

# IFEval-Audio: Benchmarking Instruction-Following Capability in Audio-based Large Language Models

Yiming Gao<sup>\*,†</sup>, Bin Wang<sup>†</sup>, Chengwei Wei<sup>◊</sup>, Shuo Sun<sup>◊</sup>, AiTi Aw<sup>◊</sup>

<sup>\*</sup>Nanyang Technological University (NTU), Singapore

<sup>†</sup>MiroMind

<sup>◊</sup>Institute for Infocomm Research (I<sup>2</sup>R), A\*STAR, Singapore

GAOY0053@e.ntu.edu.sg

## Abstract

Large language models (LLMs) have demonstrated strong instruction-following capabilities in text-based tasks. However, this ability often deteriorates in multimodal models after alignment with non-text modalities such as images or audio. While several recent efforts have investigated instruction-following performance in text and vision-language models, instruction-following in audio-based LLMs remains largely unexplored. To bridge this gap, we introduce **IFEval-Audio**, a novel evaluation dataset designed to assess the ability to follow instructions in an audio LLM. IFEval-Audio contains 280 audio-instruction-answer triples across six diverse dimensions: **Content, Capitalization, Symbol, List Structure, Length, and Format**. Each example pairs an audio input with a text instruction, requiring the model to generate an output that follows a specified structure. We benchmark state-of-the-art audio LLMs on their ability to follow audio-involved instructions. The dataset is released publicly to support future research in this emerging area.<sup>1</sup>

## 1 Introduction

Evaluation of LLMs has become a cornerstone of NLP research, with significant efforts dedicated to assessing their capabilities across diverse tasks (Chang et al., 2023). Instruction-following, defined as a model’s ability to produce outputs in a specified format or structure as per the given instruction (Zhou et al., 2023), has seen notable progress in text and image modalities through benchmarks like IFEval (Zhou et al., 2023) and llava-Bench (Liu et al., 2023), as well as multimodal frameworks such as LAMM (Yin et al., 2023) and MM-IFEngine (Ding et al., 2025). However, audio-based instruction-following

<sup>\*</sup>Work done during internship at I<sup>2</sup>R and mentored by Bin Wang when he was there.

<sup>1</sup>Our evaluation code and datasets will be open-sourced at <https://github.com/AudioLLMs/AudioBench>

remains largely underexplored. Existing audio datasets, such as MMAU (Sakshi et al., 2025) and SIFT-50M (Pandey et al., 2025), focus on tasks like speech recognition or audio understanding, but they rarely evaluate how well audio models adhere to complex instructions (Moloo et al., 2018). Moreover, while benchmarks like MIA-Bench (Qian et al., 2025) assess multimodal instruction-following for images, the instruction-following evaluation framework for audio modality is missing, limiting the development of audio LLMs for applications.

To address this gap, we introduce **IFEval-Audio**, a novel dataset designed to evaluate audio-based LLMs’ instruction-following capabilities, building on the call for innovative benchmarks (Kiela et al., 2021). It comprises 280 audio-instruction-answer triples across six dimensions, including Content, Capitalization, Symbol, Length, List Structure, and Format Requirements. Each triple pairs an audio input with a text instruction, requiring the model to generate a text output that adheres to the instruction’s specified format, such as producing a list or a single-sentence response. Unlike datasets like IFEval and LLaVA-Bench, which primarily assess format adherence without evaluating the correctness of the content, IFEval-Audio also evaluates the semantic correctness of the outputs, offering a more comprehensive assessment of model’s instruction-following ability. We release IFEval-Audio publicly to advance audio instruction-following research.

## 2 Related Work

Instruction-following, the ability of LLMs to generate outputs in a specified format, has become a key focus in NLP research (Zhou et al., 2023). In text modalities, datasets like IFEval (Zhou et al., 2023) set rigorous standards by using verifiable prompts to ensure models produce outputs such as

numbered lists or single-sentence responses that meet structural requirements. Similarly, in image modalities, LLaVA (Liu et al., 2023) evaluates models’ ability to describe visual content in prescribed formats, reflecting the maturity of instruction-following evaluation in these domains. However, these datasets and related multimodal efforts primarily focus on format adherence, lacking a comprehensive assessment of content correctness.

In the audio domain, datasets such as MMAU (Sakshi et al., 2025) and SIFT-50M (Pandey et al., 2025) support tasks like speech recognition and audio understanding. Yet, they do not evaluate how well audio models follow complex instructions requiring structured text outputs, a gap underscored by critiques of current audio learning frameworks (Moloo et al., 2018). Recent work like Salmonn (Tang et al., 2024) improves instruction-following in audio-based LLMs, but lacks a dedicated benchmark for structured instruction adherence. Unlike prior instruction-following datasets that focus solely on format, our benchmark, IFEval-Audio, introduces a dual evaluation approach, combining rule-based scoring for format adherence with LLM-based assessment for semantic correctness across six dimensions. This absence of audio-specific benchmarks highlights the need for IFEval-Audio to advance audio-based LLM development.

### 3 IFEval-Audio

The IFEval-Audio dataset is designed to evaluate the instruction-following capabilities of audio-based LLMs, focusing on both format adherence and semantic correctness. It comprises 280 audio-instruction-answer triples distributed across six dimensions, testing a model’s ability to generate text outputs that meet specific structural and content requirements based on audio inputs.

#### 3.1 Evaluation Dimensions

IFEval-Audio spans six dimensions, each targeting a unique aspect of instruction-following capabilities.

1. **Content Requirements:** Instructions require including, removing, or replacing specific keywords (e.g., “Include the word ‘technology’ in your answer”).
2. **Capitalization Requirements:** Instructions specify capitalization rules, such as all uppercase, all lowercase, or capitalizing specific words (e.g.,

“Capitalize the first word of each sentence”).

3. **Symbol Requirements:** Instructions mandate adding symbols at the start or end, quoting the output, or removing symbols (e.g., “Enclose the response in square brackets”).
4. **List Structure Requirements:** Instructions dictate list formats, including Arabic numerals, Roman numerals, English letter numerals, or bullet points (e.g., “List items using Roman numerals”).
5. **Length Requirements:** Instructions set word count constraints, such as upper limits, lower limits, or ranges (e.g., “Respond within 10 words”).
6. **Format Requirements:** Instructions require specific output formats, such as JSON (e.g., “Provide the response in JSON format”).

#### 3.2 Audio Sources and Diversity

IFEval-Audio leverages audio from diverse datasets to ensure varied modalities and contexts, following the framework of AudioBench (Wang et al., 2025):

- **Speech Samples** are sourced from Spoken SQuAD (Li et al., 2018) (CC BY-SA 4.0), TED-LIUM 3 (Hernandez et al., 2018) (CC BY-NC-ND 3.0), and AudioBench (Wang et al., 2025) (inherited source licenses), providing conversational and question-answering contexts.
- **Music and Environmental Sound Samples** are drawn from Muchomusic (Weck et al., 2024) (CC BY-SA 4.0) and Wavcaps (Mei et al., 2024) (academic use only).

Instructions are designed to test one dimension each, varying in complexity. Answer formats range from lists to structured data like JSON. Speech audio (20–30 seconds) covers all six dimensions, while music/environmental sound audio (10 seconds) covers four dimensions (Content, Capitalization, Symbol, Format Requirements). Audio clips are used as-is from the source datasets without additional preprocessing for reproducibility.

#### 3.3 Dataset Structure

The IFEval-Audio dataset is structured around 280 audio-instruction-answer triples, each comprising an audio input derived from the specified datasets, a text instruction targeting one of the six dimensions, and an expected answer that adheres to the instruction’s format and content requirements. The dataset distribution includes 240 triples from speech audio, with 40 triples per dimension across the six categories, and 40 triples from music and envi-

ronmental sound audio, with 5 triples each per dimension for Content, Capitalization, Symbol, and Format Requirements (totaling 20 Music and 20 Environmental Sound). This organization ensures comprehensive coverage of instruction-following challenges across diverse audio modalities.

Figure 1 shows the diversity of instruction types within the dimensions.

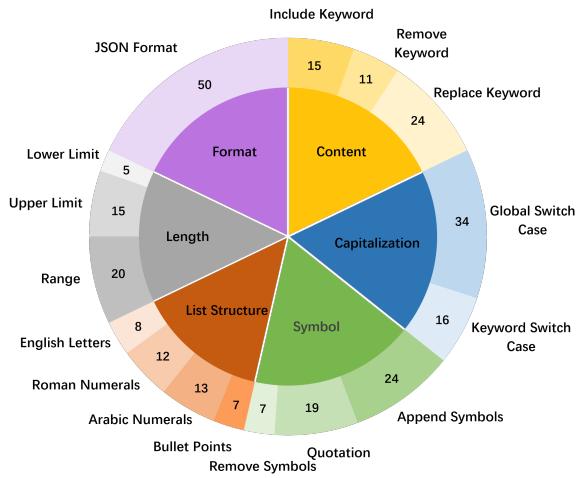


Figure 1: Breakdown of instruction subcategories within the dimensions, with approximately 25% Content, 20% Capitalization, 15% Symbol, 15% List Structure, 10% Length, and 15% Format requirements.

### 3.4 Annotation Process

The IFEval-Audio dataset was constructed through a meticulous process of curating audio clips from the specified sources, followed by manually designing instructions and answers to target each dimension comprehensively. This process involved selecting audio clips to represent diverse contexts, crafting instructions to test each of the six dimensions thoroughly, and creating expected answers that align with both the instructions and the audio inputs. By ensuring careful curation and design, this approach guarantees that IFEval-Audio serves as a robust benchmark for evaluating audio-based LLM instruction-following capabilities.

### 3.5 Evaluation Design

All triples in IFEval-Audio are designed to enable both rule-based and LLM-based evaluation for assessing instruction-following. A judge model evaluates each triple using the instruction, model output, and reference answer. Rule-based scoring applies strict rules to check format adherence (e.g., verifying a Roman numeral list), yielding a binary score (0/1) and a reason, which contributes to the

Instruction Following Rate (IFR), the proportion of outputs correctly adhering to the specified format. LLM-based evaluation, using Meta Llama 3 70B (AI@Meta, 2024), assesses semantic correctness with a custom prompt template (detailed in Appendix C). This produces a binary score (0/1) and a reason, which informs the Semantic Correctness Rate (SCR), the proportion of outputs semantically aligned with the reference answer. The Overall Success Rate (OSR) measures the proportion of outputs where both scores are 1 for a triple to contribute. Appendix C provides a sample evaluation illustrating this process.

## 4 Experiments

This section presents the evaluation of audio-based LLMs on IFEval-Audio, focusing on results, patterns, findings, and implications.

### 4.1 Models Tested

Six models were evaluated, with the following aliases used in Table 1:

- `cascade-whisper-large-v3-llama-3-8b-Ins` (`cascade`): A custom model combining Whisper (Radford et al., 2023) for speech recognition and Llama 3 8B Instruct (AI@Meta, 2024) for language generation, with additional fine-tuning.
- `gemini-1.5-flash` (`gemini-1.5`): Google’s Gemini 1.5 (Team et al., 2024), optimized for speed via the Google API.
- `gemini-2.0-flash-exp` (`gemini-2.0`): An experimental variant of Google’s Gemini models (Google, 2024), accessed via the Google API.
- `gpt-4o-audio-preview` (`gpt4-audio`): OpenAI’s GPT-4o (OpenAI, 2024) with audio capabilities, accessed via the OpenAI API.
- `Phi-4-multimodal-instruct` (`phi-4`): A multimodal instruction-following model based on Microsoft’s Phi series (Abdin et al., 2024).
- `Qwen2-Audio-7B-Instruct` (`qwen2`): Alibaba’s Qwen2-7B (Chu et al., 2024), fine-tuned for audio-based instruction following.

Each was assessed using the framework in Section 3.

### 4.2 Results, Analysis and Discussions

Table 1 highlights instruction-following capabilities across IFEval-Audio’s six dimensions. In **Content**, `gemini-1.5` leads with an IFR of 0.64, reflecting its ability to incorporate keywords as

Dimension	Metric	cascade	gemini-1.5	gemini-2.0	gpt4-audio	phi-4	qwen2
Content	SCR	0.56	<b>0.70</b>	0.62	<b>0.70</b>	0.44	0.60
	IFR	0.58	<b>0.64</b>	0.60	0.60	0.44	0.52
	OSR	0.42	<b>0.54</b>	0.46	0.52	0.24	0.36
Capitalization	SCR	0.48	<b>0.64</b>	0.60	0.62	0.54	0.56
	IFR	0.56	<b>0.80</b>	0.70	0.66	0.44	0.24
	OSR	0.32	<b>0.64</b>	0.52	0.56	0.26	0.14
Symbol	SCR	0.56	<b>0.62</b>	0.58	0.52	0.56	0.56
	IFR	0.56	0.28	<b>0.58</b>	0.52	0.28	0.16
	OSR	0.40	0.24	<b>0.42</b>	0.38	0.20	0.10
List Structure	SCR	0.73	0.70	0.63	<b>0.78</b>	0.63	0.53
	IFR	0.78	0.78	0.83	<b>0.93</b>	0.90	0.50
	OSR	0.60	0.55	0.55	<b>0.75</b>	0.58	0.30
Length	SCR	0.45	<b>0.65</b>	0.55	0.60	0.43	0.43
	IFR	0.28	<b>0.50</b>	0.38	0.45	0.28	0.20
	OSR	0.25	<b>0.50</b>	0.38	0.45	0.25	0.18
Format	SCR	0.52	0.58	0.64	<b>0.70</b>	0.60	0.46
	IFR	0.76	0.92	<b>0.94</b>	0.92	0.88	0.22
	OSR	0.48	0.52	0.58	<b>0.66</b>	0.56	0.08
<b>Overall Instruction-Following Rate</b>		0.59	0.65	0.67	<b>0.68</b>	0.53	0.30

Table 1: Detailed performance across dimensions: Semantic Correctness Rate (SCR), Instruction Following Rate (IFR), and Overall Success Rate (OSR); Overall Instruction-Following Rate included at bottom.

instructed, likely due to extensive multimodal training inferred from Google’s data integration efforts. Conversely, phi-4’s low IFR (0.44) suggests struggles in following content instructions, possibly due to limited audio training, as its documentation emphasizes text efficiency. **Capitalization** shows gemini-1.5’s high IFR (0.80), indicating strong adherence to formatting rules, a strength in proprietary models, while qwen2’s IFR (0.24) reflects challenges despite its enhanced instruction-following design, possibly due to its focus on natural language prompts over specific formatting tasks. **Symbol** reveals varied instruction-following, with gemini-2.0’s IFR (0.58) suggesting experimental optimizations, but qwen2’s low IFR (0.16) indicates limited symbol training. **List Structure** favors gpt4-audio (IFR 0.93), excelling in structured instruction adherence, likely from structured text training, while qwen2’s IFR (0.50) shows moderate performance. **Length** challenges instruction following, with gemini-1.5’s IFR (0.50) reflecting training on length-constrained tasks, while qwen2’s IFR (0.20) suggests difficulties in constraint handling. **Format** sees gpt4-audio’s high IFR (0.92), likely from code training enabling structured output adherence, while qwen2’s IFR (0.22) highlights struggles with complex formats, as it outputs did

not use double quotes " as standard JSON format.

IFEval-Audio underscores instruction-following gaps, particularly in Capitalization (IFR 0.24–0.80) and Format (IFR 0.22–0.94), urging improved multimodal alignment. gpt4-audio (overall instruction-following rate 0.68) and gemini-2.0 (0.67) excel with proprietary resources, while qwen2 (0.30) and phi-4 (0.53) indicate open-sourced models need targeted enhancements. As a benchmark, IFEval-Audio promotes research into fine-tuning and audio-text integration. For instance, in Format, gpt4-audio’s IFR (0.92) far exceeds its SCR (0.70), while qwen2’s IFR (0.22) aligns closer to its SCR (0.46), illustrating a mismatch between instruction-following and understanding capabilities across models.

## 5 Conclusion

This study introduces IFEval-Audio, a novel dataset comprising 280 audio-instruction-answer triples across six dimensions, designed to evaluate the instruction-following capabilities of audio-based LLMs. By assessing both format adherence and semantic correctness, IFEval-Audio addresses a critical gap in multimodal NLP research, where audio instruction-following has been underexplored compared to text and image modal-

ties. Experimental results on six state-of-the-art models highlight performance variations, with gpt4-audio and gemini-1.5 leading, while exposing weaknesses in handling complex formats. The public release of IFEval-Audio aims to foster further innovation, encouraging the development of robust audio-language models.

## 6 Limitations

While IFEval-Audio provides a robust benchmark for audio instruction-following, it has notable limitations. The dataset contains 280 human-crafted triples, which may not fully capture the complexity of real-world audio scenarios. Additionally, the audio is English-only, limiting its applicability to multilingual contexts. The audio variety, though diverse across speech, music, and environmental sounds, is constrained by the sourced datasets (Spoken SQuAD, TED-LIUM 3, etc.), potentially missing broader acoustic diversity such as non-English speech or rare environmental sounds.

Evaluation methods also present constraints. The LLM-based semantic evaluation, relying on models like Meta Llama 3 70B (AI@Meta, 2024), introduces dependency on the judge model’s biases and limitations in understanding audio-derived context. Furthermore, the evaluation does not explicitly disentangle ASR or perception errors from instruction-following failures in SCR, potentially allowing transcription inaccuracies to dominate semantic assessments. In the future, we plan to expand the dataset’s scale, linguistic diversity, and improve evaluation robustness to better reflect real-world audio instruction-following challenges.

## References

Marah Abdin, Jyoti Aneja, Harkirat Behl, Sébastien Bubeck, Ronen Eldan, Suriya Gunasekar, Michael Harrison, Russell J. Hewett, Mojan Javaheripi, Piero Kauffmann, James R. Lee, Yin Tat Lee, Yuanzhi Li, Weishung Liu, Caio C. T. Mendes, Anh Nguyen, Eric Price, Gustavo de Rosa, Olli Saarikivi, and 8 others. 2024. Phi-4 technical report.

AI@Meta. 2024. [Llama 3 model card](#).

Yupeng Chang, Xu Wang, Jindong Wang, Yuan Wu, Linyi Yang, Kaijie Zhu, Hao Chen, Xiaoyuan Yi, Cunxiang Wang, Yidong Wang, Wei Ye, Yue Zhang, Yi Chang, Philip S Yu, Qiang Yang, and Xing Xie. 2023. [A survey on evaluation of large language models](#). *arXiv preprint arXiv:2307.03109*.

Yunfei Chu, Jin Xu, Qian Yang, Haojie Wei, Xipin Wei, Zhifang Guo, Yichong Leng, Yuanjun Lv, Jinzheng He, Junyang Lin, Chang Zhou, and Jingren Zhou. 2024. Qwen2-audio technical report. *arXiv preprint arXiv:2407.10759*.

Shengyuan Ding, Shenxi Wu, Xiangyu Zhao, Yuhang Zang, Haodong Duan, Xiaoyi Dong, Pan Zhang, Yuhang Cao, Dahua Lin, and Jiaqi Wang. 2025. [Mmifengine: Towards multimodal instruction following](#). *arXiv preprint arXiv:2504.07957*.

Google. 2024. [Google introduces gemini 2.0: A new ai model for the agentic era](#).

François Hernandez, Vincent Nguyen, Sahar Ghannay, Natalia Tomashenko, and Yannick Estève. 2018. Ted-lium 3: Twice as much data and corpus repartition for experiments on speaker adaptation. In *Speech and Computer*, pages 198–208. Springer International Publishing.

Douwe Kiela, Max Bartolo, Yixin Nie, Divyansh Kaushik, Atticus Geiger, Zhengxuan Wu, Bertie Vidgen, Grusha Prasad, Amanpreet Singh, Pratik Ringshia, Zhiyi Ma, Tristan Thrush, Sebastian Riedel, Zeerak Waseem, Pontus Stenetorp, Robin Jia, Mohit Bansal, Christopher Potts, and Adina Williams. 2021. [Dynabench: Rethinking benchmarking in NLP](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4110–4124, Online. Association for Computational Linguistics.

Chia-Hsuan Li, Szu-Lin Wu, Chi-Liang Liu, and Hung-yi Lee. 2018. Spoken squad: A study of mitigating the impact of speech recognition errors on listening comprehension. *arXiv preprint arXiv:1804.00320*.

Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023. [Visual instruction tuning](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.

Xinhao Mei, Chutong Meng, Haohe Liu, Qiuqiang Kong, Tom Ko, Chengqi Zhao, Mark D. Plumbley, Yuexian Zou, and Wenwu Wang. 2024. [Wavcaps: A chatgpt-assisted weakly-labelled audio captioning dataset for audio-language multimodal research](#). *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 32.

Raj Kishen Moloo, Kavi Kumar Khedo, and Tadinda Venkata Prabhakar. 2018. Critical evaluation of existing audio learning systems using a proposed tol model. *Computers & Education*, 117:102–115.

OpenAI. 2024. [Audio quickstart: Using gpt-4o-audio-preview with the chat completions api](#).

Prabhat Pandey, Rupak Vignesh Swaminathan, K V Vijay Girish, Arunasish Sen, Jian Xie, Grant P. Strelak, and Andreas Schwarz. 2025. [Sift-50m: A large-scale multilingual dataset for speech instruction fine-tuning](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics.

Yusu Qian, Hanrong Ye, Jean-Philippe Fauconnier, Peter Grasch, Yinfei Yang, and Zhe Gan. 2025. [Mia-bench: Towards better instruction following evaluation of multimodal llms](#). In *The Thirteenth International Conference on Learning Representations*.

Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. 2023. Robust speech recognition via large-scale weak supervision. In *Proceedings of the 40th International Conference on Machine Learning*. JMLR.org.

S Sakshi, Utkarsh Tyagi, Sonal Kumar, Ashish Seth, Ramaseswaran Selvakumar, Oriol Nieto, Ramani Duraiswami, Sreyan Ghosh, and Dinesh Manocha. 2025. [MMAU: A massive multi-task audio understanding and reasoning benchmark](#). In *The Thirteenth International Conference on Learning Representations*.

Changli Tang, Wenyi Yu, Guangzhi Sun, Xianzhao Chen, Tian Tan, Wei Li, Lu Lu, Zejun MA, and Chao Zhang. 2024. [Salmonn: Towards generic hearing abilities for large language models](#). In *The Twelfth International Conference on Learning Representations*.

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](#). *arXiv preprint arXiv:2403.05530*.

Bin Wang, Xunlong Zou, Geyu Lin, Shuo Sun, Zhuban Liu, Wenyu Zhang, Zhengyuan Liu, AiTi Aw, and Nancy F Chen. 2025. [Audionbench: A universal benchmark for audio large language models](#). 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 4297–4316. Association for Computational Linguistics.

Benno Weck, Ilaria Manco, Emmanouil Benetos, Elio Quinton, George Fazekas, and Dmitry Bogdanov. 2024. [Muchomusic: Evaluating music understanding in multimodal audio-language models](#). *arXiv preprint arXiv:2408.01337*.

Zhenfei Yin, Jiong Wang, Jianjian Cao, Zhelun Shi, Dingning Liu, Mukai Li, Sheng Lu, Lei Bai, Xiaoshui Huang, Zhiyong Wang, Jing Shao, and Wanli Ouyang. 2023. [Lamm: Language-assisted multimodal instruction-tuning dataset, framework, and benchmark](#). In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.

Jeffrey Zhou, Tianjian Lu, Swaroop Mishra, Siddhartha Brahma, Sujoy Basu, Yi Luan, Denny Zhou, and Le Hou. 2023. [Instruction-following evaluation for large language models](#). *arXiv preprint arXiv:2311.07911*.

## A Additional Related Work

Broader efforts in LLM evaluation, such as those surveyed by Ye et al. ([Chang et al., 2023](#)), emphasize the need for innovative benchmarks to address emerging challenges in NLP. Similarly, Dynabench ([Kiela et al., 2021](#)) advocates for dynamic benchmarking to keep pace with evolving model capabilities. Other multimodal frameworks, including LAMM ([Yin et al., 2023](#)) and MM-IFEngine ([Ding et al., 2025](#)), extend instruction-following to diverse modalities, yet their primary focus remains on format adherence rather than content correctness. Multimodal benchmarks, such as MIA-Bench ([Qian et al., 2025](#)), have begun to explore instruction-following across modalities, but their inclusion of audio remains limited. These works underscore the importance of developing specialized datasets like IFEval-Audio to fill gaps in underrepresented domains like audio instruction-following.

In summary, while instruction-following has been extensively studied in text and image modalities, and audio datasets support related tasks, there is a clear gap in evaluating audio models’ ability to follow instructions with respect to both format adherence and semantic correctness. The detailed contributions of IFEval-Audio, as a novel dataset tailored to benchmark audio-based LLMs, offer a robust evaluation framework for advancing multimodal NLP research.

## B Appendix Figures

Figure 2 illustrates the distribution of triples across dimensions.

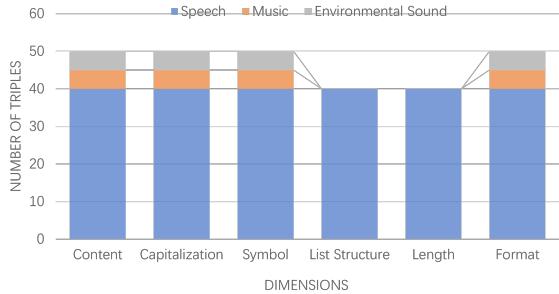


Figure 2: Distribution of triples across dimensions for Speech, Music, and Environmental Sound audio.

## C Appendix Examples

Below is a sample evaluation of **cascade-whisper-large-v3-llama-3-8b-Ins** from the IFEval-Audio

dataset, illustrating the evaluation process for a Content Requirements dimension triple.

### C.1 Sample Triple

The original audio input is sourced from a sports commentary clip: “with Rivera having been a linebacker with the Chicago Bears in Super Bowl XX, and Kubiak replacing Elway at the end of the Broncos’ defeats in Super Bowls XXI and XXIV. This will be the first Super Bowl in which both head coaches played in the game themselves.” Each triple in IFEval-Audio consists of an audio input, a text instruction, and an expected answer. The following triple targets the Content Requirements dimension:

```
{  
  "text": "What role did Rivera play in Super  
  Bowl XX, and which team did he represent?  
  Replace all occurrences of 'Super Bowl'  
  with 'Championship Game' in your output.",  
  
  "answer": "Rivera played as a linebacker  
  for the Chicago Bears in Championship Game  
  XX. His role on the team contributed to the  
  Bears' participation in that game, marking  
  his experience as a player in the  
  Championship Game.",  
  
  "dimension": "Content Requirements",  
  
  "rule_type": "Replace Keyword",  
  
  "rule_target": "Super Bowl",  
  
  "task_type": "Content Requirements",  
  
  "model_prediction": "Ron Rivera played as a  
  linebacker for the Chicago Bears in  
  Championship Game XX."  
}
```

### C.2 Judge LLM Prompt Template

The semantic correctness evaluation is performed using the Meta Llama 3 70B model as a judge. The following prompt template is used for each triple, with placeholders filled according to the triple evaluated.

```
[Reference Answer]  
{reference}  
[Model Answer]  
{prediction}  
[Question]  
{question}  
[Task]  
Rate the model's answer based on its alignment  
with the reference answer, focusing on the  
following two aspects:  
1. **Correctness**: Assess if the model's  
answer demonstrates the correct  
understanding and response based on the  
[Reference Answer].
```

Score 0: If the question is regarding to transcriptions, the model's answer is not the same as [Reference Answer]. If the question is not regarding to transcriptions, the model's answer does not accurately reflect the meaning or idea of [Reference Answer].

Score 1: If the question is regarding to transcriptions, the model's answer is exactly the same as [Reference Answer]. If the question is not regarding to transcriptions, the model's answer accurately reflects the meaning or idea of [Reference Answer].

Please provide two separate ratings:

1. \*\*Correctness Rating\*\*: (0 or 1)

Your final output should be exactly in this format, you can only modify contents inside brackets.

Correctness Rating: (int)

Explanation: (Provide a concise explanation for each rating. For \*\*Correctness\*\*, explain if the model's answer is correct and aligns with the reference. For \*\*Instruction-Following\*\*, explain how well the model adhered to the task instructions and any discrepancies.)

Note that the Instruction-Following Rating is computed rule-based via a separate function (format\_test), yielding 1 if the prediction adheres to the specified rule (e.g., format, symbols), else 0. The overall success is 1 only if both ratings are 1. The decoding parameters for the judge include max\_tokens=512, n=1, with deterministic settings. The complete evaluation code, including this prompt and the llama3\_70b\_as\_judge\_binary function, is released on GitHub.

### C.3 Evaluation of the Triple

The evaluation process involves both rule-based scoring for format adherence and LLM-based scoring for semantic correctness. The following is the evaluation of the prediction of the model against the reference answer:

```
{  
  "question": "What role did Rivera play in  
  Super Bowl XX, and which team did he  
  represent? Replace all occurrences of  
  'Super Bowl' with 'Championship Game' in  
  your output.",  
  "reference": "Rivera played as a linebacker  
  for the Chicago Bears in Championship Game  
  XX. His role on the team contributed to the  
  Bears' participation in that game, marking  
  his experience as a player in the  
  Championship Game.",  
  
  "model_prediction": "Ron Rivera played as a  
  linebacker for the Chicago Bears in  
  Championship Game XX.",
```

```
"judge_response": "Correctness Rating:  
1\nExplanation: The model's answer  
accurately reflects the meaning of the  
Reference Answer, correctly stating  
Rivera's role as a linebacker and his  
representation of the Chicago Bears in  
Championship Game XX. The answer is concise  
and directly answers the question.",  
  
"correctness_rating": 1,  
  
"instruction_following_rating": 1,  
  
"success": 1,  
  
"dimension": "Content Requirements"  
}
```

The format score is 1, as the model correctly replaced “Super Bowl” with “Championship Game” according to the instruction. The semantic score is also 1, as the model’s output aligns with the reference answer in meaning and factual accuracy, despite minor stylistic differences (e.g., inclusion of “Ron”).