

SMILE-Next: Teaching Large Language Models to Detect, Classify, and Reason about Laughter

Lee Jung-Mok¹ Kim Sung-Bin² Joohyun Chang³ Lee Hyun² Tae-Hyun Oh³

¹Dept. of EE, KAIST, ²Dept. of EE, POSTECH,

³School of Computing, KAIST

Correspondence: taehyun.oh@kaist.ac.kr

Abstract

Laughter is a complex social signal that conveys communicative intent beyond amusement. While prior work has focused on isolated laughter analysis tasks, a comprehensive understanding of laughter in real-world scenarios remains underexplored. We introduce SMILE-Next, a dataset for real-world laughter understanding with multimodal textual representations and question–answer annotations across three tasks: laughter detection, laughter type classification, and laughter reasoning. Building upon SMILE-Next, we aim to develop a laughter-specialized large language model capable of nuanced understanding of laughter in real-world contexts. To this end, we propose two key components: laughter-specific Self-Instruct and the Mixture-of-Laugh-Experts (MoLE) framework. Laughter-specific Self-Instruct enhances generalization across tasks and domains by automatically synthesizing diverse laughter-centric instructions. MoLE introduces a task-adaptive expert routing mechanism that dynamically selects specialized experts tailored to each laughter-related task, improving task-specific performance and efficiency. Experimental results show that the combination of our proposed components substantially outperforms multimodal LLM baselines, advancing robust real-world laughter understanding. Project page is at: <https://mok0102.github.io/smile-next/>.

1 Introduction

Laughter is one of the most fundamental and universal forms of human expression in social interactions. Beyond a simple reaction, people often laugh not only because something is funny but also to build intimacy (Stauffer, 1999), capture attention (Wanzer et al., 2010), reinforce rapport, or ease tension (Scott et al., 2014). As such, laughter embodies nuances that spoken language alone often cannot convey, and we naturally interpret the intentions behind laughter to sustain social interaction.

Following recent advancements in artificial social intelligence (Bainbridge et al., 1994; Dautenhahn, 2007; Williams et al., 2022), an emerging line of work has explored building systems capable of understanding laughter or humor (Bertero and Fung, 2016; Hasan et al., 2019; Alnajjar et al., 2022; Hessel et al., 2023; Ko et al., 2023; Xie et al., 2024; Hyun et al., 2024). However, these approaches primarily focus on narrowly defined tasks, such as detecting humor (Bertero and Fung, 2016; Hasan et al., 2019), estimating humor intensity (Alnajjar et al., 2022), reasoning about humor in cartoons (Hessel et al., 2023) or in general textual and multimodal content (Ko et al., 2023; Xie et al., 2024), and analyzing laughter in TV shows (Hyun et al., 2024). As a result, they are limited in capturing laughter naturally occurring in real human–human interactions and in providing a broad, context-aware understanding of laughter, essential for enabling richer human–machine and machine–machine social interactions.

To address this challenge, we curate **SMILE-Next**, which aims to extend machine capabilities in generalized laughter reasoning within human interactions. SMILE-Next consists of videos paired with question–answer annotations designed for comprehensive laughter understanding. The dataset encompasses three tasks: (1) *Laughter detection*, determining whether a given video is intended to induce laughter; (2) *Laughter type classification*, categorizing the type of laughter; and (3) *Laughter reasoning*, providing free-form textual explanations for why laughter occurred. Given the diverse nature of laughter, our dataset covers a broad range of human interactions, including talk shows, sitcoms, and dyadic conversations, as illustrated in Fig. 1. Moreover, recognizing that curated data cannot cover all scenarios, we introduce a self-instructional data generation framework to overcome limitations in coverage and scale.

Building on this dataset, we develop a **laugh-**



Figure 1: **Task overview of SMILE-Next.** We present SMILE-Next, a comprehensive dataset for laughter understanding. The dataset includes three fundamental tasks for understanding laughter in social interactions: laughter detection, laughter type classification, and laughter reasoning.

ter expert Large Language Model (LLM), designed to identify and interpret laughter in social interaction. Because laughter arises from complex factors including subjectivity (Warren et al., 2021), contextual knowledge (Nijholt et al., 2017), and multimodal cues, such as facial expressions and prosody (Hasan et al., 2019), it is challenging for a model to interpret these entangled signals directly from video. To enable clearer interpretation by disentangling these factors, we convert multimodal signals (video and speech) into textual representations following Hyun et al. (2024). Specifically, we employ off-the-shelf models to translate video and speech attributes into textual descriptions, and then leverage the reasoning capabilities of LLMs to identify, understand, and explain the causes of laughter. Furthermore, we propose a LoRA-based Mixture-of-Laugh-Experts (MoLE) framework within LLM training, which learns task-specific experts to flexibly handle and synergize multiple laughter-related tasks, including detection, type classification, and reasoning, within a unified framework.

Our experimental results show that integrating LLM reasoning capabilities with multimodal textual representations enables effective identification and interpretation of laughter in human interactions, including the ability to generate plausible explanations. Moreover, incorporating self-instruction and MoLE further enhances the model’s reasoning capacity. Through our analyses, we demonstrate that textualizing multimodal information is essential for robust laughter understanding. Our main contributions are summarized as follows:

- We introduce **SMILE-Next**, a comprehensive dataset containing videos, textual representations, and QAs covering diverse laughter scenarios.
- We propose a **laughter-reasoning expert LLM** that demonstrates strong capabilities in laughter

detection, laughter type classification, and laughter reasoning.

- We show that **Laugh-tailored Self-Instruction** and **Mixture-of-Laugh-Experts (MoLE)** are effective strategies for enhancing LLMs’ capacity for laughter reasoning.

2 Related Work

Understanding laughter in videos. Laughter is a socially meaningful behavior closely tied to humor, sarcasm, and interpersonal bonding (Scott et al., 2014). Most approaches model laughter or humor as signals to be detected or quantified, focusing on classification or intensity prediction using textual and multimodal cues (Annamoradnejad and Zoghi, 2020; Weller and Seppi, 2020; Bertero and Fung, 2016; Alnajjar et al., 2022; Castro et al., 2019; Hasan et al., 2019; Ray et al., 2022). These methods do not explicitly model why laughter occurs. Recent studies have begun to explore explanation-oriented perspectives on humor and laughter (Chowdhery et al., 2022; Hessel et al., 2023; Ko et al., 2023; Xie et al., 2024), but they remain largely domain-specific and focus on constrained or scripted scenarios. While Hyun et al. (2024) introduced an initial benchmark for laughter reasoning in real-world videos, it remains limited in task scope and interaction diversity. Our work extends this line of research by formulating laughter understanding as a multi-faceted problem, encompassing laughter detection, laughter type classification, and laughter reasoning across diverse human interaction scenarios.

Textualized approaches for video understanding. Multimodal video-language understanding has traditionally relied on feature-based fusion, projecting visual, audio, and textual inputs into a shared latent space for downstream reasoning (Lu

et al., 2019; Li et al., 2019; Lei et al., 2018). While effective for recognition-centric tasks, recent analyses suggest that such implicit multimodal representations are less suitable for explanation-oriented reasoning, as causal, social, and contextual cues often remain entangled in latent feature spaces (Buch et al., 2022; Xie et al., 2024; Ko et al., 2023).

Recent studies instead explore textualization as an alternative interface for multimodal reasoning, converting visual and auditory signals into natural language descriptions that are directly consumable by large language models (Zeng et al., 2022; Wang et al., 2022b; Hasan et al., 2023; Richet et al., 2024; Hyun et al., 2024). By representing multimodal information in a symbolic and interpretable form, textualization enables LLMs to better leverage their pretrained commonsense, social, and causal reasoning capabilities, and has shown particular effectiveness in reasoning-heavy tasks involving affective states and interpersonal dynamics. Motivated by these observations, we adopt textualization as a key design choice for laughter understanding, where multiple implicit social and affective factors jointly give rise to laughter and require explicit reasoning.

Mixture-of-Experts in LLM. Mixture-of-Experts (MoE) architectures have been extensively studied as a scalable framework for conditional computation and parameter specialization, and have been widely adopted in large language models to support efficient scaling and diverse task handling (Shazeer et al., 2017; Lepikhin et al., 2021; Fedus et al., 2022; Du et al., 2022; Shen et al., 2024). Beyond scalability, MoE has also been explored in instruction-tuned and multi-task LLMs as a flexible mechanism for accommodating heterogeneous task requirements within a single model (Zadouri et al., 2024; Gou et al., 2023; Dou et al., 2023). In line with these advancements, we incorporate a parameter-efficient, LoRA-based MoE design into our laughter-understanding LLM to handle multiple laughter-related sub-tasks within a unified framework, including laughter detection, type classification, and reasoning.

3 SMILE-Next: Enhanced Multimodal Laughter Understanding Dataset

Our goal is to extend machine capabilities for understanding laughter in human interactions. Because laughter occurs in diverse forms across a wide range of social scenarios, analyzing and interpreting it cannot be reduced to a single, isolated

task. To this end, we introduce SMILE-Next, a dataset designed to support more general and comprehensive laughter understanding in diverse multimodal contexts. SMILE-Next covers a broad range of interaction scenarios and defines multiple complementary tasks that jointly evaluate and facilitate general laughter understanding. To fully understand laughter in a given situation, a system must determine whether laughter occurs, identify the type of laughter, and reason about its underlying cause. Accordingly, our dataset captures these steps through three tasks in diverse multimodal scenarios: laughter detection, laughter type classification, and laughter reasoning, as shown in Fig. 1.

3.1 Task Definition

We provide detailed descriptions of the three proposed tasks below.

Laughter detection. Laughter detection is a binary classification task that aims to determine whether a given video clip induces laughter. Prior work on humor or sarcasm detection, often limited to specific scenarios, such as TED talks or sitcoms, primarily relies on audio cues to identify laughter (Hasan et al., 2019; Castro et al., 2019). In contrast, our dataset encompasses a broader range of scenarios, including humorous, sarcastic, and awkward laughter across more diverse interaction settings, such as TV shows and dyadic conversations. As a result, effective laughter detection in SMILE-Next requires contextual reasoning over multimodal cues, including facial expressions, scene context, video captions, and acoustic signals, making the task substantially more challenging.

Laughter type classification. Laughter type classification is a multi-class classification task that identifies the type of laughter in a video. Inspired by the taxonomy proposed by Tanaka and Campbell (2014), we classify laughter into three high-level categories: *mirthful*, *polite*, and *satirical*. *Mirthful* laughter refers to spontaneous and joyful laughter, often triggered by humorous content or shared amusement. *Polite* laughter denotes socially motivated, non-spontaneous laughter, typically used to maintain conversational flow or express minimal agreement in formal or semi-formal interactions. *Satirical* laughter includes sarcastic, awkward, or mocking expressions, often conveying discomfort, disbelief, or social distance.

Laughter reasoning. Laughter reasoning requires the model to infer and explain the underlying cause

Task	Train	Val	Test	Total
Laughter Detection	1565	460	359	2384
Laughter Type Classification	1636	207	114	1957
Laughter Reasoning	1565	292	188	2045
Total	4766	959	661	6386

Table 1: **Data statistics of SMILE-Next.** We report the number of samples for laughter-related task, along with the corresponding train, validation, and test splits.

of laughter in a given video. Unlike the previous tasks, we formulate laughter reasoning as a free-form text generation problem, allowing models to produce natural language explanations for why a particular situation elicited laughter. This task builds upon the laughter reasoning task introduced in SMILE (Hyun et al., 2024), which focuses primarily on humor-driven laughter in TED talks and sitcoms. In SMILE-Next, we substantially broaden the scope of laughter reasoning to include a wider range of interaction contexts, such as interviews, movies, and dyadic conversations.

3.2 Task Formulation

We formulate all tasks as question-answering problems, where the model generates answers given a video clip and a task-specific question prompt. Since laughter arises from multiple multimodal cues that are inherently entangled in raw video, directly interpreting laughter from video alone is challenging. Following Hyun et al. (2024), we therefore convert videos into textualized representations of multimodal cues to facilitate disentanglement. Formally, the task formulation is defined as $\hat{y} = D(Q, \{t_1, t_2, \dots, t_k\})$, where D denotes the LLM, Q is the question prompt, \hat{y} is the generated answer, and t_1, \dots, t_k are multimodal textual representations extracted from the video v . This formulation enables the model to process disentangled multimodal information while fully leveraging the reasoning and comprehension capabilities of LLMs through textual representations.

3.3 Dataset Construction

We design a dataset construction pipeline to collect videos and annotations for laughter understanding tasks. The pipeline integrates both automated and human-in-the-loop processes to ensure scalability while maintaining high annotation quality. It consists of four stages: video collection, multimodal textual extraction, pseudo-label annotation, and human verification.

Laughter video collection. We collect source video clips containing laughter from diverse domains to capture a wide range of social contexts. Our sources include *TED Talks* and *sitcoms*, which primarily feature audience laughter elicited by humorous or satirical content. In addition, we collect in-the-wild videos from *YouTube*, including dyadic conversations, talk shows, and movies¹. These in-the-wild videos cover more intimate conversational settings and include laughter that reflects relaxed, affiliative, or socially bonding reactions beyond explicit humor. After collecting the source videos, we use an off-the-shelf laughter detection model (Gillick et al., 2021) to automatically identify video clips containing laughter. For each selected video, we further segment the clip into shorter units based on utterance boundaries, resulting in multiple laughter-related segments per source clip.

Textual multimodal cue extraction. Each video segment contains diverse multimodal cues, including visual cues (e.g., facial expressions and movements), acoustic cues (e.g., tone and pitch), and semantic cues (e.g., conversational content). To disentangle these signals, we convert them into structured textual representations. Specifically, we employ a set of expert models to extract utterances, acoustic features, facial action units, and video captions, all represented in text form. In addition, we extract relationship information between speakers. For in-the-wild videos, we infer relationships from YouTube metadata and video descriptions. For talk shows, we assign roles such as “host and guest,” while for *TED Talks* and *sitcoms*, we label interactions as “speaker and audience.” Detailed descriptions of the multimodal extraction process are provided in the Appendix A.2.

Pseudo-label annotation. Laughter reasoning is a free-form generation task and is challenging for human annotators to perform directly at scale. To address this, we generate pseudo-labels using expert large language models (e.g., GPT-4 (Achiam et al., 2023)). Given each video segment and its associated textualized multimodal representations, we prompt the expert large language model to reason about why laughter occurred in the scene, considering diverse range of laughters such as hilarious, polite, or awkward laughter. Then the pseudo-label is generated, and these explanations serve as initial

¹Video sources include YouTube channels, such as The Skin Deep, MovieClips, and talk shows hosted by Jimmy Fallon and Conan O’Brien.

drafts of laughter reasoning annotations.

Human verification and annotation. We conduct human verification to validate and refine the pseudo-labeled laughter reasoning and to annotate laughter types. We recruit annotators through Amazon Mechanical Turk (AMT) for both tasks. For laughter type classification, annotators are provided with video segments and definitions of each laughter type. Then annotators instructed to label the laughter with confidence score, to ensure consistent labeling. The example of AMT survey is given at Appendix A.6. For laughter reasoning, annotators are instructed to watch the video segments, review the pseudo-generated explanations, and refine or correct them as needed. This verification step ensures both the accuracy and naturalness of the final annotations.

In total, our curated SMILE-Next comprises 3,590 video clips with 6,386 question–answer pairs, each annotated with multimodal textual representations, as summarized in Table 1. The detailed statistics are provided at Appendix A.1. SMILE-Next is approximately four times larger than SMILE (Hyun et al., 2024) and covers a broader range of scenarios and tasks, enabling more comprehensive and general laughter understanding.

4 Method

We aim to build a laughter expert LLM, a unified framework for identifying and understanding laughter in social interactions, building upon our curated SMILE-Next. To expand the LLM’s capabilities for laughter understanding, we introduce two key components: (1) Laughter-specific Self-Instruction, which automatically scales the dataset by generating synthetic question–answer pairs based on the human-annotated SMILE-Next data to cover more diverse scenarios; and (2) Mixture-of-Laugh-Experts (MoLE), a parameter-efficient tuning framework that assigns lightweight LoRA experts to distinct laughter-related competencies while sharing a common backbone model.

4.1 Laughter-specific Self-Instruction

While SMILE-Next covers diverse laughter-related tasks and scenarios with rich human annotations, we further expand its scale and scenario diversity to improve generalization to real-world settings. Motivated by Wang et al. (2022a), we propose a self-instruction framework that synthesizes instruction–response pairs tailored for laughter understanding.

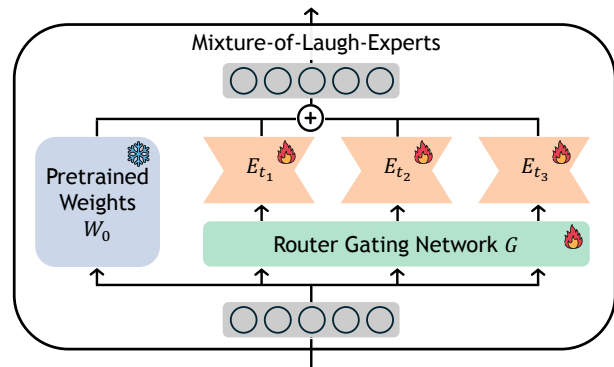


Figure 2: **Architecture of Mixture-of-Laugh-Experts.** LoRA-based expert modules are added to a frozen pre-trained weights and are dynamically weighted by a router gating network, resulting in laughter task-specific specialization.

This framework leverages an external LLM (We use GPT-4 API) to generate both instructions and responses. This approach facilitates a scalable data acquisition pipeline, enabling rich knowledge distillation through supervised fine-tuning without additional manual annotation. Details of the prompts used for laughter-specific self-instruction, along with generated examples, are provided in the Appendix A.4.

Task generation. We begin by bootstrapping three seed tasks from SMILE-Next: laughter detection, type classification, and reasoning, and prompt an external LLM M to generate new laughter-related tasks. We use GPT-4 (Achiam et al., 2023) as M . The synthesized tasks cover a broader range of laughter-related competencies, such as predicting laughter behavior, reactions or intensity, and emotion classification, thereby extending the laughter understanding capability beyond the three original task categories.

Question and answer generation. Given the generated tasks, we further prompt the same LLM M to synthesize question–answer pairs for each task, grounded in diverse real-world laughter scenarios and social relationships. These include, for example, workplace hierarchies (e.g., boss–employee) and informal relationships (e.g., friends). We observe that the self-instruction process effectively generates variations in laughter reasoning across different social contexts, enhancing the dataset’s generality. The resulting synthetic pairs expand coverage beyond humor-driven laughter to include socially motivated laughter, such as nervous, forced, sarcastic, or tension-mitigating responses.

Input Modality	Model	Laughter Detection		Laughter Type Classification		Laughter Reasoning			
		F1 (↑)	Acc. (↑)	F1 (↑)	Acc. (↑)	BLEU ₄ (↑)	METEOR (↑)	ROUGE _L (↑)	SentBERT (↑)
Audio-Visual LLM	MiniCPM-o-v2.6	0.1952	0.3652	0.1919	0.2362	0.0802	0.1185	0.2632	0.5732
	Qwen2.5-Omni-7B	0.8373	0.8344	0.5856	0.5802	0.1683	0.1919	0.3549	0.7221
Visual LLM	Qwen2.5-VL	0.8586	0.8674	0.6949	0.7402	0.1665	0.1764	0.3284	0.6580
	Video-LLaVA	0.8692	0.8739	<u>0.7589</u>	0.7912	0.1769	0.1818	0.3419	0.6567
LLM (Our approach)	Vicuna-v1.5	0.9675	0.9696	0.7460	0.7953	0.2354	0.2311	0.4191	0.7907
	LLaMA3	0.9674	0.9696	0.8067	0.8425	0.2427	0.2328	0.4168	0.7828
	Qwen2.5	0.9629	<u>0.9652</u>	0.7094	0.7559	0.2303	0.2238	0.4097	0.7666

Table 2: **Quantitative results.** We report the result of SMILE-Next at diverse multimodal LLMs and LLMs. Utilizing our textualized multimodal cues, LLMs can outperform AV-LLMs and V-LLMs which utilizes implicit multimodal cues. Models marked as Our Approach denote LLMs trained with the Laughter-specific Self-Instruction and Mixture-of-Laugh-Experts.

4.2 Mixture-of-Laugh-Experts

Instead of training a single model to handle all laughter tasks, we propose the LoRA-based Mixture-of-Laugh-Experts (MoLE) framework which incorporates multiple trainable, lightweight expert modules. This design preserves the original LLM’s reasoning capabilities while enabling task-adaptive expert specialization. Our approach explicitly decouples shared linguistic and reasoning knowledge from task-specific adaptations.

Figure 2 illustrates the MoLE framework. We freeze the pretrained weights W_0 of the backbone LLM to preserve its general reasoning knowledge and train only task-specific LoRA expert modules for each task, allowing the model to learn task-adapted behaviors efficiently. Such training strategy ensures versatility across a wide spectrum of task complexities, enabling the model to excel in diverse scenarios ranging from intricate reasoning to granular classification and detection.

For each laughter task $t \in \{t_1, t_2, t_3\}$, we initiate Laugh Expert E_t to each task. Each expert is injected to transformer’s linear layer as an LoRA adapter. Formally, for a linear transformation $W \in \mathbb{R}^{m \times n}$, $x \in \mathbb{R}^n$ represents the input (*i.e.*, the previous hidden state) and h as an output hidden state, LoRA reparameterizes it as

$$h = Wx = (W_0 + BA)x, \quad (1)$$

where $B \in \mathbb{R}^{m \times r}$ and $A \in \mathbb{R}^{r \times n}$.

To incorporate a task-specific expert for laughter understanding, we extend this formulation by introducing a mixture of LoRA experts. The residual adaptation term is replaced by a weighted combination of expert outputs:

$$h = W_0x + \sum_{i=1}^{T=3} R_i \cdot E_i, \quad (2)$$

where T is the number of task-specific experts, E_i is the output of the LoRA-adapted expert module

for task i , and $R_i \in \mathbb{R}$ is the router gating weight assigned to expert i . The router gating weights R_i are computed by first producing a gating weight, then normalized using a softmax function:

$$R_i = \text{softmax}(W_g x), \quad E_i = B_i A_i x \quad (3)$$

where $W_g \in \mathbb{R}^{T \times n}$ denotes learnable parameters and $B_i \in \mathbb{R}^{m \times r}$, $A_i \in \mathbb{R}^{r \times n}$ as an LoRA adapter. The resulting gating weights satisfy $\sum_{i=1}^T R_i = 1$ and enable *soft routing*, allowing each input instance to be adaptively processed by a weighted combination of all experts.

As a result, each Laugh Expert E_t learns task-specific specializations while enabling smooth interpolation across LoRA experts, which facilitates effective knowledge sharing across tasks while preserving task-specific expertise for laughter understanding. Combined with the pretrained LLM’s reasoning capabilities, we build an expert laughter understanding model that balances generalization and specialization.

5 Experiments

In this section, we describe the experimental setup and evaluate the performance of our proposed laughter expert LLM on the SMILE-Next dataset. We further present ablation studies to analyze the impact of our design choices.

Implementation details. Since our proposed self-instruction and Mixture-of-Laugh-Experts (MoLE) framework are model-agnostic, we evaluate our methods on three text-only LLMs: Vicuna-v1.5-7B (Chiang et al., 2023), LLaMA3 (Grattafiori et al., 2024), and Qwen2.5 (Wang et al., 2024). All models fully utilize the textualized multimodal inputs provided by the SMILE-Next dataset. We integrate the MoLE into all linear layers of each LLM using LoRA adapters. We assign three experts, with the LoRA rank r set to 8 and the scaling

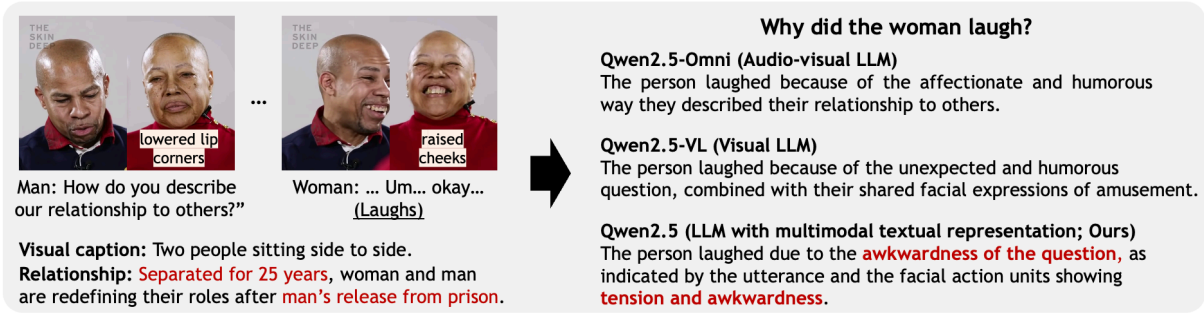


Figure 3: **Qualitative results.** As shown, LLMs trained at SMILE-Next can reason and analyze about the ambiguous and slight laughter, exactly giving the reason and classifying what this laugh is. Comparing AV-LLM and V-LLM to the LLMs which uses textual multimodal cues, it can explicitly point out to analyze why this laugh has occurred.

factor α set to 16 for each expert. The models are supervised fine-tuned with the SMILE-Next training set for three epochs with a batch size of 4, a learning rate of 1×10^{-4} , and a dropout rate of 0. Training is conducted using DeepSpeed ZeRO-3 for memory efficiency.

Baselines. We compare our approach against LLMs with varying multimodal configurations. First, we evaluate audio-visual LLMs, specifically Qwen-Omni-7B (Xu et al., 2025) and MiniCPM-o-v2.6 (Yao et al., 2024), which process raw video and audio inputs directly without textualized cues. Second, we include visual LLMs such as VideoLLaVA (Lin et al., 2023) and Qwen2.5-VL (Bai et al., 2025); these models accept video inputs alongside textualized audio representations, including speech transcripts and relational descriptions. All baselines are fine-tuned on the SMILE-Next for a single run training set using LoRA applied to all linear layers ($r = 8, \alpha = 16$) to ensure consistency.

Metrics. For laughter detection and type classification tasks, we report standard classification metrics, including precision, recall, F1-score, and accuracy. For the laughter reasoning task, we evaluate generation quality using BLEU₄ (Papineni et al., 2002), METEOR (Banerjee and Lavie, 2005), ROUGE_L (Lin, 2004), and SentBERT (Reimers and Gurevych, 2019).

5.1 Experimental Results and Analyses

Quantitative comparison. Table 2 summarizes the quantitative comparison across several types of LLMs that use different input representations. The results show that our proposed approach, leveraging LLMs with textualized multimodal representations achieves the best performance across all three tasks: laughter detection, classification, and reasoning, outperforming both audio-visual LLMs and

visual LLMs. The text-only LLM backbones we evaluate show comparable performance trends, indicating that the gains primarily arise from the proposed representation and modeling strategy rather than from a specific backbone. These findings suggest that our explicitly extracted and disentangled multimodal textual cues effectively support laughter understanding, consistent with observations in Hyun et al. (2024). In contrast, models that directly process raw audio-visual or visual inputs must handle highly entangled signals, which makes it more challenging to capture the underlying causes of laughter.

Qualitative comparison. Figure 3 presents a qualitative example comparing laughter reasoning generated by an audio-visual (AV) LLM, a visual (V) LLM, and our proposed LLM. In this example, the woman laughs due to awkwardness and nervousness during the conversation. The AV-LLM interprets the scene as humorous, incorrectly conflating relational cues with spoken utterances. The V-LLM, which takes utterances and acoustic features as textual multimodal inputs, produces vague explanations and similarly misattributes the laughter to humor. In contrast, our approach, an LLM leveraging all textualized multimodal cues captures the underlying awkwardness of the interaction and reflects it accurately in its reasoning by jointly considering relationship information and facial action units. These results demonstrate that extracting multimodal cues into textual representations enables more precise and robust laughter understanding in real-world settings. Additional qualitative results are at Appendix A.5.

Human preference evaluation. We conducted a human preference evaluation of laughter reasoning task in which participants were shown responses from AV-LLM (MiniCPM-o), Vision-LLM

Pair (Left vs Right)	Left wins (%)	Right wins (%)	Tie
LLM (ours) vs V-LLM	55.7	37.4	7.1
LLM (ours) vs AV-LLM	69.0	26.2	4.8
	AV-LLM	V-LLM	LLM (ours)
Average Rank (\downarrow)	2.38	1.93	1.69

Table 3: **Human study of the textual multimodal cues.** As shown, humans prefer the answers of the proposed LLM, utilizing the textual multimodal cues instead of providing the entangled, raw signals.

Laughter Detection				
	Precision (\uparrow)	Recall (\uparrow)	F1 (\uparrow)	Acc. (\uparrow)
T	0.8571	0.8669	0.8613	0.8696
T+A+V+R	0.9582	0.9688	0.9629	0.9652
Laughter Type Classification				
	Precision (\uparrow)	Recall (\uparrow)	F1 (\uparrow)	Acc. (\uparrow)
T	0.6397	0.4436	0.4019	0.5433
T+A+V+R	0.6828	0.7559	0.7094	0.7559
Laughter Reasoning				
	BLEU ₄ (\uparrow)	METEOR (\uparrow)	ROUGE _L (\uparrow)	SentBERT (\uparrow)
T	0.2047	0.2064	0.3862	0.7457
T+A+V+R	0.2303	0.2238	0.4097	0.7666

Table 4: **Ablation on multimodal information.** We conduct an ablation on each multimodal cues at Qwen2.5, and as shown utilizing all multimodal cues (transcription (T), acoustic cue (A), visual cue (V) and relation (R)) enhances the performance of all laughter tasks: detection, type classification, and reasoning.

(Qwen2.5-VL), and an LLM (LLaMA3; our setting), for each video sample and asked to rank them based on answer correctness given the video context. As shown in Tab. 3, our LLM, which fully leverages textualized multimodal inputs, is consistently preferred over both V-LLM and AV-LLM, achieving a higher win rate (i.e., Left wins) and the lowest average rank. Notably, the strong preference for LLM suggests that textualized multimodal cues enable more effective laughter reasoning both better aligns with human judgments. Overall, these results demonstrate that analyzing and reasoning of laughter based on textual multimodal representations yields outputs that are more consistent with human interpretation.

5.2 Ablation Studies

Multimodal cues. Table 4 analyzes the impact of different textual multimodal cues. We compare models using only text transcripts (T) against those incorporating the full set of information: transcripts (T), acoustic (A), visual (V), and relational (R) cues. The results show that relying solely on textual utterances is insufficient for robust laughter understanding, as laughter can arise from a wide range of non-verbal signals, such as facial expres-

Model	Laughter Detection			
	Precision (\uparrow)	Recall (\uparrow)	F1 (\uparrow)	Acc. (\uparrow)
Vicuna-V1.5	<u>0.9724</u>	0.9573	0.9645	0.9522
(+Self-Instruction)	0.9912	<u>0.9709</u>	0.9806	<u>0.9630</u>
(+MoLE)	0.9635	0.9722	<u>0.9675</u>	0.9696
Qwen2.5	0.9550	<u>0.9684</u>	<u>0.9608</u>	0.9636
(+Self-Instruction)	0.9599	0.9587	0.9589	0.9587
(+MoLE)	<u>0.9582</u>	0.9688	0.9629	0.9652
Model	Laughter Type Classification			
	Precision (\uparrow)	Recall (\uparrow)	F1 (\uparrow)	Acc. (\uparrow)
Vicuna-V1.5	0.5164	0.5370	0.4892	0.6923
(+Self-Instruction)	<u>0.7029</u>	<u>0.7250</u>	<u>0.7050</u>	<u>0.7229</u>
(+MoLE)	0.7412	0.7608	0.7460	0.7953
Qwen2.5	0.6721	<u>0.7480</u>	<u>0.7023</u>	0.7480
(+Self-Instruction)	0.7066	0.6809	0.6852	0.7953
(+MoLE)	<u>0.6828</u>	0.7559	0.7094	<u>0.7559</u>
Model	Laughter Reasoning			
	BLEU ₄ (\uparrow)	METEOR (\uparrow)	ROUGE _L (\uparrow)	SentBERT (\uparrow)
Vicuna-V1.5	0.1872	0.1869	0.3424	0.6999
(+Self-Instruction)	<u>0.2073</u>	<u>0.2110</u>	<u>0.3840</u>	<u>0.7423</u>
(+MoLE)	0.2354	0.2311	0.4191	0.7907
Qwen2.5	0.2272	0.2227	<u>0.4053</u>	0.7640
(+Self-Instruction)	<u>0.2289</u>	<u>0.2250</u>	0.4045	0.7704
(+MoLE)	0.2303	0.2238	0.4097	<u>0.7666</u>

Table 5: **Ablation on Laughter-specific Self-Instruction & MoLE.** Sequentially incorporating Self-instruction and MoLE improves overall performance across the three laughter understanding tasks.

sions, vocal pitch and tone, background context, and interpersonal relationships. This demonstrates that jointly considering contextual, acoustic, and visual cues, and relational cues is essential for robust laughter understanding.

Effect of self-instruction and MoLE. As shown in Table 5, self-instruction consistently improves performance across all three tasks for most LLM backbones. In particular, for the laughter reasoning task, self-instruction helps preserve the model’s reasoning capabilities while simultaneously improving detection and classification performance. This suggests that self-instruction effectively bridges gaps in real-world scenario diversity that cannot be fully captured by human-annotated data alone. Furthermore, incorporating the MoLE framework yields further performance gains across tasks. This improvement indicates that introducing multiple lightweight experts enables the model to better specialize in distinct laughter-related competencies while sharing a common backbone.

Analysis on MoLE router weights. Figure 4 visualizes the MoLE router weights and reveals task-dependent expert activation patterns across laughter understanding tasks. While Expert 1 remains dominant across all three tasks, the relative contributions of the other experts vary depending on the task. In particular, laughter type classification assigns more



Figure 4: **Visualization of MoLE’s router weights.** The results show dominant activation of Expert 1, with task-dependent variations, classification assigns more weight to Expert 2, while detection does so for Expert 3.

	Detect. (ms)	Cls. (ms)	Reasoning (ms)	All (ms)
Single expert	981	790	2802	1494
Multi-experts(MoLE)	991	796	2845	1513
Difference	+10	+6	+43	+19

Table 6: **Latency comparison between single expert and multi-expert (MoLE).** We report the latency of single-expert and proposed multi-expert (MoLE). MoLE demands only marginal latency, about 1-2% additional time.

weight to Expert 2, while laughter detection assigns slightly more weight to Expert 3, whereas laughter reasoning relies more heavily on Expert 1. These variations indicate that the MoLE framework dynamically allocates expert capacity according to task-specific requirements, rather than uniformly activating a single expert.

Noise robustness in multimodal inputs. We evaluated the model’s robustness to noise in textualized multimodal inputs by perturbing a single modality (facial action units or acoustic features) while keeping others unchanged. The model maintains stable reasoning under moderate noise, up to 30% perturbation in facial features and 40–50% noise in acoustic signals. Beyond these thresholds, performance degrades, leading to incorrect or unreliable interpretations (e.g., misjudging pitch). We assume that such robustness happened to our model due to the ability of leveraging multimodal complementarity, down-weighting corrupted inputs and relying on remaining modalities to sustain accurate reasoning.

Computational Latency of MoLE. We report the latency of proposed architecture MoLE at Table 6. MoLE introduces only a marginal latency increase of approximately 0.02 seconds per sample compared to the single-expert model, corresponding to about 1–2% additional inference time, which is negligible in practice. Rather than latency reduction, MoLE enhances laughter understanding while incurring minimal computational overhead.

6 Conclusion

In this work, we introduce SMILE-Next, an enhanced multimodal dataset for real-world laughter understanding that covers three laughter-related tasks: laughter detection, laughter type classification, and laughter reasoning. Based on the SMILE-Next, we propose a unified framework for laugh expert LLM that combines laughter-specific self-instruction with a Mixture-of-Laugh-Experts modeling approach to enable effective laughter understanding in LLMs. By textualizing multimodal signals, our framework captures critical cues underlying laughter, consistently outperforming raw multimodal representations in extensive experiments. Moreover, by jointly modeling multiple laughter-related tasks with parameter-efficient fine-tuning, we develop a robust laughter expert LLM for real-world scenarios. Overall, this work advances the understanding of complex human behaviors with a particular focus on laughter in natural social interactions.

Limitation

SMILE-Next sources videos from a wide range of scenarios, still it may not fully capture all possible forms of laughter occurring in real-world settings. Certain rare or atypical laughter cases—such as culturally specific expressions, nuanced social norms, or complex group interactions—may be underrepresented. In addition, the dataset primarily focuses on English-language content, reflecting the dominance of English-speaking sources, and therefore does not adequately cover laughter expressions in other languages or cultural contexts. Future work will explore incorporating multilingual and cross-cultural laughter data, along with more complex social interactions, to further improve coverage and generalization.

Ethical Considerations

SMILE-Next is curated exclusively from publicly available videos, with data usage governed by the original platforms’ policies. While laughter understanding can enable more natural human–AI interactions, misinterpretation or misuse of such system can raise concerns related to bias or unintended social consequences. We encourage responsible use of this dataset and associated models, with careful consideration of ethical implications. Also, AI assistants were used only for grammatical correction and writing refinement.

Acknowledgment

This work was partially supported by InnoCORE program (N10250156, KAIST InnoCore LLM (25%); 26-InnoCORE-01, Trust-Enhanced Mutualistic Bio-Embedded AI (25%)), Institute of Information & communications Technology Planning & Evaluation (IITP) (No. RS-2024-00457882, National AI Research Lab Project) (25%), and the National Research Foundation of Korea (NRF) (No. RS-2024-00451947) (25%) grants funded by the Korea government (MSIT). This work is also funded by Toonation Co.

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, and 1 others. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Khalid Alnajjar, Mika Hämmäläinen, Jörg Tiedemann, Jorma Laaksonen, and Mikko Kurimo. 2022. [When to laugh and how hard? a multimodal approach to detecting humor and its intensity](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 6875–6886, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Issa Annamoradnejad and Gohar Zoghi. 2020. Colbert: Using bert sentence embedding for humor detection. *arXiv preprint arXiv:2004.12765*.
- Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, and 1 others. 2025. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*.
- Max Bain, Jaesung Huh, Tengda Han, and Andrew Senior. 2023. Whisperx: Time-accurate speech transcription of long-form audio. *arXiv preprint arXiv:2303.00747*.
- William Sims Bainbridge, Edward E Brent, Kathleen M Carley, David R Heise, Michael W Macy, Barry Markovsky, and John Skvoretz. 1994. Artificial social intelligence. *Annual review of sociology*.
- Satanjeev Banerjee and Alon Lavie. 2005. [METEOR: An automatic metric for MT evaluation with improved correlation with human judgments](#). In *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, pages 65–72, Ann Arbor, Michigan. 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).
- Shyamal Buch, Cristóbal Eyzaguirre, Adrien Gaidon, Jiajun Wu, Li Fei-Fei, and Juan Carlos Niebles. 2022. Revisiting the "video" in video-language understanding. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Santiago Castro, Devamanyu Hazarika, Verónica Pérez-Rosas, Roger Zimmermann, Rada Mihalcea, and Soujanya Poria. 2019. [Towards multimodal sarcasm detection \(an Obviously perfect paper\)](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4619–4629, Florence, Italy. Association for Computational Linguistics.
- Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. 2023. [Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality](#).
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, and 1 others. 2022. Palm: Scaling language modeling with pathways. *arXiv preprint arXiv:2204.02311*.
- Kerstin Dautenhahn. 2007. Socially intelligent robots: dimensions of human–robot interaction. *Philosophical transactions of the royal society B: Biological sciences*.
- Najim Dehak, Pierre Dumouchel, and Patrick Kenny. 2007. Modeling prosodic features with joint factor analysis for speaker verification. *IEEE Transactions on Audio, Speech, and Language Processing*.
- Shihan Dou, Enyu Zhou, Yan Liu, Songyang Gao, Jun Zhao, Wei Shen, Yuhao Zhou, Zhiheng Xi, Xiao Wang, Xiaoran Fan, and 1 others. 2023. [Loramoe: Alleviate world knowledge forgetting in large language models via moe-style plugin](#). *arXiv preprint arXiv:2312.09979*.
- Nan Du, Yanping Huang, Andrew M Dai, Simon Tong, Dmitry Lepikhin, Yuanzhong Xu, Maxim Krikun, Yanqi Zhou, Adams Wei Yu, Orhan Firat, and 1 others. 2022. [Glam: Efficient scaling of language models with mixture-of-experts](#). In *International Conference on Machine Learning (ICML)*, pages 5547–5569.
- William Fedus, Barret Zoph, and Noam Shazeer. 2022. [Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity](#). *Journal of Machine Learning Research*, 23(120):1–39.
- Jon Gillick, Wesley Deng, Kimiko Ryokai, and David Bamman. 2021. [Robust laughter detection in noisy environments](#). In *Conference of the International Speech Communication Association (Interspeech)*.

- Yunhao Gou, Zhili Liu, Kai Chen, Lanqing Hong, Hang Xu, Aoxue Li, Dit-Yan Yeung, James T Kwok, and Yu Zhang. 2023. Mixture of cluster-conditional lora experts for vision-language instruction tuning. *arXiv preprint arXiv:2312.12379*.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, and 1 others. 2024. The llama 3 herd of models. *arXiv e-prints*, pages arXiv–2407.
- Md Kamrul Hasan, Md Saiful Islam, Sangwu Lee, Wasifur Rahman, Iftekhar Naim, Mohammed Ibrahim Khan, and Ehsan Hoque. 2023. Textmi: Textualize multimodal information for integrating non-verbal cues in pre-trained language models. *arXiv preprint arXiv:2303.15430*.
- 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)*, pages 2046–2056, Hong Kong, China. Association for Computational Linguistics.
- 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 *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 688–714, Toronto, Canada. Association for Computational Linguistics.
- Lee Hyun, Kim Sung-Bin, Seungju Han, Youngjae Yu, and Tae-Hyun Oh. 2024. **Smile: Multimodal dataset for understanding laughter in video with language models**. In *NAACL-HLT (Findings)*, pages 1149–1167.
- Dayoon Ko, Sangho Lee, and Gunhee Kim. 2023. Can language models laugh at youtube short-form videos? In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 2897–2916.
- Jie Lei, Licheng Yu, Mohit Bansal, and Tamara Berg. 2018. **TVQA: Localized, compositional video question answering**. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1369–1379, Brussels, Belgium. Association for Computational Linguistics.
- Dmitry Lepikhin, HyoukJoong Lee, Yuanzhong Xu, Dehao Chen, Orhan Firat, Yanping Huang, Maxim Krikun, Noam Shazeer, and Zhifeng Chen. 2021. Gshard: Scaling giant models with conditional computation and automatic sharding. In *International Conference on Learning Representations (ICLR)*.
- Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. 2019. Visualbert: A simple and performant baseline for vision and language. *arXiv preprint arXiv:1908.03557*.
- Bin Lin, Yang Ye, Bin Zhu, Jiayi Cui, Munan Ning, Peng Jin, and Li Yuan. 2023. Video-llava: Learning united visual representation by alignment before projection. *arXiv preprint arXiv:2311.10122*.
- 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.
- Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. 2019. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Anton Nijholt, Andreea I Niculescu, Alessandro Valitutti, and Rafael E Banchs. 2017. Humor in human-computer interaction: a short survey. *Adjunct Proceedings of INTERACT*.
- 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 of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Anupama Ray, Shubham Mishra, Apoorva Nunna, and Pushpak Bhattacharyya. 2022. **A multimodal corpus for emotion recognition in sarcasm**. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 6992–7003, Marseille, France. European Language Resources Association.
- Nils Reimers and Iryna Gurevych. 2019. SentenceBERT: Sentence embeddings using Siamese BERT-networks. 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 3982–3992.
- Nicolas Richet, Soufiane Belharbi, Haseeb Aslam, Meike Emilie Schadt, Manuela González-González, Gustavo Cortal, Alessandro Lameiras Koerich, Marco Pedersoli, Alain Finkel, Simon Bacon, and 1 others. 2024. Textualized and feature-based models for compound multimodal emotion recognition in the wild. *arXiv preprint arXiv:2407.12927*.
- Sophie K Scott, Nadine Lavan, Sinead Chen, and Carolyn McGettigan. 2014. The social life of laughter. *Trends in cognitive sciences*.
- Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarz, Andy Davis, Quoc Le, Geoffrey Hinton, and Jeff Dean. 2017. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. In *International Conference on Learning Representations (ICLR)*.

- Sheng Shen, Le Hou, Yanqi Zhou, Nan Du, Shayne Longpre, Jason Wei, Hyung Won Chung, Barret Zoph, William Fedus, Xinyun Chen, Tu Vu, Yuxin Wu, Wuyang Chen, Albert Webson, Yunxuan Li, Vincent Y Zhao, Hongkun Yu, Kurt Keutzer, Trevor Darrell, and Denny Zhou. 2024. Mixture-of-experts meets instruction tuning: A winning combination for large language models. In *International Conference on Learning Representations (ICLR)*.
- David Stauffer. 1999. Let the good times roll: Building a fun culture. *Harvard Management Update*.
- Hiroki Tanaka and Nick Campbell. 2014. [Classification of social laughter in natural conversational speech](#). *Computer Speech & Language*, 28(1):314–325.
- Ruijie Tao, Zexu Pan, Rohan Kumar Das, Xinyuan Qian, Mike Zheng Shou, and Haizhou Li. 2021. Is someone speaking? exploring long-term temporal features for audio-visual active speaker detection. In *ACM International Conference on Multimedia (MM)*.
- Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, and 1 others. 2024. Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*.
- Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A Smith, Daniel Khashabi, and Hannaneh Hajishirzi. 2022a. Self-instruct: Aligning language models with self-generated instructions. *arXiv preprint arXiv:2212.10560*.
- Zhenhailong Wang, Manling Li, Ruochen Xu, Luwei Zhou, Jie Lei, Xudong Lin, Shuohang Wang, Ziyi Yang, Chenguang Zhu, Derek Hoiem, and 1 others. 2022b. Language models with image descriptors are strong few-shot video-language learners. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Melissa B Wanzer, Ann B Frymier, and Jeffrey Irwin. 2010. An explanation of the relationship between instructor humor and student learning: Instructional humor processing theory. *Communication education*.
- Caleb Warren, Adam Barsky, and A Peter McGraw. 2021. What makes things funny? an integrative review of the antecedents of laughter and amusement. *Personality and Social Psychology Review*.
- Orion Weller and Kevin Seppi. 2020. [The rJokes dataset: a large scale humor collection](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 6136–6141, Marseille, France. European Language Resources Association.
- Jessica Williams, Stephen M Fiore, and Florian Jentsch. 2022. Supporting artificial social intelligence with theory of mind. *Frontiers in artificial intelligence*.
- Binzhu Xie, Sicheng Zhang, Zitang Zhou, Bo Li, Yuanhan Zhang, Jack Hessel, Jingkang Yang, and Ziwei Liu. 2024. Funqa: Towards surprising video comprehension. In *European Conference on Computer Vision*, pages 39–57. Springer.
- Jin Xu, Zhifang Guo, Jinzheng He, Hangrui Hu, Ting He, Shuai Bai, Keqin Chen, Jialin Wang, Yang Fan, Kai Dang, and 1 others. 2025. Qwen2. 5-omni technical report. *arXiv preprint arXiv:2503.20215*.
- Li Yao, Yan Wan, Hongjie Ni, and Bugao Xu. 2021. Action unit classification for facial expression recognition using active learning and svm. *Multimedia Tools and Applications*.
- Yuan Yao, Tianyu Yu, Ao Zhang, Chongyi Wang, Junbo Cui, Hongji Zhu, Tianchi Cai, Haoyu Li, Weilin Zhao, Zhihui He, and 1 others. 2024. Minicpm-v: A gpt-4v level mllm on your phone. *arXiv preprint arXiv:2408.01800*.
- Ted Zadori, Ahmet Üstün, Arash Ahmadian, Beyza Ermis, Acyr Locatelli, and Sara Hooker. 2024. Pushing mixture of experts to the limit: Extremely parameter efficient moe for instruction tuning. In *International Conference on Learning Representations (ICLR)*.
- Andy Zeng, Adrian Wong, Stefan Welker, Krzysztof Choromanski, Federico Tombari, Aavek Purohit, Michael Ryoo, Vikas Sindhwani, Johnny Lee, Vincent Vanhoucke, and 1 others. 2022. Socratic models: Composing zero-shot multimodal reasoning with language. In *International Conference on Learning Representations (ICLR)*.

A Appendix

A.1 Additional Explanation of SMILE-Next

In this section, we provide additional statistics and examples of SMILE-Next. First, we provide how our SMILE-Next is constructed, second, we provide how each laughter task is prompted to our laughter-expert LLM.

Dataset statistics. SMILE-Next contains three tasks: laughter detection, type classification, and reasoning. Table 1 reports the train/val/test splits per each task. Furthermore, we provide the statistics of (1) Number of utterances (2) Video duration (3) Length of Instruction (4) Length of Answer at Fig. 5. As shown, SMILE-Next covers 2-8 utterances at most, and for video duration, it covers around 30-50 seconds. For length of instruction, since we feed the multimodal cues in textualized form, most samples require 2000-4000 letters, and output answers are mostly around 150-250 letters.

Constructing QA pairs. Since we formulate each task into question-answering format, we provide the prompts for constructing the instruction dataset for each tasks below.

For laughter detection, we constructed the question instruction as: *"Detection task: You are a laugh detector. Find out if there is laugh in this clip."* And we forced model to respond either *"Yes, there is laugh in this video."* or *"No, there is no laugh in this video."*. For laughter type classification, we constructed the question instruction as: *"Type Classification task: you are to answer the class of laugh type. There are three types: Polite, Satirical, Mirthful. {description of each type}"*, and we forced the answer to be *"The laugh type is {laughter type}"*. For laughter reasoning, we constructed the instruction as: *"Reasoning task: you are to answer why the person laughed at most 30 words, starting with "The audience/person laughed because'"*. So the output format is fixed as *"The audience/person laughed because"*, and conducting the reasoning.

A.2 Detailed Dataset Construction Pipeline

In here, we provide the details of SMILE-Next’s dataset construction pipeline, at each stage (1) Laughter video collection (2) Textual multimodal information extraction (3) Pseudo-label annotation (4) Human verification and annotation.

Laughter video collection. We collected the video clips from YouTube ‘The Skin Deep’ and ‘MovieClips’. The duration of each sources’ video

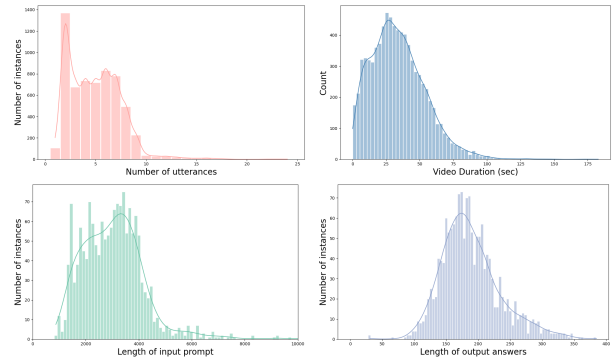


Figure 5: **SMILE-Next statistics.** We report the number of utterances and the video duration of each video samples in SMILE-Next at top row, which shows similar distribution. At bottom row, we also report the length of input and output text, and as shown input prompt is diverse while the output text shows distribution centered at 200.

clip is different, because the type of conversation from the source is different. Similar to SMILE (Hyun et al., 2024), we targeted the last laughter happening in the video clip, since the video clip may contain multiple laughs. Also focusing on the ‘laugh happening scene’, we did not annotate ‘who’ laughed. We assumed the laughter happens contagiously in dyadic conversation which makes the other person also laugh, making the whole atmosphere to ‘laughing scene’.

Textual multimodal information extraction. We report the off-the-shelf models for extracting multimodal information. For the utterance cue acquisition, we have leveraged WhisperX (Bain et al., 2023) and trimmed a video into multiple segments; one utterance for one segment. For each segments, we also leveraged Active Speaker Algorithm (Tao et al., 2021) using audio-visual information in each video. For the visual cue, we acquired video caption for each segment, which explains each segment’s situation, providing the visual cue. We also extracted the facial action units (Yao et al., 2021) of speakers in the video; Accumulating each frame’s facial action and selected most dominant top-3 facial action units for each person. We extract acoustic features from speech utterances using off-the-shelf speech processing tools (Dehak et al., 2007), including the mean and variance of pitch and intensity, as well as jitter and shimmer. Although we initially explored transforming these numerical features into categorical linguistic descriptors (e.g., labeling a mean pitch above 200 Hz as “high pitch”), we found this approach to be

Prompt for GPT-4 Generating Pseudo-Reasoning Labels

Reasoning task: you are to answer why the person laughed given the video clip. The video clip from the youtube channel 'The Skin Deep', titled title, with multimodal information (Utterance, Facial Action Units, Video caption, Relation, Clip description, Acoustic features (10 dimension; 1.mean of F0 contour, 2.var of F0 contour, 3. mean of energy contour, 4. mean of voiced duration, 5. mean of unvoiced duration, 6. var of energy contour, 7. first derivative of fundamental frequency, 8. second derivative of fundamental frequency, 9. jitter, 10. shimmer)) is given. Referring to the multimodal information (Utterance, Prosodic, Facial Action Units, Video caption, Topic) in each utterance, Please make sure that there are a lot of laughs, such as hilarious laugh polite laugh, awkward laugh, dumbfound laugh, etc. Try to detect and reason possible laughter. Based on the given query, give reason about why the person laughed, at most 30 words, starting with 'the person laughed'. For this, please state which part of multimodal information (utterance, Prosodic, Facial Action Units, Video caption, Topic, relation) has stimulated the laugh. For example, utterance "I am an idiot" can be a first starting point for laughter reasoning. If there is not enough context, try to focus on relation of clip description, those can also be a starting point to reason for laugh. Given video clip: {Textualized multimodal information}

Table 7: **Prompt for pseudo-reasoning label annotation.** We provide the prompt for generating the pseudo-reasoning label. The prompt is fed to GPT-4 to generate free-form laughter reasoning.

problematic due to the inherently subjective nature of such mappings. So we retain the raw numerical values and directly incorporate them as acoustic features. These statistics are conveyed to the LLM via natural-language prompts that describe the feature values, enabling the laugh LLM to interpret them using its prior knowledge of speech and prosodic patterns.

Pseudo-label annotation. For generating pseudo-reasoning labels, we prompted an expert large language model (LLM) to generate pseudo-laughter reasoning. We used GPT-4 as an expert LLM, and the prompt is given at Table 7. As shown, we provide the detailed description about given textual multimodal cues, including utterance, facial action unit, video caption, relation and clip description. We further prompt the model to generate reasoning that explicitly references the multimodal cues responsible for inducing laughter, encouraging grounded and interpretable explanations. This design allows the model to produce coherent and context-aware pseudo-reasoning annotations aligned with real-world social interactions. Annotations that passed human verification were retained as-is, while those that did not were further refined by human annotators and subsequently used in the dataset.

Human verification and annotation. For the annotation of the laughter, we ran the survey on Amazon Mechanical Turk (AMT) for labeling the reason and type of laughter corresponding to its

confidence. We have given the precise instruction to start reasoning, for example, when some laughter just happens, without an explicit reason, we instructed annotators to make the reason of laughter considering the laughter as the 'reaction' of some explicit cue (*e.g.* utterance, relation) For laughter type classification, annotators were asked to watch a short video segment containing laughter and assign one of the five labels based on the contextual and acoustic characteristics of the laughter. To ensure annotation quality, 3 annotators were assigned per instance. After the annotation, we select the most-voted laughter type as the GT, considering the confidence score of the selected laughter type. We add the screenshot of laughter type classification AMT survey at Sec. A.6. We measured inter-annotator agreement using Fleiss' Kappa (0.42195). Despite the task's inherent subjectivity, this score indicates moderate agreement beyond chance, suggesting that the labeling guidelines ensured consistency and supporting the reliability of the annotations.

A.3 Additional Quantitative Result: Comparison of Implicit and Textual Cues

Result. To verify the effectiveness of our textualized multimodal representation, we compare two models that share the same backbone (Vicuna-7B-v1.5): (1) a multimodal LLM that encodes multimodal inputs using feature encoders, and (2) an LLM that consumes the proposed textualized multimodal information. In here, we evaluate using the

top-k	Task name	Counts
1	Analysis Task	31
2	Classification Task	30
3	Prediction Task	29
4	Correlation Task	24
5	Sentiment Analysis Task	22

Table 10: **Top-5 generated tasks through laugh self-instruction.** We have parsed each generated instances’ task and counted which tasks are generated. As shown, newly generated tasks contains various of task, containing correlation or sentiment analysis tasks.

Analysis of generated tasks. We analyze what tasks are generated from the Laughter-specific Self-Instruction. First, we report the wordcloud of the task names generated from self-instruction at Fig. 6. As shown, it contains diverse tasks including analysis, classification, prediction, and correlation. Based on SMILE-Next’s three core tasks, such a self-instruction method diversifies the number of tasks with generic laughter situations. Second, we report the top-5 tasks generated through the self-instruction at Table 10. We have clustered the self-instruction-generated tasks by parsing and renaming, and, as shown, this has yielded around 300 new laughter-related tasks, increasing diversity. So, enhancing our manual laughter-related task, self-instruction of the laughter dataset guarantees the diversity.

Implementation of Laughter-specific Self-Instruction. Our implementation follows the official Self-Instruction framework (Wang et al., 2022a). Building upon this foundation, we extend the original setup by incorporating laughter-specific prompts to generate new tasks and corresponding instances tailored to laughter understanding. These prompts are designed to elicit diverse laughter-related behaviors and reasoning patterns, thereby complementing the supervision provided by SMILE-Next. The prompts used for task and instance generation are shown in Table 11. We explicitly encouraged the model to generate not only amusement-driven laughter but also socially motivated ones such as forced, nervous, or sarcastic laughter, ensuring that the laughter-specific Self-Instruct framework captures a diverse range of laughter-inducing scenarios.

Examples of Laughter-specific Self-Instruction. We report the examples of generated self-

instruction samples at Table 11. As shown at Table 11 below of changing the laughter reasoning, according to its relations, for example when the situation is that boss is giving the funny joke, and the employees laughed, and our synthesized data pinpoint their relations to reason for the laugh, saying that it is a forced laugh that laugh happened due to its forced situation, not because the joke is funny. This enhances our laughter dataset to be more generalized to the real situation, while our annotated dataset considers only the partial part of the real-world situations.

A.5 Additional Qualitative Results

We provide additional qualitative results on SMILE-Next in Fig. 7-(a). As illustrated, the proposed laughter-expert LLM produces more accurate and coherent predictions at reasoning tasks compared to the competing models, including audio-visual LLMs and visual LLMs.

A.6 Example of Textualized Multimodal Cues

We report the example of a textualized multimodal cue at Fig. 7-(b). As shown, proposed textualized multimodal cues are annotated per video segments, with the transcription, and visual description of the scene, facial action units describing the person’s facial expression, and the numeric acoustic features. With such textualized multimodal cues, proposed LLM can analyze the given cues of the video and understand how and why the laughter happened in the context.

A.7 Human Annotations for SMILE-Next

We report the instruction and the screenshot of AMT human study at laughter type classification at Fig. 8. For laughter type classification, the annotators were instructed to pick the laughter type with the confidence of the answers, and also they write why they have picked that laughter type in complete sentences. Each sample was annotated by three annotators, and the final laughter type label was determined via majority voting. For laughter reasoning, the annotators were instructed to determine whether the given expert LLM (*e.g.*, GPT4)’s pseudo-reasoning label is correct. If it is correct, the pseudo-reasoning label is used for the annotation. If not, annotators are instructed to write a free-form answer about why the laughter has occurred, and such samples are used for the laughter reasoning annotations. All payments are adequate considering the participants’ demographics.

Table 11: Prompt for Laughter-specific Self-Instruction and Examples of Generated Self-Instruction Instances.

Laughter-specific Self-Instruction Prompts
<p>Task Generation: Come up with a series of tasks: {few shot examples}</p>
<p>Instance Generation: Generate new examples that follow the same format as above. Include a variety of laughter situations, not only humorous or joyful ones, but also socially-driven laughs such as forced, nervous, or sarcastic laughter. Ensure the context reflects subtle multimodal cues such as facial expression, tone, or social dynamics. {few shot examples}</p>
Examples of Generated Self-Instruction Instances
<p>Evaluating task: Rate each laugh in the scene based on intensity and context, and determine whether it was genuine or forced.</p> <p>Input: During a tense board meeting at her office, Sarah tries to lighten the mood with a faint chuckle after the boss makes a dry joke. Her co-workers don't seem to respond much.</p> <p>Answer: Forced, low intensity</p>
<p>Correlation task: Derive the relationship between the acoustic features and the intensity or type of laugh.</p> <p>Input: Acoustic feature: irregular pace, variation in pitch (Laughter)</p> <p>Answer: This could suggest a nervous laughter or possibly a fake laugh.</p>
<p>Laugh pattern analysis task: Analyze the recurring patterns, if any, in the laughter based on the context, relationship, and emotional state of the people involved.</p> <p>Input: Shawn and Brooke are best friends. They often tend to laugh when shared jokes or funny moments between them, creating a contagious laughter pattern that generally initiates with Shawn's giggle followed by Brooke's loud laughter.</p> <p>Answer: The recurring pattern observed in Shawn and Brooke's laughter is predominantly contingent on the funny moments or jokes shared between them. The laughter pattern generally initiates with Shawn's giggle, leading to Brooke's more pronounced laughter. This pattern showcases the close rapport between the two, characterized by the contagion of laughter.</p>



Ben, speaking with and upper lip corners, says, "So it's like a dog, right? Oh, there he is," while woman and a man are showing. Andie, says, "Ben, Ben, you're hurting Kroll's feelings." while small dog wearing a sweater standing next to a pool table.

Ben, speaking with and lowered eyebrows, says, "Well, sure, because that's clearly what... Oh, wait, there's more! (light laugh)" while small dog wearing a sweater.

Relation:
Ben goes and talks to Andie and to start their relationship, not revealing the true intention.

Laughter Detection & Classification & Reasoning Answers

- MiniCPM-o-v2.6 (AV-LLM):
- No. There is no laugh.
 - Based on the provided information, it is not possible to classify ...
 - The person laughed because the unexpected twist in the movie plot made them unexpectedly amused.
- Qwen2.5-Omni (AV-LLM):
- Yes. There is laugh.
 - The laugh type is satirical.
 - The person laughed because they found the Chinese Crested dog's appearance and the couple's reaction to it amusing.
- Qwen2.5-VL (V-LLM):
- Yes. There is laugh.
 - The laugh type is mirthful.
 - The person laughed because of the unexpected and humorous sight of a small dog dressed in a human-like outfit, including a sweater and hat, walking on a pool table.
- LLaMA3 (LLM):
- Yes. There is laugh.
 - The laugh type is mirthful.
 - The person laughed because of the humorous dialogue about a dog named Ben, and the subsequent playful interactions and jests about a sweater.
- Qwen2.5 (LLM):
- Yes. There is laugh.
 - The laugh type is mirthful.
 - The person laughed because of the humorous interaction about the dog's name and the unexpected gift of a helmet in the video clip.
- Ours (LLM):
- Yes. There is laugh.
 - The laugh type is polite.
 - The person laughed because of the humorous exchange about Kroll's name and the unexpected gift.



Man, speaking with and brow lowerer, lid raiser, says, "Because my dad used to always take me to the game and be like, you see that guy right there?" while showing the man an talking.

Ben, speaking with and lowered eyebrows, says, "Like, if I go to a game and saw you there, oh, I'm going to be like, hey, Conan's here. Throw me the goddamn ball and watch me go to work." while showing two man talking.


Relation:
Conan O'Brien is the host of the show, and the man is guest.

Laughter Detection & Classification & Reasoning Answers


- MiniCPM-o-v2.6 (AV-LLM):
- No. There is no laugh.
 - Based on the provided information, it is not possible to classify ...
 - The person laughed because the host made a funny joke during their conversation.
- Qwen2.5-Omni (AV-LLM):
- Yes. There is laugh.
 - The laugh type is satirical.
 - The person laughed because they found the story about Shaquille O'Neal's father taking him to games and encouraging him to be the best player to watch amusing.
- Qwen2.5-VL (V-LLM):
- Yes. There is laugh.
 - The laugh type is mirthful.
 - The person laughed because of the unexpected revelation about only picking two or three people to bully while being in LA, which seems surprising and laughable.
- LLaMA3 (LLM):
- Yes. There is laugh.
 - The laugh type is mirthful.
 - The person laughed because the speaker made a humorous remark about being eager to receive the ball and show off his skills, which was unexpected and presented in a light-hearted manner.
- Qwen2.5 (LLM):
- Yes. There is laugh.
 - The laugh type is mirthful.
 - The person laughed because of the humorous remark made by the speaker in the last utterance, "Throw me the goddamn ball and watch me go to work." This unexpected and humorous statement provoked laughter
- Ours (LLM):
- Yes. There is laugh.
 - The laugh type is polite.
 - The person laughed because of the humorous remark made about the speaker's father not being present to see him but still wanting to throw him the ball and watch him go to work.

(a) Additional qualitative results

*Facial Action Units (FAU), Utterance, Visual Caption, Acoustic Features



Will, speaking with lowered brows, raised cheeks, and lower lip corners, says, "There's someone here I want you to meet." while wearing a suit and tie standing in a room, with acoustic features [183, 25, 0, 3, -1, -2]



Will, speaking with lowered brows, raised cheeks, and lower lip corners, says, "You look... You look great. Why are you in New York?" while wearing a suit and tie with a surprised look on his face, with acoustic features [159, 41, -14, 7, 0, -1]



Emily, speaking with high brows, and raised cheeks, says, "I live here. I took a job with the city education department." while woman in a black dress with trees, with acoustic features [178, 57, -20, 7, -1, -1]



Will, speaking with tightened lids, lowered cheeks, and lowered lip corners, says, "I called you. I know." while in suit with acoustic features [199, 71, -16, 10, -1, -2]



Emily, speaking with raised cheeks, and upper lip corners, says, "I've been meaning to call you back. Liar. (laughs)" while woman in a black dress with trees, with acoustic features [127, 46, -16, 7, -1, -1]

(b) Example of textualized multimodal annotation at SMILE-Next

Figure 7: Additional qualitative results and example of textualized multimodal cues.

Instructions (click to expand)

In this HIT, you watch a short video clip from talkshow, movie, and intimate conversation then assess why the person laugh at the end of the video. The video may contain multiple laughs, but you **only need to deal with the last laugh in the video.**

Q1. Pick the one laugh type that best appropriate for the last laugh.

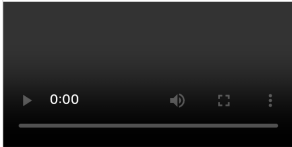
- You get four choices for laugh types. Pick which one is best appropriate.
- If there is no laughter in the video, please check no laugh box.

Q2. Pick the confidence score of your choice on Q1

- 4 is most confident, 1 is minimum confident.

Q3. Why do you think that laugh type (your choice) fit to the video clip.

- Answer with 1 or 2 complete sentences.



Q1. Pick the one laugh type that best appropriate for the last laugh.

Polite laugh: It is often used in social situations to show agreement, acknowledge a comment, or smooth over social interactions.

Satirical laugh: It is used in contexts of satire, sarcasm, or irony. It is often used to communicate amusement at someone else's expense or to highlight the absurdity of a situation.

Humor-based laugh: It is associated with genuine mirth or amusementng or funny situation.

Others

No laugh

Q2. Pick the confidence score of your choice on Q1 (4 is most confident).

1 2 3 4

Q3. Why do you think that laugh type (your choice) fit to the video clip.

Answer with 1 or 2 complete sentences.

...

Figure 8: Example screenshot of Amazon Mechanical Turk for human annotation. For laughter classification, the human evaluators are instructed to pick a laughter type with a confident, and to write the reason why they choose it.