Poster: Privacy-aware dysarthria classification

Bhawana Chhaglani*
University of Massachusetts Amherst
Amherst, USA
bchhaglani@umass.edu

Tanvi Kandepuneni†
University of Massachusetts Amherst
Amherst, USA
tkandepuneni@umass.edu

Prashant Shenoy
University of Massachusetts Amherst
Amherst, USA
shenoy@cs.umass.edu

Abstract

Dysarthria detection from speech enables scalable neurological health monitoring but raises serious privacy concerns due to leakage of spoken content. We explore whether dysarthria can be detected and classified without intelligible speech using privacy-aware signal processing. Using MFCC features and lightweight classifiers, we show that dysarthria-relevant cues persist under strong privacy filtering, enabling a favorable privacy–utility trade-off.

CCS Concepts

• **Security and privacy → Usability in security and privacy**.

Keywords

Privacy, Dysarthria, Audio Sensing, Health Monitoring

ACM Reference Format:

Bhawana Chhaglani, Tanvi Kandepuneni, and Prashant Shenoy. 2026. Poster: Privacy-aware dysarthria classification. In *The 27th International Workshop on Mobile Computing Systems and Applications (HotMobile '26), February 25–26, 2026, Atlanta, GA, USA.* ACM, New York, NY, USA, 1 page. https://doi.org/10.1145/3789514.3796257

1 Motivation and Approach

Speech is a powerful biomarker for neurological health, and prior work has demonstrated strong performance in detecting dysarthria using acoustic features and machine learning models [1, 2]. Speech-based detection enables remote and scalable monitoring, but continuous audio capture poses privacy risks due to leakage of speech [3, 4]. These concerns limit real-world deployment of always-on health sensing systems. We ask a simple question: *can dysarthria be detected without intelligible speech?* As dysarthria affects respiration, phonation, articulation, and prosody, non-speech signals capture breathing, articulation transitions, and dysfluencies associated with dysarthria. This can be leveraged to monitor health related information without leaking privacy. We study the privacy implications of these non-speech sounds and whether they can leak human speech.

Design: We use an off-the-shelf Silero VAD (with threshold of 0.5) to detect and eliminate speech segments and generate privacy filtered audio. From privacy-filtered audio, we extract MFCCs ($\pm\Delta$, $\Delta\Delta$) - computed with a 25 ms window size (400 samples at 16kHz) and a 10 ms hop length - and train an RBF kernel SVM model (C = 10, γ = scale) on Cepstral Mean and Variance Normalization (CMVN) features for dysarthria detection and severity classification. Training and testing were done using a speaker-disjoint split (train utterances = 16791, test utterances = 4590). We use UASpeech [5] dataset and design a 3-class classifier: healthy, low dysarthria, and high dysarthria. We evaluate balanced accuracy, precision, and recall using subject-independent splits. We use a pretrained Wav2Vec2 ASR model to compute Character Error Rate (CER) as a measure of intelligible speech leakage.

Table 1: Utility–privacy results for dysarthria classification Privacy filtered sensing preserves dysarthria classification performance while substantially increasing ASR CER, indicating almost 0 speech leakage.

Input	Acc.	Pre.	Rec.	CER (Dys)	CER (Control)
Full signal	0.64	0.67	0.64	0.68	0.23
Privacy filtered	**0.71**	**0.71**	**0.71**	**1.02**	**0.997**

Results: Figure 1 and Table 1 shows the result of privacy and utility comparison of full signal Vs privacy-filtered signal for dysarthia classification. The full raw signal approach shows 64% accuracy while leaking significant speech (CER=23%). Privacy-filtered sensing improves dysarthria classification performance by 7% while substantially increasing ASR CER for both dysarthric and control speech, indicating near-zero intelligible speech leakage. Thus, non-speech cues are valuable biomarkers for continuous neurological health monitoring.

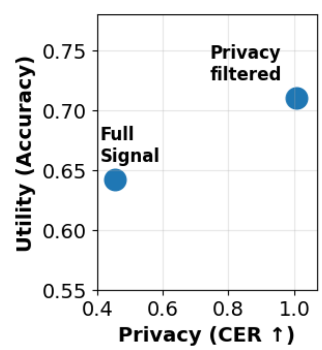

Figure 1: Privacy-utility comparison.

*These authors contributed equally to this work.
†These authors contributed equally to this work.

ACM ISBN 979-8-4007-2471-8/2026/02
https://doi.org/10.1145/3789514.3796257

References

[1] Al-Ali et al. 2024. The detection of dysarthria severity levels using AI models: A review. *IEEE Access* 12 (2024).

[2] Anuprabha et al. 2025. A Multi-modal Approach to Dysarthria Detection and Severity Assessment Using Speech and Text Information. In *ICASSP 2025*. IEEE, 1–5.

[3] Chhaglani et al. 2024. Towards Privacy-Preserving Audio Classification Systems. *arXiv preprint arXiv:2404.18002* (2024).

[4] Chhaglani et al. 2025. FeatureSense: Protecting Speaker Attributes in Always-On Audio Sensing System. *arXiv preprint arXiv:2505.24115* (2025).

[5] Kim et al. 2008. Dysarthric speech database for universal access research. *Interspeech* (2008).