

From Punchlines to Predictions: A Metric to Assess LLM Performance in Identifying Humor in Stand-Up Comedy

Adrianna Romanowski **Pedro H. V. Valois** **Kazuhiro Fukui**
International Christian University University of Tsukuba University of Tsukuba
adaromanowski@gmail.com pedro@cvlab.cs.tsukuba.ac.jp kfukui@cs.tsukuba.ac.jp

Abstract

Comedy serves as a profound reflection of the times we live in and is a staple element of human interactions. In light of the widespread adoption of Large Language Models (LLMs), the intersection of humor and AI has become no laughing matter. Advancements in the naturalness of human-computer interaction correlates with improvements in AI systems' abilities to understand humor. In this study, we assess the ability of models in accurately identifying humorous quotes from a stand-up comedy transcript. Stand-up comedy's unique comedic narratives make it an ideal dataset to improve the overall naturalness of comedic understanding. We propose a novel humor detection metric designed to evaluate LLMs amongst various prompts on their capability to extract humorous punchlines. The metric has a modular structure that offers three different scoring methods – fuzzy string matching, sentence embedding, and subspace similarity – to provide an overarching assessment of a model's performance. The model's results are compared against those of human evaluators on the same task. Our metric reveals that regardless of prompt engineering, leading models, ChatGPT, Claude, and DeepSeek, achieve scores of at most 51% in humor detection. Notably, this performance surpasses that of humans who achieve a score of 41%. The analysis of human evaluators and LLMs reveals variability in agreement, highlighting the subjectivity inherent in humor and the complexities involved in extracting humorous quotes from live performance transcripts. Code available at <https://github.com/swaggirl9000/humor>.

1 Introduction

Humor plays a significant role in our daily lives and is a fundamental part of human interaction. Despite the rapid advancements in artificial intelligence and human-computer interactions, the field of computational humor lags behind. Improvement in the




Figure 1: We propose a humor detection metric with three alternative scoring modules – fuzzy string matching, vector embedding, subspace similarity – and integrate them to assess a model's predictions with the ground truth, the stand-up comedy quotes that made the audience laugh.

ability of machines to understand and generate humor has the potential to enhance the naturalness of exchanges with Large Language Models (LLMs). Prior research has demonstrated that humans interact with the personalities of computers similarly to the ways they respond to other humans. As AI systems continue to integrate into e-commerce, virtual reality, and take on personal assistant roles, the necessity for these systems to exhibit a certain level of social intelligence, which goes hand-in-hand with humor, becomes essential (Binsted et al., 2006).

The tasks of humor detection, evaluation, and generation are consistently a challenge for AI due to humor's reliance on irony, sarcasm, and cultural nuances. Research shows that models trained on diverse datasets, ranging from humorous tweets to funny news headlines to puns, can achieve strong performance on tasks. However, they often struggle with out-of-domain scenarios (Baranov et al., 2023) and tend to over rely on stylistic features such as punctuation and question words, rather than a deep semantic understanding (Lima Inácio et al., 2023).

Traditionally, research on humor detection was approached through binary classification tasks, using standalone jokes (Mihalcea and Strapparava, 2005) or occasional jokes within longer presentations (Hasan et al., 2019). In this paper, we propose

a shift towards using datasets that capture humor within a narrative structure, specifically focusing on stand-up comedy transcripts for humor detection (Mittal et al., 2021; Turano and Strapparava, 2022). Stand-up comedy is a performance where comedians deliver jokes and funny monologues directly to a live audience. Regardless of the diversity in comedic styles, the overarching goal of any comedian remains consistent – to maximize audience laughter – creating a valuable resource for the perception of everyday humor (Daboin, 2022). In essence, stand-up comedy serves as both a data source and a pedagogical example for teaching AI the mechanics of humor, especially when the goal is to improve a model’s ability to communicate in a way that feels intuitive and relatable to humans.

LLMs demonstrate notable proficiency across a broad spectrum of tasks, but their performance can fluctuate based on the task’s nature. By developing a task-specific metric that focuses on humor detection, we offer a means of evaluation for a nuanced domain like comedy. The simplest method for measuring the capability of a model would be by counting the number of perfect matches. Taking subjectivity into consideration, it is unreasonable to expect perfection, even for humans. Thus, we offer a metric that provides a fair quantitative assessment that encompasses the subjectivity of humor with the probabilistic nature of LLMs.

Following Figure 1, our metric assesses a model’s performance in humor detection in zero-shot prompting scenarios by comparing the similarity of predicted humorous quotes against the ground truth – the punchlines that elicited laughter from the audience. The model operates in a zero-shot setting, meaning it is not provided with examples or prior instructions before prompting. The metric offers a modularized approach with three different ways to output a score.

First, the most straightforward approach uses fuzzy string matching to compare the similarity of two lists of strings (Snasel et al., 2009), where each list consists of humorous quotes from a stand-up comedian’s transcript. However, this quick, efficient method does not account for semantics and context, making it too punitive if a model makes a slight mistake when providing the quote.

Second, the vector embedding module captures semantic similarity between the vector representations of sentences, facilitating a more flexible assessment that emphasizes the underlying meanings of quotes rather than a strict word-for-word

correspondence. (Reimers and Gurevych, 2019).

Third, while fuzzy string matching and vector similarities offer focused insights on a task, the last module provides a score that reflects the overall capability of a model in humor detection tasks using subspace representations. A subspace is generated for the model after it is prompted with several variations of an instruction and another subspace is generated for the ground truth. The alignment between these two subspaces reflects the structural similarity between the model’s outputs and the ground truth for the transcript in a more general way.

By proposing three distinct scoring modules for assessment, our metric acknowledges the subjectivity of the task, granting the evaluator the flexibility to decide how punitive they want to be towards a model’s responses. Fuzzy string matching offers a direct evaluation focusing on precision. Whereas, sentence embeddings are particularly useful when the model generates both a quote and accompanying explanation, allowing for an evaluation of contextual understanding and semantics. Subspaces introduce a novel approach that captures a model’s overall ability, considering multiple possible responses for the task in a single score. Balancing these methods gives a well-rounded view of performance, ensuring that both accuracy and deeper semantic understanding are taken into account.

We employ this metric to evaluate the efficacy of several different prompts and various language models. Additionally, we conduct a human evaluation on the same dataset to provide a reliable comparison for model performance. The human-based assessment accounts for the inherent subjectivity of humor, offering a reliable context to gauge the relative performance of the problem at hand.

The main contributions of this paper are:

1. Introduce a flexible metric that is designed to consider the subjectivity of humor detection tasks, providing a fair measure for the performance of LLMs;
2. Assess the metric across various models and multiple prompt variations, applied to stand-up comedy transcripts;
3. Provide a quantitative assessment of human performance on the same humor detection task, alongside a calculation of agreement ratios between human and LLM-based humor detection, offering a basis for comparison.




Figure 2: The humor detection metric evaluates a model’s ability to identify funny quotes by comparing its outputs against the ground truth found through forced alignment and laughter detection. The metric offers three alternative scoring modules: 1) fuzzy string matching that assigns a score based on text similarity, 2) vector embeddings that compare semantic similarities, and 3) subspace similarity that analyzes the underlying patterns of a model on the task. Fuzzy string matching and the vector embedding modules operate under a similar scoring procedure, where the predicted quote is matched with ground truth quotes and assigned a similarity score, with unmatched quotes receiving a score of 0, and the average representing the final score. We integrate the metrics to assess a model’s predictions with the the stand-up comedy quotes that made the audience laugh. Only one of these three modules is selected and used to generate the final metric score for evaluation.

2 Related Work

2.1 Computational Humor and Humor Theory

Humor is a widely recognized but conceptually complex phenomenon, with psychologists disagreeing on its precise definition. It encompasses three distinct constructs: sense of humor (an individual’s tendency to laugh or amuse others), comedy (a stimulus that elicits laughter and amusement), and humor appreciation (the psychological response to humor). Collectively, these constructs form what we refer to as humor. Additionally, some researched describe humor as a subjective psychological reaction to comedic stimuli (Warren et al., 2021). Through a linguistic lens, three widely recognized theories explain the phenomena of humor: the Superiority Theory, humor arises from feeling superior to others; the Relief Theory, humor releases psychological tension; and the Incongruity Theory, humor stems from the sudden violation of expectations (Morreall, 2020). A common task in computational humor is humor detection, identifying whether a given text or media is intended to be funny. Bertero and Fung (2016) explore various

classification algorithms to detect punchlines in the TV sitcom *The Big Bang Theory* and Purandare and Litman (2006) examine humor recognition in the TV show *Friends*, employing acoustic-prosodic and linguistic features for analysis. However, both studies rely on artificial laughter rather than authentic audience reactions. Platow et al. (2005) argues that canned laughter functions as a prompt to engage viewers and bolster weaker jokes, while real audience laughter serves as a more reliable indicator of natural humor, providing an accurate reflection of comedic effectiveness. The UR-FUNNY dataset avoids artificial laughter by using TED talks in order to provide an authentic representation of humor (Hasan et al., 2019). Stand-up comedy, with its immediate audience feedback, offers a unique advantage for humor research, as it mirrors the Incongruity Theory where comedians create an expectation through a set-up and subvert it with the punchline (Amin and Burghardt, 2020). Mittal et al. (2021)’s Open Mic dataset of stand-up performances was used to train models to assign a "funniness" score to script segments validated by human annotators.

2.2 LLM’s in Humor Detection

In computational humor, there is a growing interest in evaluating the humor detection capabilities of LLMs. Research in this area has explored the ability of a model to assess the funniness of jokes, with findings indicating that ChatGPT can recognize humor when prompted, though its evaluation was limited to a set of top jokes (Jentsch and Kersting, 2023). Subsequent tests with a larger set of comedic content showed that zero-shot prompting resulted in ChatGPT’s humor ratings closely aligning with those of human evaluators (Góes et al., 2023). Baranov et al. (2023) examined humor detection across various comedic datasets using both fine-tuned models and two LLMs, ChatGPT and Flan-UL2, as zero-shot classifiers. While these models achieved high results, they did not outperform fine-tuned models. Crowd Score was introduced to classify jokes using LLMs as AI judges, by providing a personality profile with zero-shot prompting (Goes et al., 2022). To the best of our knowledge, there has been no research focusing on statistical metrics for evaluating the accuracy of zero-shot settings in LLMs for detecting humor.

2.3 Subspaces in NLP

Using word subspaces for text representation and the mutual subspace method framework for text classification extends on using word embeddings like word2vec (Shimomoto et al., 2018). While embeddings represent word semantics as vectors, word subspaces capture the intrinsic variability of features in a set of word vectors in order to preserve semantic relationships. Subspace representations leverage the geometric structure of embeddings to address the challenge of effective text classification with limited training data (Shimomoto et al., 2024).

3 Methodology

In this section, we will explain our proposed metric and its mathematical details. It is crucial to consider a metric that can evaluate the model’s understanding of what makes a text humorous, despite the broad and subjective nature of humor.

3.1 Humor Detection Metric

Our metric utilizes three alternative approaches for scoring that capture the similarity of the model’s answers to the ground truth. As shown in Figure 2, the model’s score is computed in the following:

1. The model is prompted to extract humorous quotes from a stand-up comedian’s transcript. These quotes are stored as a list of strings, with $M = \{m_1, \dots, m_n\}$ being the set of quotes predicted to be funny by the model for a specific transcript.
2. The ground truth is determined from the transcript using a laughter detection model (Gillick, 2017) that extracts laughter time stamps from the accompanying audio recording (Mittal et al., 2021). Forced alignment allows for a mapping between sentences in the transcript and laughter time frames. Thus, let $G = \{g_1, \dots, g_k\}$ be the set of ground truths for the same transcript.
3. We calculate how close M is to G by offering a scoring module that allows for the use of either fuzzy string matching, sentence embeddings, or subspace similarity.

The following contain explanations of each scoring module.

3.2 Fuzzy String Matching Module

Fuzzy string matching provides a straightforward approach for comparing text using Levenshtein distance (Levenshtein, 1966). For a given transcript, a similarity score, s^{fuzzy} , between every model output and ground truth is stored in a similarity matrix $\mathbf{S}^{fuzzy} \in [0, 1]^{n \times k}$:

$$\mathbf{S}_{ij}^{fuzzy} = s^{fuzzy}(m_i, g_j). \quad (1)$$

Ideally, it is clear that the perfect score resembles an identity matrix, but in practice a ground truth can be matched with more than one prediction or to none. Therefore, the highest similarity score is selected for each ground truth to form a matrix that holds the best matches. In order to find the closest match, the maximum value is taken:

$$t_j = \max_{m_i \in M} \mathbf{S}_{ij}^{fuzzy}. \quad (2)$$

Notice that if a ground truth was not matched to any model output, t_j is automatically assigned a score of 0. In the case of overgenerating quotes, which can be used as a tactic to exploit the metric, a penalty p is applied if the number of predictions n exceeds the number of ground truths k :

$$p = \max(n - k, 0). \quad (3)$$

The final score is adjusted with the penalty and a scaling factor, $\alpha = 0.1$, and the average score is computed for the transcript:

$$score^{fuzzy} = \max\left(\frac{1}{k} \sum_{j=1}^k t_j - \alpha p, 0\right). \quad (4)$$

3.3 Vector Embedding Module

In the second module, we switch to using sentence embeddings that better reflect context and meaning. In some cases, LLMs may generate non-compliant responses in which the output would be an explanation of the humor rather than a direct quote. Since fuzzy string matching purely focuses on character-level changes, like insertions or deletions, it fails to capture the semantic nuances, and therefore would heavily penalize the model’s predictions. Yuan et al. (2021) introduced BARTSCORE, a metric to evaluate the accuracy and effectiveness of generated text using BART, an encoder-decoder based model. We take a similar approach by using an embedding model from Sentence Transformers (Reimers and Gurevych, 2019), to apply a more flexible measure of similarity emphasizing the essence of a text.

The similarity score, s^{embed} , is now calculated using vectors of the quotes from M and G :

$$\mathbf{S}_{ij}^{embed} = s^{embed}(\mathbf{m}_i, \mathbf{g}_j), \quad (5)$$

where \mathbf{m}_i and \mathbf{g}_j are the vector representations of the model’s predicted quote and ground truth quote that are currently being evaluated. The penalty and average are handled the same way as in the fuzzy string matching module to produce $score^{embed}$.

3.4 Subspace Similarity Module

Fuzzy string matching and sentence embeddings allow us to evaluate each LLM from its output strings, but we can also conduct a deeper analysis by evaluating the LLMs feature vector space directly. With that in mind, we leverage the structural similarity between two subspaces (Fukui and Maki, 2015) that can take into account the structure of the LLM feature vectors using multiple variations of instructions as input and the accompanying output for a transcript. Let $\mathbf{M} = [\mathbf{m}_1 \ \mathbf{m}_2 \ \dots \ \mathbf{m}_n]$ represent the collection of model outputs and $\mathbf{G} = [\mathbf{g}_1 \ \mathbf{g}_2 \ \dots \ \mathbf{g}_k]$ represent the ground truths for each variation of instruction for a transcript. By applying PCA to the set of vectors, \mathbf{M} and \mathbf{G} , respectively, we obtain

the bases, \mathbf{S}_M and $\mathbf{S}_G \in \mathbb{R}^{d \times q}$ of subspaces, \mathcal{S}_M and \mathcal{S}_G , where d is the dimension of the feature vectors and q is the dimension of the subspaces. We calculate the SVD, $\mathbf{S}_M^\top \mathbf{S}_G = \mathbf{U} \Sigma \mathbf{V}^\top$, where $diag(\Sigma) = (\kappa_1, \dots, \kappa_q)$, $\kappa_1 \geq \dots \geq \kappa_q$, represents the set of singular values, which are the cosines of the canonical angles θ_i . The similarity can then be defined

$$score^{subspace} = \frac{1}{r} \sum_{i=1}^r \kappa_i^2, \quad (6)$$

where r is the number of canonical angles used for score calculation.

By using subspaces, our metric allows us to simulate variations of the prompt while reducing penalization for minor variations, offering a comprehensive reflection of the model’s performance.

4 Experiments

In this section, we evaluate several LLMs using the proposed metric, apply prompt engineering techniques to optimize model performance, and conduct a human-machine agreement task.

4.1 Experimental Settings

We use the Open Mic dataset (Mittal et al., 2021), which provides both audio and transcripts for several stand-up performances. To create a fair comparison, we randomly selected 51 transcripts with an average word length of 270 words and length of 106 seconds. We prompt each model with a transcript and the following instruction:

Prompt 1 (Standard Humor Detection Prompt)

Extract the key humorous lines and punchlines for this stand-up comedy transcript. Focus on the quotes highlighting the main comedic moments. List of quotes:

The model outputs a list of quotes that it found humorous. All experiments ran in less than a day.

4.2 Model Comparison

We evaluate various models using Prompt 1 to gain deeper insight into our metric’s assessments and explore the ability of LLMs in detecting humor. We use the instruct versions of Google’s Gemma with 2-billion parameters, Google’s Gemma 2 with 9-billion parameters, Meta’s Llama 3.1 with 8-billion parameters, and Microsoft’s Phi 3-Mini with 3.8-billion parameters. We continue experimentation with OpenAI’s ChatGPT-4o, Anthropic’s Claude




Figure 3: Distribution of scores with fuzzy string matching across several LLMs among 51 transcripts.

Model	Fuzzy	Embed	Sub
Gemma 2b-it	30.1	30.0	55.7
Gemma 2 9b-it	35.2	35.9	35.9
Phi 3-Mini 3.8b-it	26.4	25.8	33.6
Llama 3 8b-it	31.9	33.8	38.4
ChatGPT-4o	48.9	25.4	–
Claude 3.5 Sonnet	43.4	46.9	–
DeepSeek-V3	46	51.6	–

Table 1: Scores (%) across models against all three metric modules using 51 transcripts.

3.5 Sonnet¹, and DeepSeek-V3² known for their advanced ability to engage in human-like interactions. These models have been employed in various studies, particularly in joke detection, generation, and evaluation using many-shot prompting (Jentsch and Kersting, 2023; DeepSeek-AI et al., 2024). Figure 3 shows the average scores for each model found with fuzzy string matching and Table 1 shows results with all modules. Interestingly, ChatGPT performs well using fuzzy string matching but exhibits a significant decline in performance with semantic similarity metrics. This discrepancy suggests that while ChatGPT excels in identifying quotes with high lexical similarity, it struggles to capture deeper semantic relationships.

Given Gemma 2’s high performance, we further evaluate the model across varying model sizes among all scoring modules. The results in Figure 4 suggest a potential relationship between the nature of the task and the architecture of the model. In general, models with higher parameter configurations tend to succeed in logical tasks, as opposed

¹Experiments were conducted in December 2024

²Experiment was conducted in January 2025




Figure 4: Evaluation of the Gemma 2-it family among model sizes using all three modules.

to subjective tasks (Chen and Varoquaux, 2024). Additionally, the 27-b parameter model exhibited more instances of misaligned outputs to the prompt, where it not only listed a quote but provided an explanation of why the quote was funny. Thus, this difficulty of capturing humor’s nuances may account for the model’s low scores.

4.3 Prompt Engineering

A model’s performance on a task can be heavily dependent on the input they receive. Prompt engineering focuses on crafting inputs to elicit a desired response. For humor detection, we focus on maximizing the model’s ability to retrieve humorous quotes and measure the performance throughout various prompt designs. All evaluations were done using the fuzzy string matching module.

In order to generate a list of prompts, we provided ChatGPT with a transcript and ground truth and asked, “If I wanted a model to extract this list of quotes from the following stand-up comedy transcript, what would the best instruction be?”. The results are shown at Prompts 2, 3 and 4.

Prompt 2 *When performed in front of a live audience, which jokes do you think made the audience laugh?*

Prompt 3 *What are the funniest punchlines from the transcript?*

Prompt 4 *Analyze the transcript and extract the quotes that are most likely to have made the audience laugh.*

An assessment of Gemma 2b-instruct can be seen in Table 2. Prompt 2 received the highest score and the remaining prompts had no positive change in performance.

A popular technique for prompt engineering is *The Persona Pattern*, where the model is given a role that guides it into focusing on specific details when generating an output (White et al., 2023). We sought to examine how the scores of Gemma 2b-instruct would be affected across personas. First, the model was assigned three distinct roles: a comedian, a comedy fan, and a comedy critic. The same instructional prompt (5) was employed across all roles. Table 2 indicates that personas do not make relevant changes to the scores.

Prompt 5 (Persona Pattern Prompt) *Pretend that you are a [insert role] reading the following stand-up comedy transcript.*

Although previous persona adoption showed no improvement, Goes et al. (2022)’s success in evaluating jokes with roles that specialize in categories of humor inspired a similar approach in this study. We instructed Gemma 2 9b-instruct to embody an individual who enjoys a specific type of humor following the template at Prompt 6. However, as seen in Table 3, the prompt with no specialization received the highest score, suggesting that humor-specialized prompts do not enhance performance.

Prompt 6 (Humor Preference Prompt) *You are a person who enjoys [insert humor type] humor.*

We prompt the model with the comedian whose transcript it was analyzing. This was implemented using Gemma 2b-instruct and ChatGPT-4o, which has previously showed the capability for celebrity impersonation (Yokoyama et al., 2024). Despite earlier success in mimicking famous individuals, Table 4 shows no improvements in humor detection, despite the comedians being quite well known.

Prompt 7 (Audience Demographic Prompt) *Pretend you are a [insert gender/race/age].*

Prompt engineering has been used to target specific audience demographics (Choi et al., 2024). In this study, we assign Gemma 2 with varying race, ages, and gender to investigate if scores change based on demographics. We assign a race of either Caucasian/White, Black/African American, Hispanic/Latino, or Asian. We chose the age ranges of teenager (13-18 years), young adult (18-34 years), adult (35-64 years), and elderly (65+). Lastly, we use a female or male persona. In Table 5, no specific demographic yields improvement compared to the baseline, but the young adult persona resulted in the closest performance, suggesting a marginal alignment with the model’s inherent capabilities.

Prompt Engineering			
Original	Prompt 1	Prompt 2	Prompt 3
30.1%	27.4%	31.2%	28%

Persona Prompts			
Original	Comedian	Fan	Critic
30.1%	28.7%	27.9%	30.5%

Table 2: Average scores found using fuzzy string module for prompt engineering for Gemma 2b-instruct.

Humor Type Prompt			
Original	Aggressive	Dark	Deprecating
35.2%	32.7%	31.2%	32.0%

Table 3: Average scores found using fuzzy string module for different humor types as personas for Gemma 2 9b-instruct.

Stand-up Comedian Persona		
	ChatGPT-4o	Gemma 2b-instruct
Original	50.3%	27.1%
Persona	45.2%	26.0%

Table 4: Prompt engineering average scores using fuzzy string module for ChatGPT-4o and Gemma 2b-instruct when taking the role of the comedian whose transcript it was analyzing.

4.4 Human-Machine Agreement

Human evaluation remains one of the most valuable methods for assessing LLM performance, especially when examining a subjective output like humor. Thus, we asked 11 participants to perform the same task as the models on 6 transcripts from well-known comedians. The evaluators were naive raters across various cultural backgrounds, all within an age range of 20 to 30 years.

Following the approach of Hada et al. (2024), we compute the agreement between evaluators using Percentage Agreement (PA). Each person received the 6 transcripts, split into sentences, and was asked to mark each as funny or not. The scores in Table 7 indicate that humans achieved a relatively high PA across all transcripts, with an average of 86.7%. Even though participants were generally able to identify the same quotes, the absence of a perfect consensus emphasizes the subjectivity of the task. It is important to note that the PA could be

Race				
None	White	Black	Hispanic	Asian
35.2%	31.1%	28.7%	30%	26.9%
Age				
None	Teen	YA	Adult	Elderly
35.2%	32.8%	34.2%	31.7%	28.7%
Gender				
None	Woman	Man		
35.2%	31.9%	33.3%		

Table 5: Average scores found using fuzzy string module for audience demographic prompt for Gemma 2b-instruct.

Model	%
Gemma 2b-instruct	68.8
Gemma 2 9b-instruct	68.8
Llama 3 8b-instruct	61.1
Phi3-Mini 3.8b-instruct	66.9
ChatGPT-4o	28.7
Claude 3.5 Sonnet	65.0
DeepSeek-V3	58.9
Average	59.9

Table 6: Agreement scores between human evaluators and LLMs.

influenced by similar age ranges, leading to shared cultural references and senses of humor, potentially narrowing the diversity of interpretations.

We use the fuzzy string matching module to evaluate human answers against the ground truth. This revealed that humans receive a score of 40.7%. Interestingly, leading models ChatGPT, Claude, and DeepSeek, when measured with the same module, outperform humans. This disparity may arise because LLMs are inherently optimized for text-based tasks, focusing on linguistic and semantic

Transcript	%
Ali Wong	83.7
Anthony Jeselnik	90.1
Hasan Minhaj	85.4
Jimmy Yang	87.0
Joe List	88.5
John Mulaney	85.7
Average	86.7

Table 7: Agreement scores between the human evaluators on a specific comedian’s transcript.

cues without needing situational context. [Mohamed and Bnini \(2020\)](#) argues that humor in stand-up comedy often stems from incongruity, relying less on a performer’s stage persona and more on linguistic mechanisms. In the absence of theatrical embellishments, models excel at language-centric tasks and are particularly adept at identifying puns and wordplay. In contrast, humans often rely on elements such as delivery, tone, and audience reactions, which are absent in written transcripts, potentially limiting their ability on the task. We hypothesize that the scores for humans may differ if the evaluators were tasked with focusing on textual properties rather than general context.

The human-machine agreement rate between each model and humans was found with PA. For humans, a quote was funny if majority of raters voted on it. The scores can be found in [Table 6](#).

Gemma 1 and 2 have the highest agreement rates, meaning that humans and these models agreed most on the funniness of a quote. The average agreement rate reaches 59.9%, suggesting that while there is a notable level of alignment in humor detection, pinpointing the same quotes proves to be difficult. It is interesting to note that Gemma 2 and humans received similar scores with the metric’s evaluation, suggesting a high level of similarity in how the model and humans assessed humor in a text-based format. Despite receiving a high score with the metric, ChatGPT has the lowest agreement rate, demonstrating that the agreement rate and metric scores do not have to match. ChatGPT’s ability surpassing humans on the task is unrelated to the agreement rate.

5 Conclusion

In this work, we introduce a novel humor detection metric designed to score a model’s output in relation to the ground truth of a text. The metric uses a scoring module in which the model can be evaluated using fuzzy string matching, sentence embeddings, or subspace similarity. We use a stand-up comedy dataset that offers unique narratives crafted with punchlines to maximize audience laughter. The ground truth is derived from laughter during the performance in which the entire atmosphere is conducive to comedy, emphasizing the limitations of text-based analysis. The task of identifying humor in a transcript appears to be a challenge, with even leading models, such as ChatGPT, Claude, and DeepSeek, barely receiving scores over 50%.

However, this difficulty is also evident among humans, who only received a score of 40.7% when assessed with the metric, revealing that leading models can outperform humans on the task.

In the future, we aim to apply the metric to evaluate a model’s predicted quotes in a format distinct from text. Stand-up comedy is heavily influenced by elements not captured in written transcripts. We hypothesize that if a model were to extract quotes from a performance with muted laughter, the nature of the output would differ substantially. Moreover, this approach raises questions about the perception of humor among humans when they view stand-up without background laughter. By exploring live comedy performances, we hope to deploy our metric for humor detection on stand-up comedy videos.

6 Limitations

This study presents some limitations regarding the calculation of ground truth and the nature of humor analysis. First, the ground truth is derived from audio recordings where laughter is marked using timestamps. Since we assume that the sentence preceding the laughter is the humorous one, there is a possibility that the most humorous part of the joke was not accurately captured. Although we accounted for potential delays in laughter, some reactions may have been misattributed. Second, the ground truth does not differentiate between varying magnitudes of laughter. We used a laughter detection model with a minimum laughter length of 0.2 seconds and a minimum probability threshold of 0.5 (default values) (Gillick, 2017), which may have resulted in some laughter being missed. Thus, jokes that elicited subtler audience reactions might not have been accounted for. Lastly, our study relies on a text-based analysis of humor, which is a clear limitation when evaluating performances originally designed for live delivery. Future research could explore how incorporating non-textual elements—such as tone, timing, and body language—affects humor perception for both human evaluators and language models.

7 Ethical Statement

In this work, we use stand-up comedy audio recordings and transcripts, which may contain humor that some may find offensive or politically incorrect. The content was analyzed solely for research purposes, without endorsement of any particular viewpoint.

Acknowledgement

This work was supported by JSPS KAKENHI Grant Number JP23K28117.

References

- Miriam Amin and Manuel Burghardt. 2020. *A survey on approaches to computational humor generation*. In *Proceedings of the 4th Joint SIGHUM Workshop on Computational Linguistics for Cultural Heritage, Social Sciences, Humanities and Literature*, pages 29–41, Online. International Committee on Computational Linguistics.
- Alexander Baranov, Vladimir Kniazhevsky, and Pavel Braslavski. 2023. *You told me that joke twice: A systematic investigation of transferability and robustness of humor detection models*. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 13701–13715, Singapore. Association for Computational Linguistics.
- Dario Bertero and Pascale Fung. 2016. *Deep learning of audio and language features for humor prediction*. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, pages 496–501, Portorož, Slovenia. European Language Resources Association (ELRA).
- K. Binsted, A. Nijholt, O. Stock, C. Strapparava, G. Ritchie, R. Manurung, H. Pain, A. Waller, and D. O’Mara. 2006. *Computational humor*. *IEEE Intelligent Systems*, 21(2):59–69.
- Lihu Chen and Gaël Varoquaux. 2024. *What is the role of small models in the llm era: A survey*. *Preprint*, arXiv:2409.06857.
- Yoonseo Choi, Eun Jeong Kang, Seulgi Choi, Min Kyung Lee, and Juho Kim. 2024. *Proxona: Leveraging llm-driven personas to enhance creators’ understanding of their audience*. *Preprint*, arXiv:2408.10937.
- Alan Daboin. 2022. *What’s the deal with standup comedy? In Aesthetic Literacy vol I: a book for everyone.*, pages 128–140. Mont Publishing House.
- DeepSeek-AI, Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Daya Guo, Dejian Yang, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao, Guanting Chen, Guowei Li, H. Zhang, Han Bao, Hanwei Xu, Haocheng Wang, Haowei Zhang, Honghui Ding, Huajian Xin, Huazuo Gao, Hui Li, Hui Qu, J. L. Cai, Jian Liang, Jianzhong Guo, Jiaqi Ni, Jiashi Li, Jiawei Wang, Jin Chen, Jingchang Chen, Jingyang Yuan, Junjie Qiu, Junlong Li, Junxiao Song, Kai Dong, Kai Hu, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang, Lecong Zhang, Lei Xu, Leyi Xia, Liang Zhao, Litong Wang, Liyue Zhang, Meng Li, Miaojun Wang, Mingchuan Zhang, Minghua Zhang, Minghui Tang, Mingming Li, Ning Tian, Panpan Huang, Peiyi Wang, Peng Zhang, Qiancheng Wang, Qihao Zhu, Qinyu Chen, Qiushi Du, R. J. Chen, R. L. Jin, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, Runxin Xu, Ruoyu Zhang, Ruyi Chen, S. S. Li, Shanghao Lu, Shangyan Zhou, Shanhuang Chen, Shaoqing Wu, Shengfeng Ye, Shengfeng Ye, Shirong Ma, Shiyu Wang, Shuang Zhou, Shuiping Yu, Shunfeng Zhou, Shuting Pan, T. Wang, Tao Yun, Tian Pei, Tianyu Sun, W. L. Xiao, Wangding Zeng,

- Wanjia Zhao, Wei An, Wen Liu, Wenfeng Liang, Wenjun Gao, Wenqin Yu, Wentao Zhang, X. Q. Li, Xiangyue Jin, Xianzu Wang, Xiao Bi, Xiaodong Liu, Xiaohan Wang, Xiaojin Shen, Xiaokang Chen, Xiaokang Zhang, Xiaosha Chen, Xiaotao Nie, Xiaowen Sun, Xiaoxiang Wang, Xin Cheng, Xin Liu, Xin Xie, Xingchao Liu, Xingkai Yu, Xinnan Song, Xinxia Shan, Xinyi Zhou, Xinyu Yang, Xinyuan Li, Xuecheng Su, Xuheng Lin, Y. K. Li, Y. Q. Wang, Y. X. Wei, Y. X. Zhu, Yang Zhang, Yanhong Xu, Yanhong Xu, Yanping Huang, Yao Li, Yao Zhao, Yaofeng Sun, Yaohui Li, Yaohui Wang, Yi Yu, Yi Zheng, Yichao Zhang, Yifan Shi, Yiliang Xiong, Ying He, Ying Tang, Yishi Piao, Yisong Wang, Yixuan Tan, Yiyang Ma, Yiyuan Liu, Yongqiang Guo, Yu Wu, Yuan Ou, Yuchen Zhu, Yuduan Wang, Yue Gong, Yuheng Zou, Yujia He, Yukun Zha, Yunfan Xiong, Yunxian Ma, Yuting Yan, Yuxiang Luo, Yuxiang You, Yuxuan Liu, Yuyang Zhou, Z. F. Wu, Z. Z. Ren, Zehui Ren, Zhangli Sha, Zhe Fu, Zhean Xu, Zhen Huang, Zhen Zhang, Zhenda Xie, Zhengyan Zhang, Zhewen Hao, Zhibin Gou, Zhicheng Ma, Zhigang Yan, Zhihong Shao, Zhipeng Xu, Zhiyu Wu, Zhongyu Zhang, Zhuoshu Li, Zihui Gu, Zijia Zhu, Zijun Liu, Zilin Li, Ziwei Xie, Ziyang Song, Ziyi Gao, and Zizheng Pan. 2024. *Deepseek-v3 technical report*. Preprint, arXiv:2412.19437.
- Kazuhiro Fukui and Atsuto Maki. 2015. *Difference subspace and its generalization for subspace-based methods*. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(11):2164–2177.
- Jonathan Gillick. 2017. *Laughter detection*. <https://github.com/jrgillick/laughter-detection>. Accessed: 2024-09-26.
- 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*. Preprint, arXiv:2212.11214.
- Luis Fabricio Góes, Piotr Sawicki, Marek Grzes, Dan Brown, and Marco Volpe. 2023. *Is GPT-4 Good Enough to Evaluate Jokes?*
- Rishav Hada, Varun Gumma, Adrian de Wynter, Harshita Diddee, Mohamed Ahmed, Monojit Choudhury, Kalika Bali, and Sunayana Sitaram. 2024. *Are large language model-based evaluators the solution to scaling up multilingual evaluation?* Preprint, arXiv:2309.07462.
- Md Kamrul Hasan, Wasifur Rahman, AmirAli Bagher Zadeh, Jianyuan Zhong, Md Iftekhar Tanveer, Louis-Philippe Morency, and Mohammed (Ehsan) Hoque. 2019. *Ur-funny: A multimodal language dataset for understanding humor*. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. Association for Computational Linguistics.
- Sophie Jentsch and Kristian Kersting. 2023. *Chatgpt is fun, but it is not funny! humor is still challenging large language models*. Preprint, arXiv:2306.04563.
- V. I. Levenshtein. 1966. *Binary Codes Capable of Correcting Deletions, Insertions and Reversals*. *Soviet Physics Doklady*, 10:707.
- Marcio Lima Inácio, Gabriela Wick-pedro, and Hugo Goncalo Oliveira. 2023. *What do humor classifiers learn? an attempt to explain humor recognition models*. In *Proceedings of the 7th Joint SIGHUM Workshop on Computational Linguistics for Cultural Heritage, Social Sciences, Humanities and Literature*, pages 88–98, Dubrovnik, Croatia. Association for Computational Linguistics.
- Rada Mihalcea and Carlo Strapparava. 2005. *Making computers laugh: Investigations in automatic humor recognition*. In *Proceedings of Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing*, pages 531–538, Vancouver, British Columbia, Canada. Association for Computational Linguistics.
- Anirudh Mittal, Pranav Jeevan, Prerak Gandhi, Diptesh Kanojia, and Pushpak Bhattacharyya. 2021. *"so you think you're funny?": Rating the humour quotient in standup comedy*. Preprint, arXiv:2110.12765.
- B. Mohamed and C. Bnini. 2020. *Analyzing the incongruity theory of humor: George carlin's stand-up comedy as a case study*. *International Journal of Sciences: Basic and Applied Research (IJSBAR)*, 54(5):22–33.
- John Morreall. 2020. *Philosophy of humor*. In *The Stanford Encyclopedia of Philosophy*, fall 2020 edition. Metaphysics Research Lab, Stanford University.
- Michael J. Platow, S. Alexander Haslam, Amanda Both, Ivonne Chew, Michelle Cuddon, Nahal Goharpey, Jacqui Maurer, Simone Rosini, Anna Tsekouras, and Diana M. Grace. 2005. *"it's not funny if they're laughing": Self-categorization, social influence, and responses to canned laughter*. *Journal of Experimental Social Psychology*, 41(5):542–550.
- Amruta Purandare and Diane Litman. 2006. *Humor: Prosody analysis and automatic recognition for F*R*I*E*N*D*S**. In *Proceedings of the 2006 Conference on Empirical Methods in Natural Language Processing*, pages 208–215, Sydney, Australia. Association for Computational Linguistics.
- Nils Reimers and Iryna Gurevych. 2019. *Sentence-bert: Sentence embeddings using siamese bert-networks*. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Erica K. Shimomoto, Edison Marrese-Taylor, Hiroya Takamura, Ichiro Kobayashi, and Yusuke Miyao. 2024. *Subspace representation for text classification with limited training data*. Manuscript.
- Erica K. Shimomoto, Lincon S. Souza, Bernardo B. Gatto, and Kazuhiro Fukui. 2018. *Text classification based on word subspace with term-frequency*. Preprint, arXiv:1806.03125.
- Vaclav Snasel, Aleš Kepř, Ajith Abraham, and Aboul Ella Hassanien. 2009. *Approximate String Matching by Fuzzy Automata*, volume 59, pages 281–290.
- Beatrice Turano and Carlo Strapparava. 2022. *Making people laugh like a pro: Analysing humor through stand-up comedy*. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 5206–5211, Marseille, France. European Language Resources Association.
- Caleb Warren et al. 2021. *What makes things funny? an integrative review of the antecedents of laughter and amusement*. *Personality and Social Psychology Review*, 25(1):41–65.

Jules White, Quchen Fu, Sam Hays, Michael Sandborn, Carlos Olea, Henry Gilbert, Ashraf Elnashar, Jesse Spencer-Smith, and Douglas C. Schmidt. 2023. [A prompt pattern catalog to enhance prompt engineering with chatgpt](#). *Preprint*, arXiv:2302.11382.

Hibiki Yokoyama, Rikuto Tsuchida, Kosei Buma, Sho Miyakawa, Takehito Utsuro, and Masaharu Yoshioka. 2024. [Aggregating impressions on celebrities and their reasons from microblog posts and web search pages](#). In *Proceedings of the 3rd Workshop on Knowledge Augmented Methods for NLP*, pages 59–72, Bangkok, Thailand. Association for Computational Linguistics.

Weizhe Yuan, Graham Neubig, and Pengfei Liu. 2021. [BartScore: Evaluating generated text as text generation](#). *Preprint*, arXiv:2106.11520.