

EEG-to-fMRI Translation: Methodological Foundations for Scalable Multi-Modal Real-Time Neurofeedback

Saampras Ganesan^{1,4}, Zhiyin Chen^{1,5}, Masaya Misaki^{1,2,3}

¹ Laureate Institute for Brain Research, Tulsa, OK, USA

² Oxley College of Health and Natural Sciences, The University of Tulsa, Tulsa, OK, USA

³ University of Oklahoma College of Medicine School of Community Medicine, Tulsa, OK, USA

⁴ Centre for youth mental health, Faculty of Medicine, Dentistry and Health Sciences, University of Melbourne, Australia

⁵ Department of Electrical and Computer Engineering, The University of Tulsa, Tulsa, OK, US

Introduction

EEG-to-fMRI translation (predicting hemodynamic activity from scalp electrophysiology using simultaneously acquired EEG and fMRI) offers a route toward portable, high-resolution brain decoding. A compelling motivation is neurofeedback: an EEG-derived approximation of fMRI dynamics could enable circuit-specific neuromodulation from EEG alone, without the cost and immobility of an MRI scanner. Because both signals come from a single EEG recording, it could also support ecologically-valid multimodal neurofeedback from one modality, while improving spatiotemporal specificity over conventional EEG-neurofeedback. Prior EEG-to-fMRI translation work has shown feasibility but remains constrained by small datasets (typically <60 subjects), weak cross-subject generalization, or reconstructions limited to a few brain regions (Meir-Hasson et al., 2014; Li et al., 2024; Roos et al., 2025; Iwama et al., 2024). Overcoming these limitations requires solving two intertwined problems together: a modelling problem (mapping electrophysiology to whole-brain hemodynamics in a way that generalizes across people) and a data-engineering problem (denoising simultaneous EEG-fMRI at scale). The latter is often treated as routine, yet residual MRI-induced artifacts in EEG can directly limit the learnable mapping. The present work targets both, emphasizing the methodological framework rather than final performance.

Methods

We introduce EEG-to-fMRI-Translate, a framework trained on one of the largest simultaneous EEG-fMRI datasets assembled to date (~400 subjects; 32-channel EEG; 3T fMRI; two event-design cognitive tasks). Rather than learning a mapping directly between high-dimensional raw EEG and whole-brain fMRI, the approach works in the compact latent spaces of two recently pretrained neuroimaging foundation models: REVE (EEG) and NeuroSTORM (fMRI). REVE (El Ouahidi et al., 2025) was pretrained on ~25,000 subjects to encode whole-brain EEG into a compact representation, and NeuroSTORM (Wang et al., 2026) was pretrained on ~50,000 subjects to do the same for whole-brain fMRI. Because these representations already capture the underlying structure of each modality, a bridge only has to map between these compact, denoised spaces rather than between large, noisy raw signals. This is likely far more tractable than training a fresh model end-to-end on the raw signals, which would demand enormous paired EEG-fMRI datasets that are not currently feasible to acquire. A trainable Transformer-based bridge connects the EEG representation to the fMRI target. During training, one or both foundation models can be fine-tuned to adapt their pretrained representations to the cross-modal task. The fMRI latent space guides this training in one of two ways: as an auxiliary signal shaping the BOLD output, or as the prediction target itself, from which whole-brain BOLD recovers readily via the NeuroSTORM's defined latent-to-signal mapping.

A second methodological pillar concerns preprocessing. EEG recorded inside the MRI scanner is contaminated by two large, structured artifacts: the gradient (scanner) artifact from switching imaging gradients, and the ballistocardiogram (pulse) artifact locked to the cardiac cycle. Both dwarf the neural signal of interest. Correction quality directly limits what any translation model can achieve: residual artifacts either mask the neural variance the model needs or are learned as spurious signal, degrading prediction and cross-subject generalization. Established methods are largely interactive and tuned manually per-recording, creating a bottleneck at the scale of

hundreds of subjects, and even existing automated approaches are often inaccurate, leaving substantial residual noise. We are therefore developing automated, high-quality artifact-correction procedures alongside the modelling effort; EEG clean of MRI artifacts is a prerequisite, not an afterthought, for any EEG-to-fMRI translation system.

Results (preliminary)

Results are preliminary and reported here only to indicate direction rather than to make performance claims. Cross-subject decoding of the full BOLD timecourse is currently strongest in somatomotor and visual networks, with mean correlations of $r \approx 0.26$ – 0.30 and the best subjects reaching $r = 0.58$ – 0.67 . Higher-order association networks are weaker, with default-mode subnetworks averaging $r \approx 0.15$ – 0.18 , and cortical limbic networks weakest (mean $r \approx 0.08$). Subcortical structures are recovered above chance but remain challenging: the putamen, hippocampus, and caudate are strongest (mean $r \approx 0.15$ – 0.17 , best subjects up to $r \approx 0.44$ – 0.51), while the pallidum is weakest (mean $r \approx 0.06$ – 0.08). Within every network the spread across subjects is wide, often ranging from near zero to $r > 0.5$, indicating considerable inter-subject variability.

These results come from a model trained only on event-related tasks, whose rapid, transient responses may be intrinsically harder to learn from and decode. We expect performance to improve once training incorporates resting-state and block-design simultaneous EEG-fMRI datasets, which tend to be more favorable. This expansion depends on the second pillar: the pulse- and gradient-artifact correction needed to bring new datasets to usable quality. We are currently developing automated methods for this correction; early results on held-out subjects are promising, though a small but meaningful fraction of recordings remain difficult to fully denoise.

Conclusions

Scalable EEG-to-fMRI translation is as much a data-engineering challenge as a modelling one. By building on large pretrained foundation models and, in parallel, developing automated correction for scanner and pulse artifacts, EEG-to-fMRI-Translate aims to move the field beyond small-sample, region-limited demonstrations toward a generalizable system. Preliminary results are encouraging in somatomotor and visual networks, but accuracy falls off in higher-order association and subcortical regions, and varies widely across subjects, so the present model is far from deployment-ready. Closing these gaps defines our near-term roadmap: expanding the training data corpus across more subjects, tasks and datasets, refining model architectures, and maturing the automated artifact-correction needed to bring those datasets to usable quality. Removing MRI artifacts as fully as possible, so the in-scanner EEG resembles an externally collected recording, is ultimately what will enable such systems to run reliably on real-world EEG outside the scanner. The longer-term significance is a practical path to reliable, anatomically targeted multimodal real-time neurofeedback driven by EEG alone.

References

1. Meir-Hasson, Y., Kinreich, S., Podlipsky, I., Hendler, T., & Intrator, N. (2014). An EEG Finger-Print of fMRI deep regional activation. *NeuroImage*, 102(Part 1), 128-141. <https://doi.org/10.1016/j.neuroimage.2013.11.004>
2. Li, Y., Lou, A., Xu, Z., Zhang, S., Wang, S., Englot, D. J., ... & Chang, C. (2024). NeuroBOLT: Resting-state EEG-to-fMRI synthesis with multi-dimensional feature mapping. *Advances in Neural Information Processing Systems*, 37, 29c294ddc628c94cd2c636383ef106c1.
3. Roos, K. G., Fukuda, A., & Cap, Q. H. (2025). *From brainwaves to brain scans: A robust neural network for EEG-to-fMRI synthesis*. arXiv.
4. Iwama, S., Tsuchimoto, S., Mizuguchi, N., & Ushiba, J. (2024). EEG decoding with spatiotemporal convolutional neural network for visualization and closed-loop control of sensorimotor activities: A simultaneous EEG-fMRI study. *Human Brain Mapping*, 45(9), e26767. <https://doi.org/10.1002/hbm.26767>
5. El Ouahidi, Y., et al. (2025). REVE: A Foundation Model for EEG - Adapting to Any Setup with Large-Scale Pretraining on 25,000 Subjects. *Advances in Neural Information Processing Systems* (NeurIPS).
6. Wang, C., Jiang, Y., Peng, Z., Li, C., Bang, C. B., Zhao, L., Fu, W., Lv, J., Sepulcre, J., Yang, C., et al. (2026). Towards a general-purpose foundation model for functional MRI analysis. *Nature Biomedical Engineering*, 1–12