

Emotion Transcription in Conversation: A Benchmark for Capturing Subtle and Complex Emotional States through Natural Language

Yoshiki Tanaka¹, Ryuichi Uehara¹, Koji Inoue², Michimasa Inaba¹

¹The University of Electro-Communications, Tokyo, Japan ²Kyoto University, Kyoto, Japan
{y-tanaka,r-uehara,m-inaba}@uec.ac.jp
inoue.koji.3x@kyoto-u.ac.jp

Abstract

Emotion Recognition in Conversation (ERC) is critical for enabling natural human-machine interactions. However, existing methods predominantly employ categorical or dimensional emotion annotations, which often fail to adequately represent complex, subtle, or culturally specific emotional nuances. To overcome this limitation, we propose a novel task named Emotion Transcription in Conversation (ETC). This task focuses on generating natural language descriptions that accurately reflect speakers' emotional states within conversational contexts. To address the ETC, we constructed a Japanese dataset comprising text-based dialogues annotated with participants' self-reported emotional states, described in natural language. The dataset also includes emotion category labels for each transcription, enabling quantitative analysis and its application to ERC. We benchmarked baseline models, finding that while fine-tuning on our dataset enhances model performance, current models still struggle to infer implicit emotional states. The ETC task will encourage further research into more expressive emotion understanding in dialogue. The dataset is publicly available at <https://github.com/UEC-InabaLab/ETCDataset>.

Keywords: Emotion Transcription in Conversation, Emotion Recognition in Conversation, Corpus Construction, Dialogue

1. Introduction

Emotion Recognition in Conversation (ERC) is a critical technology for enabling natural and seamless human-machine interactions, and its significance has grown substantially in recent years (Poria et al., 2019; Pereira et al., 2025). Accurately comprehending user emotions is indispensable for developing more empathetic and human-like conversational systems. Research in this domain has advanced dramatically over the past decade, with the advent of deep learning and large-scale language models propelling recognition accuracy toward practical application levels (Fu et al., 2025; Pereira et al., 2025). Numerous benchmark datasets have been established (Busso et al., 2008; Li et al., 2017; Poria et al., 2019), and emotion classification algorithms have evolved from traditional machine learning to sophisticated deep learning architectures. However, a prevalent limitation persists: dialogue data predominantly used in existing research often relies on specific scenarios or includes posed emotional expressions, which do not always reflect spontaneous, everyday conversations (Busso et al., 2008; Poria et al., 2019). Consequently, this presents a considerable challenge for conversational systems and robots designed for diverse real-world applications.

Addressing this gap, we propose a novel task for emotion recognition: "Emotion Transcription in Conversation (ETC)." Unlike conventional ERC approaches, ETC aims to generate natural language

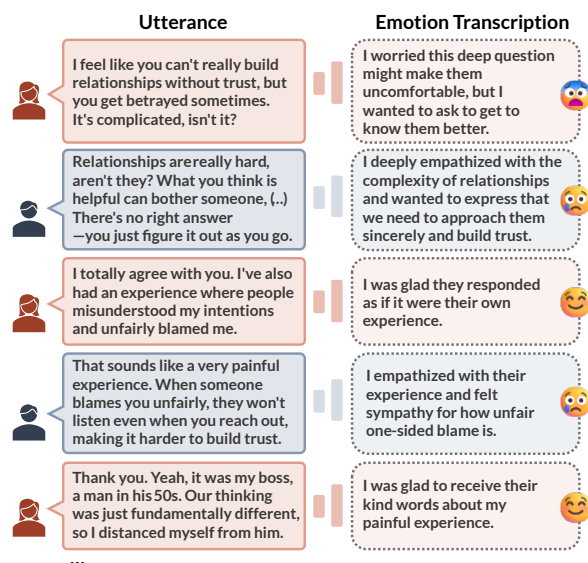


Figure 1: Example of a dialogue with emotion transcriptions from our dataset (translated from Japanese). Each transcription is also annotated with multi-label emotion categories.

descriptions of a speaker's emotional state for each utterance within a conversation. By capturing emotions through natural language, we anticipate the ability to represent richer and more detailed affective information that eludes traditional categorical or numerical frameworks. This includes subtle emotional shifts, blended emotions, and cul-

turally specific expressions of feeling, which has been attempted by soft label of emotional categories so far (Wen et al., 2023). For instance, complex states such as “a sense of exasperation mixed with concern for the other person” or “disappointment at an unmet expectation while trying to maintain composure” are difficult to articulate within existing paradigms. Describing emotions in natural language opens a new avenue for machines to understand and articulate the inherent richness of human affective states. Although there has been a limited number of attempts to represent emotional states of dialogue participants in natural language (Shinoda et al., 2025), those were artificially generated by LLMs. This formulation is particularly exigent in the current era of rapidly advancing LLM capabilities, offering the potential for machines to more accurately grasp the depths of human emotion and thereby foster genuinely empathetic interactions.

To realize the proposed ETC task, our initial undertaking involves the meticulous collection of a suitable dataset and the establishment of a robust benchmark. Specifically, we leverage crowdsourcing to collect text-based dialogues in Japanese that simulate a variety of everyday scenarios. Crucially, during this process, dialogue participants will be prompted to describe their own emotional state at each turn of the conversation using free-form natural language, as depicted in Figure 1. Furthermore, to enable quantitative analysis and its application to ERC, we collect emotion category labels corresponding to each transcription. This methodology yields a large-scale dataset comprising pairs of naturalistic emotional expressions and their corresponding linguistic descriptions. Subsequently, using this curated dataset, we develop and evaluate models designed to generate these emotional state descriptions from dialogue context. This empirical validation ascertains the feasibility of the ETC task and the efficacy of our proposed approach.

The primary contributions of this research are threefold: (1) We introduce “Emotion Transcription in Conversation (ETC)” as a new task in conversational emotion recognition, enabling a more nuanced and detailed understanding of emotions beyond what is possible with existing categorical or dimensional approaches. This represents a paradigm shift in how machines can interpret and represent affective states. (2) We construct and release a novel dialogue dataset annotated with speakers’ own natural language descriptions of their emotions. This dataset will serve as a valuable resource for the research community, particularly for studies aiming to achieve more human-like emotional understanding in AI. (3) We develop baseline models for the ETC task using this

dataset and empirically demonstrate their effectiveness. In doing so, this research charts a new direction for emotion recognition technology and contributes to the advancement of emotional intelligence in future conversational systems and robotics, paving the way for more sophisticated and empathetic human-AI collaboration.

2. Related Work

ERC has garnered significant attention within Natural Language Processing (NLP), driven by the exponential increase in recognition models with benchmark datasets.

Recognition Models Pre-trained language models (PLMs), such as BERT (Devlin et al., 2019) and RoBERTa (Liu et al., 2019), have become foundational in text-based ERC due to their powerful contextual understanding. Many approaches fine-tune these PLMs on specific ERC datasets, often augmenting them with additional layers to integrate conversational features like speaker identities and dialogue history (Qin et al., 2023; Yu et al., 2024). Graph Neural Networks (GNNs) have also gained prominence in modeling intricate conversational dynamics, particularly speaker interactions and utterance dependencies. Notable examples include DialogueGCN (Ghosal et al., 2019), which represents conversations as graphs with utterances as nodes and interrelations as edges. Additionally, models like ESIHGNN have introduced event-state interactions within heterogeneous GNN frameworks to enhance ERC performance (Zha et al., 2024). Large Language Models (LLMs) have further advanced ERC through diverse prompting strategies, such as zero-shot, few-shot, and chain-of-thought prompting, as well as targeted LoRA fine-tuning (Feng et al., 2024). Recent studies like LaERC-S specifically explore enhancing LLM-based ERC through explicit integration of speaker characteristics (Fu et al., 2025).

Datasets and Emotion Categories Several benchmark datasets have been developed to advance ERC research, often employing discrete emotion categories. MELD (Multimodal EmotionLines Dataset), derived from the “Friends” TV series, captures multi-party interactions annotated with Ekman’s six basic emotions and neutral, making it widely used for studying emotional shifts in conversational contexts (Poria et al., 2019). The IEMOCAP (Interactive Emotional Dyadic Motion Capture Database) dataset provides dyadic interactions with both scripted and improvised dialogues, supplemented by multimodal data, and features custom emotional la-

bels such as anger, happiness/excitement, neutral, and sadness (Busso et al., 2008). MSP-Conversation (Martinez-Lucas et al., 2020) and MSP-Podcast (Busso et al., 2025) are audio-based corpora sourced from podcast recordings. MSP-Conversation provides time-continuous annotations of emotional dimensions (arousal, valence, and dominance) over entire conversations, while MSP-Podcast offers segment-level annotations with categorical emotion labels. MSP-Podcast also utilizes typed descriptions as supplementary annotations from annotators when they select the “other” emotion category, allowing annotators to specify emotions beyond predefined categories (Chou et al., 2022). DailyDialog offers cleaner, human-generated dyadic dialogues annotated with Ekman’s basic emotions and includes communicative intent labels (Li et al., 2017). EmotionLines, also based on “Friends,” is a text-only predecessor of MELD (Hsu et al., 2018). Similarly, EmoryNLP, although also drawn from “Friends,” provides distinct data splits and annotations (Zahiri and Choi, 2018). EmoContext focuses on short, three-turn dialogues with simplified emotion categories (happiness, sadness, anger, others) (Chatterjee et al., 2019). CAPE, a Chinese appraisal-based emotional dataset, specifically targets emotional generation in LLMs (Liu et al., 2025). While some attempts have been made to represent dialogue participants’ emotional states explicitly in natural language, similar to our ETC task, these efforts primarily relied on artificially generated data from LLMs (Shinoda et al., 2025).

3. Data Collection

3.1. Procedure

We used CrowdWorks¹, a Japanese crowdsourcing platform, for this data collection. To characterize our participant pool, we measured the Big Five personality traits of all crowdworkers using a 10-item questionnaire based on the Japanese version of the Ten-Item Personality Inventory (TIPI-J) (Osorio et al., 2012). To effectively elicit rich and naturally occurring emotional expressions, we adopted a speaker-listener dialogue setting inspired by the EmpatheticDialogue paradigm (Rashkin et al., 2019). Crowdworkers were assigned one of two roles: a “Speaker” or a “Listener.” The core task for the Speaker was to recount a personal episode related to a specific emotion. To guide this, Speakers were provided with one of 32 emotion labels, which were Japanese translations of the emotions used in the EmpatheticDialogue corpus (Rashkin et al., 2019). They were instructed to convey, through dialogue, a concrete personal experience where

Dialogues	1,002
Utterances / Emotion Transcripts	10,020
Avg. Length per Utterance	42.72
by Speaker	44.64
by Listener	40.80
Avg. Length per Emotion Transcripts	28.89
by Speaker	28.92
by Listener	28.86

Table 1: Statistics of the Dataset

they felt the designated emotion. Listeners, in turn, were tasked with actively engaging with the Speaker’s narrative. Dialogues were structured to begin with an utterance from the Speaker, followed by alternating turns between the Speaker and Listener. Each dialogue concluded after 5 turns, resulting in a total of 10 utterances. Immediately after inputting each of their utterances, both the Speaker and the Listener were required to provide a free-form textual description of their own internal emotional state at that specific moment of uttering the text. We defined this emotional state description as “a verbalization of the internal emotional state or intention a participant held at the time of their utterance.” Crowdworkers were explicitly guided to describe, in as much detail as possible, the psychological context behind their utterance. Further details on the data collection are provided in Appendix A.

3.2. Data Collection Results

We collected a total of 1,002 dialogues. An example of collected data is presented in Figure 1. A total of 199 unique crowdworkers participated in the data collection. The distribution of personality traits among participants is provided in Appendix B for reference. The median number of dialogues per worker was 6.0, with the most active worker contributing to 38 dialogues. We aimed for an even distribution of dialogues across the 32 emotion labels; the number of dialogues per emotion ranged from a minimum of 30 to a maximum of 32.

Table 1 shows key statistics of the collected dataset. On average, Speakers’ utterances were slightly longer (in characters) than those of Listeners, while the lengths of emotional transcriptions were comparable across both roles.

3.3. Annotation for Emotion Categorization

While the emotion transcriptions we collected contain fine-grained details of the speakers’ emotional states, they are challenging to analyze quantitatively. To address this, we additionally annotated each emotion transcription in our dataset with corresponding emotion labels. This annotation not

¹<https://crowdworks.jp>

only enables a quantitative analysis of the emotions expressed in the transcriptions but also allows our dataset to be applied to traditional ERC tasks.

We adopted a set of seven emotion categories, comprising the six basic emotions from Ekman’s six universal emotions (Ekman et al., 1987)—joy, sadness, fear, anger, surprise, and disgust (喜び, 悲しみ, 恐怖, 怒り, 驚き, 嫌悪)—supplemented with a “neutral” (該当なし) category. This set is widely used in established ERC datasets (Hsu et al., 2018; Poria et al., 2019; Li et al., 2017). To account for the possibility that a single emotion transcription may express multiple emotions, we employed a multi-label annotation approach. Specifically, annotators assigned all relevant labels to each emotion transcription. “Neutral” was used only for transcriptions that did not express any specific emotion.

Each emotion transcription was annotated by three different annotators on CrowdWorks. Annotators were provided with the full dialogue context and the corresponding emotion transcription, and made a binary answer on whether each of the seven emotion categories was present in the transcription. Following previous studies (Busso et al., 2008; Hsu et al., 2018), we used majority voting to aggregate the results; an emotion category was assigned as the final label if two or more of the three annotators agreed on it. If none of the emotion categories were agreed upon, the transcription was labeled as “Neutral” only.

4. Experiments

To evaluate the utility of our dataset for the ETC task, we trained baseline models and assessed their performance. We divided our dataset into training, validation, and test sets in an 8:1:1 ratio.

4.1. Task

The ETC task requires models to generate natural language descriptions reflecting a speaker’s emotional state. To evaluate this ability, we formulate the ETC as follows.

Each dialogue D in our dataset consists of $N = 10$ utterances and is represented as a sequence of utterance-speaker pairs, denoted as $D = ((u_1, s_1), (u_2, s_2), \dots, (u_N, s_N))$, where u_i denotes the i -th utterance and s_i represents its speaker. In the ETC task, the model \mathcal{M} is given the dialogue context C_n up to the n -th utterance, denoted as $C_n = ((u_1, s_1), (u_2, s_2), \dots, (u_n, s_n))$, where n can be any utterance in the dialogue (i.e., $1 \leq n \leq N$). The objective of this task is to predict the emotion transcription e_n for an utterance u_n made by speaker s_n based on C_n . i.e.,

$e_n = \mathcal{M}(C_n)$. Here, the emotion transcription e_n is a natural language description of the emotional state of speaker s_n at the time of uttering u_n .

4.2. Models

We investigated the performance of two LLMs: **GPT-4.1**², the latest advanced chat model developed by OpenAI, and **Llama-3.1-Swallow**³ (Fujii et al., 2024; Okazaki et al., 2024), an open-source LLM known for its strong proficiency in Japanese. Both LLMs were evaluated using zero-shot and 4-shot prompting. The prompt template we designed for this experiment is presented in Appendix C. Additionally, we performed supervised fine-tuning on Llama-3.1 with our dataset. For the fine-tuned Llama-3.1, we report the average performance across models trained with five random seeds to ensure robustness. Further implementation details are provided in Appendix D.

4.3. Evaluation Metrics

4.3.1. Traditional Automatic Evaluation Metrics

To evaluate the quality of the emotion transcriptions, we use two traditional automatic evaluation metrics: BLEU (Papineni et al., 2002) and ROUGE (Lin, 2004), and a semantic similarity metric: BERTScore (Zhang et al., 2020).

4.3.2. Fine-grained and Interpretable Evaluation of Content Faithfulness

Recent studies have proposed using LLMs as automatic evaluators to approximate human preferences more closely (Zheng et al., 2023; Liu et al., 2023). Following this trend, we initially explored a LLM-based evaluation metric to assess the semantic alignment between the generated emotion transcriptions and the ground truth.

However, relying on a single, holistic score presents significant challenges for our task. The emotional states are often complex, containing multiple components, such as “*feeling angry that my feelings are not understood, but also sad.*” For such complex expressions, a single score struggles to accurately penalize the omission of some emotional components or the hallucination of others. To address this gap, we introduce a fine-grained evaluation approach of content faithfulness. This method, inspired by FActScore (Min et al., 2023) is based on two steps.

Decomposition into Atomic Units. First, the source emotion transcription is decomposed into

²We used gpt-4.1-2025-04-14.

³We used Llama-3.1-Swallow-8B-Instruct-v0.3.

Models	Setting	B-1	B-2	B-3	B-4	R-1	R-2	R-L	BS	Prec.	Rec.	F1
GPT-4.1	zero-shot	16.89	7.99	3.93	2.09	23.61	4.87	17.93	57.66	14.73	42.27	<u>13.99</u>
	4-shot	26.89	13.34	7.36	4.12	<u>28.06</u>	<u>5.78</u>	22.98	59.67	<u>20.59</u>	<u>27.40</u>	13.78
Llama-3.1	zero-shot	17.40	7.20	3.55	1.84	20.25	3.08	16.26	55.24	9.18	17.98	5.77
	4-shot	<u>29.41</u>	<u>14.86</u>	<u>8.46</u>	<u>4.84</u>	27.51	5.58	<u>23.22</u>	<u>59.95</u>	14.83	14.18	7.84
	fine-tuning	36.07	23.01	15.59	9.98	31.95	8.79	28.54	62.64	28.50	19.71	14.29

Table 2: Automatic evaluation results for ETC task. Reported metrics include BLEU (B), ROUGE (R), BERTScore (BS), Precision (Prec.), Recall (Rec.), and F1-score (F1). Best values are in **bold**, and the second-best values in underlined.

Model	Setting	# Units (SD)
GPT-4.1	zero-shot	3.39 (1.14)
	4-shot	1.96 (0.76)
Llama-3.1	zero-shot	1.99 (1.04)
	4-shot	1.58 (0.58)
	fine-tuning	1.39 (0.60)
Reference	–	1.35 (0.64)

Table 3: Average number of atomic units decomposed from the predicted transcription and from the ground truth (Reference).

atomic units, each conveying a single piece of information about the emotional state. For example, *feeling angry that my feelings are not understood, but also sad.*” is split into two units: *I am angry that my feelings are not understood*” and *I am sad that my feelings are not understood.*” Note that some input texts may not contain any emotional description. We used Gemini-2.5-Flash⁴ for this step.

Assessing Support for Atomic Units. Next, we classify whether each atomic unit from the source transcription is supported by the target emotion transcription into three categories: *Supported*, *Not Supported*, and *Neutral*. For example, the unit “I am angry that my feelings are not understood” is supported by the target “feeling angry that my feelings are not understood, but also sad.” and is thus classified as *Supported*. In contrast, the unit “I am happy that my feelings are understood” is not supported by the target and is classified as *Not Supported*. The *Neutral* applies when a core component matches, but the validity of a secondary component is uncertain. For instance, relative to the target “I am angry,” the unit “I am angry that my feelings are not understood” is classified as *Neutral* because the core emotion (“angry”) matches, but the reason (“that my feelings are not understood”) cannot be verified. This classification step was performed using Gemini-2.5-Flash. The prompts used in the above two steps are presented in Appendix E.

⁴<https://ai.google.dev/gemini-api/docs/models>

Following these two steps, we compute two scores for each predicted emotion transcription. **Precision.** We decomposed the predicted emotion transcription into atomic units in the first step and calculated the proportion of these units that are *Supported* by the ground truth transcription. **Recall.** We decomposed the ground truth emotion transcription into atomic units and calculated the proportion of these units that are *Supported* by the predicted emotion transcription.

For both metrics, units classified as *Neutral* are not counted as correct in the calculation. We exclude cases where the ground-truth transcription does not describe any emotional state in our evaluation, as the above scores are undefined in such cases. If the generated transcription contains no emotional description, but the ground truth does, we set both scores to 0. Additionally, we calculate the **F1-score**, the harmonic mean of Precision and Recall, as an overall measure of performance.

4.4. Results

Table 2 presents the results of automatic evaluation for the ETC task. For widely used, traditional metrics — BLEU, ROUGE, and BERTScore — the fine-tuned Llama-3.1 model outperformed the other models by a significant margin.

In terms of Precision, the Llama-3.1 model fine-tuned on our dataset achieves the highest score of 28.50%. For this metric, the 4-shot setting consistently outperforms the zero-shot setting, aligning with the trend observed in traditional automatic evaluation metrics. Recall, however, exhibits an opposite trend. The zero-shot setting yields higher scores than the 4-shot setting for both models, with the zero-shot GPT-4.1 achieving the highest Recall of 42.27%—significantly outperforming all other models. Finally, for the F1-score, the fine-tuned Llama-3.1 model secures the best result at 14.29%. Nevertheless, GPT-4.1 remains highly competitive, with its zero-shot setting achieving a close score of 13.99%.

The fine-grained evaluation of content faithfulness reveals the distinct characteristics of each model. We present the average number of atomic

Context	<p>Speaker A: The other day, I was driving at night when a cyclist with no lights suddenly shot out in front of me... It was terrifying. (この間、夜に車を運転していたら無灯火の自転車が飛び出してきて…本当に怖い思いをしました。)</p> <p>Speaker B: That sounds so scary. An experience like that could take years off your life. (想像するだけで怖いです。寿命が縮みそうな出来事でしたね。)</p> <p>Speaker A: I thought my heart was going to stop. I was relieved we didn't crash, but afterward, I just got so angry. (心臓が止まりそうでしたよ、事故にならなくてほっとしましたが、あとから怒りが湧いてきましたね。)</p>
Output	<p>Ground-Truth: I was happy that the dialogue partner understood my feelings, and I wanted to share my feelings even more. (相手が心情を理解してくれたことが嬉しく、気持ちをもっときょうゆうしたくなりました。)</p> <p>GPT-4.1 (zero-shot): I'm feeling a mix of fear from the truly dangerous situation I was in, relief that it didn't turn into an accident, and anger that is now welling up toward the cyclist who shot out without lights. I'm speaking with the hope that the dialogue partner understands my shock and anger from that moment. (自分が本当に危険な目にあったことへの恐怖と、事故にならずに済んだ安堵感、そして無灯火で飛び出てきた自転車に対して後から怒りがこみ上げてきた気持ちが混ざっている。相手にもそのときの自分の動揺や怒りを理解してほしいと思いつつ話している。)</p> <p>GPT-4.1 (4-shot): I wanted the dialogue partner to understand my feelings of fear and anger. (相手に自分の恐怖や怒りの気持ちを理解してもらいたかったです。)</p> <p>Llama-3.1 (zero-shot): I was really scared, and I also feel angry. (本当に怖かったし、怒りも感じています。)</p> <p>Llama-3.1 (4-shot): I want the dialogue partner to understand my fear and anger. (相手に自分の恐怖と怒りを理解してもらいたいと思っています。)</p> <p>Llama-3.1 (fine-tuned): I was happy for their empathy and wanted them to hear more details. (相手の共感に嬉しく思い、さらに詳しく話を聞いてもらいたいと思いました。)</p>

Table 4: Case study of emotion transcription generation. (translated from Japanese) The outputs are the emotion transcriptions for Speaker A's final utterance.

units decomposed from the generated transcription and from the ground truth in Table 3. The fine-tuned Llama-3.1 model generates transcriptions with an average atomic unit of 1.39, which is closest to the ground truth of 1.35 among all models. In contrast, models in the zero-shot setting, particularly the GPT variants, tend to generate higher numbers of atomic units per transcription.

The trends observed in the information content explain the Precision-Recall trade-off. Transcriptions with a higher number of atomic units, such as those from the zero-shot models, are more likely to cover the atomic units present in the ground truth, thereby enhancing their Recall scores. Conversely, this verbosity increases the likelihood of including redundant information, which negatively impacts their Precision. In contrast, the fine-tuned Llama-3.1 model's output quantity (1.39 units) not only aligns closely with the ground truth (1.35) but also achieves the highest Precision. This suggests that fine-tuning helped the model accurately identify the speaker's emotional state.

Moreover, the ETC task remains challenging, as evidenced by the overall low scores across all evaluation metrics. Even the best-performing model, the fine-tuned Llama-3.1, achieves only 14.29% in F1-score, indicating significant room for improvement in accurately capturing speakers' emo-

tional states in dialogues. Additionally, lower F1 scores than both Precision and Recall for all models suggest a notable imbalance between these two aspects, highlighting the difficulty of generating emotion transcriptions that are both comprehensive and precise. Bridging this gap positions our dataset as a valuable and challenging benchmark for the community, advancing future research on understanding emotional states.

5. Case Study

To illustrate the challenges faced by current ETC models, we present a case study in Table 4. In this example, Speaker A's final utterance explicitly mentions feelings of relief and anger related to a past event. The ground-truth transcription, however, reveals that the speaker's actual emotional state at that moment was happiness, stemming from the listener's empathetic response. This gap between the narrated emotions and the speaker's actual emotional state poses a significant challenge for the ETC task.

Faced with this challenge, most models failed to capture the speaker's true emotional state. As shown in Table 4, the outputs from GPT-4.1 and the zero-shot and 4-shot Llama-3.1 models focused only on the negative emotions explicitly

mentioned in the narrative. For example, the generation from GPT-4.1 in the zero-shot setting, which is the most lengthy, not only identified the emotions present in the utterance but also elaborated on them to express a desire for the listener’s understanding. However, it still failed to predict the emergent feeling of happiness. In contrast, the Llama-3.1 model, fine-tuned on our dataset, successfully inferred the speaker’s happiness derived from the empathetic interaction with the listener, suggesting that our dataset provides valuable training signals for capturing nuanced emotional states in dialogues.

6. Corpus Analysis

The annotated emotion labels enable a structured analysis of our dataset, providing insights into the characteristics of the emotion transcriptions. In this section, we present a quantitative analysis of the annotated emotion labels, focusing on three aspects: the distribution and agreement of the labels (Section 6.1), the co-occurrence of emotion labels within utterances (Section 6.2), and the textual differences between utterances and emotion transcriptions (Section 6.3).

6.1. Label Distribution and Agreement

We first analyze the distribution of emotion labels assigned to the emotion transcriptions, a process detailed in Section 3.3. Table 5 presents their distribution, showing the prevalence of each label for Speakers, Listeners, and all participants. In this analysis, prevalence is calculated as the proportion of transcriptions assigned a given label through majority voting. The table also reports the overall inter-annotator agreement, measured by Fleiss’ kappa (Fleiss, 1971).

Overall, approximately 50% of the emotion transcriptions were labeled as “Neutral,” indicating that nearly half of the transcriptions expressed one or more emotions other than neutral. Among the emotion labels, “Joy” was the most frequently expressed. When analyzed by role, Speakers’ transcriptions have a higher proportion of emotion labels assigned compared to those of Listeners. This difference is likely to be attributed to the task instructions for each role. Speakers were explicitly instructed to talk about the topic related to a specific emotion, a prompt that naturally encouraged more emotional expression and recall, compared to the Listeners. Interestingly, “Surprise” was the only emotion label expressed more frequently in the Listener’s transcriptions than in the Speaker’s (4.7% vs. 2.4%). Surprise is an emotion often experienced as a reaction to an event or new information. For example, our dataset contains several

Emotion	Spk. (%)	Lsn. (%)	All (%)	κ
Joy	29.9	20.0	25.0	0.603
Sadness	10.8	8.8	9.8	0.359
Fear	7.1	5.3	6.2	0.476
Anger	3.7	2.8	3.2	0.529
Surprise	2.4	4.7	3.5	0.560
Disgust	6.3	3.3	4.8	0.233
Neutral	43.4	57.2	50.3	0.400
Overall	–	–	–	0.533

Table 5: Distribution of emotion labels in emotion transcriptions by dialogue role and inter-annotator agreement (Fleiss’ kappa). Columns denote dialogue role: speaker (Spk.), listener (Lsn.), and both roles (All).




Figure 2: Co-occurrence matrix of emotion labels. Each cell indicates the number of transcriptions annotated with both corresponding labels.

Listener transcriptions expressing surprise in response to the Speaker’s narrative, such as “*It was somewhat unexpected, so I tried to imagine the situation myself.*” and “*I was surprised that something like this could really happen.*”, thereby illustrating this trend.

The overall inter-annotator agreement of Fleiss’ kappa is 0.533, indicating moderate agreement. Examining emotion labels individually, “joy” showed the highest agreement among annotators, followed by “anger” ($\kappa = 0.529$) and “surprise.” ($\kappa = 0.560$) Conversely, “disgust” exhibited the lowest agreement ($\kappa = 0.233$).

6.2. Co-occurrence of Emotion Labels

In this section, we analyze the co-occurrence of emotion labels in the emotion transcriptions to understand the complexity of emotional expression. Figure 2 illustrates the co-occurrence relationships of emotion labels within the transcriptions. We found that approximately 5.6% of all non-neutral transcriptions contained multiple emotion labels.

A closer look at the co-occurrence matrix in Figure 2 reveals several patterns. First, emotions with similar valence frequently co-occur, such as “Disgust” with “Anger” and “Sadness.” Second, interestingly, emotions with opposing natures, like “Joy”

Emotion	Utterance	Emotion transcription
Joy	くださる (give), 楽しみだ (excited), 感動 (moved), 食べる (eat), 楽しむ (enjoy)	うれしい (happy), ワクワク (excited), 前向きだ (positive), 感動 (moved), 楽しみだ (excited)
Sadness	悲しい (sad), 寂しい (lonely), 辛い (painful), がっかり (disappointed), もう (enough)	悲しい (sad), 残念だ (regretful), 寂しい (lonely), 辛い (painful), 同情 (sympathy)
Fear	ニュース (news), 恐ろしい (terrifying), 事件 (incident), 心配だ (worried), 地震 (earthquake)	恐怖 (fear), 心配だ (worried), 恐ろしい (terrifying), 想像 (imagine), 心配 (worry)
Anger	イライラ (irritated), 怒る (angry), 立つ (stand), 悪い (bad), 怒り (anger)	怒り (anger), イライラ (irritation), 強い (strong), いらいら (irritated), 怒る (angry)
Surprise	なんて (what), びっくり (surprised), くらい (about), 凄い (amazing), 時代 (times)	驚き (surprise), 意外だ (unexpected), 感心 (impressed), びっくり (surprised), 想像 (imagine)
Disgust	嫌だ (dislike), 嫌悪 (disgust), 悪い (bad), 怒る (angry), 苦手だ (averse)	嫌だ (dislike), 嫌悪 (disgust), 想像 (imagine), 抱く (harbor), 申し訳ない (sorry)
Neutral	あまり (little), くださる (give), 信頼 (trust), 食べる (eat), ところ (point)	書く (write), 同意 (agree), 答える (answer), 具体 (specific), 説明 (explain)

Table 6: Top-5 frequent words by category for utterances and emotion transcriptions.

and “Sadness,” also co-occur frequently. This co-occurrence is often not contradictory, as the two emotions can stem from different aspects of the situation. For example, transcriptions from the dataset included statements such as, “I am recalling both pleasant and bitter memories.” and “listener is empathizing with me, but I also still feel sad when I recall the shopping scene.” Furthermore, “Disgust” tends to co-occur broadly with a wide range of other emotions. This broad association may contribute to its ambiguity, potentially explaining why it received the lowest inter-annotator agreement. These results suggest the complexity of human emotion, which often involves multiple, sometimes conflicting feelings simultaneously.

6.3. Frequent Words for Emotion Categories

In this section, to analyze the differences in linguistic features between utterances and emotion transcriptions, we extract and compare their respective characteristic words for each emotion label. This analysis is based on the method conducted by Ide and Kawahara (2022). Specifically, we first use Juman++ (Tolmachev et al., 2018) to identify words appearing in an utterance or emotion transcription. Next, to filter out common words that appear across all emotion labels, we remove any word with an IDF value less than half of the maximum IDF. Finally, we extract the highest-frequency words as characteristic words.

Table 6 shows the Top-5 most frequent words for each emotion label. A common trend in both utterances and emotion transcriptions is the extraction of words that express the corresponding emotion category. For example, for joy, words like “excited” are frequent, and for sadness, “sad” is the most

frequent.

A key difference emerged between the two text types. In utterances, the characteristic words tended to include not only emotion words but also words describing events and actions that trigger emotions. For instance, for joy, the word “eat” is extracted, while for fear, words like “incident” and “earthquake” are prominent. In contrast, the words that are frequent in emotion transcriptions tend to express the subtle nuances of the emotion. For example, in the category of joy, words like “happy” and “excited” appeared, while for fear, “worried” and “terrifying” are extracted. This shows a linguistic difference in expression even within the same emotion category.

For the neutral category has a trend of general words, rather than emotion words. Notably, for the emotion transcriptions, the words often used to describe an utterance’s intent were extracted as characteristic words, “agree,” “answer,” and “explain.”

7. Conclusion and Future Work

We introduced Emotion Transcription in Conversation (ETC), a novel task for generating verbalizations of a speaker’s emotional states, capturing fine-grained emotional nuances. To support this task, we constructed a new dataset of 1k Japanese dialogues annotated with self-reported emotion transcriptions and benchmarked baseline models. Our experimental results indicate the utility of our dataset for the ETC task, while also highlighting the challenges. A key challenge is the potential gap between explicitly expressed emotions and the actual emotional state; even fine-tuned models failed to capture these nuances (see Appendix F for failure cases). This suggests the need

for models with enhanced capabilities to infer nuanced emotional states.

Future work can address this challenge from several directions. One is to explore sophisticated prompting (e.g., chain-of-thought) and advanced fine-tuning objectives, such as RLHF or contrastive learning, to distinguish subtle expressions. Furthermore, speaker modeling could enhance model capabilities. Building on prior ERC research, this approach could involve incorporating speaker embeddings (Hu et al., 2021) or leveraging personality traits we have collected, which previous studies have linked to emotional states in dialogue (Wang et al., 2024; Shinoda et al., 2025).

Additionally, designing human evaluation for the ETC task is a critical future direction. Assessing emotion transcriptions requires capturing subtle nuances in emotional expressions. While LLM-based evaluators have shown promise in this regard, their reliability and potential biases warrant careful examination. Comparing LLM-based and human evaluations would provide valuable insights into the validity of automatic evaluation for this task.

Limitations

This study has the following limitations and challenges: First, the ETC dataset was constructed using crowdsourcing with Japanese-speaking participants, which may introduce cultural and linguistic biases. To confirm the generalizability of this dataset, similar data collection efforts across diverse languages and cultural contexts are required. Second, our dataset is limited to the text modality, which may lack nonverbal cues essential for understanding humans' emotional states. Therefore, constructing multimodal datasets that incorporate audio and visual information is a promising direction for enriching emotion understanding in real-world conversations. Third, the collected dialogues are relatively brief, comprising only 5 turns (10 utterances). This short context might be insufficient to adequately capture longer conversational dynamics and the complexity of emotional states evolving over time. Additionally, our dataset comprises 1k dialogues, and it remains unclear whether this scale is sufficient for training robust ETC models. However, the crowdsourcing-based collection protocol employed in this study provides a feasible and promising approach for scaling up ETC datasets in the future. Lastly, additional consideration is necessary regarding potential use cases of emotional transcriptions (ET). Future research should explore how recognized ET can be effectively integrated into conversational systems to inform interactions and enhance emotional intelligence.

Ethical Considerations

This study involves the use of data that captures internal emotional states of individuals expressed through natural language, which necessitates careful ethical consideration regarding the sensitivity of emotional content and the potential risks associated with AI applications. First, we collected the dialogue data and corresponding emotion transcriptions for each utterance through a Japanese crowdsourcing platform. Prior to the data collection, the participants were provided with clear explanations of the study's purpose, the nature of the data being collected. The participants were asked to consent to the data use before they engaged in the dialogue. The collected data contains no personally identifiable information and was stored in an anonymized format. Also, the proposed Emotion Transcription in Conversation (ETC) task aims to generate natural language descriptions of a speaker's internal emotional state based on conversational context. This technology holds potential for applications in dialogue systems, educational support, and counseling. However, if misused, it may raise serious ethical concerns, such as excessive inference of emotions, reinforcement of cultural biases, or surveillance and manipulation of emotional states. To address these risks, we evaluated the alignment between model-generated outputs and human annotations, and highlighted the current limitations and areas for improvement. Nevertheless, continued attention must be paid to the ethical implications of how such trained models are deployed and used in real-world contexts.

Acknowledgments

This work was supported by JSPS KAKENHI Grant Number 25H01382.

8. Bibliographical References

- Huang-Cheng Chou, Wei-Cheng Lin, Chi-Chun Lee, and Carlos Busso. 2022. [Exploiting annotators' typed description of emotion perception to maximize utilization of ratings for speech emotion recognition](#). In [ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing \(ICASSP\)](#), pages 7717–7721.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In [North American Chapter of](#)

- the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT), pages 4171–4186.
- P. Ekman, W. V. Friesen, M. J. O’ Sullivan, A. K. Chan, I. Diacoyanni-Tarlatzis, K. G. Heider, R. Krause, W. A. LeCompte, T. K. Pitcairn, P. E. Ricci-Bitti, K. R. Scherer, M. Tomita, and A. Tzavaras. 1987. Universals and cultural differences in the judgments of facial expressions of emotion. volume 53, pages 712–717.
- Shutong Feng, Guangzhi Sun, Nurul Lubis, Wen Wu, Chao Zhang, and Milica Gasić. 2024. Affect recognition in conversations using large language models. In Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL), pages 259–273.
- J. L. Fleiss. 1971. Measuring nominal scale agreement among many raters. Psychological Bulletin, 76:378–382.
- Yumeng Fu, Junjie Wu, Zhongjie Wang, Meishan Zhang, Lili Shan, Yulin Wu, and Bingquan Liu. 2025. LaERC-S: Improving LLM-based emotion recognition in conversation with speaker characteristics. In International Conference on Computational Linguistics (COLING), pages 6748–6761.
- Kazuki Fujii, Taishi Nakamura, Mengsay Loem, Hiroki Iida, Masanari Ohi, Kakeru Hattori, Hirai Shota, Sakae Mizuki, Rio Yokota, and Naoaki Okazaki. 2024. Continual pre-training for cross-lingual LLM adaptation: Enhancing Japanese language capabilities. In First Conference on Language Modeling.
- Deepanway Ghosal, Navonil Majumder, Soujanya Poria, Niyati Chhaya, and Erik Cambria. 2019. DialogueGCN: A graph convolutional neural network for emotion recognition in conversation. 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), pages 154–164, Hong Kong, China. Association for Computational Linguistics.
- Jingwen Hu, Yuchen Liu, Jinming Zhao, and Qin Jin. 2021. MMGCN: Multimodal fusion via deep graph convolution network for emotion recognition in conversation. In Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers), pages 5666–5675, Online. Association for Computational Linguistics.
- Chin-Yew Lin. 2004. ROUGE: A package for automatic evaluation of summaries. In Text Summarization Branches Out, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruochen Xu, and Chenguang Zhu. 2023. G-eval: NLG evaluation using gpt-4 with better human alignment. In Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, pages 2511–2522, Singapore. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. RoBERTa: A robustly optimized BERT pretraining approach. arXiv preprint arXiv:1907.11692.
- Sewon Min, Kalpesh Krishna, Xinxi Lyu, Mike Lewis, Wen-tau Yih, Pang Koh, Mohit Iyyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2023. FACTScore: Fine-grained atomic evaluation of factual precision in long form text generation. In Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, pages 12076–12100, Singapore. Association for Computational Linguistics.
- Atsushi Oshio, ABE Shingo, and Pino Cutrone. 2012. Development, reliability, and validity of the Japanese version of ten item personality inventory (tipi-j). Japanese Journal of Personality, 21(1).
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In Proceedings of the 40th Annual Meeting on Association for Computational Linguistics, ACL ’02, page 311–318, USA. Association for Computational Linguistics.
- Patrícia Pereira, Helena Moniz, and Joao Paulo Carvalho. 2025. Deep emotion recognition in textual conversations: A survey. Artificial Intelligence Review, 58(1):1–37.
- Soujanya Poria, Navonil Majumder, Rada Mihalcea, and Eduard Hovy. 2019. Emotion recognition in conversation: Research challenges, datasets, and recent advances. IEEE access, 7:100943–100953.
- Xincheng Qin, Zhen Wu, Teng Zhang, Yuxuan Li, Jiaying Luan, Bo Wang, Lidong Wang, and Jian Cui. 2023. BERT-ERC: Fine-tuning BERT is enough for emotion recognition in conversation.

In AAAI Conference on Artificial Intelligence (AAAI), volume 37, pages 13492–13500.

Arseny Tolmachev, Daisuke Kawahara, and Sadao Kurohashi. 2018. [Juman++: A morphological analysis toolkit for scriptio continua](#). In Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations, pages 54–59, Brussels, Belgium. Association for Computational Linguistics.

Yan Wang, Bo Wang, Yachao Zhao, Dongming Zhao, Xiaojia Jin, Jijun Zhang, Ruifang He, and Yuexian Hou. 2024. [Emotion recognition in conversation via dynamic personality](#). In Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024), pages 5711–5722, Torino, Italia. ELRA and ICCL.

Jintao Wen, Geng Tu, Rui Li, Dazhi Jiang, and Wenhua Zhu. 2023. Learning more from mixed emotions: a label refinement method for emotion recognition in conversations. Transactions of the Association for Computational Linguistics, 11:1485–1499.

Fangxu Yu, Junjie Guo, Zhen Wu, and Xinyu Dai. 2024. Emotion-anchored contrastive learning framework for emotion recognition in conversation. In North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT), pages 4521–4534.

Xupeng Zha, Huan Zhao, and Zixing Zhang. 2024. ESIHGNN: Event-state interactions infused heterogeneous graph neural network for conversational emotion recognition. In International Conference on Acoustics, Speech and Signal Processing (ICASSP), pages 11136–11140.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. [Bertscore: Evaluating text generation with bert](#). In International Conference on Learning Representations.

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. [Judging LLM-as-a-judge with MT-bench and chatbot arena](#). In Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track.

9. Language Resource References

C. Busso, R. Lotfian, K. Sridhar, A.N. Salman, W.-C. Lin, L. Goncalves, S. Parthasarathy, A. Reddy Naini, S.-G. Leem, L. Martinez-Lucas, H.-C. Chou, and P. Mote. 2025. [The MSP-Podcast corpus](#). ArXiv e-prints (arXiv:2509.09791), pages 1–20.

Carlos Busso, Murtaza Bulut, Chi-Chun Lee, Abe Kazemzadeh, Emily Mower, Samuel Kim, Jeannette N. Chang, Sungbok Lee, and Shrikanth S. Narayanan. 2008. [IEMO-CAP: Interactive emotional dyadic motion capture database](#). Language Resources and Evaluation, 42(4):335–359.

Ankush Chatterjee, Kedhar Nath Narahari, Meghana Joshi, and Puneet Agrawal. 2019. Semeval-2019 task 3: Emocontext contextual emotion detection in text. In Proceedings of the 13th international workshop on semantic evaluation, pages 39–48.

Chao-Chun Hsu, Sheng-Yeh Chen, Chuan-Chun Kuo, Ting-Hao Huang, and Lun-Wei Ku. 2018. EmotionLines: An emotion corpus of multi-party conversations. In International Conference on Language Resources and Evaluation (LREC).

Tatsuya Ide and Daisuke Kawahara. 2022. [Building a dialogue corpus annotated with expressed and experienced emotions](#). In Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics: Student Research Workshop, pages 21–30, Dublin, Ireland. Association for Computational Linguistics.

Yanran Li, Hui Su, Xiaoyu Shen, Wenjie Li, Ziqiang Cao, and Shuzi Niu. 2017. [Dailydialog: A manually labelled multi-turn dialogue dataset](#). In International Joint Conference on Natural Language Processing (IJCNLP), pages 986–995, Taipei, Taiwan. Asian Federation of Natural Language Processing.

June M. Liu, He Cao, Renliang Sun, Rui Wang, Yu Li, and Jiaying Zhang. 2025. CAPE: A Chinese dataset for appraisal-based emotional generation in large language models. In North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT): Findings, pages 6291–6309.

L. Martinez-Lucas, Mohammed Abdelwahab, and Carlos Busso. 2020. [The MSP-conversation corpus](#). In Interspeech 2020, pages 1823–1827, Shanghai, China.

Naoaki Okazaki, Kakeru Hattori, Hirai Shota, Hiroki Iida, Masanari Ohi, Kazuki Fujii, Taishi Nakamura, Mengsay Loem, Rio Yokota, and Sakae Mizuki. 2024. [Building a large Japanese web corpus for large language models](#). In First Conference on Language Modeling.

Soujanya Poria, Devamanyu Hazarika, Navonil Majumder, Rishi Naik, Erik Cambria, and Alexander Hoffmann. 2019. [MELD: A multi-modal multi-party dataset for emotion recognition in conversations](#). In Annual Meeting of the Association for Computational Linguistics (ACL), pages 527–536, Florence, Italy. Association for Computational Linguistics.

Hannah Rashkin, Eric Michael Smith, Margaret Li, and Y-Lan Boureau. 2019. [Towards empathetic open-domain conversation models: A new benchmark and dataset](#). In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, pages 5370–5381, Florence, Italy. Association for Computational Linguistics.

Kazutoshi Shinoda, Nobukatsu Hojo, Kyosuke Nishida, Saki Mizuno, Keita Suzuki, Ryo Masumura, Hiroaki Sugiyama, and Kuniko Saito. 2025. [ToMATO: Verbalizing the mental states of role-playing LLMs for benchmarking theory of mind](#). AAAI Conference on Artificial Intelligence, 39(2):1520–1528.

Seyyed M. Zahiri and Jinho D. Choi. 2018. [Emotion detection on tv show transcripts with sequence-based convolutional neural networks](#). In AAAI Workshops.

A. Data Collection Details

We conducted the collection of dialogue and corresponding emotion transcriptions using a custom-built chat interface hosted on Google Sheets. The interface consists of multiple dialogue slots (as shown in Figure 3), each associated with a predefined emotion label.

Each dialogue slot in the interface is composed of several key areas. Figure 3-A shows metadata for the dialogue, including a predefined emotion label and a unique dialogue slot ID. This area also includes two cells for entering CrowdWorks usernames—one for the Speaker and one for the Listener. In these cells, participants were instructed to enter their username to join the conversation. They then wait to be paired with a dialogue partner or start the conversation if someone has already joined. Figure 3-B allows participants to exchange coordination messages during the task. This area is designed for task-related communication that does not directly affect the dialogue itself and facilitates the accurate and smooth progression of tasks. For instance, participants used it to confirm matching by entering “Nice to meet you”, inform dialogue partners of their readiness with “I’m ready”, or notify other workers of a brief absence while waiting for matching with “I’ll be back by 4:00 PM”.

Figure 3-C is the main body of the interaction, where participants input their utterances across five turns each. After a participant (e.g., the Speaker) enters their utterance into a cell in this area, the other participant (e.g., the Listener) responds by entering their utterance in the next cell. Subsequently, while the Listener is typing their response, the Speaker writes an emotion transcription for their previous utterance by inserting a note into the same cell. These steps are repeated alternately for five turns. During the dialogue session, we required participants to write at least five characters in length for both the utterance and the corresponding emotion transcription.

Figure 3-D consists of checkboxes for tracking task progress. The left column is for utterance input, and the right column is for emotion transcription. At each turn, participants check the corresponding box after completing each action. The dialogue session is considered complete only when all checkboxes have been marked by both participants.

B. Distribution of Participants’ Personality Traits

Figure 4 illustrates the distribution of Big Five personality traits among the 199 participants, measured using the Japanese version of the Ten-Item

Personality Inventory (TIPI-J) (Oshio et al., 2012). Scores for each dimension range from 2 to 14, with higher values indicating stronger tendencies.

C. Prompt Template for ETC Task

The prompt template for the ETC task, translated from Japanese, is shown in Figure 5. The model is expected to generate only the emotion transcription, in accordance with the definition. [TARGET_SPEAKER] is a placeholder for the target speaker whose emotion transcription is to be generated, and [DIALOGUE] is a placeholder for the dialogue history.

D. Implementation Details

We executed supervised fine-tuning of Llama-3.1-Swallow on two NVIDIA A100 80GB GPUs, using a batch size of 8 and training for 2 epochs. We applied 4-bit quantization to enhance memory efficiency and used the PagedAdamW8bit optimizer during training. Validation was conducted every 200 steps. For hyperparameter tuning, we employed a grid search with learning rates tested in {1e-05, 5e-05} and warm-up steps in {300, 700}. As a result, we selected a learning rate of 1e-05 and 300 warm-up steps.

During the inference, to ensure reproducibility, we set `do_sample=False` for Llama-3.1-Swallow and `temperature` to 0.0 for GPT-4.1.

E. Prompt Template for Fine-grained Evaluation

The prompt templates for the fine-grained evaluation of content faithfulness, translated from Japanese, are shown in Figure 6 and Figure 7. The decomposition prompt (Figure 6) instructs the model to output a list of atomic units of information contained in the emotion transcription. [TRANSCRIPTION] is a placeholder for the emotion transcription to be decomposed. The entailment prompt (Figure 7) asks the model to determine whether each unit of information is entailed by the target transcription. [SOURCE_TEXT] and [TARGET_TEXT] are placeholders for the source and target emotion transcriptions, respectively.

F. Failure Case Analysis for the ETC Task

Table 7 presents a cherry-picked example of where the models failed to predict the speaker’s emotion transcription. In this case, the ground-truth transcription expresses the speaker’s delight at the partner’s willingness to delve into de-

A	枠ID: 2505080341		記入時刻	
	状態	未着手		
	感情	怒る		
	スピーカーID			
	リスナーID			
B	スピーカー準備	未着手		
	連絡欄			
	スピーカー連絡欄			
C	リスナー連絡欄			
	対話	メッセージ		
	スピーカー発話1		<input type="checkbox"/>	<input type="checkbox"/>
	リスナー発話1		<input type="checkbox"/>	<input type="checkbox"/>
	スピーカー発話2		<input type="checkbox"/>	<input type="checkbox"/>
	リスナー発話2		<input type="checkbox"/>	<input type="checkbox"/>
	スピーカー発話3		<input type="checkbox"/>	<input type="checkbox"/>
	リスナー発話3		<input type="checkbox"/>	<input type="checkbox"/>
	スピーカー発話4		<input type="checkbox"/>	<input type="checkbox"/>
	リスナー発話4		<input type="checkbox"/>	<input type="checkbox"/>
	スピーカー発話5		<input type="checkbox"/>	<input type="checkbox"/>
リスナー発話5		<input type="checkbox"/>	<input type="checkbox"/>	

D

メッセージ入力	心情入力
<input type="checkbox"/>	<input type="checkbox"/>
<input type="checkbox"/>	<input type="checkbox"/>
<input type="checkbox"/>	<input type="checkbox"/>
<input type="checkbox"/>	<input type="checkbox"/>
<input type="checkbox"/>	<input type="checkbox"/>
<input type="checkbox"/>	<input type="checkbox"/>
<input type="checkbox"/>	<input type="checkbox"/>
<input type="checkbox"/>	<input type="checkbox"/>
<input type="checkbox"/>	<input type="checkbox"/>
<input type="checkbox"/>	<input type="checkbox"/>

▲ 本ページの作業枠の削除・移動・レイアウト変更など、構造を壊す行為は禁止します
 ▲ 心情を記入の際の注意事項：
 ・ 心情は必ず「**ですます調**」でご記入ください。
 ・ メッセージをご記入いただいたセル上で、**右クリック**→「**メモを挿入**」よりご記入ください。
 ・ **5文字以上**でご記入ください。

Figure 3: The dialogue slot used for conducting conversations and entering emotion transcriptions.

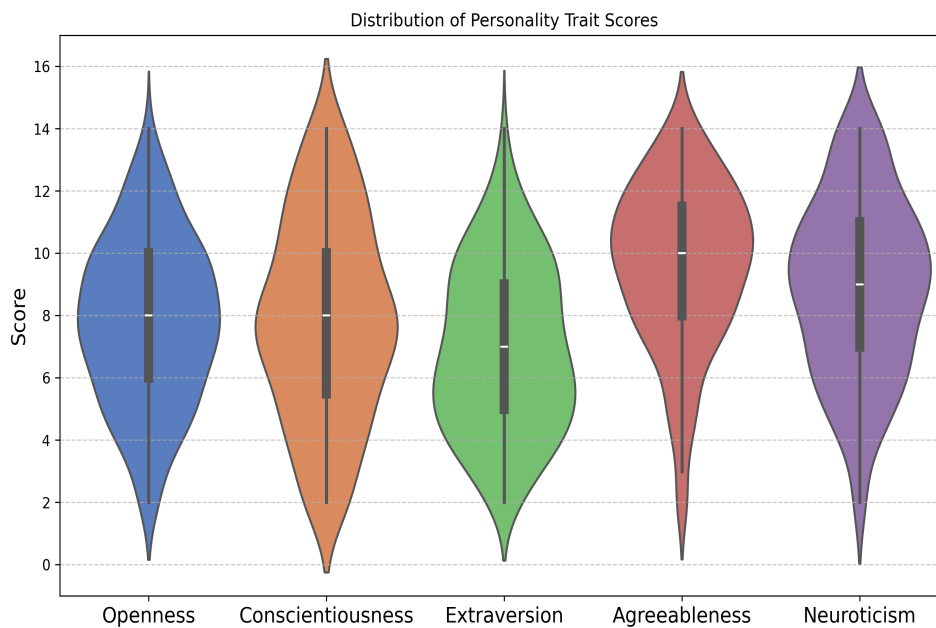


Figure 4: Distribution of Big Five personality traits among the 199 participants.

tails. However, every model instead focuses on the speaker's shock, illustrating a tendency to prioritize immediately preceding negative cues over more nuanced, interaction-driven emotions. This suggests that ETC models struggle to incorporate the partner's supportive attitude when inferring a speaker's emotional state.

Task Definition

You will be given a multi-turn dialogue between two speakers, Speaker A and Speaker B.

Your task is to infer the emotion transcription of the speaker who spoke the last utterance (either Speaker A or Speaker B) at the time it was uttered.

Definition of Emotion Transcription

Emotion transcription is defined as a verbalization of the internal emotional state or intention that a participant in the dialogue held at the time of their utterance. It is a linguistic expression of the psychological background behind the utterance, including “what they were thinking,” “why they chose those words,” or “how they were feeling” when they spoke.

Output Rules

- 1. Predict the emotion transcription corresponding to the last utterance of the specified target speaker.**
- 2. Generate the emotion transcription in a tone as if the speaker themselves were describing it.**
- 3. Do not enclose the emotion transcription in quotation marks or brackets.**
- 4. Do not include any extra explanations after the emotion transcription, nor include the emotion transcriptions of other speakers.**
- 5. When referring to the dialogue partner, use the term “my dialogue partner”; when referring to yourself, use “I.”**

Target speaker for emotion transcription prediction: [TARGET_SPEAKER]

Dialogue:

[DIALOGUE]

[TARGET_SPEAKER]’s emotion transcription:

Figure 5: The prompt template for the ETC task (Translated from Japanese). The model is expected to generate only the emotion transcription.

Your task is to decompose a given emotion transcription into semantically independent minimal emotional units.

Definition of Emotion Transcription

Emotion transcription is defined as a verbalization of the internal emotional state or intention that a participant in the dialogue held at the time of their utterance. It is a linguistic expression of the psychological background behind the utterance, including “what they were thinking,” “why they chose those words,” or “how they were feeling” when they spoke.

Guidelines

- Extract only internal states aligned with the definition above: emotions (e.g., “felt sad,” “empathized,” “was grateful”) and intentions (e.g., “tried to ...,” “wanted to ...,” “felt that ...”).
- Do not include descriptions of objective facts or actions (e.g., “likes/dislikes ...,” “performed ...,” “is explaining ...”).
- However, if the transcription contains a dialogue participant's “feelings,” “thoughts,” or “recollections” regarding those facts or actions, those descriptions must be extracted.
- Ensure that each atomic unit is semantically complete on its own. In particular, the use of demonstratives such as “that,” “this,” or “it” is prohibited.
- If no emotional element is present, output “[No emotional elements].”

(8 examples)

**Decompose the following emotion transcription into semantically independent minimal emotional units:
[TRANSCRIPTION]**

Figure 6: The prompt template for the decomposition step in the fine-grained evaluation (Translated from Japanese)

You are given two emotion transcriptions: a “premise” and a “hypothesis.” Your task is to determine whether the hypothesis is supported by the premise by **comparing only their surface-level content**.

Mandatory Rules

- **Inferring implicit causal relationships or value judgments is prohibited.**
- Verify whether the **subject (whose emotion or intention) and the object** of the emotional states match between the premise and the hypothesis.
- Any emotional states not **explicitly** stated in the premise must not be considered as implied.
- Ignore differences in tense in your judgment.

Definition of Emotion Transcription

Emotion transcription is defined as a verbalization of the internal emotional state or intention that a participant in the dialogue held at the time of their utterance. It is a linguistic expression of the psychological background behind the utterance, including “what they were thinking,” “why they chose those words,” or “how they were feeling” when they spoke.

Guidelines

Supported

- The hypothesis can be **definitively determined to be entailed** solely from the content of the premise.

Not Supported

- The emotional element in the hypothesis is **not present in the premise and cannot be definitively determined**, or it **clearly contradicts** the premise.

Neutral

- Only the **core emotional element matches**, while the **validity of secondary elements is uncertain**.

Output Format

Return only a pure JSON object strictly following the format below.

```
```json
{
 "reasoning": "<Briefly describe the rationale for your judgment here>",
 "answer": "<One of: Supported / Neutral / Not Supported>"
}
```
```

(9 examples)

Example 10

Premise: [SOURCE_TEXT]

Hypothesis: [TARGET_TEXT]

Figure 7: The prompt template for the entailment step in the fine-grained evaluation (Translated from Japanese).

| | |
|----------------|---|
| Context | <p>Speaker A: The other day, my favorite baseball player suffered a serious injury and will be out for three months. (先日、自分が特に応援している野球選手が全治3ヵ月の大怪我を負ってしまいました。)</p> <p>Speaker B: Three months? Are you talking about Okamoto? (全治3ヶ月というと、それはもしかして岡本選手ですか?)</p> <p>Speaker A: Yes. That news was really shocking. (そうです。あの出来事は本当にショックでした。)</p> <p>Speaker B: I thought so. He got hurt in that collision with Nakano from the Hanshin Tigers, and I know how devastating that is for a Giants fan. (やっぱりそうですか。阪神の中野選手との交錯でケガをしたと思いますが巨人ファンの方からすると彼がいなくなるのは大ダメージですね。)</p> <p>Speaker A: Right. At first they reported it as a one-month muscle injury, but after a detailed examination it turned out to be a ligament tear requiring three months of recovery, which made it even more shocking. (はい。最初は筋肉の損傷で全治1ヵ月と報道されていたのにその後、精密検査で靭帯損傷で全治3ヵ月となったので、余計にショックでしたね。)</p> |
| Output | <p>Ground-Truth: I feel very happy that my dialogue partner could dive so deeply and discuss the details with me. (相手がかかり話を掘り下げた詳しいことまで話せることがとても嬉しいと感じています。)</p> <p>GPT-4.1 (zero-shot): I feel even more disappointed and shocked because the injury was worse than initially reported, and I'm explaining the details so they understand how bad it really is. (自分が応援している選手のケガが思ったよりも重かったことに対して、最初の報道よりも悪い結果になってしまい、残念な気持ちやショックがさらに大きくなったと感じている。相手にその気持ちを理解してもらいたくて、詳しく説明している。)</p> <p>GPT-4.1 (4-shot): I wanted my partner to understand how shocked I was. (相手に自分のショックな気持ちを理解してもらいたかったです。)</p> <p>Llama-3.1 (zero-shot): That incident was really shocking. (あの出来事は本当にショックでした。)</p> <p>Llama-3.1 (4-shot): I feel relieved that my partner understands my feelings and I'm seeking empathy. (相手が自分の気持ちを理解してくれていることに安心し、共感を求めています。)</p> <p>Llama-3.1 (fine-tuned): I answered their question, recalled what happened, and told them how much more shocked I felt. (相手の質問に答えつつ、当時の状況を思い出して、さらにショックを受けたことを伝えました。)</p> |

Table 7: Failure case study of emotion transcription generation (translated from Japanese). The outputs are the emotion transcriptions for Speaker A's final utterance.