Differential Privacy Enabled Dementia Classification: An Exploration of the Privacy-Accuracy Trade-off in Speech Signal Data

Suhas BN [1] Sarah Rajtmajer [1] Saeed Abdullah [1]

[1]College of Information Sciences & Technology, Penn State, USA

Introduction

- Dementia impacts over 55 million individuals [1]. The global cost resulting from dementia is estimated to be $2 trillion in 2030 [2]. Given its increasing prevalence and the resultant personal and societal burden, dementia is one of the most critical public health issues of our time.
- *Early detection* of dementia is critical for effective illness management [3]. However, current diagnosis methods can be resource-intensive and time-consuming. As a result, assessment and monitoring of dementia at scale can be challenging, specifically for individuals living in remote regions.
- Speech and language changes can signal early dementia onset. Recent ML models using speech show promise for automated assessment at scale [4, 5].
- However, prior work has not adequately addressed the privacy-accuracy trade-off of these models. For example, features used in some of these models can be used to reconstruct utterances [6, 7]. This is a serious concern given the models are using data from a vulnerable population.
- **Our objective** is to evaluate this tradeoff for dementia detection using speech data and differential privacy (DP) techniques.

Our contributions

1. Establish benchmark on tradeoff using DP with varying ϵ budget.
2. Demonstrate the feasibility of DP for privacy-preserving dementia classification.
3. Provide insights on optimizing models for intended use case.

Experimental Setup

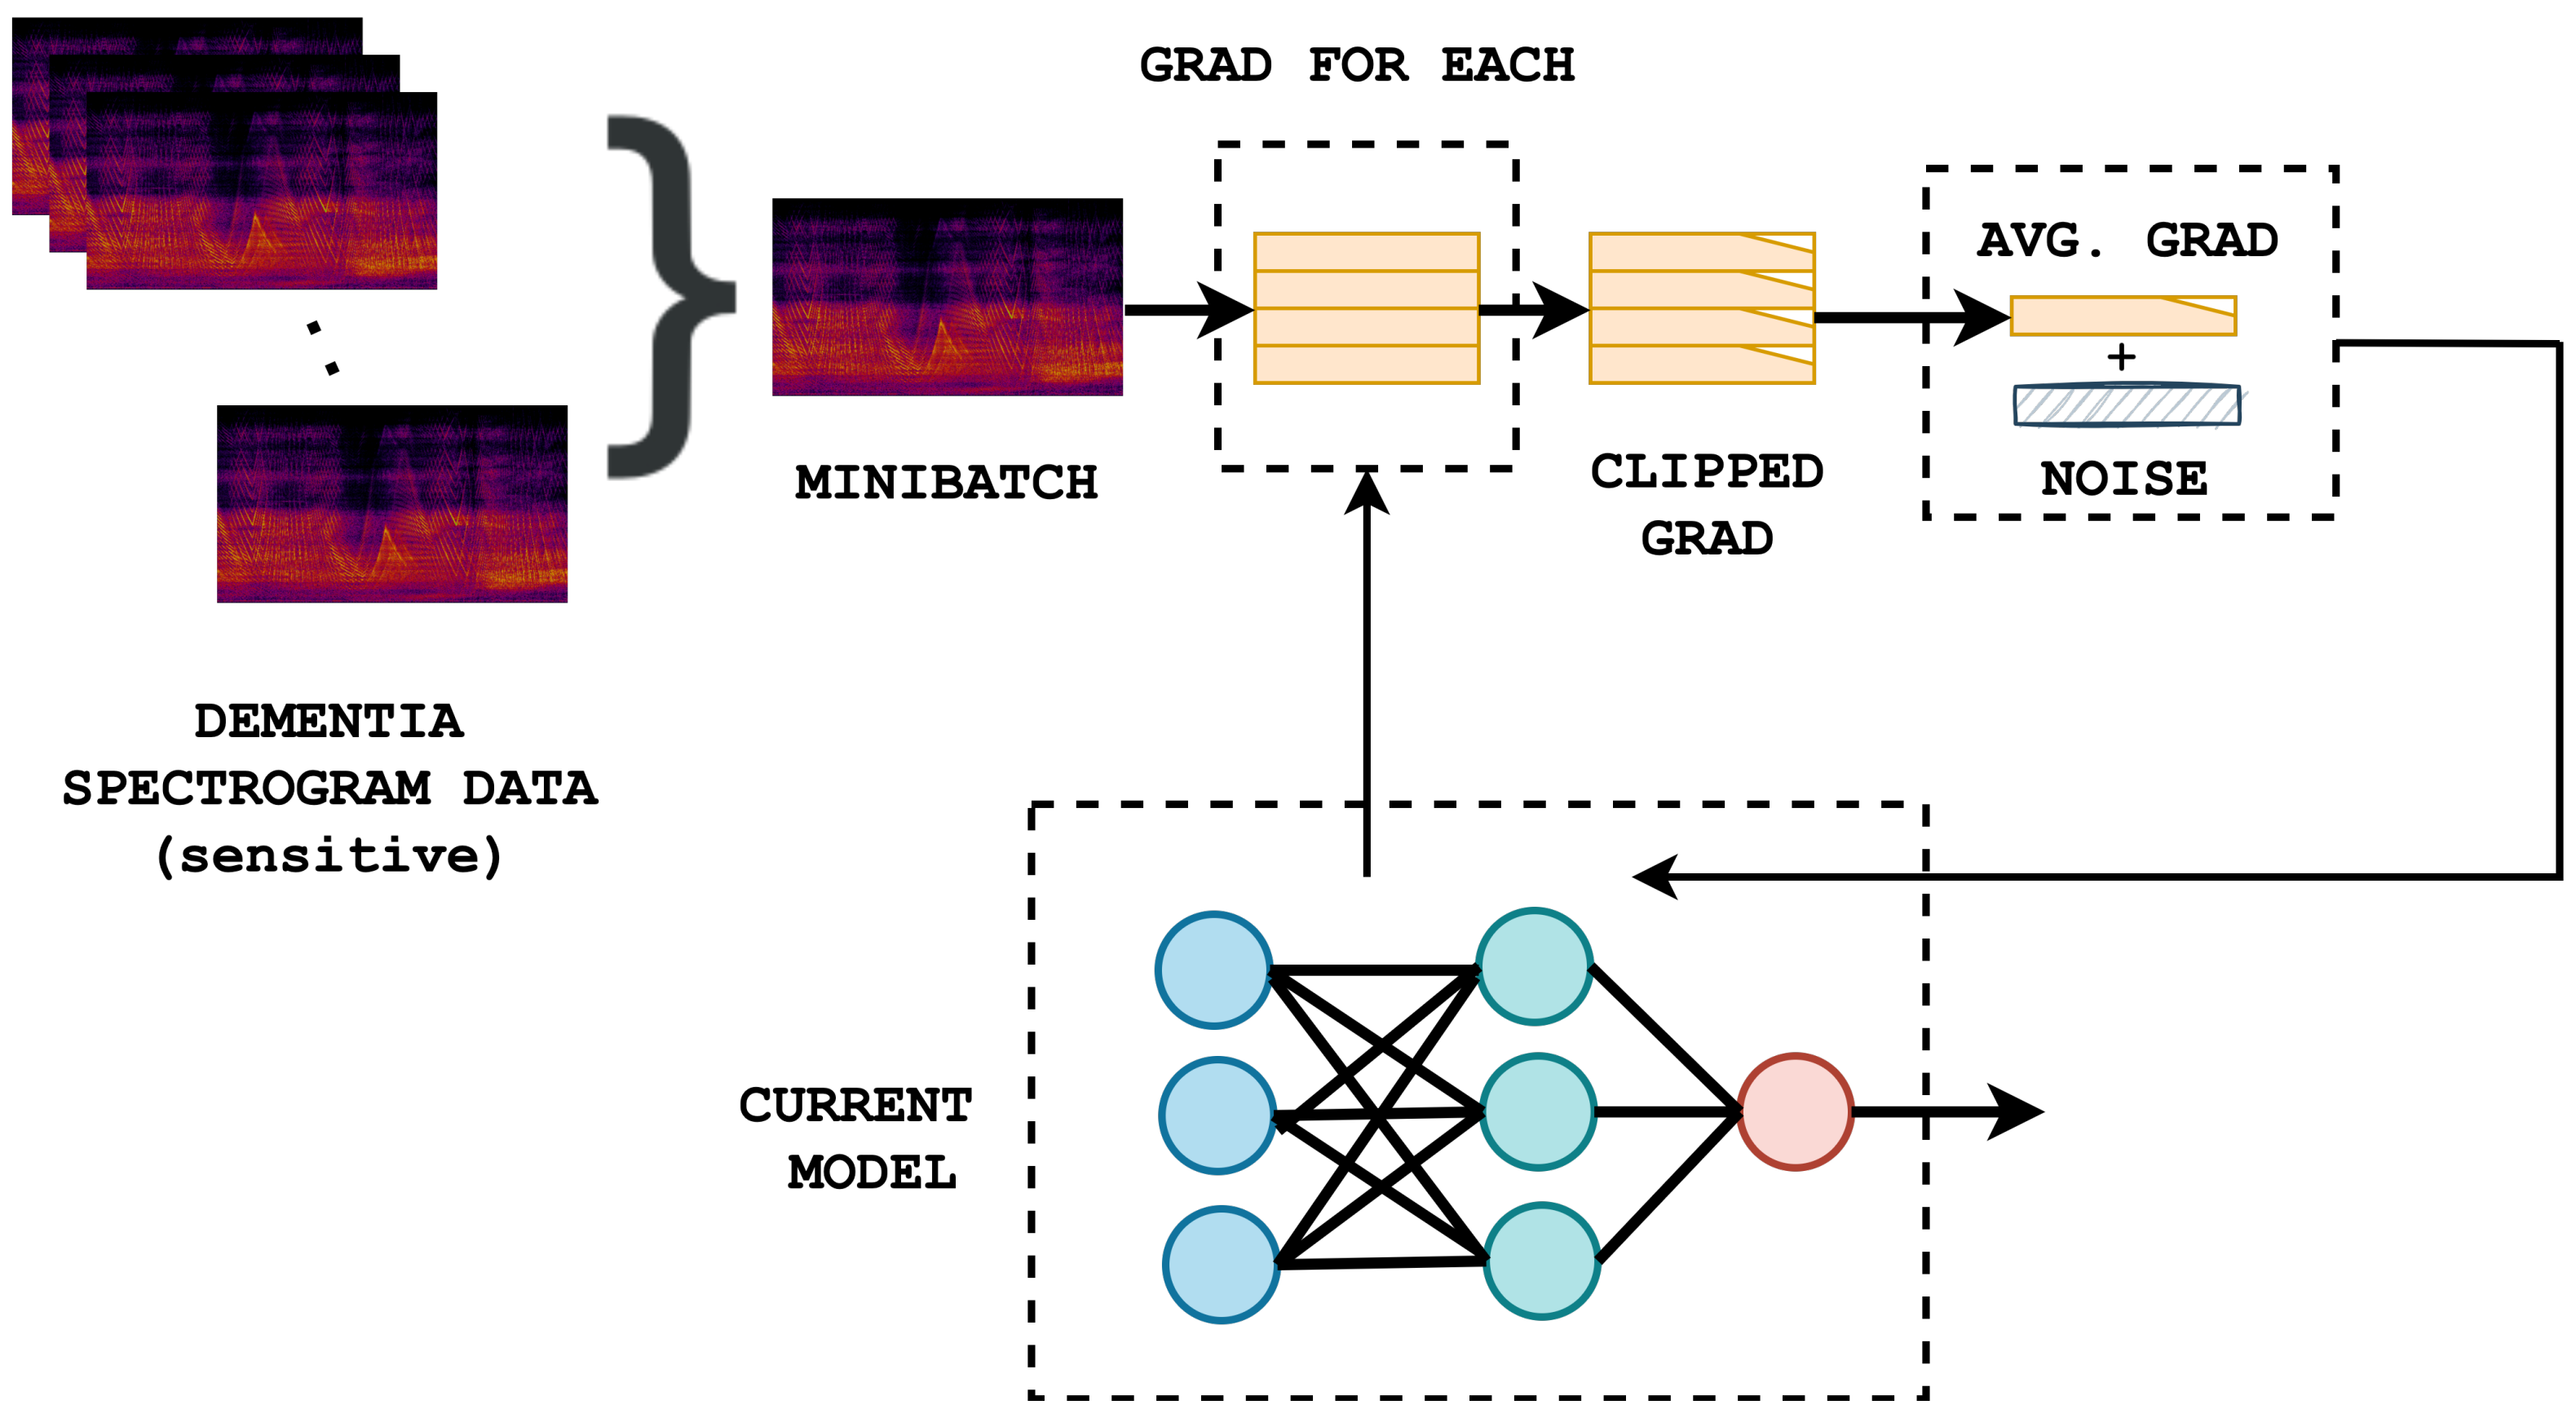

Figure 1. The Differential Privacy methodology used in this work selects minibatches from the training data, clips the calculated gradients and adds gaussian noise. This is repeated for the duration of the training phase.

- **Data**: DementiaBank Pitt and WLS speech corpora
  1. Pitt: 459 recordings from 292 subjects, cookie theft task
  2. WLS: 116 recordings, verbal fluency task
  3. Labels using MMSE scores and diagnostic criteria
- **Preprocessing**:
  1. Convert MP3 to WAV
  2. Downsample to 16kHz, trim audio to participant speech
  3. Extract log-mel spectrogram features (librosa)
- **Feature**: Reshaped spectrogram images (224x224)
- **Model Architecture**: ResNet-18 finetuned for binary dementia classification
- **Method**: Five-fold cross-validation setup and early stopping criteria. Prevents overfitting
- **Platform**: PyTorch on an NVIDIA Tesla V100-SXM2-32GB GPU
- **Optimizer**: Stochastic Gradient Descent (SGD) optimizer for standard model, and DP-SGD optimizer for private model (Opacus library) see paper for optimizer parameters
- **Loss criterion**: Cross Entropy
- **Parameters**:
  - Delta (δ): Set to 3e-4 (based on dataset)
  - Epsilon (ϵ): [0.1, 0.5, 1, 5, 10, 50, 100]. When $\epsilon \rightarrow \infty$, we get non-DP case.
  - Max Grad Norm (C): The maximum L2 norm of per-sample gradients before the averaging step aggregates them. We have chosen the following values: [0.1, 0.5, 1, 5, 10].
- **Experiments**:
  - **Classification Task**: 2-class Dementia classification.
  - **Comparison**: Evaluate DP against non-DP models.
  - **Hyperparameter Tuning**: Vary the ϵ privacy budget to investigate its effect on accuracy and select the optimal ϵ value for performance.

Results

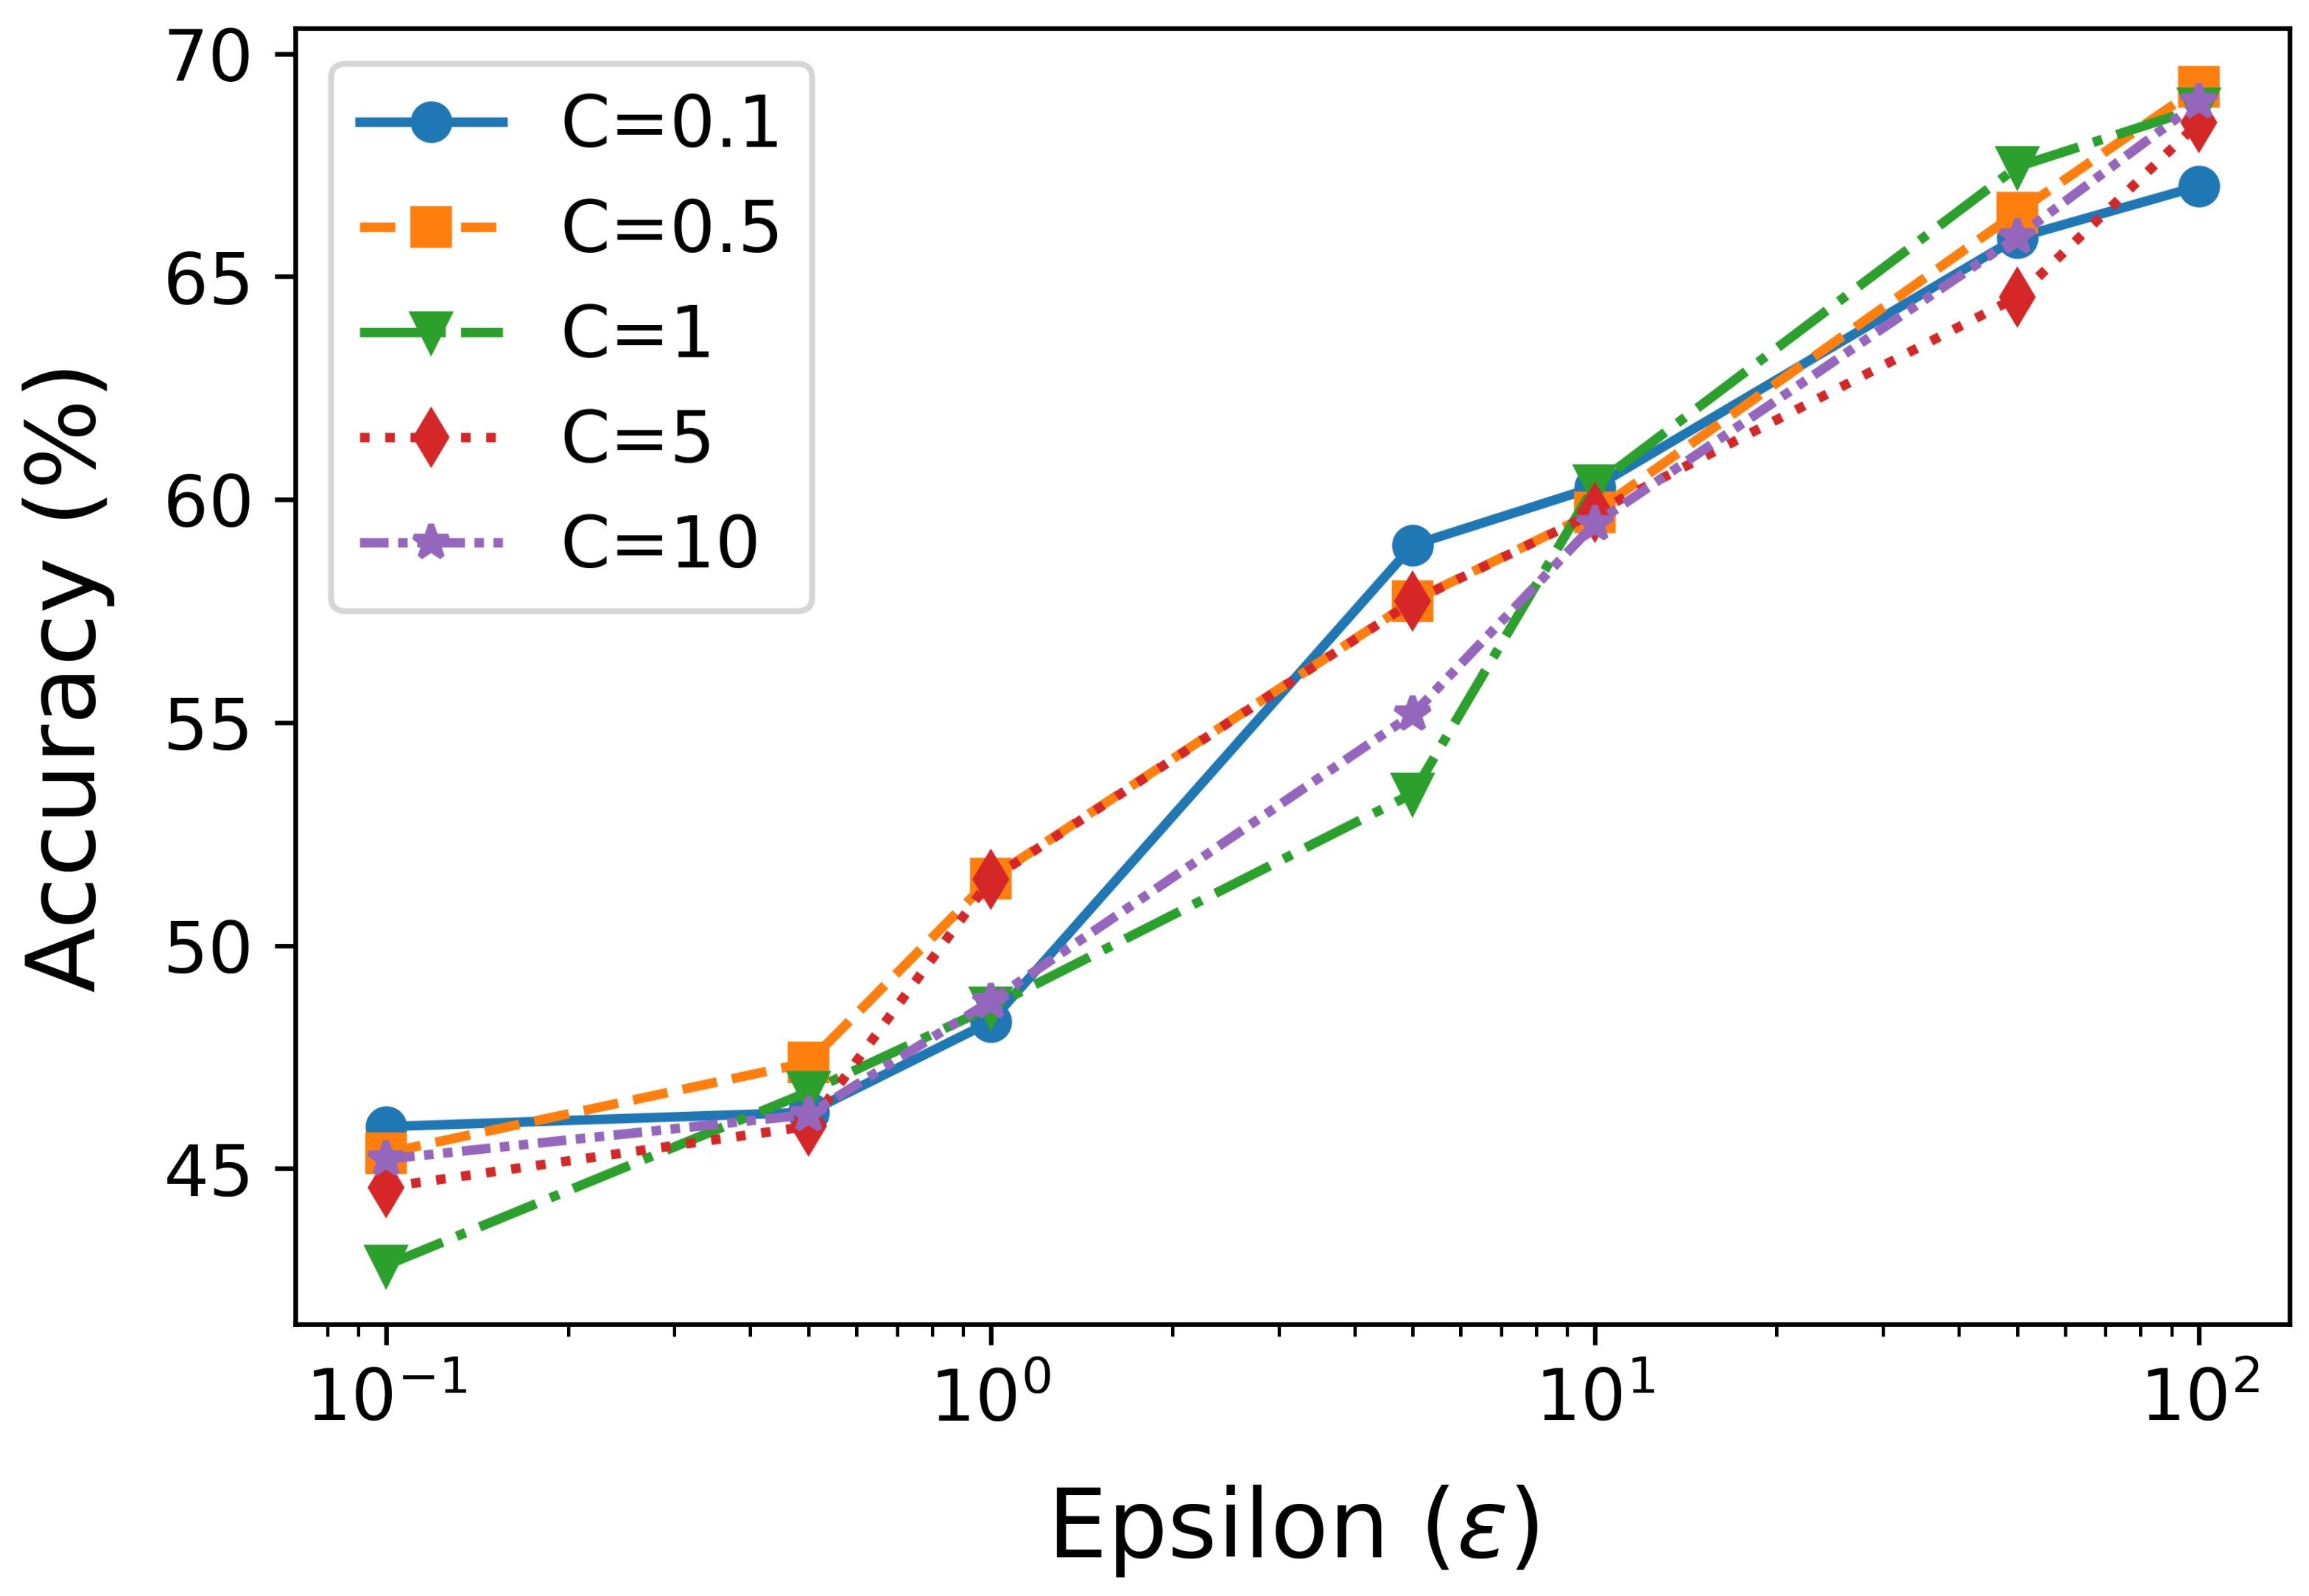

Figure 2. Privacy-Accuracy Tradeoff in Dementia Prediction using DP-SGD: Investigating Hyperparameter Impact on Combined WLS and Pitt Corpus Datasets

Performance Comparison

Table 1. Accuracy for varying values of ϵ and C with constant δ. As ϵ increases, accuracy is observed to improve. Bold numbers show the best performance per ϵ row.

↓ϵ\C→	C= 0.1	C= 0.5	C= 1	C= 5	C= 10
ϵ = 0.1	**45.94**	45.35	42.83	44.58	45.34
ϵ = 0.5	46.26	**47.38**	46.72	45.95	46.21
ϵ = 1	48.29	**51.5**	48.62	51.48	48.75
ϵ = 5	**58.98**	57.73	53.41	57.72	55.19
ϵ = 10	60.28	59.72	**60.32**	59.71	59.46
ϵ = 50	65.87	66.44	**67.46**	64.54	65.89
ϵ = 100	67.03	**69.25**	68.78	68.46	68.92

Table 2. ϵ v/s σ values for the experiments

ϵ	0.1	0.5	1	5	10	50	100
σ	130	10.78	16.25	2.22	1.37	0.77	0.37
$\epsilon \cdot \sigma$	13	5.39	16.25	11.1	13.7	38.5	37.2

In a nutshell

Increasing ϵ improves accuracy; ϵ and σ tuning is context-specific.

Discussion

- **Privacy-Accuracy Trade-off**: As ϵ increases, accuracy improves to 94.2%.
- Regularization parameter C has a less pronounced impact.
- **Need for better regulation & methods**: Recent works [8] have shown inconsistencies in privacy parameters across organizations, with some exceeding recommended levels of ϵ, posing risks to user privacy [9].
- **Future work**: DP-SGD shows a significant drop in accuracy compared to non-DP methods. Approaches like dynamic privacy budgeting or combining multiple privacy-enhancing techniques could offer better trade-offs. What, then, is a good privacy budget?

Data & Code Availability

- **Data Source:** DementiaBank speech corpus [10].
- **Code Repository:** `https://github.com/suhasbn/SpeechDP`.

References

[1] World Health Organization. Dementia Fact Sheet kernel description. https://www.who.int/news-room/fact-sheets/detail/dementia, 2021. Accessed: 2022-04-20.

[2] Arindam Nandi, Nathaniel Counts, Simiao Chen, Benjamin Seligman, Daniel Tortorice, Daniel Vigo, and David E Bloom. Global and regional projections of the economic burden of alzheimer's disease and related dementias from 2019 to 2050: A value of statistical life approach. *EClinicalMedicine*, 51, 2022.

[3] Robert Briggs, Sean P Kennelly, and Desmond O'Neill. Drug treatments in alzheimer's disease. *Clinical medicine*, 16:247, 2016.

[4] Aparna Balagopalan, Benjamin Eyre, et al. To bert or not to bert: comparing speech and language-based approaches for alzheimer's disease detection. *arXiv preprint arXiv:2008.01551*, 2020.

[5] Benjamin Eyre et al. Fantastic features and where to find them: detecting cognitive impairment with a subsequence classification guided approach. *arXiv preprint arXiv:2010.06579*, 2020.

[6] Daniel Griffin and Jae Lim. Signal estimation from modified short-time fourier transform. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 32(2):236–243, 1984.

[7] Paul Magron and Tuomas Virtanen. Online spectrogram inversion for low-latency audio source separation. *IEEE Signal Processing Letters*, 27:306–310, 2020.

[8] Damien Desfontaines. A list of real-world uses of differential privacy. https://desfontain.es/privacy/real-world-differential-privacy.html, 2022. Accessed: 2022-10-20.

[9] WIRED. How One of Apple's Key Privacy Safeguards Falls Short. https://www.wired.com/story/apple-differential-privacy-shortcomings/. Accessed: 2023-03-07.

[10] James T Becker et al. The natural history of alzheimer's disease: description of study cohort and accuracy of diagnosis. *Archives of neurology*, 51(6):585–594, 1994.