

Why Do We Laugh? Annotation and Taxonomy Generation for Laughable Contexts in Spontaneous Text Conversation

Koji Inoue, Mikey Elmers, Divesh Lala, Tatsuya Kawahara

Graduate School of Informatics, Kyoto University, Japan,

Correspondence: inoue@sap.ist.i.kyoto-u.ac.jp

Abstract

Laughter serves as a multifaceted communicative signal in human interaction, yet its identification within dialogue presents a significant challenge for conversational AI systems. This study addresses this challenge by annotating laughable contexts in Japanese spontaneous text conversation data and developing a taxonomy to classify the underlying reasons for such contexts. Initially, multiple annotators manually labeled laughable contexts using a binary decision (laughable or non-laughable). Subsequently, an LLM was used to generate explanations for the binary annotations of laughable contexts, which were then categorized into a taxonomy comprising ten categories, including “Empathy and Affinity” and “Humor and Surprise,” highlighting the diverse range of laughter-inducing scenarios. The study also evaluated GPT-4o’s performance in recognizing the majority labels of laughable contexts, achieving an F1 score of 43.14%. These findings contribute to the advancement of conversational AI by establishing a foundation for more nuanced recognition and generation of laughter, ultimately fostering more natural and engaging human-AI interactions.

1 Introduction

In human dialogue, laughter serves as a communicative signal conveying humor, empathy, surprise, or social bonding (Norrick, 1993; Glenn, 2003; Attardo, 2009). However, its mechanisms are complex and multifaceted, and understanding them remains a long-term challenge for dialogue systems aiming to achieve human-like interaction (Tian et al., 2016; Türker et al., 2017; Mazzocconi et al., 2020; Inoue et al., 2022; Ludusan and Wagner, 2023; Perkins Booker et al., 2024). Furthermore, traditional approaches to modeling laughter and humor have often been limited to scenarios involving explicit auditory or visual stimuli, with few addressing the subtle contextual nuances present in sponta-

Table 1: Annotation example for laughable context (majority voting, translated from Japanese)

	Utterance	Laughable?
A:	I think that’s a wonderful attitude. I always end up talking about myself, so I should follow your example.	NO
B:	Is that so? But does your husband listen to your stories?	NO
A:	Yes, yes, he listens to me. I wonder if I’m putting too much on him?	NO
B:	I don’t think so! He’s so kind. My husband doesn’t seem to listen to me. Huh, that’s strange.	YES

neous dialogue (Bertero and Fung, 2016; Choube and Soleymani, 2020; Jentsch and Kersting, 2023; Ko et al., 2023; Hessel et al., 2023). Therefore, elucidating the underlying reasons for laughter in spontaneous dialogue data can contribute to making large language model (LLM)-based dialogue more natural and empathetic. However, annotating the reasons for laughter in any formalized manner has been prohibitively time- and labor-intensive, leaving the field largely reliant on qualitative approaches through conversational analysis.

In this study, we address the question of “why do we laugh?” from an informatics perspective by proposing a semi-automated approach to constructing taxonomy labels for the reasons of laughter. First, to identify target segments, multiple annotators were asked to perform a simple binary classification on each utterance in dialogue data, determining whether it was “laughable” or not, as shown in Table 1. Subsequently, for contexts labeled as “laughable” based on the majority voting, we used an LLM (GPT-4o) to generate the reasoning sentence behind this judgment and further classified these generated reasons into distinct categories (taxonomy labels). This semi-automated taxonomy generation approach is generalizable and can be particularly effective in scenarios where manual annotation is limited to simpler labels, such as emo-

Table 2: Number of samples in each ratio of annotators judged as laughable (laughable agreement)

Laughable agreement	# sample
1.0 (5/5)	163 (0.64%)
0.8 (4/5)	845 (3.34%)
0.6 (3/5)	2731 (10.80%)
0.4 (2/5)	8143 (32.20%)
0.2 (1/5)	11928 (47.17%)
0.0 (0/5)	1479 (5.85%)

tion labeling.

The purpose of this research is to contribute toward more nuanced conversational AI systems that can recognize and even anticipate moments for laughter, ultimately fostering more natural interactions between humans and machines. Ideally, such systems should be able to respond with the correct acoustics, delay, and consider group size for different laughter types (Truong and Trouvain, 2012). Our findings reveal that AI can improve our understanding of laughter and offer a foundation for future research in AI context-sensitive recognition.

2 Annotation of Laughable Context

We annotated laughable contexts in the RealPersonaChat dataset (Yamashita et al., 2023). This textual data contains one-on-one Japanese spontaneous conversation where participants chat without assuming assigned personas. It includes approximately 30 utterances per conversation, totaling around 14,000 dialogues. We annotated 900 dialogues, with plans to annotate the remainder in future work.

During the annotation process, each annotator reviewed each dialogue and, after the initial two greeting utterances, made a binary decision for whether the next person would laugh (*laughable*) or not. Five annotators assigned these binary labels to each utterance. Table 2 summarizes the agreement amongst annotators for laughable labels, which we refer to as “laughable agreement”. While some samples showed clear agreement (either all or none of the annotators marked them as laughable), there were also numerous split samples, highlighting the subjectivity and complexity of the task. If we applied a majority voting process, 3,739 contexts (14.8%) were labeled as laughable, and 21,550 contexts (85.2%) as non-laughable.

Table 1 illustrates a laughable context example. In this dialogue, person A’s final utterance is self-contradictory, requiring high-level comprehension

of the dialogue context. These annotations underscore the significance of cultural context and conversational flow in interpreting laughter cues.

3 Generating Taxonomy of Laughable Reason

Towards developing human-like laughter behaviors in LLMs, we investigated the reasons behind human annotators’ recognition of laughable contexts. In this section, we used only samples with majority labels marked as laughable (3,739 samples). Since manual annotations are costly, we utilized GPT-4o to generate explanations for the human judgments. For example, a generated reason for the example context in Table 1 is:

If we were to speculate on the reasons a third party might judge that Person A laughed at Person B’s final remark in this conversation, the following points can be considered:

Element of Humor: Person B’s comment, “My husband doesn’t seem to listen to me. Huh, that’s strange.,” contains a touch of self-deprecating humor. This lighthearted tone, making fun of their own situation, can be amusing to the listener. (...)

Relaxed Atmosphere: The overall tone of the conversation seems light and relaxed, and Person B’s comment might have been perceived as a playful joke in line with this mood.

A combination of these factors may have led the third party to interpret that Person A laughed in response to Person B’s remark.

We then aimed to summarize the generated reasoning texts for laughable contexts by applying a taxonomy generation approach using LLMs (Wan et al., 2024). First, we randomly divided the generated reason samples into smaller subsets, each comprising roughly 5% of the samples. Starting with the first subset, we gave the reason sentences GPT-4o to generate initial taxonomy labels and those explanations, which we manually validated when necessary. We then iteratively refined the taxonomy by having the LLM update it based on the previous taxonomy and the reason sentence data from each new subset, continuing this process until all data were processed. This resulted in ten taxonomy labels, summarized in Table 3, including categories such as (1) Empathy and Affinity and (2) Humor and Surprise.

After generating these taxonomy labels, we used the LLM to assign them to each reason sample,

Table 3: Generated taxonomy labels for laughable context reasoning, number of samples assigned to each taxonomy label, and related references for each taxonomy label

	Label name	Explanation	#sample	Reference
(1)	Empathy and Affinity	Situations where a sense of closeness and laughter is generated by sharing common experiences or emotions in a conversation. This includes empathy for shared hobbies or everyday relatable situations.	3013 (80.6%)	(Hay, 2001; Garbarski et al., 2016)
(2)	Humor and Surprise	Cases where humor or an element of surprise in the statement triggers laughter. This includes unexpected twists, wordplay, and exaggeration.	3233 (86.5%)	(Dyner, 2009; Martin and Ford, 2018)
(3)	Relaxed Atmosphere	Situations where the conversation progresses in a calm, relaxed atmosphere, naturally leading to laughter. Lighthearted exchanges and conversations with jokes fall into this category.	2955 (79.0%)	(Vettin and Todt, 2004)
(4)	Self-Disclosure and Friendliness	Situations where sharing personal stories or past mistakes creates a sense of approachability and triggers laughter. Self-disclosure that makes the other person feel at ease is also included.	475 (12.7%)	(Gelkopf and Kreitler, 1996)
(5)	Cultural Background and Shared Understanding	Laughter based on specific cultural backgrounds or shared understandings. This includes jokes related to a particular region or culture or remarks based on common superstitions or folklore.	176 (4.7%)	(Bryant and Bainbridge, 2022; Kamiloğlu et al., 2022)
(6)	Nostalgia and Fondness	Situations where past memories or nostalgic topics trigger laughter. This includes shared past experiences and the enjoyment of recalling familiar events.	204 (5.5%)	(Bazzini et al., 2007)
(7)	Self-Deprecating Humor	Situations where turning one’s flaws or mistakes into humor makes the other person feel closer and triggers laughter. Slightly putting oneself down can also give the other person a sense of ease.	404 (10.8%)	(Atkinson, 2015)
(8)	Defying Expectations	Situations where intentionally defying the flow of conversation or the other person’s expectations creates an element of surprise and triggers laughter. This includes unexpected responses or developments.	323 (8.6%)	(Ginzburg et al., 2020; Xu, 2022)
(9)	Positive Energy	Situations where positive emotions or energy in the conversation bring a smile to the other person. Enjoyable topics and positive comments fall into this category.	338 (9.0%)	(Wang et al., 2024)
(10)	Exaggeration	Situations where exaggerating things gives a comical impression and triggers laughter. Exaggerated expressions can be amusing to the listener.	478 (12.8%)	(McCarthy and Carter, 2004)

allowing for multiple labels per sample. The labeling results are shown on the right side of Table 3. While some categories, such as (1) Empathy and Affinity, were predominant, many samples were also assigned to other categories, including (4) Self-Disclosure and Friendliness and (5) Cultural Background and Shared Understanding. This broad distribution across categories reinforces the validity of the generated taxonomy. A correlation matrix showing relationships between the taxonomy labels is provided in Appendix A. Finally, we reviewed related studies in conversational analysis, as listed on the right side of Table 3. These studies further substantiate the explanatory power of our taxonomy within the context of conversational analysis research.

4 LLM’s Performance on Laughable Context Recognition

We then examined how much LLMs, specifically GPT-4o, can recognize the laughable contexts in spontaneous text conversation. The model was

tested in a zero-shot setting, instructed to first analyze the conversational context and then determine its laughability as a binary. The provided prompt included a task description for laughable context recognition, followed by a Chain-of-Thought (CoT) reasoning approach to encourage the model to consider the reasoning behind its decision step by step. We evaluated GPT-4o’s performance against the majority labels, achieving an F1 score of 43.14%, with a precision of 41.66% and recall of 44.72%. While this score was significantly above the chance level (14.8%), capturing the nuanced subtleties of conversational humor remains challenging.

We then further examined the LLM’s performance on each generated taxonomy label. Table 4 shows the distribution of binary outputs by GPT-4o and its accuracy within each label. First, the primary labels, from (1) to (3), showed similar accuracy rates, ranging from 40% to 50%. Additionally, we observed comparatively higher scores for (5) *Cultural Background and Shared Understanding*, (7) *Self-Deprecating Humor*, and (8) *Defying Ex-*

Table 4: Laughable context recognition performance of GPT-4o on each taxonomy label

	Label name	# output by GPT-4o	
		Laughable (correct)	Non-laughable (incorrect)
(1)	Empathy and Affinity	1226 (40.69%)	1787 (59.31%)
(2)	Humor and Surprise	1571 (48.59%)	1662 (51.41%)
(3)	Relaxed Atmosphere	1257 (42.54%)	1698 (57.46%)
(4)	Self-Disclosure and Friendliness	232 (48.84%)	243 (51.16%)
(5)	Cultural Background and Shared Understanding	102 (57.95%)	74 (42.05%)
(6)	Nostalgia and Fondness	62 (30.39%)	142 (69.61%)
(7)	Self-Deprecating Humor	255 (63.12%)	149 (36.88%)
(8)	Defying Expectations	227 (70.28%)	96 (29.72%)
(9)	Positive Energy	50 (14.79%)	288 (85.21%)
(10)	Exaggeration	239 (50.00%)	239 (50.00%)

Table 5: Example context for “Nostalgia and Fondness” (translated from Japanese)

Utterance
A: Do you also consume milk or yogurt for calcium?
B: I drink milk with Milo in it. I also eat yogurt as a snack.
A: That’s really well-balanced!
B: Yes, health is important.
A: It’s been a while since I last heard about Milo.

Table 6: Example context for “Positive Energy” (translated from Japanese)

Utterance
A: Oh, as they grow up, that kind of help really makes a difference, doesn’t it?
B: Absolutely! It’s such a joy, isn’t it? So reassuring.
A: When they’re little, it’s like a never-ending story of challenges, isn’t it?
B: Haha, so true. All we have now are funny memories of those times.
Once you get through it, those challenges become stories
A: you can laugh about, and you feel glad you went through them.

pectations, suggesting that the current LLM may effectively capture these contexts. In contrast, categories like (6) *Nostalgia and Fondness* and (9) *Positive Energy* displayed lower accuracy, potentially highlighting limitations in the LLM’s understanding.

Table 5 presents an example dialogue context where the LLM marked non-laughable for the final utterance from person A, despite a positive majority label with a (6) *Nostalgia and Fondness* reason. This context was also assigned the (2) *Humor and Surprise* and (3) *Relaxed Atmosphere* labels. In this example, the participants discuss a nostalgic memory of drinking a powdered beverage with milk. The last utterance evokes nostalgia, implicitly inviting laughter. Here, capturing person A’s sentiment seems to be difficult for the current LLM, but is

essential for appropriate laughter response.

Table 6 provides an example for (9) *Positive Energy* label. This context was also assigned the (1) *Empathy and Affinity* and (2) *Humor and Surprise* labels. The participants discussed a challenging experience with childcare, but in the final utterance, person A reflects positively on the experience after some time has passed. Although the story itself recounts a difficult time, it is now viewed positively, making it laughable. This example suggests that the LLM needs to comprehend the temporal structure of the story and the person’s current feelings to accurately interpret the context as laughable.

5 Conclusion

This study investigated laughter in the context of conversational AI by annotating laughable contexts within a Japanese text dialogue dataset. A taxonomy of ten distinct reasons for laughter was generated by an LLM, providing valuable insights into the multifaceted nature of laughter. Subsequently, this study evaluated the ability of GPT-4o to recognize those laughable contexts. While the model’s performance surpassed chance levels, it highlighted the inherent challenges in capturing the nuances of conversational humor.

This automated approach employed for reasoning and taxonomy generation with LLMs can be applied in other scenarios where only binary (or simplified) decision labels from human annotators are available, yet more fine-grained explanations are required. Future work will focus on expanding the dataset to cover other languages and cultural contexts, validating the generated taxonomy by incorporating additional linguistic research perspectives, exploring multimodal approaches, and including spoken dialogue to enhance AI’s understanding of humor and social interaction.

Acknowledgement

This work was supported by JST PREST JP-MJPR24I4 and JSPS KAKENHI JP23K16901.

References

- Camille Atkinson. 2015. Self-deprecation and the habit of laughter. *Florida Philosophical Review*, 15(1).
- Salvatore Attardo. 2009. *Linguistic theories of humor*. Walter de Gruyter.
- Doris G Bazzini, Elizabeth R Stack, Penny D Martincin, and Carmen P Davis. 2007. The effect of reminiscing about laughter on relationship satisfaction. *Motivation and Emotion*, 31(1):25–34.
- Dario Bertero and Pascale Fung. 2016. A long short-term memory framework for predicting humor in dialogues. In *Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pages 130–135.
- Gregory A Bryant and Constance M Bainbridge. 2022. Laughter and culture. *Philosophical Transactions of the Royal Society B*, 377(1863):20210179.
- Akshat Choube and Mohammad Soleymani. 2020. Punchline detection using context-aware hierarchical multimodal fusion. In *International Conference on Multimodal Interaction (ICMI)*, pages 675–679.
- Marta Dynel. 2009. Beyond a joke: Types of conversational humour. *Language and linguistics compass*, 3(5):1284–1299.
- Dana Garbarski, Nora Cate Schaeffer, and Jennifer Dykema. 2016. Interviewing practices, conversational practices, and rapport: Responsiveness and engagement in the standardized survey interview. *Sociological methodology*, 46(1):1–38.
- Marc Gelkopf and Shulamith Kreitler. 1996. Is humor only fun, an alternative cure or magic? the cognitive therapeutic potential of humor. *Journal of Cognitive Psychotherapy*, 10(4).
- Jonathan Ginzburg, Chiara Mazzocconi, and Ye Tian. 2020. Laughter as language. *Glossa: a journal of general linguistics*, 5(1).
- Phillip Glenn. 2003. *Laughter in interaction*. Cambridge University Press.
- Jennifer Hay. 2001. *The pragmatics of humor support*. Walter de Gruyter.
- Jack Hessel, Ana Marasovic, Jena D. Hwang, Lillian Lee, Jeff Da, Rowan Zellers, Robert Mankoff, and Yejin Choi. 2023. Do androids laugh at electric sheep? humor “understanding” benchmarks from the new yorker caption contest. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 688–714.
- Koji Inoue, Divesh Lala, and Tatsuya Kawahara. 2022. Can a robot laugh with you?: Shared laughter generation for empathetic spoken dialogue. *Frontiers in Robotics and AI*, 9.
- Sophie Jentsch and Kristian Kersting. 2023. ChatGPT is fun, but it is not funny! humor is still challenging large language models. In *Workshop on Computational Approaches to Subjectivity, Sentiment, & Social Media Analysis (WASSA)*, pages 325–340.
- Roza G Kamiloğlu, Akihiro Tanaka, Sophie K Scott, and Disa A Sauter. 2022. Perception of group membership from spontaneous and volitional laughter. *Philosophical Transactions of the Royal Society B*, 377(1841):20200404.
- Dayoon Ko, Sangho Lee, and Gunhee Kim. 2023. Can language models laugh at YouTube short-form videos? In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2897–2916.
- Bogdan Ludusan and Petra Wagner. 2023. The effect of conversation type on entrainment: Evidence from laughter. In *Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL)*, pages 168–174.
- Rod A Martin and Thomas Ford. 2018. *The psychology of humor: An integrative approach*. Academic press.
- Chiara Mazzocconi, Ye Tian, and Jonathan Ginzburg. 2020. What’s your laughter doing there? A taxonomy of the pragmatic functions of laughter. *IEEE Transactions on Affective Computing*, 13(3):1302–1321.
- Michael McCarthy and Ronald Carter. 2004. “there’s millions of them”: hyperbole in everyday conversation. *Journal of pragmatics*, 36(2):149–184.
- Neal R Norrick. 1993. Conversational joking: Humor in everyday talk.
- Nynaeve Perkins Booker, Michelle Cohn, and Georgia Zellou. 2024. Linguistic patterning of laughter in human-socialbot interactions. *Frontiers in Communication*, 9.
- Ye Tian, Chiara Mazzocconi, and Jonathan Ginzburg. 2016. When do we laugh? In *Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL)*, pages 360–369.
- Khiet P. Truong and Jürgen Trouvain. 2012. On the acoustics of overlapping laughter in conversational speech. In *INTERSPEECH*, pages 851–854.
- Bekir Berker Türker, Zana Buçinca, Engin Erzin, Yücel Yemez, and T Metin Sezgin. 2017. Analysis of engagement and user experience with a laughter responsive social robot. In *INTERSPEECH*, pages 844–848.




Figure 1: Correlation matrix of assigned taxonomy labels

latter explicit. To gain deeper insight into the relationships between these labels, further qualitative analysis will be conducted in future work.

Julia Vettin and Dietmar Todt. 2004. Laughter in conversation: Features of occurrence and acoustic structure. *Journal of Nonverbal Behavior*, 28:93–115.

Mengting Wan, Tara Safavi, Sujay Kumar Jauhar, Yujin Kim, Scott Counts, Jennifer Neville, Siddharth Suri, Chirag Shah, Ryen W White, Longqi Yang, Reid Andersen, Georg Buscher, Dhruv Joshi, and Nagu Rangan. 2024. TnT-LLM: Text mining at scale with large language models. In *SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*, pages 5836–5847.

Kexin Wang, Carlos Ishi, and Ryoko Hayashi. 2024. A multimodal analysis of different types of laughter expression in conversational dialogues. In *INTER-SPEECH*, pages 4673–4677.

Ge Xu. 2022. An analysis of humor discourse in friends from the perspective of the cooperative principle. *Open Journal of Modern Linguistics*, 12(4):460–470.

Sanae Yamashita, Koji Inoue, Ao Guo, Shota Mochizuki, Tatsuya Kawahara, and Ryuichiro Higashinaka. 2023. Realpersonachat: A realistic persona chat corpus with interlocutors’ own personalities. In *Pacific Asia Conference on Language, Information and Computation (PACLIC)*, pages 852–861.

A Correlation Among Taxonomy Labels

Figure 1 presents the correlation matrix of the assigned labels discussed in Section 3, where multiple labels can be assigned to the same laughable context. For instance, “Empathy and Affinity” shows a weak positive correlation with “Relaxed Atmosphere.” Conversely, “Empathy and Affinity” exhibits a negative correlation with “Defying Expressions.” We also find a negative correlation between “Humor and Surprise” and “Positive Energy,” despite both being associated with positive sentiment. This may be attributed to different expressive styles, with the former implicit and the