

Latent Representation Learning for Multimodal Brain Activity Translation

Arman Afrasiyabi^{1,4,9}, Dhananjay Bhaskar^{1,4,9}, Erica L. Busch^{3,8,9}, Laurent Caplette^{3,8,9}, Rahul Singh^{1,8,9},
Guillaume Lajoie^{5,6,7}, Nicholas B. Turk-Browne^{3,8,9,*}, Smita Krishnaswamy^{1,4,2,8,9,*}

Department of {¹Computer Science, ²Applied Mathematics, ³Psychology, ⁴Genetics, ⁵Mathematics and Statistics}

⁶Mila - Quebec AI Institute, ⁷Université de Montréal ⁸Wu Tsai Institute, ⁹Yale University *Jointly Supervised

Abstract—Neuroscience employs diverse neuroimaging techniques, each offering distinct insights into brain activity, from electrophysiological recordings such as EEG, which have high temporal resolution, to hemodynamic modalities such as fMRI, which have increased spatial precision. However, integrating these heterogeneous data sources remains a challenge, which limits a comprehensive understanding of brain function. We present the Spatiotemporal Alignment of Multimodal Brain Activity (SAMBA) framework, which bridges the spatial and temporal resolution gaps across modalities by learning a unified latent space free of modality-specific biases. SAMBA introduces a novel attention-based wavelet decomposition for spectral filtering of electrophysiological recordings, graph attention networks to model functional connectivity between functional brain units, and recurrent layers to capture temporal autocorrelations in brain signal. We show that the training of SAMBA, aside from achieving translation, also learns a rich representation of brain information processing. We showcase this classify external stimuli driving brain activity from the representation learned in hidden layers of SAMBA, paving the way for broad downstream applications in neuroscience research and clinical contexts.

I. INTRODUCTION

Non-invasive techniques such as electroencephalography (EEG) and magnetoencephalography (MEG) provide high temporal resolution, capturing the rapid dynamics of neural activity. In contrast, hemodynamic methods, such as functional magnetic resonance imaging (fMRI), offer rich spatial resolution [23]. As neuroscience advances towards more sophisticated models of cognition, integrating these diverse data types becomes increasingly critical [3]. Combining the complementary strengths of these modalities could offer a more comprehensive understanding of brain function, but this remains a challenging task.

While substantial progress has been made in utilizing multimodal data consisting of image stimuli and brain activity pairs – particularly with Generative Adversarial Networks (GANs), transformers, and diffusion models to reconstruct images from brain activity [6, 10, 12, 13, 16–19] – the same is not true for the integration of multiple brain imaging modalities. Most of the work in this area has focused on leveraging information from EEG to enhance the fidelity of fMRI signals [1, 5, 7, 15]. These efforts, while valuable in improving fMRI’s localization and signal-to-noise ratio with temporally rich EEG signals, often fall short of addressing the more complex task of multimodal fusion and do not address the complexities of spatiotemporal upsampling and downsampling between modalities.

To bridge this gap, we propose a novel multi-modal neural network framework, Spatiotemporal Alignment of Multimodal Brain Activity (SAMBA), designed to generalize the translation

between electrophysiological and hemodynamic signals. SAMBA addresses both spatial and temporal disparities through graph attention and wavelet-based modules. Our objectives are threefold: (1) to create a unified latent space that captures spatiotemporal dynamics without modality-specific biases, enabling its application across a broad set of downstream tasks, such as brain state classification, cognitive assessment, and diagnosis of neurological disorders; (2) to develop data-driven models of hemodynamic response and functional connectivity in the brain; and (3) to combine smaller unimodal datasets into larger multimodal cohorts, laying the groundwork for training foundational models. SAMBA incorporates (1) temporal upsampling and downsampling modules based on learnable hemodynamic response functions (HRFs) and attention-based wavelet decomposition for spectral filtering; (2) spatial upsampling and downsampling modules powered by graph attention networks (GATs) to model functional connectivity across brain regions; and (3) recurrent layers to capture autocorrelations in the temporal domain.

We demonstrate the efficacy of SAMBA in several key tasks. First, the framework enables precise translation between electrophysiological and hemodynamic modalities, allowing for accurate cross-modal mapping. We also perform ablation studies to confirm the essential roles of all SAMBA components. Next, we show that SAMBA’s unified latent representations can accurately classify scenes in a movie shown to the subjects during data acquisition, demonstrating that the translation task allows SAMBA to capture rich representations of cognitive activity. Finally, we also show that the wavelet decomposition module in SAMBA filters specific EEG/MEG frequencies during translation for denoising, while the learnable HRF module models heterogeneity in neurovascular coupling across brain regions.

II. METHODS

Electrophysiological recordings, denoted as $X(t) = \{x_1(t), \dots, x_N(t)\}$, represent the neural activity across N parcels of the brain. Hemodynamic responses, represented as $Y(\tau) = \{y_1(\tau), \dots, y_M(\tau)\}$, capture the blood oxygenation and flow changes across M parcels, where $M \gg N$ due to the finer spatial resolution offered by fMRI. However, the temporal resolution of X is higher than that of Y .

A. Electrophysiological Activity to Hemodynamic Response

We elaborate on the translation from $X(t)$ to $Y(\tau)$.

1) *Temporal Smoothing with HRF learning*: The HRF is designed to model the latency and variability of blood flow in response to neural activity. Due to significant variations in

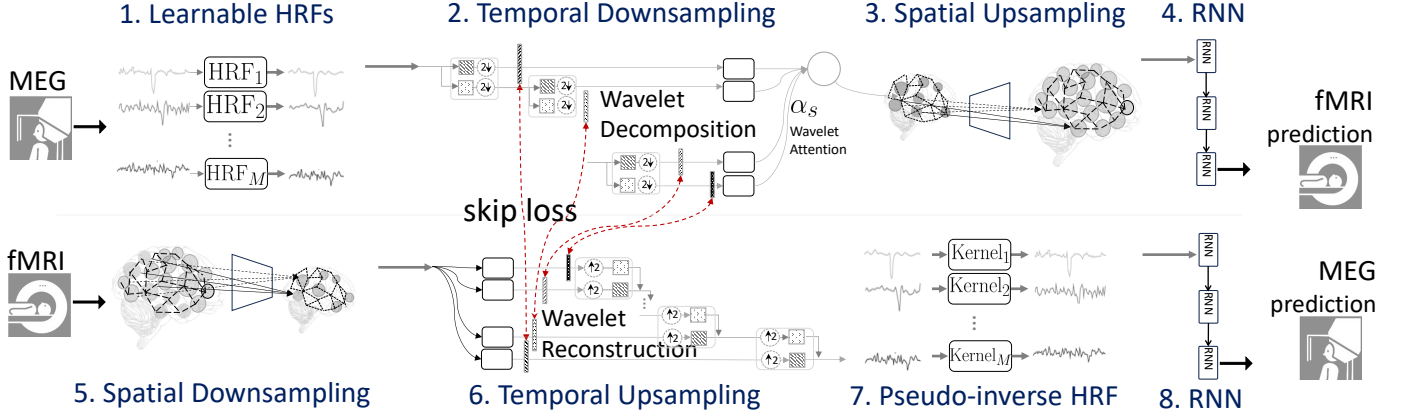


Fig. 1: SAMBA translates between MEG and fMRI modalities by upsampling and downsampling using wavelet decomposition and graph-attention modules in the temporal and spatial domains respectively. The upper and bottom parts show the fMRI-to-MEG and MEG-to-fMRI prediction modules respectively.

neuronal density and metabolic demand across regions of the brain, the HRF responses also vary across the brain [2, 11]. To address this, we employ a parcel-specific $\text{HRF}_n(t)$, parameterized by six learnable parameters $\text{HRF}_n(t) =$

$$\theta_1 \left(\frac{t}{p_r} \right)^{\theta_2} \exp \left(-\frac{t-p_r}{\theta_3} \right) - \theta_4 \left(\frac{t}{p_u} \right)^{\theta_5} \exp \left(-\frac{t-p_u}{\theta_6} \right) \quad (1)$$

where θ_1 and θ_4 are the amplitude of the response and undershoot components, respectively, modulating the increase and decrease in blood flow and oxygenation to the brain area activated following neural activity. θ_2 and θ_5 represent the time-to-peak of the response and undershoot components, respectively. θ_3 and θ_6 are the dispersion factors, influencing the width of the response and undershoot curves. $p_r = (\theta_2 \cdot \theta_3)$, and $p_u = (\theta_5 \cdot \theta_6)$ denote the peak times of the respective components. The learnable parameters of the HRF model are inferred via a three-layer MLP for each brain parcel. For each parcel n , the HRF is convolved with the electrophysiological signal $x_n(t)$ to produce: $\tilde{x}_n(t) = \text{HRF}_n(t) * x_n(t)$, where $*$ denotes the convolution operation. This convolution process smooths the electrophysiological signal into a representation of the blood flow dynamics resulting from neural activity.

2) *Temporal Downsampling*: To perform temporal downsampling, we propose a unique architecture that compresses temporal signals via a rich wavelet transform and then uses attention to select the appropriate signal bands for the translation tasks. The process involves constructing daughter wavelets by scaling and translating the mother wavelet, ψ , by s and u respectively: $\psi_{s,u}(t) = \psi((t-u)/s)$. Wavelet coefficients are computed by convolving $\tilde{x}_n(t)$ with daughter wavelets $c_n(s,u) = \tilde{x}_n(t) * \psi_{s,u}(t)$. At smaller scales, where higher frequencies are analyzed, more translations u are required to perform the convolution, resulting in a larger number of coefficients. Conversely, fewer translations are necessary at larger scales, yielding fewer coefficients. Next, we concatenate the scale-specific embeddings to form a multiscale representation, expressed as $z_n = \parallel_s \alpha_s c_n(s)$. Here, \parallel denotes concatenation of features, $z_n \in \mathbb{R}^d$ and α_s represents the learnable attention weight allocated to the embedding at scale s normalized by the Softmax function, indicating the significance of features captured at that scale relative to others in the final multiscale representation. The attention weights are normalized

using the Softmax function, transforming them into a probabilistic distribution that identifies the most salient frequency bands in the electrophysiological data.

3) *Spatial Upsampling Module*: In this module, we outline our approach for translating data from a coarse-grained graph of brain regions, denoted as $G_X = (V^X, E^X, W^X)$, derived from electrophysiological measurements in the source modality, to a fine-grained graph, $G_Y = (V^Y, E^Y, W^Y)$, which features a higher spatial resolution using hemodynamic data from the target modality (Fig. 1c). Recall that our task is to translate N time-lapse electrophysiological signals represented as $X(t) = \{x_1(t), \dots, x_N(t)\}$, to M time-lapse hemodynamic signals $Y(\tau) = \{y_1(\tau), \dots, y_M(\tau)\}$. To achieve this, our source graph contains N nodes ($|V_X| = N$) and our target graph contains M nodes ($|V_Y| = M$), where $M \gg N$. Here, the edge weights, W^X , in the source graph, are assigned based on the cosine similarity between timelapse electrophysiological signals: $W_{pq}^X = (x_p(t) - \bar{x}_p) \cdot (x_q(t) - \bar{x}_q) / \|x_p(t) - \bar{x}_p\| \|x_q(t) - \bar{x}_q\|$, where \bar{x}_p is the mean of the signal $x_p(t)$. We input the latent representations $\{z_j\}_{j=1}^N$ as node features into a GAT layer, which computes hidden features of nodes

$$h_n^X(\tau) = \sigma \left(\frac{1}{K} \sum_{k=1}^K \sum_{j \in \mathcal{N}(n)} \beta_{nj}^{(k)} W^{(k)} z_j(\tau) \right), \quad (2)$$

where K is the number of attention heads, $\beta^{(k)}$ are the attention coefficients, and $W^{(k)}$ are the head-specific weight matrices. We then follow the standard GAT implementation [4, 22]. Edge weights in the target graph G_Y are based on the cosine similarity of hemodynamic signals: $W_{pq}^Y = (y_p(\tau) - \bar{y}_p) \cdot (y_q(\tau) - \bar{y}_q) / \|y_p(\tau) - \bar{y}_p\| \|y_q(\tau) - \bar{y}_q\|$ where, \bar{y}_p denotes the mean of the hemodynamic signal in parcel p . The node features in G_Y are defined using single-layer feed-forward networks, $\{\phi_m\}_{m=1}^M$, which map the hidden representations $\{h_n^X\}_{n=1}^N$ in G_X to the nodes in G_Y . Each network ϕ_m takes the aggregated representations $\{h_i^X\}_{i \in \mathcal{X}_m}$ as input, where \mathcal{X}_m is the subset of nodes from the same neuroanatomical region in the source graph. For example, to obtain the node features of a visual cortex parcel in the target graph, G_Y , we use hidden representations of all available visual cortex parcels in the source graph, G_X . We then used a GAT layer to aggregate the features in the target graph:

$$h_m^Y(\tau) = \sigma \left(\frac{1}{K} \sum_{k=1}^K \sum_{j \in \mathcal{N}(m)} \gamma_{mj}^{(k)} W^{(k)} \phi_m(\{h_i^X\}_{i \in \mathcal{X}_m}) \right), \quad (3)$$

TABLE I: Evaluation of translation using Spearman correlation for minute and second predictions between fMRI, MEG, and EEG modalities.

		MEG → fMRI		EEG → fMRI	
		Minute	Second	Minute	Second
a) Electrophysiological to Hemodynamic	MLP	0.05	0.12	0.04	0.10
	1D-CNN	0.07	0.14	0.09	0.14
	LSTM	0.26	0.39	0.18	0.31
	Transformer	0.34	0.60	0.19	0.28
	No Wavelet	0.14	0.27	0.12	0.24
	No LSTM	0.18	0.30	0.11	0.23
	No LSTM: Avg. 2 samples	0.36	0.65	0.26	0.36
	HRF-Wavelet-MLP-LSTM	0.37	0.66	0.28	0.39
	Transformer instead of LSTM	0.33	0.63	0.23	0.37
	Fixed HRF	0.36	0.60	0.28	0.41
	MSE-Loss instead of Cosine	0.36	0.58	0.25	0.41
	SAMBA	0.38	0.63	0.29	0.43
	Transformer	0.33	0.62	0.14	0.30
	SAMBA	0.39	0.67	0.28	0.44
b) Hemodynamic to Electro.	MLP	0.05	0.11	0.05	0.10
	1D-CNN	0.06	0.16	0.07	0.15
	LSTM	0.13	0.25	0.11	0.22
	Transformer	0.15	0.30	0.11	0.28
	No Pseud HRF	0.21	0.34	0.09	0.20
	No Skip Loss	0.15	0.24	0.10	0.19
	SAMBA	0.21	0.35	0.15	0.33
	Transformer	0.11	0.26	0.10	0.26
	SAMBA	0.19	0.31	0.13	0.27

We built a baseline using five methods between withheld time intervals for all subjects and subject pairs (blue-coded).

where, $\gamma^{(k)}$ are normalized attention coefficients, $\mathcal{N}(m)$ is neighboring nodes of m , and $W^{(k)}$ are unique weight matrices for each attention head. Ultimately, this module generates a series of high-resolution node representations, $\{h_m^Y\}_{m=1}^M$, which produce the desired output, $Y(\tau)$.

4) *Hemodynamic Sequence Generation via RNNs*: Upon spatially upscaling, the refined high-resolution node representations, denoted as h_m^Y , are fed into a recurrent model in the final stage. To this end, we employ a LSTM network, since it is well-suited for modeling the autoregressive characteristics inherent in these temporal sequences. The LSTM processes the sequence of node representations, $h_m^Y(\tau)$, to predict hemodynamic activity, $\hat{Y}(\tau) = \{\hat{y}_1(\tau), \dots, \hat{y}_M(\tau)\}$, as follows:

$$\hat{y}_m(\tau_o + 1) = \text{LSTM}(\hat{y}_m(\tau_o), h_m^Y(\tau_o + 1)), \quad (4)$$

where $m = 1, \dots, M$ and $\hat{y}_m(\tau_o + 1)$ is the estimated hemodynamic activity in the m -th parcel at time $\tau = \tau_o + 1$. This estimation relies on the previously predicted τ_o , denoted as $\hat{y}_m(\tau_o)$, and the current node representation, $h_m^Y(\tau_o + 1)$.

B. Hemodynamic Response to Electrophysiological Activity

Here, we describe our methodology for converting hemodynamic activity, $Y(\tau)$, to electrophysiological activity, $X(\tau)$.

1) *Spatial Downsampling Module*: To perform spatial downsampling, we invert and adapt the methodology detailed in the graph upsampling section, converting a fine-grained hemodynamic graph, G_Y , containing M nodes, into a coarse-grained electrophysiological graph, G_X , containing N nodes, where $M \gg N$. Here, a GAT layer aggregates node features from the brain activity graph G_Y , which are then mapped to a coarser target graph G_X using linear layers.

2) *Temporal Upsampling Module*: Given $h_n^X(\tau)$, as the spatially downsampled hemodynamic data, we now aim to perform temporal upsampling. We first model the reverse process of wavelet decomposition by estimating the wavelet coefficients at various wavelet coefficient scales and performing the inverse wavelet decomposition. We achieve this in two steps. First, we estimate the wavelet coefficients using a set of linear layers $\{f_s\}_{s=1}^S$. Each layer f_s maps the input signal to the wavelet coefficient space at a specific scale: $\hat{c}(s, u) = f_s(h_n^X(\tau))$, where $\hat{c}(s, u)$ represents the estimated wavelet coefficient at scale s and

position u . To reconstruct n -th HRF smoothed signal, we then perform wavelet reconstruction using the estimated coefficients:

$$\tilde{x}_n(t) = \sum_{s \in \mathcal{S}} \sum_{u \in \mathcal{U}} \hat{c}(s, u) \psi_{s,u}(t), \quad (5)$$

where $\psi_{s,u}(t)$ denotes the daughter wavelets obtained by scaling and translating the mother wavelet ψ by factors of s and u , respectively. However, to ensure accurate wavelet coefficient estimation, we employ a regularization strategy using wavelet coefficient skip losses (between blocks 1 and 6 in Fig. 1). This function penalizes the network for discrepancies between the true wavelet coefficients $c(s, u)$ from the electrophysiological, and the estimated coefficients $\hat{c}(s, u)$:

$$L_{\text{reg}} = \frac{1}{|\mathcal{S}|} \frac{1}{|\mathcal{U}|} \sum_{s \in \mathcal{S}} \sum_{u \in \mathcal{U}} (c(s, u) - \hat{c}(s, u))^2. \quad (6)$$

3) *Deconvolution using Pseudo-inverse HRF*: We now aim to build a pseudo-inverse HRF function to estimate the original neural signals from smoothed HRF. Since the double gamma form of the HRF function is not invertible, we estimate the original temporal dimension of MEG or EEG (at 200 Hz) using per-parcel single kernel learning via 1D transpose convolution. The reconstruction is mathematically represented as: $\hat{x}_n(t) = \text{DeConv1D}_n(\tilde{x}_n(t))$, where DeConv1D_n is the parcel-specific transpose convolution with the single learnable kernel.

4) *Electrophysiological Sequence Generation with RNNs*: Upon temporal reconstruction, the refined low-resolution node representations, denoted as h_n^X , are fed into a recurrent model in the final stage of translation from hemodynamic activity to electrophysiological signals in the brain. To this end, we employ an LSTM to process the sequence of node representations, $h_n^X(t)$, to predict electrophysiological activity, $\hat{X}(t) = \{\hat{x}_1(t), \dots, \hat{x}_N(t)\}$, akin to Eq. 4.

C. Loss Formulation

We employed the cosine similarity loss function to train the model to align the predicted signal with the target signal. In hemodynamic mapping to electrophysiological, for example, given the predicted m -th parcel \hat{y}_m , the loss is defined as:

$$L_{\text{match}} = \sum_{m=1}^M \left(1 - \frac{\hat{y}_m \cdot y_m}{\|\hat{y}_m\|_2 \|y_m\|_2} \right), \quad (7)$$

where, M is the number of parcels, $\|\hat{y}_m\|_2$, and $\|y_m\|_2$ are the L2 norms of \hat{y}_m and y_m , respectively. Here, in addition to the cosine




Fig. 2: PyCortex [8] visualizations of fMRI activity on the unfolded brain surface compare ground truth (top row) with translations via SAMBA (middle) and the SOTA transformer model (bottom), as quantitatively shown in Table I. Timestamps (mm:ss) correspond to the Forrest Gump movie.

loss we also regularized the network using skip loss, as in Eq. 6: $\lambda L_{\text{match}} + (1 - \lambda) L_{\text{reg}}$. To map MEG/EEG to fMRI we only train the model with the cosine loss Eq. 7, given \hat{y}_n and y_n .

III. RESULTS

We conducted experiments using two datasets: (1) StudyForest [9, 14], comprising MEG and fMRI data, and (2) Naturalistic Viewing [20], containing EEG and fMRI recordings. Both datasets include fMRI data aligned to the fsaverage surface (10,242 vertices per hemisphere) and collected at 0.5 Hz (2s TR). MEG and EEG data were source-reconstructed onto the same surface, initially sampled at 600 Hz and resampled to 200 Hz. We used Schaefer¹ parcellations, with 200 regions for MEG and 500 for fMRI, to model spatial connectivity as a graph. Data from each subject were segmented into randomly selected one-minute samples, yielding 106 MEG-fMRI and 100 EEG-fMRI samples per subject. For classification tasks, we utilized movie labels from the EEG-fMRI dataset. To this end, we evaluate SAMBA on four translation tasks: (1) fMRI-to-MEG, (2) fMRI-to-EEG, (3) MEG-to-fMRI, and (4) EEG-to-fMRI. We then explore our SAMBA model’s evaluation of the classification task to detect eight distinct movies in the Naturalistic Viewing dataset.

In Table I, we compare SAMBA’s performance against several baseline architectures, including convolutional, transformer, recurrent, and feed-forward networks. We also include ablation studies of the SAMBA architecture, where key components such as wavelet decomposition, the learnable HRF, and the recurrent layer are systematically removed or replaced. Specifically, in Table Ia we assess performance in translating electrophysiological data to hemodynamic data, and in Table Ib, we report results for the reverse task. The primary evaluation metric is Spearman correlation, averaged across all Schaefer parcels, between the predicted and ground truth time-lapse signals in both long (1 min) and short (15 sec) intervals of withheld timepoints. The evaluate SAMBA when trained across all fMRI-EEG/MEG subject pairs (black text), as well as a subject-specific SAMBA model, where a separate model is trained for each subject pair (blue text), and the reported Spearman correlations are averaged across all withheld timepoints for each subject. SAMBA outperforms all baseline models across all tasks, with the transformer model by Vaswani et al. [21] achieving the second-best performance.

Fig. 2 illustrates SAMBA’s performance in translating MEG to fMRI data, using pycortex[8] from the StudyForest dataset. The results, highlighted specifically by red circles, indicate that SAMBA effectively recovers fMRI signals from MEG data compared to SOTA (transformer), even in the test set.

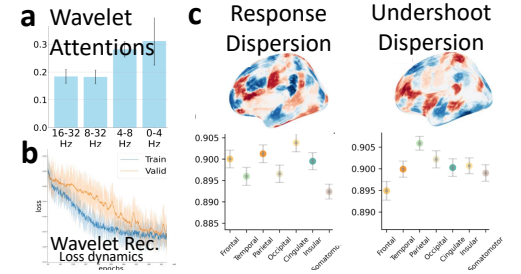


Fig. 3: Wavelet attention in a, and loss dynamics in b. c) Inferred HRF undershoot and response dispersion.

Fig. 3a illustrates the dynamics of wavelet decomposition attention and wavelet reconstruction skip loss in our model. Based on the attention intensity values, our models primarily focus on lower frequencies (4-8 Hz and 0-4 Hz), likely due to the higher signal-to-noise ratio at these frequencies compared to higher frequencies. Fig. 3b presents variations in the details of the skip-loss dynamics during wavelet reconstruction. To highlight SAMBA’s learned representations, we added a classification head to identify eight movies from the Naturalistic Viewing dataset [20], comparing performance with baseline methods in Table II. Using K-fold cross-validation, we withheld random samples per fold, applying a 65-15-20 train-validation-test split. Notably, our model achieves a 10.54% improvement in the EEG to fMRI classification tasks over the baseline.

TABLE II: Movie classification accuracy results.

	EEG-to-fMRI	fMRI-to-EEG
1D-CNN	48.83%	30.69%
LSTM	53.71%	37.09%
Transformer	51.04%	38.24%
SAMBA	61.58%	46.50%

Our model also offers neuroscientific interpretations. Here, we outline key findings from the best-performing MEG-to-fMRI model. Fig. 3c displays the inferred HRF parameters for each brain parcel. This figure shows the variation in HRF response and undershoot dispersion across different brain regions, highlighting the diversity in oxygenation and deoxygenation levels [11]. The left somatomotor network exhibits minimal response dispersion compared to the cingulate, whereas the parietal lobe regions show greater undershoot dispersion than the right somatomotor.

IV. CONCLUSIONS

This paper introduces SAMBA, a framework designed to address spatiotemporal trade-offs in multimodal brain activity translation. Using wavelet-attention-based temporal encoding and decoding with context-aware graph upsampling and downsampling, SAMBA outperforms baseline methods like transformers. The framework’s translation task yields rich representations useful for downstream tasks like classification.

Acknowledgements This work was supported in part by the National Science Foundation (NSF Career Grant 2047856, NSF DMS 2327211, and NSF IIS 2473317) and the National Institute of Health (NIH 1R01GM130847-01A1, NIH 1R01GM135929-01). D.B. was supported by the Kavli Institute for Neuroscience Postdoctoral Fellowship. E.L.B. was supported by an NSF Graduate Research Fellowship (Award no. 2139841). L.C. was supported by the Fonds de Recherche du Québec – Nature et Technologies Postdoctoral Scholarship. R.S. was supported by the Wu Tsai Institute Fellowship. N.B.T.B. was supported by CIFAR.

REFERENCES

- [1] Rodolfo Abreu, Alberto Leal, and Patrícia Figueiredo. Eeg-informed fmri: a review of data analysis methods. *Frontiers in human neuroscience*, 12:29, 2018.
- [2] David Attwell, Alastair M Buchan, Serge Charpak, Martin Lauritzen, Brian A MacVicar, and Eric A Newman. Glial and neuronal control of brain blood flow. *Nature*, 468(7321):232–243, 2010.
- [3] Melanie Boly, Olivia Gosseries, Marcello Massimini, and Mario Rosanova. Functional neuroimaging techniques. In *The Neurology of Consciousness*, pages 31–47. Elsevier, 2016.
- [4] Shaked Brody, Uri Alon, and Eran Yahav. How attentive are graph attention networks? *International Conference on Learning Representations*, 2022.
- [5] David Calhas and Rui Henriques. Eeg to fmri synthesis: Is deep learning a candidate? *arXiv preprint arXiv:2009.14133*, 2020.
- [6] Thirza Dado, Yağmur Güçlütürk, Luca Ambrogioni, Gabriëlle Ras, Sander Bosch, Marcel van Gerven, and Umut Güçlü. Hyperrealistic neural decoding for reconstructing faces from fmri activations via the gan latent space. *Scientific reports*, 12(1):141, 2022.
- [7] Calhas David. Eeg-to-fmri neuroimaging cross modal synthesis in python. *Proceedings of the 22nd Python in Science Conference*, 36:53 – 58, 2023.
- [8] James S Gao, Alexander G Huth, Mark D Lescroart, and Jack L Gallant. Pycortex: an interactive surface visualizer for fmri. *Frontiers in neuroinformatics*, 9:23, 2015.
- [9] Michael Hanke, Nico Adelhöfer, Daniel Kottke, Vittorio Iacovella, Ayan Sengupta, Falko R Kaule, Roland Nigbur, Alexander Q Waite, Florian Baumgartner, and Jörg Stadler. A studyforrest extension, simultaneous fmri and eye gaze recordings during prolonged natural stimulation. *Scientific data*, 3(1):1–15, 2016.
- [10] James V Haxby, M Ida Gobbini, Maura L Furey, Alumit Ishai, Jennifer L Schouten, and Pietro Pietrini. Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science*, 293(5539):2425–2430, 2001.
- [11] Suzanaerculano-Houzel. The human brain in numbers: a linearly scaled-up primate brain. *Frontiers in human neuroscience*, page 31, 2009.
- [12] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- [13] Sikun Lin, Thomas Sprague, and Ambuj K Singh. Mind reader: Reconstructing complex images from brain activities. *Advances in Neural Information Processing Systems*, 35:29624–29636, 2022.
- [14] Xingyu Liu, Yuxuan Dai, Hailun Xie, and Zonglei Zhen. A studyforrest extension, meg recordings while watching the audio-visual movie “forrest gump”. *Scientific data*, 9(1):206, 2022.
- [15] Xueqing Liu and Paul Sajda. A convolutional neural network for transcoding simultaneously acquired EEG-fMRI data. In *2019 9th International IEEE/EMBS Conference on Neural Engineering (NER)*, pages 477–482. IEEE, 2019.
- [16] Furkan Ozelik, Bhavin Choksi, Milad Mozafari, Leila Reddy, and Rufin VanRullen. Reconstruction of perceived images from fmri patterns and semantic brain exploration using instance-conditioned gans. In *2022 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE, 2022.
- [17] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022.
- [18] Katja Seeliger, Umut Güçlü, Luca Ambrogioni, Yagmur Güçlütürk, and Marcel AJ van Gerven. Generative adversarial networks for reconstructing natural images from brain activity. *NeuroImage*, 181:775–785, 2018.
- [19] Guohua Shen, Tomoyasu Horikawa, Kei Majima, and Yukiyasu Kamitani. Deep image reconstruction from human brain activity. *PLoS computational biology*, 15(1):e1006633, 2019.
- [20] Qawi K Telesford, Eduardo Gonzalez-Moreira, Ting Xu, Yiwen Tian, Stanley J Colcombe, Jessica Cloud, Brian E Russ, Arnaud Falchier, Maximilian Nentwich, Jens Madsen, et al. An open-access dataset of naturalistic viewing using simultaneous eeg-fmri. *Scientific Data*, 10(1):554, 2023.
- [21] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [22] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. *International Conference on Learning Representations*, 2018.
- [23] Hongtu Zhu, Tengfei Li, and Bingxin Zhao. Statistical learning methods for neuroimaging data analysis with applications. *Annual Review of Biomedical Data Science*, 6:73–104, 2023.