Persistent Homology Reveals Meditation Expertise-Related Regularity in Auditory Oddball EEG Dynamics

Chandan Srivastava[a,*], Tapas Rath[b,*], Jyotiranjan Beuria[c], Rashmi Gupta[a]

[a] *Department of Humanities and Social Sciences, Indian Institute of Technology Bombay, India*
[b] *Department of Cognitive Science, Indian Institute of Technology Kanpur, India*
[c] *Indian Knowledge System and Mental Health Center, Indian Institute of Technology Mandi, India*

Abstract

Mismatch negativity (MMN) is widely used to probe automatic auditory change detection, but conventional event-related potential (ERP) amplitudes may miss differences in the organization of auditory responses that give rise to MMN difference waveforms. We re-analysed Fz-channel EEG from an intensity oddball paradigm comparing experienced mantra meditators and novice practitioners (27 per group). Conventional MMN amplitudes showed no reliable expertise effects, consistent with the original analysis. We then embedded raw single-trial Fz waveforms from five oddball bins using Takens delay coordinates and quantified Vietoris-Rips persistent homology in dimensions H0 and H1. The topology analysis was therefore applied to condition-specific auditory responses rather than to deviant-minus-standard MMN difference waveforms. After exploratory screening, we locked a restricted

*Corresponding authors: Chandan Srivastava (srivastava.chandan@iitb.ac.in) and Tapas Rath (tapasr@iitk.ac.in).

H0 entropy family and reran 20 tests with 10,000 permutations. Experts showed higher mean H0 entropy and lower across-trial entropy variability in key deviant and standard bins after false-discovery-rate correction across this locked family (minimum $q = 0.000444$). Exploratory classification remained secondary. These findings suggest that meditation expertise may be associated with more stable frontocentral auditory response geometry, even when average MMN amplitudes are comparable. The results support persistent homology as a complementary tool for interrogating neural response regularity in auditory cognitive neuroscience.
Keywords: persistent homology, topological data analysis, EEG, mismatch negativity, auditory oddball, mantra meditation, neural variability

1. Introduction

Focused-attention meditation is proposed to alter attentional control and sensory processing, but EEG evidence from mismatch negativity (MMN) paradigms has been mixed [1–3]. Broader meditation-neuroscience syntheses similarly note heterogeneous electrophysiological effects across paradigms and practice styles [4–6]. The original analysis of the present intensity-oddball dataset found no statistically significant differences between experienced mantra meditators and novice practitioners in conventional Fz MMN amplitude. This null amplitude result is scientifically informative, but it leaves open whether meditation expertise is reflected in other properties of auditory response dynamics.

Topological data analysis (TDA) offers one way to test this possibility. By reconstructing a waveform as a delay-coordinate point cloud and summarizing its persistent homology, TDA can capture geometric organization that is not reducible to mean voltage within a fixed component window [7–10]. In EEG, this is especially attractive when the hypothesis concerns trial-level stability, stereotypy, or dynamical regularity rather than a simple increase or decrease in ERP amplitude.

Here we ask whether persistent-homology summaries of Fz auditory oddball responses distinguish experienced mantra meditators from novices in the same dataset where conventional ERP amplitudes did not. We treat Fz as an a priori frontocentral channel for MMN-style auditory responses [11, 12]. This framing follows core MMN literature emphasizing predictive-auditory context effects and careful interpretation in oddball designs [13–15]. We focus the interpretation on response geometry and neural regularity.

2. Materials and Methods

2.1. Participants and auditory oddball task

The present analysis reused the curated dataset reported in the original study [16] and did not recruit new participants. The source study was approved by the Ethical Review Board of the Indian Institute of Technology Bombay, and informed consent was obtained from all participants. The cohort included 54 participants (27 experienced mantra meditators; 27 novice practitioners). Experts had at least 10 years of practice (about 120 minutes/day), whereas novices had 1 year or less of practice (about 10–20 minutes/day). Group ages were comparable (experts: mean 29.9 years, SD 3.8; novices: mean 28.1 years, SD 6.4). The sample was predominantly male (25 males and 2 females in each group, as reported in the source dataset), with normal hearing and no reported psychiatric or neurological disorder history.

Participants completed an auditory intensity oddball paradigm with 80 dB standard tones, 90 dB louder deviants, and 70 dB quieter deviants. Five bins were used for the topology analysis: D90 deviants, S80 standards following standards, D70 deviants, S90 standards following 90 dB deviants, and S70 standards following 70 dB deviants.

2.2. EEG preprocessing and conventional ERP context

Preprocessed EEG set files from the original manuscript were used without altering the preprocessing pipeline. The source preprocessing steps (including filtering, artifact handling, referencing, baseline correction, and epoch construction) were adopted as reported in the original dataset paper [16]. The present re-analysis therefore applies TDA and ERP summaries to those preprocessed waveforms without additional preprocessing changes. The analysis focused on the Fz channel. Fz indexing was consistent across the available montage files and the extracted feature table, and no participant-level Fz/Cz channel-swap history was documented in the available source records; no per-subject channel remapping was therefore applied. Consistent with the source manuscript, conventional MMN amplitude refers to the mean amplitude of the deviant-minus-standard difference waveform (D90-S80 and D70-S80) in the 125-225 ms measurement window [17].

2.3. Topological feature extraction

Unless explicitly stated otherwise, TDA was applied to raw condition-specific Fz waveforms, not to deviant-minus-standard MMN difference waveforms. For each subject/bin, waveforms were restricted to -100 to 496 ms.

For each Fz waveform, we constructed a Takens delay embedding with embedding dimension 5 and delay 3 samples. Point clouds were deterministically landmarked to a maximum of 40 delay-coordinate vectors selected across the embedded trajectory; this cap does not refer to taking the first 40 voltage samples. Pairwise Euclidean distance matrices were normalized by the 95th percentile of positive distances. These settings were chosen as computationally practical defaults for short EEG waveforms and were subsequently evaluated in sensitivity analyses. Vietoris-Rips persistence was computed for H0 and H1. Features included Betti-curve summaries, total and maximum persistence, persistence entropy, normalized entropy, and persistence-image summaries [18].

Two raw-waveform signal levels were analysed. First, ERP-level TDA was applied to each subject/bin average waveform, using all retained epochs in that bin. Second, trial-level TDA was applied to a deterministic cap of 80 single trials per subject/bin and then summarized within subject/bin by mean, standard deviation, and median. This cap kept trial-level point clouds comparable across bins and computationally tractable. When more than 80 trials were available, trial indices were selected at evenly spaced positions across the available epochs; otherwise all available trials were used. No random trial sampling was used. The primary expertise-related effects emerged from trial-level H0 entropy summaries of raw auditory response waveforms.

2.4. Statistical analysis

Expert-vs-novice group comparisons were tested within each bin using Welch tests and label-permutation tests. Within-subject condition contrasts used paired tests and sign-flip permutation tests. P-values were adjusted using the Benjamini-Hochberg false-discovery-rate procedure [19]. Bootstrap confidence intervals were computed for selected effects [20], and leave-one-subject-out summaries assessed effect-size sign stability. Logistic-regression classifiers were used only as exploratory secondary probes.

The analysis workflow included an exploratory first pass over ERP and TDA feature summaries. The first-pass screening table used for feature-family selection contained 169 numeric ERP/TDA feature summaries, yielding 843 condition-specific expert-vs-novice group comparisons across the five auditory bins. After that screening identified trial-level H0 entropy features as the most consistently group-separating signal across auditory bins, we froze a restricted primary family consisting of four H0 entropy summaries (trial mean entropy, trial standard deviation of entropy, and their normalized forms) across five auditory bins. We then reran these 20 locked tests with 10,000 label permutations per test and FDR correction across the locked family. Because this family was locked after exploratory screening rather than prospectively registered, all other feature families, condition-contrast modulation tests, and classifier analyses are treated as exploratory, and the locked-family effect magnitudes require independent replication.

3. Results

3.1. Conventional MMN amplitudes provided the null-effect context

Consistent with the original manuscript, conventional Fz MMN amplitudes did not show reliable expert-vs-novice differences (Fig. 1). In the source analysis, D90MMN amplitudes did not differ significantly between novices

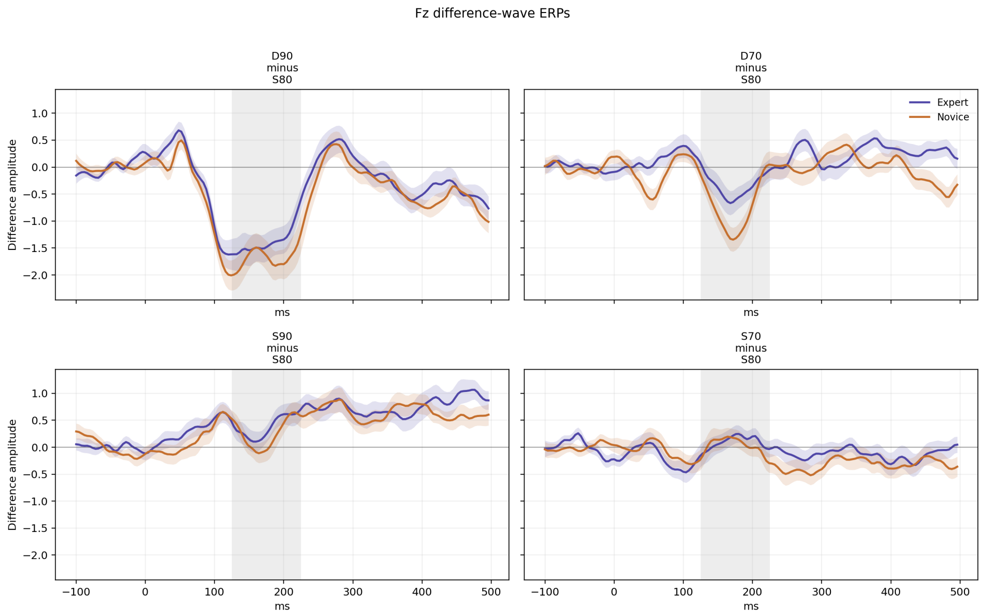

Figure 1: Fz difference-wave ERPs for primary condition contrasts. Shading indicates the MMN measurement window. Conventional MMN amplitudes provide physiological context but did not reliably distinguish expertise groups. The topology statistics reported below were computed from raw condition-specific waveforms rather than from these difference waves.

(M = -2.16, SD = 1.26) and experts (M = -1.77, SD = 1.07), $t(52) = -1.22$, $p = 0.229$, $d = -0.33$. D70MMN amplitudes were similarly non-significant for novices (M = -0.67, SD = 0.86) versus experts (M = -0.39, SD = 0.99), $t(52) = -1.12$, $p = 0.267$, $d = -0.31$. Likewise, individual-condition ERP mean amplitudes in the topology output table showed no reliable expertise effects for D90 (d = 0.32, permutation p = 0.225, q = 0.506) or D70 (d = 0.42, permutation p = 0.127, q = 0.360). These conventional MMN and individual-condition ERP amplitude tests are separate from the topology tests below, which were applied to raw condition-specific waveforms.

3.2. Trial-level H0 topology showed the clearest expertise-related pattern

The clearest group effects involved trial-level H0 persistence entropy extracted from raw condition-specific waveforms. In the locked rerun (10,000 permutations; 20 tests), experts showed higher mean H0 entropy than novices for D70 deviants (expert minus novice = 0.01061, d = 0.863, q = 0.00122) and D90 deviants (expert minus novice = 0.01092, d = 0.939, q = 0.000444). Experts also showed lower across-trial standard deviation of H0 entropy for D90 deviants (expert minus novice = -0.00999, d = -0.966, q = 0.000444). Similar effects appeared in S80 standards and S90 standards following D90 deviants. These deviant-bin results indicate expertise-related organization of responses to oddball tones, but they are not direct TDA measures of the MMN difference waveform. Effect-size confidence intervals are shown in Fig. 2, and subject-level distributions are shown in Fig. 3. For clarity, we treat these as two distinct feature categories throughout the manuscript: (i) **Mean-entropy features** (trial mean H0 entropy, including normalized form) and (ii) **Entropy-variability features** (trial SD H0 entropy, including normalized form).

3.3. Condition effects validated sensitivity to auditory structure

Across all subjects, within-subject contrasts of separately extracted condition features produced robust Fz ERP and TDA effects. D90-minus-S80 showed a large conventional mean-amplitude effect (paired d = -1.02, sign-flip q = 0.00661). D70-minus-S80 and S70-minus-S80 also showed H1 topology-feature differences, including beta1 area-under-curve and total-persistence differences. For TDA, these contrasts compare features extracted separately from each raw condition waveform; they are not persistent-homology features

Table 1: Selected expert-vs-novice topology effects, organized into two categories: Mean-entropy features (rows with trial mean H0 entropy) and Entropy-variability features (rows with trial SD H0 entropy). Positive d indicates higher values in experts. Values are from the locked-primary 10,000-permutation group tests, with q values reflecting Benjamini-Hochberg FDR correction across the 20 locked tests.

Condition	Feature	Direction	Cohen's d	FDR q
D70 deviant	Mean H0 entropy	Expert > novice	0.863	0.00122
D90 deviant	Mean H0 entropy	Expert > novice	0.939	0.000444
D90 deviant	SD H0 entropy	Expert < novice	-0.966	0.000444
S80 standard	Mean H0 entropy	Expert > novice	0.808	0.0018
S80 standard	SD H0 entropy	Expert < novice	-0.829	0.000444
S90 standard	Mean norm. H0 entropy	Expert > novice	0.881	0.000444

computed on deviant-minus-standard difference waveforms. These condition effects suggest that the TDA pipeline was sensitive to auditory response structure, even though the primary expertise interpretation should rely on group comparisons.

3.4. Exploratory classifier benchmark

Classifier outputs are retained in the Supplementary Material only. Nested cross-validation suggested that TDA-only features sometimes exceeded ERP-only features for expert-vs-novice classification, but permutation testing and FDR correction did not support classifier performance as confirmatory evidence. We therefore treat classification as exploratory and secondary.

3.5. Robustness and sensitivity analyses

Robustness analyses completed include locked-primary 10,000-permutation inference, trial-count covariate checks, a checkpoint embedding-parameter grid, and matched-parameter Fz/Cz channel-sensitivity reruns. The Fz/Cz rerun showed directionally consistent expert-minus-novice effects for the locked H0 entropy features at Cz, although effect sizes were generally smaller than at Fz (Supplementary Fig. S3).

4. Discussion

The main finding is that persistent-homology features of raw single-trial Fz delay embeddings revealed expertise-related response regularity in a dataset where conventional MMN amplitudes did not reliably separate experienced mantra meditators from novices. The pattern was not a direct topology measure of deviant-minus-standard MMN amplitude; instead, it appeared across standards and deviants, suggesting a trait-like difference in auditory response geometry.

The most consistent effects involved H0 persistence entropy. In this context, H0 summarizes how connected components in the delay-embedded point cloud merge across filtration radii. Higher H0 entropy can be interpreted as a more evenly distributed component-merging profile, while lower across-trial H0 entropy variability indicates greater trial-to-trial consistency. This supports a cautious interpretation in terms of response-geometry regularity, rather than larger MMN amplitude. The mapping between H0 entropy and specific neurophysiological mechanisms remains indirect, and stability or regularity language should be regarded as a heuristic description of the observed geometry.

This interpretation aligns with the original null MMN-amplitude conclusion. It does not imply that mantra meditation enhances or suppresses MMN magnitude, nor does it show that the MMN difference waveform itself has a different topology. Instead, it suggests that expertise-related differences may be expressed in the geometry and variability of the raw auditory responses that contribute to ERP and MMN estimates.

The effect sizes observed in the locked H0 entropy family were large relative to the sample size and should be interpreted cautiously until replicated in independent cohorts. Additional limitations include the cross-sectional design, reliance on curated preprocessed data from the source study, and the fact that the locked primary family was selected after exploratory screening rather than prospective registration. The present findings should therefore be viewed as hypothesis-generating evidence for topology-based markers of meditation expertise until replicated in an independent dataset.

5. Conclusion

Persistent homology revealed a consistent expertise-related topology signal in Fz auditory oddball EEG responses. Experts showed higher trial-level H0 entropy and lower H0 entropy variability than novices, despite comparable conventional MMN amplitudes. These results motivate the use of topology-based analyses as complementary tools for studying neural response regularity in auditory cognitive neuroscience.

CRediT Author Contribution Statement

Chandan Srivastava: Conceptualization, Methodology, Investigation, Data curation, Formal analysis, Writing – original draft, Writing – review & editing. Tapas Rath: Formal analysis, Methodology, Writing – original draft, Writing – review & editing. Jyotiranjan Beuria: Formal analysis, Methodology, Writing – review & editing. Rashmi Gupta: Supervision, Conceptualization, Writing – review & editing.

Declaration of Competing Interest

The authors declare no competing interests.

Funding

No specific funding was received for this work.

Data and Code Availability

The data and code that support the findings of this study are available from the corresponding authors upon reasonable request, subject to institutional ethics and participant-consent constraints.

Declaration of Generative AI and AI-Assisted Technologies in the Manuscript Preparation Process

During the preparation of this work, the authors used ChatGPT and Grammarly to assist with language editing and readability. The authors reviewed and edited the content as needed and take full responsibility for the content of the article.

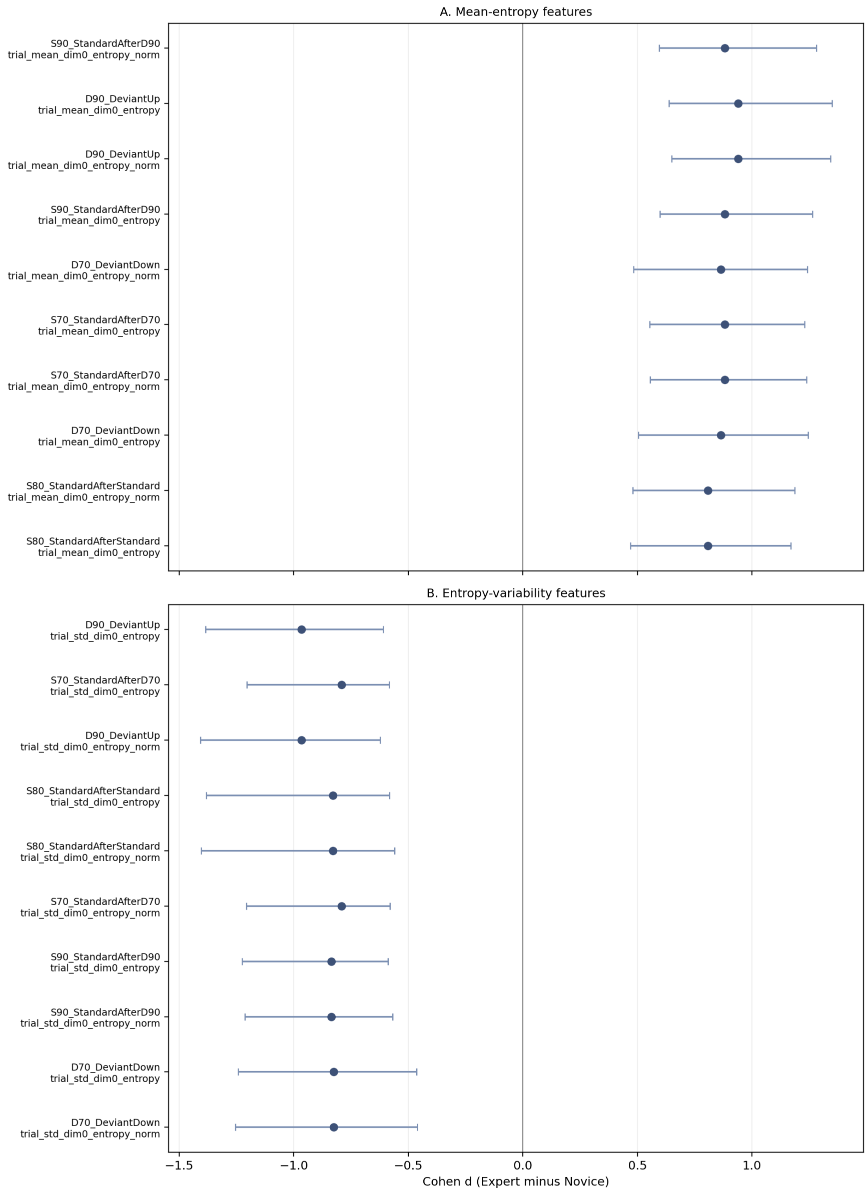

Figure 2: Bootstrap confidence intervals for selected expert-vs-novice Fz topology effects, split into two subpanels: (A) Mean-entropy features and (B) Entropy-variability features. Effects are shown as Cohen's d for experts minus novices.

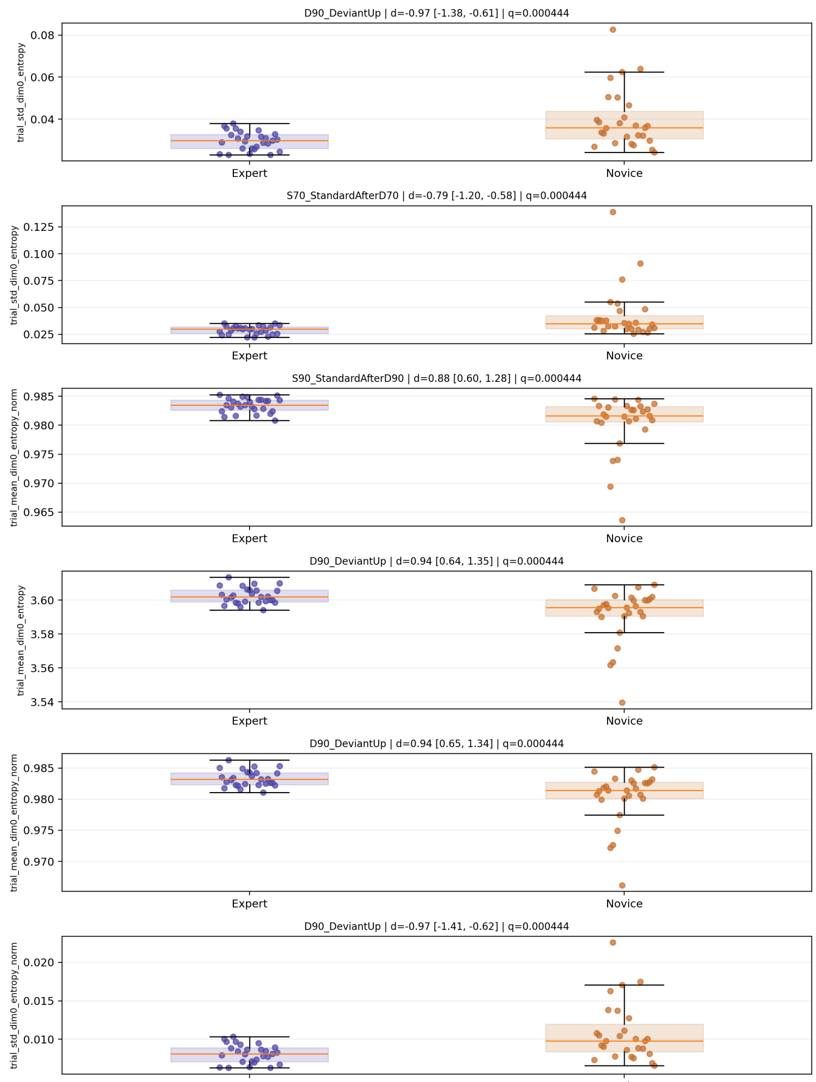

Figure 3: Subject-level distributions for the primary group-separating Fz topology features, split into two subpanels: (A) Mean-entropy features and (B) Entropy-variability features. The strongest effects indicate higher mean H0 entropy and lower across-trial H0 entropy variability among experts. Displayed primary group effects were significant under permutation testing with FDR correction; exact effect sizes and corrected p-values are reported in Table S1.

References

[1] B. R. Cahn, J. Polich, Meditation states and traits: Eeg, erp, and neuroimaging studies., Psychological bulletin 132 (2) (2006) 180.

[2] A. Lutz, H. A. Slagter, J. D. Dunne, R. J. Davidson, Attention regulation and monitoring in meditation, Trends in cognitive sciences 12 (4) (2008) 163–169.

[3] E. Fucci, A. Poublan-Couzardot, O. Abdoun, A. Lutz, No effect of focused attention and open monitoring meditation on eeg auditory mismatch negativity in expert and novice practitioners, International Journal of Psychophysiology 176 (2022) 62–72.

[4] Y.-Y. Tang, B. K. Hölzel, M. I. Posner, The neuroscience of mindfulness meditation, Nature reviews neuroscience 16 (4) (2015) 213–225.

[5] T. Brandmeyer, A. Delorme, H. Wahbeh, The neuroscience of meditation: classification, phenomenology, correlates, and mechanisms, Progress in brain research 244 (2019) 1–29.

[6] P. Sedlmeier, J. Eberth, M. Schwarz, D. Zimmermann, F. Haarig, S. Jaeger, S. Kunze, The psychological effects of meditation: a meta-analysis., Psychological bulletin 138 (6) (2012) 1139.

[7] F. Takens, Detecting strange attractors in turbulence, in: Dynamical Systems and Turbulence, Warwick 1980, Springer, 1981, pp. 366–381.

[8] H. Edelsbrunner, D. Letscher, A. Zomorodian, Topological persistence

and simplification, Discrete & Computational Geometry 28 (4) (2002) 511–533.

[9] R. Ghrist, Barcodes: The persistent topology of data, Bulletin of the American Mathematical Society 45 (1) (2008) 61–75.

[10] J. A. Perea, J. Harer, Sliding windows and persistence: An application of topological methods to signal analysis, Foundations of Computational Mathematics 15 (3) (2015) 799–838.

[11] R. Näätänen, P. Paavilainen, T. Rinne, K. Alho, The mismatch negativity (mmn) in basic research of central auditory processing: a review, Clinical neurophysiology 118 (12) (2007) 2544–2590.

[12] C. C. Duncan, R. J. Barry, J. F. Connolly, C. Fischer, P. T. Michie, R. Näätänen, J. Polich, I. Reinvang, C. Van Petten, Event-related potentials in clinical research: guidelines for eliciting, recording, and quantifying mismatch negativity, p300, and n400, Clinical neurophysiology 120 (11) (2009) 1883–1908.

[13] R. Näätänen, T. Jacobsen, I. Winkler, Memory-based or afferent processes in mismatch negativity (mmn): A review of the evidence, Psychophysiology 42 (1) (2005) 25–32.

[14] P. J. May, H. Tiitinen, Mismatch negativity (mmn), the deviance-elicited auditory deflection, explained, Psychophysiology 47 (1) (2010) 66–122.

[15] E. S. Sussman, S. Chen, J. Sussman-Fort, E. Dinces, The five myths of mmn: redefining how to use mmn in basic and clinical research, Brain Topography 27 (2014) 553–564.

[16] C. Srivastava, J. A. O'Reilly, R. Gupta, Investigating the effects of focused attention (mantra) meditation on mismatch negativity: Insights into sensory and cognitive processing using an intensity oddball paradigm, Neuroscience 571 (2025) 62–73. doi:10.1016/j.neuroscience.2025.02.017.
URL https://doi.org/10.1016/j.neuroscience.2025.02.017

[17] E. S. Kappenman, J. L. Farrens, W. Zhang, A. X. Stewart, S. J. Luck, Erp core: An open resource for human event-related potential research, NeuroImage 225 (2021) 117465.

[18] H. Adams, T. Emerson, M. Kirby, R. Neville, C. Peterson, P. Shipman, S. Chepushtanova, E. Hanson, F. Motta, L. Ziegelmeier, Persistence images: A stable vector representation of persistent homology, Journal of Machine Learning Research 18 (8) (2017) 1–35.

[19] Y. Benjamini, Y. Hochberg, Controlling the false discovery rate: A practical and powerful approach to multiple testing, Journal of the Royal Statistical Society: Series B 57 (1) (1995) 289–300.

[20] B. Efron, R. J. Tibshirani, An Introduction to the Bootstrap, Chapman and Hall/CRC, 1994.