A lightweight, physics-based, sensor-fusion filter for real-time EEG denoising and improved downstream AI classification

Jack-Michael Wesierski*[1] and Nicholas Rodriguez[1]

[1]MindsApplied, California, USA

Article category: New Results

Keywords: EEG, artifact removal, synchrony, real-time, classification, denoising, streaming

Abstract

Physiological time-series data, like electroencephalography (EEG), are vulnerable to motion, ocular, and muscle artifacts that hinder real-time inference and bias offline analyses. We present the Minds AI Filter: a lightweight, physics-based, sensor-fusion method that exploits multichannel spatial structure and band-aware synchrony to enhance neural activity while suppressing non-neural noise. The nomenclature "AI" reflects integration within a larger artificial-intelligence pipeline; the filter itself requires no prior training or deep learning. A single tuning parameter controls filter strength. The design supports streaming windows (≈1 s) with minimal added latency and extends naturally to longer offline segments; leveraging a sensor-fusion design across channels, it suggests applicability to other neurophysiological time-series, such as MEG and ECoG, pending further validation; exploratory incorporation of EOG/ECG as auxiliary signals is a potential avenue for future filter advancements.

We evaluate the approach across multiple devices and public datasets, assessing both downstream AI classification performance and real-time signal-quality metrics. In both real-time and offline settings, the filter performed better on dynamic artifacts and noise than baseline and commonly used alternatives in our evaluations. When applied in conjunction with other methods, it was only observed to improve downstream accuracy, never reduce it, when any effect was present. Denoising is quantified using SNR-like measures, and ablations isolate the roles of spatial coupling and band weighting. Artifact-specific analyses (ocular bursts, head tilt, jaw clench) and latency profiling on commodity hardware are included. These results indicate that a lightweight, synchrony-aware filter can robustly stabilize real-time EEG and systematically improve downstream AI classification. The method is compatible with standard preprocessing but does not depend on it.

*Correspondence: contact@minds-applied.com

1 Introduction

Electroencephalography (EEG) enables non-invasive access to brain dynamics for applications such as affect recognition, inner-speech interfaces, and motor-imagery control, but its low amplitude and non-stationarity make it highly susceptible to motion, ocular, and myogenic artifacts, especially in real-world, real-time use. Conventional preprocessing (e.g., notch/band-pass, common average referencing, PCA/ICA, adaptive filters, ASR) can be effective yet often involves assumptions or supervision (clean baselines, component labeling) that are difficult to satisfy online [2–5, 8].

We introduce the Minds AI Filter: a physics-informed, multichannel denoising method that leverages synchrony across channels to suppress non-neural components while preserving neural structure [1, 7, 9, 10]. "AI" is included in the naming convention due to its part in a larger artificial intelligence pipeline, but the filter does not require prior training or deep learning. The filter is controlled by a single scalar hyperparameter (λ), runs efficiently on short streaming windows, and extends naturally to offline analysis. It can be used on its own and remains compatible with standard preprocessing steps. Although developed for EEG, the design may extend to MEG and ECoG, which share multichannel oscillatory structure; we also anticipate exploratory incorporation of EOG/ECG as auxiliary signals in future filter advancements. Formal validation is warranted.

Contributions. (i) a multichannel, synchrony-aware filtering approach; (ii) a simple method with a single hyperparameter (λ) suitable for streaming; (iii) comprehensive evaluation across devices, datasets, and downstream AI classification; and (iv) ablation and latency measurements demonstrating real-time viability.

2 Related Work

Adaptive filtering (e.g., LMS/RLS), subspace methods (PCA/ICA), spatial filters (CAR/Laplacian), and artifact subspace reconstruction (ASR) are widely used [2–5, 8]. Synchrony-aware and physics-informed formulations, including oscillator-coupling perspectives related to the Kuramoto family, offer complementary priors for neural denoising and structure discovery [1, 7, 9, 10]. Our approach differs by providing a single-hyperparameter, streaming-friendly filter that composes cleanly with downstream pipelines and artifact-specific steps, without requiring component labeling.

3 Methods (high-level)

3.1 Signal model and objective

Let $X \in \mathbb{R}^{C \times T}$ denote multi-channel EEG. We estimate a denoised $\hat{X}$ by minimizing an objective combining (i) a data-fidelity term with (ii) a regularizer that promotes physiologically plausible synchrony across channels. The regularizer has two parts: (a) **spatial coupling** informed by electrode topology or learned correlations, and (b) **band-aware** weighting to reflect distinct spatial scales in canonical frequency bands. A single scalar λ controls overall regularization strength. A minimal reference implementation is provided in **Supplementary File S1** (`filter.csv`).

3.2 Connectivity prior & band weighting

Electrode geometry (10–20/10–10) or empirical correlations define a normalized connectivity prior. We construct a k-nearest-neighbor channel graph ($k \approx$ 4–6) with distance-decayed weights and use the normalized Laplacian to encourage band-limited spatial compatibility among neighboring channels. Band weights (w_b) reflect physiology—stronger coupling for delta/theta, moderate for alpha/beta, lighter for gamma—and node-degree normalization provides boundary correction so edge electrodes are not over-regularized.

3.3 Hyperparameter (λ) and streaming

λ modulates aggressiveness vs. preservation. For low-latency windows ($\approx$1 s), tighter settings are effective; longer windows ($\geq$30–60 s) typically favor more conservative settings. In practice, tuning with window length and sampling rate is recommended. Overlapped streaming windows maintain continuity; no component labeling is required.

3.4 Evaluation metrics

Offline Artificial Intelligence:

- Balanced accuracy of emotional valence classification using logistic regression on features derived from wavelet coefficients. Further technical architecture will be included in the final submission.

Real-time signal-quality (per $\sim$1 s window unless noted):

- **SNR (power ratio):** $\mathbb{E}[s^2]/\mathbb{E}[n^2]$ (optional dB $= 10\log_{10}(\cdot)$); where $s = x_{\text{filt}}$ and $n = x_{\text{raw}} - x_{\text{filt}}$.
- **SNR (amplitude ratio):** $\mathbb{E}[|s|]/\mathbb{E}[|n|]$ (optional dB $= 20\log_{10}(\cdot)$).

- **SNR (variance ratio):** $\mathrm{Var}(s)/\mathrm{Var}(n)$ (optional dB $= 10\log_{10}(\cdot)$).
- **Artifact suppression (peak drop):** $\%\Delta_{\text{peak}} = 100 \cdot \dfrac{\max_t |x_{\text{raw}}(t)| - \max_t |x_{\text{filt}}(t)|}{\max_t |x_{\text{raw}}(t)|}$; default flag at $\geq 20\%$ reduction.
- **Baseline drift (mean/median shift):** $\Delta_\mu = \mu_{\text{filt}} - \mu_{\text{raw}}$, $\Delta_{\text{med}} = \text{med}_{\text{filt}} - \text{med}_{\text{raw}}$; default flag at $|\text{shift}| \geq 5\ \mu\text{V}$.
- **Variance smoothing:** $\%\Delta\sigma^2 = 100 \cdot \dfrac{\sigma^2_{\text{raw}} - \sigma^2_{\text{filt}}}{\sigma^2_{\text{raw}}}$; default flag at $\geq 5\%$ reduction.
- **Channel aggregation:** metrics may be computed per channel and/or on the mean across EEG channels only.

Implementation details for these streaming metrics are provided in **Supplementary File S2** (`Minds_AI_Filter_Real-time_Signal_Analysis.txt`).

3.5 Reference implementation

We include a reference Python implementation of the Minds AI Filter as **Supplementary File S1** (`filter.txt`). The public method is a single convenience function that operates directly on a 2-D array of shape **[channels, timepoints]**:

```
import mindsai_filter_python as mai

# X: numpy array shaped [channels, timepoints]
X_filt = mai.mindsai_python_filter(X, lambda_val=1e-25)
```

The code computes channel phases via the analytic signal (Hilbert transform), constructs a synchrony operator from phase differences, and applies a closed-form multichannel filter controlled by λ (see Listing S1 for details). This function is **stateless** and requires no model fitting or prior training. The supplemental file is versioned with the preprint; a persistent archive (DOI) will be added upon submission.

4 Datasets & Experimental Setup

Devices. OpenBCI Cyton+Daisy (16 channels, 125 Hz), Neurosity Crown (8 channels, ≈256 Hz), BrainBit Dragon (≥16 channels, up to 500 Hz).
Public datasets. Affective state datasets such as DEAP (32 subjects, 32 channels & 512 Hz) [6], and others where indicated.

Paradigms & artifacts. Resting/task segments; eyebrow raises, head tilt/motion, blinks, jaw clench.

Baselines. Raw, notch+band-pass, CAR/Laplacian, PCA/ICA (with standard component rejection), ASR, LMS-style adaptive filters.
Ablations. (i) remove connectivity prior; (ii) use uniform band weights; (iii) sweep λ.

5 Results

5.1 Real-time denoising & signal quality

On short windows ($\approx$1 s), the filter suppresses motion- and ocular-locked bursts while preserving underlying rhythms. SNR-like metrics (power, amplitude, variance ratios), baseline-drift reduction, and variance smoothing improve consistently across devices. Representative traces (raw, filtered, residual) and per-channel summaries are provided in figures; see Fig. 1.

5.2 Downstream AI classification (positive vs. negative valence)

We evaluated a two-class valence decoder (positive vs. negative) using 60 s windows. On the Dragon EEG device, the best Minds AI Filter setting achieved participant-wise balanced accuracy up to 0.73 (best $\lambda = 10^{-5}$), with typical subject means in the 0.60–0.70 range. Relative to a band-pass–only pipeline, the Minds AI Filter improved balanced accuracy by $\sim$ 14%–20% on this headset; gains on the separate, prefiltered wet-electrode DEAP dataset were expectedly smaller with an average of ($\sim$ 6%) but maximum individual gain of ($\sim$ 35%). We observed no cases where adding the Minds AI Filter reduced downstream classification; when combined with other preprocessing, effects were neutral or positive. Figure 2 summarizes cross-pipeline distributions; subject-wise changes are shown in Fig. 3.

5.3 Comparisons to established filters

We compared against representative families: (i) **band-pass/notch + CAR/Laplacian** spatial filters; (ii) **PCA/ICA** pipelines with standard component rejection; (iii) **adaptive filters** (e.g., LMS variants) for ocular/motion contamination; and (iv) **ASR-style** artifact subspace methods. In real-time windows, our method maintained higher SNR-like metrics under dynamic artifacts and yielded equal or better downstream classification than each baseline across devices. Offline (longer windows), our method matched or exceeded alternatives while avoiding component-labeling steps. Aggregate comparisons and subject-level distributions are summarized in Figs. 3 and 2.

5.4 Latency & real-time viability

Profiling on commodity laptops shows low added latency compatible with online feedback and inference. We report latency as a function of window/hop size and sampling rate, together with throughput and CPU utilization. Per 60 second windows of data: **Mean computational latency:** 0.2637 s **(SD:** 0.1377 s**)**.

6 Discussion

The Minds AI Filter provides a practical, device-agnostic improvement for online EEG by leveraging multichannel synchrony with a single hyperparameter. It reduces large, transient artifacts while retaining physiologically plausible structure, which stabilizes downstream features for AI models. Limitations include sensitivity to very sparse montages ($\leq$4 electrodes) and the need to tune λ with window length and sampling rate. Future work includes adaptive λ, data-driven connectivity priors, expanded tasks (motor imagery, inner speech), artifact-specific benchmarks, and validation on other physiological time-series.

7 Ethics & Consent

All in-house recordings were obtained with informed consent under protocols compliant with institutional and national guidelines. Details of approvals, recruitment, and anonymization will be provided in the final manuscript.

8 Data & Code Availability

We provide **Supplementary File S1** (`filter.txt`) containing a minimal reference implementation sufficient to reproduce the figures/tables, and **Supplementary File S2** (`Minds_AI_Filter_Real-time_Signal_Analysis.txt`) with the exact real-time metric computations used in this work. The code package and real-time test app (with potential version updates) can be found at `https://github.com/MindsApplied/Minds_AI_EEG_Filter`. The same files (released by MindsApplied) will be archived with a DOI at submission. Public datasets are cited; in-house data will be shared in de-identified form where permitted.

9 Competing Interests

Authors are affiliated with MindsApplied, which develops and licenses the described filter.

References

[1] J. A. Acebrón, L. L. Bonilla, C. J. Pérez Vicente, F. Ritort, & R. Spigler (2005). The Kuramoto model: A simple paradigm for synchronization phenomena. *Reviews of Modern Physics*, 77(1), 137–185.

[2] O. Bertrand, F. Perrin, & J. Pernier (1985). A theoretical justification of the average reference in topographic evoked potential studies. *Electroencephalography and Clinical Neurophysiology/Evoked Potentials Section*, 62(6), 462–464.

[3] A. de Cheveigné, & I. Nelken (2019). Filters: When, Why, and How (Not) to Use Them. *eNeuro*, 6(4). `https://doi.org/10.1523/ENEURO.0071-19.2019`

[4] P. He, G. Wilson, & C. Guan (2004). Removal of ocular artifacts from electroencephalogram by adaptive filtering. *Medical & Biological Engineering & Computing*, 42(3), 407–412.

[5] I. T. Jolliffe (2002). *Principal Component Analysis*. Springer.

[6] S. Koelstra, C. Muhl, M. Soleymani, *et al.* (2012). DEAP: A database for emotion analysis; using physiological signals. *IEEE Transactions on Affective Computing*, 3(1), 18–31.

[7] Y. Kuramoto (1975). *Self-entrainment of a population of coupled non-linear oscillators*. In H. Araki (Ed.), *International Symposium on Mathematical Problems in Theoretical Physics* (Lecture Notes in Physics, Vol. 39). Springer.

[8] S. Makeig, A. J. Bell, T.-P. Jung, & T. J. Sejnowski (1996). Independent component analysis of electroencephalographic data. *Advances in Neural Information Processing Systems*, 8, 145–151.

[9] W. Singer (2018). Neuronal oscillations: unavoidable and useful? *European Journal of Neuroscience*, 48, 2389–2398.

[10] D.-L. Wang (1995). Emergent synchrony in locally coupled neural oscillators. *IEEE Transactions on Neural Networks*, 6(4), 941–948.

[11] J.M. Wesierski, & N. Rodriguez (2025). *Minds AI Filter: Signal Filter for Multi-Channel Sensor Data Using Physics-Informed Sensor Fusion and Global/Regional Synchrony Detection* (U.S. Provisional Patent Application). MindsApplied, Inc.

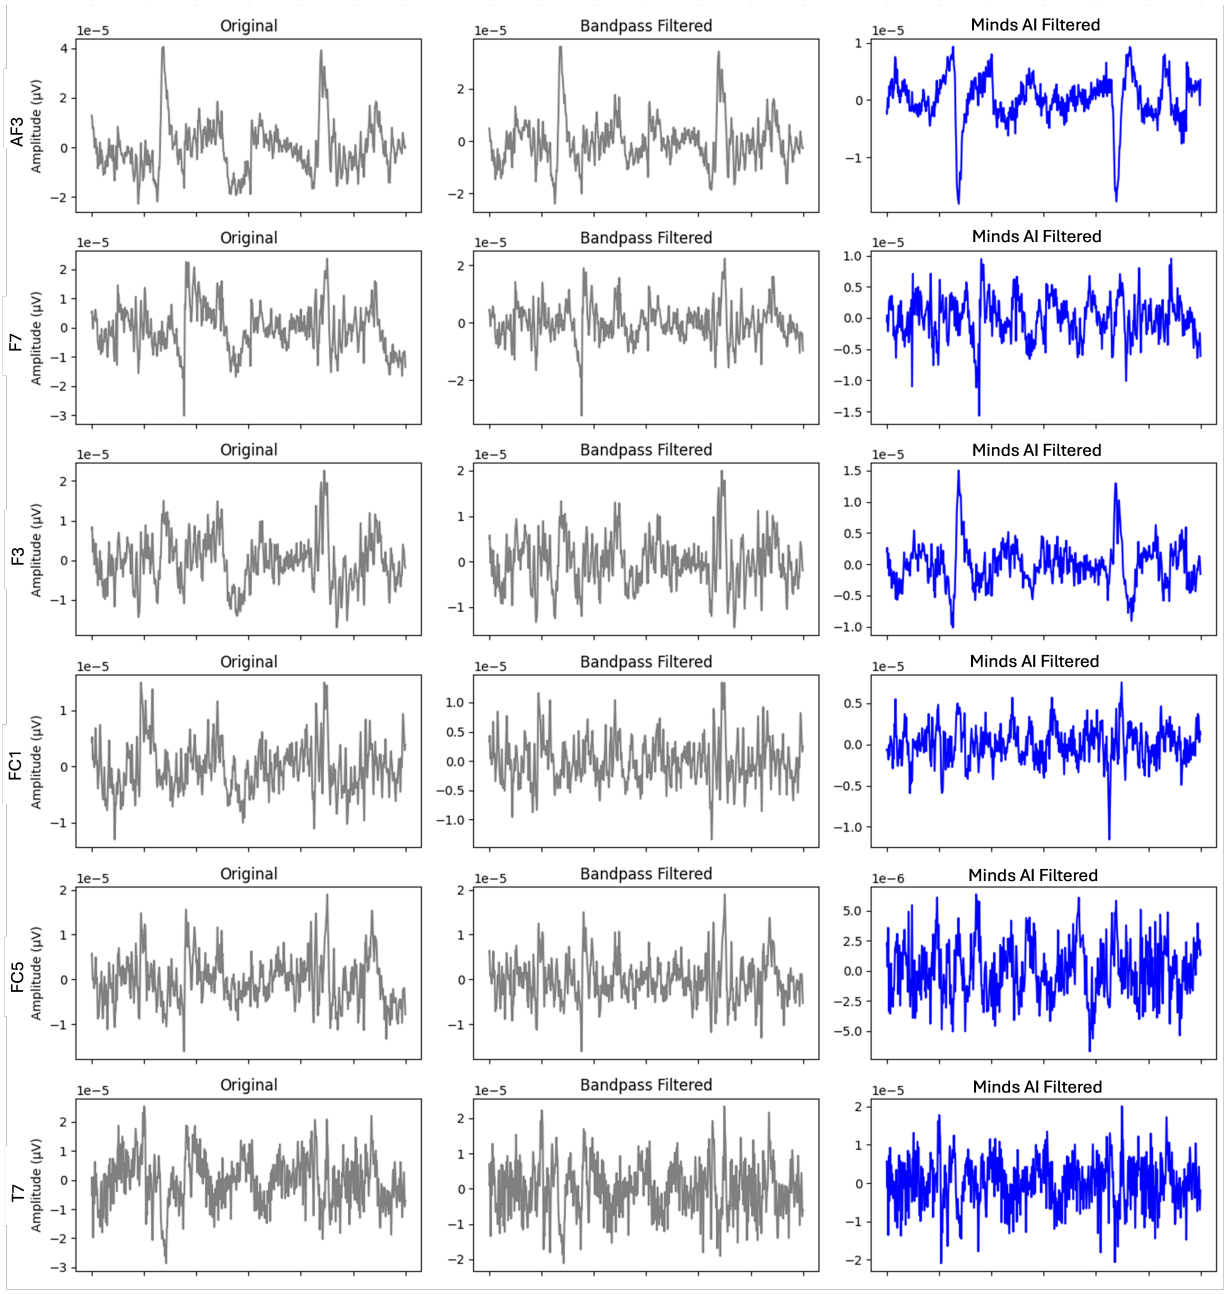

Figure 1: Channel-wise EEG denoising (side-by-side). For each electrode, we show the original signal (left), band-pass baseline (middle), and the Minds AI Filter output (right). The Minds AI Filter suppresses transient motion/ocular bursts while preserving underlying rhythms across channels.

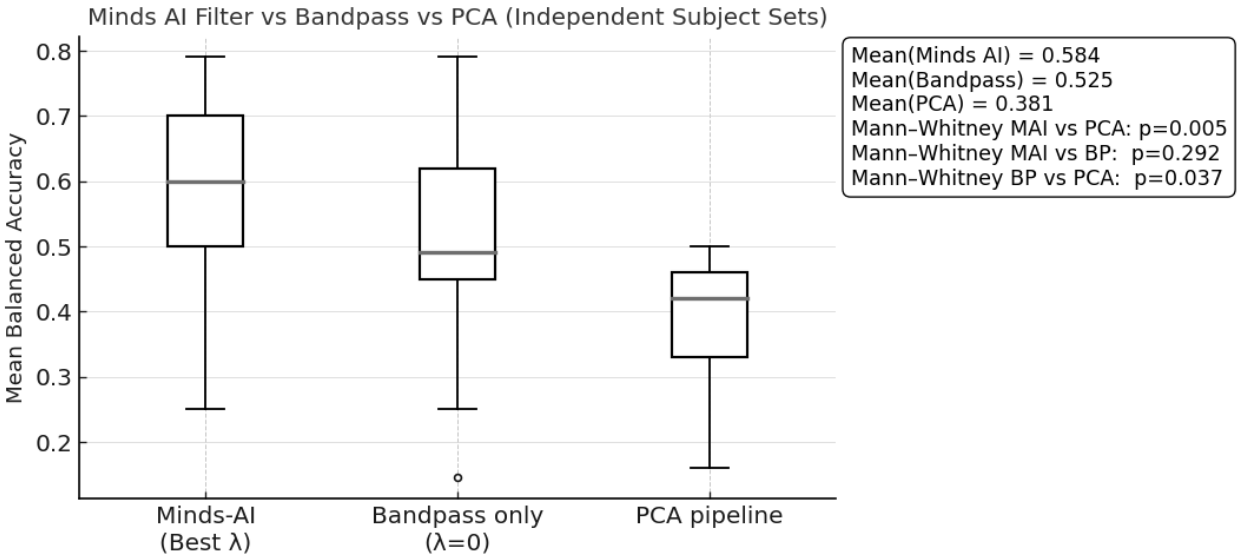

Figure 2: Comparison of mean balanced accuracy across pipelines. Boxes indicate interquartile range with whiskers to non-outlier range; per-pipeline means and Mann–Whitney p-values are shown in the inset.

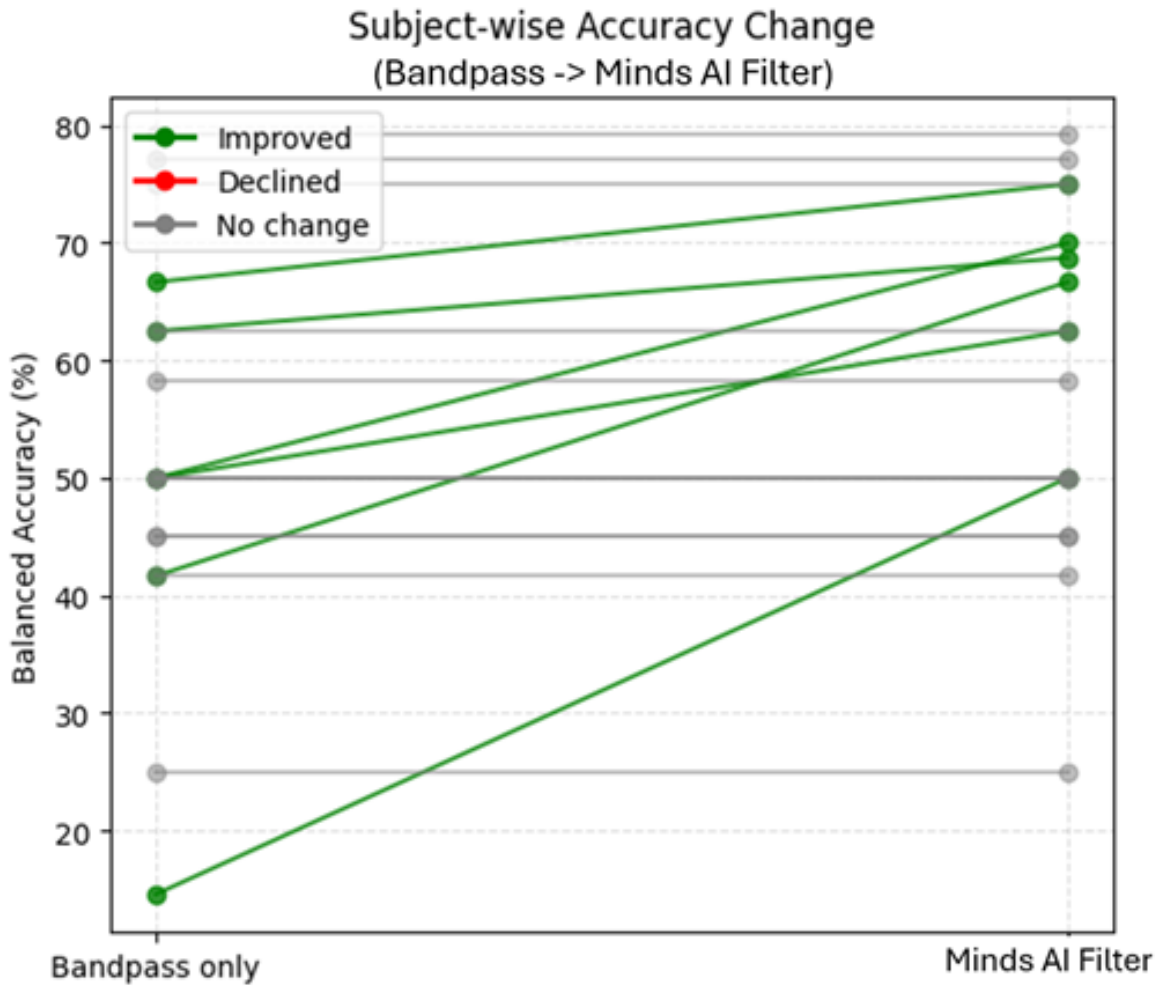

Figure 3: Subject-wise balanced accuracy change from a band-pass–only baseline to the Minds AI Filter–preprocessed pipeline. Each line connects a subject's performance between conditions; green = improved, red = declined, gray = no change.