

# MindLLM: A Subject-Agnostic and Versatile Model for fMRI-to-Text Decoding

Weikang Qiu<sup>1</sup> Zheng Huang<sup>\*2</sup> Haoyu Hu<sup>\*3</sup> Aosong Feng<sup>1</sup> Yujun Yan<sup>2</sup> Rex Ying<sup>1</sup>

## Abstract

Decoding functional magnetic resonance imaging (fMRI) signals into text has been a key challenge in the neuroscience community, with the potential to advance brain-computer interfaces and uncover deeper insights into brain mechanisms. However, existing approaches often struggle with suboptimal predictive performance, limited task variety, and poor generalization across subjects. In response to this, we propose MindLLM, a model designed for subject-agnostic and versatile fMRI-to-text decoding. MindLLM consists of an fMRI encoder and an off-the-shelf LLM. The fMRI encoder employs a neuroscience-informed attention mechanism, which is capable of accommodating subjects with varying input shapes and thus achieves high-performance subject-agnostic decoding. Moreover, we introduce Brain Instruction Tuning (BIT), a novel approach that enhances the model’s ability to capture diverse semantic representations from fMRI signals, facilitating more versatile decoding. We evaluate MindLLM on comprehensive fMRI-to-text benchmarks. Results demonstrate that our model outperforms the baselines, improving downstream tasks by 12.0%, unseen subject generalization by 24.5%, and novel task adaptation by 25.0%. Furthermore, the attention patterns in MindLLM provide interpretable insights into its decision-making process. Code is available at <https://github.com/Graph-and-Geometric-Learning/MindLLM>.

## 1. Introduction

Decoding human brain activity (fMRI) to text has sparked significant interest within the neuroscience community (Xia et al., 2024; Chen et al., 2023a; Luo et al., 2023;

<sup>\*</sup>Equal contribution <sup>1</sup>Yale University <sup>2</sup>Dartmouth College <sup>3</sup>University of Cambridge. Correspondence to: Weikang Qiu <weikang.qiu@yale.edu>.

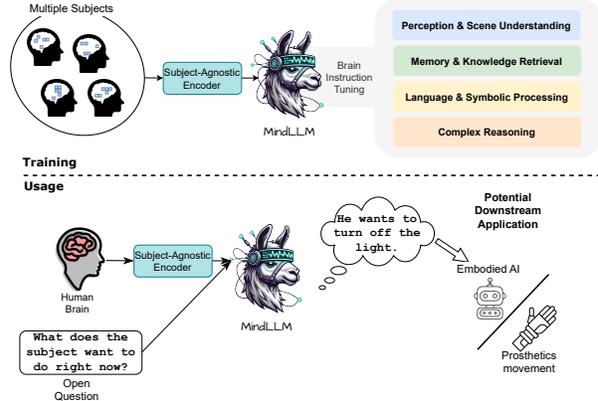


Figure 1. The overview of our method. MindLLM is equipped with a subject-agnostic fMRI encoder and an off-the-shelf LLM. MindLLM is trained on multiple subjects with varying input shapes and an instruction-tuning dataset, aiming to encode different facets of semantic information in fMRI. After training, MindLLM is capable of various text decoding tasks. One application is that the decoded contents can be used to achieve neural control of existing systems that are not designed for it.

Hmamouche et al., 2024). The ability to translate brain activity patterns into natural language carries both academic and societal importance. For neuroscientists, it provides deeper and novel insights into cognition, behavior, and consciousness (Qiu et al., 2023; Luo et al., 2023). On a societal level, it presents opportunities for medical applications and improves human-computer interaction (HCI) (Bernal et al., 2022; Du et al., 2022). For example, for individuals with speech impairments, this technology could restore communication capabilities, enabling them to express their thoughts effortlessly (Card et al., 2024). Moreover, as shown in Figure 1, it benefits healthy individuals by allowing neural control of digital devices, such as embodied AIs or prosthetic limbs, allowing for more intuitive and precise movements.

Despite its potential, decoding brain activity to language still faces significant challenges. One major obstacle is the need for versatile decoding tailored to specific applications. For example, decoding may aim to translate a subject’s movement intention to control a prosthesis, or to interpret abstract thoughts or memories. Traditional models fail to accommodate such diverse requirements. To address this, UMBRAE (Xia et al., 2024) integrates a Visual Language Model (VLM) (Chen et al., 2023b) and learns to map from

fMRI data to corresponding stimulus images. While this approach achieves versatility to some extent, it remains constrained to tasks directly tied to the current stimulus image and cannot address broader tasks, such as retrieving memories of past visual experiences.

Another critical challenge lies in designing a unified and subject-agnostic architecture. Current methods of brain multimodal decoding mostly rely on a preprocessing step: selecting responsive voxels by comparing the task-based fMRI to the resting-state fMRI. The selection typically results in higher performance compared to using whole-brain data. However, the varying number and irregular spatial distribution of selected voxels across subjects pose significant challenges for developing a unified architecture. To this end, recent studies (Wang et al., 2024a;b) have proposed pooling or sampling voxels to standardize input dimensions. However, as illustrated in Figure 3, these methods still suffer from the loss of spatial information and uneven representations of certain areas, ultimately compromising performance.

**Present Work** Here we present MindLLM, a subject-agnostic and versatile model for fMRI-to-text decoding. Our approach consists of a subject-agnostic fMRI encoder and an off-the-shelf LLM. The subject-agnostic fMRI encoder incorporates a neuroscience-informed attention layer with learnable *queries*, enabling dynamic feature extraction by leveraging both spatial information and neuroscientific priors of voxels, thereby significantly enhancing prediction accuracy. The design of *values* and *keys* separates the voxel’s functional information—which is largely consistent across individuals—from its fMRI value, allowing the model to benefit from shared priors across subjects and enhancing generalization to novel subjects. Moreover, to address the challenge of versatile decoding, we propose Brain Instruction Tuning (BIT). BIT trains the model using a diverse dataset that employs images as intermediaries, encompassing tasks designed to capture diverse aspects of semantic information encoded in fMRI data, including perception & scene understanding, memory & knowledge retrieval, language & symbolic processing, and complex reasoning. Figure 1 illustrates the corresponding components.

We evaluate our model on comprehensive benchmarks. Results reveal it outperforms baselines with 12.0% average improvement in various downstream tasks and 24.5% improvement in generalization on unseen subjects. Additionally, we show that our model adapts effectively to novel tasks, demonstrating high customizability and flexibility in real-world applications. Furthermore, our analysis of attention weights offers valuable insights into the working mechanism of our fMRI encoder.

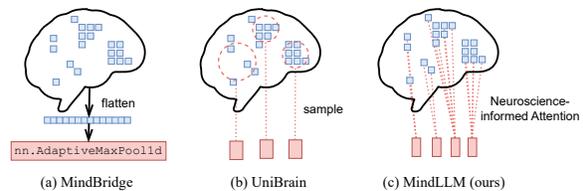


Figure 2. Comparison between our model and previous unified models. MindBridge (Wang et al., 2024a) flattens the voxels and adaptively pools them to a fixed dimension, which overlooks the rich information in positions. UniBrain (Wang et al., 2024b) uniformly samples a subset of voxels and aggregates their neighbors. Different from these methods, we propose neuroscience-informed attention, where each query token attends to all voxels, which minimizes potential information loss in pooling or sampling.

## 2. Related Works

### Brain-Conditioned Text Generation

This line of research mostly focuses on decoding perceived visual stimuli into natural language from fMRI signals. MindGPT (Chen et al., 2023a), UniBrain (Mai & Zhang, 2023) and BrainCap (Ferrante et al., 2023) employ an fMRI encoder guided by CLIP (Radford et al., 2021) and use a language model (Radford et al., 2019; Wang et al., 2022) to decode natural language from the encoded representations. BrainChat (Huang, 2024) utilizes multiple pretraining strategies (Devlin, 2018; He et al., 2022; Yu et al., 2022) to align fMRI with image and text embeddings. These methods fall short in performance and versatility. UMBRAE (Xia et al., 2024) proposes to learn a mapping from fMRI to stimulus images, which later serves as a proxy input for an off-the-shelf visual language model (VLM). Although they achieve performance improvements, the strategy prevents the model from performing tasks that are not directly related to the stimulus images (e.g., answering memory-related questions). In contrast, our end-to-end Brain Instruction Tuning (BIT) ensures seamless and versatile fMRI-to-text decoding, offering the potential to tackle tasks beyond vision-related ones.

**Cross-subjects Decoding** In voxel-level machine learning for brain decoding, the number of voxels varies between subjects (Allen et al., 2022). Most prior works (Scotti et al., 2024a;b) use an MLP for each subject individually. However, due to the fixed input size required by MLP architectures, these models cannot handle varying input shapes. As illustrated in Figure 2, MindBridge (Wang et al., 2024a) proposed to use an adaptive max pooling layer to standardize the input shapes. However, unlike images, which are considered translation invariance, positions in fMRI carry specific bio-informative significance that pooling operations may overlook. UniBrain (Wang et al., 2024b) proposed to sample groups of voxels. Such a sampling strategy, on the one hand, may lead to information loss if some voxels are

not included in any group. On the other hand, the irregular spatial distribution of 3D voxels with varying density and curvature may result in underrepresentation or overrepresentation of certain areas. Different from these methods, our model employs a neuroscience-informed attention mechanism that accounts for every single voxel while preserving their bio-informative positional information, ensuring a more comprehensive and precise representation.

**Multi-Modal Large Language Model** Aiming to augment the perceptual capacities of Large Language Models (LLMs), there has been a growing interest in extending them to handle multiple modalities within a unified model. Numerous studies have attempted to incorporate modalities such as images (Alayrac et al., 2022; Zhang et al., 2023b; Wang et al., 2023b), videos (Cheng et al., 2024; Kondratyuk et al., 2023; Zhang et al., 2023a), and point clouds (Xu et al., 2023; Qi et al., 2025). OneLLM (Han et al., 2024a) stands out by aligning eight different modalities, including fMRI, with language. However, their approach employs an individual convolution network for each subject instead of a unified architecture for fMRI encoding across subjects, which restricts its applicability to new subjects in real-world scenarios. Furthermore, the approach solely relies on captions as textual annotations, which limits the model’s capability for versatile fMRI decoding.

**Multimodal fMRI Decoding** Beyond text, recent studies have explored decoding brain signals into other modalities, including images (Scotti et al., 2024b; Wang et al., 2024a; Chen et al., 2023c), videos (Sun et al., 2024; Chen et al., 2023d), speech (Chen et al., 2024) and music (Denk et al., 2023), as well as in reverse directions (Toneva & Wehbe, 2019). The encoder architectures in these works typically either adopt transformers for time-series data, which is distinct from our voxel-based setting, or fall into one of the architectures we have discussed in previous paragraphs. Therefore, focusing on text generation, our work improves the fMRI encoder architecture for voxel-level inputs, and designs novel objectives (i.e., brain instruction tuning) for our target modality.

### 3. Method

In this section, we propose a neuroscience-informed fMRI encoder designed to achieve high-performance, subject-agnostic decoding. To further enable versatile decoding, we introduce the construction of a brain instruction tuning dataset, which captures diverse semantic representations encoded in fMRI data.

#### 3.1. Method Overview

As illustrated in Figure 3, our model consists of an fMRI encoder  $f_\theta$  and an off-the-shelf LLM. In practice, we use

Vicuna-7b (Zheng et al., 2023) as our LLM to maintain consistency with our baseline (Xia et al., 2024). For each sample, let  $\mathbf{v} = [v_1, v_2, \dots, v_N] \in \mathbb{R}^N$  be the fMRI signals of input voxels, where  $N$  is the number of voxels. Note that  $N$  varies between different subjects, ranging from 12,682 to 17,907 in the dataset we use (Allen et al., 2022).

The fMRI encoder  $f_\theta$ , featuring a neuroscience-informed attention layer, encodes  $\mathbf{v}$  to fMRI tokens  $X_v = [\mathbf{x}_{v,1}, \mathbf{x}_{v,2}, \dots, \mathbf{x}_{v,L}] \in \mathbb{R}^{d \times L}$ , where  $L$  is the number of tokens and  $d$  is the dimension of token embeddings. We then prepend these learned fMRI tokens to the language tokens in the BIT dataset we propose.

#### 3.2. fMRI Encoder

As mentioned before, currently most models for fMRI decoding can not handle varying input shapes and are not subject-agnostic, with only a few exceptions (Wang et al., 2024b). However, these exceptions still suffer from information loss and uneven representations of certain brain areas. To this end, we propose a novel neuroscience-informed attention mechanism to accommodate varying voxel numbers across subjects, enabling a subject-agnostic encoding strategy. Below we talk about the design of *queries*  $\{\mathbf{q}_i\}$ , *keys*  $\{\mathbf{k}_i\}$  and *values*  $\{\mathbf{v}_i\}$  in the attention layer. For *values*, we directly use the fMRI signal of each voxel, which means  $\mathbf{v}_i = v_i \in \mathbb{R}$ . Making each voxel a *value* token maximally prevents information loss compared to pooling- (Wang et al., 2024a) or sampling-based (Mai & Zhang, 2023) methods. The *queries* are randomly initialized and learnable. We expect each *query* to represent a certain pattern of the brain (refer to visualizations in Section 4.8). The design of *keys* will be discussed below.

**Exclude fMRI values from keys** The vanilla cross attention (Zhu et al., 2020; Vaswani, 2017) derives both *keys* and *values* from the same input source. However, we found this would lead to poor performance in fMRI. We argue the reason: different from images or text, which are usually considered translation-invariant, the positions of voxels carry specific brain *functional information*, as voxels in different areas are associated with distinct brain functions. Consequently, a voxel’s position alone can theoretically serve as effective *keys* for attention weight computation (McManus, 2002; Zhang et al., 2023c; Liu et al., 2024). Including fMRI values into *keys*, however, introduces additional noise instead of valuable information, thus resulting in poorer performance. Moreover, since brain regions tend to serve similar functions across individuals, decoupling voxel positions from fMRI signals can facilitate the sharing of priors across subjects, potentially improving generalization to unseen subjects.

In light of this, instead of the vanilla cross attention, which derives the *keys* and *values* from the same inputs, we ex-

clude the fMRI value of each voxel and use its positional information alone as its *key* embedding. The positional information is encoded from the coordinates of each voxel, i.e.  $k_i^{\text{pos}} = \text{PE}(c_i)$  for the  $i$ -th voxel, where  $c_i \in \mathbb{R}^3$  denotes the coordinates of the voxel. In practice, we use the Fourier positional encoding proposed in (Tancik et al., 2020) due to its superiority in encoding coordinate information.

**Incorporation of Brain Parcellations** While positional encoding alone improves performance, it lacks inherent neuroscientific grounding, potentially making it challenging for the model to efficiently learn representations aligned with established principles of brain function. To overcome this, we incorporate existing brain region parcellations (Glasser et al., 2016; Rolls et al., 2020) into the *key* embeddings. Formally, given a parcellation  $\mathcal{P}$ , with regions indexed by  $1, \dots, N_{\mathcal{P}}$ . Let  $\mathcal{P}(i) \in [1, 2, \dots, N_{\mathcal{P}}]$  be the region that the  $i$ -th voxel belongs to, and  $E[\mathcal{P}(i)] \in \mathbb{R}^d$  be the corresponding learnable embedding of the region, which will be incorporated in the *key* embeddings as  $k_i^{\text{reg}, \mathcal{P}} = E[\mathcal{P}(i)] \in \mathbb{R}^d$ .

**Combining Multiple Parcellations** It is crucial to choose an appropriate brain region parcellation. Previous region-based methods (Qiu et al., 2023; Li et al., 2021; Kan et al., 2022) can usually only choose one arbitrarily. In contrast, our model design allows us to combine multiple parcellations  $\mathcal{P}^1, \mathcal{P}^2, \dots$  by concatenating their respective region encodings to the *key* embeddings. In conclusion, the final *key* embeddings are the concatenation of the positional encoding and multiple region encodings,

$$k_i = k_i^{\text{pos}} \parallel k_i^{\text{reg}, \mathcal{P}^1} \parallel k_i^{\text{reg}, \mathcal{P}^2} \parallel \dots \quad (1)$$

where  $\parallel$  denotes the concatenation operation. This process is illustrated in Figure 3’s lower right part.

The positional and region encodings complement each other: The region encodings serve as coarse-scale features, providing a neuroscientific-grounded basis, while the fine-scale positional encoding allows our model to learn finer-grained information directly from the data.

This attention design separates a voxel’s *functional information*—which is largely consistent across individuals—from its fMRI value, thereby enhancing generalization. Instead of relying on pooling or sampling, the attention mechanism employs learnable aggregation, while the integration of positional encoding and neuroscientifically informed region encodings further ensures high performance.

After the attention layer, we obtain the hidden representations  $z_q \in \mathbb{R}^{N_q}$  where  $N_q$  is the number of query embeddings. We then employ an MLP and a reshape operation to map the hidden representations to  $L$  fMRI tokens, i.e.,  $X_v = \text{reshape}(\text{MLP}(\{z_q\})) \in \mathbb{R}^{L \times d}$ .

The process of the fMRI encoder is illustrated in Figure 3. The obtained fMRI tokens are then prepended to the lan-

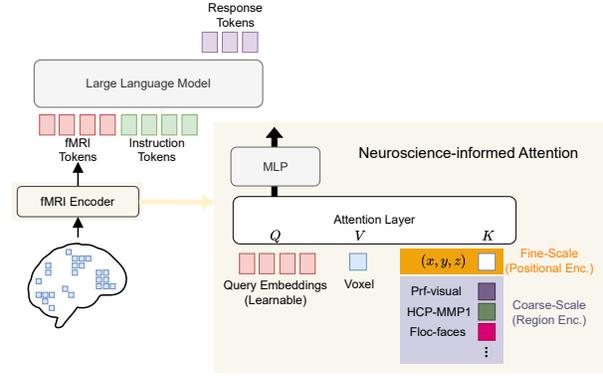


Figure 3. Model Architecture. The fMRI encoder maps fMRI to a series of fMRI tokens through our proposed neuroscience-informed attention. The large language model, with both fMRI and text tokens, will be trained by brain instruction tuning.

guage tokens in conversations.

In addition, we provide a side-by-side comparison between our encoder and existing ones in Appendix F.

### 3.3. Brain Instruction Tuning (BIT)

To enable versatile fMRI-to-text decoding, an appropriate BIT dataset is required, yet no such dataset currently exists. To bridge this gap, we construct one based on the fact: MSCOCO images (Chen et al., 2015) serve as stimuli for fMRI recordings in the fMRI study (Allen et al., 2022), and an abundance of datasets provide text annotations (e.g., VQA) for MSCOCO images. Using the images as intermediaries, we select those relevant to brain functions and pair the fMRI data with corresponding text annotations. For example, given an image of a billboard with annotated textual content, we can reasonably infer that when a subject perceives textual information (e.g., contents on the billboard), corresponding representations are encoded in the brain. This suggests the possibility of extracting such information from fMRI signals. We select datasets to fulfill various purposes, enabling the model to capture diverse aspects of semantic information embedded in fMRI signals, including visual perception & scene understanding, language & symbolic processing, memory & knowledge retrieval and complex reasoning, which are considered among most fundamental and essential properties of human brains (Robertson, 2002; Stenning & Van Lambalgen, 2012; Wade & Swanston, 2013; Friederici, 2017).

**Perception & Scene Understanding** As illustrated in Figure 4, we begin by using caption tasks at both coarse and fine-grained levels to train the model’s ability to understand and summarize what the subject perceives visually (Chen et al., 2015; Krause et al., 2017). Additionally, we incorporate QA tasks (Ren et al., 2015; Krishna et al., 2017;

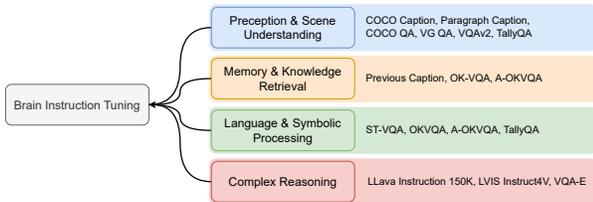


Figure 4. Dataset Taxonomy in Brain Instruction Tuning.

Acharya et al., 2019) to enhance the model’s ability to retrieve and reason about visually perceived content.

**Memory & Knowledge Retrieval** To go beyond tasks directly related to present visual perception, we construct the *previous captioning* task, a memory-oriented task that challenges the model to caption images that the subject previously viewed, simulating memory recall processes. Furthermore, we aim to encode knowledge structures in human brains. The OK-VQA (Marino et al., 2019) and A-OKVQA (Schwenk et al., 2022) datasets include questions requiring external knowledge that is not present in the image but resides in human brains. For example, A photo of a hydrant may prompt the answer “firetruck,” even though the firetruck is absent in the image. This association also reflects the way human cognition operates through a network of interconnected meanings, where one concept unconsciously triggers another. Such a process, which is called “slippage of the signifier” (Lacan, 2001; 1988; Miller & Lacan, 2018), highlights the symbolic processes through which the brain constructs and retrieves meaning.

**Language & Symbolic Processing** In addition to the aforementioned OK-VQA and A-OKVQA datasets, which are also related to symbolic process, we further combine datasets of text recognition (Biten et al., 2019) and numerical reasoning (Acharya et al., 2019) to facilitate this aspect.

**Complex Reasoning** Finally, we try to approximate the reasoning process that happens in human brains with datasets (Liu et al., 2023; Wang et al., 2023a; Li et al., 2018) that require intricate logical and inferential processes. We expect these datasets to challenge the model to extract the reasoning process, drawing upon both visual understanding and abstract problem-solving, thus bridging perception, memory, and knowledge into a cohesive cognitive framework.

We ended up with a brain instruction tuning dataset consisting of 980,610 conversations associated with fMRI recordings from 15 datasets. Appendix A lists the instructions and other details for each dataset. The instruction tuning enables versatile fMRI-to-text decoding. In particular, the introduction of tasks like *previous caption* empowers the model to perform a broader range of tasks beyond vision-related ones, which the previous model (Xia et al., 2024) fails.

To train the model with the BIT dataset, for each

sample  $v$ , we sample a multi-run conversation  $X_t = (X_u^1, X_a^1, \dots, X_u^T, X_a^T)$  from all conversations associated with it, where  $T \geq 1$  represents the number of turns.  $a$  indicates the message from the assistant and  $u$  indicates the message is from the user. The training objective is to maximize the probability of the assistant’s response only

$$\arg \max_{\theta} p(X_a | X_v, X_{\text{inst}}) = \prod_{t=1}^T p(X_a^t | X_u^{\leq t}, X_a^{\leq t}, X_{\text{inst}}, X_v)$$

Figure 5 illustrates the chat template and the training objective. We freeze the weights of the LLM and only train the fMRI encoder since we want to preserve the LLM’s language modeling prior and ensure a fair comparison with baselines such as Xia et al. (2024).

**Computational Complexity** According to the analysis in Appendix C, our model does not introduce additional complexity compared to previous methods (Scotti et al., 2024b; Wang et al., 2024a).



Figure 5. The chat template used during instruction tuning, illustrating two turns of conversations. Two turns of conversation are shown. Tokens highlighted in magenta are used for next-token prediction loss computation.

## 4. Experiments

In this section, we first evaluate our model on various downstream tasks, demonstrating its versatile decoding capabilities. Next, we assess its generalizability to novel subjects and its adaptability to real-world applications. Finally, we analyze the functions of queries in our neuroscience-informed attention mechanism.

### 4.1. Settings

**fMRI Datasets** We use the widely used Natural Scenes Dataset (NSD) (Allen et al., 2022), a large-scale dataset consisting of fMRI measurements of 8 healthy adult subjects. During data collection, subjects viewed images from the MS-COCO dataset (Lin et al., 2014) and were instructed to press buttons to indicate whether they had previously seen each image.

**Downstream Datasets** The downstream dataset will be discussed within each experiment section. See examples and a short description for all datasets we will use in Appendix A. Implementation details could be found in Appendix B.

## 4.2. Brain Captioning

To evaluate the model’s performance on downstream tasks, we start with the widely used brain captioning benchmark (Xia et al., 2024). The task, built upon COCO Caption (Chen et al., 2015) requires the model to predict captions of given images as fMRI stimuli.

**Baselines** The following baselines are considered in this experiment: SDRecon (Takagi & Nishimoto, 2023), UniBrain (Mai & Zhang, 2023), and BrainCap (Ferrante et al., 2023) employs a linear regression, mapping the fMRI to the inputs of an image caption model (Li et al., 2023). OneLLM (Han et al., 2024a) is a multimodal large language models that align 8 modalities (including fMRI) with language all in one model. For fair and efficient comparison, we only finetune the encoder, given that we freeze the LLM in our method as well. UMBRAE learns an encoder that maps fMRIs to images through an encoder similar to the MLP mixer (Tolstikhin et al., 2021). BrainChat (Huang, 2024) segments the flattened voxels into 16 patches and employs a transformer to decode text conditioned on the patches. It is worth noting that all of these baselines require subject-specific layers or parameters. In contrast, our model is subject-agnostic, thus with the potential to generalize on novel subjects.

**Metric** Following previous works, we use five standard metrics for text generation: BLEU- $k$  (Papineni et al., 2002), ROUGE-L (Lin, 2004), CIDEr (Vedantam et al., 2015), SPICE (Anderson et al., 2016), METEOR (Banerjee & Lavie, 2005).

Table 1 shows that our model outperforms baselines in terms of most metrics, with an average improvement of 3.32%, even if our model does not have any subject-specific layers. We argue that this is attributed to both the novel architecture design and the introduction of BIT, which will be evident in the next experiment.

## 4.3. Versatile Decoding

The purpose of experiments in this section is two-fold: 1) To investigate the impact of our model design and the introduction of BIT on performance improvement. 2) To evaluate

the versatility of the model, i.e., its performance on various downstream tasks.

**Baselines** Besides baselines that could be adapted to this experiment from the previous one, we further consider the following subject-agnostic models as baselines. 1) Mind-Bridge (Wang et al., 2024a) flattens the voxels and adaptively adjusts the padding and stride to pool the voxels into a fixed dimension. The original implementation of Mind-Bridge has subject-specific parameters. However, since those parameters are of the same size, we make them shared across subjects and thus make the model subject-agnostic. 2) UniBrain (Wang et al., 2024b) samples voxels into a fixed number of groups and employs a transformer where groups are treated as tokens. This UniBrain is unrelated to the UniBrain in the previous section; they just share the same name.

**Datasets & Metric** We use the test split of all QA & caption datasets in the BIT dataset. We strictly adhere to the official metrics on all datasets. In summary, for sentence generation, we use BLEU- $k$  (Papineni et al., 2002), ROUGE-L (Lin, 2004), CIDEr (Vedantam et al., 2015), SPICE (Anderson et al., 2016), METEOR (Banerjee & Lavie, 2005). For QA-related tasks, we use VQA accuracy (Antol et al., 2015) as well as special metrics proposed in the original paper (e.g. ANLS for ST-VQA (Biten et al., 2019)).

The results on subject 1 are shown in Table 2. Our model outperforms baselines, with an average improvement of 12.0%. Further, by comparing instruction tuning and from-scratch models, we find that instruction tuning has a significant positive effect, with an average improvement of 28.0%. The results remain stable across different random seeds; for instance, according to our observations, the BLEU-1 score for paragraph captioning exhibits a maximum of  $\pm 0.3$  variance. We report experiments on other subjects in Appendix G.

## 4.4. Unseen Subject Generalization

Our neuroscience-informed, subject-agnostic design enhances generalization to novel subjects, a crucial factor in real-world applications where training a model for each

Table 1. Results of brain captioning. The CIDEr metric is scaled by a factor of 100 for consistency with Table 1 and baselines.

Method	subject agnostic	BLEU-1 $\uparrow$	BLEU-2 $\uparrow$	BLEU-3 $\uparrow$	BLEU-4 $\uparrow$	METEOR $\uparrow$	ROUGE $\uparrow$	CIDEr $\uparrow$	SPICE $\uparrow$
SDRecon (Takagi & Nishimoto, 2023)	$\times$	36.21	17.11	7.22	3.43	10.03	25.13	13.83	5.02
OneLLM (Han et al., 2024a)	$\times$	47.04	26.97	15.49	9.51	13.55	35.05	22.99	6.26
UniBrain (Mai & Zhang, 2023)	$\times$	–	–	–	–	16.90	22.20	–	–
BrainCap (Ferrante et al., 2023)	$\times$	55.96	36.21	22.70	14.51	16.68	40.69	41.30	9.06
BrainChat (Huang, 2024)	$\times$	52.30	29.20	17.10	10.70	14.30	45.70	26.10	–
UMBRAE (Xia et al., 2024)	$\times$	59.44	40.48	27.66	19.03	19.45	43.71	<b>61.06</b>	<b>12.79</b>
UniBrain (Wang et al., 2024b)	$\checkmark$	59.08	39.64	26.36	17.68	17.49	43.48	48.20	9.38
MindLLM (Ours)	$\checkmark$	<b>61.75</b>	<b>42.84</b>	<b>29.86</b>	<b>21.24</b>	<b>19.54</b>	<b>45.82</b>	60.97	11.79

Table 2. Versatile decoding. A dash — means the model could not perform this task. The superscript <sup>◦</sup> means the model is trained from scratch, in contrast to their BIT version. The CIDEr metric is scaled by a factor of 100 for consistency with Table 1 and baselines.

		OneLLM	UMBRAE	BrainChat	MindBridge <sup>◦</sup>	UniBrain <sup>◦</sup>	MindLLM <sup>◦</sup>	MindBridge	UniBrain	MindLLM
subj-agnostic		✗	✗	✗	✓	✓	✓	✓	✓	✓
COCO-QA	Accuracy↑	11.09%	22.23%	39.44%	40.19%	38.38%	42.09%	<u>45.33%</u>	42.00%	<b>48.19%</b>
VG-QA	Accuracy↑	8.76%	19.67%	21.00%	20.84%	21.27%	21.68%	23.53%	<u>24.02%</u>	<b>24.06%</b>
VQA-v2	Accuracy ↑	33.68%	<u>51.23%</u>	40.02%	43.25%	46.04%	44.13%	47.91%	48.58%	<b>52.14%</b>
A-OKVQA	Accuracy ↑	25.23%	43.24%	20.52%	22.12%	19.47%	29.20%	<u>50.44%</u>	43.36%	<b>52.21%</b>
ST-VQA	ANLS ↑	5.74%	5.46%	9.58%	10.20%	7.01%	<u>12.76%</u>	11.64%	8.76%	<b>12.92%</b>
OK-VQA	Accuracy ↑	22.98%	10.35%	17.22%	27.63%	18.63%	27.70%	32.13%	<u>32.30%</u>	<b>33.33%</b>
TallyQA	Accuracy ↑	8.34%	44.10%	43.22%	43.49%	44.83%	43.75%	49.46%	<u>53.77%</u>	<b>54.76%</b>
	RMSE ↓	7.45	3.94	1.90	2.03	1.83	2.04	1.86	<b>1.67</b>	<u>1.76</u>
Paragraph Caption	BLEU-1↑	0.26	<b>29.82</b>	22.21	21.82	25.69	26.49	25.69	28.28	<u>29.43</u>
	BLEU-2↑	0.08	14.26	10.23	10.47	12.62	12.48	13.00	<u>15.47</u>	<b>15.78</b>
	BLEU-3↑	0.03	6.52	6.38	5.58	6.70	6.43	7.10	<u>8.90</u>	<b>9.14</b>
	BLEU-4↑	0.01	2.95	2.12	3.14	3.81	3.63	4.22	<b>5.60</b>	<u>5.51</u>
	METEOR ↑	2.36	12.60	9.10	10.95	11.13	10.71	11.39	<b>13.50</b>	<u>13.18</u>
	CIDEr ↑	0.00	7.39	6.02	<u>7.50</u>	3.92	2.44	3.55	1.82	<b>7.80</b>
VQA-E	Accuracy ↑	19.60%	47.84%	46.20%	45.40%	44.42%	44.55%	<u>48.48%</u>	48.39%	<b>50.95%</b>
	BLEU-1 ↑	17.32	29.83	35.99	35.63	35.30	35.08	36.18	<u>37.26</u>	<b>37.70</b>
	BLEU-2 ↑	7.44	14.76	18.33	18.27	18.04	17.82	19.38	<u>20.41</u>	<b>20.56</b>
	BLEU-3 ↑	3.62	8.17	10.01	10.32	10.20	10.05	11.30	<u>12.25</u>	<b>12.34</b>
	BLEU-4 ↑	1.82	4.87	6.60	6.27	6.14	6.00	7.00	<u>7.83</u>	<b>7.92</b>
	CIDEr ↑	19.32	63.26	78.33	79.05	77.31	76.80	86.62	<u>92.09</u>	<b>93.60</b>
	METEOR ↑	6.69	12.25	13.64	14.13	13.89	13.96	14.81	<u>15.51</u>	<b>15.62</b>
	ROUGE ↑	16.84	28.38	32.82	33.78	33.25	33.11	34.56	<u>35.87</u>	<b>35.88</b>
FSVQA	VQA Acc. ↑	31.44%	40.67%	36.30%	42.00%	37.05%	42.53%	<u>45.95%</u>	44.58%	<b>48.03%</b>
	FSVQA Acc. ↑	21.02%	0.00%	30.22%	37.40%	32.30%	38.50%	<u>40.97%</u>	37.87%	<b>43.00%</b>
	BLEU-1 ↑	37.42	23.11	83.99	85.68	83.84	85.88	<u>86.52</u>	85.10	<b>87.10</b>
	BLEU-2 ↑	31.72	5.86	78.50	81.27	78.81	81.62	<u>82.28</u>	80.01	<b>83.03</b>
	BLEU-3 ↑	26.95	2.10	73.00	77.10	73.97	77.62	<u>78.34</u>	75.49	<b>79.27</b>
	BLEU-4 ↑	22.48	1.04	69.73	72.89	68.91	73.56	<u>74.35</u>	70.73	<b>75.50</b>
	METEOR ↑	26.35	8.93	44.76	47.59	45.94	47.96	<u>48.63</u>	46.89	<b>49.05</b>
	CIDEr ↑	312.75	4.07	600.00	636.40	609.00	646.26	<u>657.02</u>	628.83	<b>666.26</b>
Previous Caption	BLEU-1 ↑	41.86	—	21.19	21.17	24.84	<u>44.52</u>	42.45	43.01	<b>47.20</b>
	BLEU-2 ↑	19.44	—	8.00	7.57	9.70	<u>22.46</u>	20.04	20.03	<b>25.16</b>
	BLEU-3 ↑	9.25	—	1.98	2.85	3.40	<u>10.39</u>	9.61	9.19	<b>12.95</b>
	BLEU-4 ↑	3.67	—	1.02	1.28	1.46	<u>5.45</u>	5.31	4.58	<b>7.49</b>
	METEOR ↑	10.14	—	6.55	6.46	7.20	<u>11.00</u>	10.83	10.81	<b>11.96</b>
	ROUGE ↑	30.19	—	21.23	20.88	23.04	<u>33.20</u>	32.38	31.99	<b>34.58</b>
	CIDEr ↑	6.65	—	9.21	8.83	<u>11.73</u>	9.39	7.89	7.53	<b>16.02</b>
	SPICE ↑	2.49	—	2.44	2.56	2.78	<u>3.07</u>	2.80	2.92	<b>3.93</b>

individual is impractical. To evaluate it, we perform instruction tuning on 7 out of the 8 subjects in the natural scene dataset (Allen et al., 2022), and evaluate generalization on the held-out subject. Table 3 shows our model outperforms two other subject-agnostic baselines in most cases, with an average improvement of 24.5% compared to the second-best model. We anticipate further performance gains through the application of domain adaptation techniques (Xiao et al., 2023; Ganin & Lempitsky, 2015; Ganin et al., 2016; Gong et al., 2012), which we leave as a direction for future work.

#### 4.5. Adapting to New Tasks

It is common that users want to adapt the MindLLM to their own specific use cases. To this end, we aim to assess our model’s adaptability to new tasks.

**Dataset & Metrics** We use TDIUC (Kafle & Kanan, 2017), a QA dataset consisting of 12 types of questions, as a benchmark to evaluate the model’s various capabilities comprehensively. Additionally, we further select 2 task types—*sentiment understanding* and *utility/affordance* tasks, that are particularly relevant to BCI applications as sub-datasets. The *utility/affordance* task, for instance, enables the model to identify useful objects in a given scene and autonomously decide whether to utilize them. Following their paper, we compute the accuracy of each type and report the arithmetic mean-per-type (A-MPT) and the harmonic mean-per-type (H-MPT). For the 2 selected types, we report the accuracy respectively.

Table 4 shows our model achieves balanced (high harmonic mean) and consistently improved performances with an

Table 3. Model generalization, compared with subject-agnostic model. We train the models on subject 1 – 7 and evaluate on subject 8, which is the held-out subject.

		MindBridge	UniBrain	MindLLM
COCO-QA	Accuracy ↑	35.88	24.95	<b>38.96</b>
VG-QA	Accuracy ↑	<b>20.56</b>	16.23	19.29
VQA-v2	Accuracy ↑	42.80	40.16	<b>46.04</b>
A-OKVQA	Accuracy ↑	44.55	28.71	<b>45.54</b>
ST-VQA	ANLS ↑	9.33	9.30	<b>13.63</b>
OK-VQA	Accuracy ↑	21.94	17.09	<b>25.44</b>
TallyQA	Accuracy ↑	38.92	32.51	<b>42.29</b>
	RMSE ↓	2.12	<b>2.02</b>	2.11
COCO-Caption	BLEU-1 ↑	39.84	41.90	<b>44.60</b>
	BLEU-2 ↑	19.55	19.67	<b>24.04</b>
	BLEU-3 ↑	9.29	8.89	<b>12.79</b>
	BLEU-4 ↑	5.24	4.33	<b>7.52</b>
	METEOR ↑	10.39	10.80	<b>11.16</b>
	ROUGE ↑	31.10	31.54	<b>33.47</b>
	CIDEr ↑	8.70	6.40	<b>13.22</b>
	SPICE ↑	2.67	2.39	<b>3.82</b>
Paragraph Caption	BLEU-1 ↑	23.18	21.73	<b>26.07</b>
	BLEU-2 ↑	10.71	8.94	<b>12.64</b>
	BLEU-3 ↑	4.61	3.72	<b>6.27</b>
	BLEU-4 ↑	2.22	1.92	<b>3.46</b>
	METEOR ↑	9.99	9.47	<b>11.18</b>
CIDEr ↑	0.71	1.56	<b>5.61</b>	
VQA-E	Accuracy ↑	41.78	38.53	<b>46.19</b>
	BLEU-1 ↑	32.54	32.86	<b>34.54</b>
	BLEU-2 ↑	16.13	15.48	<b>17.54</b>
	BLEU-3 ↑	8.82	7.98	<b>9.84</b>
	BLEU-4 ↑	5.16	4.42	<b>5.81</b>
	CIDEr ↑	68.13	58.79	<b>73.27</b>
	ROUGE ↑	12.74	12.26	<b>13.56</b>
FSVQA	VQA Acc. ↑	42.33	37.92	<b>43.92</b>
	FSVQA Acc. ↑	37.16	30.83	<b>39.07</b>
	BLEU-1 ↑	75.81	82.94	<b>86.19</b>
	BLEU-2 ↑	71.03	77.24	<b>81.83</b>
	BLEU-3 ↑	66.40	71.94	<b>77.77</b>
	BLEU-4 ↑	61.55	66.54	<b>73.58</b>
	METEOR ↑	45.84	45.24	<b>47.94</b>
CIDEr ↑	428.39	587.78	<b>646.30</b>	

average of 13.5%. We could also observe the performance benefits from BIT, with 25.0% absolute improvement.

#### 4.6. Performance scale with the number of subjects

We conducted experiments to evaluate how model performance scales with the number of subjects and report performances on the COCO caption task. We examined both the in-distribution (seen subjects) setting in Figure 5 and the out-of-distribution (held-out subjects) setting in Figure 6. Our results show significant performance improvements as the number of training subjects increases, demonstrating that the model benefits from exposure to more subjects during pre-training.

#### 4.7. Ablation Study

We conduct ablation studies on the design of key embeddings in the neuroscience-informed attention module in Fig-

Table 4. Model adaptation to new tasks. *sentiment understanding* and *utility/affordance* are sub-datasets from TDIUC that are particularly relevant to BCI applications.

Method	Overall		Sentiment Understanding	Utility/Affordance
	A-MPT	H-MPT	Accuracy	Accuracy
MindLLM <sup>o</sup>	41.09%	19.38%	70.00%	0.00%
MindBridge	49.77%	39.88%	80.00%	14.29%
UniBrain	51.50%	36.76%	80.00%	28.57%
MindLLM	<b>54.08%</b>	<b>45.43%</b>	<b>80.77%</b>	<b>50.00%</b>

Table 5. Performance scaling with number of subjects (in-distribution). Evaluation is performed on Subject 1, who is included during training.

number of subjects	COCO Caption							
	BLEU-1 ↑	BLEU-2 ↑	BLEU-3 ↑	BLEU-4 ↑	METEOR ↑	ROUGE ↑	CIDEr ↑	SPICE ↑
1	58.05	37.95	24.40	16.14	16.62	42.03	43.04	8.81
3	56.81	37.16	24.12	16.30	16.44	41.93	40.94	9.12
5	58.17	38.65	25.49	17.30	16.97	42.36	46.54	9.19
7	58.52	39.20	25.95	17.51	17.12	42.87	47.45	9.48

Table 6. Performance scaling with number of subjects (out-of-distribution). Evaluation is performed on Subject 8, who is held out during training.

number of subjects	BLEU-1 ↑	BLEU-2 ↑	BLEU-3 ↑	BLEU-4 ↑	METEOR ↑	ROUGE ↑	CIDEr ↑	SPICE ↑
	1	42.40	19.51	8.61	4.19	9.48	30.68	5.41
3	43.95	22.45	10.37	5.53	10.37	32.49	5.79	2.59
5	45.00	22.61	10.64	5.60	10.53	32.53	6.22	2.79
7	47.30	25.35	13.61	8.15	11.40	34.64	6.41	3.61

ure 6. The results strongly validate our design. The vanilla cross attention (*Pos Enc. +fMRI*) leads to poor performance, while removing fMRI values from the key embeddings (*Pos Enc.*) yields a significant improvement. Replacing positional encoding with region encodings (*Reg. Enc.*) accelerates convergence in the early stages since it is grounded by neuroscientific principles. However, it is eventually outperformed by *Pos Enc.* due to the lack of finer-grained information. Combining the positional encoding and region encodings (*Pos Enc. +Reg Enc.*) achieves the best model design. In addition, replacing positional encoding with an MLP that maps coordinates to embeddings results in poor performance (*(x,y,z)+MLP*), which indicates the amount of high-frequency spatial information in fMRI signals. Additional ablation results on evaluation metrics are provided in Appendix E.

#### 4.8. Visualizations and Interpretations

Unlike previous deep learning models (Scotti et al., 2024b; Mai & Zhang, 2023), our model allows interpretations by investigating how *queries* work in the neuroscience-informed attention layer. We inspect the attention weights between queries and voxels in Figure 7.

We found that some queries primarily focus on processing

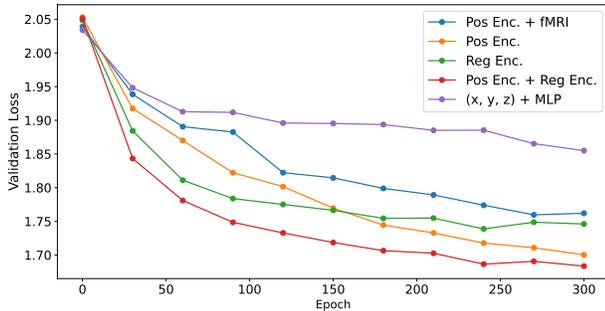


Figure 6. Ablation study of the key embedding design. Pos Enc. stands for positional encoding. Reg Enc. stands for multiple region encodings.  $(x, y, z)+MLP$  means we employ an MLP to map the coordinates to the embeddings instead of positional encoding.

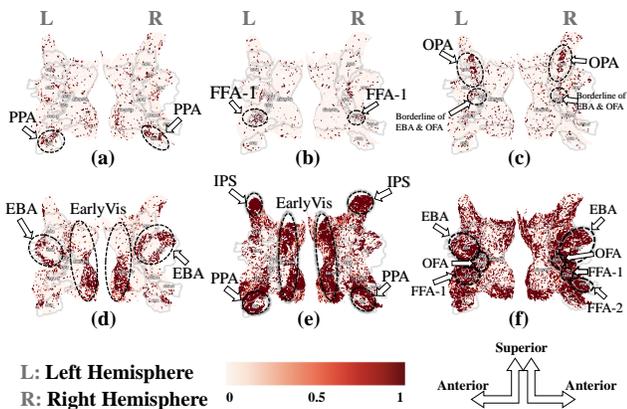


Figure 7. Each subfigure corresponds to a specific query of subject 1. We visualize the attention map between that query token and all voxels. Here we randomly select 6 query tokens, each of which exhibits a distinct spatial focus.

single brain regions, such as Parahippocampal Place Area (PPA) (Figure 7a) and Fusiform Face Area (FFA) (7b). As previous research has shown, PPA is related to conceptual association, semantic processing and environmental memory (Epstein et al., 1999; Köhler et al., 2002; Bar et al., 2008; Epstein & Ward, 2010) and FFA is known for its critical role in expertise recognition, social cognition and identity memory (Schultz et al., 2003; Tsantani et al., 2021; Xu, 2005). Both are important brain regions for the conceptualization of visual information and are responsible for the interaction between real-time stimulus and past memory (Brewer et al., 1998; Ranganath et al., 2004; Golarai et al., 2007).

Moreover, there are some queries that attend to multiple brain regions, revealing the information transmission between low- and high-level brain regions. For instance, interactions between early visual areas and higher-level regions like PPA and IntraParietal Sulcus (IPS) (Figure 7e), revealing a potential pattern for human attention-guided actions

(Tunik et al., 2007; Connolly et al., 2016). Additionally, queries are also found responsible for communications between high-level brain regions (Figure 7c,7f). Together, these findings indicate that the learnable queries may reflect the dynamics of human brain activities in the visual task, from seeing and thinking about the image to pressing the button for the visual recall task in NSD (Allen et al., 2022).

We provide attention map visualizations of subjects other than subject 1 in Appendix I. We also visualize latent embeddings during the model forward pass in Appendix H. Qualitative analysis of model responses could be found in Appendix D.

## 5. Conclusion

In this work, we propose MindLLM, a subject-agnostic and versatile fMRI-to-text decoding model. Our neuroscience-informed attention mechanism in the fMRI encoder ensures subject generalization, while brain instruction tuning enhances versatility. Comprehensive benchmarks demonstrate our model’s state-of-the-art performance, and visualizations offer insights into the patterns it leverages. We envision that these advancements will contribute to medical treatments, neuroscience research, and brain-computer interfaces in the future. **Limitations:** This work focuses on static fMRI, without incorporating temporal dynamics. Future research could explore integrating temporal information and investigating its relationship with other modalities such as videos.

## Acknowledgments

We are grateful for funding and computing resource support from Amazon Research, Snap Research, National Science Foundation (NSF) IIS Div Of Information & Intelligent Systems 2403317, Yale AI Engineering Research Seed Grants, Yale Office of the Provost, and Yale Wu Tsai Institute.

## Impact Statement

This paper presents work whose goal is to advance the field of Machine Learning. There are many potential societal consequences of our work, none of which we feel must be specifically highlighted here.

## References

- Acharya, M., Kaffe, K., and Kanan, C. Tallyqa: Answering complex counting questions. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pp. 8076–8084, 2019.
- Alayrac, J.-B., Donahue, J., Luc, P., Miech, A., Barr, I., Hasson, Y., Lenc, K., Mensch, A., Millican, K., Reynolds, M., et al. Flamingo: a visual language model for few-

- shot learning. *Advances in neural information processing systems*, 35:23716–23736, 2022.
- Allen, E. J., St-Yves, G., Wu, Y., Breedlove, J. L., Prince, J. S., Dowdle, L. T., Nau, M., Caron, B., Pestilli, F., Charest, I., et al. A massive 7t fMRI dataset to bridge cognitive neuroscience and artificial intelligence. *Nature neuroscience*, 25(1):116–126, 2022.
- Anderson, P., Fernando, B., Johnson, M., and Gould, S. Spice: Semantic propositional image caption evaluation. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part V 14*, pp. 382–398. Springer, 2016.
- Antol, S., Agrawal, A., Lu, J., Mitchell, M., Batra, D., Zitnick, C. L., and Parikh, D. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*, pp. 2425–2433, 2015.
- Banerjee, S. and Lavie, A. Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In *Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, pp. 65–72, 2005.
- Bar, M., Aminoff, E., and Schacter, D. L. Scenes unseen: the parahippocampal cortex intrinsically subserves contextual associations, not scenes or places per se. *Journal of Neuroscience*, 28(34):8539–8544, 2008.
- Bernal, S. L., Pérez, M. Q., Beltrán, E. T. M., Pérez, G. M., and Celdrán, A. H. When brain-computer interfaces meet the metaverse: Landscape, demonstrator, trends, challenges, and concerns. *arXiv preprint arXiv:2212.03169*, 2022.
- Biten, A. F., Tito, R., Mafla, A., Gomez, L., Rusinol, M., Valveny, E., Jawahar, C., and Karatzas, D. Scene text visual question answering. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 4291–4301, 2019.
- Brewer, J. B., Zhao, Z., Desmond, J. E., Glover, G. H., and Gabrieli, J. D. Making memories: brain activity that predicts how well visual experience will be remembered. *Science*, 281(5380):1185–1187, 1998.
- Card, N. S., Wairagkar, M., Iacobacci, C., Hou, X., Singer-Clark, T., Willett, F. R., Kunz, E. M., Fan, C., Vahdati Nia, M., Deo, D. R., et al. An accurate and rapidly calibrating speech neuroprosthesis. *New England Journal of Medicine*, 391(7):609–618, 2024.
- Chen, J., Qi, Y., Wang, Y., and Pan, G. Mindgpt: Interpreting what you see with non-invasive brain recordings. *arXiv preprint arXiv:2309.15729*, 2023a.
- Chen, K., Zhang, Z., Zeng, W., Zhang, R., Zhu, F., and Zhao, R. Shikra: Unleashing multimodal llm’s referential dialogue magic. *arXiv preprint arXiv:2306.15195*, 2023b.
- Chen, X., Fang, H., Lin, T.-Y., Vedantam, R., Gupta, S., Dollár, P., and Zitnick, C. L. Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*, 2015.
- Chen, X., Du, C., Liu, C., Wang, Y., and He, H. Open-vocabulary auditory neural decoding using fmri-prompted llm. *arXiv preprint arXiv:2405.07840*, 2024.
- Chen, Z., Qing, J., Xiang, T., Yue, W. L., and Zhou, J. H. Seeing beyond the brain: Conditional diffusion model with sparse masked modeling for vision decoding, 2023c. URL <https://arxiv.org/abs/2211.06956>.
- Chen, Z., Qing, J., and Zhou, J. H. Cinematic mindscapes: High-quality video reconstruction from brain activity, 2023d. URL <https://arxiv.org/abs/2305.11675>.
- Cheng, Z., Leng, S., Zhang, H., Xin, Y., Li, X., Chen, G., Zhu, Y., Zhang, W., Luo, Z., Zhao, D., et al. Videollama 2: Advancing spatial-temporal modeling and audio understanding in video-llms. *arXiv preprint arXiv:2406.07476*, 2024.
- Connolly, J. D., Kentrige, R. W., and Cavina-Pratesi, C. Coding of attention across the human intraparietal sulcus. *Experimental brain research*, 234:917–930, 2016.
- Denk, T. I., Takagi, Y., Matsuyama, T., Agostinelli, A., Nakai, T., Frank, C., and Nishimoto, S. Brain2music: Reconstructing music from human brain activity, 2023. URL <https://arxiv.org/abs/2307.11078>.
- Devlin, J. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- Du, B., Cheng, X., Duan, Y., and Ning, H. fMRI brain decoding and its applications in brain-computer interface: A survey. *Brain Sciences*, 12(2):228, 2022.
- Epstein, R., Harris, A., Stanley, D., and Kanwisher, N. The parahippocampal place area: recognition, navigation, or encoding? *Neuron*, 23(1):115–125, 1999.
- Epstein, R. A. and Ward, E. J. How reliable are visual context effects in the parahippocampal place area? *Cerebral Cortex*, 20(2):294–303, 2010.
- Ferrante, M., Ozcelik, F., Boccato, T., VanRullen, R., and Toschi, N. Brain captioning: Decoding human brain activity into images and text. *arXiv preprint arXiv:2305.11560*, 2023.

- Friederici, A. D. *Language in our brain: The origins of a uniquely human capacity*. MIT Press, 2017.
- Ganin, Y. and Lempitsky, V. Unsupervised domain adaptation by backpropagation. In *International conference on machine learning*, pp. 1180–1189. PMLR, 2015.
- Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., March, M., and Lempitsky, V. Domain-adversarial training of neural networks. *Journal of machine learning research*, 17(59):1–35, 2016.
- Glasser, M. F., Coalson, T. S., Robinson, E. C., Hacker, C. D., Harwell, J., Yacoub, E., Ugurbil, K., Andersson, J., Beckmann, C. F., Jenkinson, M., et al. A multi-modal parcellation of human cerebral cortex. *Nature*, 536(7615): 171–178, 2016.
- Golarai, G., Ghahremani, D. G., Whitfield-Gabrieli, S., Reiss, A., Eberhardt, J. L., Gabrieli, J. D., and Grill-Spector, K. Differential development of high-level visual cortex correlates with category-specific recognition memory. *Nature neuroscience*, 10(4):512–522, 2007.
- Gong, B., Shi, Y., Sha, F., and Grauman, K. Geodesic flow kernel for unsupervised domain adaptation. In *2012 IEEE conference on computer vision and pattern recognition*, pp. 2066–2073. IEEE, 2012.
- Goyal, Y., Khot, T., Summers-Stay, D., Batra, D., and Parikh, D. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 6904–6913, 2017.
- Han, J., Gong, K., Zhang, Y., Wang, J., Zhang, K., Lin, D., Qiao, Y., Gao, P., and Yue, X. Onellm: One framework to align all modalities with language. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 26584–26595, 2024a.
- Han, X., Wang, Y., Zhai, B., You, Q., and Yang, H. Coco is” all”you need for visual instruction fine-tuning. *arXiv preprint arXiv:2401.08968*, 2024b.
- He, K., Chen, X., Xie, S., Li, Y., Dollár, P., and Girshick, R. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 16000–16009, 2022.
- Hmamouche, Y., Chihab, I., Kdouri, L., and Seghrouchni, A. E. F. A multimodal llm for the non-invasive decoding of spoken text from brain recordings. *arXiv preprint arXiv:2409.19710*, 2024.
- Huang, W. Brainchat: Decoding semantic information from fMRI using vision-language pretrained models. *arXiv preprint arXiv:2406.07584*, 2024.
- Kafle, K. and Kanan, C. An analysis of visual question answering algorithms. In *ICCV*, 2017.
- Kan, X., Dai, W., Cui, H., Zhang, Z., Guo, Y., and Yang, C. Brain network transformer. *Advances in Neural Information Processing Systems*, 35:25586–25599, 2022.
- Köhler, S., Crane, J., and Milner, B. Differential contributions of the parahippocampal place area and the anterior hippocampus to human memory for scenes. *Hippocampus*, 12(6):718–723, 2002.
- Kondratyuk, D., Yu, L., Gu, X., Lezama, J., Huang, J., Schindler, G., Hornung, R., Birodkar, V., Yan, J., Chiu, M.-C., et al. Videopoet: A large language model for zero-shot video generation. *arXiv preprint arXiv:2312.14125*, 2023.
- Krause, J., Johnson, J., Krishna, R., and Fei-Fei, L. A hierarchical approach for generating descriptive image paragraphs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 317–325, 2017.
- Krishna, R., Zhu, Y., Groth, O., Johnson, J., Hata, K., Kravitz, J., Chen, S., Kalantidis, Y., Li, L.-J., Shamma, D. A., et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International journal of computer vision*, 123:32–73, 2017.
- Lacan, J. *The Seminar of Jacques Lacan*. WW Norton & Company, 1988.
- Lacan, J. *Ecrits: A selection*. Routledge, 2001.
- Li, J., Li, D., Savarese, S., and Hoi, S. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, pp. 19730–19742. PMLR, 2023.
- Li, Q., Tao, Q., Joty, S., Cai, J., and Luo, J. Vqa-e: Explaining, elaborating, and enhancing your answers for visual questions. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 552–567, 2018.
- Li, X., Zhou, Y., Dvornek, N., Zhang, M., Gao, S., Zhuang, J., Scheinost, D., Staib, L. H., Ventola, P., and Duncan, J. S. Braingnn: Interpretable brain graph neural network for fMRI analysis. *Medical Image Analysis*, 74:102233, 2021.
- Lin, C.-Y. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pp. 74–81, 2004.
- Lin, T.-Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., and Zitnick, C. L. Microsoft coco: Common objects in context. In *Computer Vision–ECCV*

- 2014: *13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V 13*, pp. 740–755. Springer, 2014.
- Liu, C., Jing, J., Jiang, J., Wen, W., Zhu, W., Li, Z., Pan, Y., Cai, X., Liu, H., Zhou, Y., et al. Relationships between brain structure-function coupling in normal aging and cognition: A cross-ethnicity population-based study. *NeuroImage*, 299:120847, 2024.
- Liu, F., Emerson, G., and Collier, N. Visual spatial reasoning. *Transactions of the Association for Computational Linguistics*, 11:635–651, 2023.
- Luo, A. F., Henderson, M. M., Tarr, M. J., and Wehbe, L. Brainscuba: Fine-grained natural language captions of visual cortex selectivity. *arXiv preprint arXiv:2310.04420*, 2023.
- Mai, W. and Zhang, Z. Unibrain: Unify image reconstruction and captioning all in one diffusion model from human brain activity. *arXiv preprint arXiv:2308.07428*, 2023.
- Marino, K., Rastegari, M., Farhadi, A., and Mottaghi, R. Ok-vqa: A visual question answering benchmark requiring external knowledge. In *Proceedings of the IEEE/cvf conference on computer vision and pattern recognition*, pp. 3195–3204, 2019.
- McManus, I. C. *Right hand, left hand: The origins of asymmetry in brains, bodies, atoms, and cultures*. Harvard University Press, 2002.
- Miller, J. A. and Lacan, J. *The four fundamental concepts of psycho-analysis*. Routledge, 2018.
- Murahari, V., Chattopadhyay, P., Batra, D., Parikh, D., and Das, A. Improving generative visual dialog by answering diverse questions. *arXiv preprint arXiv:1909.10470*, 2019.
- Papineni, K., Roukos, S., Ward, T., and Zhu, W.-J. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pp. 311–318, 2002.
- Qi, Z., Dong, R., Zhang, S., Geng, H., Han, C., Ge, Z., Yi, L., and Ma, K. Shapellm: Universal 3d object understanding for embodied interaction. In *European Conference on Computer Vision*, pp. 214–238. Springer, 2025.
- Qiu, W., Chu, H., Wang, S., Zuo, H., Li, X., Zhao, Y., and Ying, R. Learning high-order relationships of brain regions. In *Forty-first International Conference on Machine Learning*, 2023.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pp. 8748–8763. PMLR, 2021.
- Ranganath, C., DeGutis, J., and D’Esposito, M. Category-specific modulation of inferior temporal activity during working memory encoding and maintenance. *Cognitive Brain Research*, 20(1):37–45, 2004.
- Ren, M., Kiros, R., and Zemel, R. Exploring models and data for image question answering. *Advances in neural information processing systems*, 28, 2015.
- Robertson, L. T. Memory and the brain. *Journal of dental education*, 66(1):30–42, 2002.
- Rolls, E. T., Huang, C.-C., Lin, C.-P., Feng, J., and Joliot, M. Automated anatomical labelling atlas 3. *Neuroimage*, 206:116189, 2020.
- Schultz, R. T., Grelotti, D. J., Klin, A., Kleinman, J., Van der Gaag, C., Marois, R., and Skudlarski, P. The role of the fusiform face area in social cognition: implications for the pathobiology of autism. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 358(1430):415–427, 2003.
- Schwenk, D., Khandelwal, A., Clark, C., Marino, K., and Mottaghi, R. A-okvqa: A benchmark for visual question answering using world knowledge. In *European conference on computer vision*, pp. 146–162. Springer, 2022.
- Scotti, P., Banerjee, A., Goode, J., Shabalin, S., Nguyen, A., Dempster, A., Verlinde, N., Yundler, E., Weisberg, D., Norman, K., et al. Reconstructing the mind’s eye: fMRI-to-image with contrastive learning and diffusion priors. *Advances in Neural Information Processing Systems*, 36, 2024a.
- Scotti, P. S., Tripathy, M., Villanueva, C. K. T., Kneeland, R., Chen, T., Narang, A., Santhirasegaran, C., Xu, J., Naselaris, T., Norman, K. A., et al. Mindeye2: Shared-subject models enable fMRI-to-image with 1 hour of data. *arXiv preprint arXiv:2403.11207*, 2024b.
- Shin, A., Ushiku, Y., and Harada, T. The color of the cat is gray: 1 million full-sentences visual question answering (fsvqa). *arXiv preprint arXiv:1609.06657*, 2016.
- Stenning, K. and Van Lambalgen, M. *Human reasoning and cognitive science*. MIT Press, 2012.

- Sun, J., Li, M., Chen, Z., and Moens, M.-F. Neurocine: Decoding vivid video sequences from human brain activities, 2024. URL <https://arxiv.org/abs/2402.01590>.
- Takagi, Y. and Nishimoto, S. High-resolution image reconstruction with latent diffusion models from human brain activity. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14453–14463, 2023.
- Tancik, M., Srinivasan, P., Mildenhall, B., Fridovich-Keil, S., Raghavan, N., Singhal, U., Ramamoorthi, R., Barron, J., and Ng, R. Fourier features let networks learn high frequency functions in low dimensional domains. *Advances in neural information processing systems*, 33:7537–7547, 2020.
- Tolstikhin, I. O., Houlsby, N., Kolesnikov, A., Beyer, L., Zhai, X., Unterthiner, T., Yung, J., Steiner, A., Keysers, D., Uszkoreit, J., et al. Mlp-mixer: An all-mlp architecture for vision. *Advances in neural information processing systems*, 34:24261–24272, 2021.
- Toneva, M. and Wehbe, L. Interpreting and improving natural-language processing (in machines) with natural language-processing (in the brain). *Advances in neural information processing systems*, 32, 2019.
- Tsantani, M., Kriegeskorte, N., Storrs, K., Williams, A. L., McGettigan, C., and Garrido, L. Ffa and ofa encode distinct types of face identity information. *Journal of Neuroscience*, 41(9):1952–1969, 2021.
- Tunik, E., Rice, N. J., Hamilton, A., and Grafton, S. T. Beyond grasping: representation of action in human anterior intraparietal sulcus. *Neuroimage*, 36:T77–T86, 2007.
- Vaswani, A. Attention is all you need. *Advances in Neural Information Processing Systems*, 2017.
- Vedantam, R., Lawrence Zitnick, C., and Parikh, D. Cider: Consensus-based image description evaluation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4566–4575, 2015.
- Wade, N. and Swanston, M. *Visual perception: An introduction*. Psychology Press, 2013.
- Wang, J., Yang, Z., Hu, X., Li, L., Lin, K., Gan, Z., Liu, Z., Liu, C., and Wang, L. Git: A generative image-to-text transformer for vision and language. *arXiv preprint arXiv:2205.14100*, 2022.
- Wang, J., Meng, L., Weng, Z., He, B., Wu, Z., and Jiang, Y.-G. To see is to believe: Prompting gpt-4v for better visual instruction tuning. *arXiv preprint arXiv:2311.07574*, 2023a.
- Wang, S., Liu, S., Tan, Z., and Wang, X. Mindbridge: A cross-subject brain decoding framework. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11333–11342, 2024a.
- Wang, W., Lv, Q., Yu, W., Hong, W., Qi, J., Wang, Y., Ji, J., Yang, Z., Zhao, L., Song, X., et al. Cogvlm: Visual expert for pretrained language models. *arXiv preprint arXiv:2311.03079*, 2023b.
- Wang, Z., Zhao, Z., Zhou, L., and Nachev, P. Unibrain: A unified model for cross-subject brain decoding. *arXiv preprint arXiv:2412.19487*, 2024b.
- Xia, W., de Charette, R., Öztireli, C., and Xue, J.-H. Umbrae: Unified multimodal decoding of brain signals. *arXiv preprint arXiv:2404.07202*, 2024.
- Xiao, Z., Wang, H., Jin, Y., Feng, L., Chen, G., Huang, F., and Zhao, J. Spa: A graph spectral alignment perspective for domain adaptation. *Advances in Neural Information Processing Systems*, 36:37252–37272, 2023.
- Xu, R., Wang, X., Wang, T., Chen, Y., Pang, J., and Lin, D. Pointllm: Empowering large language models to understand point clouds. *arXiv preprint arXiv:2308.16911*, 2023.
- Xu, Y. Revisiting the role of the fusiform face area in visual expertise. *Cerebral Cortex*, 15(8):1234–1242, 2005.
- Yu, J., Wang, Z., Vasudevan, V., Yeung, L., Seyedhosseini, M., and Wu, Y. Coca: Contrastive captioners are image-text foundation models. *arXiv preprint arXiv:2205.01917*, 2022.
- Zhang, H., Li, X., and Bing, L. Video-llama: An instruction-tuned audio-visual language model for video understanding. *arXiv preprint arXiv:2306.02858*, 2023a.
- Zhang, P., Dong, X., Wang, B., Cao, Y., Xu, C., Ouyang, L., Zhao, Z., Duan, H., Zhang, S., Ding, S., et al. Internlm-xcomposer: A vision-language large model for advanced text-image comprehension and composition. *arXiv preprint arXiv:2309.15112*, 2023b.
- Zhang, X., Liang, C., Wang, N., Wang, Y., Gao, Y., Sui, C., Xin, H., Feng, M., Guo, L., and Wen, H. Abnormal whole-brain voxelwise structure-function coupling and its association with cognitive dysfunction in patients with different cerebral small vessel disease burdens. *Frontiers in Aging Neuroscience*, 15:1148738, 2023c.
- Zheng, L., Chiang, W.-L., Sheng, Y., Zhuang, S., Wu, Z., Zhuang, Y., Lin, Z., Li, Z., Li, D., Xing, E., et al. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in Neural Information Processing Systems*, 36:46595–46623, 2023.

Zhu, X., Su, W., Lu, L., Li, B., Wang, X., and Dai, J. Deformable detr: Deformable transformers for end-to-end object detection. *arXiv preprint arXiv:2010.04159*, 2020.

## A. Dataset Details

### A.1. Details of each dataset in brain instruction tuning

In this section, we give a brief description of each source of our brain instruction datasets as well as examples from them.

Table 7: Dataset details and examples.

Dataset	Description	Example	
Previous Caption	Generating a one-sentence caption of the image that the subject previously saw.		A neat bedroom pairs modern chairs with a glass table.
COCO Caption (Chen et al., 2015)	Generate a one-sentence caption of the image the subject currently sees.		The pedestrian is walking down the side of the highway by the bus.
Image Paragraph Captioning (Krause et al., 2017)	Generate a one-paragraph caption of the image the subject currently sees		An elephant with a harness, and a seat on his back is seen in a dirt field on a sunny day. The seat on the back of the elephant is brown and is tied with ropes. The sun is shining through the trees onto the ground below. Behind the elephant there is a cinder block wall with grass growing in front of the wall. There are trees behind the wall.
COCO QA (Ren et al., 2015)	Answer questions according to the image.		[Q]: what eats leaves from a basket in an enclosure? [A]: Giraffe.
Visual Genome QA (Krishna et al., 2017)	Answer image-based questions that require richer semantic understanding of the image than COCO-QA.		[Q]: Where was the photo taken? [A]: In an office.
VQAv2 (Goyal et al., 2017)	Answer image-based questions with better equality and diversity than COCO-QA.		[Q]: What are the two white letters? [A]: hu
OK-VQA (Marino et al., 2019)	Answer image-based questions that require external knowledge beyond the image itself.		[Q]: What part of the body do you wear the rightmost objects on? [A]: Neck.
ST-VQA (Biten et al., 2019)	Answer questions of high-level semantic information present in images as the textual cue		[Q]: What is written on the front of the shirt? [A]: Marine.
TallyQA (Acharya et al., 2019)	Count objects present in images.		[Q]: How many bats on the wall? [A]: 2
VQA-E (Li et al., 2018)	Answer questions and generate corresponding explanations for an image-based question.		[Q]: Are the people going for a walk in the forest? [A]: Yes. Here is the explanation: A picture of the land, trees, and people passing by as they ride in a vehicle.
A-OKVQA (Schwenk et al., 2022)	Answer multiple-choice questions.		[Q]: What season is up next? Multiple Choices: A. autumn B. spring C. summer D. winter [A]: B

Continued on next page

## MindLLM: A Subject-Agnostic and Versatile Model for fMRI-to-Text Decoding

Dataset	Description	Example	
FSVQA (Shin et al., 2016)	Answer the questions in full sentences.		[Q]: Is this an area that is more populated with pedestrians than cars? [A]: Yes, this is an area that is more populated with pedestrians than cars.
VisDial (Murahari et al., 2019)	Generate answers for image-based questions in a multi-turn dialogue.		[Q1]: Is the photo in color? [A1]: No. [Q2]: Is the man wearing glasses [A2]: Can't tell. [Q3]: How many horses are there? [A3]: 2
LLava Instruction 150K (Liu et al., 2023)	Generate answers for object-level and scene-level answering or reasoning questions for single or multi-round conversations.		[Q1]: What is the main activity of the person in the fMRI? [A1]: The main activity of the person in the fMRI is brushing her teeth. [Q2]: What should be considered when maintaining oral hygiene in a public setting? [A2]: When maintaining oral hygiene in a public setting (as depicted in the fMRI by a woman brushing her teeth with a green toothbrush while carrying a shoulder bag), factors such as cleanliness, privacy, time management, and personal hygiene supplies should be taken into account...
LVIS Instruct4V (Wang et al., 2023a)	Similar to LLava Instruct4V but is finer-grained and of higher quality.		[Q1]: What is the primary object in the fMRI? [A1]: The primary object in the fMRI is a bird. [Q2]: Can you describe the color and appearance of the bird? [A2]: The bird has a deep brown body with a darker head. Its eyes are distinctively surrounded by a yellowish hue. The bird also possesses a blue-tinted beak. [Q3]: How many birds are in the picture? [A3]: There is only one bird in the picture...
TDIUC (only used in downstream) (Kafle & Kanan, 2017)	Generate answers for image-based questions from 12 different task-directed question types such as object presence and utility/affordance.		[Q]: What object can be thrown? [A]: Baseball

### A.2. Dataset statistics

Table 8 summarizes the statistics of each subject in the natural scene dataset (Allen et al., 2022).

Table 8. Number of voxels for each subject.

subject	1	2	3	4	5	6	7	8
#(input voxels)	15724	14278	15226	13153	13039	17907	126682	14386
#(samples)	30000	30000	24000	22500	30000	24000	30000	22500

### A.3. Instructions

Table 9 lists instructions for each dataset (i.e.  $X_{inst}$ ). Some instructions are inspired by (Han et al., 2024b). For question-answering tasks, the instructions are appended after the question in a new line.

## B. Implementation details

We choose Vicuna-7b (Zheng et al., 2023) as our backbone LLM. During the brain instruction tuning stage, We use AdamW as the optimizer, with the learning rate of  $1 \times 10^{-3}$ , weight decay of 0.01 and  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$ . We do not use a learning rate scheduler. We set the batch size to 64. The instruction tuning is conducted on a machine with  $8 \times$  L40S GPUs for 8 days. And each downstream fine-tuning is conducted on a single L40S GPU with a  $1 \times 10^{-4}$  learning rate and 48 batch size. For generations, we have adopted the greedy decoding strategy.

We use 128 as the number of fMRI tokens. We use a 4-layer MLP in the fMRI encoder. We set the number of queries to 1024. The dimension of the query embeddings is 128.

Table 9. Instruction template and statistics of the BIT dataset. "—" indicates that the instruction is embedded within the conversation and does not require an additional one. Note that not all conversations are associated with fMRI recordings as only a subset of MSCOCO images were used as stimuli in the study. Consequently, the number of usable conversations in practice will be lower.

Dataset	#(conversations)	Instruction
system prompt	/	You are a helpful agent that decodes the brain activity of a person looking at an image.
Previous Caption	149, 875	Please describe the image the subject saw previously.
COCO Caption	616, 767	Please describe the fMRI as simply as possible.
Image Paragraph Captioning	9, 598	Describe the fMRI in one paragraph.
COCO QA	117, 684	Answer the question with a short phrase.
Visual Genome QA	676, 116	Answer the question with a short phrase.
VQAv2	6, 581, 110	Answer the question with a short phrase.
OK-VQA	140, 550	Answer the question with a short phrase.
ST-VQA	29, 751	Answer the question with a short phrase.
TallyQA	238, 056	Answer the question with a number.
VQA-E	2, 697, 860	Answer with a short phrase and provide explanation for your answer.
A-OKVQA	18, 201	Answer with the option's letter from the given choices and provide explanation for your choice.
FSVQA	369, 861	Answer the question in a full sentence.
VisDial	125, 351	—
Llava Instruct 150K	157, 712	—
LVIS Instruct4V	222, 711	—

### C. Computational Complexity and Runtime Analysis

In the neuroscience-informed attention, the complexity of the dot product between queries and keys is  $O(dNN_q)$ . The complexity of the aggregation of values is  $O(NN_q)$ , which is negligible. The MLP maps the hidden representation of dimension  $N_q$  to  $L \times d$ , therefore its complexity is  $dLN_q$ . Therefore, the complexity of the fMRI encoder is  $O(dNN_q + dLN_q) = O(dN_q(L + N)) = O(dN_qN)$  given that  $L \ll N$ .

Table 10 summarizes the computational complexity and runtime comparison of different models. Despite its superior performance, our model introduces only a marginal increase in encoder-side computational cost compared to other baselines. Notably, the majority of inference time and complexity arises from the LLM component, which is shared across all models. This highlights that the design choices in the encoder, while crucial for performance, do not significantly affect runtime.

### D. Qualitative Analysis

In this section, we present a qualitative analysis of our model on COCO Captioning, COCO-QA, and OK-VQA, comparing its performance against MindBridge (Wang et al., 2024a) and UniBrain (Wang et al., 2024b). As shown in Figure 8 and 9, our results demonstrate significant improvements in visual understanding across multiple tasks. The model shows strength in the following areas: **Static Object Recognition.** The model demonstrates superior accuracy in identifying stationary objects. In comparison with baseline models (MindBridge and UniBrain), our approach shows improvement in spatial context understanding. For example, when analyzing aircraft imagery (e.g., (a) of Figure 8), our model correctly identifies "airplane sitting on the runway" while baselines incorrectly interpret the scene as "flying in the sky" or "flying over a city", demonstrating better state-space recognition. **Action Recognition.** Our proposed model exhibits enhanced capability in distinguishing between similar actions. In sports scenarios (e.g., (f) of Figure 9), our model correctly identifies "catch ball" while both baselines incorrectly predict "serve", indicating improved action-state discrimination. **Potential of neuroscience application.** The demonstrated improvements in object understanding and action recognition suggest the potential for advancing brain-computer interface technology and neural processing research. The model's enhanced capabilities in distinguishing object states and actions could lead to more effective neural prosthetics and improved assistive technologies for individuals with visual or motor impairments.

Table 10. Computational complexity and runtime of each encoder and the LLM. For the first 4 lines, we only time the encoder part, excluding the LLM. The runtime of the LLM forward pass is reported separately in the last row.

Module	Forward Pass (ms/sample)	Complexity	Notation Explanation	Approximate Values
UMBRAE	0.0109	$O(dN + dN_q^2)$	$d$ : hidden size $N$ : number of voxels $N_q$ : number of query tokens	$d = 100$ $N = 15000$ $N_q = 200$
MindBridge	0.0048	$O(dN)$	$d$ : hidden size $N$ : number of voxels	$d = 1000$ $N = 15000$
UniBrain	0.0175	$O(dGK + dN_q^2)$	$d$ : hidden size $G$ : number of groups $K$ : voxels in each group $N_q$ : number of query tokens	$d = 1000$ $G = 500$ $K = 30$ $N_q = 200$
MindLLM	0.0181	$O(dN_qN)$	$d$ : hidden size $N$ : number of voxels $N_q$ : number of query tokens	$d = 100$ $N = 15000$ $N_q = 100$
LLM	0.0621	-	-	-

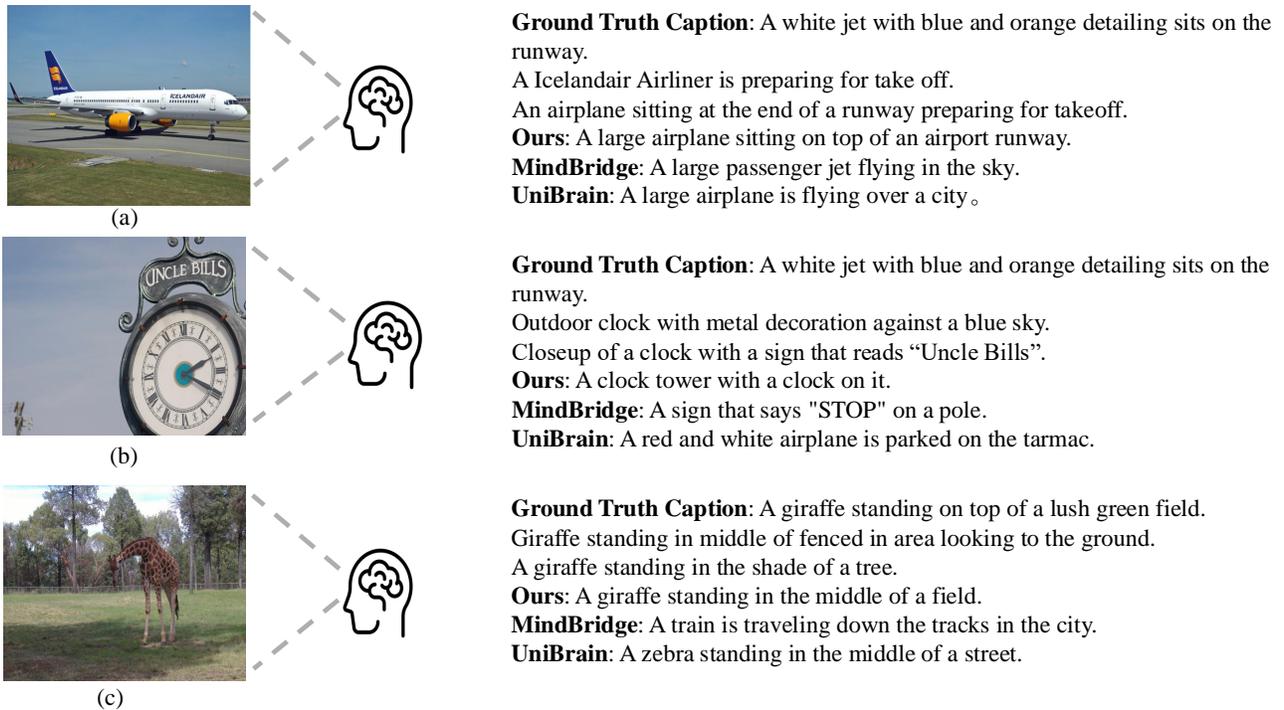


Figure 8. Qualitative Analysis of COCO Captioning.

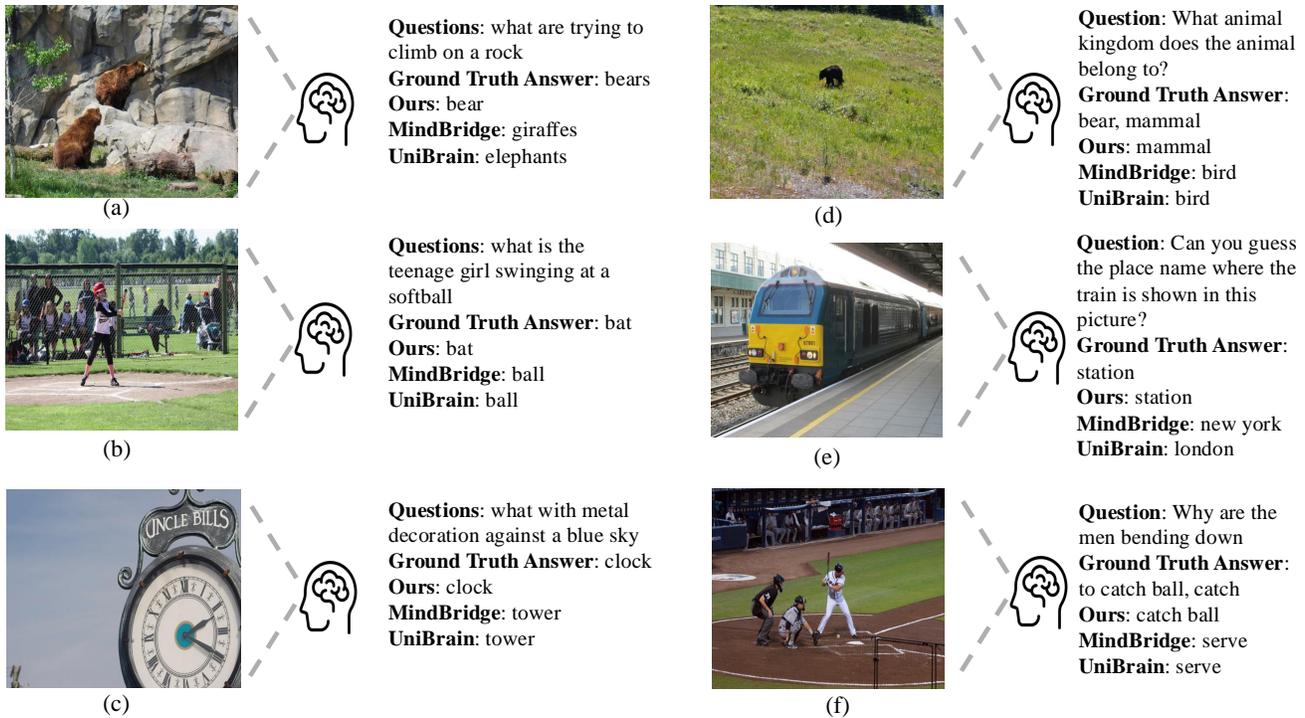


Figure 9. Qualitative Analysis of COCO QA (left column) and OK-VQA (right column).

## E. Ablation Studies on Downstream Tasks

Table 11 reports the results of the ablation study on each downstream task. The results validate the design of neuroscience-informed attention.

## F. Architecture Comparison

We make a side-by-side comparison between our model and related baselines in Table 12. Our model is the only one that achieves the four properties listed in the table’s head.

## G. Experiments on Other Subjects

We report experiments on subjects 2, 5, 7 in Table 13. We chose these subjects because they are commonly used in previous works. We observe consistent performance gains over the baselines in the majority of scenarios.

## H. Latent Space Visualization

To assess how subject-specific and subject-agnostic information evolves through the model, we visualize the latent embeddings at various stages of the encoder in Figure 10. The original inputs exhibit clear subject-specific clusters. After passing through the neuroscience-informed attention layer, the embeddings remain separable by subject, indicating that subject identity is still preserved. As the data flows through successive layers of the MLP, the subject-specific patterns become increasingly mixed, ultimately forming a subject-agnostic representation. This transformation facilitates generalization on held-out subjects in a zero-shot manner.

## I. Brain Attention Maps of Other Subjects

We visualize the flat brain attention maps (corresponding to Figure 7a) in Figure 11. We observe that the attention maps exhibit highly similar spatial patterns across subjects, indicating that the model captures shared structural-functional

Table 11. Ablation study of the key embedding design. Pos Enc. stands for positional encoding. Reg Enc. stands for multiple region encodings.  $(x, y, z)+MLP$  means we employ an MLP to map the coordinates to the embeddings instead of positional encoding. Note that unlike in Table 2, where we fine-tune the model on each task individually, here we do not fine-tune the model on each individual task and report the performances after brain instruction tuning.

		$(x, y, z) + MLP$	Pos Enc. + fMRI	Pos Enc.	Reg Enc.	Pos Enc. + Reg Enc.
COCO-QA	Accuracy $\uparrow$	40.57	43.14	<b>46.29</b>	43.90	<u>44.76</u>
VG-QA	Accuracy $\uparrow$	20.29	20.39	<b>21.12</b>	20.83	<b>21.12</b>
VQA-v2	Accuracy $\uparrow$	45.25	46.06	47.30	<u>47.44</u>	<b>47.86</b>
A-OKVQA	Accuracy $\uparrow$	<u>46.02</u>	44.25	45.13	42.48	<b>46.90</b>
ST-VQA	ANLS $\uparrow$	<b>13.70</b>	9.70	<u>12.79</u>	10.61	11.64
Tally QA	Accuracy $\uparrow$	44.63	49.77	<u>52.29</u>	48.84	<b>53.01</b>
	RMSE $\downarrow$	2.15	1.83	<u>1.75</u>	1.96	<b>1.73</b>
COCO-Caption	BLEU-1 $\uparrow$	49.39	54.86	<u>57.03</u>	52.53	<b>58.98</b>
	BLEU-2 $\uparrow$	26.30	33.84	<u>37.12</u>	33.00	<b>39.08</b>
	BLEU-3 $\uparrow$	14.06	20.67	<u>23.63</u>	20.69	<b>25.56</b>
	BLEU-4 $\uparrow$	8.19	13.42	<u>15.51</u>	13.93	<b>17.36</b>
	METEOR $\uparrow$	12.31	14.97	<u>16.41</u>	14.81	<b>17.12</b>
	ROUGE $\uparrow$	35.45	40.17	<u>42.14</u>	39.62	<b>43.34</b>
	CIDEr $\uparrow$	18.26	32.15	<u>41.32</u>	34.25	<b>46.17</b>
	SPICE $\uparrow$	3.94	7.14	<u>8.56</u>	6.76	<b>9.24</b>
Paragraph Caption	BLEU-1 $\uparrow$	24.13	24.57	<u>25.90</u>	24.74	<b>27.17</b>
	BLEU-2 $\uparrow$	11.88	<u>13.03</u>	12.41	12.61	<b>14.50</b>
	BLEU-3 $\uparrow$	6.12	7.16	6.10	<u>7.20</u>	<b>7.99</b>
	BLEU-4 $\uparrow$	3.55	4.24	3.22	<u>4.35</u>	<b>4.67</b>
	METEOR $\uparrow$	11.19	11.78	9.05	<u>12.12</u>	<b>14.43</b>
	CIDEr $\uparrow$	3.32	2.35	<b>7.46</b>	<u>4.97</u>	4.30
VQA-E	Accuracy $\uparrow$	46.74	48.17	<u>48.60</u>	48.48	<b>48.95</b>
	BLEU-1 $\uparrow$	35.64	36.19	<u>36.46</u>	35.06	<b>36.65</b>
	BLEU-2 $\uparrow$	18.65	19.38	<u>19.53</u>	18.32	<b>19.58</b>
	BLEU-3 $\uparrow$	10.85	11.40	<b>11.62</b>	10.54	<u>11.57</u>
	BLEU-4 $\uparrow$	6.79	7.12	<u>7.29</u>	6.38	<b>7.31</b>
	CIDEr $\uparrow$	83.18	88.14	<b>89.24</b>	85.34	88.96
	METEOR $\uparrow$	14.20	14.81	<b>15.08</b>	14.54	<b>15.08</b>
	ROUGE $\uparrow$	33.81	<u>35.15</u>	<b>35.17</b>	34.28	34.76
Average	Perplexity $\downarrow$	6.39	5.82	<u>5.47</u>	5.73	<b>5.38</b>

Method	cross-subject	unified architecture	shared parameters	using all voxels
MindEye (Scotti et al., 2024a)	$\times$	$\times$	$\times$	$\checkmark$
MindEye2 (Scotti et al., 2024b)	$\checkmark$	$\times$	$\times$	$\checkmark$
UMBRAE (Xia et al., 2024)	$\checkmark$	$\times$	$\times$	$\checkmark$
MindBridge (Wang et al., 2024a)	$\checkmark$	$\checkmark$	$\times$	$\times$
UniBrain (Wang et al., 2024b)	$\checkmark$	$\checkmark$	$\checkmark$	$\times$
MindLLM	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$

Table 12. Comparison of our model and baselines. Concepts of properties: *cross-subject*: the model performs across subjects; *unified architecture*: the architecture for each subject is the same (the parameters might be different, e.g., MLPs with the same size but different parameters); *shared parameters*: The model uses the same parameters for all subjects; *using all voxels*: the model uses all voxels as inputs, thereby mitigating the risk of information loss from voxel omission. Refer to Section 2.2 and Figure 2 for details of this property.

Table 13. Breakdown performances on subjects 2, 5, 7 respectively.

		subject 2				subject 5				subject 7			
		UMBRAE	UniBrain	MindBridge	MindLLM	UMBRAE	UniBrain	MindBridge	MindLLM	UMBRAE	UniBrain	MindBridge	MindLLM
COCO-QA	Accuracy ↑	26.85	47.14	49.04	<b>51.71</b>	29.86	45.52	50.57	<b>53.14</b>	25.78	45.05	50.19	<b>50.76</b>
	Accuracy ↑	18.65	23.53	24.97	<b>25.96</b>	19.38	23.65	26.11	<b>26.89</b>	18.44	23.02	25.16	<b>25.92</b>
VQA-v2	Accuracy ↑	49.96	48.17	51.28	<b>53.17</b>	51.49	48.81	51.78	<b>53.09</b>	50.13	47.97	50.60	<b>52.70</b>
A-OKVQA	Accuracy ↑	40.54	43.36	44.25	<b>49.56</b>	43.24	45.13	50.44	<b>50.74</b>	40.54	46.02	45.13	<b>46.90</b>
ST-VQA	ANLS ↑	6.61	9.34	10.49	<b>10.85</b>	4.31	9.77	<b>10.49</b>	9.63	5.46	8.26	13.48	<b>13.71</b>
OK-VQA	Accuracy ↑	20.70	32.82	33.50	<b>35.73</b>	22.98	30.43	34.36	<b>35.04</b>	17.89	34.19	34.54	<b>34.87</b>
TallyQA	Accuracy ↑	42.21	55.99	0.20	<b>57.84</b>	42.69	54.65	58.71	<b>59.3</b>	40.95	53.98	56.20	<b>58.41</b>
	RMSE ↓	4.85	1.75	<b>1.65</b>	1.66	4.86	1.64	<b>1.60</b>	<b>1.60</b>	3.65	1.72	<b>1.66</b>	1.73
Paragraph Caption	BLEU-1 ↑	29.29	27.53	27.19	<b>30.13</b>	29.21	29.30	28.70	<b>34.54</b>	29.61	25.54	27.45	<b>29.54</b>
	BLEU-2 ↑	13.82	14.73	14.75	<b>16.29</b>	14.37	15.38	15.21	<b>19.11</b>	14.21	13.26	14.44	<b>16.19</b>
	BLEU-3 ↑	6.44	8.42	8.70	<b>11.69</b>	6.84	8.73	8.87	<b>10.11</b>	6.48	7.42	8.05	<b>9.34</b>
	BLEU-4 ↑	3.08	5.10	5.46	<b>5.96</b>	3.28	5.23	5.53	<b>7.60</b>	2.95	4.52	4.84	<b>5.73</b>
	METEOR ↑	11.94	12.31	12.36	<b>13.06</b>	12.68	11.63	12.63	<b>13.09</b>	12.47	11.31	12.57	<b>13.54</b>
	CIDEr ↑	5.40	7.53	0.21	<b>11.90</b>	7.46	6.00	8.63	<b>12.97</b>	<b>8.84</b>	4.03	5.62	8.53
VQA-E	Accuracy ↑	47.62	50.03	52.30	<b>54.92</b>	49.06	50.18	54.21	<b>55.75</b>	48.95	49.48	53.84	<b>54.68</b>
	BLEU-1 ↑	30.05	36.63	37.48	<b>38.13</b>	30.55	37.41	38.07	<b>38.40</b>	29.58	36.71	38.33	<b>38.41</b>
	BLEU-2 ↑	14.87	19.60	20.16	<b>20.77</b>	14.99	19.96	20.90	<b>21.04</b>	14.43	19.78	20.90	<b>20.98</b>
	BLEU-3 ↑	8.31	11.49	11.76	<b>12.34</b>	8.30	11.66	12.54	<b>12.60</b>	7.97	11.80	12.49	<b>12.56</b>
	BLEU-4 ↑	4.99	7.18	7.26	<b>7.74</b>	4.93	7.15	7.94	<b>8.17</b>	4.68	7.53	<b>7.99</b>	7.90
	CIDEr ↑	64.60	89.50	91.93	<b>97.68</b>	66.47	89.91	95.30	<b>99.48</b>	61.34	90.49	98.23	<b>99.36</b>
	METEOR ↑	12.37	15.12	15.56	<b>16.02</b>	12.51	15.37	15.88	<b>16.28</b>	12.09	15.02	16.07	<b>16.16</b>
	ROUGE ↑	28.49	35.10	35.86	<b>36.52</b>	28.75	35.42	36.43	<b>36.82</b>	28.04	35.24	36.61	<b>36.73</b>
FSVQA	VQA Acc. ↑	16.37	45.62	48.62	<b>50.19</b>	16.40	45.66	49.45	<b>50.76</b>	16.37	46.11	49.19	<b>51.13</b>
	FSVQA Acc. ↑	0.00	40.13	43.57	<b>45.21</b>	0.00	45.66	44.60	<b>45.53</b>	0.00	40.07	44.70	<b>46.07</b>
	BLEU-1 ↑	18.98	86.03	87.28	<b>87.84</b>	19.00	85.99	87.76	<b>87.89</b>	18.98	85.48	87.29	<b>88.05</b>
	BLEU-2 ↑	9.11	81.66	83.14	<b>83.94</b>	11.94	81.52	83.25	<b>83.87</b>	10.89	81.00	83.42	<b>84.16</b>
	BLEU-3 ↑	1.98	77.61	79.31	<b>80.28</b>	3.42	77.34	79.64	<b>80.21</b>	2.82	76.84	79.70	<b>80.60</b>
	BLEU-4 ↑	0.00	73.45	75.42	<b>76.51</b>	1.65	73.05	75.99	<b>76.45</b>	0.00	72.56	76.11	<b>77.01</b>
	METEOR ↑	5.65	48.06	49.26	<b>49.88</b>	5.67	48.01	49.57	<b>50.10</b>	5.66	47.86	49.74	<b>50.18</b>
	CIDEr ↑	3.19	647.62	671.26	<b>681.94</b>	3.26	648.21	675.99	<b>684.12</b>	3.37	643.44	676.50	<b>684.72</b>

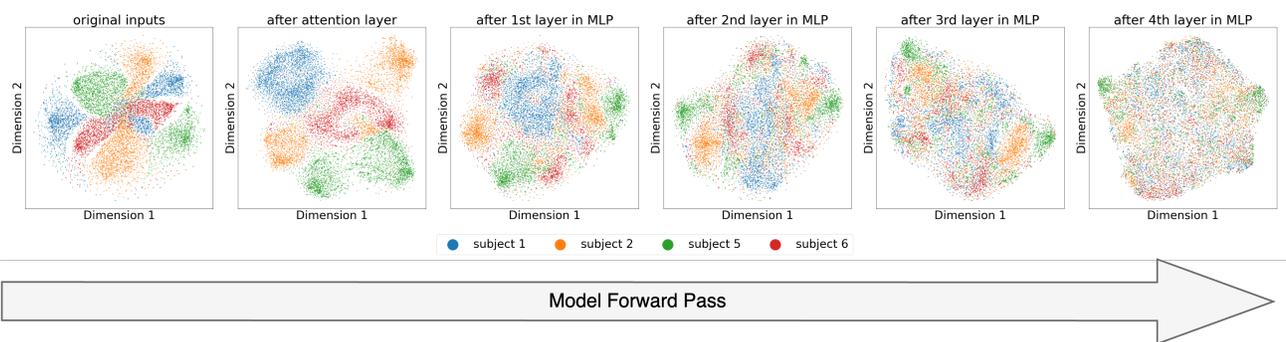


Figure 10. Visualization of latent embeddings across model layers using T-SNE.

correspondences across human brains, which is expected due to the overall anatomical and functional similarity among individuals. At the same time, we do observe moderate subject-specific variations in the attention maps. These reflect differences in voxel-level functional signals and individual variability in precise voxel locations. The model is able to accommodate these differences through flexible attention mechanisms.

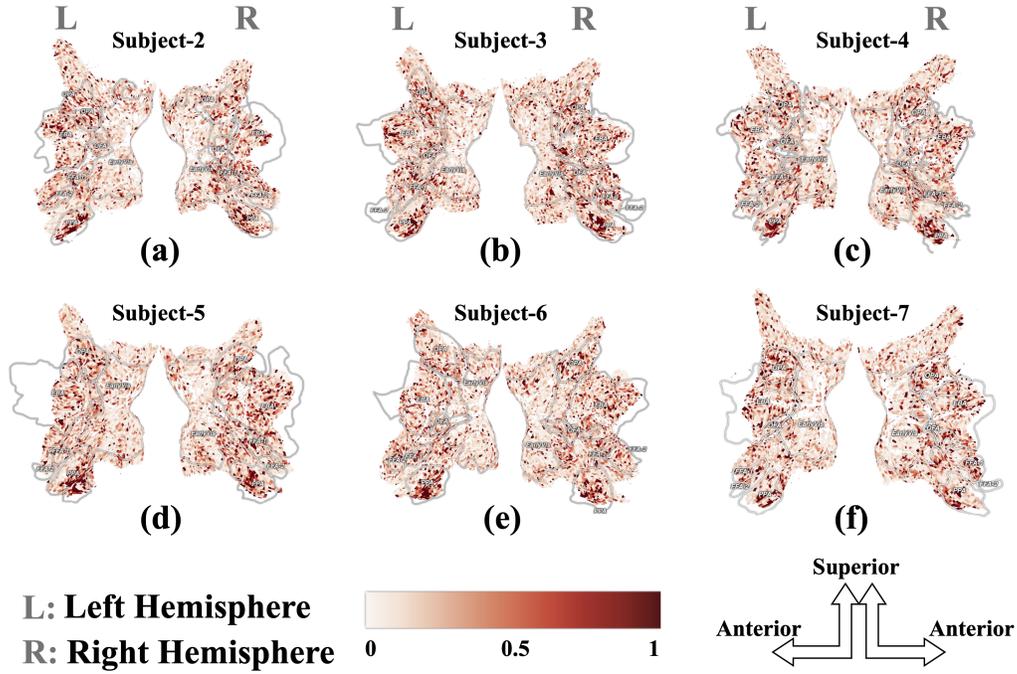


Figure 11. Figure 7a’s query token on different subjects. It shows similar patterns across subjects, suggesting shared brain organization, while also capturing some subject-specific variances.