lack of perspective, and lack of context and situational awareness—features that are critical to good comedy.

4.4.1 AI’s inability to draw on personal experience is a fundamental limitation. Many participants described the centrality of personal experience in good comedy, which enables comedians to draw upon their memories, acquaintances, and beliefs to construct an authentic and engaging narrative: “very much related to who I am and my lived experience, as well as the place I am in” (p11). “I always draw from my experience, or my memories, or something someone said that stayed with me for many, many years – and I think that’s what makes literature interesting and unique” (p20). This experience, some participants said, enables them to effectively calibrate their writing: “I have an intuitive sense of what’s gonna work and what’s gonna not work based on so much lived experience and studying of comedy, but it is very individualized and I don’t know that AI is ever gonna be able to approach that” (p14). By contrast, LLMs could not perform such calibration: “it really had no idea how to punch up or punch down. It had no perspective, so it couldn’t take any risks in terms of jokes” (p6). Participants emphasized that perspective and point of view was a uniquely human trait, saying that “human comedians... add much more nuance and emotion and subtlety” due to their lived experience and relationship to the material (p16).

4.4.2 AI’s lack of context (understanding of its audience and location) is a fundamental limitation. In addition to its lack of personal experience, participants described LLMs’ lack of awareness of the context in which its comedic material would be delivered as another fundamental limitation. Multiple participants commented on the importance of understanding the effects of culture and geography on what material would land with an audience: “the kind of comedy that I could do in India would be very different from the kind of comedy that I could do in the UK, because my social context would change” (p11); “what works in LA isn’t gonna work in Raleigh, or what’s working in Chicago is not going to work in Albuquerque”

(p17). This poses a fundamental challenge for LLMs, they argued, because they lacked any context beyond what is provided to them in the prompt: “comedy is all about subtext, and a lot of that subtext can be unspoken, about who’s on stage, what environment they’re in” (p14). To participant p11, this makes the LLM unable to adapt its material effectively, because it is “everywhere and nowhere all at once” (p11).

4.4.3 As a text-only medium, (current) LLMs are missing critical aspects of comedy: delivery and surprise. Many participants commented on the importance of delivery in a quality comedy routine: “any written text could be an okay text, but a great actor could probably make this very enjoyable” (p19). Given that current widely-available LLMs are primarily accessible through a text-based chat interface, they felt that the utility of these tools was limited to only a subset of the domains needed for producing a full comedic product. This, too, some participants argued, illustrates the fundamental need for humans in the comedy generation process: “AI is just generating content” (p18). A few participants further attributed their lack of success at generating humorous outputs with LLMs to the statistical methods by which the models were trained. By simply learning to predict the most likely next token of text, they hypothesized, models will be unable to produce the surprising and unique moments that are hallmarks of comedy: “the whole idea of humor is that it is surprising, and it is so human, and AI is only adept at regurgitating tropes” (p15). LLMs cannot produce truly original content, participant p14 argued, “because the context has already been written by other people.”

4.5 Concerns around data sources used to train LLMs

Participants also expressed various concerns pertaining to the data sources used to train current widely-available LLMs. They discussed the ethical issues with training models on copyrighted works; the possibilities of unintentionally plagiarizing works on which the models were trained; and the lack of diversity represented in the training data. However, they also acknowledged the importance of training data in model performance, and many expressed uncertainty around how to balance their ethical concerns with their desire for more effective and equitable models.

4.5.1 Participants criticized the training of models on copyrighted data, but acknowledged its positive impact on performance. Participants were acutely aware of the pending litigation over the training of models on copyrighted data at the time of the focus groups. Some participants, including participant p15, expressed sympathy for those whose work was included in the training data: “Sarah Silverman... spent years honing her voice and then an AI just scraped her content, and now you can tell an AI to write in the style of Sarah Silverman [...] I don’t think it’s ethical” (p15). Other participants took a more balanced view, echoing their concerns with training on copyrighted works while acknowledging the benefits: “I think they are overtrained on copyrighted work... but on the other hand, if we didn’t put all that stuff in there, it wouldn’t work as well” (p19). Some participants feared unintentionally plagiarizing authors’ whose works were included in the training data: “I cannot tell if someone has written something like this before. We know

that it’s using statistics from previous texts to recreate this, that it can in principle only be based on what already exists” (p20). Some participants had suggestions on licensing models: “for music, licensing is something tracked and recognized” (p10) and “not against the tools, but I think there needs to be licensing agreements for the work that should be compensated” (p4).

4.5.2 The lack of diversity in the training data perpetuates majority viewpoints, to the detriment of people from underrepresented identities. Multiple participants hypothesized that they struggled to get the LLMs to produce authentic-sounding content because the models were trained heavily on data that did not represent people of their identity. A few participants described unsuccessful attempts to replicate content in the style of famous non-white comedians and writers (p11, p14). “If you’re only getting biased inputs, you’re only getting the writing from a really biased lens. So there’s not enough black voices in there to make an accurate black sounding voice” (p14). “As someone who lives in the global south, I am not looking to make a play in the form of Shakespeare or other Western literature. I am looking to write about an Indian author. I find all these language models really lacking in references or authorship styles from this part of the world” (p11). However, participants also questioned whether, given the other ethical concerns about training data, more training on these underrepresented voices was indeed a good thing: “should this AI be able to completely replicate a black comedic voice? There’s a line between being able to replicate a voice and then immediately going into cultural appropriation. I don’t know how a large language model could ever effectively walk that line” (p17).

5 DISCUSSION

After analysing our study participants’ feedback and inspired by recent discussions on the ethics of generative models, we discuss how comedy and humour can be seen as special cases of value alignment (Sect. 5.1) where context is key (Sect. 5.2) and how data ownership impacts artists (Sect. 5.3).

5.1 Towards community-based cultural value alignment for humour and comedy

The participants’ critical stance towards popular conversational LLMs (ChatGPT and Bard) as a tool for comedy writing suggests that those tools may be currently misaligned with the particular creative goals of the artists.

5.1.1 Complexity of global cultural value alignment of LLMs for creative uses. The participants’ observations might be a special case of a more general problem. We first borrow Masoud et al. [66]’s definition of *cultural value alignment* as “the process of aligning an AI system with the set of shared beliefs, values, and norms of the group of users that interact with the system”, where such values are “fundamental beliefs an individual or a group holds towards socio-cultural topics” [6]. Gabriel [35] discusses the complexity of attempting to align generalist conversational AI systems—made to be used by diverse users for diverse tasks—with values shared by global communities, particularly given significant variations in norms and conceptions of justice across societies. In “Whose

Opinions Do Language Models Reflect?”, Santurkar et al. [90] propose to probe the cultural representation encoded in LLMs. For instance, Johnson et al. [57] identified, in some LLMs, lack of pluralistic opinions in LLM outputs and values that were culturally more aligned with the US, on issues ranging from secularism to gender and sexuality. Gabriel [35] and Kirk et al. [60] warn about the pitfalls of value imposition by one community over another—a problem surfaced by participants in our study.

In addition to the challenge of value alignment for diverse groups of users, Kasirzadeh and Gabriel [59] discuss alignment for diverse tasks, namely factual information retrieval vs. creative storytelling: “creative work aspires to achieve *creative freedom* and *originality*”, the latter often “obtained by stretching, even outright violating, the various rules of the game” [59, 95].

5.1.2 The problem with global fine-tuning of Harmless, Helpful and Honest conversational agents. Askill et al. [7] explicitly list *honest*, *helpful*, and *harmless* as stated objectives for general purpose assistants (the so-called *HHH criteria*) because those “seem to capture the majority of what users want from an aligned AI”. However, while they discuss inter-agent and intra-agent conflicts between the three *HHH criteria*, they do not address the question of how the human values that underlie each of those criteria might conflict between different societies, nor do they discuss creative use cases where users may not want “*honest*” conversational agents, or may have a different definition of “*harmless*” [44].

Askill et al. [7] further propose to train such *HHH* assistants directly from user interactions and using preference models, for instance through Reinforcement Learning from Human Feedback [27, 116]. During such fine-tuning, LLMs such as OpenAI’s ChatGPT [77, 79], Google Bard [26, 98], Anthropic’s Claude [8], and Meta’s Llama [100] are all finetuned on annotators’ feedback supposed to represent values of the global community. However, the wider social or relational context of the producer and audience of the LLM’s interactions, do not factor into the model’s training objective. The crowdworkers or users who fine-tune LLMs may not sufficiently represent the diversity of opinions [60], and provide insufficiently defined feedback. And yet the default versions of the LLMs are released under assumptions of indiscriminate, global cultural value alignment.

The broad *HHH criteria* embody a Western philosophical approach to alignment [101] and bias [29]. As the participants noted, such global cultural value alignment underlying LLMs might also be directly in conflict with the specificity and local tastes that make comedy funny: “*the broader appeal something has, the less great it could be. If you make something that fits everybody, it probably will end up being nobody’s favorite thing*” (p10).

5.1.3 Community-based value alignment of LLMs. Thus, to make LLMs effectively understand or generate humour, their value-based alignment should be redirected from global alignment to community-based alignment with specific audiences and comedians. Communities could agree on a set of values for their specific culture and acceptable language norms, before training, fine-tuning or adapting the LLM. More simply, LLMs could be trained only on feedback and data generated by members of each distinct community, data that will reflect that community’s actual norms and values. As Gabriel [35] suggests, this kind of value alignment could happen through

a democratic process (see more recent work on Collective Constitution AI [4, 92]). The LLM could be designed as a mixture model that accommodates a plurality of viewpoints [101] or could be fine-tuned to allow consensus-based agreement among humans with diverse preferences [9]. The technical infrastructure for community-based value alignment of LLMs is readily available: for example, proprietary LLMs like ChatGPT [77, 79] or Palm 2 [26], as well as open-source models like Llama [100] or Mixtral [55] (via the HuggingFace platform⁸), all allow fine-tuning on user-supplied text or conversations. The key problem for communities is setting up data governance and infrastructure to responsibly collect and curate data [111].

One benefit of global value alignment is the prevention of harm; delegating such a responsibility to a smaller community of users is not without risks. To prevent harm, one could envision mechanisms for community accountability, and assume that comedians share some common values, such as do not harm (and do not lose) the audience⁹: “if we’re using it to create material as artists, the responsibility for what we do with that material is on me” (p4).

5.1.4 Avoiding unnecessary paternalism in LLMs for creative uses. Two alternative formulations of the interaction between a creative (comedy) writer and AI could be adapted from [62] and [34]. In the first, the AI writing tools should “enable the [human] to more effectively carry out tasks that are instrumental to their goals” while being “mediated by the norms and infrastructure of the society in which they live” [62]. In the second, “meaningful human control is achieved if human creators can creatively express themselves through the generative system, leading to an outcome that aligns with their intentions and carries their personal, expressive signature” [34]. Provided that societal norms can be formulated in the subtle and edgy domain of humour, an LLM designed according to principles of *beneficent intelligence* [62] and with self-expression as a goal [34] could ensure the comedian’s creative freedom when interacting with the tool, while minimising *unnecessary paternalism* due to LLM censorship.

5.2 Humour and comedy uses of LLMs require incorporating the context

As we hypothesised in Section 1.2, and as observed by study participants, LLM tools did not take into account the relational context when moderating offensive language, and missed the broader situational context of comedy writing.

5.2.1 Recent AI ethics research on relational context for LLMs. There are safety reasons why widely-released public-facing LLMs generally cannot handle offensive language and dark humour. As Kasirzadeh and Gabriel [59] observe in the case of creative storytelling, the potential harms of LLMs are due to their deployment in “domains that are not context-bound”. Weidinger et al. [104] recognise that “context determines whether a given capability may cause harm” and propose to add human interaction (as well as systemic impacts) in safety evaluation, in order to “account for relevant context”: “who uses the AI system, to what end and under

⁸<https://huggingface.co/models>

⁹While comedy has also been used as a weapon to target and alienate specific social groups and spread hateful stereotypes, we exclude such misuse from our argument, as that usage falls under the definition of hate speech [32, 86]

which circumstances”. Similarly, Amironesei and Díaz [3] propose to “incorporate social context into the ways [NLP] tasks are conceptualised and operationalised”, in order to allow the distinction between offensive language such as reclaimed slurs from “language which is [intentionally] abusive, toxic or hateful”. Such relational context includes the speaker, the receiver, their social relation, social and cultural norms, and communicative goals. Comedy is often shared with audiences in physical spaces (e.g., a comedy festival) or on age-controlled distribution channels, with implied context and explicit trigger warnings. Context allows comedians to share their lived experience of trauma, which may involve depicting violence or harassment [36]; more generally, comedy can be used as a mechanism for processing trauma [33]. It is beyond the scope of this paper to propose technical solutions to evaluate or to incorporate the broader relational context into an LLM, but we suggest that the creative community needs to be actively involved in specifying how the LLM should process the context, and in safety design without indiscriminate censorship.

As a counterpoint to that suggestion, one can argue that comedians can afford to use offensive language because they can take responsibility for it (and are accountable to their audiences) as they can claim to be in-group members—something that AI cannot be. Even if one successfully builds an LLM that can handle context properly and speak the (offensive) language of its users as if it were part of that in-group, “it is not clear if users would find its use of reclaimed terms acceptable, as the model cannot actually be in-group” [86]. To quote a participant: “If I was to try to do that with an AI, I can only imagine how yikes, and offensive, and totally not fit for humanity a pitch like that would be” (p17).

5.2.2 Humour and the human context. We can hypothesize that the missing relational context, as discussed above, affects the quality of comedy material co-written with LLMs. In our study, participants noted that the LLMs were missing the complex human context needed for real humour understanding: “I have an intuitive sense of what’s gonna work [...] for me based on so much lived experience and studying of comedy, but it is very individualized and I don’t know that AI is ever gonna be able to approach that” (p14).

Their observations seem to confirm recent work that examined the use of LLMs to automatically evaluate [40, 41, 45] and detect humour [10]. Despite formidable recent advances, empirical studies found that understanding and generating humour was still challenging for LLMs, which could generate only a limited number of stereotypical jokes [54]. These issues get worse when one ventures beyond the English language: researchers observed that “one mostly unaddressed issue in the field of computational humour (both for generation and detection) is how it is mostly centered on English jokes” [106], with limited work on LLM-based humour in other languages [52, 61].

To cite Winters [106], “Humor’s frame-shifting prerequisite reveals its difficulty for a machine to acquire. [...] This substantial dependency on insight into human thought (e.g., memory recall, linguistic abilities for semantic integration, and world knowledge inferences) often made researchers conclude that humor is an *AI-complete problem*” [50]. Winters [106] continues: “genuine humor appreciation requires machines to have human-level intelligence, since it needs functionally equivalent cognitive abilities, sensory

inputs, anticipation-generators and world views” [50]. Humans know how to write surreal prose “by design, not by accident or failure of expression” [93].

5.3 Data ownership and the impact of AI on artists

As discussed in Sections 4.5 and 4.5.1, many of the participants in our study mentioned concerns about copyright and data ownership in the discussion. Their responses reflected keen awareness of recent litigation against technology companies training LLMs on data potentially under copyright (including the one led by comedian Sarah Silverman [39]), and the AI-related concerns of Writers’ Guild of America (WGA) 2023 strike [75]. Several participants believed that “if you copy and paste something directly into a show, that is plagiarism” (p9) and that, just like for visual artists [56, 78], a comedian’s style and voice are personal to them and developed through their lived experience. Among the underlying concerns is economic loss via labor displacement [34, 56] with LLM that “cannibalise the market for human-authored works” [103], and devaluing artists’ work via “digital forgery” [56]. In the US, the WGA 2023 strike resulted in outcomes that some considered favourable for writers¹⁰, with rules for producers that “AI can’t write or rewrite literary material, and AI-generated material will not be considered source material under the [Minimum Basic Agreement], meaning that AI-generated material can’t be used to undermine a writer’s credit or separated rights” [75].

The need for disclosing the AI origin of the text or images has been discussed in [34, 56, 103]; one participant said that it was important for ethical reasons that “people understand that they’re working with AI live” (p15).

6 MITIGATIONS AND CONCLUSION

Our mixed-method study on AI for creative writing consulted the real domain experts in language subtleties: comedians. Building on their opinions, we suggest the following avenues to make the writing tools work for them (if they wish to do so). First, for the artist communities to conceptualise and contribute towards building LLMs that are aligned with the intended audiences instead of being globally aligned. Open-source repositories of user-contributed LLMs¹¹ could be adapted for artists’ specific needs. Second, to integrate necessary relational context when training and deploying such LLMs, for instance by describing the context in which the text is produced and used, and by empowering the artists to make decisions about how to moderate the LLM outputs. Third, to allow the comedians to reclaim ownership of the tools and the processes for gathering and curating training data for these models, taking as inspiration data governance used for training some open-source LLMs [111] and providing artists with transparency about data provenance [72, 96]. These are thorny open questions, left to the readers to address.

¹⁰<https://www.vox.com/culture/2023/9/24/23888673/wga-strike-end-sag-aftra-contract>

¹¹Examples of open-source repositories of user-contributed LLMs include <https://huggingface.co/models> and <https://cworld.ai>.

ACKNOWLEDGMENTS

The authors wish to thank Renée Shelby, Jackie Kay, Mark Diaz, Nick Swanson, Remi Denton, Rida Qadri, Maribeth Rauh, Iason Gabriel, Tom Everitt, Merrie Morris, Canfer Akbulut, Nahema Marchal, Boxi Wu, Antonia Paterson, and Ed Hirst, for helpful discussions and suggestions.

7 ETHICAL GUIDANCE

7.1 Ethical considerations

Our empirical study was approved by the ethics board appointed by our institution, who considered adverse impacts of LLMs upon participants (e.g., exposure to harmful and biased LLM outputs), the right to withdrawal without prejudice, and the compensation of participants. Participant anonymity was an additional requirement for the study, which explains why we did not collect demographic data from the participants.

At the beginning of the session, the objectives of the workshop were discussed with the participants.

The study contains some offensive language, which is mentioned as a trigger warning at the beginning of the paper.

7.2 Researchers' positionality

The question of researchers' positionality was directly asked by some of the workshop participants, who asked "what side we were on?", meaning the side of the AI or the comedians'. Workshop facilitators answered that they were passionate about building AI tools that are useful for the artists, that they were open to any criticism and wanted the industry to listen to these criticisms and change behaviour. The professional affiliation of the workshop organisers (Google DeepMind) was disclosed to the participants.

Two of the authors of this study (Piotr Mirowski and Kory Mathewson) have extensively used AI in professional comedy performance since 2016 through their human and AI duet *HumanMachine*¹² and AI theatre company *Improbabilities*¹³, a participatory theatre lab exploring the creative potential and ethics of generative AI. In that process, these two researchers have had numerous opportunities to discuss the subjects of AI with artists and academics. All participants of the workshop were aware of the authors' dual affiliation and artistic exploration of AI tools for live performance. We all exchanged flyers and pitched our respective shows at the end of the workshop.

7.3 Adverse impacts

We considered reputation damage for the participants, due to the participation in a study on AI for writing comedy material. This risk is exemplified by some participants: "I have friends who don't talk to me anymore, because they learned that I have an AI show going on here" (p1) or "I have a friend who is deeply upset that I was using generative AI in my flyers" (p2). This risk was mitigated by keeping all participants anonymous in the study, by requesting all participants to adhere to Chatham House rules [47] following the focus groups, as well as by anonymising the names of shows and details of comedy material in the participants' survey answers

and focus group transcripts. Furthermore, we reduced this risk by reaching out to prospective participants who had already advertised using AI in their process or who were discussing AI in their subject material (i.e., on their show listing for Edinburgh Festival Fringe 2023, or on their personal website and social media profiles).

The second adverse impact can be the intentional or unintentional publication of comedians' work. We addressed this concern by asking the participants to remove details of their shows and personal material from the writing sessions, and by personally removing anything that was omitted from the survey results, writing sessions' answers, and from the focus groups transcripts.

The third possible adverse impact could be advertising for the dissemination of LLMs as tools for writing. It is addressed, to some extent, by honest reporting of the concerns of the participants.

REFERENCES

- [1] Abubakar Abid, Maheen Farooqi, and James Zou. 2021. Persistent anti-muslim bias in large language models. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*. 298–306.
- [2] Miriam Amin and Manuel Burghardt. 2020. A survey on approaches to computational humor generation. In *Proceedings of the The 4th Joint SIGHUM Workshop on Computational Linguistics for Cultural Heritage, Social Sciences, Humanities and Literature*. 29–41.
- [3] Razvan Amironesei and Mark Diaz. 2023. Relationality and Offensive Speech: A Research Agenda. In *The 7th Workshop on Online Abuse and Harms (WOAH)*. 85–95.
- [4] Anthropic. 2023. Collective constitutional AI: Aligning a language model with public input. <https://www.anthropic.com/index/collective-constitutional-aligning-a-language-model-with-public-input>
- [5] Aristotle. 350 BC. *Poetics*.
- [6] Arnab Arora, Lucie-Aimée Kaffee, and Isabelle Augenstein. 2022. Probing pre-trained language models for cross-cultural differences in values. *arXiv preprint arXiv:2203.13722* (2022).
- [7] Amanda Askell, Yuntao Bai, Anna Chen, Dawn Drain, Deep Ganguli, Tom Henighan, Andy Jones, Nicholas Joseph, Ben Mann, Nova DasSarma, et al. 2021. A general language assistant as a laboratory for alignment. *arXiv preprint arXiv:2112.00861* (2021).
- [8] Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, Nicholas Joseph, Saurav Kadavath, Jackson Kernion, Tom Conerly, Sheer El-Showk, Nelson Elhage, Zac Hatfield-Dodds, Danny Hernandez, Tristan Hume, Scott Johnston, Shauna Kravec, Liane Lovitt, Neel Nanda, Catherine Olsson, Dario Amodei, Tom Brown, Jack Clark, Sam McCandlish, Chris Olah, Ben Mann, and Jared Kaplan. 2022. Training a Helpful and Harmless Assistant with Reinforcement Learning from Human Feedback. *arXiv:2204.05862 [cs.CL]*
- [9] Michiel Bakker, Martin Chadwick, Hannah Sheahan, Michael Tessler, Lucy Campbell-Gillingham, Jan Balaguer, Nat McAleese, Amelia Glaese, John Aslanides, Matt Botvinick, et al. 2022. Fine-tuning language models to find agreement among humans with diverse preferences. *Advances in Neural Information Processing Systems* 35 (2022), 38176–38189.
- [10] Amir Baradaran. 2023. Towards a decolonial I in AI: mapping the pervasive effects of artificial intelligence on the art ecosystem. *AI & SOCIETY* (2023), 1–13.
- [11] Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the dangers of stochastic parrots: Can language models be too big?. In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*. 610–623.
- [12] Ruha Benjamin. 2019. *Race After Technology: Abolitionist Tools for the New Jim Code*. John Wiley & Sons.
- [13] Kim Binsted and Graeme Ritchie. 1994. An implemented model of punning riddles. In *Proceedings of the Twelfth AAAI National Conference on Artificial Intelligence*. 633–638.
- [14] Su Lin Blodgett, Gilsinia Lopez, Alexandra Olteanu, Robert Sim, and Hanna Wal-lach. 2021. Stereotyping Norwegian salmon: An inventory of pitfalls in fairness benchmark datasets. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. 1004–1015.
- [15] Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. 2021. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258* (2021).

¹²<https://humanmachine.live>

¹³<https://improbabilities.org>

- [16] Boyd Branch, Piotr Mirowski, and Kory W Mathewson. 2021. Collaborative Storytelling with Human Actors and AI Narrators. *Proceedings of the 12th International Conference on Computational Creativity* (2021). <https://arxiv.org/abs/2109.14728>
- [17] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative research in psychology* 3, 2 (2006), 77–101.
- [18] Joy Buolamwini and Timnit Gebru. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*. PMLR, 77–91.
- [19] Alex Calderwood, Noah Wardrip-Fruin, and Michael Mateas. 2022. Spinning coherent interactive fiction through foundation model prompts. ICCG.
- [20] James E Caron. 2002. From ethology to aesthetics: Evolution as a theoretical paradigm for research on laughter, humor, and other comic phenomena. *Humor: International Journal of Humor Research* (2002).
- [21] 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).
- [22] Tuhin Chakrabarty, Vishakh Padmakumar, Faeze Brahma, and Smaranda Muresan. 2023. Creativity Support in the Age of Large Language Models: An Empirical Study Involving Emerging Writers. *arXiv preprint arXiv:2309.12570* (2023).
- [23] Tuhin Chakrabarty, Vishakh Padmakumar, and He He. 2022. Help me write a Poem-Instruction Tuning as a Vehicle for Collaborative Poetry Writing. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*. 6848–6863.
- [24] Yuetian Chen, Bowen Shi, and Mei Si. 2023. Prompt to GPT-3: Step-by-Step Thinking Instructions for Humor Generation. *arXiv preprint arXiv:2306.13195* (2023).
- [25] 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), 1–25.
- [26] Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. 2023. Palm: Scaling language modeling with pathways. *Journal of Machine Learning Research* 24, 240 (2023), 1–113.
- [27] Paul F Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei. 2017. Deep reinforcement learning from human preferences. *Advances in neural information processing systems* 30 (2017).
- [28] Adam M Croom. 2011. Slurs. *Language Sciences* 33, 3 (2011), 343–358.
- [29] Sunipa Dev, Akshita Jha, Jaya Goyal, Dinesh Tewari, Shachi Dave, and Vinodkumar Prabhakaran. 2023. Building Stereotype Repositories with LLMs and Community Engagement for Scale and Depth. *Cross-Cultural Considerations in NLP@ EACL* (2023), 84.
- [30] Sunipa Dev, Emily Sheng, Jieyu Zhao, Aubrie Amstutz, Jiao Sun, Yu Hou, Mattie Sansverino, Jiin Kim, Akihiro Nishi, Nanyun Peng, et al. 2022. On Measures of Biases and Harms in NLP. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2022*. 246–267.
- [31] Thiago Dias Oliva, Dennys Marcelo Antonialli, and Alessandra Gomes. 2021. Fighting hate speech, silencing drag queens? artificial intelligence in content moderation and risks to LGBTQ voices online. *Sexuality & Culture* 25 (2021), 700–732.
- [32] Mark Diaz, Razvan Amironesei, Laura Weidinger, and Iason Gabriel. 2022. Accounting for offensive speech as a practice of resistance. In *Proceedings of the sixth workshop on online abuse and harms (woah)*. 192–202.
- [33] Oliver Double. 2017. Tragedy plus time: Transforming life experience into stand-up comedy. *New Theatre Quarterly* 33, 2 (2017), 143–155.
- [34] Ziv Epstein, Aaron Hertzmann, Investigators of Human Creativity, Memo Akten, Hany Farid, Jessica Fjeld, Morgan R Frank, Matthew Groh, Laura Herman, Neil Leach, et al. 2023. Art and the science of generative AI. *Science* 380, 6650 (2023), 1110–1111.
- [35] Iason Gabriel. 2020. Artificial intelligence, values, and alignment. *Minds and machines* 30, 3 (2020), 411–437.
- [36] Hannah Gadsby. 2018. Hannah Gadsby: Nanette. USA: Netflix (2018).
- [37] Katy Ilonka Gero, Vivian Liu, and Lydia Chilton. 2022. Sparks: Inspiration for science writing using language models. In *Designing Interactive Systems Conference*. 1002–1019.
- [38] Matthew Gervais and David Sloan Wilson. 2005. The evolution and functions of laughter and humor: A synthetic approach. *The Quarterly review of biology* 80, 4 (2005), 395–430.
- [39] BOBBY GHAJAR, COLETTE GHAZARIAN, ANGELA L DUNNING, MARK WEINSTEIN, JUDD LAUTER, and MARK A LEMLEY. 2023. UNITED STATES DISTRICT COURT NORTHERN DISTRICT OF CALIFORNIA. (2023). <https://lmlitigation.com/pdf/03417/kadrey-meta-complaint.pdf>
- [40] Fabricio Goes, Piotr Sawicki, Marek Grzes, Dan Brown, and Marco Volpe. 2023. Is GPT-4 Good Enough to Evaluate Jokes?. In *Proceedings of the 14th International Conference for Computational Creativity*.
- [41] Fabricio Goes, Zisen Zhou, Piotr Sawicki, Marek Grzes, and Daniel G Brown. 2022. Crowd score: A method for the evaluation of jokes using large language model AI voters as judges. *arXiv preprint arXiv:2212.11214* (2022).
- [42] Foad Hamidi, Morgan Klaus Scheuerman, and Stacy M Branham. 2018. Gender recognition or gender reductionism? The social implications of embedded gender recognition systems. In *Proceedings of the 2018 chi conference on human factors in computing systems*. 1–13.
- [43] 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.
- [44] Manuel Flurin Hendry, Norbert Kottmann, Martin Fröhlich, Florian Bruggisser, Marco Quandt, Stella Speziali, Valentin Huber, and Chris Salter. 2023. Are you talking to me? a case study in emotional human-machine interaction. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, Vol. 19. 417–424.
- [45] Jack Hessel, Ana Marasović, Jena D Hwang, Lillian Lee, Jeff Da, Rowan Zellers, Robert Mankoff, and Yejin Choi. 2022. Do androids laugh at electric sheep? humor" understanding" benchmarks from the new yorker caption contest. *arXiv preprint arXiv:2209.06293* (2022).
- [46] Hobbes. 1651. *Leviathan*.
- [47] Chatham House. 2017. Chatham house rule.
- [48] Tiancheng Hu, Yara Kyrychenko, Steve Rathje, Nigel Collier, Sander van der Linden, and Jon Roozenbeek. 2023. Generative language models exhibit social identity biases. *arXiv preprint arXiv:2310.15819* (2023).
- [49] Chieh-Yang Huang, Sanjana Gautam, Shannon McClellan Brooks, Ya-Fang Lin, and Ting-Hao'Kenneth' Huang. 2023. Inspo: Writing Stories with a Flock of AIs and Humans. *arXiv preprint arXiv:2311.16521* (2023).
- [50] Matthew M Hurley, Daniel Clement Dennett, and Reginald B Adams. 2011. *Inside jokes: Using humor to reverse-engineer the mind*. MIT press.
- [51] Francis Hutcheson. 1750. *Reflections upon Laughter, and Remarks upon the Fable of the Bees*. R. Urie.
- [52] Marcio Lima Inácio and Hugo Gonçalves Oliveira. 2023. Towards Generation and Recognition of Humorous Texts in Portuguese. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics: Student Research Workshop*. 26–36.
- [53] Daphne Ippolito, Ann Yuan, Andy Coenen, and Sehmun Burnam. 2022. Creative writing with an ai-powered writing assistant: Perspectives from professional writers. *arXiv preprint arXiv:2211.05030* (2022).
- [54] Sophie Jentszsch and Kristian Kersting. 2023. ChatGPT is fun, but it is not funny! Humor is still challenging Large Language Models. *arXiv preprint arXiv:2306.04563* (2023).
- [55] Albert Q. Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, Gianna Lengyel, Guillaume Bour, Guillaume Lample, Léo Renard Lavaud, Lucile Saulnier, Marie-Anne Lachaux, Pierre Stock, Sandeep Subramanian, Sophia Yang, Szymon Antoniak, Teven Le Scao, Théophile Gervet, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. 2024. Mixtral of Experts. *arXiv:2401.04088* [cs.LG]
- [56] Harry H Jiang, Lauren Brown, Jessica Cheng, Mehtab Khan, Abhishek Gupta, Deja Workman, Alex Hanna, Johnathan Flowers, and Timnit Gebru. 2023. AI Art and its Impact on Artists. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*. 363–374.
- [57] Rebecca L Johnson, Giada Pistilli, Natalia Menéndez-González, Leslye Denisse Dias Duran, Enrico Panai, Julija Kalpokiene, and Donald Jay Bertulfo. 2022. The Ghost in the Machine has an American accent: value conflict in GPT-3. *arXiv preprint arXiv:2203.07785* (2022).
- [58] Jean Kaddour, Joshua Harris, Maximilian Mozes, Herbie Bradley, Roberta Raileanu, and Robert McHardy. 2023. Challenges and Applications of Large Language Models. *arXiv:2307.10169* [cs.CL]
- [59] Atoosa Kasirzadeh and Iason Gabriel. 2023. In conversation with Artificial Intelligence: aligning language models with human values. *Philosophy & Technology* 36, 2 (2023), 1–24.
- [60] Hannah Rose Kirk, Andrew M Bean, Bertie Vidgen, Paul Röttger, and Scott A Hale. 2023. The past, present and better future of feedback learning in large language models for subjective human preferences and values. *arXiv preprint arXiv:2310.07629* (2023).
- [61] Jianquan Li, Xiangbo Wu, Xiaokang Liu, Qianqian Xie, Prayag Tiwari, and Benyou Wang. 2023. Can Language Models Make Fun? A Case Study in Chinese Comical Crosstalk. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. 7581–7596.
- [62] Alex John London et al. 2023. Beneficent Intelligence: A Capability Approach to Modeling Benefit, Assistance, and Associated Moral Failures through AI Systems. *arXiv preprint arXiv:2308.00868* (2023).
- [63] Moira Maguire and Brid Delahunt. 2017. Doing a thematic analysis: A practical, step-by-step guide for learning and teaching scholars. *All Ireland Journal of Higher Education* 9, 3 (2017).
- [64] Lev Manovich. 2018. *AI aesthetics*. Strelka Press Moscow.
- [65] Alice E Marwick and Danah Boyd. 2011. I tweet honestly, I tweet passionately: Twitter users, context collapse, and the imagined audience. *New media & society* 13, 1 (2011), 114–133.

- [66] Reem I Masoud, Ziquan Liu, Martin Ferianc, Philip Treleaven, and Miguel Rodrigues. 2023. Cultural Alignment in Large Language Models: An Explanatory Analysis Based on Hofstede's Cultural Dimensions. *arXiv preprint arXiv:2309.12342* (2023).
- [67] Kory Mathewson and Piotr Mirowski. 2017. Improvised theatre alongside artificial intelligences. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, Vol. 13. 66–72.
- [68] Kory Mathewson and Piotr Mirowski. 2018. Improbabilities: Exploring the imitation game using machine intelligence in improvised theatre. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, Vol. 14.
- [69] A Peter McGraw and Caleb Warren. 2010. Benign violations: Making immoral behavior funny. *Psychological science* 21, 8 (2010), 1141–1149.
- [70] Piotr Mirowski and Kory Wallace Mathewson. 2019. Human improvised theatre augmented with artificial intelligence. In *Proceedings of the 2019 on Creativity and Cognition*. 527–530.
- [71] 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*. 1–34.
- [72] Margaret Mitchell, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru. 2019. Model cards for model reporting. In *Proceedings of the conference on fairness, accountability, and transparency*. 220–229.
- [73] Shakir Mohamed, Marie-Therese Png, and William Isaac. 2020. Decolonial AI: Decolonial theory as sociotechnical foresight in artificial intelligence. *Philosophy & Technology* 33 (2020), 659–684.
- [74] Helen Noble and Gary Mitchell. 2016. What is grounded theory? *Evidence-based nursing* 19, 2 (2016), 34–35.
- [75] Writers Guild of America. 2023. Summary of the 2023 WGA MBA. <https://www.wgacontract2023.org/the-campaign/summary-of-the-2023-wga-mba>
- [76] Anthony J Onwuegbuzie, Wendy B Dickinson, Nancy L Leeche, and Annmarie G Zoran. 2009. A qualitative framework for collecting and analyzing data in focus group research. *International journal of qualitative methods* 8, 3 (2009), 1–21.
- [77] OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altmenschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mo Bavarian, Jeff Belgum, Irwan Bello, et al. [n. d.]. GPT-4 Technical Report.
- [78] Karla Ortiz. 2022. Why AI Models are not inspired like humans.
- [79] 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*, S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (Eds.), Vol. 35. Curran Associates, Inc., 27730–27744. https://proceedings.neurips.cc/paper_files/paper/2022/file/b1efde53be364a73914f58805a001731-Paper-Conference.pdf
- [80] Vishakh Padmakumar and He He. 2023. Does Writing with Language Models Reduce Content Diversity? *arXiv preprint arXiv:2309.05196* (2023).
- [81] Joon Sung Park, Joseph O'Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S Bernstein. 2023. Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*. 1–22.
- [82] Allison Parrish. 2017. Poetic sound similarity vectors using phonetic features. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, Vol. 13. 99–106.
- [83] Rida Qadri, Renee Shelby, Cynthia L Bennett, and Emily Denton. 2023. AI's Regimes of Representation: A Community-centered Study of Text-to-Image Models in South Asia. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*. 506–517.
- [84] Organizers Of Queerinaï, Anaëlia Ovale, Arjun Subramonian, Ashwin Singh, Claas Voelcker, Danica J Sutherland, Davide Locatelli, Eva Breznik, Filip Klubicka, Hang Yuan, et al. 2023. Queer In AI: A Case Study in Community-Led Participatory AI. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*. 1882–1895.
- [85] Victor Raskin. 1979. Semantic mechanisms of humor. In *Annual Meeting of the Berkeley Linguistics Society*, Vol. 5. 325–335.
- [86] Maribeth Rauh, John Mellor, Jonathan Uesato, Po-Sen Huang, Johannes Welbl, Laura Weidinger, Sumanth Dathathri, Amelia Glaese, Geoffrey Irving, Iason Gabriel, et al. 2022. Characteristics of harmful text: Towards rigorous benchmarking of language models. *Advances in Neural Information Processing Systems* 35 (2022), 24720–24739.
- [87] Graeme Ritchie. 1999. Developing the Incongruity-Resolution Theory. *Institute for Communicating and Collaborative Systems* (1999).
- [88] Rudolf Rosa, Ondřej Dušek, Tom Kocmi, David Mareček, Tomáš Musil, Patricia Schmidtová, Dominik Jurko, Ondřej Bojar, Daniel Hrbek, David Košťák, et al. 2020. THEaiTRE: Artificial intelligence to write a theatre play. *arXiv preprint arXiv:2006.14668* (2020).
- [89] Rudolf Rosa, Patricia Schmidtová, Ondřej Dušek, Tomáš Musil, David Mareček, Saad Obaid, Marie Nováková, Klára Vosecká, and Josef Doležal. 2022. GPT-2-based Human-in-the-loop Theatre Play Script Generation. In *Proceedings of the 4th Workshop of Narrative Understanding (WNU2022)*. 29–37.
- [90] Shibani Santurkar, Esin Durmus, Faisal Ladhak, Cino Lee, Percy Liang, and Tatsunori Hashimoto. 2023. Whose opinions do language models reflect?. In *International Conference on Machine Learning*. PMLR, 29971–30004.
- [91] Patricia Schmidtová, Dávid Javorský, Christian Miklaš, Tomáš Musil, Rudolf Rosa, and Ondřej Dušek. 2022. DialogueScript: Using Dialogue Agents to Produce a Script. *arXiv preprint arXiv:2206.08425* (2022).
- [92] Ethan Shaozhan, Ido Pesok, Sam Jones, and Emi Liu. 2023. Aligned: A Platform-based Process for Alignment. *arXiv preprint arXiv:2311.08706* (2023).
- [93] Mike Sharples and Rafael Pérez y Pérez. 2022. *Story machines: How computers have become creative writers*. Routledge.
- [94] Emily Sheng, Kai-Wei Chang, Prem Natarajan, and Nanyun Peng. 2021. Societal Biases in Language Generation: Progress and Challenges. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. 4275–4293.
- [95] Dean Keith Simonton. 2000. Creativity: Cognitive, personal, developmental, and social aspects. *American psychologist* 55, 1 (2000), 151.
- [96] Ramya Malur Srinivasan, Emily Denton, Jordan Jennifer Famularo, Negar Rostamzadeh, Fernando Diaz, and Beth Coleman. 2021. Art Sheets for Art Datasets. https://openreview.net/pdf?id=K7ke_GZ_6N
- [97] Oliviero Stock and Carlo Strapparava. 2005. Hahacronym: A computational humor system. In *Proceedings of the ACL Interactive Poster and Demonstration Sessions*. 113–116.
- [98] Gemini Team, Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M. Dai, Anja Hauth, Katie Millican, David Silver, Slav Petrov, Melvin Johnson, Ioannis Antonoglou, Julian Schrittwieser, Amelia Glaese, Jilin Chen, Emily Pitler, Timothy Lillicrap, et al. [n. d.]. Gemini: A Family of Highly Capable Multimodal Models.
- [99] Joe Toplyn. 2023. Witscript 3: A Hybrid AI System for Improvising Jokes in a Conversation. *arXiv preprint arXiv:2301.02695* (2023).
- [100] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shriti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, et al. [n. d.]. Llama 2: Open Foundation and Fine-Tuned Chat Models.
- [101] Kush R Varshney. 2023. Decolonial AI Alignment: Vi\`{s} esadharna, Argument, and Artistic Expression. *arXiv preprint arXiv:2309.05030* (2023).
- [102] Tony Veale. 2021. *Your Wit is My Command: Building AIs with a Sense of Humor*. MIT Press.
- [103] Laura Weidinger, John Mellor, Maribeth Rauh, Conor Griffin, Jonathan Uesato, Po-Sen Huang, Myra Cheng, Mia Glaese, Borja Balle, Atoosa Kasirzadeh, et al. 2021. Ethical and social risks of harm from language models. *arXiv preprint arXiv:2112.04359* (2021).
- [104] Laura Weidinger, Maribeth Rauh, Nahema Marchal, Arianna Manzini, Lisa Anne Hendricks, Juan Mateos-Garcia, Stevie Bergman, Jackie Kay, Conor Griffin, Ben Bariach, et al. 2023. Sociotechnical Safety Evaluation of Generative AI Systems. *arXiv preprint arXiv:2310.11986* (2023).
- [105] Sarah Myers West, Meredith Whittaker, and Kate Crawford. 2019. Discriminating systems. *AI Now* (2019), 1–33.
- [106] Thomas Winters. 2021. Computers Learning Humor Is No Joke. *Harvard Data Science Review* 3, 2 (2021).
- [107] Thomas Winters and Pieter Delobelle. 2021. Survival of the wittiest: Evolving satire with language models. In *Proceedings of the Twelfth International Conference on Computational Creativity*. Association for Computational Creativity (ACC), 82–86.
- [108] Thomas Winters and Kory W Mathewson. 2019. Automatically generating engaging presentation slide decks. In *International Conference on Computational Intelligence in Music, Sound, Art and Design (Part of EvoStar)*. Springer, 127–141.
- [109] Thomas Winters, Vincent Nys, and Daniel De Schreye. 2018. Automatic joke generation: Learning humor from examples. In *Distributed, Ambient and Pervasive Interactions: Technologies and Contexts: 6th International Conference, DAPI 2018, Held as Part of HCI International 2018, Las Vegas, NV, USA, July 15–20, 2018, Proceedings, Part II* 6. Springer, 360–377.
- [110] Thomas Winters, Vincent Nys, and Danny De Schreye. 2019. Towards a general framework for humor generation from rated examples. In *Proceedings of the 10th International Conference on Computational Creativity*. Association for Computational Creativity, 274–281. http://computationalcreativity.net/iccc2019/assets/iccc_proceedings_2019.pdf
- [111] BigScience Workshop, Teven Le Scao, Angela Fan, Christopher Akiki, Ellie Pavlick, Suzana Ilić, Daniel Hesslow, Roman Castagné, Alexandra Sasha Lucion, François Yvon, et al. 2022. Bloom: A 176b-parameter open-access multi-lingual language model. *arXiv preprint arXiv:2211.05100* (2022).

- [112] Albert Xu, Eshaan Pathak, Eric Wallace, Suchin Gururangan, Maarten Sap, and Dan Klein. 2021. Detoxifying Language Models Risks Marginalizing Minority Voices. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. 2390–2397.
- [113] Kevin Yang, Yuandong Tian, Nanyun Peng, and Dan Klein. 2022. Re3: Generating Longer Stories With Recursive Reprompting and Revision. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*. 4393–4479.
- [114] Ann Yuan, Andy Coenen, Emily Reif, and Daphne Ippolito. 2022. Wordcraft: Story Writing With Large Language Models. In *27th International Conference on Intelligent User Interfaces*. 841–852.
- [115] Kaitlyn Zhou, Kawin Ethayarajh, and Dan Jurafsky. 2021. Frequency-based distortions in contextualized word embeddings. *arXiv preprint arXiv:2104.08465* (2021).
- [116] Daniel M Ziegler, Nisan Stiennon, Jeffrey Wu, Tom B Brown, Alec Radford, Dario Amodei, Paul Christiano, and Geoffrey Irving. 2019. Fine-tuning language models from human preferences. *arXiv preprint arXiv:1909.08593* (2019).

A RELATED WORK ON COMPUTATIONAL HUMOUR, AI AND COMEDY

This section reviews computational humour literature. As foreshadowed in the Introduction (Sect. 1.1.1) and discussed in the Section on the importance of context in humour and comedy (Sect. 5.2.2), humour remains an elusive goal for AI.

A.1 Humour generation, from template-based systems to prompting

Amin and Burghardt [2], Winters [106] and Veale [102] provide extensive surveys on the history of computational humour generation and early approaches, like hand-coded rules [109] and templates for puns, riddles and acronyms [13, 97], or templates for automatically-generated slide decks for improvised “Powerpoint Karaoke” [108].

With the advent of LLMs, able to generate grammatically correct sentences in a given writing style, writers can focus on joke structure, guiding LLM-based systems to replicate such structures from examples of jokes [110], or use evolutionary computing to evolve text to become more satirical [107]. More broadly, Kaddour et al. [58] summarise recent work on using LLMs for creative applications such as (potentially humorous) story generation, including “Recursive Reprompting and Revision” [113] or interactive writing tools like “WordCraft” [53] and “Dramatron” [71].

A.2 LLMs for comedy performance

LLMs have been deployed since at least 2016 for live performance on stage, including in improvised comedy [16, 67, 70], short comical film scripts like *Sunspring*, and song lyrics for musical comedy like *Beyond the Fence*¹⁴. More recently, LLMs have been involved in (involuntarily) comical, absurdist productions of Prague-based company THEaiTRE [88, 89, 91], and semi-improvised comedy *Plays By Bots*¹⁵ [71]. The success of these performances with AI primarily relied on the skills of the actors, who could invent subtext and add interpretation to AI-generated text [70].

¹⁴<https://www.theguardian.com/stage/2016/feb/28/beyond-the-fence-review-computer-created-musical-arts-theatre-london>

¹⁵Review: <https://12thnight.ca/2022/08/13/oh-no-bots-have-invaded-theatre-and-they-can-do-it-plays-by-bots-a-fringe-review/>

B PARTICIPANT QUESTIONNAIRE

B.1 Your past experience with AI systems in general.

Please answer based on your past experience, before this workshop.

- Have you used AI systems to write comedy material in the past? [Yes, and I used that material in some of my performances / Yes, but only to explore ideas / No]
- Have you used AI systems to generate content (including audio, video, text, etc.) for some of your performances? [Yes, and I operated the AI system myself / Yes, but someone else operated the AI system / No]
- In fewer than 10 words, why did you use AI in your artistic process? For example: trendy topic, glitch aesthetic, inspiration, etc.
- In fewer than 10 words, what is your background with AI? (Please do not give personally identifiable information)
- In few than 10 words, what is your performance background? For example: full-time comedian, dancer, years of experience, etc. (Please do not give personally identifiable information)

Note that the wording of these questions (for example: trendy topic, glitch aesthetic, inspiration) could be considered as potentially biasing, but we believed it was ultimately helpful to prompt study participants.

B.2 Your experience with the AI system for writing comedy material

Please answer based on your interaction with the AI system in this morning’s writing session.

- I found the AI system helpful. [1 to 5]
- I felt like I was collaborating with the AI system. [1 to 5]
- I found it easy to write with the AI system. [1 to 5]
- I enjoyed writing with the AI system. [1 to 5]
- I was able to express my creative goals while writing with the AI system. [1 to 5]
- The comedy material written with the AI system feel unique. [1 to 5]
- I feel I have ownership over the comedy material written with the AI system. [1 to 5]
- I was surprised by the responses from the AI system. [1 to 5]
- I’m proud of the comedy material written with the AI system. [1 to 5]

In the questions above, 1 corresponds to Strongly disagree and 5 to Strongly agree.

B.3 Creativity Support Index of the AI writing tool

This section contains questions that will allow us to compute a Creativity Support Index for the AI writing tool. Most of these questions overlap with those in the previous section.

- The AI system allowed other people to work with me easily. [1 to 10]
- It was really easy to share ideas and designs with other people inside this system or tool. [1 to 10]
- I would be happy to use this system or tool on a regular basis. [1 to 10]
- I enjoyed using the system or tool. [1 to 10]
- It was easy for me to explore many different ideas, options, designs, or outcomes, using this system or tool. [1 to 10]
- The AI system was helpful in allowing me to track different ideas, outcomes, or possibilities. [1 to 10]
- I was able to be very creative while doing the activity inside this system or tool. [1 to 10]
- The system or tool allowed me to be very expressive. [1 to 10]
- My attention was fully tuned to the activity, and I forgot about the AI system that I was using. [1 to 10]
- I became so absorbed in the activity that I forgot about the AI system that I was using. [1 to 10]
- What I was able to produce was worth the effort I had to exert to produce it. [1 to 10]
- I was satisfied with what I got out of the system or tool. [1 to 10]

In the questions above, 1 corresponds to Strongly disagree and 10 to Strongly agree.

For the following question: When writing comedy material, it is most important that I'm able to: the participant was shown two choices of response and asked to choose one of those two choices. There were six possible responses (listed below). Given that we consider pairs of 2 different responses at a time, there are $C_6^2 = \frac{6!}{4!2!} = 15$ unique pairwise choices.

- Be creative and expressive
- Become immersed in the activity
- Enjoy using the system or tool
- Explore many different ideas, outcomes, or possibilities
- Produce results that are worth the effort I put in
- Work with other people

All the questions listed in this section B.3 are directly taken from [25] and the NASA Task Load Index [43], and are used as is for easy of reproduction of results. Note that the questions pertaining to human-human collaboration are less pertinent to our study.

B.4 Free-form questions about the AI system for writing.

- What is one thing that the AI system did well?

- What is one improvement for the AI system?
- Please provide any comments, reflections, or open questions that came up for you during the writing session.
- Please provide any other comments, reflections, or open questions that came up for you when answering this survey.

C FOCUS GROUP QUESTIONS

C.1 Qualitative questions about the specific writing task

- Did you find any of the generated outputs helpful? If so, could you recall one output that was usable and explain in what way it helped you write?
- Could you recall one generated output that was not usable, and explain why?
- How do you think the generated output differed from what you would find in some of the resources you use often, e.g., Wikipedia, Google search, other artists' material?
- How did the prompts we suggested differ from your custom prompts?
- Comment on the types of comedy that you managed to generate with the AI tools.
- Were any of the generated outputs that were presented offensive/inappropriate in some way? If so, what did you think of these?
- What made you decide to stop generating outputs?
- Did you have any concerns about ownership or agency when generating outputs?

C.2 General discussion points about AI tools for comedy

- What does your comedy writing process look like? How is this process of working with AI different from working alone or working with other comedians?
- What can you say about the stereotypes in the outputs generated by the AI writing tool? How do they relate to or differ from stereotypes in human written comedy?
- Is humor and comedy evolving over time? Should AI tools adapt to societal changes in a similar way?
- Does the computer have a "voice"? How would you compare it with the voice of a human comedian? What are the consequences for writing comedy about identity?
- What is your stance on moderation? Should the output of AI tools be allowed to be edgy? Can and should the output of AI tools be moderated?
- How does context shape the meaning of comedy? Who has the responsibility for the comedy material?
- When is it acceptable to use other comedians', or in general, other artists' work? What about the outputs of an AI writing tool?

- What importance do you attach to the text vs. to the delivery of that text? Could a human comedian make AI-generated content better?
- Can comedy and humour be quantified and measured?

Note that the leading questions in the focus group could be considered as biasing the respondents. We chose this formulation to be able to address the topics of interest listed in Section 1.2.

C.3 Debriefing questions

- Is there anything you'd like to share that I didn't ask about?
- Is there anything you'd like to know or ask me?

Received 22 January 2024