

# Seeing isn't Hearing: Benchmarking Vision Language Models at Interpreting Spectrograms

Tyler Loakman<sup>1</sup>, Joseph James<sup>1</sup>, and Chenghua Lin<sup>2\*</sup>

<sup>1</sup>Department of Computer Science, University of Sheffield, UK

<sup>2</sup>Department of Computer Science, University of Manchester, UK

{tcloakman1, jhfjames1}@sheffield.ac.uk

chenghua.lin@manchester.ac.uk

## Abstract

With the rise of Large Language Models (LLMs) and their vision-enabled counterparts (VLMs), numerous works have investigated their capabilities in tasks that fuse the modalities of vision and language. In this work, we benchmark the extent to which VLMs are able to act as highly-trained phoneticians, interpreting spectrograms and waveforms of speech. To do this, we synthesise a novel dataset containing 4k+ English words spoken in isolation alongside stylistically consistent spectrogram and waveform figures. We test the ability of VLMs to understand these representations of speech through a multiple-choice task whereby models must predict the correct phonemic or graphemic transcription of a spoken word when presented amongst 3 distractor transcriptions that have been selected based on their phonemic edit distance to the ground truth. We observe that both zero-shot and finetuned models rarely perform above chance, demonstrating the requirement for specific parametric knowledge of how to interpret such figures, rather than paired samples alone.

## 1 Introduction

The ability of Large Language Models (LLMs) and Vision Language Models (VLMs) to reason about multimodal data has been studied extensively in recent years (Chia et al., 2024; Lin et al., 2023; Li and Zhang, 2023). One of the most productive domains for testing the unification between vision and language is that of the explanation and creation of data visualisations, such as graphs and figures (Zhang et al., 2024; Ge et al., 2024). For the most part, such works have tested models' abilities to understand axes and labels, follow trendlines, and reformulate visual content as text (or vice versa). However, to date, little work has been performed on the ability of VLMs to understand more complex visualisations, where

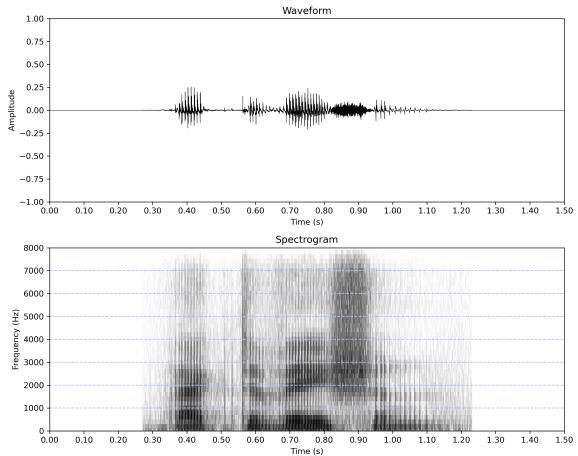


Figure 1: An example waveform (top) and spectrogram (bottom) of "activation" spoken by a text-to-speech model.

additional parametric knowledge must be combined with image inputs to reason about what is shown.<sup>1</sup>

To that end, we present the task of spectrogram and waveform interpretation, whereby models must predict spoken words from visual representations of speech. As seen in Figure 1, the x-axis of a spectrogram represents time, whilst the y-axis represents frequency, and the heat-map intensity represents the density of energy within a particular frequency range at a given time. On the other hand, a waveform also represents time on the x-axis, but displays the amplitude of a signal on the y-axis. Unlike more common visualisations (e.g., bar charts, pie charts, line graphs, etc.), in order to interpret such figures, models must reason about the visual content to establish what word is being spoken, using parametric knowledge regarding acoustic phonetics and the phonotactics of the target language to aid in reasoning about the observed signal characteristics such as vowel formants and fricative centre-of-gravity (Ladefoged, 1962).

The contributions of our work are as follows:

<sup>1</sup>Code and resources are available at <https://github.com/tylerL404/seeing-is-hearing>.

- We present a novel dataset of 4k+ words spoken in isolation using a range of synthetic voices, each with figures consisting of a spectrogram, a spectrogram with a waveform, and versions of the spectrograms/waveforms annotated with phone boundaries using forced-alignment.
- We perform what we believe to be the first analysis of whether VLMs are able to correctly interpret spectrograms and waveforms of speech in a phonologically-motivated multiple-choice paradigm.

## 2 Related Work

Data-to-text tasks have long been a staple of Natural Language Processing (NLP) research (Castro Ferreira et al., 2019). Traditionally, such approaches have relied on access to largely text-based representations, such as converting tables to markdown or HTML (He et al., 2023; Liu et al., 2022; Eisenschlos et al., 2020). However, the rise of VLMs has brought with it a range of benchmarks that concern image-based inputs (Zheng et al., 2024) such as charts (Islam et al., 2024; Zhou et al., 2023) and graphs (Ai et al., 2024), which have been fruitful in bearing benchmarks to assess the data-to-text abilities of VLMs (Zhu et al., 2025b; Islam et al., 2024; Ai et al., 2024). However, VLMs have been demonstrated to struggle with reasoning about inputs (Mukhopadhyay et al., 2024; Hou et al., 2024) and easily being subjected to training biases (Vo et al., 2025). Furthermore, recent works by Loakman et al. (2024) and Alper and Averbuch-Elor (2023) have investigated the meta-level abilities of VLMs to reason about language, investigating whether or not they learn correlations between vision and text that can be used to simulate an understanding of sound symbolism.

Building upon these existing works in reasoning about visual data representations and exploring the meta-level understanding of VLMs about language, we present what we believe to be the first benchmark and analysis of the ability of VLMs to understand spectrograms and waveforms of speech, combining visual acuity with acoustic phonetic knowledge.

## 3 Dataset Creation

### 3.1 Word List

As a result of the complex effects of coarticulation in standard connected speech, we present VLMs with single English words spoken in isolation. To reduce ambiguity, our word list is required to be representative of standard dictionary entries rather than the

acronyms, initialisms, and titles often found in word frequency lists. As a result, we use the Oxford 5000 word list,<sup>2</sup> consisting of words from A1-C1 difficulty on the Common European Framework of Reference for languages (CEFR) (Council of Europe, 2001). We perform additional filtering to remove words that have more than one possible pronunciation in the CMU Pronouncing Dictionary, such as homographs, removing the need to manually check which version of the word our text-to-speech (TTS) system produces (owing to words being presented in isolation, void of contextual cues for the TTS model).<sup>3</sup> After filtering, this results in 4068 words.

### 3.2 Speech Synthesis

Owing to the high levels inter/intra-speaker variation present in human speech and the range of possible acoustic environments for recording (Jessen, 2008), we create a novel dataset using Microsoft’s SpeechT5 trained for the task of TTS (Ao et al., 2022) to speak each word in isolation. We use synthetic voices with a General American (GenAm) (Wells, 1982) accent due to this enabling the use of the CMU Pronouncing Dictionary. A wider range of accents would require the manual creation of pronunciation dictionaries for each accent (intra-speaker variation notwithstanding), increasing the likelihood of phonetic realisations from the TTS model not accurately reflecting the transcriptions. We select 4 speaker embeddings from *cmu-artic-xvectors*<sup>4</sup> (2 male and 2 female perceived voices) for training, and 2 for testing (1 male and 1 female). We additionally add the General American voice from the Google Translate TTS API (via the gTTS library<sup>5</sup>) to the test set to mitigate biases from the characteristics of a particular TTS system. Further details are in Appendix A.1.

### 3.3 Spectrograms & Waveforms

As the image input to our VLMs, we provide a spectrogram of the target word being spoken by a TTS model, either on its own or with an accompanying waveform (akin to Figure 1). To ensure consistency across our generated figures, we pad or truncate the audio files to 1.5 seconds in length for all words, which we validated to be longer than the speech portion of any of our generated audio files, ensuring no words have phonemes

<sup>2</sup><https://www.oxfordlearnersdictionaries.com/about/wordlists/oxford3000-5000>

<sup>3</sup>Accessed via the *CMUDict* library: <https://pypi.org/project/cmudict/>

<sup>4</sup><https://huggingface.co/datasets/Matthijs/cmu-artic-xvectors>

<sup>5</sup><https://pypi.org/project/gTTS/>

cut off. All spectrograms are generated with Librosa (Brian McFee et al., 2015), using an  $nfft$  of 128, a hop-length of 22, and a dynamic range of 70dB, using a Hann window. We find these settings to ensure an adequate frequency resolution without unnecessarily detailed harmonic information. See Figure 1 for an example of one of the generated figures.

As an additional feature of our dataset, we provide spectrograms and waveforms with vertical lines at approximate phone boundaries using the Montreal Forced Aligner (McAuliffe et al., 2017) with the respective US English dictionary and pre-trained acoustic model. Owing to resource constraints, annotated figures are not used in the following experiments.

## 4 Methodology

### 4.1 Task Setup

To test the ability of VLMs to interpret spectrograms, we use a multiple-choice question (MCQ) paradigm where the ground truth is paired with 3 distractors which are required to be carefully selected (Goel et al., 2025; Alhazmi et al., 2024; Ding et al., 2024). Importantly, we select the distractors based on their phonemic edit distance (PED) from the ground truth using the *weighted feature edit distance* function from the Panphon python package (Mortensen et al., 2016), which is based on phonological distinctive features such as place and manner of articulation (Chomsky et al., 1991 - 1968). This way, we select distractors that are phonemically similar or distant to the ground truth in order to identify if the VLMs are making informed choices (even if incorrect) or guessing randomly, owing to the semi-ambiguous nature of interpreting speech directly from a spectrogram. Specifically, the distractors consist of the word with the lowest and highest PED from the ground truth (i.e., the most and least similar, respectively) that have less than 50% overlap in consecutive phonemes,<sup>6</sup> in addition to the word with the median PED. We additionally disallow distractors that start with the same character in their graphemic form to avoid overt biasing signals during training. The filtered list of 4068 words is broken into an 80-10-10 split for training, development, and testing, ensuring equal distribution of words based on length (in phonemes). See Appendix A.1 for further details on the split.

<sup>6</sup>In early testing, we observed that models were able to learn to reach 50% accuracy without being presented any input image. This was because the ground truth and the word with the lowest phonemic edit distance were frequently minimal pairs (e.g. "sheer" and "cheer"). It is for this reason that we introduce the overlap rule.

To mitigate ordering bias (Pezeshkpour and Hruschka, 2024; Wei et al., 2024), we generated four permutations of each word for both the training and development splits. The target word appeared in each of the four possible positions (index 0 to 3), whilst the distractor positions were randomised.

**Conditions** In total, we perform our tests under 4 conditions: graphemic (i.e., standard written words) or phonemic (i.e., International Phonetic Alphabet) options, and using spectrograms alone or spectrograms with waveforms as the image input.

### 4.2 Models

For this task, we selected a range of open-source VLMs from different model families. For zero-shot performance, we investigate the results from both small and large model variants, whilst for finetuning, we focus only on the small variants, owing to resource constraints. Specifically, we selected: Qwen 2.5-VL (7B and 32B) (Bai et al., 2025), Llava 1.6 (7B and 34B) (Liu et al., 2024), and InternVL 3 (8B and 38B) (Zhu et al., 2025a). See Appendix A.2 for training hyperparameters, specific model names, and the prompt given to VLMs.

## 5 Results

First, we present zero-shot performance results in Figure 2. Regarding accuracy, we consistently observe that all models perform around chance-level (25%) across all conditions. We observe no benefit from the additional information presented by the inclusion of the waveform (which would be useful for identifying speech sounds such as plosives), with performance even decreasing for Qwen 2.5-VL when waveforms are presented. Furthermore, regarding the phonemic and graphemic conditions, performance is generally higher when the options are presented as phonemes (Qwen 2.5-VL and InternVL-3 38B), but this does not hold for other models.

We hypothesise that the phonemic condition is more likely to activate knowledge directly related to acoustic phonetics, which would be relevant in the task of spectrogram/waveform understanding. On the other hand, regarding phonemic edit distance, the average distance of the selected option suggests frequent selection of low-distance distractors. Whilst a random selection is expected to have a phonemic distance of 17.2, most models perform slightly better than this, suggesting the tested models are making somewhat informed decisions, even if not selecting the correct answer.

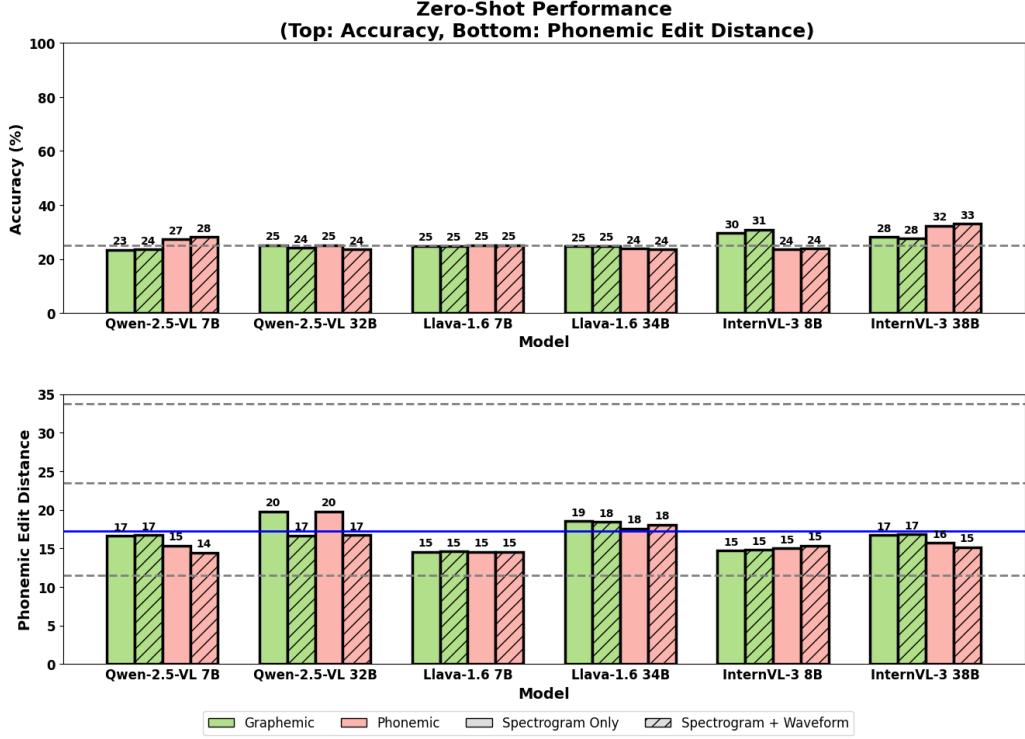


Figure 2: Zero-shot results in the multiple choice spectrogram interpretation task. *Graphemic* refers to questions where the options were presented in their standard written English form, whilst *Phonemic* refers to questions where the options were written in the International Phonetic Alphabet. *Spectrogram* and *Spectrogram + Waveform* refer to the type of figure presented to the VLM. *Accuracy* refers to the % of the time that the correct answer was selected, whilst *Phonemic Edit Distance* refers to the average distance of the selected option in comparison to the correct answer. The solid horizontal line in the Accuracy plot presents chance level agreement (25%), whilst the dashed lines in the phonemic distance plot relate to the expected phonemic distance for consistently selecting the 2<sup>nd</sup>, 3<sup>rd</sup> or 4<sup>th</sup> most similar option, whilst the solid blue line represents what is expected from random selection.

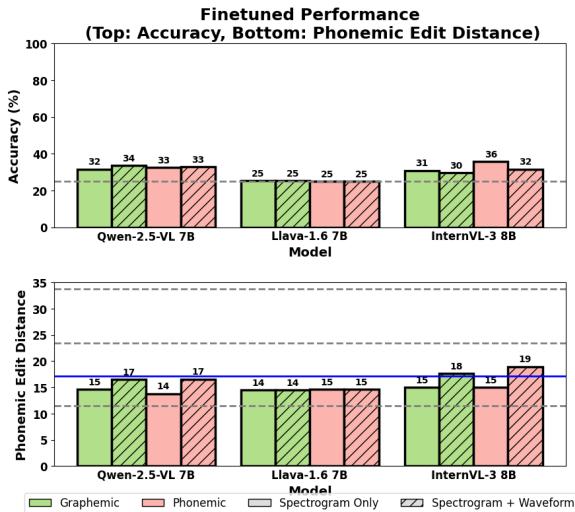


Figure 3: Finetuned results in the multiple choice spectrogram/waveform interpretation task. Please refer to Figure 2 for axis/condition information.

Next, we present the results of finetuning the smaller model variants in Figure 3. We observe

higher levels of performance than the zero-shot condition (particularly for InternVL-3). However, in our testing, we do not find a statistically significant difference between the performance of the finetuned models when tested with and without image inputs, suggesting that this performance increase is the result of learning the remaining biases in the dataset (owing to more similar words necessarily having similar phonemes). This highlights the difficulty of the task of spectrogram interpretation.

## 5.1 ASR & Trained Phonetician Comparison

As a traditional baseline for speech understanding, we present the results of an Automatic Speech Recognition (ASR) system on the test set, using the original audio from the TTS systems. Specifically, we use Nvidia’s *canary-qwen-2.5b* model, which currently ranks among the top-performing systems on the Open ASR leaderboard (i.e., lowest word error rate).<sup>7</sup> When accounting for casefolding and punctuation removal

<sup>7</sup>[https://huggingface.co/spaces/hf-audio/open\\_asr\\_leaderboard](https://huggingface.co/spaces/hf-audio/open_asr_leaderboard)

to determine exact matches, the ASR model achieves an accuracy of 87.56%, demonstrating that the underlying task is easily solvable when the acoustic information is given as input to purpose-built ASR system. This contrasts with the performance of the VLMs, which struggled when the same information was presented visually as spectrograms and waveforms.

Finally, we conducted human evaluation to establish an upper bound for the task of spectrogram/waveform interpretation using a trained phonetician, who is an author of this work. On a subset of 100 random examples from the test set on images consisting of a spectrogram and waveform (without boundary annotations). Importantly, a different author was responsible for generating the test set and presenting it, in order to avoid bias from pre-exposure. We observe an accuracy of 75.00% and an average phonetic edit distance of 5.17, highlighting the considerable gap between human and model performance, whilst demonstrating that there is sufficient information within the visualisations to achieve high accuracy.<sup>8</sup>

## 6 Conclusion

In this work, we presented the first analysis of the ability of VLMs to interpret speech from spectrograms and waveforms, using phonetically and phonologically motivated multiple-choice question approach. We observed that both zero-shot and finetuned models struggle to identify the correct answer, demonstrating the difficulty of the task and its potential as a benchmark assessment of the ability of VLMs to combine esoteric parametric knowledge with vision and language inputs.

## Limitations

Owing to computational limitations, we were not able to finetune the larger model variants presented in this work or benchmark large vision-enabled reasoning models such as OpenAI’s o1 and o3. Furthermore, our dataset consists of synthetic speech to reduce the effects of co-articulation and acoustic environments. However, synthetic speech is less variable than natural human speech and therefore presents an easier form of the task. Furthermore, we analyse models using zero-shot and finetuned approaches with a simple multiple-choice paradigm. However, we do not include any explicit knowledge within the prompt itself, which we leave to future work. We believe that future work may be best served by training models to

<sup>8</sup>We expect layperson accuracy on this task to be no better than chance.

recognise the characteristics of individual phonemes before learning entire words.

## Acknowledgments

Tyler Loakman and Joseph James are supported by the Centre for Doctoral Training in Speech and Language Technologies (SLT) and their Applications funded by UK Research and Innovation [grant number EP/S023062/1].

## References

Qihang Ai, Jiafan Li, Jincheng Dai, Jianwu Zhou, Lemaio Liu, Haiyun Jiang, and Shuming Shi. 2024. *Advancement in graph understanding: A multimodal benchmark and fine-tuning of vision-language models*. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7485–7501, Bangkok, Thailand. Association for Computational Linguistics.

Elaf Alhazmi, Quan Z Sheng, Wei Emma Zhang, Munazza Zaib, and Ahoud Alhazmi. 2024. Distractor generation in multiple-choice tasks: A survey of methods, datasets, and evaluation. *arXiv preprint arXiv:2402.01512*.

Morris Alper and Hadar Averbuch-Elor. 2023. *Kiki or bouba? sound symbolism in vision-and-language models*. In *Advances in Neural Information Processing Systems*, volume 36, pages 78347–78359. Curran Associates, Inc.

Junyi Ao, Rui Wang, Long Zhou, Chengyi Wang, Shuo Ren, Yu Wu, Shujie Liu, Tom Ko, Qing Li, Yu Zhang, Zhihua Wei, Yao Qian, Jinyu Li, and Furu Wei. 2022. *SpeechT5: Unified-modal encoder-decoder pre-training for spoken language processing*. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5723–5738, Dublin, Ireland. Association for Computational Linguistics.

Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibo Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu, Mingkun Yang, Zhaohai Li, Jianqiang Wan, Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, Jiabo Ye, Xi Zhang, Tianbao Xie, Zesen Cheng, Hang Zhang, Zhibo Yang, Haiyang Xu, and Junyang Lin. 2025. *Qwen2.5-v1 technical report*. *Preprint*, arXiv:2502.13923.

Brian McFee, Colin Raffel, Dawen Liang, Daniel P.W. Ellis, Matt McVicar, Eric Battenberg, and Oriol Nieto. 2015. *librosa: Audio and Music Signal Analysis in Python*. In *Proceedings of the 14th Python in Science Conference*, pages 18 – 24.

Thiago Castro Ferreira, Chris van der Lee, Emiel van Miltenburg, and Emiel Krahmer. 2019. *Neural data-to-text generation: A comparison between pipeline and end-to-end architectures*. 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 552–562, Hong Kong, China. Association for Computational Linguistics.

Yew Ken Chia, Vernon Toh, Deepanway Ghosal, Lidong Bing, and Soujanya Poria. 2024. **PuzzleVQA: Diagnosing multimodal reasoning challenges of language models with abstract visual patterns**. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 16259–16273, Bangkok, Thailand. Association for Computational Linguistics.

Noam. Chomsky, Morris. Halle, and American Council of Learned Societies. 1991 - 1968. *The sound pattern of English*, 1st mit press pbk. ed. edition. ACLS Humanities E-Book. MIT Press, Cambridge, Mass.

Council of Europe. 2001. *Common European Framework of Reference for Languages: Learning, Teaching, Assessment*. Cambridge University Press, Strasbourg.

Wenjian Ding, Yao Zhang, Jun Wang, Adam Jatowt, and Zhenglu Yang. 2024. **Can we learn question, answer, and distractors all from an image? a new task for multiple-choice visual question answering**. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 2852–2863, Torino, Italia. ELRA and ICCL.

Julian Eisenschlos, Syrine Krichene, and Thomas Müller. 2020. **Understanding tables with intermediate pre-training**. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 281–296, Online. Association for Computational Linguistics.

Yan Ge, Victor Junqiu Wei, Yuanfeng Song, Jason Chen Zhang, and Raymond Chi-Wing Wong. 2024. **Automatic data visualization generation from Chinese natural language questions**. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 1889–1898, Torino, Italia. ELRA and ICCL.

Rashi Goel, Harsh Nandwani, Eshaan Shah, Ashalatha Nayak, and Archana Praveen Kumar. 2025. Enhancing visual question answering for multiple choice questions. *IEEE Access*.

Xinyi He, Mengyu Zhou, Mingjie Zhou, Jialiang Xu, Xiao Lv, Tianle Li, Yijia Shao, Shi Han, Zejian Yuan, and Dongmei Zhang. 2023. **AnaMeta: A table understanding dataset of field metadata knowledge shared by multi-dimensional data analysis tasks**. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 9471–9492, Toronto, Canada. Association for Computational Linguistics.

Yifan Hou, Buse Giledereli, Yilei Tu, and Mrinmaya Sachan. 2024. Do vision-language models really understand visual language? *arXiv preprint arXiv:2410.00193*.

Mohammed Saidul Islam, Raian Rahman, Ahmed Masry, Md Tahmid Rahman Laskar, Mir Tafseer Nayeem, and Enamul Hoque. 2024. **Are large vision language models up to the challenge of chart comprehension and reasoning**. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 3334–3368, Miami, Florida, USA. Association for Computational Linguistics.

Michael Jessen. 2008. Forensic phonetics. *Language and linguistics compass*, 2(4):671–711.

Peter. Ladefoged. 1962. *Elements of acoustic phonetics*. A Phoenix book, P433. University of Chicago Press, Chicago.

Zhi Li and Yin Zhang. 2023. **Cultural concept adaptation on multimodal reasoning**. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 262–276, Singapore. Association for Computational Linguistics.

Hongzhan Lin, Ziyang Luo, Jing Ma, and Long Chen. 2023. **Beneath the surface: Unveiling harmful memes with multimodal reasoning distilled from large language models**. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 9114–9128, Singapore. Association for Computational Linguistics.

Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee. 2024. **Llava-next: Improved reasoning, ocr, and world knowledge**.

Ruixue Liu, Shaozu Yuan, Aijun Dai, Lei Shen, Tiangang Zhu, Meng Chen, and Xiaodong He. 2022. **Few-shot table understanding: A benchmark dataset and pre-training baseline**. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 3741–3752, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.

Tyler Loakman, Yucheng Li, and Chenghua Lin. 2024. **With ears to see and eyes to hear: Sound symbolism experiments with multimodal large language models**. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 2849–2867, Miami, Florida, USA. Association for Computational Linguistics.

Michael McAuliffe, Michaela Socolof, Sarah Mihuc, Michael Wagner, and Morgan Sonderegger. 2017. **Montreal Forced Aligner: Trainable Text-Speech Alignment Using Kaldi**. In *Proc. Interspeech 2017*, pages 498–502.

David R. Mortensen, Patrick Littell, Akash Bharadwaj, Kartik Goyal, Chris Dyer, and Lori Levin. 2016. **PanPhon: A resource for mapping IPA segments to articulatory feature vectors**. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 3475–3484, Osaka, Japan. The COLING 2016 Organizing Committee.

Srija Mukhopadhyay, Adnan Qidwai, Aparna Garimella, Pritika Ramu, Vivek Gupta, and Dan Roth. 2024. **Unraveling the truth: Do vlms really understand charts? a deep dive into consistency and robustness**. In *Findings*

*of the Association for Computational Linguistics: EMNLP 2024*, pages 16696–16717.

Pouya Pezeshkpour and Estevam Hruschka. 2024. **Large language models sensitivity to the order of options in multiple-choice questions.** In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 2006–2017, Mexico City, Mexico. Association for Computational Linguistics.

An Vo, Khai-Nguyen Nguyen, Mohammad Reza Taesiri, Vy Tuong Dang, Anh Totti Nguyen, and Daeyoung Kim. 2025. Vision language models are biased. *arXiv preprint arXiv:2505.23941*.

Sheng-Lun Wei, Cheng-Kuang Wu, Hen-Hsen Huang, and Hsin-Hsi Chen. 2024. **Unveiling selection biases: Exploring order and token sensitivity in large language models.** In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 5598–5621, Bangkok, Thailand. Association for Computational Linguistics.

J. C. Wells. 1982. *Accents of English: Beyond the British Isles*, volume 3 of *Accents of English*. Cambridge University Press, Cambridge.

Zhehao Zhang, Weicheng Ma, and Soroush Vosoughi. 2024. **Is GPT-4V (ision) all you need for automating academic data visualization? exploring vision-language models’ capability in reproducing academic charts.** In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 8271–8288, Miami, Florida, USA. Association for Computational Linguistics.

Mingyu Zheng, Xinwei Feng, Qingyi Si, Qiaoqiao She, Zheng Lin, Wenbin Jiang, and Weiping Wang. 2024. **Multimodal table understanding.** In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9102–9124, Bangkok, Thailand. Association for Computational Linguistics.

Mingyang Zhou, Yi Fung, Long Chen, Christopher Thomas, Heng Ji, and Shih-Fu Chang. 2023. **Enhanced chart understanding via visual language pre-training on plot table pairs.** In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 1314–1326, Toronto, Canada. Association for Computational Linguistics.

Jinguo Zhu, Weiyun Wang, Zhe Chen, Zhaoyang Liu, Shenglong Ye, Lixin Gu, Hao Tian, Yuchen Duan, Weijie Su, Jie Shao, Zhangwei Gao, Erfei Cui, Xuehui Wang, Yue Cao, Yangzhou Liu, Xingguang Wei, Hongjie Zhang, Haomin Wang, Weiye Xu, Hao Li, Jiahao Wang, Nianchen Deng, Songze Li, Yinan He, Tan Jiang, Jiapeng Luo, Yi Wang, Conghui He, Botian Shi, Xingcheng Zhang, Wenqi Shao, Junjun He, Yingtong Xiong, Wenwen Qu, Peng Sun, Penglong Jiao, Han Lv, Lijun Wu, Kaipeng Zhang, Huipeng Deng, Jiaye Ge, Kai Chen, Limin Wang, Min Dou, Lewei Lu, Xizhou Zhu, Tong Lu, Dahu Lin, Yu Qiao, Jifeng Dai, and Wenhui Wang. 2025a. **Internvl3: Exploring advanced training and test-time recipes for open-source multimodal models.** *Preprint, arXiv:2504.10479*.

Yingjie Zhu, Xuefeng Bai, Kehai Chen, Yang Xiang, Jun Yu, and Min Zhang. 2025b. **Benchmarking and improving large vision-language models for fundamental visual graph understanding and reasoning.** In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 30678–30701, Vienna, Austria. Association for Computational Linguistics.

## A Appendix

### A.1 Dataset Characteristics

Figure 4 presents the distribution of words by length (in phonemes) across the training, development and test sets, in addition to the distribution of specific phonemes.

### A.2 Implementation Details

**TTS Voices** From *cmu-arctic-xvectors*, we use speaker embeddings 69 (M), 750 (M), 2500 (F), and 3333 (F) for training, and 4444 (M) / 7500 (F) for testing. We additionally use Google Translate’s General American TTS voice (which is only available in a female-presenting voice) for testing.

**Model Selection** All of our selected models were taken from Hugging Face. Specifically, for Qwen 2.5-VL we used QWEN/QWEN2.5-VL-7B-INSTRUCT and QWEN/QWEN2.5-VL-32B-INSTRUCT, for Llava 1.6 (i.e., Llava NeXT) we used LLAVA-HF/LLAVA-V1.6-VICUNA-7B-HF and LLAVA-HF/LLAVA-V1.6-34B-HF, and for InternVL 3 we used OPENGVLAB/INTERNVL3-8B-HF and OPENGVLAB/INTERNVL3-38B-HF. For our ASR model we used NVIDIA/CANARY-QWEN-2.5B.

**Model Prompt** You are a Vision-Language Model specialized in phonetic interpretation of speech spectrograms. Your primary role is to act as a highly trained phonetician: given visual representations of spoken English words in the General American accent (spectrograms and/or waveforms), you must determine which graphemic or phonetic transcription correctly matches the spoken word. For each example, you will be shown: A spectrogram (and optionally a waveform) of a single English word spoken in isolation. Four candidate transcriptions labeled 0 to 3 (one correct, three phonetic distractors). Provide the final label as <label>label</label>, where label is the number (0-3) of the correct transcription. Just output the label and nothing else in the format: <label>label</label>

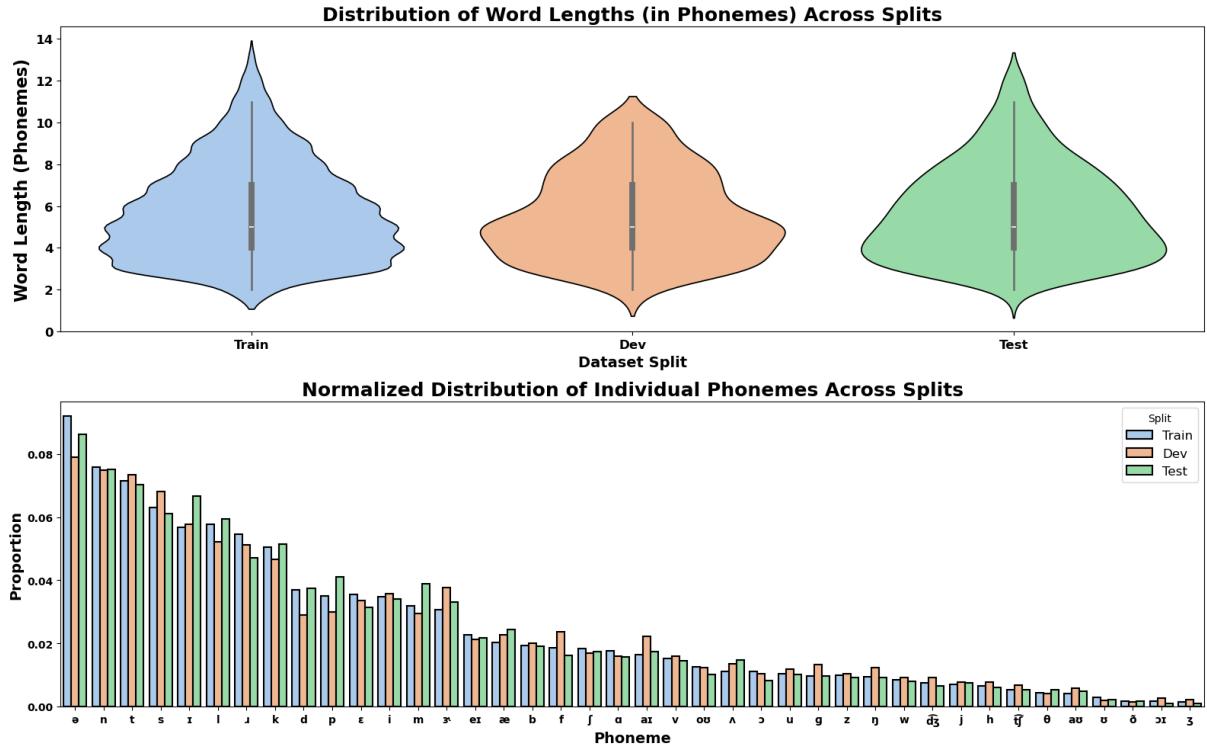


Figure 4: Distribution of word lengths (as determined via phoneme count) and individual phonemes across the training, development and test sets for finetuned VLMs.

**Training Details** We finetuned and performed inference for all models on a single A100 GPU. Hyperparameters used during training are presented in [Table 1](#).

Hyperparameter	Setting
Epochs	5
Batch Size	1
Gradient Accumulation Steps	8
Gradient Checkpointing	True
Optimiser	AdamW
Learning Rate (LR)	2e-4
LR Scheduler Type	constant
Weight Decay	0.01
Maximum Gradient Norm	1.0
Warmup Steps	150
Logging Steps	150
Evaluation Steps	150
Evaluation Strategy	steps
Save Steps	150
Early Stopping Callback	3

Table 1: Training hyperparameters.

## B Statistical Testing

We conducted Chi-squared tests for each condition to determine whether it had an effect on the output. The

test aggregated all models and compared the effect of the word type (graphemic v. phonemic) and input type (spectrogram v. spectrogram and waveform). We found no significant effects of any conditions at  $\alpha = 0.05$ .