NeurReps 2024 Proceedings Paper

Geometric Machine Learning on EEG Signals

Benjamin J. Choi
Harvard University

NEURAL INFORMATION PROCESSING SYSTEMS

Kempner INSTITUTE For the Study of Natural & Artificial Intelligence

Abstract

This study focuses on elucidating the underlying signal and low-dimensional geometric structure present in high-dimensional brainwave data. We demonstrate EEG signal processing pipelines: (1) a preliminary pipeline for denoising individual EEG channels, and (2) a downstream manifold learning pipeline uncovering geometric structure across networks of EEG channels. We conduct validation using two EEG datasets for denoising and imagined digit decoding. Our preliminary pipeline uses an attention-based EEG filtration network to extract clean signal from individual EEG channels. Our primary manifold learning pipeline averages a discrete analog of Ricci flow and a graph convolutional network to perform dimensionality reduction on multi-channel EEG data. Our system achieves competitive performance with existing signal processing and classification benchmarks; we demonstrate a mean test correlation coefficient of >0.95 at 2 dB on semi-synthetic neural denoising and a downstream EEG-based classification accuracy of 0.97 on distinguishing digit- versus non-digit thoughts. Results are preliminary and further validation on broader use cases will be crucial.

Background

Our study is performed using the EEGdenoiseNet dataset (Zhang et al., 2021), a benchmark dataset for EEG denoising, and the *MindBigData* (MBD) dataset (Vivancos et al., 2022). MBD consists of multi-channel EEG recordings corresponding to a wide variety of thoughts, including imagined digits and non-digits. Digital recordings are segmented into 2-second intervals, exist in high-dimensional spaces (d > 2000) with hypothetical low-dimensional structure, and are noisy. In order to denoise the incoming EEG channels, we apply a novel autoencoder-targeted adversarial transformer (AT-AT). Time series transformers have previously achieved success working in high-dimensional settings with intrinsic geometric structure (Cho et al., 2023), and our novel method of autoencoder-based tokenization (building on methods from Choi et al., 2024) and masking enables competitive filtration performance with minimal model footprint.

Detailed Overview of the AT-AT Filtration System

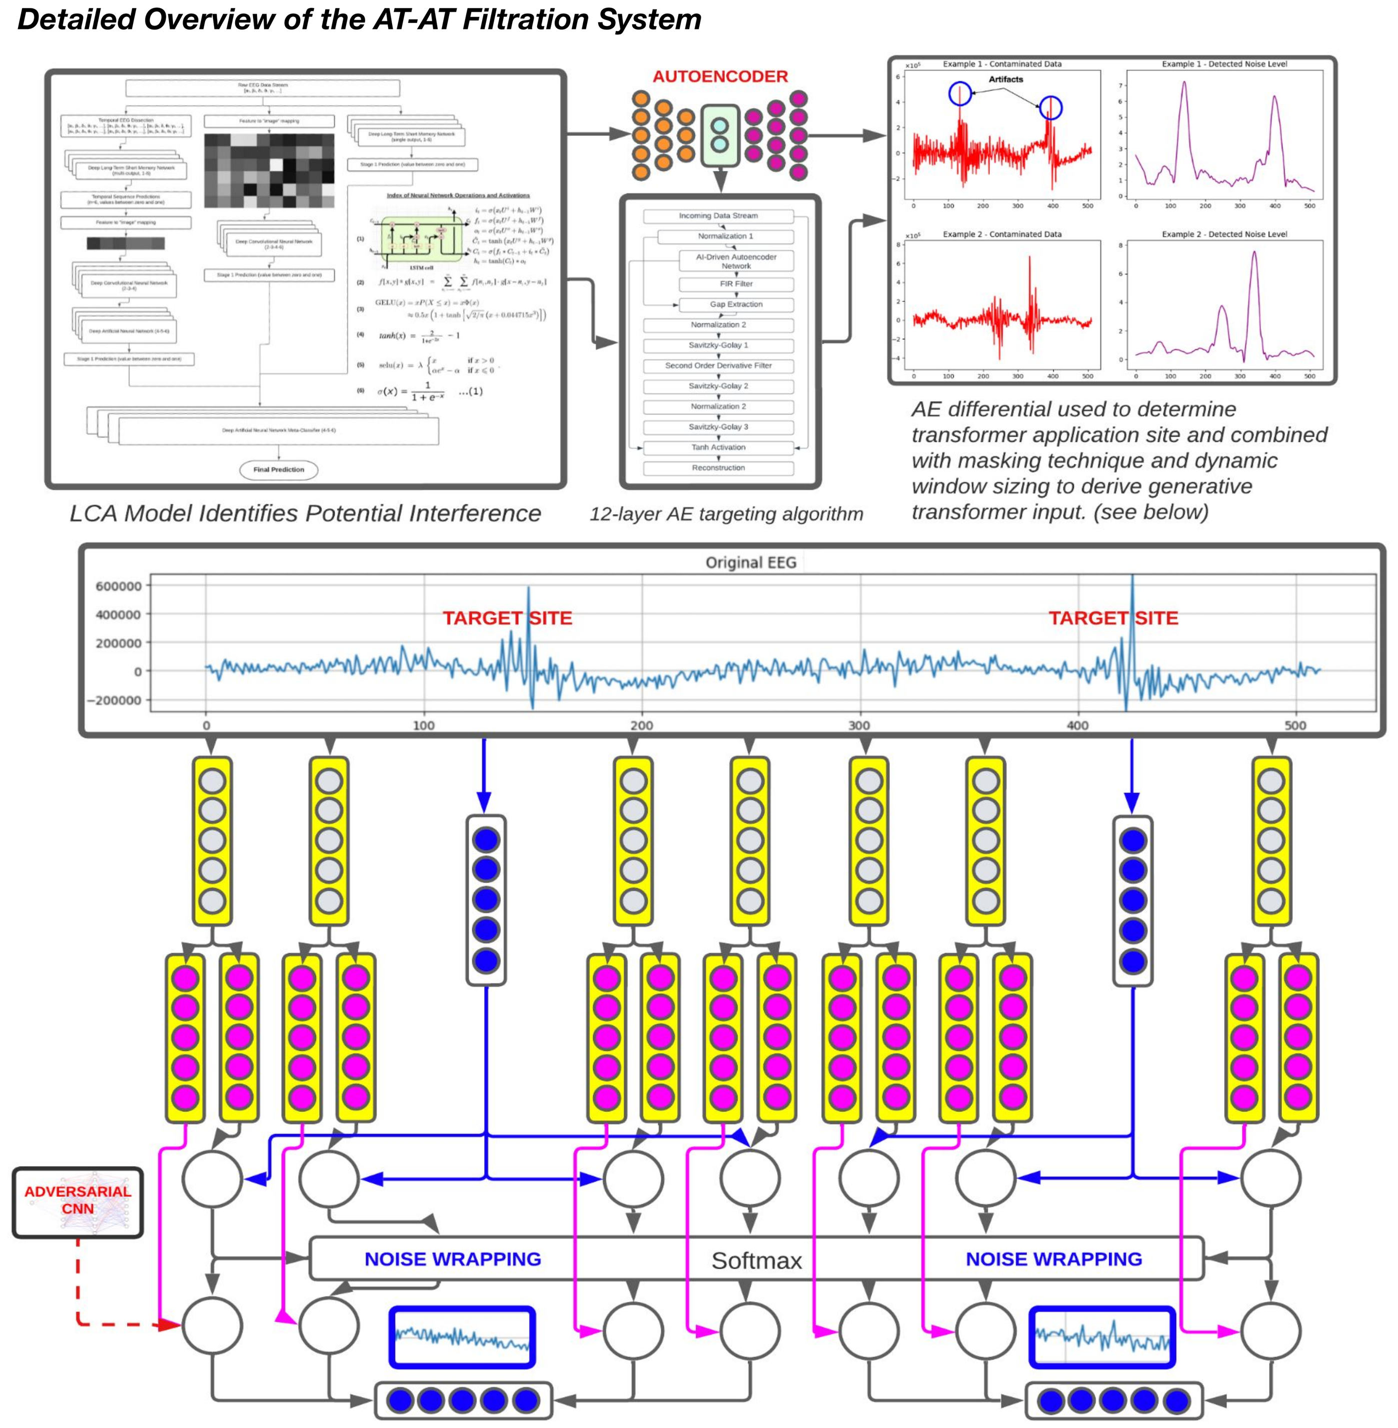

LCA Model Identifies Potential Interference

12-layer AE targeting algorithm

AE differential used to determine transformer application site and combined with masking technique and dynamic window sizing to derive generative transformer input. (see below)

Methods

Key Manifold Learning Pipeline Stages

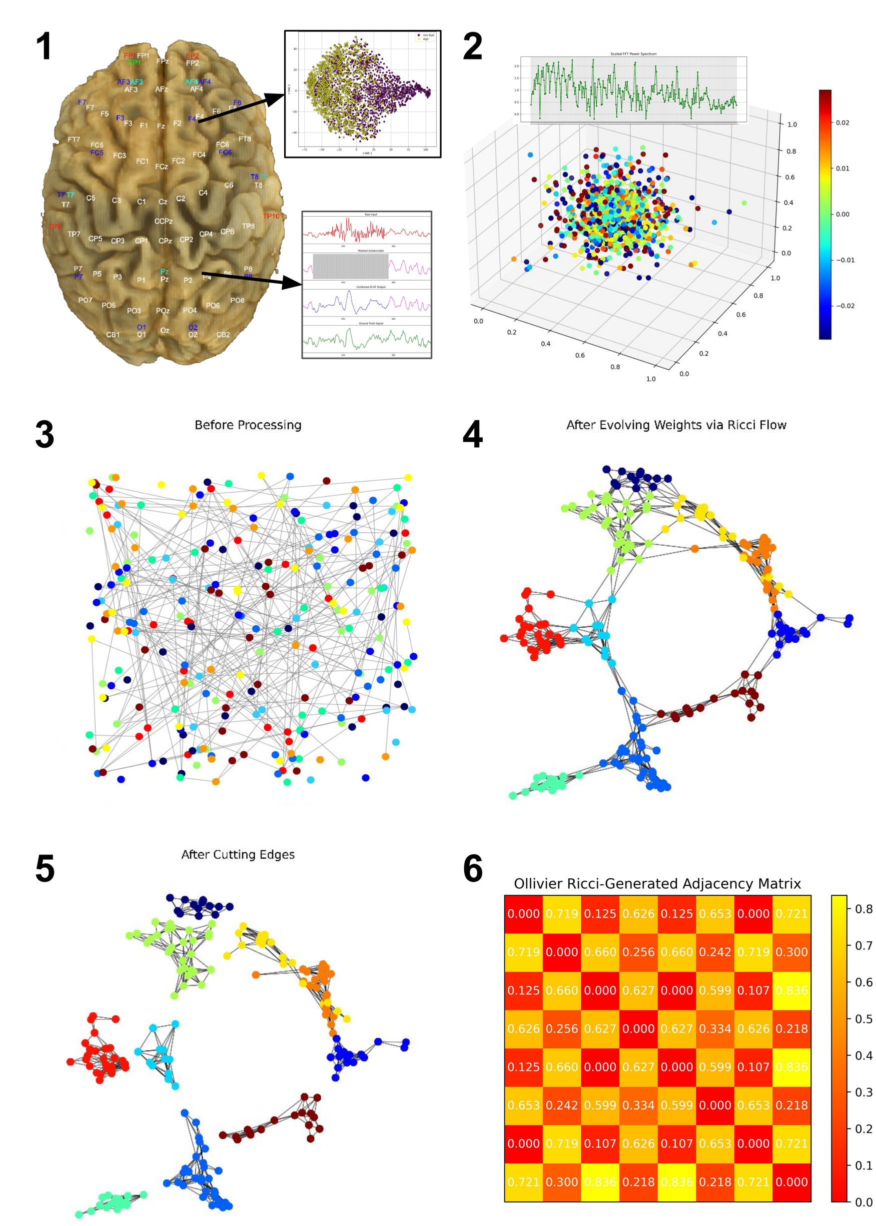

Our AT-AT denoising model was trained and validated on the semisynthetic EEGdenoiseNet dataset (Zhang et al., 2021). We then leveraged the trained AT-AT model to denoise the raw EEG channels in the MindBigData dataset. We found that AT-AT denoising not only removed artifacts in unseen data but also facilitated latent space separability of neural signal classes. We then turned to solving the curse-of-dimensionality problem in multi-channel EEG by uncovering underlying low-dimensional network structure via our manifold learning pipeline. This pipeline consists of four stages: (1) a fast Fourier transform and Laplacian eigenmap are used to obtain spectral and temporal representations of each channel, (2) a discrete analog of Ricci flow (via Ollivier-Ricci curvature) is leveraged to determine community structure within networks of EEG channels across neural regions, (3) a graph convolutional network (GCN) performs subsequent dimensionality reduction via the Ricci flow-computed adjacency matrix, and (4) a downstream classifier generates the final prediction.

Implementation of Ricci Flow

Graph Definition: $G = (V, E, w), \quad w : E \to \mathbb{R}^+$

Distance Function: $d(u,v) = \begin{cases} \frac{1}{2}(\|f(u) - f(v)\| + \mathrm{FFT}\Delta(u,v)) \\ \|f(u) - f(v)\| \\ \frac{1-\rho(\mathrm{FFT}(f(u)), \mathrm{rs}(f(v),256))}{2} \end{cases}$

where $\mathrm{FFT}\Delta(\mathrm{u,v}) = 0.5(1 - \rho(\mathrm{FFT(f(u))}, \mathrm{FFT(f(v))}))$

Graph Construction: $w(u,v) = \dfrac{1}{d(u,v)}, \quad \forall (u,v) \in E \text{ with } d(u,v) \neq 0$

Ollivier-Ricci Curvature: $\kappa(u,v) = 1 - \dfrac{W_1(\mu_u, \mu_v)}{d(u,v)}$

Ricci Flow Evolution: $w_{uv}(t+1) = w_{uv}(t) \cdot e^{-\alpha\kappa(u,v)}, \quad \alpha > 0$

Edge Cutting: $E_{\text{cut}} = \{(u,v) \in E : w_{uv}(t) \text{ in top } \rho\% \text{ of } E\}$

Post-Processing: $\bar{A} = \dfrac{1}{N}\sum_{i=1}^{N} A_i, \quad A_i \text{ adjacency matrix of } G_i$

Stages (1), (2), and (3) are all preceded by careful graph construction leveraging spatial, temporal, and frequency similarities in the multi-channel EEG setup. In the diagram (*left*), panel 1 depicts AT-AT processing, panel 2 depicts FFT and Laplacian eigenmap transformation, and panels 3-6 depict use of the Ricci flow. The Ricci flow algorithm for community detection involves evolving edge weights (with initial distances determined by temporal and spectral similarity) via Ollivier's notion of Ricci curvature; surgery is then performed on the generated graph by cutting heavily-weighted edges to partition the network into communities of neural channels.

Key Ollivier-Ricci Curvature Metrics

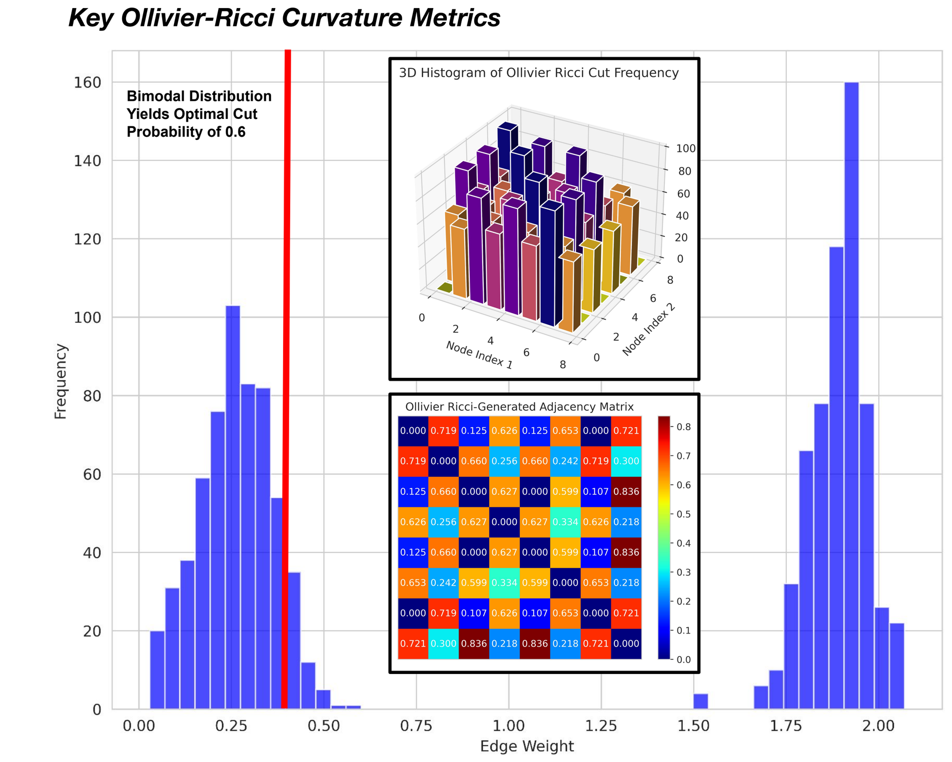

The subsequent Ricci flow-generated graph was used to train a downstream GCN with 271K+ parameters performing 8-to-1 dimensionality reduction; Skip connections were implemented to prevent oversmoothing. The post-GCN embedding was fed into a downstream lightweight 1D-CNN classifier to perform final digit vs. non-digit classification.

Results

AT-AT Filtration Stage Pipeline

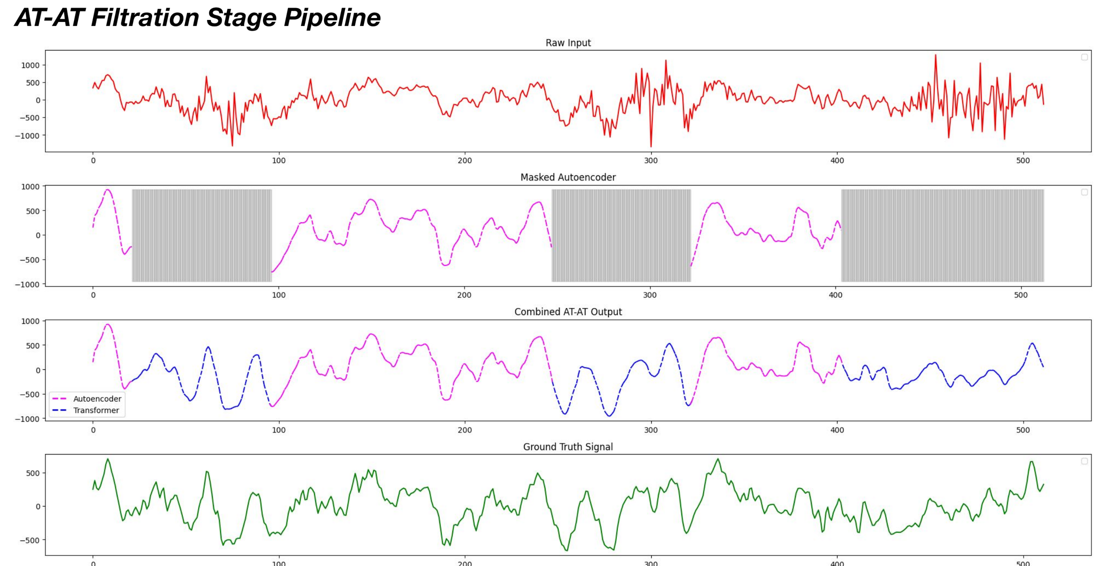

t-SNE Visualization BEFORE Processing via AT-AT

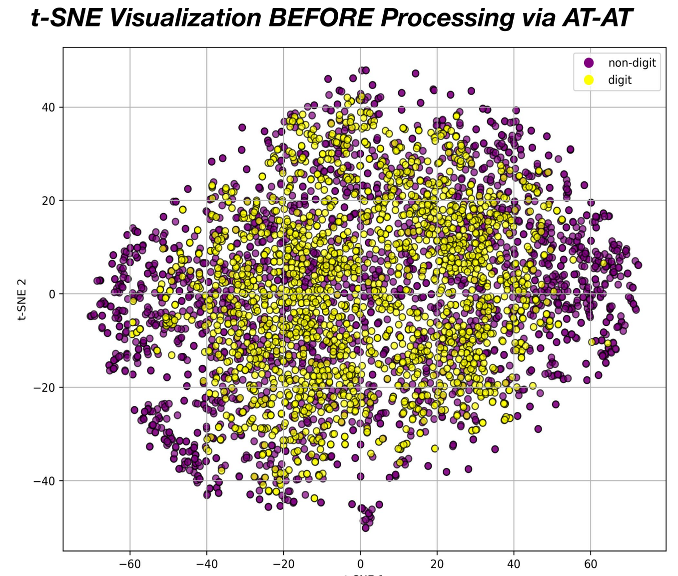

t-SNE Visualization AFTER Processing via AT-AT

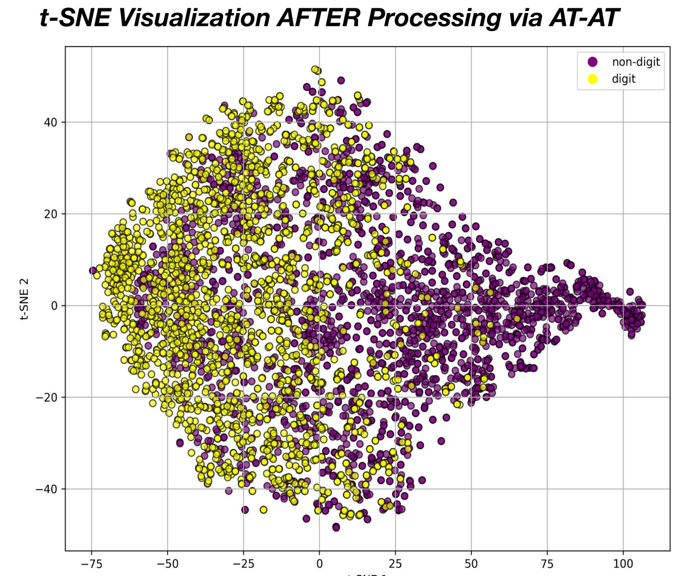

Classification Results on Digit vs. Non-Digit Detection Problem

Metric	Score
Accuracy	0.970
F1 Score	0.968
AUCROC	0.971
Sensitivity	0.978
Specificity	0.963

Confusion Matrix

	TP 0.52	FP 0.02
	FN 0.01	TN 0.45

t-SNE of the Post-GCN Embedding

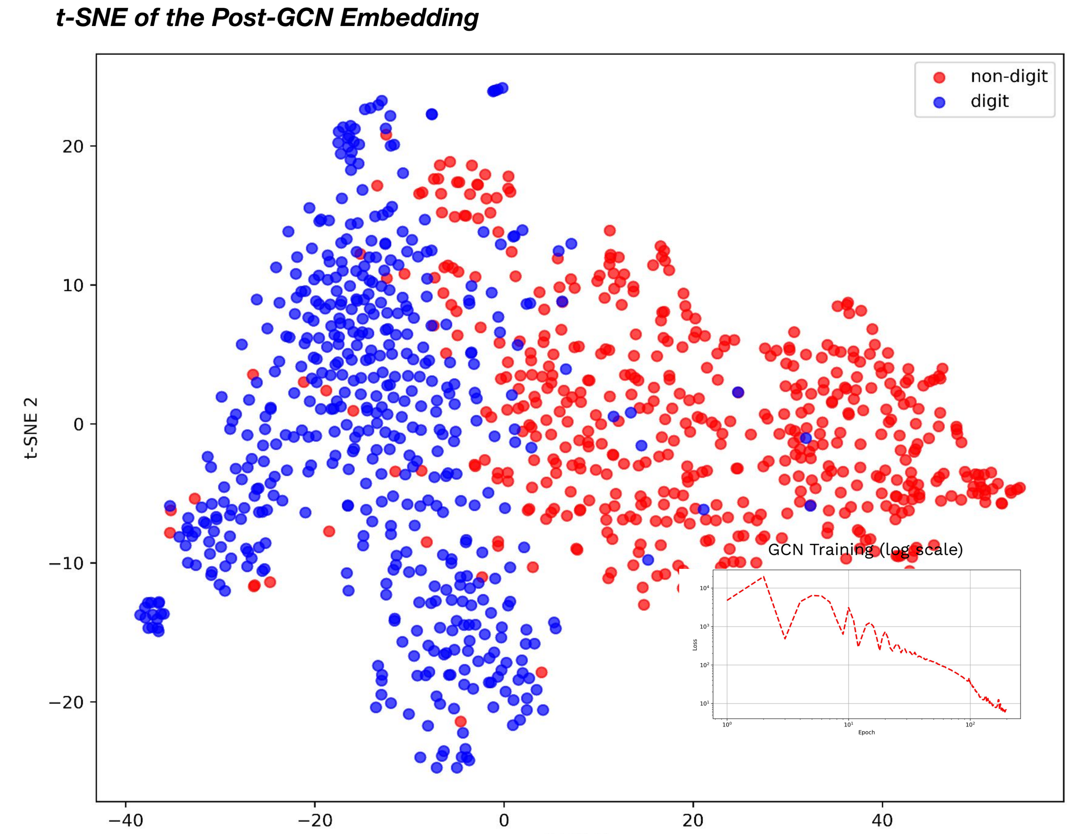

AT-AT model performance on the EEGdenoiseNet task was competitive (*below*) with existing models while maintaining a reduced model footprint. Moreover, AT-AT reduced representational space class ambiguity (*shown above*) when applied to denoise EEG channels in the imagined digit- vs non-digit classification problem. The subsequent geometric machine learning pipeline further fostered increased latent space separability (*left*), enabling high downstream performance on the classification task (*right*).

AT-AT Performance Compared with Existing Models

Model	C-T-S (-7 dB)	C-T-S (2 dB)	Est. Size	TT
Novel CNN (A)	0.69-0.72-0.65	0.92-0.33-0.30	58.7M	unk.
EEGIFNet (B)	*did not test*	0.95-0.32-*unk.*	5.9M	*unk.*
GCTNet (C)	*did not test*	0.94-0.28-*unk.*	~10M	*unk.*
AT-AT (Ours)	0.70-0.76-0.80	0.95-0.32-0.27	438K+	249s

Discussion and Future Direction

Our geometric machine learning pipeline (and AT-AT denoising model) demonstrate promise for uncovering signal within neural data. Beyond direct neural interfacing applications in the realm of imagined digit classification, the manifold learning pipeline presented in this study could open doors to further research on the underlying semantic geometric structure present in high-dimensional brainwave data. The t-SNE visualizations (*left*) offer a preliminary indication that raw high-dimensional EEG signals could contain an underlying geometric structure corresponding to semantic thoughts. More extensive validation of this early idea could open the door to future studies on EEG-based thought modeling with implications for uncovering fundamental properties of intelligence. However, it is crucial to note that more extensive studies will be needed to generalize these results to larger sample sizes, broader user studies, different hardware systems, and wider-ranging use cases.