Integrating Slow Neural Oscillations and Physiological Burden for Trait Anxiety Prediction

Jungyoun Janice Min[1] **Jiong Chen**[2] **Jingxuan Bao**[3] **Shu Yang**[2] **Yize Zhao**[4]
Li Shen[2] **Duy Duong-Tran**[2]*

[1]Seoul National University [2]University of Pennsylvania [3]Brown University [4]Yale University

Abstract

Effective modeling of health outcomes from biomedical time series requires methods that capture both temporal and frequency dynamics. Trait anxiety, a stable disposition characterized by heightened anticipatory stress across contexts, manifests through both neural and systemic physiological burden, yet existing approaches rarely integrate these modalities. We present a graph-attention framework that models brain functional dynamics over structural connectivity and integrates them with allostatic-load-related blood biomarkers via cross-modal attention. In 120 young adults from the LEMON dataset, we systematically evaluated four feature extraction strategies and demonstrated that preserving temporal order in anxiety-relevant slow-4/slow-5 oscillations (0.01–0.073 Hz) was critical for stable prediction, while temporal order-discarding approaches consistently underperformed. Multimodal integration provided additional gains over brain-only models. Model interpretability analyses revealed that limbic and visual networks, along with metabolic and immune markers (creatinine, glucose, C-reactive protein), served as the most informative features. Our results show that temporal dynamics in neural oscillations are essential for modeling psychiatric vulnerability and establish a framework for integrating brain–body signals into interpretable digital biomarkers for mental health.

1 Introduction

Trait anxiety is a stable disposition characterized by heightened anticipatory stress across contexts. Neuroimaging has linked it to altered functional connectivity in the default mode and limbic networks [1, 2] from resting state fMRI (rsfMRI), and disrupted white-matter pathways [3]. However, existing brain-centric approaches rarely account for systemic physiological states that shape neural activity. Elevated allostatic load, the cumulative physiological burden from chronic stress adaptation, is associated with trait anxiety and may induce inflammation and excitatory/inhibitory neurotransmission imbalance [4, 5]. These findings suggest the need for models that integrate brain–body associations, as systemic physiological states shape neural activity patterns.

Regional activity metrics such as amplitude of low-frequency fluctuations implicate slow-4/slow-5 oscillations in anxiety disorders [6, 7, 8], yet they capture only amplitude and fail to characterize temporal-frequency dynamics. Prior multimodal efforts have focused on either brain or physiology in isolation, leaving their interaction underexplored.

We propose a multimodal two-stage attention-based framework that jointly models rs-fMRI time-series features on a structural scaffold using graph attention network and fuses them with allostatic-

*Corresponding author: DuyAnh.Duong-Tran@pennmedicine.upenn.edu

39th Conference on Neural Information Processing Systems (NeurIPS 2025) Workshop: Learning from Time Series for Health (TS4H).

load biomarkers via cross-modal attention. A central contribution is the systematic comparison of feature extraction strategies, showing that preserving temporal order in slow-4/slow-5 oscillations is critical for prediction, whereas temporal dynamics-discarding approaches underperform. This framework also enables us to highlight anxiety-relevant brain connectivity patterns, important physiological markers, providing a computational lens on brain–body pathways of anxiety vulnerability.

2 Methods

2.1 Dataset and Participants

We used the open LEMON dataset [9], with multimodal measures (rs-fMRI, diffusion MRI, blood biomarkers, behavior). From 132 young adults, 120 participants (84 Males, age 20–30) with complete neuroimaging, trait anxiety (STAI-20 items[10, 11]), and blood marker data were included. Data were de-identified and relied on the dataset's released preprocessing[12].

2.2 Neuroimaging and Allostatic Load Data

The rs-fMRI blood oxygen level dependent (BOLD) time series signal (TR=1.4s, timepoints = 522) and diffusion MRI were preprocessed by the standard preprocessing pipelines (motion correction, distortion correction, denoising, spatial normalization to MNI152) and parcellated using 183 data-driven regions via voxel-wise clustering of rs-fMRI signals within macro-anatomical regions [12]. Structural connectivity was quantified as streamline counts from diffusion MRI. Allostatic load markers spanned cardiovascular, metabolic, and immune systems (10 biomarkers including, Systolic/Diastolic Blood Pressure, Body Mass Index, Total Cholesterol, Low-Density Lipoprotein Cholesterol, High-Density Lipoprotein Cholesterol, Total Cholesterol to HDL-C Ratio, and Glucose, Creatinine, C-Reactive Protein), following prior anxiety work [13, 14, 15]. Each biomarker was treated as an individual feature without constructing a composite AL marker.

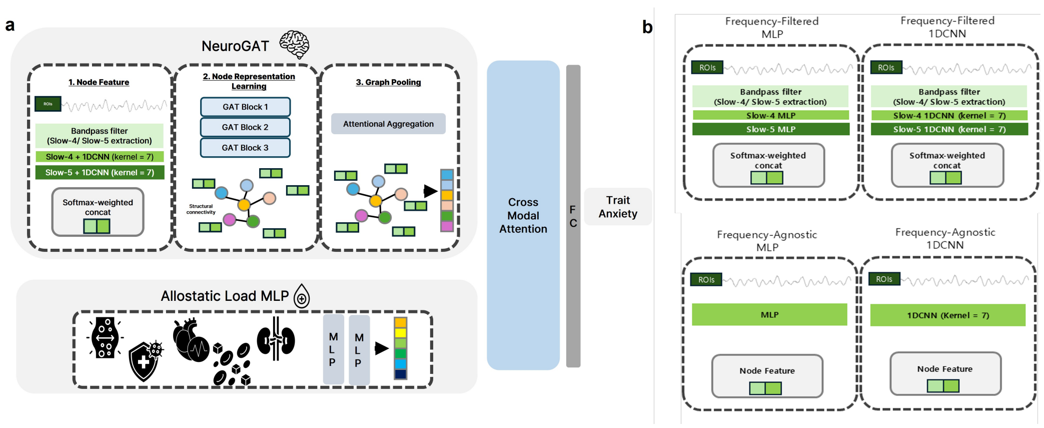

Figure 1: Overview of the model.

2.3 Model Architecture

We propose a multimodal framework (AlloNeuroGAT) combining brain and physiology (Fig. 1a).

NeuroGAT : Three GATv2Conv layers with residual connections, each followed by LayerNorm, LeakyReLU, and dropout (0.2). Node features were represented with features extracted from rs-fMRI time series, and diffusion MRI–derived structural connectivity provides the binary adjacency matrix (Fig.1b). Strategies for node feature extraction are as follows:

- **Frequency-Filtered MLP:** Per-band MLP on slow-4 (0.027–0.073 Hz) and slow-5 (0.01–0.027 Hz) band-pass filtered signals.
- **Frequency-Filtered 1D-CNN:** Per-band 1D-CNN (kernel=7) on slow-4 and slow-5 band-pass filtered signals.
- **Frequency-Agnostic MLP:** A single MLP on unfiltered signals.

- **Frequency-Agnostic 1D-CNN:** A single 1D-CNN (kernel=7) on unfiltered signals.

  For the frequency-filtered models, we used softmax-weighted concatenation to combine the outputs of the per-band models, generating a 64-dimensional node feature vector for each brain region.

Allostatic load markers Projection: A two-layer MLP (Leaky ReLU, dropout=0.2) embeds ten biomarkers into a 64-d representation. Inputs are z-scored within each train/val/test splits.

Cross-Modal Attention: Brain representation (query) attends to AL representation (key/value) via 2-head cross-modal attention, producing a fused embedding passed to a linear head for trait anxiety prediction.

2.4 Training and Evaluation

Models were trained with AdamW and MSE loss with early stopping (patience = 20). The learning rate and weight decay were selected via a grid search over learning rate = [0.001, 0.0005, 0.0001, 0.00005] and weight decay = [0.001, 0.01] based on the validation set performance. Trait anxiety scores were z-scored within each data split. Performance was evaluated using nested 3-fold subject-level cross-validation, repeated five times with different splits and reported on held-out tests using r^2, MSE, and Pearson correlation coefficient (Pearson r).

3 Experimental Results

Table 1: Performance Comparison of Feature Extraction Strategies in AlloNeuroGAT

Model	MSE (↓)	r^2 (↑)	Pearson r (↑)
AlloNeuroGAT-Freq_filtered_MLP	1.008 ± 0.06	-0.034 ± 0.061	0.115 ± 0.056
AlloNeuroGAT-Freq_agnostic_MLP	1.005 ± 0.051	-0.030 ± 0.052	0.105 ± 0.082
AlloNeuroGAT-Freq_agnostic_1DCNN	0.940 ± 0.035	0.036 ± 0.036	0.206 ± 0.072
AlloNeuroGAT-Freq_filtered_1DCNN	$\mathbf{0.924 \pm 0.028}$	$\mathbf{0.052 \pm 0.028}$	$\mathbf{0.257 \pm 0.05}$

3.1 Performance of Node Feature Extraction Strategies

We observed significant performance differences across four temporal feature extraction strategies (Repeated Measures ANOVA; MSE: $p = 0.0091$; r^2: $p = 0.0209$; $pearsonr$: $p = 0.0009$). The frequency-filtered 1D-CNN consistently outperformed others, achieving significantly higher Pearson correlation coefficients than both frequency-filtered ($p\_corrected = 0.0379$) and frequency-agnostic ($p\_corrected = 0.0369$) MLP models (post-hoc Paired T-tests). This confirms that explicitly modeling temporal dynamics within anxiety-relevant slow oscillation bands slow-4 and slow-5 bands [6] is superior to disregarding moment-to-moment temporal dynamics.

Table 2: Multimodal Model Performance Comparison

Model	MSE (↓)	r^2 (↑)	Pearson r (↑)
NeuroGAT_Only	0.986 ± 0.022	-0.012 ± 0.023	-0.068 ± 0.081
ALMLP_Only	0.948 ± 0.051	0.028 ± 0.052	0.187 ± 0.095
AlloNeuroGAT-Node_coordinate	0.951 ± 0.026	0.025 ± 0.026	0.198 ± 0.055
AlloNeuroGAT-Edge_FC	0.936 ± 0.049	0.040 ± 0.050	0.204 ± 0.091
AlloNeuroGAT-QAL/KV_BrainRep	0.938 ± 0.028	0.038 ± 0.029	0.218 ± 0.048
AlloNeuroGAT (Freq_filtered_1DCNN)	$\mathbf{0.924 \pm 0.028}$	$\mathbf{0.052 \pm 0.028}$	$\mathbf{0.257 \pm 0.05}$

3.2 Ablation Study

To disentangle the contribution of brain and physiology, we compared unimodal baselines (NeuroGAT-only, using brain features only; AL-only, using biomarkers only). Significant performance differences were found (Repeated Measures ANOVA $p = 0.0008$); the multimodal model significantly surpassed the brain-only variant (Paired T-test: $p\_corrected = 0.0046$) but performed better than the AL-only baseline across repeated settings, although the difference did not reach statistical significance.

Within the brain representation, substituting temporal node features with static ROI coordinates marginally degraded performance (Paired T-Test: $r\, p = 0.087$), and replacing the structural graph with a functional connectivity graph also showed no benefit. Reversing the cross-modal attention roles (Q=allostatic load, K/V=brain) likewise impaired performance. Overall, these findings indicate that both brain and physiological modalities contribute to prediction, though their relative contributions differ. The strong performance of the AL-only baseline suggests that systemic physiological markers capture substantial variance in trait anxiety, while brain features provide additional, albeit modest, predictive value. The sensitivity to temporal feature representation and attention directionality suggests that how neural information is encoded and integrated matters for extracting this complementary signal.

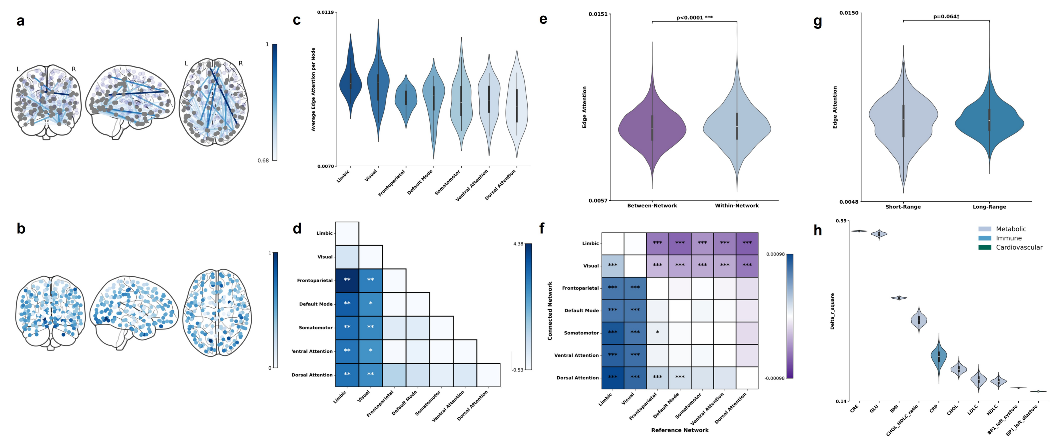

Figure 2: Model Interpretations.

3.3 Model Interpretation

We aggregated attention maps from five runs of the best AlloNeuroGAT to derive edge attention (averaged over layers) and node importance (mean incident–edge attention; Fig. 2a–b). Node importance significantly differed across Yeo networks [16] (Kruskal–Wallis $H = 31.03$, $p < 10^{-4}$); with Limbic and Visual network consistently ranked above others after FDR-BH correction (Fig. 2c-d). This pattern indicates that affective and sensory systems carried greater predictive weight for trait anxiety.

At the edge level, within-network connections received higher attention than between-network ones (Mann–Whitney U, $p < 10^{-4}$; Fig. 2e). Limbic network showed significantly higher within-network attention than edges to any other network and Visual network also favored within-network edges except for the Limbic–Visual pairing (all $p\_corrected < 10^{-3}$). By contrast, edges projecting from Frontoparietal, Default Mode, or Ventral/Dorsal Attention to Limbic or Visual exceeded those networks' own within-network attention (all $p\_corrected < 10^{-3}$; Fig 2f). Together with node importance, this highlights Limbic and Visual as anxiety-relevant hubs that integrate both local and cross-network signals. Model attention also showed a trend toward favoring long-range over short-range connections (Mann–Whitney U, $p = 0.064$; Fig. 2g), suggesting distributed pathways are relevant for prediction.

Finally, permutation importance within the allostatic-load branch (1,000 shuffles per marker) highlighted metabolic and immune axes: shuffling creatinine, glucose, body mass index, and C-reactive protein produced the largest drops in performance, underscoring that systemic load in these domains interacts with brain dynamics to shape anxiety vulnerability (Fig. 2h).

4 Conclusions

We introduced a multimodal graph-attention framework for modeling health time series that integrates rs-fMRI temporal features, structural connectivity, and allostatic-load biomarkers. Experiments showed that preserving temporal dynamics of slow-4/slow-5 oscillations was essential for prediction, while approaches discarding temporal fluctuation underperformed. Multimodal fusion with physiology provided modest but consistent gains. Attention-based analyses suggested that limbic and visual network and metabolic–immune markers contributed most to model decisions. Although predictive accuracy was modest, the reproducibility of these signals indicates that capturing temporal dynamics can yield stable and interpretable patterns from heterogeneous biomedical time series, positioning brain–body integration as a promising direction. This approach aligns with recent work on multimodal fingerprinting (e.g., brain fingerprint [17, 18, 19, 20, 21]) and could extend to other dimensions of psychiatric vulnerability.

References

[1] Liron Rabany, Gretchen J Diefenbach, Laura B Bragdon, Brian P Pittman, Luis Zertuche, David F Tolin, John W Goethe, and Michal Assaf. Resting-state functional connectivity in generalized anxiety disorder and social anxiety disorder: evidence for a dimensional approach. *Brain connectivity*, 7(5):289–298, 2017.

[2] Till Langhammer, Kevin Hilbert, Dirk Adolph, Volker Arolt, Sophie Bischoff, Joscha Böhnlein, Jan C Cwik, Udo Dannlowski, Jürgen Deckert, Katharina Domschke, et al. Resting-state functional connectivity in anxiety disorders: a multicenter fmri study. *Molecular psychiatry*, 30(4):1548–1557, 2025.

[3] Chunlan Yang, Yining Zhang, Min Lu, Jiechuan Ren, and Zhimei Li. White matter structural brain connectivity of young healthy individuals with high trait anxiety. *Frontiers in neurology*, 10:1421, 2020.

[4] Hernando Santamaría-García, Joaquin Migeot, Vicente Medel, Jessica L Hazelton, Vanessa Teckentrup, Roman Romero-Ortuno, Olivier Piguet, Brian Lawor, George Northoff, and Agustin Ibanez. Allostatic interoceptive overload across psychiatric and neurological conditions. *Biological psychiatry*, 97(1):28–40, 2025.

[5] Bruce S McEwen. Allostasis, allostatic load, and the aging nervous system: role of excitatory amino acids and excitotoxicity. *Neurochemical research*, 25(9):1219–1231, 2000.

[6] Youxue Zhang, Chunyan Zhu, Heng Chen, Xujun Duan, Fengmei Lu, Meiling Li, Feng Liu, Xujing Ma, Yifeng Wang, Ling Zeng, et al. Frequency-dependent alterations in the amplitude of low-frequency fluctuations in social anxiety disorder. *Journal of affective disorders*, 174:329–335, 2015.

[7] Wei Wang, Jingming Hou, Shaowen Qian, Kai Liu, Bo Li, Min Li, Zhaohui Peng, Kuolin Xin, and Gang Sun. Aberrant regional neural fluctuations and functional connectivity in generalized anxiety disorder revealed by resting-state functional magnetic resonance imaging. *Neuroscience Letters*, 624:78–84, 2016.

[8] Changjian Qiu, Yuan Feng, Yajing Meng, Wei Liao, Xiaoqi Huang, Su Lui, Chunyan Zhu, Huafu Chen, Qiyong Gong, and Wei Zhang. Analysis of altered baseline brain activity in drug-naive adult patients with social anxiety disorder using resting-state functional mri. *Psychiatry Investigation*, 12(3):372, 2015.

[9] A Babayan, M Erbey, D Kumral, JD Reinelt, AMF Reiter, J Röbbig, HL Schaare, M Uhlig, A Anwander, PL Bazin, et al. A mind-brain-body dataset of mri, eeg, cognition, emotion, and peripheral physiology in young and old adults. sci. data 6, 180308, 2018.

[10] Charles D Spielberger, Fernando Gonzalez-Reigosa, Angel Martinez-Urrutia, Luiz FS Natalicio, and Diana S Natalicio. The state-trait anxiety inventory. *Revista Interamericana de Psicologia/Interamerican journal of psychology*, 5(3 & 4), 1971.

[11] Lothar Laux, Peter Glanzmann, Paul Schaffner, and Charles D Spielberger. Das state-trait-angstinventar (stai): theoretische grundlagen und handanweisung. 1981.

[12] Antonio Jimenez-Marin, Ibai Diez, Asier Erramuzpe, Sebastiano Stramaglia, Paolo Bonifazi, and Jesus M Cortes. Open datasets and code for multi-scale relations on structure, function and neuro-genetics in the human brain. *Scientific Data*, 11(1):256, 2024.

[13] Carlos A Soria, Carolina Remedi, Daniel A Núñez, Luciana D'Alessio, and Emilio JA Roldán. Impact of alprazolam in allostatic load and neurocognition of patients with anxiety disorders and chronic stress (gema): observational study protocol. *BMJ open*, 5(7):e007231, 2015.

[14] Carlos A Soria, Carolina Remedi, Luciana DAlessio, and Emilio JA Roldán. Sex and age-related differences in neuroticism and allostatic load index in urban patients with general anxiety disorder treated with alprazolam. 2018.

[15] Luciana D'Alessio, Guido Pablo Korman, Mercedes Sarudiansky, Laura Ruth Guelman, Laura Scévola, Alejandra Pastore, Amilcar Obregón, and Emilio JA Roldán. Reducing allostatic load in depression and anxiety disorders: physical activity and yoga practice as add-on therapies. *Frontiers in Psychiatry*, 11:520121, 2020.

[16] BT Thomas Yeo, Fenna M Krienen, Jorge Sepulcre, Mert R Sabuncu, Danial Lashkari, Marisa Hollinshead, Joshua L Roffman, Jordan W Smoller, Lilla Zöllei, Jonathan R Polimeni, et al. The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *Journal of neurophysiology*, 2011.

[17] Kausar Abbas, Enrico Amico, Diana Otero Svaldi, Uttara Tipnis, Duy Anh Duong-Tran, Mintao Liu, Meenusree Rajapandian, Jaroslaw Harezlak, Beau M Ances, and Joaquín Goñi. Geff: Graph embedding for functional fingerprinting. *NeuroImage*, 221:117181, 2020.

[18] Kausar Abbas, Mintao Liu, Michael Wang, Duy Duong-Tran, Uttara Tipnis, Enrico Amico, Alan D Kaplan, Mario Dzemidzic, David Kareken, Beau M Ances, et al. Tangent functional connectomes uncover more unique phenotypic traits. *Iscience*, 26(9), 2023.

[19] Benjamin Chiêm, Kausar Abbas, Enrico Amico, Duy Anh Duong-Tran, Frédéric Crevecoeur, and Joaquín Goñi. Improving functional connectome fingerprinting with degree-normalization. *Brain Connectivity*, 12(2):180–192, 2022.

[20] Duy Duong-Tran, Mark Magsino, Joaquín Goñi, and Li Shen. Preserving human large-scale brain connectivity fingerprint identifiability with random projections. In *2024 IEEE International Symposium on Biomedical Imaging (ISBI)*, pages 1–5. IEEE, 2024.

[21] Dachuan Song, Li Shen, Duy Duong-Tran, and Xuan Wang. Causality-based subject and task fingerprints using fmri time-series data. In *Proceedings of the 15th ACM International Conference on Bioinformatics, Computational Biology and Health Informatics*, pages 1–10, 2024.