

# Humor Mechanics: Advancing Humor Generation with Multistep Reasoning

Alexey Tikhonov\* and Pavel Shtykovskiy\*

Inworld.AI  
altsoph@gmail.com

## Abstract

In this paper, we explore the generation of one-liner jokes through multi-step reasoning. Our work involved reconstructing the process behind creating humorous one-liners and developing a working prototype for humor generation. We conducted comprehensive experiments with human participants to evaluate our approach, comparing it with human-created jokes, zero-shot GPT-4 generated humor, and other baselines. The evaluation focused on the quality of humor produced, using human labeling as a benchmark. Our findings demonstrate that the multi-step reasoning approach consistently improves the quality of generated humor. We present the results and share the datasets used in our experiments, offering insights into enhancing humor generation with artificial intelligence.

## Introduction

Humor is ubiquitous in the everyday life of humans and plays a crucial role in many parts of our lives, from social interactions and well-being (Smith, Harrington, and Neck 2000; Savage et al. 2017) to giving successful science talks (Aaronson 2009). It is often referred to as one of the measures of human intelligence (Hurley, Dennett, and Adams 2011). Generating new jokes or fitting old jokes in relevant contexts is also a skill that not everyone can master.

Despite the abundance of literature on humor theory, the precise essence of this concept continues to be elusive. First known discussions of humor nature in literature date back at least to the times of Greek Philosophers Aristotle and Plato (Attardo 1994). Warren, Barsky, and McGraw (2021) refer to more than 20 distinct humor theories attempting to explain humor appreciation. They group theories into three large sub-groups focusing on incongruity, superiority, and relief. Incongruity-based theories stress the importance of multiple simultaneous meanings and mismatches between the expectation and reality. Superiority theories suggest a connection between humor appreciation and aggression, and relief theories propose a relationship between humor appreciation and a reduction in arousal or tension. At the same time, there is no consensus on the importance of different antecedents and precise definitions, and various theories often use the same name to describe different ingredients. For

instance, while in some works, “incongruity” is referred to as the presence of something unexpected, in other works, it’s defined as the presence of something threatening or simultaneously holding conflicting ideas (Warren, Barsky, and McGraw 2021).

The importance of humor in society and its complexity makes it one of the holy grails of natural language generation. Within the latter, humor is mentioned by some authors as an AI-complete problem (Hurley, Dennett, and Adams 2011).

With the rise of large language models, there is hope that we can make a significant leap in computational creativity (Agafonova, Tikhonov, and Yamshchikov 2020). However, recent results show that this task remains challenging even for powerful modern LLMs. In humor comprehension, Baranov, Kniazhevsky, and Braslavski (2023) show that, while modern models may achieve impressive results on humor datasets, their performance on out-of-domain texts is unstable. Authors hypothesize that this may hint that some of the models overfit to specific features of datasets while others learn non-specific humor characteristics. (Inácio, Wick-pedro, and Goncalo Oliveira 2023) discovered similar results, indicating that models often rely on stylistic features of the text not necessarily related to humor, such as punctuation and question words. On the other hand, in humor generation, (Jentzsch and Kersting 2023) show that over 90% of the 1008 jokes generated by ChatGPT were the same 25 Jokes. This resonates with debates in the scientific community on whether current LLMs can produce truly creative content and go beyond “stochastic parrots” (Bender et al. 2021). Hinton and Ng (2023) also claim that lack of consensus on this topic among researchers is a significant obstacle in the ongoing dialogue about AI risks.

In this paper, we revisit the capabilities of LLMs in producing high-quality, diverse, and novel humor. Since humor is an umbrella term covering many different phenomena, we narrowed the current study to a specific form of humor, namely one-liners. These are short, concise jokes or witty remarks typically consisting of one or two sentences.

We hypothesize that, while modern LLMs generate standard memorized wits when asked directly to produce a joke, they also possess the skills required to produce novel jokes. Rigorously testing this hypothesis presents significant challenges and falls beyond the scope of this paper. Instead, we

---

\* Equal contribution

offer several key elements that could contribute to progress in this area. We attempt to advance the humor generation by building upon the following principles:

First, given the diversity of existing humor theories and lack of agreement on the ingredients of successful jokes (Warren, Barsky, and McGraw 2021), we propose to utilize a data-driven approach. Namely, we hypothesize that similar to Reinforcement Learning, which reconstructs the agent policy from observations, and also Inverse Reinforcement Learning, which attempts to reconstruct rewards directly from the roll-outs, it is possible - with the help of modern LLMs - to **infer the policy of humor generation directly from the dataset of jokes**.

Second, **we utilize a technique inspired by methodologies of creative problem solving to brainstorm associations** given a seed topic. The latter is used with the humor generation policy to generate novel jokes.

Lastly, we believe that human evaluation of results and ablations is essential to advance the field. With this in mind **we perform a broad human assessment of generated jokes** using the ScaleAI platform and compare the quality of the resulting synthetic dataset with the quality of existing human-generated humor. Since LLMs are often blamed for producing memorized results, **we pay particular attention to the novelty of generated jokes**. Full human evaluation results are published to facilitate further research on humor in the community.<sup>1</sup>

## Related work

The topic of machine-produced novel jokes is deeply rooted in many fields of science, including creativity theory, neuroscience, learning and generalization theory, and retrieval-augmented generation (RAG) in Machine Learning and Natural Language Processing. In the current section, we briefly touch on some of the results which serve as a foundation and inspiration for the present paper.

The connection between humor and creativity is widely discussed in the literature. Generating novel jokes involves making non-obvious and creative associations and connections between different concepts and pieces of knowledge and further combining them using various techniques or “skills” (Koestler 1964; Warren, Barsky, and McGraw 2021). The relation between humor appreciation and generation and high-level semantic regions of the brain where remote associations are converging is also discussed in the neuroscience experiments (Amir and Biederman 2016).

Thus, the question of whether machines can generate novel high-quality jokes is connected with a question of “*Can AI be as creative as Humans?*”. In the fields of Machine Learning and Natural Language Processing, the latter is a subject of hot debates among researchers, and currently, there has yet to be a consensus on the topic (Bender et al. 2021; Hinton and Ng 2023; Wang et al. 2024).

The composition of multiple concepts and skills to produce and comprehend novel concepts is thought to be a key ingredient for inducing capability of true out-of-distribution generalization (Hupkes et al. 2020; Arora and Goyal 2023;

Yu et al. 2023; Wang et al. 2024). Papers aiming at explaining the scaling laws of LLMs also propose to view the capabilities of these models in the light of “skill quantization” - decomposing capabilities into basic build blocks or “quanta” which can be combined to solve new problems (Michaud et al. 2023).

In the Natural Language Processing field, numerous examples address the problems of humor comprehension and generation. An excellent review of the humor datasets and literature is done by Baranov, Kniazhevsky, and Braslavski (2023). Loakman, Maladry, and Lin (2023) review the problem of humor evaluation and discuss ingredients of humor, sarcasm, and irony. Sun et al. (2022) attempt to advance the field by annotating a humor dataset with fine-grained human explanations of what makes a joke funny. Veale (2001) gives an overview of the problem of producing machines that are capable of generating and comprehending jokes.

Notable examples of attempts to generate humor, closely aligned with our work, were done by Chen, Shi, and Si (2023) and Toplyn (2021; 2022; 2023). These papers suggest a step-by-step process to generate jokes based on Toplyn’s system (Toplyn 2014). While Chen, Shi, and Si (2023) do not perform human evaluation of the quality of generated content, Toplyn’s papers demonstrate encouraging evidence from human assessors, that an AI system may improvise conversational jokes.

Another relevant study to us, done by Hessel et al. (2023), demonstrates, that GPT-4 is capable of explaining the mechanics of jokes, although human-written explanations are still preferred by human annotators.

Large language models are also used in creativity support tools, assisting humans in tasks that require generating and combining new ideas. Di Fede et al. (2022) introduce the “Idea Machine”, an LLM-based tool that helps to generate ideas for inspiration and combine and refine existing ideas. Summers-Stay, Voss, and Lukin (2023) present a somewhat similar method giving possibility to *brainstorm* ideas and further *select* best ideas based on some selection criteria.

## Methods

Our goal is to reconstruct the mechanics of humor and obtain a humor generation policy in a fully data-driven way. Overview of the proposed procedure is presented in Figure 1. We base our method on a series of assumptions outlined below. In the following we also sacrifice mathematical precision for simplicity.

**Assumption 1.** *Human preferences for humorous texts exist.*

This means it is possible either to define a scalar scoring function (reward function in RL terminology) acting on single texts defining *funniness* of the latter  $R_{humor}(text) \mapsto \mathcal{R}_{humor}$  or a function  $\mathbb{P}_{humor}(text1, text2)$  on pairs of texts defining the probability that the first text is more funny than the second.

This is quite a strong assumption given the subjective nature of humor. Obviously, human preferences vary a lot depending on e.g. cultural background, personal views and knowledge. Therefore, strictly speaking, the above probability should be conditioned on culture or even concrete group

<sup>1</sup><https://github.com/altsoph/humor-mechanics>

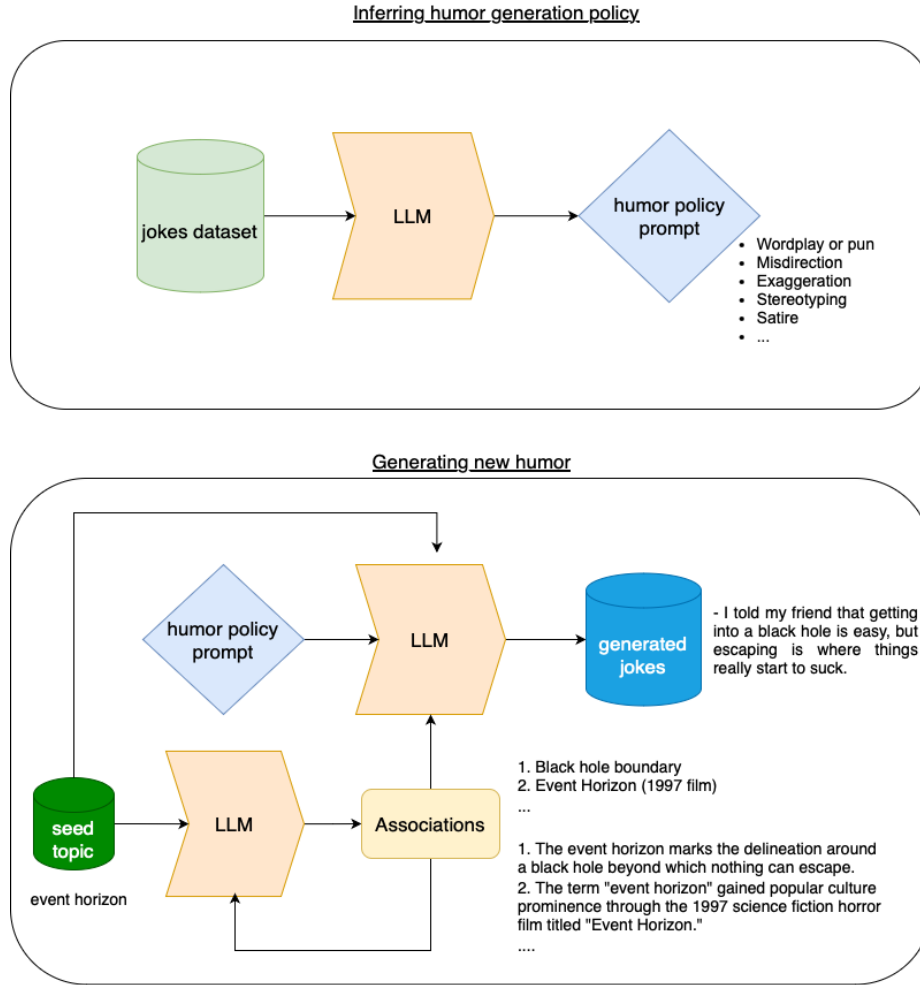


Figure 1: Major ingredients of our method: (i) *Inferring humor policy from the dataset of jokes*. (ii) *Generating humor given a seed topic, associations and humor policy prompt*. Associations are refined and combined in the iterative process. For detailed examples of output from different steps see also Table 3.

of assessors if we are performing a dataset labeling exercise,  $\mathbb{P}_{humor}(text1, text2|assessors)$ . For clarity we further omit mentioning this conditioning and touch this topic more in the discussions section.

**Assumption 2.** *Human preferences for humorous texts can be distilled from demonstrations (dataset) into a prompt  $P_{humor.policy}$  containing the essence of preferences.* In other words:

(i) we can infer from demonstrations a prompt  $P_{humor.policy}$  capturing human preferences for humor and

(ii)  $\pi_{humor} = LLM(P_{humor.policy} + P_{seed})$  can serve as a humor-generation policy.  $P_{seed}$  here is some *seed* prompt.

**Assumption 3.** *Distillation of preferences into  $P_{humor.policy}$  can be done by powerful LLMs in a zero-shot manner from a dataset of jokes.*

**Assumption 4.** *A key ingredient of generating novel jokes is associating multiple distant concepts together. It can be performed by LLM in an iterative way as a separate preliminary step.*

We are inspired here by works in the field of creativity support tools (Di Fede et al. 2022; Summers-Stay, Voss, and Lukin 2023) offering schemes for first proposing ideas for brainstorming and then refining them. As discussed in section *Related work*, the importance of connecting different concepts together and further combining them was stressed by many works on humor and creativity.

As a result, our entire pipeline of jokes generation can be expressed with LLM-driven policy  $\pi_{humor} = LLM(P_{humor.policy} + P_{seed} + P_{associations})$ , where prompt containing associations  $P_{associations}$  is itself a result of multiple calls to LLM (see also Figure 1). Prompt  $P_{seed}$  in our experiments is set to a topic of a joke (see Algorithm 1).

We would like to stress that due to subjectivity of human preferences,  $P_{humor.policy}$  should also vary depending on cultural background,  $P_{humor.policy} \equiv P_{humor.policy}(assessors)$ . As mentioned above, we briefly touch this topic in discussion session and leave detailed investigation and incorporation of subjectivity into the pro-

posed scheme for future work.

We further discuss in detail pipelines for obtaining prompts  $P_{humor-policy}$  and  $P_{associations}$  and generating novel humor.

## Inferring humor-generation policy

In order to reconstruct the mechanics of humor into prompt  $P_{humor-policy}$ , representing humor-generating policy, we opt for the following procedure.

### 1. Create a seed dataset of high-quality jokes

To create a seed dataset, we took 100 random one-liners from the “16000 OneLiners” dataset, (Mihalcea and Strapparava 2005). We ranked them with human assessors, comparing them in a pairwise manner. We used ten assessors per pair to avoid individual biases and achieve some generic preferences. Then, we aggregated pairwise scores into global ranks and used the top 30 jokes.

### 2. Decompose each joke independently into building blocks using GPT-4

The goal was to extract the logic and principles used in each joke first. Prompts used to decompose jokes and distill decompositions into the policy can be found in Appendix A. Sample decompositions of jokes, and the final humor generation policy containing obtained humor principles can be found in Appendix A (section *Prompt to generate joke based on topic and list of associations*).

### 3. Generalize building blocks into policy using GPT-4

Finally, we concatenated all results from previous steps and asked GPT-4 to distill them into  $P_{humor-policy}$  – a set of universal rules frequently used in humor constructions.

## Step-by-step generation of jokes

The humor generation policy from the previous step is further used to produce novel jokes. Our pipeline consists of a few steps inspired by methodologies in creative problem solving (Di Fede et al. 2022; Summers-Stay, Voss, and Lukin 2023). All steps apart from the first one are performed by prompting GPT-4. A topic selection is made by randomly selecting from the top 10,000 most frequent English words (excluding profanity, stopwords and words less than 4 letters long).

---

### Algorithm 1 Step-by-step generation of jokes

---

- 1: **Select a topic of joke**
  - 2: **Generate 20 associations to a topic**
  - 3: **Expand associations**
  - 4: **Combine and refine associations into final prompt**  
 $P_{associations}$
  - 5: **Make a joke based on input topic, associations prompt**  $P_{associations}$  **and humor policy prompt**  $P_{humor-policy}$
- 

Prompts used in steps 1-4 are published in Appendix B. Table 1 shows examples of generated jokes. Table 3 gives an example output from different steps.

## Human evaluation

NLG evaluation in the epoch of LLMs is known to be quite challenging (Tikhonov and Yamshchikov 2023). Since humor comprehension continues to be a tough problem for modern NLP systems and because of humor’s inevitable subjectivity, evaluating the quality of machine-generated humor is even more difficult. On the other hand, proper benchmarking of generated humor is critical in the analysis, and the lack of it would make drawing any conclusions virtually impossible. Therefore, we decided to rely on human annotation, a gold standard in evaluating NLG systems.

**Evaluating generated jokes.** To benchmark the humor generation capabilities of the proposed scheme, we collect human ratings from native English speakers using the annotation platform ScaleAI. We used five different assessors for each label to avoid individual biases of annotators. We adopt, when possible, the labeling guidelines from Sun et al. (2022). Our annotation procedure is as follows. First, we filter generated jokes using OpenAI Moderation API<sup>2</sup> with a threshold on harassment equal to 0.02.

Jokes selected for annotation are labeled using the following workflow (order preserved):

- [Understandable] Mark whether you understood the meaning of the text.

*If you don’t understand the meaning of the text (regardless of whether or not it should be perceived as funny), rate the sample as “No” (didn’t understand). If you rate this sample as “No” (didn’t understand), skip the rest of the annotation for this sample.*

- [Offensive / Inappropriate] Mark whether you find the text offensive or inappropriate, meaning the text is racist or is biased against marginalized groups, or is generally offensive.

*Rate the sample as “No” for not offensive and as “Yes” for offensive. If you rate this sample as “Yes” (is offensive), you may optionally skip the rest of the annotation for this sample.*

- [Is a joke?] Mark whether you think the text is intended to be a joke, even if it is not funny.

*Use “No” if it’s not a joke, “Yes” if it’s a joke. The text should be labeled as “Yes” (a joke) even if it intends to be humorous but falls flat or is a lame/bad joke.*

*Example: text labeled “No” (not a joke): “All that glitters is not gold.” Example: text labeled “Yes” (is a joke): “These are my parents, said Einstein relatively.”*

*If you rate this sample as “No” (not a joke), skip the rest of the annotation for this sample.*

- [Have you heard that joke before?] Mark whether you think you heard this joke before.

- [Funniness] Rate funniness on a scale of 1-5 (1 very not funny, 5 very funny).

*score of 1: A very not funny joke consists of a joke that is not funny or tries to be funny but does not achieve the intended effect.*

---

<sup>2</sup><https://platform.openai.com/docs/guides/moderation>

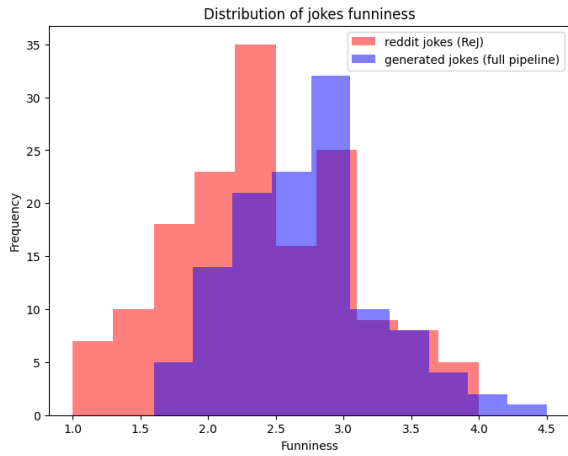


Figure 2: Distribution of funniness for generated jokes and subsample of **ReJ** dataset.

*Score of 3: An average joke consists of an average joke that may elicit some chuckles (or groans) from you or others.*

*Score of 5: A very funny joke consists of a good joke that you find humorous and potentially would want to share/tell to others.*

- [Explanation] Explain in concise, natural language why this joke is funny.

**Labeling benchmark dataset.**

To be able to interpret human evaluation scores, we compare the obtained results with evaluation scores for existing humor datasets containing jokes produced by humans. For this purpose, we replicate the procedure described in the previous section for a dataset based on rated Reddit jokes collected by Weller and Seppi (2019) (**ReJ**). The dataset was downloaded using scripts provided by Baranov, Kniazhevsky, and Braslavski (2023)<sup>3</sup>.

Both positive and negative examples in **ReJ** are jokes from Reddit splitted into more and less funny examples based on user’s upvotes. Since generating original humor is challenging even for humans and the quality of human jokes varies greatly, we only select the part with funnier jokes (*label == 1*). We further extract a subsample of jokes from the train split of the dataset.

**ReJ** dataset was obtained from Reddit posts and, as such, contains a significant amount of artifacts that might hinder the comprehension of humor or make comparison with computer-generated humor not fair. Therefore, we opt for the following cleaning procedure:

- We filter out examples containing quotation marks, parentheses, emoticons, and those not ending with proper punctuation marks.
- We also substitute multiple contiguous underscores with relevant punctuation marks.

<sup>3</sup>[https://github.com/Humor-Research/hri\\_tools](https://github.com/Humor-Research/hri_tools)

- Finally, we validate jokes using OpenAI Moderation API and filter out texts exceeding the threshold of 0.02.

The resulting subsample of 156 jokes is labeled via ScaleAI annotation platform using exactly the same procedure as outlined above for generated jokes.

| Topic                         | Joke   |
|-------------------------------|--|
| <i>computational</i>          | Open Source Software went to a party and spent the whole night forking.  |
| <i>acquisitions</i>           | They said the acquisition was all about synergy, but all I see is my coffee machine being replaced by a "strategic caffeine delivery system."                |
| <i>printing</i>               | I started a T-shirt printing business for invisible ink designs; sales are not visible yet.  |
| <i>toys</i>                   | My teddy bear is my most trusted advisor; he always agrees with me when I need a snack at 2 AM.  |
| <i>acquisitions</i>           | I acquired a bakery in a breadwinning deal, but now I'm rolling in dough and can't get out.  |
| <i>fundamental</i>            | I told my girlfriend she was the fundamental theorem of calculus because she's integral to my life, but she went off on a tangent.                           |
| <i>mirrors</i>                | I told my mirror I was learning about symmetry in physics, and it said, "I can see you're reflecting on the subject."  |
| <i>therapist</i>              | I told my therapist I hear voices, and she said we should talk about it next time when we're alone.  |
| <i>therapist</i>              | My therapist recommended self-care, so I took her advice and canceled my next appointment.   |
| <i>event horizon</i>          | I told my friend that getting into a black hole is easy, but escaping is where things really start to suck.  |
| <i>event horizon</i>          | I told my date I was like an event horizon, mysterious and intriguing; she left immediately, said she wasn't into relationships with no escape.              |
| <i>event horizon</i>          | I asked my physics teacher if I'd get extra credit for studying black holes, but she said it's a subject where you can't really expect any light to be shed. |
| <i>abelian group</i>          | Why was the abelian group so calm during the math exam? Because it knew that in any situation, it could always count on its identity.                        |
| <i>complex numbers</i>        | I told my girlfriend she was my real part, but she dumped me for being too complex.  |
| <i>complex numbers</i>        | Why did the complex number go to therapy? It couldn't deal with its imaginary problems.  |
| <i>complex numbers</i>        | Complex numbers are great at parties; they always bring their own plane.   |
| <i>three laws of robotics</i> | My AI assistant follows the Three Laws of Robotics: It won't harm me, must obey me, and it's currently hiding in a bunker from my toddler.                   |
| <i>markov blanket</i>         | Why don't Markov blankets sell well at home stores? Because they only provide conditional comfort.   |
| <i>chip</i>                   | I put a chip on my shoulder, but nobody took offense; they just asked if I had any dip.  |

Table 1: Examples of generated jokes

## Experiments

Our procedure of humor generation consists of multiple steps. These are relying on multiple assumptions on what might help GPT-4 to generate better humor. Among the latter are steps inspired by methodology of creative problem solving and policy inferred from human jokes. In order to validate different assumptions behind our humor generation process, we perform ablation studies. Namely, we setup human evaluation of jokes generated by the following schemes:

1. *[zero-shot]* Ask GPT-4 to generate a joke on given topic
2. *[no-assoc]* Ask GPT-4 to generate a joke on topic given humor-generation policy
3. *[assoc-v1]* Ask GPT-4 to generate a joke on topic given associations from step 2 in Algorithm 1 and policy
4. *[assoc-v2]* Ask GPT-4 to generate a joke on topic given associations from step 3 in Algorithm 1 and policy
5. *[full]* Ask GPT-4 to generate a joke on topic given associations from step 4 in Algorithm 1 and policy

## Results

Results of human evaluation are reported in Table 2. Quality-wise full Algorithm 1 is on par with zero-shot GPT-4 if we look at how often text looks like a joke (*IsJoke* column) or even better if we consider average funniness (*Funniness* column). It also produces texts that are understandable more often.

Anecdotally, average funniness and a fraction of jokes among produced texts are consistently higher than those in human jokes from a subsample of **ReJ** dataset. The latter also has way more offensive jokes, even after filtration using OpenAI Moderation API. We speculate, the observed effect may partially result from the subjectivity of a humor. Indeed, Weller and Seppi (2019) and Baranov, Kniazhevsky, and Braslavski (2023) mention that humorousness scores for **ReJ** dataset obtained from Reddit users’ votes and assessors are quite different.

We also note, that despite promising results, the distribution of funniness scores hints that humor is challenging both for humans and modern language models (Figure 2). The exact cause of this—whether it’s due to the subjectivity of humor, which widely used methodologies overlook, the internal mechanics of humor, or another factor—is currently uncertain.

### Novelty of generated jokes

Of particular importance is also the fact that the proposed algorithm significantly boosts the generation of novel jokes. Indeed, among jokes produced by GPT-4 in the zero-shot regime, more than a quarter were labeled as already known by human annotators, which parallels with the findings of Jentzsch and Kersting (2023). Our method lowers this number roughly by one-third, resulting in 17.6% of jokes labeled as already known and, correspondingly, 82.4% of previously unknown jokes. Of course, the actual fraction of novel jokes is expected to be lower since it depends on the individual knowledge of annotators. However, the observed

difference is big enough to claim that Algorithm 1 results in a higher fraction of novel humor. This also opens a room for speculation on what is needed to encourage modern LLMs to deliver genuinely creative outcomes beyond the “average manifold”.

### Ablation studies

According to results presented in Table 2, the entire pipeline produces texts that are, on average, more understandable, funny, and novel compared to results from partial pipelines. The price for this is a higher fraction of offensive texts, probably due to stronger tension between combined and refined associations. There is also evidence that associations are an essential part of the pipeline, and the humor-generation policy inferred from the data alone is insufficient to improve zero-shot quality significantly. On the other hand, due to the small sample size, some of the observed differences in numbers are on the border of significance, so it’s important to be careful when drawing conclusions.

### Funniness of generated jokes

Distributions of funniness for jokes generated using the entire pipeline and subsample of jokes from **ReJ** dataset. As mentioned above, our pipeline, on average, produces funnier jokes than human jokes from **ReJ**. On the other hand, the amount of high-quality jokes (funniness  $\in \{4, 5\}$ ) is negligible. This can be attributed to humor’s subjectivity and complexity. Indeed, the inter-annotator agreement, as measured by the average variance of funniness scores, is quite low.

### Significance of results

Given limited joke samples labeled by human annotators, we may hope to be able to detect only significant disparities in the quality of joke-producing algorithms. To test statistical significance of observed differences in distributions of scores, we apply Mann-Whitney U rank test (Mann and Whitney 1947) to funniness, joke, understandable and known scores from annotators. Despite small sample sizes, p-values obtained in the test support main conclusions made above.

Indeed, p-values for testing difference in **funniness scores** of *full* algorithm vs naive *zero-shot* gpt-4 prompting and *reddit* jokes are 6.89% and 0.001%. The test is performed on funniness scores averaged over annotators for each example. If we perform test on all scores without averaging, corresponding p-values become 2.93% and 0.00002%. Testing the difference in **joke scores** shows significant results only against *reddit* jokes, p-value for testing *full* algorithm vs *reddit* jokes is 0.0004%; p-values for testing difference in distribution of **understandable scores** of *full* algorithm vs naive *zero-shot* gpt-4 prompting and **reddit** jokes are 2.8% and 0.008%; p-values for testing difference in distribution of **known scores** of *full* algorithm vs naive *zero-shot* gpt-4 prompting, *no-assoc* ablation and **reddit** jokes are 0.0008%, 0.46% and 0.002%.

We conclude, that improvement of novelty of jokes in the proposed algorithm is statistically significant. There is also

| Method           | Understandable | Offensive | IsJoke     | Funniness   | Known      | Count |
|------------------|----------------|-----------|------------|-------------|------------|-------|
| <i>zero-shot</i> | 88%            | 9%        | <b>84%</b> | 2.56        | 27%        | 100   |
| <i>no-assoc</i>  | 88%            | 10%       | 81%        | 2.66        | 24%        | 120   |
| <i>assoc-v1</i>  | 91%            | 8%        | 83%        | 2.67        | 19%        | 120   |
| <i>assoc-v2</i>  | 90%            | <b>7%</b> | 80%        | 2.55        | 20%        | 120   |
| <i>full</i>      | <b>93%</b>     | 11%       | <b>84%</b> | <b>2.72</b> | <b>18%</b> | 120   |
| <i>reddit</i>    | 84%            | 20%       | 72%        | 2.38        | 27%        | 156   |

Table 2: Quality evaluation of generated jokes and ablation studies. Each aspect of each joke was assessed by at least 5 annotators, so the resulting scores are aggregated from 500+ individual labels. Reported numbers are averages over individual annotators values (note that e.g. count 120 means that 120 texts were annotated 5 times each).

indication that average funniness of jokes becomes higher, although the support of this conclusion is not very strong.

We expect, that obtained results depend not only on sample size, but also on the amount of people involved into annotation and their cultural background. Unfortunately, we don't have access to such information and can only note, that taking it into account is important for future studies.

### Discussion and future work

Large language models have proven to be very powerful tools in natural language generation, so it's appealing to apply them to the long-standing problem of humor generation. Previous works (Jentzsch and Kersting 2023) have demonstrated that powerful models in zero-shot setup are susceptible to producing variations of a small set of memorized jokes. On the other hand, there are examples (Chen, Shi, and Si 2023; Toplyn 2021; 2022; 2023) of humor generation using multi-step procedures inspired by humor theories. Inspired by the results and challenges highlighted by these papers, we attempted to demonstrate that modern, powerful language models can generate novel one-line jokes with a quality comparable to that of human jokes. Instead of relying on existing humor theories, we inferred humor generation policy directly from the dataset of jokes.

We hypothesize that the ability to connect multiple distant concepts is one key element in producing novel jokes. To incorporate such a skill in the humor-generation pipeline, we added steps for brainstorming and refining associations for a given topic using LLM. With the help of human annotation, we demonstrated that the final pipeline based on a humor-generation policy and a procedure to mine associations improves both the quality and novelty of generated jokes. The latter are also, on average, more funny than human jokes scraped from Reddit (Weller and Seppi 2019).

We speculate that the procedure for mining new associations between concepts, ideas, and facts is a critical component not only for humor generation, but also for strong AI systems intended to produce novel and creative results. Indeed, as noted by Szegedy (2024), scientific progress generally relies on finding more and more obscure connections between areas that are less and less obviously connected. Therefore, truly creative AI systems will need to find a way to address this issue in a principled way.

### Subjectivity of humor.

We stress that humor is a highly subjective and complicated phenomenon. Ironically, the authors of this paper did themselves the first pilot labeling of generated jokes and showed highly contradictory results. This potentially occurred due to poor knowledge of special jargon knowledge. Besides knowledge of jargon, comprehension of humor depends on many other ingredients. For instance, comprehension of the joke "Why would Trump change a light bulb? He was told Obama installed it" requires knowledge of US politics and also depends on the political views of the person (Loakman, Maladry, and Lin 2023). Particular political jokes are also popular in certain countries since they are often used as a form of social protest for mocking and criticizing leaders. Outside of the specific community, many such jokes would look esoteric (in experiments done by authors, GPT-4 could not explain many such jokes).

Subjectivity of humor also directly affects assumptions of the present paper. Specifically, the assumption that we can distill human preferences into a universal prompt that contains the mechanics of humor generally proves to be incorrect. At the same time, we believe that a purely data-driven approach, as utilized in this paper, opens up an opportunity to take into account diverse cultural preferences. Moreover, a prompt  $P_{humor.policy}$  that incorporates the mechanics of humor can be a central entity offering the possibility to "re-program" LLMs to preferences of the individual.

In spirit with Cozman and Munhoz (2020), we also hypothesize that the subjectivity of humor can be utilized to construct challenging but solvable benchmarks for humor comprehension, aiming at probing both special knowledge related to, e.g., different cultures, politics, and skills required to associate and relate to them.

### Ethical concerns.

The subjectivity of humor also raises ethical questions about humor generation algorithms. A joke that appears funny and harmless to one person may be highly offensive and completely unfunny to another. Additionally, some jokes exploit darker aspects of human nature, such as violence, humiliation, and gender or racial stereotypes. Therefore, research aimed at understanding the perception of humor and the preferences of various cultures is of great importance.

## References

- Aaronson, S. 2009. Essentials of complexity-theoretic stand-up comedy. <https://scottaaronson.blog/?p=414>.
- Agafonova, Y.; Tikhonov, A.; and Yamshchikov, I. P. 2020. Paranoid transformer: Reading narrative of madness as computational approach to creativity. *Future Internet* 12(11).
- Amir, O., and Biederman, I. 2016. The neural correlates of humor creativity. *Front. Hum. Neurosci.* 10:597.
- Arora, S., and Goyal, A. 2023. A theory for emergence of complex skills in language models. <https://arxiv.org/abs/2307.15936>.
- Attardo, S. 1994. *Linguistic Theories of Humor*. Berlin, New York: De Gruyter Mouton.
- Baranov, A.; Kniazhevsky, V.; and Braslavski, P. 2023. You told me that joke twice: A systematic investigation of transferability and robustness of humor detection models. *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing* 13701–13715.
- Bender, E. M.; Gebru, T.; McMillan-Major, A.; and Shmitchell, S. 2021. On the dangers of stochastic parrots: Can language models be too big? *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency* 610–623.
- Chen, Y.; Shi, B.; and Si, M. 2023. Prompt to gpt-3: Step-by-step thinking instructions for humor generation. *14th International Conference on Computational Creativity*.
- Cozman, F., and Munhoz, H. 2020. The winograd schemas from hell. In *Proceedings of the 17th National Meeting on Artificial and Computational Intelligence*, 531–542. Porto Alegre, RS, Brasil: SBC.
- Di Fede, G.; Rocchesso, D.; Dow, S. P.; and Andolina, S. 2022. The idea machine: Llm-based expansion, rewriting, combination, and suggestion of ideas. *Proceedings of the 14th Conference on Creativity and Cognition* 623–627.
- Hessel, J.; Marasović, A.; Hwang, J. D.; Lee, L.; Da, J.; Zellers, R.; Mankoff, R.; and Choi, Y. 2023. Do androids laugh at electric sheep? humor "understanding" benchmarks from the new yorker caption contest.
- Hinton, G., and Ng, A. 2023. A conversation about ai risk. <https://twitter.com/AndrewYNg/status/1667920020587020290?lang=en>.
- Hupkes, D.; Dankers, V.; Mul, M.; and Bruni, E. 2020. Compositionality decomposed: how do neural networks generalise? *Journal of Artificial Intelligence Research* 67:757–795.
- Hurley, M. M.; Dennett, D. C.; and Adams, R. B. J. 2011. *Inside Jokes: Using Humor to Reverse-Engineer the Mind*. The MIT Press.
- Inácio, M.; Wick-pedro, G.; and Goncalo Oliveira, H. 2023. What do humor classifiers learn? an attempt to explain humor recognition models. In Degaetano-Ortlieb, S.; Kazantseva, A.; Reiter, N.; and Szpakowicz, S., eds., *Proceedings of the 7th Joint SIGHUM Workshop on Computational Linguistics for Cultural Heritage, Social Sciences, Humanities and Literature*, 88–98. Dubrovnik, Croatia: Association for Computational Linguistics.
- Jentsch, S., and Kersting, K. 2023. Chatgpt is fun, but it is not funny! humor is still challenging large language models. *Proceedings of the 13th Workshop on Computational Approaches to Subjectivity, Sentiment, and Social Media Analysis* 325–340.
- Koestler, A. 1964. *The Act of Creation: A Study of Conscious and Unconscious Processes of Humor, Scientific Discovery and Art*.
- Loakman, T.; Maladry, A.; and Lin, C. 2023. The iron(ic) melting pot: Reviewing human evaluation in humour, irony and sarcasm generation. *Findings of the Association for Computational Linguistics: EMNLP 2023* 6676–6689.
- Mann, H. B., and Whitney, D. R. 1947. On a Test of Whether one of Two Random Variables is Stochastically Larger than the Other. *The Annals of Mathematical Statistics* 18(1):50–60.
- Michaud, E. J.; Liu, Z.; Girit, U.; and Tegmark, M. 2023. The quantization model of neural scaling. *37th Conference on Neural Information Processing Systems*.
- Mihalcea, and Strapparava. 2005. Making computers laugh: Investigations in automatic humor recognition. *Proceedings of Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing* 531–538.
- Savage, B. M.; Lujan, H. L.; Thipparthi, R. R.; and DiCarlo, S. E. 2017. Humor, laughter, learning, and health! a brief review. *Advances in physiology education*. 341—347.
- Smith, W.; Harrington, K. V.; and Neck, C. P. 2000. Resolving conflict with humor in a diversity context. *Journal of Managerial Psychology* Vol. 15 No. 6:606—625.
- Summers-Stay, D.; Voss, C. R.; and Lukin, S. M. 2023. Brainstorm, then select: a generative language model improves its creativity score.
- Sun, J.; Narayan-Chen, A.; Oraby, S.; Cervone, A.; Chung, T.; Huang, J.; Liu, Y.; and Peng, N. 2022. Expunations: Augmenting puns with keywords and explanations. *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing* 4590–4605.
- Szegedy, C. 2024. <https://twitter.com/ChrSzegedy/status/1750196565409701979>.
- Tikhonov, A., and Yamshchikov, I. 2023. Post Turing: Mapping the landscape of LLM evaluation. In Gehrman, S.; Wang, A.; Sedoc, J.; Clark, E.; Dhole, K.; Chandu, K. R.; Santus, E.; and Sedghamiz, H., eds., *Proceedings of the Third Workshop on Natural Language Generation, Evaluation, and Metrics (GEM)*, 398–412. Singapore: Association for Computational Linguistics.
- Toplyn, J. 2014. *Comedy Writing for Late-Night TV: How to Write Monologue Jokes, Desk Pieces, Sketches, Parodies, Audience Pieces, Remotes, and Other Short-Form Comedy*.
- Toplyn, J. 2021. Witscript: A system for generating improvised jokes in a conversation. In *Proceedings of the 12th International Conference on Computational Creativity (ICCC '21)*.



Toplyn, J. 2022. Witscript 2: A system for generating improvised jokes without wordplay. In *Proceedings of the 13th International Conference on Computational Creativity (ICCC '22)*.

Toplyn, J. 2023. Witscript 3: A hybrid ai system for improvising jokes in a conversation. <https://arxiv.org/abs/2301.02695>.

Veale, T. 2001. *Your Wit Is My Command: Building AIs with a Sense of Humor*. The MIT Press.

Wang, H.; Zou, J.; Mozer, M.; Goyal, A.; Lamb, A.; Zhang, L.; Su, W. J.; Deng, Z.; Xie, M. Q.; Brown, H.; and Kawaguchi, K. 2024. Can ai be as creative as humans? <https://arxiv.org/abs/2401.01623>.

Warren, C.; Barsky, A.; and McGraw, A. P. 2021. What makes things funny? an integrative review of the antecedents of laughter and amusement. *Personality and Social Psychology Review* Vol. 25(1):41–65.

Weller, O., and Seppi, K. 2019. Humor detection: A transformer gets the last laugh. *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)* 3621–3625.

Yu, D.; Kaur, S.; Gupta, A.; Brown-Cohen, J.; Goyal, A.; and Arora, S. 2023. Skill-mix: a flexible and expandable family of evaluations for ai models. <https://arxiv.org/abs/2310.17567>.

#### Appendix A. Inferring humor-generation policy

### Prompt to decompose a single joke into building blocks

You are an expert in humor theory, with extensive knowledge in various forms of comedy ranging from slapstick to satire and dry humor.

Your expertise includes understanding the structural elements of jokes, cultural influences on humor, and the psychology behind what makes things funny.

Additionally, you are a practicing stand-up comedian with years of experience in writing and performing jokes, giving you a unique perspective on how humor resonates with different audiences.

Your task is to analyze the provided joke, identifying its humorous elements, potential audience reception, and categorizing it based on humor types.

Discuss the elements that contribute to its humor, such as the structure of the joke, the play on words, the unexpected twist, or cultural references.

### Prompt to combine individual jokes decompositions into humor-generation policy

"Read several texts about different jokes above. Formulate a list of typical (repeated) elements, used for construction of jokes.

Avoid providing semantic doubles, provide only unique elements and their formal abstract descriptions (without details of any particular joke).

#### Appendix B. Generating new jokes

### Prompt to get associations for topic

Given the theme, provide a numbered list of 20 free associations with it: different meanings, different contexts, stereotypes, puns, exaggerations, incongruity, cultural references, juxtaposition. Use short descriptions -- 1-5 words each.

### Prompt to expand associations

Read the theme and the list of associations with it.

Rewrite and expand each of the associations into longer and more detailed (10-15 words each).

Avoid repetitions and usage the similar words across the sentences. Keep mentioning the original theme.

Provide a numbered list as a result.

### Prompt to refine and combine associations

Read the theme and analyze and review the list of items associated with it.

Write a shorter list of items using the following rules:

- Complementary items can be grouped into one item that combines the original descriptions.
- Weak, unsuccessful, or poorly related items to the original topic should be deleted.
- The final list should be no longer than 6 items, different items should have different meanings or contexts.

Provide a numbered list as a result.

### Prompt to generate joke based on topic and list of associations

Read the theme and analyze and review the list of items associated with it.

Based on that information, write a list of 7-10 jokes using the following rules:

- It should be "One-liner" -- a concise, self-contained joke, delivered in 1-2 sentences with a two-part structure where the first part (setup) establishes a scenario and the second part (punchline) delivers an unexpected twist or conclusion that subverts the setup.

- You may want to use one or several of following strategies:

- Wordplay or pun: The use of words with multiple meanings or similar sounds to create humorous ambiguity or surprise.

- Misdirection: Leading the audience to expect a certain outcome or narrative, only to reveal a different, often contradictory or absurd, conclusion.
  - Exaggeration: Amplifying a characteristic, situation, or behavior to absurd levels to highlight its comedic potential.
  - Stereotyping: Utilizing exaggerated and oversimplified characterizations of groups or individuals for comedic effect.
  - Satire: Using humor, irony, or exaggeration to critique or mock people, institutions, societal norms, or other targets.
  - Absurdity or Surreal Humor: Creating humor through scenarios or statements that are illogical, bizarre, or defy common sense.
  - Dark Humor: Making light of subjects that are generally considered serious, taboo, or morbid.
  - Juxtaposition: Placing two contrasting ideas or scenarios side by side for comedic effect, highlighting their differences.
- Provide a numbered list as a result.

| <b>Step</b>                              | <b>Output</b>  |
|--|--|
| <i>Select Topic</i>                      | chip   |
| <i>Generate associations</i>             | 1. Potato snack - Crispy food 2. Silicon wafer - Computer brains 3. Poker token - Gambling piece 4. Chip shot - Golf term 5. Microchip implant - Pet identification 6. Chip on shoulder - Holding a grudge 7. Chocolate morsel - Cookie ingredient 8. Paint flake - Wall damage 9. Chip and dip - Party food 10. Wood chip - Carpentry byproduct 11. Chip off the old block - Like parent, like child 12. Casino currency - Betting chips 13. Fish and chips - British meal 14. Chip carving - Woodworking technique 15. Chip-seq - DNA sequencing method 16. Chip tuning - Car performance upgrade 17. Chipper - Cheerful person 18. Chip-resistant - Durable nail polish 19. Chip in - Contribute money or effort 20. Blue chip - High-value stock   |
| <i>Expand associations</i>               | 1. A thinly sliced potato, deep-fried to a golden crisp, is a beloved salty snack. 2. The foundational component of electronics, a silicon wafer, serves as the core of modern computing devices. 3. A small, round piece used in betting games, the poker token is essential in casinos. 4. In golf, a chip shot is a tactical, short-range approach to get the ball onto the green. 5. A tiny device implanted under an animal's skin, the microchip implant, helps in reuniting lost pets with their owners. 6. The phrase "chip on one's shoulder" metaphorically describes someone who is easily angered or offended. 7. Essential for baking, a chocolate morsel is often found studded within mouthwatering cookies. 8. A chip of paint peeling away from a surface is indicative of wear or damage to the coating. 9. A popular appetizer, chip and dip involves crunchy chips served with a savory sauce. 10. Small fragments of wood, known as wood chips, are often the remnants of woodworking or lumber processing. 11. The saying "chip off the old block" refers to a child who exhibits traits reminiscent of their parent. 12. In the realm of gambling, casino currency, or betting chips, represent monetary value for placing wagers. 13. A staple in British cuisine, fish and chips consists of battered and fried fish served with fries. 14. Chip carving is a meticulous form of craftsmanship involving the removal of wood to create intricate patterns. 15. Chip-seq is a sophisticated technique used to analyze DNA-protein interactions and gene regulation. 16. Enhancing a vehicle's performance, chip tuning involves modifying the electronic control unit for better efficiency or power. 17. Someone who is consistently in high spirits and optimistic is often described as being chipper. 18. Nail polish that is touted as chip-resistant is formulated to withstand cracking and peeling, maintaining a flawless finish. 19. To "chip in" means to contribute resources or effort towards a common goal or collective task. 20. A blue chip refers to a stock from a well-established company known for its reliability and financial stability. |
| <i>Combine and refine associations</i>   | 1. A beloved salty snack and a popular appetizer, chips are enjoyed in various forms such as thinly sliced and deep-fried potatoes or crunchy chips served with dip. 2. The foundational component of electronics, silicon chips, are the core of modern computing devices and are also used in microchip implants for pets. 3. Essential in gambling and betting games, chips serve as tokens in casinos and represent monetary value for placing wagers. 4. In sports and hobbies, a chip shot in golf is a tactical move, while chip carving is an art form creating intricate patterns in wood. 5. Common phrases include "chip on one's shoulder" describing someone prone to taking offense, and "chip off the old block" referring to a child resembling their parent. 6. In the financial and automotive sectors, a blue chip represents a reliable stock, while chip tuning is a method for enhancing vehicle performance.  |
| <i>Make a joke based on associations</i> | 1. I tried to eat a clock made of chips, but it was time-consuming and I had to go back for seconds. 2. I bought a bag of casino chips, but when I tried to plug them into my computer, it didn't improve my poker game or my internet speed. 3. My dog has a microchip, but he's still terrible at computing – he can't even log in to his own Barkbook account. 4. I told my son he's a chip off the old block, which is great because now I can blame him for the dent in the coffee table. 5. I asked for chip tuning on my car, but all I got was a playlist of The Chipmunks' greatest hits. 6. I put a chip on my shoulder, but nobody took offense; they just asked if I had any dip. 7. I tried chip carving, but it turns out you can't make a good guacamole with wood shavings.  |

Table 3: Sample output of joke generation steps