

MindTuner: Cross-Subject Visual Decoding with Visual Fingerprint and Semantic Correction

Zixuan Gong¹, Qi Zhang¹, Guangyin Bao¹,
Lei Zhu¹, Rongtao Xu², Ke Liu³, Liang Hu¹, Duoqian Miao^{1,*}

¹Tongji University

²Chinese Academy of Sciences

³Beijing Anding Hospital

{gongzx,zhangqi_cs,baogy,dqmiao}@tongji.edu.cn

Abstract

Decoding natural visual scenes from brain activity has flourished, with extensive research in single-subject tasks and, however, less in cross-subject tasks. Reconstructing high-quality images in cross-subject tasks is a challenging problem due to profound individual differences between subjects and the scarcity of data annotation. In this work, we proposed MindTuner for cross-subject visual decoding, which achieves high-quality and rich semantic reconstructions using only 1 hour of fMRI training data benefiting from the phenomena of visual fingerprint in the human visual system and a novel fMRI-to-text alignment paradigm. Firstly, we pre-train a multi-subject model among 7 subjects and fine-tune it with scarce data on new subjects, where LoRAs with Skip-LoRAs are utilized to learn the visual fingerprint. Then, we take the image modality as the intermediate pivot modality to achieve fMRI-to-text alignment, which achieves impressive fMRI-to-text retrieval performance and corrects fMRI-to-image reconstruction with fine-tuned semantics. The results of both qualitative and quantitative analyses demonstrate that MindTuner surpasses state-of-the-art cross-subject visual decoding models on the Natural Scenes Dataset (NSD), whether using training data of 1 hour or 40 hours.

Introduction

Do our brains form unified perceptions when we observe similar objects? Do our unique understandings influence these perceptions differently? The human brain exhibits substantial anatomical similarities in terms of functional organization, including shared attributes like memory, functional connectivity, and visual cortex functions (Chen et al. 2017; Fingelkurts, Fingelkurts, and Kähkönen 2005; Stringer et al. 2019). However, individual neural biases always exist due to inherent differences (Wang, Murai, and Whitney 2020). Understanding both the similarities and gaps in perception has profound implications for the fields of Artificial Intelligence (AI) (Nie et al. 2023; Zhao et al. 2014) and Brain-Computer Interface (BCI) research (Wang et al. 2009). Visual decoding is a straightforward way to understand the brain, where functional magnetic resonance imaging (fMRI) is a widely embraced non-invasive tool used to decode natural visual stimuli, revealing intricate perceptual and semantic details

in the cerebral cortex (Linden 2021). Consequently, fMRI has garnered considerable attention in image retrieval and reconstruction tasks.

Open-source large-scale fMRI datasets, such as the Natural Scenes Dataset (NSD) (Allen et al. 2022), advance deep learning models to shine in fMRI decoding. Pre-trained cross-modality models like CLIP (Radford et al. 2021) and Stable Diffusion (Rombach et al. 2022) offer effective representation space and models for high-quality visual reconstruction. A large body of literature demonstrates the feasibility of training single-subject decoding models to reconstruct high-fidelity images (Lin, Sprague, and Singh 2022; Scotti et al. 2023; Chen et al. 2023; Takagi and Nishimoto 2023; Lu et al. 2023). However, single-subject decoding has drawbacks, including the need to train a unique model for each subject, making it challenging to generalize to new subjects and requiring a substantial amount of fMRI data for training. As is widely recognized, acquiring a large amount of fMRI data for each subject is time-consuming, labor-intensive, and impractical in practical scenarios. Unfortunately, most current research focuses on single-subject visual decoding rather than exploring the challenging commonalities of the human brain. Consequently, there is an urgent need for cross-subject decoding models that can be effectively transferred to new subjects and perform well in the few-shot setting, as depicted in Figure 1.

The key to cross-subject few-shot decoding lies in effectively utilizing extensive prior knowledge from other subjects or additional modalities. On the one hand, one successful strategy for leveraging knowledge from other subjects involves aligning them to a shared space. Ridge regression is commonly employed for this purpose, aligning voxel inputs from different subjects (Scotti et al. 2024; Ferrante, Bocato, and Toschi 2023). This approach is preferred due to the low signal-to-noise ratio in fMRI data, where complex non-linear models tend to overfit noise. Nonetheless, the process of visual information perception and generating brain activity in each individual incorporates unique components, referred to as the visual fingerprint (Wang, Murai, and Whitney 2020) (refer to Preliminary for more detailed analysis). Current linear alignment methods only enable new subjects to conform to the shared components across subjects, neglecting the perception difference derived from their distinctive visual fingerprint and resulting in limited performance.

*Corresponding author

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: Cross-Subject Visual Decoding and Image Reconstruction. Subjects with adequate fMRI data are aligned to decode visual stimuli via learning a shared network. A new subject, even with scarce visual stimulus, is aligned to the common space of the shared network, which perceives the subject’s unique visual fingerprint to ensure precise visual decoding.

On the other hand, an additional strategy involves leveraging multimodal data. Previous alignment methods focused solely on the visual modality, as a means of adapting to the inputs of Stable Diffusion. However, this alignment approach is susceptible to slight disturbances, leading to semantic errors in the generated images. For visual decoding tasks, the textual modality is highly relevant and is verified effective in enhancing visual decoding semantically (Scotti et al. 2024). However, previous approaches incorporating text have placed excessive emphasis on directly aligning fMRI with detailed textual descriptions to facilitate the reconstruction process (Takagi and Nishimoto 2023), which intuitively lacks rationality and yields poor performance. Considering that subjects in visual stimulation experiments lack direct interaction with the textual modality and that individual understanding of visual stimuli varies, the relationship between fMRI and text should be considered implicit.

In this paper, we propose **MindTuner**: a cross-subject visual decoding framework. Inspired by visual fingerprint in Brain Science (Wang, Murai, and Whitney 2020), and with the help of Low-Rank Adaptation (LoRA) for large model lightweight fine-tuning. In correspondence to the above two strategies, we first propose the combination of non-linear **Skip-LoRAs** and **LoRAs** to learn the visual fingerprint of new subjects, which are injected into the fMRI encoding network to correct visual perception difference. In addition, we design a **Pivot** module that uses images as the central modality to bridge fMRI and text. The Pivot helps to correct the reconstructed image with fine-tuned semantics. Our contributions are summarized as follows:

- MindTuner makes the first attempt to introduce **LoRA** as a subject-level fine-tuning module in cross-subject decoding and further elaborately design non-linear **Skip-LoRA**. Their combination shows excellent capability to learn subjects’ visual fingerprint.
- We introduce a novel fMRI-to-text retrieval paradigm with a **Pivot** using the image modality. The Pivot conducts semantic correction with label prompts to enhance fMRI-to-image reconstruction.
- We evaluate our method on the NSD dataset, and it establishes SOTA decoding performance whether using train-

ing data of 40 hours (full) or 1 hour(2.5% of full).

Related Work

Cross-Subject Functional Alignment

In visual decoding, single-subject models have exposed the issue of excessive reliance on the data volume of individual subjects, leading researchers to shift toward cross-subject studies. Functional alignment of different brains is considered to be more effective than anatomical alignment. Previous functional alignment methods were mainly divided into two perspectives: fMRI data itself and downstream tasks. Among a series of methods starting from the fMRI perspective, Bazeille et al. (Bazeille et al. 2019) and Thual et al. (Thual et al. 2023) minimize an optimal transport cost between voxels of different brains. The methods based on fMRI self-supervision (Chen et al. 2023; Qian et al. 2023) emphasize obtaining common latent representations between subjects through autoencoders. Wills Aligner (Bao et al. 2024) aligns data from different subjects using the fsaverage template. On another cross-subject alignment research path, most of these methods focus on the shared knowledge among subjects and overlook the complex nonlinear relationships between subjects due to concerns about overfitting. Linear fitting (Ferrante, Boccato, and Toschi 2023) was performed on the responses of subjects to common images, achieving high-quality cross-subject visual reconstruction, but requiring subjects to see the same images. Mindeye2 (Scotti et al. 2024), MindBridge (Wang et al. 2024) and UMBRAE (Xia et al. 2024b) improved this and achieved impressive results, but still losing subject-specific features and ignoring the more complex relationships between subjects, which is not flexible.

Text Modality in Visual Decoding

Decoding visual stimuli from fMRI has been a long-standing endeavor, primarily focusing on the image modality (Takagi and Nishimoto 2023; Lu et al. 2023; Gong et al. 2024b; Ozcelik and VanRullen 2023; Scotti et al. 2023). However, an increasing number of studies have highlighted the role of text. UniBrain (Mai and Zhang 2023) directly aligns fMRI

to the text representation via ridge regression and then completes the brain caption task with the text generation model. MindEye2 (Scotti et al. 2024) obtains a better caption by placing the generated image representation into the image captioning model, which is used to smooth the generated image. NeuroClips utilizes the pretrained BLIP-2 (Li et al. 2023) model to achieve text-assisted video reconstruction (Gong et al. 2024a). However, these methods either do direct fMRI-to-text alignment or direct image-text alignment, ignoring the indirect relationship between fMRI and text from the perspective of intuitive understanding.

Preliminary on Visual Fingerprint

Visual stimuli are processed by the Human Visual System(HVS). Therefore, as stated in previous research (Xia et al. 2024a), brain responses such as fMRI are closely linked to our visual system. Although presented in a systematic manner, there are still significant differences in the visual systems of different individuals. Research of visual fingerprint emphasizes that idiosyncratic biases exist in the underlying representation of visual space propagate across varying levels of visual processing (Wang, Murai, and Whitney 2020). Using a position-matching task will find stable subject-specific compressions and expansions within local regions throughout the visual field. As shown on the left of Figure 2, in experiments, subjects were fixated at the center, and a target was displayed briefly at one of five possible eccentricities (depicted by dashed lines, which were not visible in the experiment). After the target disappeared, subjects moved the cursor to match the target’s location. Linear models were used to fit the Distortion Indices(DI), which measured the degree of spatial distortion between subjects:

$$\begin{aligned} DI_{self} &\sim \beta_0 + \beta_1 \times self, \\ DI_{others} &\sim \beta'_0 + \beta'_1 \times others. \end{aligned} \quad (1)$$

As shown on the right of Figure 2, the experiment results show that subjects have their own visual fingerprint, with both linear and non-linear components. Within-subject similarity ($r=0.71$) is significantly higher than between-subject similarity ($r=0.22$), suggesting that each individual subject has their own unique spatial distortions that are consistent within themselves and distinguished from others.

Method

The task of cross-subject visual decoding consists of two parts: multi-subject pre-training and new-subject fine-tuning. Given a neural dataset with brain activities of subject s , it involves the reconstruction of visual images $I \in \mathbb{R}^{3 \times H \times W}$ through a core pipeline whose inputs are flattened and aligned fMRI voxels $V \in \mathbb{R}^{1 \times d_s}$ after ROI extraction, where d_s denotes voxel numbers. To simplify notation, each (I, V) denotes data from the original subjects for this section. Our goal is to optimize the $\mathcal{F}(\cdot)$, so that $\mathcal{F}(V) = \hat{I}$, where \hat{I} best approximates I . New-subject fine-tuning adheres to the same path, albeit with scarce data. For the multi-subject pre-training, we follow the pipeline of MindEye2, while for the new-subject fine-tuning, we plug into visual fingerprint and Pivot as shown in Figure 3.




Figure 2: Visual fingerprint experiments across subjects. 'Within' denotes Pearson correlation coefficient of Distortion Indices in within-subject experiments, while 'between' denotes between-subject.

Multi-Subject Pre-training

Multi-Subject Functional Alignment It should be noted that the brain structures of different subjects vary, as do the number of voxels obtained. Consequently, a mapping model is required to align the voxel inputs from different subjects to the same dimension as the inputs to the model in the phase of multi-subject pre-training. Here, we employ linear function alignment to learn shared-subject fMRI latent space $M \in \mathbb{R}^{1 \times d_0}$ (d_0 denotes shared input dimension). This is achieved through subject-specific ridge regression, as detailed below:

$$M = \text{Ridge}^{(s)}(V) \quad (2)$$

MLP Backbone In order to achieve high-fidelity reconstruction, it is necessary to utilize a substantial quantity of CLIP image embedding. The OpenCLIP ViT-bigG/14 space is employed for alignment, with an image embedding dimension of 256×1664 . Mapped inputs M are fed into an MLP backbone comprising four residual blocks and a tokenization layer which transforms the input from a dimension of d_0 to 256×1664 . The MLP Backbone $\epsilon(\cdot)$ serves to convert the fMRI to the intermediate backbone embedding space: $Z = \epsilon(M)$. It should be noted that all subjects shared the same MLP backbone following multi-subject functional alignment. Subsequently, the backbone embeddings are conveyed into three submodules for retrieval, high-level reconstruction, and low-level reconstruction.

Retrieval Submodules A straightforward approach to performing the retrieval task is to conduct a shallow mapping of the backbone embeddings and supervise it with an fMRI-to-image CLIP contrastive loss. In the case of limited fMRI data, the application of appropriate data augmentation techniques can facilitate the convergence of the model. A recently proposed voxel mixture paradigm, based on MixCo (Kim et al. 2020), has demonstrated effectiveness. Two raw fMRI voxels V_i and V_j are mixed into $V_{mix_{i,j}}$ using a factor λ sampled from the Beta distribution:

$$\begin{aligned} V_{mix_{i,j}} &= \lambda_i V_i + (1 - \lambda_i) V_j, \\ M_{mix_{i,j}} &= \text{Ridge}^{(s)}(V_{mix_{i,j}}), \\ Z_{mix_{i,j}} &= \epsilon(M_{mix_{i,j}}) \end{aligned} \quad (3)$$

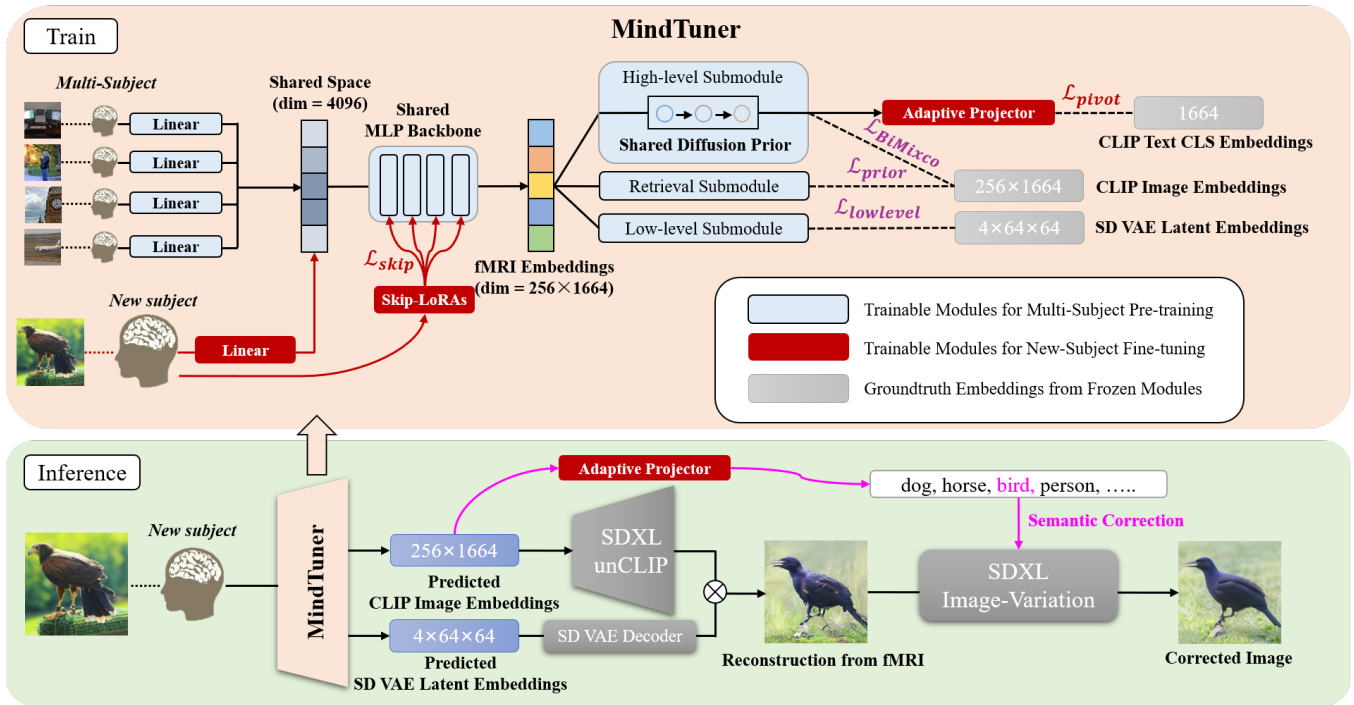


Figure 3: Schematic diagram of MindTuner. The training process was split into two phases: Multi-subject Pre-training and Cross-subject Fine-tuning, in which the corresponding modules were trained. The predicted embeddings are first obtained through MindTuner, and then the preliminary reconstructed image was obtained by SDXL unCLIP. The final reconstructed image is obtained by text retrieval and semantic correction by SDXL Image-Variation.

where j denotes an arbitrary mixing index in the batch. The forward mixed contrastive loss MixCo is formulated as:

$$\mathcal{L}_{MixCo} = -\frac{1}{|B|} \sum_{i=1}^{|B|} [\lambda_i \log \frac{\exp(Z_{mix_{i,j}} \cdot g_i / \tau)}{\sum_{m=1}^{|B|} \exp(Z_{mix_{i,j}} \cdot g_m / \tau)} + (1 - \lambda_i) \log \frac{\exp(Z_{mix_{i,j}} \cdot g_j / \tau)}{\sum_{m=1}^{|B|} \exp(Z_{mix_{i,j}} \cdot g_m / \tau)}] \quad (4)$$

where g denotes ground-truth CLIP image embeddings, τ denotes a temperature hyperparameter, and B is the batch size. Here we use the bidirectional loss $\mathcal{L}_{BiMixCo}$.

Low-level and High-level Submodules The low-level pipeline is a widely utilized technique for the enhancement of low-level visual metrics in reconstruction images. This involves the mapping of voxels to the latent space of Stable Diffusion’s Variational AutoEncoder (SDVAE), which serves as a surrogate for the reconstruction. The pipeline comprises an MLP and a CNN upsampler with L1 loss in Stable Diffusion’s latent embeddings z .

$$\mathcal{L}_{lowlevel} = \frac{1}{|B|} \sum_{i=1}^{|B|} |z_i - \hat{z}_i| \quad (5)$$

Conversely, a high-level pipeline places greater emphasis on semantic alignment. Inspired by DALLE-2 (Ramesh et al. 2022), a diffusion prior is recognized as an effective means of transforming backbone embeddings into CLIP ViT image

embeddings, in which mean square error loss is used (further insights on the deployment of diffusion priors can be found in Appendix B.4):

$$\mathcal{L}_{prior} = \frac{1}{|B|} \sum_{i=1}^{|B|} \|g_i - \hat{g}_i\|_2^2 \quad (6)$$

Thus, the end-to-end loss for multi-subject pre-training is:

$$\mathcal{L}_{multi} = \mathcal{L}_{prior} + \alpha_1 \mathcal{L}_{lowlevel} + \alpha_2 \mathcal{L}_{BiMixCo} \quad (7)$$

Where α denotes the weight to balance multiple losses.

New-Subject Fine-tuning

Low-Rank Adaptation Previous work has demonstrated the effectiveness of low-rank adaptation in fine-tuning large language models with significantly fewer trainable parameters. It is well suited to our multi-subject decoding task for two reasons. First, most of the current mainstream models for fMRI decoding are MLP-based models that contain a large number of linear layers, whereas LoRA has been shown to achieve good fine-tuning results in the linear layers. Second, in cross-subject scenarios, fMRI data from new subjects are usually scarce, and full-tuning the whole model is usually difficult to grasp, leading to a certain degree of overfitting. For each pre-trained weight matrix $\in \mathbb{R}^{d_{in} \times d_{out}}$ in the multi-subject model, where d_{in} denotes input dimension and d_{out} denotes output dimension, gradient update is

constrained with a low-rank decomposition for new-subject adapter matrix ΔW :

$$W + \Delta W = W + BA \quad (8)$$

where W is kept frozen, $B \in \mathbb{R}^{d_{out} \times r}$, $A \in \mathbb{R}^{r \times d_{in}}$, r is the rank and $r \ll \min(d_{in}, d_{out})$. At the beginning of the training phase, the parameters of the matrix B are randomly initialized, and A is initialized to zero, which ensures that the initial output of the LoRA block is all zeros.

non-linear Skip-LoRAs The LoRA model is notably lightweight, effectively circumvents complexity, and could avoid the aforementioned fMRI overfitting problem. Although LoRA can be an effective means of fine-tuning multi-subject models, simple linear LoRA models are insufficient for the capture of visual fingerprints. Indeed, there exists a non-linear relationship between subjects (discussed in the Preliminary Section). Therefore, we inject a non-linear design into it. Here, we design our non-linear LoRAs by adding the activation function and the nonlinear constraints. Inspired by the skip-connection of Unet in Computer Vision (Zhang, Rao, and Agrawala 2023), we built a brand new Skip-LoRAs to bootstrap the initial non-linearity of fMRI between subjects directly affecting the entire MLP backbone (see Appendix B.2 for Skip-LoRAs). Skip-LoRAs can be defined as **Skip-LoRA** = $Activation(BA)$. Assuming that $\epsilon_h(\cdot)$ is the h -th layer of MLP backbone $\epsilon(\cdot)$, the output of the new-subject backbone M_h could be as follows:

$$\begin{aligned} M_0 &= \mathbf{Ridge}^{(s)}(V), \\ M_h &= M_{h-skip} + M_{h-linear}, \\ M_{h-linear} &= \epsilon_h(M_{h-1}) + \mathbf{LoRA}_h(M_{h-1}), \\ M_{h-skip} &= \mathbf{Skip-LoRA}_h(V) \end{aligned} \quad (9)$$

Here, we use the Pearson correlation coefficient to define the non-linear correlation loss:

$$\mathcal{L}_{skip} = \frac{1}{|B|} \frac{1}{|H|} \sum_{i=1}^{|B|} \sum_{h=1}^{|H|} |Pearson(M_{h-linear}, M_{h-skip})| \quad (10)$$

Pivot with Adaptive Projector In addition to pixel-level alignment, semantic-level matching is equally important for reconstructing the semantic integrity of an image. In visual decoding tasks, fMRI is directly associated with the image stimulus, and semantic information is implicit in fMRI. More details on the Adaptive Projector can be found in Appendix B.3. For new subjects, we add the additional image-text loss to the model to constrain the semantic information:

$$\mathcal{L}_{pivot} = -\frac{1}{|B|} \sum_{i=1}^{|B|} \log \frac{\exp(p_i \cdot t_i / \tau)}{\sum_{m=1}^{|B|} \exp(p_i \cdot t_m / \tau)} \quad (11)$$

where p denotes the projector's output from image tokens g , and t denotes the CLS embeddings (CLS denotes the classification token) of paired text. During new-subject fine-tuning, the adaptive projector is trainable. Thus, the end-to-end loss for new-subject fine-tuning is:

$$\mathcal{L}_{new} = \mathcal{L}_{multi} + \alpha_3 \mathcal{L}_{skip} + \alpha_4 \mathcal{L}_{pivot} \quad (12)$$



Figure 4: SDXL unCLIP reconstructions and SDXL Image-Variation by MindEye2's refinement or our correction.

Semantic Correction MindEye2 has found that the reconstructed images from SDXL unCLIP after token alignment have fuzzy semantics, so the aligned embeddings are fed into the image captioning model GIT (Wang et al. 2022) to get the semantics for refinement. We have found that single category words are sufficient to accomplish the refinement task. In addition, the captions generated by MindEye2 are completely dependent on aligned embeddings, which can only make the image semantics clearer and cannot correct the semantics of the reconstructed images, as shown in Figure 4. Taking advantage of our adaptive projector, we can define category description as an fMRI-to-text retrieval task using a simple text prompt $[class]$. In this way, we achieve more accurate category reconstruction.

Experiment

Datasets

Natural Scenes Dataset (NSD) (Allen et al. 2022) is an extensive 7T fMRI dataset gathered from 8 subjects viewing images from the MSCOCO-2017 dataset, which contains images of complex natural scenes. Participants viewed three repetitions of 10,000 images with a 7-Tesla fMRI scanner over 30–40 sessions, with one session including 750 fMRI trials lasting for 1 hour. Details of the dataset information can be found in Appendix A. In this paper, we conducted 4 types of experiments, image retrieval and reconstruction in the next section, and text retrieval and brain correlation experiments in Appendix C.

Implementation Details

All of our fine-tuning experiments were run for 150 epochs on two Tesla v100 32GB GPUs with a batch size of 10. The experimental parameter settings for 1-hour and 40-hour data are consistent. For fine-tuning experiments, the loss weight is set to $[\alpha_1, \alpha_2, \alpha_3, \alpha_4] = [0.5, 1.0, 1.5, 0.5]$ and the rank r is set to 8 for all LoRA blocks, including Skip-LoRAs.

Method	Low-Level				High-Level				Retrieval	
	PixCorr \uparrow	SSIM \uparrow	Alex(2) \uparrow	Alex(5) \uparrow	Incep \uparrow	CLIP \uparrow	Eff \downarrow	SwAV \downarrow	Image \uparrow	Brain \uparrow
Takagi...	0.246	0.410	78.9%	85.6%	83.8%	82.1%	0.811	0.504	-	-
Ozcelik...	0.273	0.365	94.4%	96.6%	91.3%	90.9%	0.728	0.422	18.8%	26.3%
MindEye1	0.319	0.360	92.8%	96.9%	94.6%	93.3%	0.648	0.377	90.0%	84.1%
MindEye2	0.322	0.431	96.1%	98.6%	<u>95.4%</u>	<u>93.0%</u>	<u>0.619</u>	<u>0.344</u>	98.8%	98.3%
MindTuner(Ours)	0.322	<u>0.421</u>	<u>95.8%</u>	98.8%	95.6%	93.8%	0.612	0.340	98.9%	98.3%
MindEye2(1 hour)	0.195	0.419	84.2%	90.6%	81.2%	79.2%	0.810	0.468	79.0%	57.4%
MindTuner(1 hour)	0.224	0.420	87.8%	93.6%	84.8%	83.5%	0.780	0.440	83.1%	76.0%

Table 1: Quantitative comparison of MindTuner’s performance against other methods. All results are averaged across subjects 1, 2, 5, and 7 from the Natural Scenes Dataset. The results of other methods are taken from MindEye2(Scotti et al. 2024), using either 40-hour data or 1-hour data. Some methods were not included in the comparison because they used past 37-hour data or were not open-source. Missing values occur when metrics are not applicable. Bold font signifies the best performance, while underlined text indicates the second-best performance (See Appendix B.1 for more details about the metrics).

We use the AdamW (Loshchilov and Hutter 2017) for optimization, with a learning rate set to $3e-4$, to which the OneCircle learning rate schedule (Smith and Topin 2019) was set. During the training process, we use data augmentation from images and blurry images and replace the BiMixCo with SoftCLIP (Scotti et al. 2023) loss in one-third of the training phase. In the inference stage, we only generated a reconstructed image once and did not make multiple selections. The final high-level reconstructed image and low-level blurry image are simply weighted and averaged in a ratio of 3:1. In the inference stage of semantic correction, we use 80 category nouns from MSCOCO as correction texts and classify them in the form of text retrieval.

Results and Analysis

Image and Brain Retrieval

Image retrieval refers to retrieving the image embeddings with the highest cosine similarity based on fMRI embeddings on the test set. In Table 1, there is only a slight improvement in retrieval accuracy within 40 hours of data. This can be attributed to the fact that the retrieval accuracy is already close to the upper limit (only about 1% short of reaching 100%), which is affected by the noise of the fMRI dataset itself. However, when only 1 hour of data was available, MindTuner’s retrieval accuracy was significantly higher than MindEye2, 4.1% higher for image retrieval and 18.6% higher for brain retrieval. As MindEye2 benefits from multi-subject pre-training and new-subject fine-tuning with linear heads, this suggests that the visual fingerprint we introduced significantly improves the performance of the model when fMRI data are scarce.

Image Reconstruction

Image reconstruction aims to restore the original image as seen by the subjects, and two levels of evaluation metrics assess the quality of the reconstructions. Previous research has improved the performance of these above metrics in single-subject models by various means. The cross-subject task discussed in this paper tests the ability of the model to exploit




Figure 5: MindTuner vs MindEye2 reconstructions from fMRI brain activity with only 1 hour of data.



Figure 6: MindTuner vs MindEye2 reconstructions from fMRI brain activity with 40 hours of data.

the commonalities of multi-subject and to migrate new subjects, ultimately realizing the goal of using less fMRI data for few-shot learning. Here, we report the reconstruction results of MindTuner using 1-hour and 40-hour training data. Table 1 reflects the quantized performance comparisons, under two different training data sizes. It can be seen that when using the full NSD dataset, MindTuner achieves better performance on high-level metrics; when only 1 hour of data is available for training, MindTuner outperforms MindEye2 on all metrics. The visualization results of 1-hour and 40-

Method	Trainable Parameters	Low-Level				High-Level				Retrieval	
		PixCorr \uparrow	SSIM \uparrow	Alex(2) \uparrow	Alex(5) \uparrow	Incep \uparrow	CLIP \uparrow	Eff \downarrow	SwAV \downarrow	Image \uparrow	Brain \uparrow
MindEye2	64.4M(-)	0.235	0.428	88.0%	93.3%	83.6%	80.8%	0.798	0.459	94.0%	77.6%
with Adaptive Projector	66.5M(+2.1M)	0.233	0.426	87.8%	93.0%	84.0%	81.2%	0.794	0.454	93.8%	77.3%
with only LoRAs	74.6M(+9.4M)	0.261	0.427	90.3%	94.2%	84.8%	84.0%	0.784	0.441	93.7%	85.6%
LoRAs+Skip-LoRAs	74.6M(+10.2M)	0.264	0.427	90.8%	94.8%	85.1%	84.2%	0.780	0.437	94.5%	87.7%
MindTuner(non-linear)	76.7M(+12.3M)	0.191	0.383	82.7%	88.0%	77.3%	76.0%	0.848	0.492	27.4%	56.7%
MindTuner	76.7M(+12.3M)	0.262	0.422	90.6%	94.9%	85.8%	84.6%	0.774	0.433	<u>94.2%</u>	<u>87.4%</u>

Table 2: Ablations on the modules of MindTuner. Note that all model parameters for MindEye2 are 2.2B, and we only reported the trainable parameters required for new-subject fine-tuning. MinTuner(non-linear) denotes adding an activation function to the ridge regression head.

hour reconstructions can be seen in Figure 5 and Figure 6. It can be seen that the quality of the reconstructed images at 40 hours is significantly higher than at 1 hour as the training data increases. Meanwhile, our MindTuner is better than MindEye2 in both semantic completeness and category accuracy, demonstrating the superiority of the overall model. Our visualization excludes the other three methods in Table 1 because the image semantics generated by these three methods are very unclear. Interestingly, we also found that images generated directly from SDXL unCLIP outperformed MindTuner’s corrected images at a high level, albeit visibly distorted. Further results are in Appendix D.1.

Ablations

In this section, we explored the effectiveness of each component of our method through ablation experiments. All results were pre-trained on Subjects 2-8, and fine-tuning was performed on Subject 1 using 1 hour of data.

MindTuner’s modules. Further experiments were conducted to assess the efficacy of each module. As can be observed in Table 2, the incorporation of LoRAs markedly enhances the model’s capacity. The incorporation of Skip-LoRAs resulted in a convergence of all image reconstruction metrics towards the performance of the complete model, with the retrieval accuracy even exceeding that of the complete model. This can be attributed to the balance of additional image-to-text aligned losses, which has the effect of slightly reducing the performance of the retrieval submodule in the complete model. Furthermore, the incorporation of an adaptive projector has enhanced the performance of high-level reconstruction, and additional visualizations in Figure 4 have also demonstrated the efficacy of semantic correction. However, this approach resulted in the compromise of certain low-level effects. In comparison to the 2.2B multi-subject pre-trained main model, MindTuner resulted in a mere 12.3M increase in parameters (0.56%), yet achieved superior outcomes. Furthermore, an activation function was added to the ridge head to make the entire alignment head non-linear. However, this resulted in a serious overfitting phenomenon, which indicates that the Skip-LoRAs in MindTuner offer a more suitable alternative for non-linear fMRI alignment of diverse subjects. More ablations about the rank r in LoRA refer to Appendix.

Neuroscience Interpretability

To investigate the interpretability of MindTuner in neuroscience, we conduct experiments to explore where the subject’s non-linear relationship comes from in the visual cortex. We utilize pycortex (Gao et al. 2015) to project the weights of each voxel in the first layer of Ridge regression, LoRA, or Skip-LoRA onto the corresponding 2D flat map of the NSD dataset. The results are presented in Figure 7. For linear heads, the importance of visual cortical voxels is similar for both ridge and LoRA. However, for Skip-LoRA, a small portion of the early visual cortex and more advanced visual cortex are valued more. This indicates that at the fMRI level, a greater concentration of non-linear relationships is found within the higher visual cortex. It also demonstrates that the design of Skip-LoRAs is capable of capturing non-linear relationships in fMRI data that are not discernible with standard LoRAs. More visualization results, as well as brain region ROI templates, can be found in Appendix E.




Figure 7: The importance of different ROIs of subject 1 with fitted weights of the first layer. The weights of each module are averaged and normalized between 0 and 1.

Conclusion

In this paper, we propose MindTuner, a new cross-subject decoding method. We introduced the phenomenon of visual fingerprint in the human visual system and utilized the combination of Skip-LoRAs and LoRAs to learn each subject’s visual fingerprint. Meanwhile, we innovatively propose a method for enhancing reconstruction by indirectly connecting fMRI with text in visual decoding tasks. Experimental results have shown that we have achieved better performance, especially when the fMRI data is insufficient. Our work has relaxed the conditions for fMRI acquisition, helping to achieve a universal brain decoding model in the future.

Acknowledgements

We have included more experimental results and analyses in the appendix, which you can access through the extended version (<https://arxiv.org/abs/2404.12630>). This research is supported by the National Key Research and Development Program of China (No. 2022YFB3104700), the National Natural Science Foundation of China (No. 61976158, No. 62376198), Shanghai Baiyulan Pujiang Project (No. 08002360429).

References

- 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. 2022. A massive 7T fMRI dataset to bridge cognitive neuroscience and artificial intelligence. *Nature neuroscience*, 25(1): 116–126.
- Bao, G.; Gong, Z.; Zhang, Q.; Zhou, J.; Fan, W.; Yi, K.; Naseem, U.; Hu, L.; and Miao, D. 2024. Wills Aligner: A Robust Multi-Subject Brain Representation Learner. *arXiv preprint arXiv:2404.13282*.
- Bazeille, T.; Richard, H.; Janati, H.; and Thirion, B. 2019. Local optimal transport for functional brain template estimation. In *Information Processing in Medical Imaging: 26th International Conference, IPMI 2019, Hong Kong, China, June 2–7, 2019, Proceedings 26*, 237–248. Springer.
- Chen, J.; Leong, Y. C.; Honey, C. J.; Yong, C. H.; Norman, K. A.; and Hasson, U. 2017. Shared memories reveal shared structure in neural activity across individuals. *Nature neuroscience*, 20(1): 115–125.
- Chen, Z.; Qing, J.; Xiang, T.; Yue, W. L.; and Zhou, J. H. 2023. Seeing beyond the brain: Conditional diffusion model with sparse masked modeling for vision decoding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 22710–22720.
- Ferrante, M.; Boccato, T.; and Toschi, N. 2023. Through their eyes: multi-subject Brain Decoding with simple alignment techniques. *arXiv preprint arXiv:2309.00627*.
- Fingelkurts, A. A.; Fingelkurts, A. A.; and Kähkönen, S. 2005. Functional connectivity in the brain—is it an elusive concept? *Neuroscience & Biobehavioral Reviews*, 28(8): 827–836.
- Gao, J. S.; Huth, A. G.; Lescroart, M. D.; and Gallant, J. L. 2015. Pycortex: an interactive surface visualizer for fMRI. *Frontiers in neuroinformatics*, 9: 23.
- Gong, Z.; Bao, G.; Zhang, Q.; Wan, Z.; Miao, D.; Wang, S.; Zhu, L.; Wang, C.; Xu, R.; Hu, L.; et al. 2024a. NeuroClips: Towards High-fidelity and Smooth fMRI-to-Video Reconstruction. *arXiv preprint arXiv:2410.19452*.
- Gong, Z.; Zhang, Q.; Bao, G.; Zhu, L.; Zhang, Y.; Liu, K.; Hu, L.; and Miao, D. 2024b. Lite-Mind: Towards Efficient and Robust Brain Representation Learning. In *Proceedings of the 32nd ACM International Conference on Multimedia*, MM ’24, 4014–4023. New York, NY, USA: Association for Computing Machinery. ISBN 9798400706868.
- Kim, S.; Lee, G.; Bae, S.; and Yun, S.-Y. 2020. Mixco: Mix-up contrastive learning for visual representation. *arXiv preprint arXiv:2010.06300*.
- Li, J.; Li, D.; Savarese, S.; and Hoi, S. 2023. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, 19730–19742. PMLR.
- Lin, S.; Sprague, T.; and Singh, A. K. 2022. Mind reader: Reconstructing complex images from brain activities. *Advances in Neural Information Processing Systems*, 35: 29624–29636.
- Linden, D. 2021. Section 3 - Introduction. In Hampson, M., ed., *fMRI Neurofeedback*, 161–169. Academic Press. ISBN 978-0-12-822421-2.
- Loshchilov, I.; and Hutter, F. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Lu, Y.; Du, C.; Zhou, Q.; Wang, D.; and He, H. 2023. MindDiffuser: Controlled Image Reconstruction from Human Brain Activity with Semantic and Structural Diffusion. In *Proceedings of the 31st ACM International Conference on Multimedia*, 5899–5908.
- Mai, W.; and Zhang, Z. 2023. Unibrain: Unify image reconstruction and captioning all in one diffusion model from human brain activity. *arXiv preprint arXiv:2308.07428*.
- Nie, X.; Hu, B.; Gao, X.; Li, L.; Zhang, X.; and Xiao, B. 2023. BMI-Net: A Brain-inspired Multimodal Interaction Network for Image Aesthetic Assessment. In *Proceedings of the 31st ACM International Conference on Multimedia*, MM ’23, 5514–5522. New York, NY, USA: Association for Computing Machinery. ISBN 9798400701085.
- Ozcelik, F.; and VanRullen, R. 2023. Brain-diffuser: Natural scene reconstruction from fmri signals using generative latent diffusion. *arXiv preprint arXiv:2303.05334*.
- Qian, X.; Wang, Y.; Huo, J.; Feng, J.; and Fu, Y. 2023. fmri-pte: A large-scale fmri pretrained transformer encoder for multi-subject brain activity decoding. *arXiv preprint arXiv:2311.00342*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.
- Ramesh, A.; Dhariwal, P.; Nichol, A.; Chu, C.; and Chen, M. 2022. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.
- Scotti, P. S.; Banerjee, A.; Goode, J.; Shabalin, S.; Nguyen, A.; Cohen, E.; Dempster, A. J.; Verlinde, N.; Yundler, E.; Weisberg, D.; et al. 2023. Reconstructing the Mind’s Eye: fMRI-to-Image with Contrastive Learning and Diffusion Priors. *arXiv preprint arXiv:2305.18274*.
- Scotti, P. S.; Tripathy, M.; Villanueva, C. K. T.; Kneeland, R.; Chen, T.; Narang, A.; Santhirasegaran, C.; Xu, J.; Nessler, T.; Norman, K. A.; et al. 2024. MindEye2: Shared-Subject Models Enable fMRI-To-Image With 1 Hour of Data. *arXiv preprint arXiv:2403.11207*.

- Smith, L. N.; and Topin, N. 2019. Super-convergence: Very fast training of neural networks using large learning rates. In *Artificial intelligence and machine learning for multi-domain operations applications*, volume 11006, 369–386. SPIE.
- Stringer, C.; Pachitariu, M.; Steinmetz, N.; Carandini, M.; and Harris, K. D. 2019. High-dimensional geometry of population responses in visual cortex. *Nature*, 571(7765): 361–365.
- Takagi, Y.; and Nishimoto, S. 2023. 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*, 14453–14463.
- Thual, A.; Benchetrit, Y.; Geilert, F.; Rapin, J.; Makarov, I.; Banville, H.; and King, J.-R. 2023. Aligning brain functions boosts the decoding of visual semantics in novel subjects. *arXiv preprint arXiv:2312.06467*.
- Wang, J.; Pohlmeyer, E.; Hanna, B.; Jiang, Y.-G.; Sajda, P.; and Chang, S.-F. 2009. Brain state decoding for rapid image retrieval. In *Proceedings of the 17th ACM International Conference on Multimedia*, MM '09, 945–954. New York, NY, USA: Association for Computing Machinery. ISBN 9781605586083.
- Wang, J.; Yang, Z.; Hu, X.; Li, L.; Lin, K.; Gan, Z.; Liu, Z.; Liu, C.; and Wang, L. 2022. Git: A generative image-to-text transformer for vision and language. *arXiv preprint arXiv:2205.14100*.
- Wang, S.; Liu, S.; Tan, Z.; and Wang, X. 2024. Mindbridge: A cross-subject brain decoding framework. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11333–11342.
- Wang, Z.; Murai, Y.; and Whitney, D. 2020. Idiosyncratic perception: a link between acuity, perceived position and apparent size. *Proceedings of the Royal Society B*, 287(1930): 20200825.
- Xia, W.; de Charette, R.; Öztireli, C.; and Xue, J.-H. 2024a. Dream: Visual decoding from reversing human visual system. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 8226–8235.
- Xia, W.; de Charette, R.; Öztireli, C.; and Xue, J.-H. 2024b. UMBRAE: Unified Multimodal Decoding of Brain Signals. *arXiv preprint arXiv:2404.07202*.
- Zhang, L.; Rao, A.; and Agrawala, M. 2023. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 3836–3847.
- Zhao, S.; Jiang, X.; Han, J.; Hu, X.; Zhu, D.; Lv, J.; Zhang, T.; Guo, L.; and Liu, T. 2014. Decoding Auditory Saliency from fMRI Brain Imaging. In *Proceedings of the 22nd ACM International Conference on Multimedia*, MM '14, 873–876. New York, NY, USA: Association for Computing Machinery. ISBN 9781450330633.