

TOWARDS NEURAL FOUNDATION MODELS FOR VISION: ALIGNING EEG, MEG AND fMRI REPRESENTATIONS TO PERFORM DECODING, ENCODING AND MODALITY CONVERSION

Matteo Ferrante

University of Rome, Tor Vergata
Department of Biomedicine and Prevention
matteo.ferrante@uniroma2.it

Tommaso Boccato

University of Rome, Tor Vergata
Department of Biomedicine and Prevention

Nicola Toschi

University of Rome, Tor Vergata
Department of Biomedicine and Prevention
A.A. Martinos Center for Biomedical Imaging
Harvard Medical School/MGH, Boston (US)

ABSTRACT

This paper presents a novel approach towards to the creation of a foundational model for aligning neural data and visual stimuli representations by leveraging the power of contrastive learning. We worked with EEG, MEG and fMRI. The capabilities of our framework are showcased through three key experiments: decoding visual information from neural data, encoding images into neural representations, and converting between neural modalities. The results demonstrate the model’s ability to accurately capture semantic information across different brain imaging techniques, illustrating its potential in decoding, encoding, and modality conversion tasks.

1 INTRODUCTION

The non-invasive measurement of neural activity is crucial to understand the human brain. The advent of Artificial Intelligence has propelled the field of neuroscience into the use of novel paradigms, including a wide array of encoding/decoding models. These models have shown remarkable proficiency in interpreting various sensory inputs, encompassing vision, auditory processing, and motor imagery, among others. Key to this endeavor are non-invasive neural correlates like Electroencephalography (EEG), Magnetoencephalography (MEG), and functional Magnetic Resonance Imaging (fMRI). Each of these modalities offers a unique window into brain activity, capturing complementary aspects of its response to external stimuli and providing insights into the perceptual and representational processes within. In this context, our study introduces a step forward in the realm of neural foundation models for vision. We aim to harmonize disparate modalities drawn from diverse EEG, MEG, and fMRI datasets acquired during a vision task, creating a unified representation that transcends the limitations of individual modalities. Our approach leverages the power of contrastive learning to align representations across these varied neural correlates, and is anchored in the image representations provided by CLIP (Contrastive Language–Image Pretraining) model (Radford et al., 2021) as depicted in Fig 1. We assess the capabilities of our framework in performing an array of tasks through information retrieval as depicted in Fig 2. We demonstrate the model’s capability in a) decoding, wherein it can discern and retrieve images corresponding to neural activity recorded during experiments; b) encoding, where it exhibits its potential to predict neural activity patterns from visual stimuli; c) modality conversion, demonstrating the model’s ability to translate semantic content across different neural measurement modalities. This approach not only bridges the gap between neural activity and visual perception but also paves the way for a deeper understanding of how the brain processes and internalizes the visual world. Our work, therefore, stands at the




Figure 1: Schematic representation of our proposed model, illustrating the alignment of various neural datasets from different modalities into a unified representation space utilizing a frozen CLIP Image encoder.

intersection of neuroscience and artificial intelligence, offering a novel lens through which we can view and interpret the complex narrative of neural activity. It represents a step toward the direction of a foundation model for the neuroscience of vision, providing a framework for exploring and understanding the ways in which our brains engage with and make sense of the visual stimuli that permeate our environment.

1.1 RELATED WORK

Encoding and decoding in vision neuroscience have evolved from classical methods to advanced neural network-based models. fMRI has emerged as a promising tool to extract information with deep learning connecting biological hypotheses and computational models. In this context, classical encoding predicts brain activity from stimuli, while decoding anticipates stimuli from brain activity, and it has been shown that both tasks can benefit from a combined approach (Naselaris et al., 2011). Several models have been used for a wide variety of encoding and decoding approaches (Zafar et al., 2015) to analyze fMRI time series to decode visual stimuli, aiming to reconstruct the images linked to observed fMRI patterns or brain activity. Approaches like VAE-GAN have been applied to map fMRI activity to latent representations of human faces using linear models (VanRullen & Reddy, 2019). Additionally, sparse linear regression has predicted CNN features for natural images from fMRI data (Horikawa & Kamitani, 2017). Recently, diffusion models, noted for their excellent image generation, have become integral to decoding, often employing semantic techniques and multi-step decoding processes (Takagi & Nishimoto, 2023; Chen et al., 2022; Ferrante et al., 2023c; Ozcelik & VanRullen, 2023; Ferrante et al., 2023b). In general, recent advances leverage deep neural networks and large datasets to model complex visual and language representations, enhancing the accuracy of both encoding and decoding models (LeBel et al., 2023; Tang et al., 2023; Antonello et al., 2023; Caucheteux & King, 2022; Caucheteux et al., 2023; Défossez et al., 2023; Oota et al., 2023; Conwell et al., 2023). As a noteworthy example, the Algonauts challenge and subsequent studies emphasize the effectiveness of both pretrained multimodal transformers and tailored models for specific brain regions (Gifford et al., 2023; Adeli et al., 2023; Nguyen et al., 2023; Yang et al., 2023; Choksi et al., 2022). It is also worth mentioning tools like MindEye, which maps brain activity to multimodal latent spaces for precise image retrieval and reconstruction using a contrastive method (Scotti et al., 2023), and DREAM, which replicates image reconstruction from fMRI data, emulating the human visual system (Xia et al., 2023).

Advancements in high temporal resolution modalities have contributed significantly to the progress in encoding and decoding research within neuroscience. Previous efforts have leveraged linear models for tasks such as image classification from brain activity, prediction of brain activity based on image representations, and inter-modal comparison using representational similarity analysis (Robinson et al., 2019; Gifford et al., 2022). Although these studies incorporated image embeddings, their classification scope remained constrained. Progressing beyond linear approaches, deep neural networks have been employed to classify diverse data types, including speech, mental load, and images, from EEG signals (Défossez et al., 2023; Palazzo et al., 2020). Electroencephalography (EEG), Magnetoencephalography (MEG), and functional Magnetic Resonance Imaging (fMRI) are

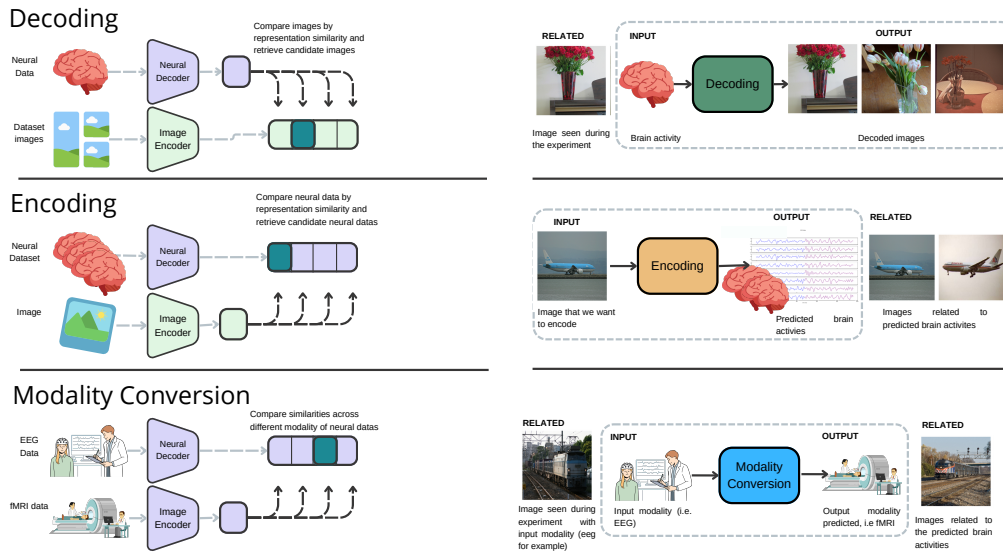


Figure 2: The top panel illustrates the 'Decoding' experiment, where neural data is processed to 'decode' and retrieve visually related images from a dataset. The middle panel depicts the 'Encoding' experiment, where an image is used to predict and retrieve neural data that could be associated with the visual perception of that image. The bottom panel shows the 'Modality Conversion' experiment, demonstrating the translation of neural data from one modality, such as EEG, into another, such as fMRI, aiming to find semantically similar brain activity across modalities

important in non-invasive brain research, each characterized by distinct advantages and limitations. EEG is esteemed for its exceptional temporal resolution, enabling precise temporal tracking of the brain's electrical activity and the observation of rapid neural dynamics. However, its spatial resolution is limited, posing challenges in accurately localizing neural activity. MEG, while offering temporal resolution comparable to EEG, provides slightly better spatial resolution by measuring neural activity's magnetic fields. Nonetheless, MEG's higher costs and lower accessibility limit its widespread use. In contrast, fMRI offers superior spatial resolution, allowing for detailed mapping of brain activity through blood flow changes, albeit with a temporal resolution inferior to that of EEG and MEG, restricting its ability to capture quick neural changes.

Our model extends beyond conventional constraints by using the principles of contrastive learning, a technique that has yielded promising outcomes in recent fMRI (Scotti et al., 2023) and MEG decoding investigations (Défossez et al., 2023). These studies, while they bear methodological resemblance to our work, predominantly concentrate on decoding within a singular modality, employing contrastive learning for data retrieval only and in conjunction with other generative methods for stimulus reconstruction. Contrastive learning differentiates between analogous (positive) and non-analogous (negative) data pairs, thereby facilitating the discernment and alignment of semantically congruent representations. In our approach, this methodology is applied not solely to decoding but also to encoding. Through the application of contrastive learning, our model establishes a bidirectional linkage between the visual and neural domains. Moreover, we introduce the concept of neural modality conversion, enabling the translation of semantic content from one neural measurement modality, such as EEG, into another, such as fMRI or MEG. This innovation opens new pathways for comprehensive neural analysis, fostering a more integrated understanding of brain functionality by capitalizing on the synergistic strengths of each modality. The primary contribution of our work is the development of a unified framework adept at managing decoding, encoding, and neural modality conversion, representing a significant advancement beyond existing models that are limited to a single task and modality. By aligning EEG, MEG, and fMRI data through contrastive learning, our model surmounts the inherent limitations of these modalities, offering a versatile infrastructure for the interpretation of neural signals.

2 MATERIAL AND METHODS

Our goal is to make a step towards a shared representation of neural data, i.e. a sort of "foundation model of neural representation". To solve this task we leveraged a powerful and well-established pretrained model for obtaining image representations - the CLIP Image encoder. We focus on vision processing, hence selecting a set of human vision datasets where neural activity is measured with different techniques like EEG, MEG and fMRI.

2.1 DATA

EEG: The EEG data for this study were sourced from the ImageNetEEG dataset (Spampinato et al., 2019), involving six participants and 40 ImageNet categories (Deng et al., 2009), totaling 2,000 images recorded at 1000 Hz. The recording protocol involved multiple sessions and sequences, resulting in 11,466 EEG sequences after quality filtering. Preprocessing included notch and band-pass filtering, normalization, and segmentation into 40 ms windows for time-frequency decomposition, producing EEG spectrogram images for model training. To avoid overestimated performances highlighted in (Li et al., 2018), a conservative data splitting approach was adopted as described in (Palazzo et al., 2020), ensuring more accurate performance assessments.

MEG: For Magnetoencephalography, Our methodology was evaluated using the "THINGS-MEG" dataset (Hebart et al., 2023). This dataset involved four participants (two female and two male, average age 23.25 years) who participated in 12 MEG sessions. During these sessions, participants were shown 22,448 distinct images from the THINGS database (Hebart et al., 2019) spanning 720 different categories. Out of this extensive collection, a smaller group of 200 images (each from a unique category) was repeatedly presented to the subjects. Each image was displayed for 500 milliseconds, followed by a variable fixation period ranging from 800 to 1200 milliseconds. To enhance our retrieval set and demonstrate our method's robustness, we also incorporated an additional 3,659 images from the THINGS dataset that were not shown to the participants. For MEG data preprocessing, we adopted a straightforward approach. The initial step involved downsampling the raw data from the 272 MEG radial gradiometer channels from 1,200 Hz to 120 Hz, followed by centering and standardizing the data. The MEG data was then segmented into epochs extending from 500 ms before to 1000 ms after the onset of each stimulus. The final preprocessing step involved baseline correction, achieved by deducting the average signal value, recorded from the beginning of each epoch to the stimulus onset, for every channel.

fMRI: In our study, the Natural Scenes Dataset (NSD) (Allen et al., 2022) was utilized. This extensive fMRI dataset includes data from eight individuals who were shown images from the COCO dataset. Our focus was on four of these subjects (consistent with subjects used in comparable decoding studies), resulting in a dataset comprising 8,859 images and 24,980 fMRI trials for training, and 982 images and 2,770 fMRI trials for testing per participant. To enhance the signal-to-noise ratio, images shown up to three times had their trials averaged. The spatial dimensionality of the fMRI data, recorded at a resolution of 1.8mm, was reduced to about 15,000 voxels. This reduction was achieved by applying the NSDGeneral ROI mask, which encompasses several visual areas. Selecting this ROI was crucial for improving the signal-to-noise ratio and reducing the complexity of the data, allowing for a more focused analysis of both low-level and high-level visual features. Temporally, dimensionality reduction was accomplished by using precalculated betas from a General Linear Model (GLM) with a fitted Hemodynamic Response Function (HRF), as detailed in the NSD paper, and further denoised as described therein.

2.2 NEURAL VISION ALIGNMENT

We focused on aligning the neural representations of different modalities with image representations derived from the CLIP model, specifically aiming to approximate the CLS (Classification) embeddings of images using the pretrained CLIP Image encoder, denoted as h . For each neural modality, we designed a distinct neural module, represented as f_n . This module is essentially a composite function, $f_n = g_n \circ a_n$, consisting of two primary components. The first component, a_n , is an alignment layer tasked with harmonizing the neural data from various subjects into a unified representation space. Once aligned, these representations are processed by g_n , a shared network that further refines them to closely match the visual representations produced by the image encoder. To illus-

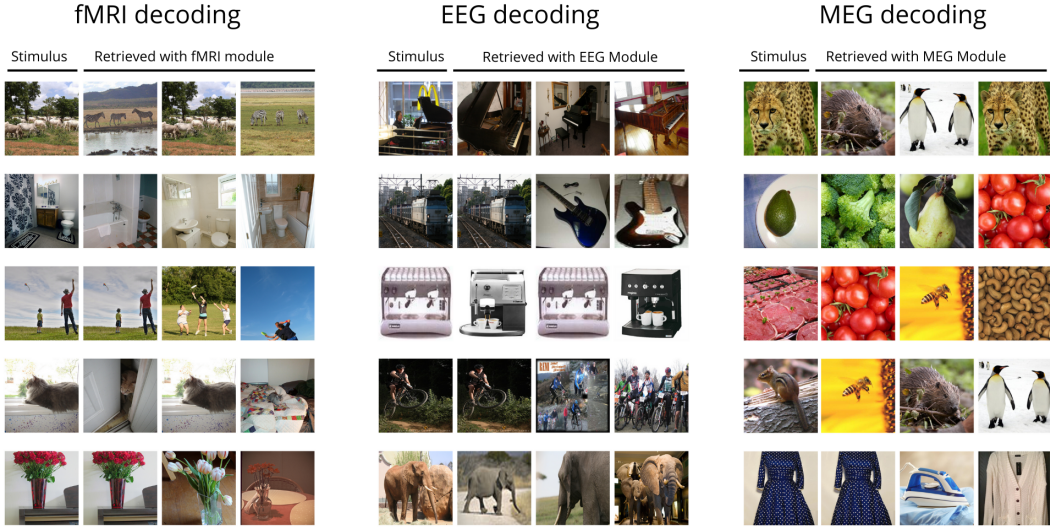


Figure 3: Comparative Results of Multimodal Neural Decoding. The image showcases the original visual stimuli and the images retrieved using decoding modules for fMRI, EEG, and MEG data. Each block corresponds to a different modality, illustrating the model’s ability to identify and retrieve images that closely resemble or are semantically related to the original stimulus

trate, consider a subject s who observes an image img while their neural activity n is being recorded. Using our neural model, we generate a representation $z_i = f(n, s) = g(a(n, s))$. Concurrently, we derive the corresponding image representation z_j through the image encoder: $z_j = h(img)$. Following their generation, these representations, are normalize and the contrastive CLIP loss is calculated forming the basis of our training regimen. The neural networks for MEG and EEG data are structured as convolutional neural networks (CNNs), offering an architecture for processing spatial and temporal patterns in the data. In contrast, the network for fMRI data is configured as a Multilayer Perceptron (MLP), suited for handling the high-dimensional and spatially complex nature of fMRI data.

All networks were implemented using the PyTorch framework and trained using the AdamW optimizer. The training parameters were a learning rate of 3×10^{-4} , weight decay set at 1×10^{-3} , and a batch size of 256. The training process spanned over 30 epochs, ensuring adequate learning while preventing overfitting, leaving room for performance improvement through hyperparameter and neural architecture search in future works. The contrastive loss function in our model aligns the representations of different modalities with the CLIP model’s image representations. Given two normalized representation vectors z_i and z_j , the function proceeds as follows: First, both vectors are normalized using the L2 norm to ensure they lie on a unit hypersphere $z_i = \frac{z_i}{\|z_i\|_2}$ and $z_j = \frac{z_j}{\|z_j\|_2}$.

Then, the similarities (logits) are computed by taking the dot product of z_i and the transpose of z_j , scaled by a temperature parameter τ , so logits = $\frac{z_i z_j^T}{\tau}$. Instead, the targets are defined as a sequence of indices, representing the matching pairs in the batch: $targets = [0, 1, 2, \dots, N - 1]$, then transformed into one-hot encoded vectors to compute the cross-entropy.

where N is the batch size. The loss is then calculated using the cross-entropy between the logits and the targets in their one-hot encoded version, so they’re 1 at index i and 0 everywhere else, thus:

$$\mathcal{L} = - \sum_{i=1}^N \log \left(\frac{\exp(\text{logits}_{ii}/\tau)}{\sum_{j=1}^N \exp(\text{logits}_{ij}/\tau)} \right) \tag{1}$$

This formulation of the contrastive loss function efficiently aligns the neural representations z_i with the corresponding image representations z_j by maximizing cosine similarity for positive pairs and

minimizing it for negative pairs, modulated by the temperature parameter τ , which is a hyperparameter set as 1 in our experiments but can be also be learned or modulated during the training.

2.3 EXPERIMENTS

We demonstrate the versatility of our model through three distinct experiments, each showcasing a unique application in neural data analysis.

Decoding Visual Information: This involves starting with neural data (EEG, MEG, or fMRI) and projecting it into a common representation space using the corresponding neural model. Concurrently, we process all images in the test set specific to each modality (comprising 337 for EEG, 2400 for MEG, and 982 for fMRI) and compute their similarities. The goal is to retrieve the top- n images that most closely match the neural signal representation. This process effectively allows us to "decode" the neural data into potential visual images that the subject might have been perceiving.

Encoding Experiment: Here, we begin with an image, pass it through the image encoder, and simultaneously process all neural data of a particular modality through the neural encoder. We then search for the top- n neural representations in the shared space and retrieve the corresponding neural data in the test set. This approach enables us to obtain "encoded" neural representations corresponding to the given image.

Modality Conversion: Here we capitalize on the alignment of all modalities in the same representation space. For instance, given the fMRI representation of a subject who has viewed a specific image, we might ask: what could be the EEG or MEG activity resulting in viewing a semantically similar image? To answer this, we encode the sample from our input modality and the target search set from the desired output modality, selecting the top- n matches based on similarity. To validate the effectiveness of this modality conversion, we compare the images associated with the source modality during data acquisition with those linked to the target modality.

These experiments collectively illustrate the robustness and multifaceted capabilities of our model, offering significant advancements in the fields of neural encoding, decoding, and inter-modality translation.

2.4 EVALUATION

To assess the performance of our model, we employed various metrics that gauge its proficiency in extracting relevant semantic information from neural data.

Decoding Performance Evaluation: For decoding, the evaluation methodology is straightforward. The ImageNetEEG and THINGS MEG datasets have distinct classes (40 and 720, respectively), allowing us to calculate and compare top1 and top5 accuracy directly against chance levels and established baselines like (Palazzo et al., 2020). In contrast, the Natural Scene Dataset (NSD) used for fMRI comprises complex scenes from the COCO dataset without distinct classes. To evaluate decoding performance in this context, we employed the CLIP 2-way accuracy metric, comparing with the state of the art (Scotti et al., 2023). For consistency and ease of comparison, we extended this measure to the EEG and MEG settings as well.

Encoding Performance Measurement: In the encoding scenario, where images are input to retrieve neural data, we faced the challenge of neural data being inherently difficult to interpret and visualize. To ascertain whether relevant information was captured, we relied on an indirect metric. Each sample of neural data is associated with a specific image viewed during the vision experiments. Encoding these images, we get candidate neural representations of stimuli. We then compute the CLIP 2-way accuracy between the related neural activity relative encoded image (ground truth) and the retrieved activities with the encoding model. This process involves starting with an image, obtaining candidate neural data, and then visually comparing the representations of the related images of these candidates with the original image. A successful encoding process would typically result in the selection of neural data associated with the observation of semantically similar images.

Modality Conversion Evaluation: The approach for evaluating modality conversion is similar. We initiate the process with neural data from one modality and obtain representations in another. Of course, given that all datasets are from different subjects, we also get activity on another subject, the one that in principle has the closest representation in the common space. The performance of

Table 1: Various metrics to assess model performance in the decoding experiment. "Baselines" columns refer to top1 and top5 chance level, while "retrieval dataset size" is useful to put CLIP 2-way accuracy in context with other works

Neural Module	top1 accuracy	top5 accuracy	CLIP 2 way	baseline accuracy (%)	baseline top5 (%)	Retrieval dataset size	Number of classes
EEG (ImageNet)	40,0	54,3	79,4	2,5	12,5	332	40
MEG (THINGS)	1,2	6,1	60,1	0,13	0,65	2400	720
fMRI (NSD)	-	-	90,3	-	-	982	-

this aspect of our model is gauged using the CLIP 2-way accuracy between the images related to the source and target modalities. This metric effectively measures how well our model translates semantic information across different neural modalities.

Overall, these diverse evaluation strategies offer a comprehensive understanding of our model’s capabilities in decoding, encoding, and modality conversion, ensuring a thorough validation of its effectiveness in handling complex neural datasets.

3 RESULTS

The model’s performance in decoding neural data into corresponding visual stimuli is quantified and presented in Table 1. The EEG module achieved a top1 accuracy of 40,0% and a top5 accuracy of 54,3%, which is a substantial improvement over the baseline chance level accuracy figures of 2,5% for top1 and 12,5% for top5, Notably, this module’s CLIP 2-way accuracy reached 79,4%, indicating the model’s capability to decode EEG data with high reliability. Unfortunately, directly comparing these performances with literature could be difficult, because some recent work with apparently impressive performances (Bai et al., 2023; Palazzo et al., 2017; Kavasidis et al., 2017) on this dataset have been seen to rely on contamination between train and test data due to an incorrect use of preprocessing choices on raw data for this event design as stated in (Li et al., 2018). When comparing results with work that explicitly preprocess data in order to avoid this confounding factor, we found performances that are on par with recent works (Palazzo et al., 2020; Ferrante et al., 2023a), reporting top1 accuracy within range (39-45%) for multisubject network trained for classification. For the MEG module, which faced the greater complexity of the THINGS dataset (720 classes), the model achieved a top1 accuracy of 1.2% and a top5 accuracy of 6.1%, which corresponds to 10 times the baseline chance level accuracies of 0.13% for top1 and 0.65% for top5. The CLIP 2-way accuracy stood at 60.1%, This demonstrates the model’s adeptness in handling a more extensive and varied set of images, comprising a total of 2,400 images. Again, these results are comparable with recent work on the same dataset which focused only on decoding of MEG images (Benchetrit et al., 2023) where top5 accuracy was tested in several settings with performances in range [1-8 %]. The fMRI module’s decoding performance was particularly notable, achieving a CLIP 2-way accuracy of 90.3% with the NSD dataset. Our results align well with findings from recent literature (Takagi & Nishimoto, 2023; Ferrante et al., 2023b; Ozcelik & VanRullen, 2023; Chen et al., 2022; Scotti et al., 2023), though direct comparisons are challenging due to the varied focus and methodologies of these studies. Many of these works concentrate on the detailed reconstruction of stimuli using complex pipelines that involve regressing fMRI data to a latent space, generating images, and then computing CLIP 2-way accuracy between the generated and actual images. These processes typically involve training individual models for each subject. Despite these differences, the performance metrics reported in these studies, which generally fall within the range of [77-95 %], are comparable to our results. This situates our approach within the high-performance spectrum for multisubject settings. This high accuracy, obtained with a retrieval set of 982 images, underscores the model’s proficiency in decoding complex scene representations from fMRI data. Fig 3 provides details on the model’s decoding capabilities with fMRI, EEG, and MEG data. In fMRI Decoding, the model’s quality is evident. For example, when presented with a stimulus of grazing animals, the fMRI module accurately retrieves other images of animals in similar pastoral settings. This level of detail suggests the model’s strong semantic grasp, as it not only identifies the subject of the image but also the context

in which it is situated. The EEG Decoding column showcases a broad understanding of the visual stimulus categories. Good examples are showcased by the piano and elephant cases (first and last row of EEG panel in Fig 3), indicating its capacity to capture the broader concepts of objects and animals. MEG Decoding presents a blend of moderately related and thematically similar images. A notable example is the retrieval of images of leopards and penguins in response to a stimulus of the jaguar (first row), demonstrating the model’s nuanced semantic retrieval. However, the module also retrieves images that are thematically related but not identical, such as different fruits and vegetables in response to a stimulus of a specific type, suggesting a wider semantic reach of the MEG module compared to fMRI.

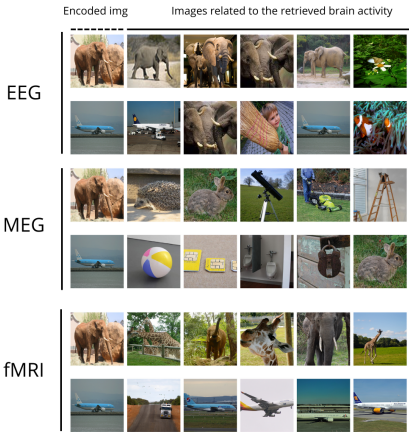


Figure 4: Encoding Experiment Results Displaying Image-to-Brain Activity Correlation. Rows illustrate the results for EEG, MEG, and fMRI modalities. The leftmost images are the encoded stimuli, and the subsequent images represent images related to the brain activities retrieved by the model.

related scenes, affirming the model’s high precision in encoding complex visual information. Overall, the figure illustrates the model’s capability to accurately encode visual stimuli into neural representations, suggesting its potential application in predicting brain activity from visual inputs across EEG, MEG, and fMRI data modalities.

Lastly, the modality conversion experiment results, as detailed in Table 2, indicated the model’s capability to accurately transform neural information across modalities. The normalized accuracies for conversions like fMRI to MEG were exceptionally high (95.40%), highlighting the model’s ability to preserve semantic content through the conversion process. These conversions demonstrate the model’s potential to provide a harmonized representation of brain activity, bridging the gap between different brain imaging techniques.

Collectively, these results underscore the capabilities of our model that can solve several tasks, reaching high performances on par with recent literature in neural decoding tasks, showing also its potential in encoding visual content into neural patterns, and its proficiency in converting neural information across modalities, setting a foundation for future breakthroughs in multimodal neural data analysis and interpretation.

4 DISCUSSION

The development of our neural foundation model stands as a pioneering stride towards an integrated understanding of the brain’s mechanisms through neural data. This work signifies the first step in

Table 2: Conversion modality CLIP-2way accuracies and their normalized values respect of the decoding performances.

Conversion	Clip 2 way decoding accuracy	Normalized Clip 2 way decoding accuracy
fMRI to EEG	0.6710	0.8370
MEG to EEG	0.6790	0.8470
fMRI to MEG	0.5679	0.9540
EEG to MEG	0.5594	0.9396
EEG to fMRI	0.7648	0.8598
MEG to fMRI	0.7928	0.8912

creating a foundational framework akin to what has been seen with large language models in the field of natural language processing. It encapsulates a multi-modal approach that not only decodes but also aligns neural representations from a variety of datasets and modalities, bringing us closer to a cohesive neural language. A pivotal aspect of this model is its capacity for multi-modal (and subject) representation alignment, effectively creating a shared representation space that harmonizes individual variability. This is particularly reminiscent of the convergence of different languages and dialects into a singular, coherent narrative—where the model serves as an interpreter of the brain’s complex ‘dialects’ of activity. However, aligning data from disparate neural recording modalities comes with several challenges, ranging from technical discrepancies to the resolution of spatiotemporal data. This work has navigated some of these complexities, yet the integration process remains a sophisticated and elaborate task. However, the model presented is not without its limitations. It has currently non-generative nature and the reliance on diverse, pre-existing datasets indicate that it remains a proof-of-concept. Looking to the future, our goal is to turn this model into a generative tool that could revolutionize data augmentation and facilitate virtual experiments. The addition of further modalities such as language and audio, alongside more extensive fMRI, EEG, and MEG data, could pave the way for a comprehensive “latent brain representation.” This representation would transcend individual modalities, offering a more holistic view of brain activity. In addition to technical advancements in the field, one must not forget the critical issue of neural data privacy. As our models become increasingly capable of decoding detailed information from neural signals, the imperative for privacy safeguards grows. All datasets used in this study are in the public domain, and participant consent was meticulously obtained. Nevertheless, future developments could lead to models that decode personal data from minimal scanning, necessitating minimal cooperation from participants. This raises the specter of privacy concerns, as well as the ethical use of such technology. It is crucial that we engage in proactive discussions on these topics to avoid potential misuse of neural decoding technologies and ensure that model-generated content can be distinguished from true subject experiences, preventing the propagation of harmful material. In sum, while this first step towards a neural foundation model marks a possible advancement in our approach to understanding and interpreting neural data, it also beckons us to contemplate the ethical framework within which such technology should operate. As we enhance the model’s capabilities and expand its applications, we must concurrently fortify the ethical boundaries that will preserve the privacy and integrity of individual neural data.

5 CONCLUSION

This paper introduces a new step towards a neural foundation model that aligns representations of multi-modal neural datasets using contrastive learning, marking a significant advance in the field of neuroscience. Our model has demonstrated considerable success in decoding, encoding, and converting neural signals, showing its potential to unravel the complex semantics of brain activity. While promising, the model’s current non-generative nature and reliance on diverse datasets indicate areas for future enhancement. The next steps involve expanding the model’s capabilities to include generative abilities and additional modalities, moving towards a comprehensive representation of brain activity. Crucially, this research also highlights the emergent issue of neural data privacy, necessitating a collaborative effort to establish ethical guidelines for future advancements. As we continue to explore the depths of neural data interpretation, we remain committed to advancing scientific understanding while upholding the utmost respect for individual privacy and data integrity.

REFERENCES

- Hossein Adeli, Sun Minni, and Nikolaus Kriegeskorte. Predicting brain activity using transformers. *bioRxiv*, 2023. doi: 10.1101/2023.08.02.551743. URL <https://www.biorxiv.org/content/early/2023/08/05/2023.08.02.551743>.
- Emily J. Allen, Ghislain St-Yves, Yihan Wu, Jesse L. Breedlove, Jacob S. Prince, Logan T. Dowdle, Matthias Nau, Brad Caron, Franco Pestilli, Ian Charest, J. Benjamin Hutchinson, Thomas Naselaris, and Kendrick Kay. A massive 7t fmri dataset to bridge cognitive neuroscience and artificial intelligence. *Nature Neuroscience*, 25(1):116–126, Jan 2022. ISSN 1546-1726. doi: 10.1038/s41593-021-00962-x. URL <https://doi.org/10.1038/s41593-021-00962-x>.
- Richard Antonello, Aditya Vaidya, and Alexander G. Huth. Scaling laws for language encoding models in fmri, 2023.
- Yunpeng Bai, Xintao Wang, Yan pei Cao, Yixiao Ge, Chun Yuan, and Ying Shan. Dreamdiffusion: Generating high-quality images from brain eeg signals, 2023.
- Yohann Benchetrit, Hubert Banville, and Jean-Rémi King. Brain decoding: toward real-time reconstruction of visual perception, 2023.
- C. Caucheteux and JR. King. Brains and algorithms partially converge in natural language processing. *Communications Biology*, 5:134, 2022. doi: 10.1038/s42003-022-03036-1.
- C. Caucheteux, A. Gramfort, and JR. King. Evidence of a predictive coding hierarchy in the human brain listening to speech. *Nature Human Behaviour*, 7:430–441, 2023. doi: 10.1038/s41562-022-01516-2.
- Zijiao Chen, Jiaxin Qing, Tiange Xiang, Wan Lin Yue, and Juan Helen Zhou. Seeing beyond the brain: Conditional diffusion model with sparse masked modeling for vision decoding, 2022.
- Bhavin Choksi, Milad Mozafari, Rufin VanRullen, and Leila Reddy. Multimodal neural networks better explain multivoxel patterns in the hippocampus. *Neural Networks*, 154:538–542, 2022. ISSN 0893-6080. doi: <https://doi.org/10.1016/j.neunet.2022.07.033>. URL <https://www.sciencedirect.com/science/article/pii/S0893608022002982>.
- Colin Conwell, Jacob S. Prince, Kendrick N. Kay, George A. Alvarez, and Talia Konkle. What can 1.8 billion regressions tell us about the pressures shaping high-level visual representation in brains and machines? *bioRxiv*, 2023. doi: 10.1101/2022.03.28.485868. URL <https://www.biorxiv.org/content/early/2023/07/01/2022.03.28.485868>.
- A. Défossez, C. Caucheteux, J. Rapin, et al. Decoding speech perception from non-invasive brain recordings. *Nature Machine Intelligence*, 5:1097–1107, 2023. doi: 10.1038/s42256-023-00714-5.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pp. 248–255. Ieee, 2009.
- Matteo Ferrante, Tommaso Boccato, Stefano Bargione, and Nicola Toschi. Decoding visual brain representations from electroencephalography through knowledge distillation and latent diffusion models, 2023a.
- Matteo Ferrante, Tommaso Boccato, Furkan Ozcelik, Rufin VanRullen, and Nicola Toschi. Multimodal decoding of human brain activity into images and text. In *UniReps: the First Workshop on Unifying Representations in Neural Models*, 2023b. URL <https://openreview.net/forum?id=rGCabZfV3d>.
- Matteo Ferrante, Tommaso Boccato, and Nicola Toschi. Semantic brain decoding: from fmri to conceptually similar image reconstruction of visual stimuli, 2023c.
- A. T. Gifford, B. Lahner, S. Saba-Sadiya, M. G. Vilas, A. Lascelles, A. Oliva, K. Kay, G. Roig, and R. M. Cichy. The algonauts project 2023 challenge: How the human brain makes sense of natural scenes, 2023.

- Alessandro T Gifford, Kshitij Dwivedi, Gemma Roig, and Radoslaw M Cichy. A large and rich EEG dataset for modeling human visual object recognition. *Neuroimage*, 264(119754):119754, December 2022.
- Martin N Hebart, Adam H Dickter, Alexis Kidder, Wan Y Kwok, Anna Corriveau, Caitlin Van Wicklin, and Chris I Baker. THINGS: A database of 1,854 object concepts and more than 26,000 naturalistic object images. *PLoS One*, 14(10):e0223792, October 2019.
- Martin N Hebart, Oliver Contier, Lina Teichmann, Adam H Rockter, Charles Y Zheng, Alexis Kidder, Anna Corriveau, Maryam Vaziri-Pashkam, and Chris I Baker. Things-data, a multi-modal collection of large-scale datasets for investigating object representations in human brain and behavior. *eLife*, 12:e82580, feb 2023. ISSN 2050-084X. doi: 10.7554/eLife.82580. URL <https://doi.org/10.7554/eLife.82580>.
- Tomoyasu Horikawa and Yukiyasu Kamitani. Generic decoding of seen and imagined objects using hierarchical visual features. *Nature Communications*, 8(1):15037, August 2017. ISSN 2041-1723. doi: 10.1038/ncomms15037. URL <http://www.nature.com/articles/ncomms15037>.
- Isaak Kavasidis, Simone Palazzo, Concetto Spampinato, Daniela Giordano, and Mubarak Shah. Brain2image: Converting brain signals into images. In *Proceedings of the 25th ACM International Conference on Multimedia*, MM '17, pp. 1809–1817, New York, NY, USA, 2017. Association for Computing Machinery. ISBN 9781450349062. doi: 10.1145/3123266.3127907. URL <https://doi.org/10.1145/3123266.3127907>.
- A. LeBel, L. Wagner, S. Jain, et al. A natural language fmri dataset for voxelwise encoding models. *Scientific Data*, 10:555, 2023. doi: 10.1038/s41597-023-02437-z.
- Ren Li, Jared S. Johansen, Hamad Ahmed, Thomas V. Ilyevsky, Ronnie B Wilbur, Hari M Bhadravaj, and Jeffrey Mark Siskind. Training on the test set? an analysis of spampinato et al. [31], 2018.
- Thomas Naselaris, Kendrick N. Kay, Shinji Nishimoto, and Jack L. Gallant. Encoding and decoding in fMRI. *NeuroImage*, 56(2):400–410, May 2011. ISSN 10538119. doi: 10.1016/j.neuroimage.2010.07.073. URL <https://linkinghub.elsevier.com/retrieve/pii/S1053811910010657>.
- Xuan-Bac Nguyen, Xudong Liu, Xin Li, and Khoa Luu. The algonauts project 2023 challenge: Uark-ualbany team solution, 2023.
- Subba Reddy Oota, Manish Gupta, Raju S. Bapi, Gael Jobard, Frederic Alexandre, and Xavier Hinaut. Deep Neural Networks and Brain Alignment: Brain Encoding and Decoding (Survey), July 2023. URL <http://arxiv.org/abs/2307.10246>. arXiv:2307.10246 [cs, q-bio].
- Furkan Ozcelik and Rufin VanRullen. Brain-diffuser: Natural scene reconstruction from fmri signals using generative latent diffusion, 2023.
- Simone Palazzo, Concetto Spampinato, Isaak Kavasidis, D. Giordano, and M. Shah. Generative adversarial networks conditioned by brain signals. pp. 3430–3438, 10 2017. doi: 10.1109/ICCV.2017.369.
- Simone Palazzo, Concetto Spampinato, Joseph Schmidt, Isaak Kavasidis, Daniela Giordano, and Mubarak Shah. Correct block-design experiments mitigate temporal correlation bias in eeg classification, 2020.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision, 2021.
- Amanda K Robinson, Tijl Grootswagers, and Thomas A Carlson. The influence of image masking on object representations during rapid serial visual presentation. *Neuroimage*, 197:224–231, August 2019.

- Paul S. Scotti, Atmadeep Banerjee, Jimmie Goode, Stepan Shabalin, Alex Nguyen, Ethan Cohen, Aidan J. Dempster, Nathalie Verlinde, Elad Yundler, David Weisberg, Kenneth A. Norman, and Tanishq Mathew Abraham. Reconstructing the mind’s eye: fmri-to-image with contrastive learning and diffusion priors, 2023.
- Concetto Spampinato, Simone Palazzo, Isaak Kavasidis, Daniela Giordano, Mubarak Shah, and Nasim Souly. Deep learning human mind for automated visual classification, 2019.
- Yu Takagi and Shinji Nishimoto. High-resolution image reconstruction with latent diffusion models from human brain activity. *bioRxiv*, 2023. doi: 10.1101/2022.11.18.517004. URL <https://www.biorxiv.org/content/early/2023/03/11/2022.11.18.517004>.
- J. Tang, A. LeBel, S. Jain, et al. Semantic reconstruction of continuous language from non-invasive brain recordings. *Nature Neuroscience*, 26:858–866, 2023. doi: 10.1038/s41593-023-01304-9.
- Rufin VanRullen and Leila Reddy. Reconstructing faces from fmri patterns using deep generative neural networks. *Communications Biology*, 2(1):193, May 2019. ISSN 2399-3642. doi: 10.1038/s42003-019-0438-y. URL <https://doi.org/10.1038/s42003-019-0438-y>.
- Weihao Xia, Raoul de Charette, Cengiz Öztireli, and Jing-Hao Xue. Dream: Visual decoding from reversing human visual system, 2023.
- Huzheng Yang, James Gee, and Jianbo Shi. Memory Encoding Model, August 2023. URL <http://arxiv.org/abs/2308.01175>. arXiv:2308.01175 [cs].
- Raheel Zafar, Aamir Saeed Malik, Nidal Kamel, Sarat C. Dass, Jafri M. Abdullah, Faruque Reza, and Ahmad Helmy Abdul Karim. Decoding of visual information from human brain activity: A review of fMRI and EEG studies. *Journal of Integrative Neuroscience*, 14(02):155–168, June 2015. ISSN 0219-6352, 1757-448X. doi: 10.1142/S0219635215500089. URL <http://www.worldscientific.com/doi/abs/10.1142/S0219635215500089>.