

# Can LLMs Understand Unvoiced Speech? Exploring EMG-to-Text Conversion with LLMs

Payal Mohapatra<sup>\*†</sup> Akash Pandey<sup>\*</sup> Xiaoyuan Zhang<sup>\*</sup> Qi Zhu<sup>†</sup>  
Northwestern University, USA

## Abstract

Unvoiced electromyography (EMG) is an effective communication tool for individuals unable to produce vocal speech. However, most prior methods rely on paired voiced and unvoiced EMG signals, along with speech data, for unvoiced EMG-to-text conversion, which is not practical for these individuals. Given the rise of large language models (LLMs) in speech recognition, we explore their potential to understand unvoiced speech. To this end, we address the challenge of *learning from unvoiced EMG alone* and propose a novel EMG adaptor module that maps EMG features to an LLM’s input space, achieving an average word error rate of 0.49 on a closed-vocabulary unvoiced EMG-to-text task. Even with a conservative data availability of just six minutes, our approach improves performance over specialized models by nearly 20%. While LLMs have been shown to be extendable to new language modalities—such as audio—understanding articulatory biosignals, like unvoiced EMG, is more challenging. This work takes a crucial first step toward enabling LLMs to comprehend unvoiced speech using surface EMG.

## 1 Introduction and Related Works

Speech impairments affect around 4 million people in the U.S. alone (NICD, 2024). Silent speech interfaces (Zhang et al., 2021; Mohapatra et al., 2024a; Gonzalez-Lopez et al., 2020; Srivastava et al., 2024) have emerged as transformative solutions, enabling communication for individuals who cannot rely on spoken language. One such instrument is surface electromyography (EMG) (Schultz et al., 2017), which captures muscle activations crucial for speech production, even during unvoiced articulation. Inspired by the immense success of the speech recognition abilities of text-to-text large language models (LLMs) (Tang et al., 2023; Yu

et al., 2024; Dighe et al., 2024), we study an important question—*what is the potential of these LLMs to understand unvoiced speech?* Specifically, can they convert silent speech to text without access to audio or voiced versions of the EMG signals? This question is particularly relevant for assisting individuals who can no longer produce audible speech (Meltzner et al., 2017), where no corresponding voiced EMG or speech data exists. Additionally, given the highly personal nature of these signals (Diener et al., 2020a; Wand et al., 2009), it is crucial to develop methods that learn effectively from limited unvoiced EMG data.

Prior works (Jou et al., 2006; Meltzner et al., 2018; Schultz and Wand, 2010) on EMG-to-text conversion focused primarily on voiced signals. Other studies (Gaddy and Klein, 2021, 2020; Benster et al., 2024) on unvoiced EMG-to-text conversion leveraged auxiliary tasks to align unvoiced EMG with audio from voiced sessions or applied strategic transfer learning from voiced EMG-audio models, both of which rely on vocal data. However, we consider a scenario where no voiced signals are available for a speaker and explore a technique to communicate with LLMs, the modern workhorses for language understanding (Dubey et al., 2024; OpenAI, 2023). Recent research successfully expanded LLMs to other language modalities, such as speech (Tang et al., 2023; Yu et al., 2024) and silent video (Maaz et al., 2023; Yeo et al., 2024). A key approach involves adaptor modules—ranging from simple trainable linear layers (Ma et al., 2024) to dedicated projector networks (Kang et al., 2024) and explicit alignment strategies (Li et al., 2023; Tan et al., 2024)—to map new language modalities to LLMs’ input embedding space. While Benster et al. (2024) incorporated LLMs as a post-processing step after multimodal EMG model predictions, it largely remains unexplored whether effective unvoiced EMG-to-text conversion can be achieved by directly leveraging LLMs.

<sup>\*</sup>Equal contribution

<sup>†</sup>Corresponding authors:{payal.mohapatra, qzhu}@northwestern.edu

**Our Contributions.** We study a practical setting of converting unvoiced EMG to text without access to voiced EMG or audio, by expanding LLMs to understand this new language modality. We propose a novel trainable EMG adaptor module to map EMG features into the LLM’s input space. Our approach, focused on a closed vocabulary, demonstrates promising results, achieving an average word error rate of 0.49. With just six minutes of training data, LLMs outperform specialized EMG-to-text models by 20%. We analyze the EMG adaptor’s design, performance across varying data amounts and features, and broader challenges in learning from unvoiced EMG. Our work paves the way for integrating unvoiced EMG with large language models (LLMs), improving their text conversion accuracy and enabling non-vocal users to fully benefit from LLM-based assistants.

## 2 Approach

**Adaptor Network Design.** The input unvoiced EMG signals are represented as  $\mathbf{X}^E \in \mathbb{R}^{T \times C}$ , where  $C$  is the number of EMG channels with  $T$  discrete time steps. The EMG signals undergo standard minimal preprocessing, similar to past works (Gaddy and Klein, 2021, 2020). Since the original sampling rate is high ( $>800$  Hz), we leverage a temporal 1D convolutional layer with a stride of  $N$  ( $N = 6$  in our case) to facilitate downsampling to  $T/6$ . Similar to Gaddy and Klein (2021), we use residual blocks with 1D convolutional layers to extract EMG features, employing a stack of two residual blocks. Next, we explicitly facilitate the learning of sequential dependencies in the extracted features, finding that a bidirectional long-short-term memory (BiLSTM) network effectively captures complex temporal dependencies (further design choices for this sequential block are compared in Section 3.2). This is followed by another 1D convolutional layer with stride  $N = 2$  for further downsampling, resulting in the embedding  $\tilde{\mathbf{E}} \in \mathbb{R}^{(T/48) \times \tilde{F}}$ , which is then projected using fully connected linear layers to match the input embedding dimension of the LLM, generating the EMG embeddings  $\mathbf{E} \in \mathbb{R}^{\hat{T} \times F}$ , where  $T > \hat{T} \approx T/48$  and  $F \in \{4096^1, 3072^2\}$  in our case. We use GeLU activation function (Baevski et al., 2020).

**Data Preparation for Large Language Models.** Our EMG adaptor network is defined as

<sup>1</sup>LLaMA 2-7B

<sup>2</sup>LLaMA 3.2-3B

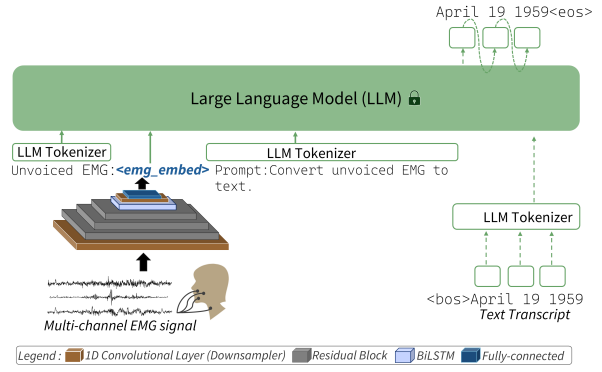


Figure 1: Our trainable EMG adaptor with frozen LLMs to transcribe text from only unvoiced EMG.

$\mathcal{G} : \mathbf{X}^E \rightarrow \mathbf{E}$ . We contextualize the embeddings by prepending them with a text identifier  $\mathbf{P}_1 = \text{Unvoiced EMG:}$ , and appending a prompt describing the task as  $\mathbf{P}_2 = \text{Prompt: Convert unvoiced EMG embeddings to text}$ . The LLM tokenizer converts the text identifier and prompt into text embeddings using the mapping  $\mathcal{M} : \mathbf{X}^P \rightarrow \mathbf{H}^P$ , where  $\mathbf{X}^P = [\mathbf{P}_1, \mathbf{P}_2]$ . To prepare the input for the LLM, the embeddings from the prompts are concatenated with the EMG embeddings.

**Training Framework.** For each unvoiced EMG signal, we have a corresponding text transcription  $\mathbf{X}^S$ . Following the recommendations of Gaddy and Klein (2021), we simplify our target by removing punctuation and converting all text to lowercase. We extract embeddings from the final LLM layer and compute the predicted logits  $z_{t_s, y'}$  for each vocabulary item  $y'$  at position  $t_s$ . The cross-entropy loss over time steps is given by:

$$\mathcal{L} = - \sum_{t_s=1}^{T_s} \sum_{y' \in \mathcal{V}} y'_{t_s} \log \frac{\exp(z_{t_s, y'} / \tau)}{\sum_{v \in \mathcal{V}} \exp(z_{t_s, v} / \tau)},$$

where  $\tau = 0.8$  is the temperature parameter and  $y'_{t_s}$  is the true class at  $t_s$ . We follow standard recommendations for fine-tuning LLMs, employing the AdamW optimizer (Loshchilov and Hutter, 2019) with a maximum learning rate of  $5 \times 10^{-5}$  and weight decay. During inference, we autoregressively generate (Dubey et al., 2024) the predicted target sequence with a beam-width (Freitag and Al-Onaizan, 2017) of 4. More implementation details are provided in Appendix B. Our codebase is release

## 3 Experimental Results and Discussion

**Datasets.** We primarily used the single-speaker, 8-channel, closed-vocabulary dataset from Gaddy

and Klein (2021), which comprises 67 words and approximately 26 minutes of unvoiced EMG data across 500 utterances. More details are provided in Appendix A. We accessed only the unvoiced EMG modality.

**Baselines and Experimental Setup.** We use the transducer model proposed by Gaddy and Klein (2021) as an application-specific baseline. Our analysis leverages two LLMs: Llama2-7B and Llama3-3B. We also fine-tune Llama3-3B using low-rank adaptation (LoRA) (Hu et al., 2022), training 0.13% of its parameters. We use three-fold validation and report word error rate (WER) statistics. Following standard practice, data is split 8:1:1 into training, validation, and test sets. All baselines are trained solely on unvoiced EMG. A minimal code implementation and sample predictions are included in the supplementary materials, with further details in Appendix B.

### 3.1 Key Findings

**LLMs boost closed-vocabulary unvoiced EMG-to-text conversion by 30% with minimal data processing.** As shown in Table 1, compared to the application-specific (App-Specific) model with transformers (54M trainable parameters), the proposed EMG adaptor (EMG-Ad) with frozen LLMs, using only 6M trainable parameters, achieves a 0.52 WER with raw EMG signals as input, outperforming the App-Specific model’s 0.75 average WER. While the closed-vocabulary setting poses a challenge due to limited training data, it supports the hypothesis that LLMs, through large-scale training, have likely learned universal language representations that help understand unvoiced EMG with limited datasets. Fine-tuning offered only a 17% improvement over the App-Specific baseline in this setting, possibly due to overparameterization.

**Handcrafted EMG features improve LLM performance for unvoiced EMG-to-text conversion in closed vocabulary.** Leveraging recommendations from previous works on temporal (Jou et al., 2006) and spectral (Gaddy and Klein, 2020) features, we extract 112 time-varying EMG features and use them as inputs to the baselines. As shown in Table 1, these handcrafted features consistently outperform raw features across both LLMs used as inputs to the EMG adaptor, showing an average improvement of 15%. However, for the App-Specific baseline, raw features perform better, similar to Gaddy and Klein (2021)’s findings.

Table 1: Comparison of App-Specific models and EMG adaptors (EMG-Ad) with frozen and fine-tuned LLMs. Frozen parts are shown in *green*, and the best performance in each setting is in **bold**. Lower WER is better.

	Model	WER
Raw EMG	App-Specific(Gaddy and Klein, 2021)	0.75 ± 0.06
	EMG-Ad + <i>Llama2-7B</i>	0.65 ± 0.01
	EMG-Ad + <i>Llama3-3B</i>	<b>0.52 ± 0.05</b>
	EMG-Ad + Fine-tuned Llama3-3B	0.62 ± 0.04
Hand- crafted Features	App-Specific(Gaddy and Klein, 2021)	0.84 ± 0.06
	EMG-Ad + <i>Llama2-7B</i>	<b>0.49 ± 0.06</b>
	EMG-Ad + <i>Llama3-3B</i>	<b>0.49 ± 0.04</b>
	EMG-Ad + Fine-tuned Llama3-3B	0.55 ± 0.02

**LLMs enable data-efficient learning for closed-vocabulary silent speech.** To further evaluate the effectiveness of LLMs in facilitating learning from a limited number of samples, we randomly subsampled the training data from approximately 26 minutes to 6 minutes, as illustrated in Figure 2. Although the WER increases with the reduced training set, our LLM-based approach still outperforms the App-Specific baseline by an average of 26%. Evidenced by prior work (Diener et al., 2020a) and our pilot study in Section 3.3, surface-EMG signals exhibit distinct person-specific phenotypes. Thus, learning from a limited of samples facilitates building personalized silent-speech interfacing models for LLM assistants.




Figure 2: Performance of EMG adaptor with Llama3-3B model vs. App-Specific model across training dataset sizes for unvoiced raw EMG-to-text conversion.

**Expanding LLMs to EMG is harder than audio.** To demonstrate the potential of additional language modalities in expanding text-based LLMs, we adopt Ma et al. (2024)’s strategy of incorporating a speech encoder—both an end-to-end trained speech encoder using mel-frequency spectrogram features and a pretrained encoder (Baevski et al., 2020)—with a linear-projection head into

an LLM to convert speech to text (more details in Appendix B.3). LLMs learn 33% better from audio even with this simple approach, highlighting the overall task complexity of unvoiced EMG-to-text, as shown in Figure 3.

### Augmenting training dataset with voiced EMG benefits specialized models more than LLMs.

In an augmented setting, where we trained using voiced and unvoiced raw EMG, the additional modality led to a 20% improvement in the App-Specific model (expected behavior as Gaddy and Klein (2021)), while it offered little benefit to the LLM-based approach in this closed vocabulary setting as shown in Figure 3. This suggests that more dedicated efforts in instruction tuning or explicit pairing of voiced and unvoiced EMG (Xu et al., 2022) may be needed to learn improved representations from LLMs. However, in this paper, our focus remains on converting unvoiced EMG to text.




Figure 3: Performance comparison in (1) expanding LLMs to the audio vs. EMG modality, and in (2) training our LLM-based approach and the specialized model using voiced vs. unvoiced EMG data.

### 3.2 Ablation Analysis

Table 2 summarizes key variants of the EMG adaptor network design with raw unvoiced EMG, omitting downsampling via high-stride 1D CNN layers. Unlike specialized EMG-to-text models (Gaddy and Klein, 2021) that benefit from transformers, we find LSTMs perform better in this setting (Mohapatra et al., 2023c). This may be due to the shorter sequence length (average of four words per utterance with more details in Appendix A) in the closed-vocabulary dataset and our ability to leverage the language-pretrained transformer layers in LLMs. Appendix B.2 presents additional ablation results on the sequential backbone architectures.

Previous works with specialized models (Gaddy and Klein, 2021; Benster et al., 2024) generally optimize connectionist temporal classifica-

Table 2: Ablation Study of EMG Adaptor training.

Component	Variants	WER
EMG-Adaptor w/ Llama3-3B	Fully Connected	0.70
	ResBlock (2)	0.64
	ResBlock (2) + Transformer	0.79
	ResBlock (2) + LSTM	0.53
Objective w/ Llama2-7B	Cross-Entropy(Section 2)	0.65
	CTC (Graves et al., 2006)	0.70

tion (CTC) (Graves et al., 2006) loss with a high beam width (>100). However, most LLMs are decoder-only architectures and are trained with cross-entropy (CE) loss. One challenge in optimizing these embeddings using CTC loss is ensuring that their temporal length exceeds the target sequence length (Sudo et al., 2025) for stable optimization. To achieve this, we leverage 1D convolution layers to dilate the embeddings. However, we find that optimizing the embeddings extracted from these LLMs with CTC loss remains suboptimal compared to using CE loss with temperature and small beam widths of just 4 for inference.

### 3.3 Further Explorations

**Person-identification from unvoiced EMG with 96% accuracy.** Physiological signals often carry person-specific phenotypes (Zlatintsi et al., 2023; Mohapatra et al., 2023b). To validate that EMG signals also encode such individualized traits, we conduct a pilot analysis on a public multi-subject dataset (Diener et al., 2020b), containing 1,000 unique utterances from four participants. We collapse the LLM’s embeddings along the temporal dimension and train a simple classification head (details in Appendix C), achieving an average accuracy of 0.96. This strong performance, consistent with findings in other EMG settings (Diener et al., 2020a; Wand et al., 2009), reaffirms that unvoiced EMG exhibits distinct user-specific patterns—even when users speak the same text segment. To further support this, we also train a fully end-to-end (non-LLM) model for person identification, which achieves 0.99 accuracy, confirming that these signals are highly discriminative. Importantly, our goal is not to propose a state-of-the-art method for user identification using LLMs, but rather to highlight that unvoiced EMG signals inherently carry identifiable traits. This motivates the need for personalized modeling, which typically requires learning from limited data. In this context, our design choice—to use a lightweight trainable adaptor

with frozen LLMs—offers a practical and efficient solution. It enables learning user-specific representations with minimal data, a key requirement in personalized EMG-based systems. Our results further emphasize the importance of data-efficient learning, as demonstrated in Figure 2.

**Exploration of data augmentation for unvoiced EMG.** We explored two data augmentation schemes: (1) random temporal shifts per EMG channel (Gaddy and Klein, 2021, 2020) and (2) Hilbert-transform-based phase augmentation from limb-based EMG gesture recognition (Mohapatra et al., 2024c; Wang et al., 2025), but neither significantly improved performance. Developing tailored augmentation methods for unvoiced EMG—balancing diversity and phonetic coherence—could enhance learning from limited data.

## Conclusion

Our approach demonstrates the potential of using LLMs to convert unvoiced EMG signals to text, achieving a 0.49 WER without any voiced data. In data-poor settings, it outperforms specialized models by 26%. Our experiments also highlight the value of hand-crafted features as input to LLMs for this task. This work<sup>3</sup> helps enable users who cannot produce vocal speech to interact with LLMs through unvoiced commands, especially as LLM-based assistants become ubiquitous.

## Acknowledgments

We gratefully acknowledge the support from National Science Foundation grants 2038853, 2324936, and 2328032. We would like to thank the creators of the open-source dataset used in our study. Special thanks to Neil Zhang for his thoughtful feedback and discussions on language modeling methods.

## Limitations

**Exploration with open vocabulary.** In most LLMs, the heavy lifting of supporting large vocabulary sizes (32,000 in Llama2 and 128,000 in Llama3) is carried out in the final layer (Wijmans et al., 2024). This makes learning even a closed vocabulary (67 words, approximately 4 words per utterance, with the potential for 50,000 unique utterances) from a new language modality a challenging task, which is the focus of this paper. Supporting a

larger, in-the-wild open vocabulary using unvoiced EMG is even more complex, with its difficulty exacerbated by the overall lack of data and the need for personalized models. However, our current explorations lay the groundwork for exploring this direction next. Prior specialized approaches in this domain heavily rely on the availability of audio and voiced EMG, typically using CTC loss for optimization. This also presents an additional challenge in training LLMs in low-resource modalities with a new objective that does not effectively utilize their learned embeddings from large-scale data, leading to suboptimal training. These challenges can be addressed by exploring multimodal LLMs to leverage implicitly aligned embeddings or by reformulating the open vocabulary as a larger closed set and using target-steering methods to train on this restructured vocabulary (Han et al., 2024).

**Broader Extension to more datasets and languages.** Specialized EMG-to-text models are often designed for a predetermined EMG configuration or are deeply tied to phonetic auxiliary tasks, limiting them to the English language and a specific dataset. Our method has the potential to be extended to multilingual and multi-configuration instrumentation. However, due to the lack of publicly available closed-vocabulary datasets in such settings, we limit our investigation to English corpora. A systematic multilingual and diverse instrumentation closed- and open-set recording setup can accelerate exploration in this challenging direction of converting unvoiced EMG to text.

Additionally, our current approach is reliant on the embedding layers of LLMs, so it cannot work with language model APIs that do not provide direct access to these embeddings.

## Ethical Concerns and Potential Risks

In this work, we utilize pretrained LLMs, specifically Llama, in accordance with their usage license, solely for academic research purposes. We do not foresee any immediate ethical concerns arising from our work. However, as an LLM application for interpreting biosignals, appropriate measures must be taken to preserve user privacy. Our techniques help make LLM-based assistants accessible to speech-impaired users, thereby encouraging socially beneficial outcomes.

<sup>3</sup>Our processed dataset, codebase and sample predictions are available at <https://github.com/payalmohapatra/SilentSpeechLLM>.

## References

- Alexei Baevski, Yuhao Zhou, Abdelrahman Mohamed, and Michael Auli. 2020. Wav2vec 2.0: A framework for self-supervised learning of speech representations. *Advances in neural information processing systems*, 33:12449–12460.
- Tyler Benster, Guy Wilson, Reshef Elisha, Francis R Willett, and Shaul Druckmann. 2024. A cross-modal approach to silent speech with LLM-enhanced recognition. *arXiv preprint arXiv:2403.05583*.
- Lorenz Diener, Shahin Amiriparian, Catarina Botelho, Kevin Scheck, Dennis Küster, Isabel Trancoso, Björn W. Schuller, and Tanja Schultz. 2020a. Towards silent paralinguistics: Deriving speaking mode and speaker id from electromyographic signals. In *Interspeech 2020*, pages 3117–3121.
- Lorenz Diener, Mehrdad Roustay Vishkasougeh, and Tanja Schultz. 2020b. CSL-EMG\_Array: An Open Access Corpus for EMG-to-Speech Conversion. In *Proceedings of Interspeech 2020*, pages 3745–3749.
- Pranay Dighe, Yi Su, Shangshang Zheng, Yunshu Liu, Vineet Garg, Xiaochuan Niu, and Ahmed Tewfik. 2024. Leveraging large language models for exploiting asr uncertainty. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 12231–12235. IEEE.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The Llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Markus Freitag and Yaser Al-Onaizan. 2017. Beam search strategies for neural machine translation. *arXiv preprint arXiv:1702.01806*.
- David Gaddy and Dan Klein. 2020. Digital voicing of silent speech. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5521–5530. Association for Computational Linguistics.
- David Gaddy and Dan Klein. 2021. An improved model for voicing silent speech. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 175–181, Online. Association for Computational Linguistics.
- Jose A Gonzalez-Lopez, Alejandro Gomez-Alanis, Juan M Martín Doñas, José L Pérez-Córdoba, and Angel M Gomez. 2020. Silent speech interfaces for speech restoration: A review. *IEEE access*, 8:177995–178021.
- Alex Graves, Santiago Fernández, Faustino Gomez, and Jürgen Schmidhuber. 2006. Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks. In *Proceedings of the 23rd international conference on Machine learning*, pages 369–376.
- Chi Han, Jialiang Xu, Manling Li, Yi Fung, Chenkai Sun, Nan Jiang, Tarek Abdelzaher, and Heng Ji. 2024. Word embeddings are steers for language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 16410–16430.
- Wei-Ning Hsu, Benjamin Bolte, Yao-Hung Hubert Tsai, Kushal Lakhotia, Ruslan Salakhutdinov, and Abdelrahman Mohamed. 2021. Hubert: Self-supervised speech representation learning by masked prediction of hidden units. *IEEE/ACM transactions on audio, speech, and language processing*, 29:3451–3460.
- 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 *ICLR*. OpenReview.net.
- Szu-Chen Jou, Tanja Schultz, Matthias Walliczek, Florian Kraft, and Alex Waibel. 2006. Towards continuous speech recognition using surface electromyography. In *Ninth International Conference on Spoken Language Processing*.
- Wonjune Kang, Junteng Jia, Chunyang Wu, Wei Zhou, Egor Lakomkin, Yashesh Gaur, Leda Sari, Suyoun Kim, Ke Li, Jay Mahadeokar, et al. 2024. Frozen large language models can perceive paralinguistic aspects of speech. *arXiv preprint arXiv:2410.01162*.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, pages 19730–19742. PMLR.
- Ilya Loshchilov and Frank Hutter. 2019. Decoupled weight decay regularization. 7th int. In *Conf. Learn. Represent. ICLR*.
- Ziyang Ma, Guanrou Yang, Yifan Yang, Zhifu Gao, Jiaming Wang, Zhihao Du, Fan Yu, Qian Chen, Siqi Zheng, Shiliang Zhang, et al. 2024. An embarrassingly simple approach for LLM with strong ASR capacity. *arXiv preprint arXiv:2402.08846*.
- Muhammad Maaz, Hanoona Rasheed, Salman Khan, and Fahad Shahbaz Khan. 2023. Video-chatgpt: Towards detailed video understanding via large vision and language models. *arXiv preprint arXiv:2306.05424*.
- Geoffrey S Meltzner, James T Heaton, Yunbin Deng, Gianluca De Luca, Serge H Roy, and Joshua C Kline. 2017. Silent speech recognition as an alternative communication device for persons with laryngectomy. *IEEE/ACM transactions on audio, speech, and language processing*, 25(12):2386–2398.

- Geoffrey S Meltzner, James T Heaton, Yunbin Deng, Gianluca De Luca, Serge H Roy, and Joshua C Kline. 2018. Development of sEMG sensors and algorithms for silent speech recognition. *Journal of neural engineering*, 15(4):046031.
- Payal Mohapatra, Ali Aroudi, Anurag Kumar, and Morteza Khaleghimeybodi. 2024a. Non-verbal hands-free control for smart glasses using teeth clicks. *arXiv preprint arXiv:2408.11346*.
- Payal Mohapatra, Bashima Islam, Md Tamzeed Islam, Ruochen Jiao, and Qi Zhu. 2023a. Efficient stuttering event detection using siamese networks. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE.
- Payal Mohapatra, Shamika Likhite, Subrata Biswas, Bashima Islam, and Qi Zhu. 2024b. Missingness-resilient video-enhanced multimodal disfluency detection. In *Interspeech 2024*, pages 5093–5097.
- Payal Mohapatra, Akash Pandey, Bashima Islam, and Qi Zhu. 2022. Speech disfluency detection with contextual representation and data distillation. In *Proceedings of the 1st ACM international workshop on intelligent acoustic systems and applications*, pages 19–24.
- Payal Mohapatra, Akash Pandey, Sinan Keten, Wei Chen, and Qi Zhu. 2023b. Person identification with wearable sensing using missing feature encoding and multi-stage modality fusion. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–2. IEEE.
- Payal Mohapatra, Akash Pandey, Yueyuan Sui, and Qi Zhu. 2023c. Effect of attention and self-supervised speech embeddings on non-semantic speech tasks. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 9511–9515.
- Payal Mohapatra, Lixu Wang, and Qi Zhu. 2024c. Phase-driven domain generalizable learning for nonstationary time series. *arXiv preprint arXiv:2402.05960*.
- Vimal Mollyn, Riku Arakawa, Mayank Goel, Chris Harrison, and Karan Ahuja. 2023. Imposer: Full-body pose estimation using imus in phones, watches, and earbuds. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, pages 1–12.
- NICD. 2024. [Quick statistics about voice, speech, language, and swallowing](#). Accessed February 13, 2025.
- Yuqi Nie, Nam H Nguyen, Phanwadee Sinthong, and Jayant Kalagnanam. 2022. A time series is worth 64 words: Long-term forecasting with transformers. *arXiv preprint arXiv:2211.14730*.
- R OpenAI. 2023. GPT-4 technical report. arxiv 2303.08774. *View in Article*, 2(5).
- Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. 2023. Robust speech recognition via large-scale weak supervision. In *International conference on machine learning*, pages 28492–28518. PMLR.
- Tanja Schultz and Michael Wand. 2010. Modeling coarticulation in EMG-based continuous speech recognition. *Speech Communication*, 52(4):341–353.
- Tanja Schultz, Michael Wand, Thomas Hueber, Dean J Krusienski, Christian Herff, and Jonathan S Brumberg. 2017. Biosignal-based spoken communication: A survey. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 25(12):2257–2271.
- Tanmay Srivastava, Prerna Khanna, Shijia Pan, Phuc Nguyen, and Shubham Jain. 2024. Poster unvoiced: Designing an unvoiced user interface using earables and llms. In *Proceedings of the 22nd ACM Conference on Embedded Networked Sensor Systems*, pages 871–872.
- Yui Sudo, Muhammad Shakeel, Yosuke Fukumoto, Brian Yan, Jiatong Shi, Yifan Peng, and Shinji Watanabe. 2025. Joint beam search integrating ctc, attention, and transducer decoders. *IEEE Transactions on Audio, Speech and Language Processing*.
- Weiting Tan, Hirofumi Inaguma, Ning Dong, Paden Tomasello, and Xutai Ma. 2024. SSR: Alignment-aware modality connector for speech language models. *arXiv preprint arXiv:2410.00168*.
- Changli Tang, Wenyi Yu, Guangzhi Sun, Xianzhao Chen, Tian Tan, Wei Li, Lu Lu, Zejun Ma, and Chao Zhang. 2023. Salmonn: Towards generic hearing abilities for large language models. *arXiv preprint arXiv:2310.13289*.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Michael Wand, Szu-Chen Stan Jou, Arthur R. Toth, and Tanja Schultz. 2009. Impact of different speaking modes on emg-based speech recognition. In *Interspeech 2009*, pages 648–651.
- Lixu Wang, Bingqi Shang, Yi Li, Payal Mohapatra, Wei Dong, Xiao Wang, and Qi Zhu. 2025. Split adaptation for pre-trained vision transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.

Erik Wijmans, Brody Huval, Alexander Hertzberg, Vladlen Koltun, and Philipp Krähenbühl. 2024. Cut your losses in large-vocabulary language models. *arXiv preprint arXiv:2411.09009*.

Haixu Wu, Jiehui Xu, Jianmin Wang, and Mingsheng Long. 2021. Autoformer: Decomposition transformers with auto-correlation for long-term series forecasting. *Advances in neural information processing systems*, 34:22419–22430.

Zhiyang Xu, Ying Shen, and Lifu Huang. 2022. Multiinstruct: Improving multi-modal zero-shot learning via instruction tuning. *arXiv preprint arXiv:2212.10773*.

Jeonghun Yeo, Seunghee Han, Minsu Kim, and Yong Man Ro. 2024. Where visual speech meets language: VSP-LLM framework for efficient and context-aware visual speech processing. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 11391–11406, Miami, Florida, USA. Association for Computational Linguistics.

Wenyi Yu, Changli Tang, Guangzhi Sun, Xianzhao Chen, Tian Tan, Wei Li, Lu Lu, Zejun Ma, and Chao Zhang. 2024. Connecting speech encoder and large language model for ASR. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 12637–12641. IEEE.

Ailing Zeng, Muxi Chen, Lei Zhang, and Qiang Xu. 2023. Are transformers effective for time series forecasting? In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, pages 11121–11128.

Ruidong Zhang, Mingyang Chen, Benjamin Steeper, Yaxuan Li, Zihan Yan, Yizhuo Chen, Songyun Tao, Tuochao Chen, Hyunchul Lim, and Cheng Zhang. 2021. Speechin: a smart necklace for silent speech recognition. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 5(4):1–23.

Athanasia Zlatintsi, Panagiotis Paraskevas Filintisis, Niki Efthymiou, Christos Garoufis, George Retsinas, Thomas Sounapoglou, Ilias Maglogiannis, Panayiotis Tsanakas, Nikolaos Smyrnis, and Petros Maragos. 2023. E-prevention: The icassp-2023 challenge on person identification and relapse detection from continuous recordings of biosignals. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–2. IEEE.

## APPENDIX

### A Dataset Statistics

Table 3 provides a summary of the dataset statistics, while Figures 4 and 5 illustrate the distribution of the input EMG sequence lengths and target lengths.

Table 3: Comparison of dataset statistics from Gaddy and Klein (2020) and Diener et al. (2020b). Here,  $N$  represents the number of participants, and "Train" includes both training and validation samples.

Statistic	Gaddy	CSL
Dataset Size	500 utterances	1000 utterances
Target Type	Closed Vocabulary Template	Fixed Open Vocabulary
Number of Individuals ( $N$ )	1	4
Train Set (Train + Val)	450	800
Evaluation Set	50	200
k-fold Evaluations	3	3
Sampling Rate	1000 Hz → 800 Hz	2048 Hz




Figure 4: Histogram of the distribution of the sequence lengths and the target lengths for Gaddy and Klein (2020) closed-vocabulary dataset.




Figure 5: Histogram of the distribution of the sequence lengths and the target lengths for Diener et al. (2020b)'s Person 1 Block1 Initial segment of the data which is the superset for all the utterances.

### B Implementation Details

We provide more implementation-specific details of the baselines here and summarize the sizes of all models in Table 4.

#### B.1 Application-specific Baseline

To test the application-specific baseline model under fair settings, we ran it with beam widths of 4 (same as our LLM-based model) and 100. As noted in Table 5, the WER for both the beam widths are nearly the same, indicating that beam width has a negligible effect on the prediction in the case of



Table 4: Trainable parameters of different models. Values are reported in millions (M) and thousands (k).

Model	Trainable Parameters
Application-Specific Model	54M
Application-Specific (Modified)	8.1M
LLaMA-7B + EMG Adaptor, LLaMA-3B + EMG Adaptor	6.4M
Fine-tuned LLaMA-3B + EMG Adaptor	10.3M
Fine-tuned LLaMA-3B + EMG Adaptor (Alternative)	5.94M
LLaMA-3B + Audio Adaptor	590k

the closed vocabulary. We also train the baseline model with smaller trainable parameters (8M) by reducing the feature size, size of feedforward linear layers, and number of layers in the transformers to 512, 512, and 2 respectively. As shown in Table 5, the WER with a smaller baseline model is also in the same range as with 54M parameters. This highlights the fact that the difference in performance between our model and the baseline model in Table 1 is due to the choice of LLM for the prediction.

Table 5: WER performance comparison of the original (54M) and reduced (8M) Gaddy models under different beam-width settings.

Model	Beam-width	WER
Original (54M)	$n = 4$	0.70
	$n = 100$	0.72
Reduced (8M)	$n = 4$	0.87

## B.2 Additional Ablation: EMG-Adaptor Backbone Variants

We conducted additional ablation studies to explore different sequential backbones for the EMG-Adaptor (EMG-Ad) when used with the LLaMA 3B LLM. The results are summarized in Table 6.

Table 6: Performance of different sequential backbones for the EMG-Adaptor with LLaMA 3B. Number of transformer layers are denoted as L.

Sequential Backbone	Params	WER
BiLSTM	5.94M	0.52
LSTM	5.5M	0.58
Transformer (6L + Sinusoidal) (Vaswani et al., 2017)	6.7M	0.79
Transformer (6L + RoPE) (Touvron et al., 2023)	6.7M	0.75
Transformer (2L + RoPE)	5.3M	0.72

We experimented with Rotary Position Encoding (RoPE) (Touvron et al., 2023), motivated by its effectiveness in the LLaMA models, with the intuition that it might produce better input tokens for

the LLM. Although RoPE-based transformers improved performance over the vanilla sinusoidal variant, they still underperformed compared to LSTM-based models.

This aligns with findings from prior time-series literature. For example, Zeng et al. (2023) have argued that the permutation-invariant nature of self-attention leads to temporal information loss. Similarly, practical studies such as IMUPoser (Molyn et al., 2023) empirically corroborate that LSTMs outperform transformers in fine-grained time-series tasks like human activity recognition, e.g., stating in Section 4.1 of (Molyn et al., 2023): “*Although we did experiment with newer architectures such as transformers, we found these models did not perform well in practice.*”

While advanced positional encoding schemes and specialized architectures (Wu et al., 2021; Nie et al., 2022) have enhanced transformer performance on time series, we restrict our analysis to vanilla transformers for simplicity and focus on their ability to generate suitable EMG tokens for LLMs. Our findings suggest that LSTMs are better suited for this task. This observation is currently limited to short sequences drawn from a closed vocabulary. Future work will investigate more specialized transformer-based architectures for unvoiced EMG modeling.

## B.3 Expanding LLMs to Audio

In this experiment, we primarily leverage the implementation of Tang et al. (2023) and the idea proposed by Ma et al. (2024) to employ an embarrassingly simple approach for speech recognition with LLMs, using a linear projector from a frozen speech encoder. We use the wav2vec 2.0 (Baevski et al., 2020; Mohapatra et al., 2023a, 2024b, 2022) BASE architecture as our speech encoder, which produces a 768-dimensional feature vector. This vector is then passed through two linear layers to generate the 3072-dimensional input required for Llama3-3B. While the speech encoder can be replaced with alternatives such as HuBERT (Hsu et al., 2021) or Whisper (Radford et al., 2023), our goal is not to optimize speech-to-text conversion. Instead, we aim to demonstrate that while both audio and EMG involve expanding an LLM’s capability to a new language modality, integrating EMG signals poses significantly greater challenges.

#### B.4 Additional Reproducibility Information

All experiments were conducted using NVIDIA A100 GPUs (3 available, 40GB CUDA memory) and TITAN RTX GPUs (4 available, 24GB CUDA memory) with a maximum runtime of 12 hours in PyTorch. Hyperparameter tuning combined manual and random search, typically requiring fewer than five runs, with selection based solely on validation loss and WER.

The batch size for LLM experiments was 8, and the maximum number of epochs was 500. For the App-Specific model, we retain the original settings (Gaddy and Klein, 2021), where the authors re-batchified the input by rolling the temporal dimensions to support training on longer sequences.

#### C Person Identification using Unvoiced EMG : Implementation Details

For the person identification task, we use LLaMA 3.2-3B. Overall the model architecture remains the same as shown in Figure 1 except we do not use any prompt in this case. The logits output from the LLM,  $Z \in \mathbb{R}^{T \times V}$ , is reduced to  $\bar{Z} \in \mathbb{R}^{1 \times V}$  by taking mean along the time axis. The  $\bar{Z}$  is then fed into several linear layers to predict logits  $y^p \in \mathbb{R}^{1 \times 4}$  as there are 4 distinct persons in the dataset (Diener et al., 2020b). Due to the linear layers, the number of trainable parameters is 38M, which is higher than the numbers mentioned in Table 4 for LLM-based models. Using  $y^p$  and actual person ID, the softmax loss is calculated to fine-tune 32M parameters.

#### D Additional Ethical Statements

In preparing this work, we only used AI assistants in the capacity to polish the language in the manuscript.