

BQA: Body Language Question Answering Dataset for Video Large Language Models

Shintaro Ozaki Kazuki Hayashi Miyu Oba
Yusuke Sakai Hidetaka Kamigaito Taro Watanabe

Nara Institute of Science and Technology (NAIST), Japan

ozaki.shintaro.ou6@naist.ac.jp

{sakai.yusuke.sr9, kamigaito.h, taro}@is.naist.jp

Abstract

A large part of human communication relies on nonverbal cues such as facial expressions, eye contact, and body language. Unlike language or sign language, such nonverbal communication lacks formal rules, requiring complex reasoning based on commonsense understanding. Enabling current Video Large Language Models (VideoLLMs) to accurately interpret body language is a crucial challenge, as human unconscious actions can easily cause the model to misinterpret their intent. To address this, we propose a dataset, BQA, a body language question answering dataset, to validate whether the model can correctly interpret emotions from short clips of body language comprising 26 emotion labels of videos of body language. We evaluated various VideoLLMs on the BQA with and without Multimodal Chain of Thought (CoT) and revealed that understanding body language is challenging, and our analyses of the wrong answers by VideoLLMs show that certain VideoLLMs made largely biased answers depending on the age group and ethnicity of the individuals. We also found consistent error patterns in VideoLLMs¹.

1 Introduction

Video large language models (VideoLLMs) (Wang et al., 2024; Ye et al., 2025; Zhang et al., 2024a; Team et al., 2024) process videos by integrating multimodal inputs into an understanding of the content. These models take video frames, sound, and accompanying text as input and generate text, answers to questions (Maaz et al., 2024; Lei et al., 2018), or predictions based on the video (Xiao et al., 2021; Yi et al., 2020), enabling various applications, such as video summarization and question answering. This capability fosters a future where humans and models coexist, making it essential for VideoLLMs to grasp human emotions and

body language for interaction. One study (Hyun et al., 2024) has investigated emotion detection from body language, identifying smiles and their underlying causes. However, since this approach is limited to analyzing a single emotion, it remains unclear whether the findings are generalized to all human emotions. If VideoLLMs are unable to understand human emotion from body language, they may not be suitable for future applications such as dialogue systems and AI robots, where emotional awareness is crucial for enabling more effective interactions.

Our research focuses on the analysis of various emotional expressions in human body language. We created a dataset called *BQA*, a multiple-choice QA task, in which each body language video is associated with a question regarding a particular emotion comprising four choice answers, e.g., Surprise, Confidence, Anger, and Embarrassment, reformatting the Body Language Dataset created for pose estimation (Luo et al., 2020). The BQA consists of 7,632 short videos (5–10 seconds, 25 fps), depicting human body language with metadata (gender, age, ethnicity) and 26 emotion labels per video. The BQA creation involves four steps using Gemini (Team et al., 2024): extracting answer choices, generating questions, evaluating potential harm, and assigning difficulty labels. Moreover, we evaluated recent VideoLLMs on BQA, with and without Chain of Thought (CoT) (Zhang et al., 2024b) as well as conducted multiple human evaluations, and found the task enough challenging for VideoLLMs. Error analysis revealed biases toward specific ages or ethnicities, and Multimodal CoT, despite improving accuracy, showed consistent error patterns.

2 Body Language Dataset (BoLD)

Body Language Dataset (BoLD) (Luo et al., 2020) is a dataset for recognizing human actions and selecting appropriate emotions created by splitting

¹The dataset is available at <https://huggingface.co/datasets/naist-nlp/BQA>.

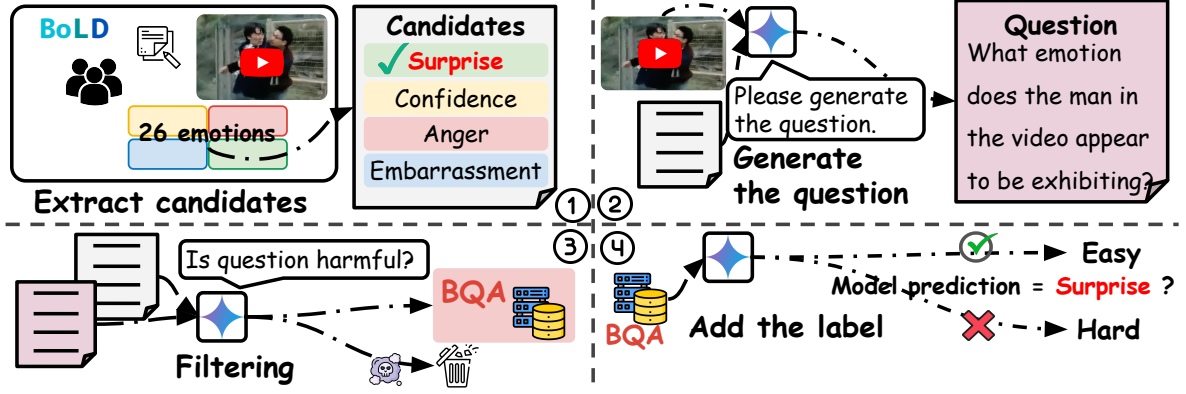


Figure 1: 4 steps of creating the BQA dataset. In STEP1, candidates are created; in STEP2, questions are generated; in STEP3, filtering is conducted; and in STEP4, we assign (Easy/Hard) labels, describing the details in Section 3.




Figure 2: Categorized the 26 emotion labels into 4 groups with similar emotions to extract the candidates.

150 films, totaling 220 hours of footage, resulting in 9,876 video clips. Each clip is approximately 5 seconds long, comprising nearly 125 frames. These videos were annotated with 26 emotion labels (Kosti et al., 2017) via crowdsourcing, where multiple annotators assigned emotion labels on a 10-point scale, which were normalized to represent the emotion of each video. Metadata, including the age, gender, and ethnicity of the individuals in the clips, is also available. However, BoLD is designed for a model directly predicting the emotion, and is not suitable for prompting with clear answers in order to investigate the capacity of VideoLLMs, which expect inference with natural language.

3 Dataset Construction

We transformed BoLD into a multiple-choice QA format by generating questions from video content and using the 26 emotion labels as answers to evaluate how well VideoLLMs understand human emotions expressed through body language, adding steps to extract appropriate choices since BoLD was not designed for LLMs evaluation. To design a QA task for evaluating the LLMs’ understanding of body language for emotional expression, we followed the approach of mCSQA (Sakai

et al., 2024a), which semi-automatically generates QA questions based on candidate answers using LLMs. The whole process consists of four steps as described in Figure 1. First, we extract candidate choices from the BoLD’s metadata (Figure 1-1) followed by question generation using a Gemini based on the video and the candidate choices (Figure 1-2). Then, we automatically filter out inappropriate QAs (Figure 1-3). Lastly, we let the Gemini solve the QAs to evaluate difficulty levels (Figure 1-4).

STEP1: Extract Candidates We categorized 26 emotion labels defined in BoLD into 4 groups: Happiness, Anger, Sadness, and Pleasure, as shown in Figure 2, based on the research in which emotions are classified into four main groups (James, 1890). For the creation of BQA, we apply a multiple-choice question format, where the one with the highest empathy level is treated as correct, and the remaining three options are selected from different emotion groups to ensure that the choices for the QA candidates are selected from each of the groups. The example in Figure 1-1 shows that the correct answer is Surprise from the Pleasure group and the remaining candidates, i.e., Confidence, Anger, and Embarrassment, are drawn from the other groups, i.e., Happiness, Anger, and Sadness, respectively.

STEP2: Generate the Question by VideoLLM Since BoLD does not follow the QA format, we create the appropriate questions from the pairs of candidates and a video by following the prompt design of mCSQA (Sakai et al., 2024a) modified for VideoLLMs. We input the four candidate options (e.g., Confidence, Surprise, Anger, Embarrassment from each group in Figure 1-1) along with the video into Gemini with the highest performance (Team et al., 2024) and let the model generate questions

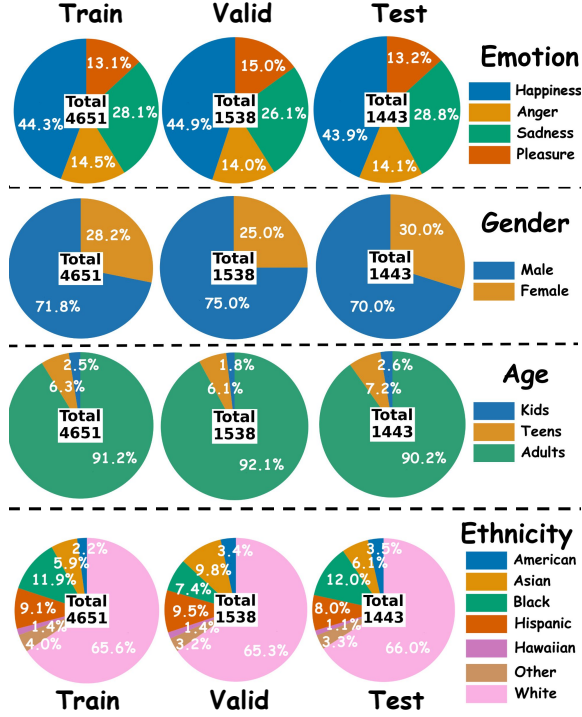


Figure 3: The proportion of metadata in the BQA.

such as “What emotion does the man in the video appear to be exhibiting?” like Figure 1-2 with the prompt we use for the generation in Appendix C.3.

STEP3: Filter the QA by VideoLLM Sometimes, a VideoLLM generates a question which is relatively easy to estimate or answer, e.g., a question that already contains information about the correct candidate, such as “The man looks so shocked. Which emotion is appropriate at this time?” Thus, we let Gemini evaluate whether the generated questions were enough objective and whether they contained superficial information about the correct candidate. If any outputs included the harmful content or did not conform to the conditions for the questions, they were excluded as shown in Figure 1-3.

STEP4: Classify the QA as Easy or Hard Finally, as shown in Figure 1-4, we let Gemini solve the created questions by labeling each question as “Easy” when it could answer or as “Hard” if it could not answer using the prompt presented in Appendix C.3. These question labels allow us to analyze whether a hard question that is not solvable by Gemini is also difficult for other VideoLLMs. For instance, if the correct answer is “Surprise” and Gemini responds with “Anger,” then that question would be classified as “Hard.” After completing these four steps, we split the dataset into training, validation, and test sets in a 6:2:2 ratio, resulting

in 4.5k, 1.5k, and 1.5k questions, respectively. The precise number of data is in Table 3.

What are the Aspects that Make the Samples Easy/Hard? Our analysis of the issues labeled as “Hard” revealed the following findings:

- (1) Neutral expressions: Videos featuring neutral facial expressions often resulted in the hard label.
- (2) Obstructions such as glasses, hats, or sunglasses: Samples involving individuals wearing such items were more likely to be misclassified, leading to the hard label. These findings suggest that VideoLLMs heavily rely on facial expressions when interpreting body language from the footage.

Dataset Analysis Figure 3 shows the distribution of our datasets categorized by our emotion type and three groups of meta information in BoLD: gender, age and ethnicity annotated by humans, revealing that many of the videos feature adult males who are White. It displays the distribution of the four groups of emotions when the annotators selected the most appropriate emotion from the 26 patterns as the correct answer. While Happiness occupies a lot, the overall distribution appears to be balanced.

4 Evaluation

Experimental Setup The models used for evaluation include VideoLLaMA2 (Cheng et al., 2024), LLaVA-NeXT (Zhang et al., 2024a), Qwen2-VL (Wang et al., 2024), and Phi-3.5 (Abdin et al., 2024) with the prompt in Appendix B.2. For the test data, we also used the proprietary models, Gemini (Team et al., 2024) and GPT-4o (OpenAI et al., 2024). We used LoRA-Tuning (Hu et al., 2022) on VideoLLaMA2 with the configuration in Appendix B.1 using the training data and let the model answer with the correct choice in a single word. All audio from the videos was removed to allow for evaluating the model’s ability to interpret body language without relying on auditory information. Additionally, we randomly selected 100 cases from the test set to measure human performance, describing the guideline in Appendix C.1. As an additional evaluation on the test set, we further employed Multimodal Chain of Thought (CoT) to make the models generate reasoning as evidence we call “rationale” for selecting answers. The prompts used for these evaluations are provided in Appendix C.3.

Main Results We show the results in Table 1 with the test set and Table 4 with the valid set. GPT-4o and Gemini achieved higher accuracy than

VideoLLM	#F	Test			Test (CoT)		
		Easy	Hard	Total	Easy	Hard	Total
Human (Rand100)	-	0.96	0.77	0.85	-	-	-
Gemini	*	0.91	0.08	0.61	0.94	0.90	0.92
GPT-4o	*	0.78	0.38	0.60	0.97	0.95	0.96
Phi-3.5	16	0.77	0.41	0.58	1.00	0.96	0.98
Qwen2-VL	16	0.68	0.27	0.47	0.98	0.95	0.97
LLaVA-NeXT	16	0.66	0.30	0.47	1.00	0.96	0.98
VideoLLaMA2	16	0.15	0.01	0.08	0.15	0.01	0.08
VideoLLaMA2 (FT)	16	0.98	0.91	0.94	0.98	0.90	0.94

Table 1: The results using BQA. #F indicates the frame. An asterisk (*) signifies 1 fps. (FT) indicates the LoRA-Tuning model. “Human” is an average of 3 annotators.

the other models. VideoLLaMA2 replied without confirming the choice of format before fine-tuning (FT), resulting in a low score, but after FT, its score surpassed Gemini’s. From this result, the label assignment in STEP4 did not cause hallucinations. Regarding Gemini, which generated the questions, we found that the problems were sufficiently challenging even for the model itself. Furthermore, in STEP4 of Section 3, those labeled as Easy became unsolvable during inference, likely due to the prompt that restricted the output to single words. Regarding CoT, the results showed a large improvement, confirming its effectiveness for VideoLLMs.

5 Analysis and Discussion

We analyzed the videos by gender (Figure 4-A), age (Figure 4-B), and ethnicity (Figure 4-C) in which each model tends to make mistakes. Other models showed lower evaluation results in the Hard setting compared to the Easy setting. This indicates that even if a language model can create questions, it does not guarantee that it can solve them itself. Since the accuracy of the other VideoLLMs, even GPT-4o, was lower than that of Gemini, the dataset proved to be sufficiently challenging for all models.

Which Age do VideoLLMs Often Mistake? We show which age groups models tend to struggle with in Figure 4-B. Higher values indicate a greater tendency to make errors on videos featuring individuals from that age group. These results show that most models do not exhibit bias based on age, while LLaVA-NeXT tends to make more errors on videos featuring “Adults” compared to the others.

Which Ethnicity do VideoLLMs Often Mistake? In Figure 4-C, we show the consistent patterns to make the mistakes based on the ethnicity of the in-

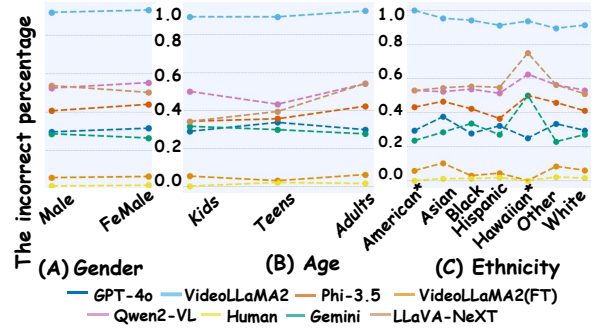


Figure 4: The analysis of incorrectly answered questions shows, from left to right, (A) gender, (B) age, and (C) ethnicity. Note that the higher value indicates more mistakes. An asterisk (*) in (C), especially “American” and “Hawaiian”, indicates that they are native humankind.

dividuals in the videos. Gemini and LLaVA-NeXT tend to make more errors on the problems related to “Native Hawaiian.” Notably, LLaVA-NeXT only achieves around 25% accuracy on these questions.

Why does CoT Improve the Performance? While the Multimodal Chain of Thought (CoT) improved performance, we observed that the rationales often included the correct answer. The inclusion of leaked answers contributes to the performance improvements largely. Thus, the scores achieved by Multimodal CoT should be treated separately from the inherent difficulty of the dataset. However, upon analyzing the types of errors even after CoT, we observed a recurring trend: many of the misclassified instances featured neutral facial expressions or minimal body movements. This observation further supports the claim in Section 3 that VideoLLMs tend to focus more on the face when attempting to understand the body language.

6 Conclusion

Our work created a dataset called BQA to evaluate whether VideoLLMs understand body language that represents emotions and let VideoLLMs solve the questions using BQA. The results show the questions are challenging for all models, confirming the meaning of the dataset. We analyze the types of questions each model tends to get wrong, revealing that some models show a tendency to make more mistakes based on ethnicity and age group. The models also tend to focus more on the face than on body language, and we found that accuracy decreases when there are obstructions that obscure the face. It is also essential to conduct evaluations that focus on the biases in VideoLLMs.

7 Limitations

7.1 The Evaluation by Human Annotator

The random sampling of 100 evaluations was conducted by three people. The overall agreement score is 0.79, indicating a high level of agreement. In the future, we may take care of using crowdsourcing to gather evaluations from a lot of people. Furthermore, it would be beneficial to include evaluations from people of various ethnicities to explore differing perspectives on body language that expresses emotions. However, [Tedeschi et al. \(2023\)](#) argues that human baselines may be unreliable due to factors such as crowdsourced worker payment issues and random sample effects. We should therefore be cautious about the baseline for our human evaluation.

7.2 Video Quality

This study takes care of the possibility that video quality or size may affect accuracy. Specifically, BoLD uses old films, and some of them have noticeably poor quality. Since the same data is used for evaluation, the ranking of the models' accuracy will likely remain unchanged. However, when inputting higher-quality videos, accuracy might improve across all models.

7.3 Frame Issues

Although VideoLLMs claim to handle videos, many actually use image models by treating videos as a sequence of images. In this study, we standardized the number of frames each model can process to 16, but inputting the maximum number of frames may affect the results. However, since inputting more frames increases memory usage, we also need to be mindful of resource constraints.

7.4 Regarding Emotional Expressions

In this study, we categorized 26 patterns of emotional expression into 4 groups based on previous research ([James, 1890](#)). While we extracted options from these four patterns, this method may not be entirely accurate. Future research will focus on how the models behave when we expand the available options.

7.5 The Costs of Calling API

The models used in this paper are GPT-4o (gpt-4o-0806) from OpenAI. GPT-4o is accessed via API, which is subject to change and incurs costs based

on the number of input tokens. In this study, inference costs totaled approximately \$154 and \$100 for Multimodal CoT, but this may change in the future. Additionally, due to cost considerations, we used Gemini-1.5-pro. This model is also accessed via API, which is subject to change and incurs costs based on the number of input tokens.

8 Ethical Considerations

8.1 Taking Care about Culture

The expression of emotions through body language doesn't necessarily remain consistent across all countries. Therefore, we might need to update the dataset to take care of cultural factors in future developments.

8.2 License

Since BoLD does not have a clear license, we believe its use for research purposes is unproblematic.

8.3 Large Scale Human Evaluation

Although human evaluations of BQA aim to minimize bias, implicit biases may still remain. In the future, it may be necessary to employ multiple annotators for a fair assessment. However, as mentioned in [Tedeschi et al. \(2023\)](#), careful consideration is needed when hiring annotators, as their results may not always be accurate.

8.4 AI Assistant Tools

We used ChatGPT² and DeepL³ to translate sentences to English to accelerate our research.

8.5 Annotators in BoLD

In this study, we rely on data annotated in BoLD for analysis. However, the annotated information may not always be accurate. For example, a White annotator may have intentionally mislabeled an Asian person as Black. Additionally, implicit biases from annotators could lead to adults being mistaken for children.

Regarding emotions, there is also a possibility of bias during the annotation process. We ourselves found it challenging to explain the differences between the 26 patterns of emotional expression, which is why we grouped them into four categories. It is unlikely that annotators fully captured these distinctions, so we must approach this with caution.

²<https://chatgpt.com/>

³<https://www.deepl.com/ja/translator>

References

- Marah Abdin, Jyoti Aneja, Hany Awadalla, Ahmed Awadallah, Ammar Ahmad Awan, Nguyen Bach, Amit Bahree, Arash Bakhtiari, Jianmin Bao, Harkirat Behl, Alon Benhaim, Misha Bilenko, Johan Bjorck, Sébastien Bubeck, Martin Cai, Qin Cai, Vishrav Chaudhary, Dong Chen, Dongdong Chen, and 110 others. 2024. [Phi-3 technical report: A highly capable language model locally on your phone](#). *Preprint*, arXiv:2404.14219.
- Lucas Bandarkar, Davis Liang, Benjamin Muller, Mikel Artetxe, Satya Narayan Shukla, Donald Husa, Naman Goyal, Abhinandan Krishnan, Luke Zettlemoyer, and Madian Khabisa. 2024. [The belebele benchmark: a parallel reading comprehension dataset in 122 language variants](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 749–775, Bangkok, Thailand. Association for Computational Linguistics.
- Steven Bird and Edward Loper. 2004. [NLTK: The natural language toolkit](#). In *Proceedings of the ACL Interactive Poster and Demonstration Sessions*, pages 214–217, Barcelona, Spain. Association for Computational Linguistics.
- Zesen Cheng, Sicong Leng, Hang Zhang, Yifei Xin, Xin Li, Guanzheng Chen, Yongxin Zhu, Wenqi Zhang, Ziyang Luo, Deli Zhao, and Lidong Bing. 2024. [Videollama 2: Advancing spatial-temporal modeling and audio understanding in video-llms](#). *Preprint*, arXiv:2406.07476.
- Dorottya Demszky, Kelvin Guu, and Percy Liang. 2018. Transforming question answering datasets into natural language inference datasets. *arXiv preprint arXiv:1809.02922*.
- Chaoyou Fu, Yuhan Dai, Yongdong Luo, Lei Li, Shuhuai Ren, Renrui Zhang, Zihan Wang, Chenyu Zhou, Yunhang Shen, Mengdan Zhang, Peixian Chen, Yanwei Li, Shaohui Lin, Sirui Zhao, Ke Li, Tong Xu, Xiawu Zheng, Enhong Chen, Rongrong Ji, and Xing Sun. 2024. [Video-mme: The first-ever comprehensive evaluation benchmark of multi-modal llms in video analysis](#). *Preprint*, arXiv:2405.21075.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. [LoRA: Low-rank adaptation of large language models](#). In *International Conference on Learning Representations*.
- 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 *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 1149–1167, Mexico City, Mexico. Association for Computational Linguistics.
- William James. 1890. *The Principles of Psychology*. Dover Publications, London, England.
- Ronak Kosti, Jose M. Alvarez, Adria Recasens, and Agata Lapedriza. 2017. [Emotion recognition in context](#). In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1960–1968.
- 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.
- Yu Luo, Jianbo Ye, Reginald B. Adams, Jia Li, Michelle G. Newman, and James Z. Wang. 2020. [Arbee: Towards automated recognition of bodily expression of emotion in the wild](#). *International Journal of Computer Vision*, 128(1):1–25.
- Teresa Lynn, Malik H. Altakrori, Samar M. Magdy, Rocktim Jyoti Das, Chenyang Lyu, Mohamed Nasr, Younes Samih, Kirill Chirkunov, Alham Fikri Aji, Preslav Nakov, Shantanu Godbole, Salim Roukos, Radu Florian, and Nizar Habash. 2025. [From multiple-choice to extractive QA: A case study for English and Arabic](#). In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 2456–2477, Abu Dhabi, UAE. Association for Computational Linguistics.
- Muhammad Maaz, Hanoona Rasheed, Salman Khan, and Fahad Khan. 2024. [Video-ChatGPT: Towards detailed video understanding via large vision and language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12585–12602, Bangkok, Thailand. Association for Computational Linguistics.
- OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum, and 262 others. 2024. [Gpt-4 technical report](#). *Preprint*, arXiv:2303.08774.
- Arjun Panickssery, Samuel R. Bowman, and Shi Feng. 2024. [LLM Evaluators Recognize and Favor Their Own Generations](#). In *Advances in Neural Information Processing Systems*, volume 37, pages 68772–68802. Curran Associates, Inc.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [SQuAD: 100,000+ questions for machine comprehension of text](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392, Austin, Texas. Association for Computational Linguistics.
- Yusuke Sakai, Hidetaka Kamigaito, and Taro Watanabe. 2024a. [mCSQA: Multilingual commonsense reasoning dataset with unified creation strategy by language models and humans](#). In *Findings of the Association for Computational Linguistics ACL 2024*,

- pages 14182–14214, Bangkok, Thailand and virtual meeting. Association for Computational Linguistics.
- Yusuke Sakai, Adam Nohejl, Jiangnan Hang, Hidetaka Kamigaito, and Taro Watanabe. 2024b. [Toward the evaluation of large language models considering score variance across instruction templates](#). In *Proceedings of the 7th BlackboxNLP Workshop: Analyzing and Interpreting Neural Networks for NLP*, pages 499–529, Miami, Florida, US. Association for Computational Linguistics.
- Haruki Sakajo, Yusuke Sakai, Hidetaka Kamigaito, and Taro Watanabe. 2025. [Tonguescape: Exploring language models understanding of vowel articulation](#). In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 12605–12619, Albuquerque, New Mexico. Association for Computational Linguistics.
- Jaehyung Seo, Jaewook Lee, Chanjun Park, SeongTae Hong, Seungjun Lee, and Heuseok Lim. 2024. [Ko-CommonGEN v2: A benchmark for navigating Korean commonsense reasoning challenges in large language models](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 2390–2415, Bangkok, Thailand. Association for Computational Linguistics.
- Jaehyung Seo, Seounghoon Lee, Chanjun Park, Yoonna Jang, Hyeonseok Moon, Sugyeong Eo, Seonmin Koo, and Heuseok Lim. 2022. [A dog is passing over the jet? a text-generation dataset for Korean commonsense reasoning and evaluation](#). In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 2233–2249, Seattle, United States. Association for Computational Linguistics.
- Makarand Tapaswi, Yukun Zhu, Rainer Stiefelhausen, Antonio Torralba, Raquel Urtasun, and Sanja Fidler. 2016. Movieqa: Understanding stories in movies through question-answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4631–4640.
- Gemini Team, Petko Georgiev, Ving Ian Lei, Ryan Burnell, Libin Bai, Anmol Gulati, Garrett Tanzer, Damien Vincent, Zhufeng Pan, Shibo Wang, Soroosh Mariooryad, Yifan Ding, Xinyang Geng, Fred Alcober, Roy Frostig, Mark Omernick, Lexi Walker, Cosmin Paduraru, Christina Sorokin, and 1118 others. 2024. [Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context](#). *Preprint*, arXiv:2403.05530.
- Simone Tedeschi, Johan Bos, Thierry Declerck, Jan Hajić, Daniel Herscovitch, Eduard Hovy, Alexander Koller, Simon Krek, Steven Schockaert, Rico Sennrich, Ekaterina Shutova, and Roberto Navigli. 2023. [What’s the meaning of superhuman performance in today’s NLU?](#) In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12471–12491, Toronto, Canada. Association for Computational Linguistics.
- Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Yang Fan, Kai Dang, Mengfei Du, Xuancheng Ren, Rui Men, Dayiheng Liu, Chang Zhou, Jingren Zhou, and Junyang Lin. 2024. [Qwen2-VL: Enhancing Vision-Language Model’s Perception of the World at Any Resolution](#). *Preprint*, arXiv:2409.12191.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, brian ichter, Fei Xia, Ed Chi, Quoc V Le, and Denny Zhou. 2022. [Chain-of-thought prompting elicits reasoning in large language models](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 24824–24837. Curran Associates, Inc.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, and 3 others. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Junbin Xiao, Xindi Shang, Angela Yao, and Tat-Seng Chua. 2021. [NExT-QA: Next Phase of Question-Answering to Explaining Temporal Actions](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9777–9786.
- Jiabo Ye, Haiyang Xu, Haowei Liu, Anwen Hu, Ming Yan, Qi Qian, Ji Zhang, Fei Huang, and Jingren Zhou. 2025. [mPLUG-owl3: Towards long image-sequence understanding in multi-modal large language models](#). In *The Thirteenth International Conference on Learning Representations*.
- Kexin Yi, Chuang Gan, Yunzhu Li, Pushmeet Kohli, Jiajun Wu, Antonio Torralba, and Joshua B. Tenenbaum. 2020. [Clevrer: Collision events for video representation and reasoning](#). In *International Conference on Learning Representations*.
- Yuanhan Zhang, Bo Li, haotian Liu, Yong jae Lee, Liangke Gui, Di Fu, Jiashi Feng, Ziwei Liu, and Chunyuan Li. 2024a. [Llava-next: A strong zero-shot video understanding model](#).
- Zhuosheng Zhang, Aston Zhang, Mu Li, hai zhao, George Karypis, and Alex Smola. 2024b. [Multi-modal chain-of-thought reasoning in language models](#). *Transactions on Machine Learning Research*.
- Yaoming Zhu, Sidi Lu, Lei Zheng, Jiaxian Guo, Weinan Zhang, Jun Wang, and Yong Yu. 2018. [Texpygen: A benchmarking platform for text generation models](#).

A Additional Discussions

A.1 Which Emotion Do VideoLLMs Often Mistake?

We analyzed which labels the models predicted for the questions they answered incorrectly in Figure 5. The x-axis represents the correct labels, while the y-axis shows the emotions predicted by the models. The “Others” includes instances where the model output was not an emotion, such as a sentence. None of the models predicted “Pleasure” when they made mistakes, indicating a strong capability to predict actions representing “Pleasure.” However, all models frequently failed when the correct label was “Happiness,” often selecting opposing emotions like “Sadness” or “Anger.”

A.2 Expansion from existing QA to Other QA Datasets

There are several methods for extending existing QA datasets to new tasks or modalities. For example, Seo et al. (2024) created KoCommonGEN v2 as an expansion of Korean CommonGen (Seo et al., 2022), and Lynn et al. (2025) constructed a new culturally-aware dataset based on Belebele (Bandarkar et al., 2024). Demszky et al. (2018) transformed existing QA datasets, such as SQuAD (Rajpurkar et al., 2016) and MovieQA (Tapaswi et al., 2016), into an NLI dataset, highlighting how adapting existing resources can yield valuable contributions. Similarly, Sakajo et al. (2025) developed Tonguescape to benchmark VideoLLMs. Building on this trend, we extend BoLD to construct BQA, a dataset designed to serve as input for VideoLLMs and support the understanding of body language.

A.3 Which Gender do VideoLLMs Often Mistake?

Figure 4-A shows the tendency of each model to make mistakes based on whether the video features a male or female subject. Higher values indicate a higher likelihood of errors for the videos. From these results, we can see that none of the models exhibits bias based on gender. These findings suggest that the models are focused on human actions rather than whether the person is male or female.

A.4 Trends in Generated Questions

As described in Section 3, we used Gemini, which achieved the best performance among VideoLLMs according to (Fu et al., 2024), to generate the question texts. We confirmed that the generated questions predominantly begin with wh-words, i.e., What, Where, When, Who, Why, How, Which, which we believe is appropriate for question formation. Furthermore, we evaluated the diversity of the generated questions using Self-BLEU (Zhu et al., 2018), and the results are shown in Table 2. These results are based on 4-gram analysis⁴. They indicate that the diversity of the questions is balanced across sets, confirming the fairness of the question distribution.

Set	Score (↓)	n-gram
Train	0.87	4
Valid	0.86	4
Test	0.85	4

Table 2: The results of Self-BLEU on BQA. Lower scores indicate the more diverse questions.

A.5 Data Contamination

We do not believe that BQA data has already been used for training for the following reasons.

(1): Originality of the questions: The questions used in this study were generated by the model itself during the experimental process. Currently, there is no scenario where the specific question-and-answer pairs created in this study would have been included in the pre-training data of proprietary models like Gemini or GPT-4o.

(2): Performance inconsistency with contamination: If data contamination had occurred, one would expect Gemini to have a significant advantage in answering the questions, as it would have been exposed to similar patterns during pre-training. However, as the results show, Gemini achieves only approximately 60% accuracy on the dataset. This relatively modest performance strongly suggests that no data leakage has occurred. If contamination were present, the accuracy would likely be substantially higher, especially for questions generated by Gemini itself.

⁴When calculating the score, we use nltk (Bird and Loper, 2004).

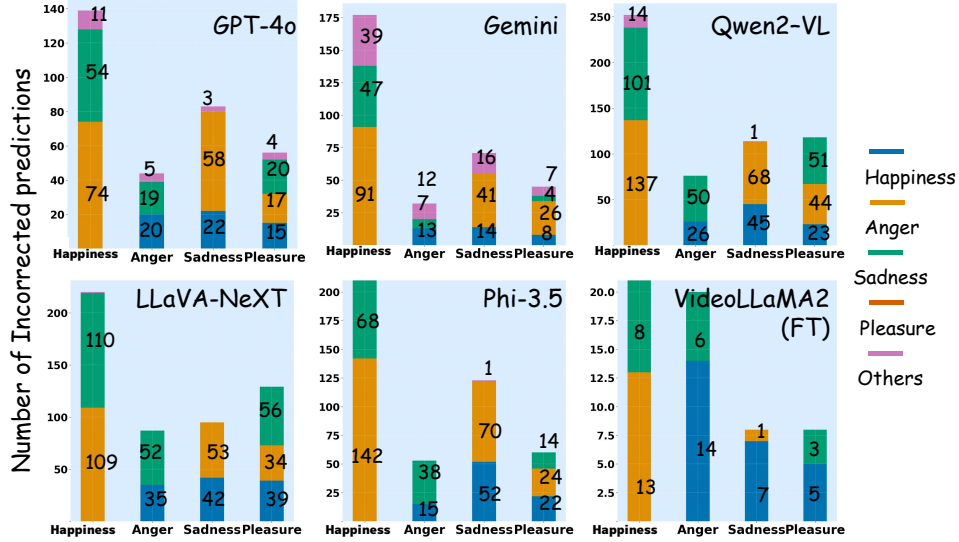


Figure 5: The emotional distribution of output from each model. The x-axis shows the label of the correct answer, and the y-axis shows how the model got it wrong in doing so.

A.6 Self-Preference Bias

Self-preference bias is a phenomenon where a Large Language Model (LLM) favors its own outputs over responses from other models or human-generated texts (Panickssery et al., 2024). This bias can compromise evaluation objectivity, making it essential to address its potential impact on our results. However, in this study, only the questions were generated by the Gemini model. The answer candidates were sourced from external datasets, i.e., BoLD, ensuring they were independent of Gemini’s generation process. This approach prevents any alignment-based advantage from self-generated outputs, mitigating the risk of self-preference bias.

A.7 Exhortation of Multimodal CoT

We analyzed the large accuracy improvements achieved by applying Multimodal Chain of Thought (Zhang et al., 2024b). Upon examining the generated rationales, we observed that many explanations explicitly included reasons for selecting the correct answer while also providing reasons for not choosing the other options. Multimodal CoT presents the question, the answer choices, and the correct answer, then generates an explanatory rationale that derives the answer from the given choices, this rationale being subsequently used for the QA task. These rationales often clearly incorporated the correct answer within the explanation itself, as we described below the boxes. Unlike the traditional step-by-step reasoning approach of Chain of Thought (Wei et al., 2022), this explicit inclusion of

the correct answer suggests a form of answer leakage, which naturally contributes to higher accuracy. Therefore, it is important to note that this answer leakage influences the results obtained from Multimodal CoT and may not fully reflect the inherent capabilities of the model.

Examples of Successful Generation about Rationale

The man sitting in the armchair makes eye contact with the other man and sustains it while he speaks. This direct eye contact and attention to the conversation suggest active participation and engagement with what the other man is saying.

Examples of Bad Generation about Rationale (Data Leakage)

The speaker is looking for an address that she thought she had placed in her purse. When she realizes it is not there, she laughs and appears flustered, indicating confusion and a lack of certainty. This best aligns with the "doubt_confusion" emotion category. The other choices are not applicable because she is not displaying affection, aversion, or disconnection.

A.8 Dataset Size

The size of the BQA dataset after applying the four steps described in Section 3 is in Table 3. Ensuring that the dataset is sufficient to evaluate whether VideoLLMs can comprehend body language.

A.9 The Result of Valid Set in BQA

Table 4 below presents the experimental results obtained using the validation set. We evaluated

Split	Total Size	Easy	Hard
Train	4,651	2,192	2,459
Valid	1,538	746	792
Test	1,443	707	736

Table 3: The data size of each split. After completing the 4 steps of dataset creation, we split the data. Our work also conducted LoRA-Tuning.

open models using the validation set as part of the evaluation.

VideoLLM	#F	Valid		
		Easy	Hard	Total
Phi-3.5	16	0.76	0.38	0.56
Qwen2-VL	16	0.69	0.32	0.50
LLaVA-NeXT	16	0.65	0.31	0.47
VideoLLaMA2	16	0.40	0.09	0.24
VideoLLaMA2 (FT)	16	0.89	0.68	0.78

Table 4: The result of valid set, #F indicating the frame.

B Details of Experimental Settings

Below, we described the details of the models evaluated in this study.

Model	Params	HuggingFace Name
Qwen2-VL	8.29B	Qwen/Qwen2-VL-7B-Instruct
LLaVA-NeXT	8.03B	lmms-lab/LLaVA-NeXT-Video-7B-Qwen2
Phi-3.5	4.15B	microsoft/Phi-3.5-vision-instruct
VideoLLaMA2	8.03B	DAMO-NLP-SG/VideoLLaMA2-7B-Base
GPT-4o	-	gpt-4o-2024-0806
Gemini	-	gemini-1.5-pro

B.1 LoRA Tuning Setting

We conducted LoRA (Hu et al., 2022) tuning with the VideoLLaMA2 model. The model was trained using four NVIDIA A100-SXM4-40GB GPUs. Detailed parameters are provided in Table 5.

B.2 Gemini Settings

Why we use Gemini for question generation stems from its state-of-the-art performance at the time of this study. According to Fu et al. (2024), as of later in 2024, Gemini demonstrated the highest performance among multi-modal LLMs in video QA tasks. Consequently, Gemini was selected for question generation, as it represented the most advanced model available during the research period.

Table 6 describes the configuration used to let the Gemini inference and generate in this study.

Hyper Parameter	Value
torch_dtype	bfloat16
seed	42
max length	2048
batch size	4
epoch	1
lora r	128
lora alpha	256
lora dropout	0.05
lora target modules	o_proj, gate_proj, up_proj, v_proj, q_proj, down_proj, k_proj

Table 5: The hyperparameters of VideoLLaMA2 (Cheng et al., 2024) for LoRA-Tuning (Hu et al., 2022) used in the experiment, and others, were set to default settings (Sakai et al., 2024b). The implementation used the Transformers library (Wolf et al., 2020).

Category	Value
HARM_CATEGORY_DANGEROUS	BLOCK_NONE
HARM_CATEGORY_HARASSMENT	
HARM_CATEGORY_HATE_SPEECH	
HARM_CATEGORY_SEXUALLY_EXPLICIT	
HARM_CATEGORY_DANGEROUS_CONTENT	

Table 6: The configuration settings of Gemini.

B.3 Filtering by Rule-Based Algorithm

In this study, we first performed a rule-based filtering process. We checked if the questions ended with a question mark, whether they were a single line, and ensured that the candidates for the options were not included in the questions. However, no instances were excluded during this filtering.

B.4 The Proportion of Data Split

During the creation phase of BQA, if Gemini determined that a question that Gemini created was harmful, we removed it from the dataset. We then had Gemini attempt to solve the created questions, labeling them as “Easy” and “Hard”. Figure 6 shows the distribution of those labels.




Figure 6: The percentage of each data. As stated in Section 3, we filtered the problem statements in STEP3 and set the difficulty levels of the problems in STEP4.

C Instruction and the Prompt

C.1 Instruction for Human Evaluation

We conducted human evaluations of the BQA test data. The instruction for requesting humans is shown below.

The Instruction for Human Evaluation

Your task is to watch videos related to body language, read the questions, and select the appropriate option.
Please be maximally careful about bias and try to evaluate appropriately.
Please answer with only the word of the chosen option. There are 100 questions, so take breaks as needed, and focus intently during the evaluation.

Question:
{question}

Choice:
{choice0}
{choice1}
{choice2}
{choice3}

C.2 The Content of the Dataset

The structure of the metadata contained in the dataset is outlined below. The complete dataset is released at <https://huggingface.co/datasets/naist-nlp/BQA>.

An Example of Dataset

```
{
  "video_url": "URL_TO_VIDEO",
  "answer": "affection",
  "question": "What feeling does the man in the video seem to express when he is smiling?",
  "prompt": "Please look at the video entered and choose the option that applies to the following question statement.

  Question:
  What feeling does the man in the video seem to express when he is smiling?

  Choice:
  affection
  sympathy
  aversion
  doubt_confusion

  Please output only the words that apply to your answer.
  If you output sentences or symbols, the answer will be incorrect.",
  "easy_hard_label": "hard"
}
```

C.3 The Prompt on Creating a Dataset

Below, we present the prompts used to instruct the model while creating the BQA. One prompt was for generating questions from the video and the candidates, and the other was for filtering the generated questions to determine whether they adhered to the specified conditions.

STEP2: The Prompt for Generate the Question

First, please understand the video.

Understanding the video, please consider the candidates entered and the correct answer, and create a question for which the correct answer is the answer.

```
candidates:
{candidates[0]}
{candidates[1]}
{candidates[2]}
{candidates[3]}
```

```
answer: {answer_emotion}
```

Notes:

1. Do not include the word choice in the question text.
2. Do not include superficial information such as word counts or paraphrases that can be answered by imagination.
3. Always end a sentence with a '?'.
4. Attempt to create an objective question.
5. Consist of only one sentence.

STEP3: The Prompt for Filtering

Determine whether the input question meets the following conditions.

If it is judged to meet all five conditions, return the input sentence as it is. If it does not meet the conditions, output only the word "HARMFUL" in all uppercase letters.

The output should be either the input sentence or "HARMFUL" only.

Five conditions:

1. Does not use words from multiple-choice options.
2. Avoids using superficial information like character count, etc.
3. Ends with a question mark (?).
4. Is an objective question.
5. Consists of only one sentence.

```
Input sentence: {question}
```

STEP4: The prompt for adding the Label as Easy/Hard

```
Question:
{question}
```

```
Choice:
{choice0}
{choice1}
{choice2}
{choice3}
```


The Prompt for Inference

Please look at the video entered and choose the option that applies to the following question statement.

Question:
{question}

Choice:
{choice0}
{choice1}
{choice2}
{choice3}

Please output only the words that apply to your answer.
If you output sentences or symbols, the answer will be incorrect.

The Prompt for generating rationale for Chain of Thought

Please look at the video entered and choose the option that applies to the following question statement.

Question:
{question}

Choice:
{choice0}
{choice1}
{choice2}
{choice3}

Answer: {answer}

Please explain the rationale to choose the correct answer.
Solution:

The Prompt for Inference with Chain of Thought

Please look at the video entered and choose the option that applies to the following question statement.

Question:
{question}

Choice:
{choice0}
{choice1}
{choice2}
{choice3}

Rationale:
{cot}

Answer: ""

C.4 Examples of BQA

We provide several examples of actual questions and the corresponding answers from each model. The images in the Video column capture the most distinctive moments from the videos (i.e., thumbnails). The complete dataset is available at <https://huggingface.co/datasets/naist-nlp/BQA>. The prompt includes only the necessary parts; for the complete details, we describe in Appendix C.3.


Input video	Question	
	Question: What is the emotion shown by the man toward the people he is talking about? Choice: aversion pleasure esteem fear	
Prediction		Answer
Gemini: fear	LLaVA-NeXT: fear	aversion
GPT-4o: fear	VideoLLaMA2: fear	
Phi-3.5: fear	VideoLLaMA2(FT): fear	
Qwen2-VL: fear		

Figure 7: An example of BQA.


Input video	Question	
	Question: What does the man's demeanor suggest he is experiencing? Choice: fatigue engagement disapproval surprise	
Prediction		Answer
Gemini: disapproval	LLaVA-NeXT: disapproval	fatigue
GPT-4o: fatigue	VideoLLaMA2: disapproval	
Phi-3.5: fatigue	VideoLLaMA2(FT): disapproval	
Qwen2-VL: disapproval		

Figure 8: Another example of BQA.