

# BRAINGUARD: Privacy-Preserving Multisubject Image Reconstructions from Brain Activities

Zhibo Tian<sup>1</sup>, Ruijie Quan<sup>2</sup>, Fan Ma<sup>3</sup>, Kun Zhan<sup>1\*</sup>, Yi Yang<sup>3</sup>

<sup>1</sup>School of Information Science and Engineering, Lanzhou University

<sup>2</sup>College of Computing and Data Science, Nanyang Technological University

<sup>3</sup>College of Computer Science and Technology, Zhejiang University  
kzhan@lzu.edu.cn

## Abstract

Reconstructing perceived images from human brain activity forms a crucial link between human and machine learning through Brain-Computer Interfaces. Early methods primarily focused on training separate models for each individual to account for individual variability in brain activity, overlooking valuable cross-subject commonalities. Recent advancements have explored multisubject methods, but these approaches face significant challenges, particularly in data privacy and effectively managing individual variability. To overcome these challenges, we introduce BRAINGUARD, a privacy-preserving collaborative training framework designed to enhance image reconstruction from multisubject fMRI data while safeguarding individual privacy. BRAINGUARD employs a collaborative global-local architecture where individual models are trained on each subject’s local data and operate in conjunction with a shared global model that captures and leverages cross-subject patterns. This architecture eliminates the need to aggregate fMRI data across subjects, thereby ensuring privacy preservation. To tackle the complexity of fMRI data, BRAINGUARD integrates a hybrid synchronization strategy, enabling individual models to dynamically incorporate parameters from the global model. By establishing a secure and collaborative training environment, BRAINGUARD not only protects sensitive brain data but also improves the image reconstructions accuracy. Extensive experiments demonstrate that BRAINGUARD sets a new benchmark in both high-level and low-level metrics, advancing the state-of-the-art in brain decoding through its innovative design.

## 1 Introduction

Deciphering brain activities and recovering encoded information represents a fundamental objective in the field of cognitive neuroscience (Chen et al. 2023; Bullmore and Sporns 2009; Parthasarathy et al. 2017; Belyi et al. 2019). Functional Magnetic Resonance Imaging (fMRI) data (Kwong et al. 1992) is commonly used to recover visual information and plays a crucial role in delineating brain activity by observing fluctuations in blood oxygenation level. Initially, methods for interpreting brain activity from fMRI mainly involved image classification (Kamitani and Tong 2005; Cox and Savoy 2003). With the integration of deep generative

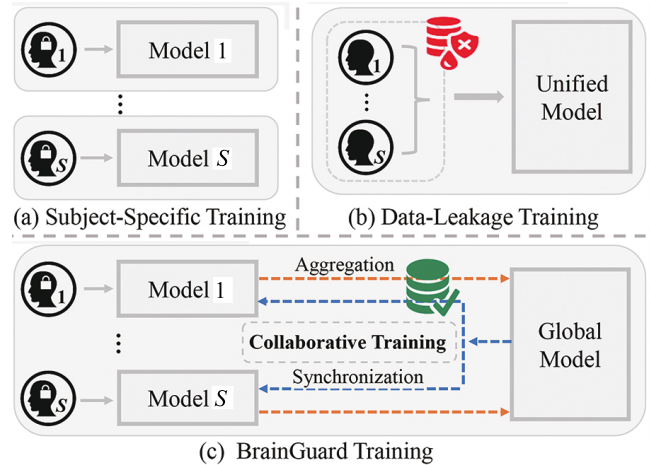


Figure 1: (a) Early subject-specific methods require separate training for each individual using their respective fMRI, overlooking intersubject commonalities. (b) Recent multisubject methods that combine all subjects’ fMRI for training pose substantial privacy concerns. (c) BRAINGUARD captures intersubject commonalities while preserving data privacy.

models, the focus has shifted towards more sophisticated fMRI-to-Image (f2I) reconstruction approaches (Shen et al. 2019; Belyi et al. 2019). However, the inherent variability in fMRI across individuals presents significant challenges, as individuals exhibit unique brain activation patterns in response to the same visual stimulus (Chen et al. 2023).

Earlier f2I reconstruction methods employed a subject-specific approach, where distinct models are developed for each individual’s fMRI to address the significant variance in brain activity across different subjects (Fig. 1(a)) (Lin, Sprague, and Singh 2022; Scotti et al. 2023; Gu et al. 2023; Chen et al. 2023; Takagi and Nishimoto 2023). However, this approach is inherently limited by its focus on individual differences, which results in the neglect of valuable insights that could be obtained from commonalities across multisubject fMRI. Recent advancements increasingly focus on multisubject or cross-subject methods (Fig. 1(b)) (Quan et al. 2024; Scotti et al. 2024; Wang et al. 2024), which aims to train a unified model capable of reconstructing images from the fMRI of multiple individuals. These methods incorporate simple

\*Corresponding Author.

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




Figure 2: (a) t-SNE visualization of the embeddings output from the subject-specific layers in MindBridge (Wang et al. 2024). (b) represents ones from BRAINGUARD models. Both are visualized upon the NSD (Allen et al. 2022) test set.

subject-specific layers, such as ridge regression, to project fMRI of different subjects into a common embedding space, thus accounting for individual differences while training on combined dataset.

However, two primary challenges remain. **First**, exiting multisubject methods that combine all subjects’ fMRI for training poses substantial privacy concerns. Brain data is highly personal and confidential, and aggregating it without robust privacy protections is problematic, especially considering that the data often originates from multiple organizations. **Second**, the subject-specific layers employed in these methods are inadequate for fully addressing the complex differences in fMRI among individuals. As illustrated in Fig. 2(a), the embedding output from the subject-specific layers in MindBridge (Wang et al. 2024) fails to effectively map the fMRI to a common embedding space. These simplistic mappings struggle to capture the distinctive neural patterns of each subject, thus limiting the model’s ability to generalize effectively across diverse data.

In response to these challenges, we introduce BRAINGUARD, the first privacy-preserving collaborative training framework designed for multisubject f2I reconstruction (Fig. 1(c)). BRAINGUARD adopts a global-local model, consisting of individual models for each subject and a global model that captures and leverages shared patterns across subjects. During the training phase, the global model integrates the updated parameters from each individual model, which are independently trained on their respective subject’s fMRI, and then the individual models dynamically synchronize the merged parameters from the global. This design promotes a distributed and collaborative learning mechanism that balances personalization with generalization, thereby enhancing reconstruction quality while maintaining data privacy.

Furthermore, unlike typical image and text, fMRI reflects complex neural activity and is typically preprocessed into a 1D vector of high-dimensional voxels. This complexity renders standard data processing methods less effective. To address this, we introduce a **hybrid synchronization strategy** in BRAINGUARD, specifically designed to meet the unique challenges of fMRI. This strategy synchronizes different parts of the model layers in a hybrid manner, effectively capturing both the distinct neural activity of individual subjects and the shared patterns across subjects by leveraging the specialized

capabilities of different model layers. Specifically, it consists of three parts: ❶ retention for foundational layers, ❷ global alignment for intermediate layers, and ❸ adaptive tuning for advanced layers. The foundational layers retain the subject-specific features inherent in fMRI, preserving the individual variability that is crucial for accurate modeling. Then, the intermediate layers are synchronized with the global model to integrate shared insights, which helps generalize across different subjects. Finally, the advanced layers are adaptively tuned using a Dynamic Fusion Learner (DFL) module, which selectively incorporates global model parameters into the individual models. This hybrid strategy is specifically designed to balance individual uniqueness and commonalities, thereby enhancing the model’s generalization ability across diverse subjects while safeguarding the privacy of sensitive data.

BRAINGUARD enjoys several attractive qualities: **First**, BRAINGUARD adeptly balances individual uniqueness and commonalities across subjects, enhancing the model’s generalization capabilities (Fig. 2(b)) while preserving the specificity of individual subject data. **Second**, the hybrid synchronization strategy within BRAINGUARD ensures that individual models dynamically incorporate broader dataset patterns, thereby improving the precision and personalization of image reconstructions. **Third**, BRAINGUARD ensures the privacy and security of sensitive fMRI by facilitating a collaborative learning framework, allowing for collaborative training without the need for direct data sharing, thus upholding stringent data confidentiality standards.

In summary, our contributions are threefold:

- To the best of our knowledge, this is the first method to propose a privacy-preserving approach for reconstructing images from fMRI. We introduce BRAINGUARD, a framework that captures both intersubject commonalities and individual variabilities while preserving data privacy.
- BRAINGUARD incorporates a hybrid synchronization strategy, enabling each individual model to align with its subject’s unique objectives and dynamically assimilate and apply relevant patterns identified by the global model.
- Experiments show BRAINGUARD’s efficacy and adaptability, highlighting its potential to significantly advance the field of brain decoding while preserving privacy.

## 2 Related Work

### 2.1 Image Reconstruction from fMRI

Earlier f2I reconstruction studies relied on handcrafted features to translate fMRI into images (Kay et al. 2008; Miyawaki et al. 2008; Nishimoto et al. 2011). Subsequent advances utilized paired fMRI images and sparse linear regression, significantly improving the association between brain activity and visual imagery. In recent years, researchers made notable progress in reconstructing images from fMRI by mapping brain signals to the latent space of generative adversarial networks (GANs) (Lin, Sprague, and Singh 2022; Karras et al. 2020). The introduction of multimodal vision-language models (Radford et al. 2021; Lu et al. 2019; Ramesh et al. 2021; Li et al. 2023; Zhao et al. 2024; Yang et al. 2024), diffusion models (Ho, Jain, and Abbeel 2020; Sohl-Dickstein et al.

2015; Song and Ermon 2019; Song et al. 2020; Rombach et al. 2022), and large-scale fMRI datasets (Van Essen et al. 2013; Horikawa and Kamitani 2017; Chang et al. 2019; Allen et al. 2022) elevated the quality of image reconstruction from fMRI to unprecedented levels (Mai et al. 2023). These diffusion model-based methods (Ozcelik and VanRullen 2023; Scotti et al. 2023) map fMRI signals to CLIP text and image embeddings using subject-specific ridge regression or MLPs, followed by a diffusion model that integrates multiple inputs. Recent advances shifted towards the cross-subject method (Quan et al. 2024; Scotti et al. 2024; Wang et al. 2024; Gong et al. 2024), which aims to generalize models across various subjects.

Although effective, current multisubject methods primarily utilize simple subject-specific layers, such as ridge regression, which are inadequate for capturing individual variability. Moreover, these methods aggregate data from all subjects, raising significant privacy concerns. Consequently, there is a critical need for models that achieve a balance between personalization and shared representations while ensuring privacy through decentralized approaches. Addressing these challenges serves as the motivation for this study.

## 2.2 Federated Learning

Federated learning (FL) addresses the challenges of statistical heterogeneity in distributed systems. Traditional methods like FedAvg (McMahan et al. 2017), aim to build a single global model for all clients. However, these methods often fail to meet the specific needs of individual clients due to the diverse nature of their data. Personalized Federated Learning (pFL) shifts the focus to creating models tailored to each client, thereby enhancing performance and relevance (Zhang et al. 2023). Recent advancements in pFL include regularizing local loss functions to prevent overfitting (Li et al. 2020; Yao and Sun 2020; Li, He, and Song 2021), fine-tuning and transfer learning (Fallah, Mokhtari, and Ozdaglar 2020; Chen et al. 2020; Yang et al. 2020), and applying meta-learning techniques (Li et al. 2020; Yao and Sun 2020; Li, He, and Song 2021). Knowledge distillation has also been used to transfer knowledge from a global model to personalized models, improving accuracy and efficiency (Li and Wang 2019; Zhu, Hong, and Zhou 2021; Lin et al. 2020; He, Annavaram, and Avestimehr 2020). In addition, clustering groups clients with similar data distributions, facilitating more effective personalization and reducing computational overhead (Sattler, Müller, and Samek 2020; Briggs, Fan, and Andras 2020; Ghosh et al. 2020). These methodologies ensure that individual client needs are met while maintaining the fundamental principles of data privacy and security in FL.

This work represents the initial exploration of privacy and security protection for sensitive fMRI data in the context of brain activities image reconstruction. Unlike conventional image or text data, fMRI captures intricate neural activity and is typically preprocessed into high-dimensional representations. To address these challenges, we propose a hybrid synchronization strategy tailored to the unique characteristics of fMRI data, improving the model’s capacity to learn both individual-specific and common features.

## 3 Methodology

### Task Setup and Notations

Our goal is to improve image reconstruction by leveraging commonalities across multisubject fMRI while maintaining subject privacy. The original acquired fMRI is four-dimensional (3D space+time). Typically, the data is averaged across the time dimension, resulting in a three-dimensional representation, which is then flattened into a one-dimensional vector of voxels corresponding to the visual stimuli presented to a healthy subject. Let  $\mathbf{X}_{s,n} \in \mathbb{R}^d$  represent the extracted fMRI when a RGB image (visual stimuli) is presented to the subject  $s \in \{1, \dots, S\}$ , where  $d$  is the number of voxels, and  $n \in \{1, \dots, N\}$  denotes the number of images. During the training phase,  $\mathbf{X}_{s,n}$  is transformed into a well-aligned CLIP (Radford et al. 2021) embedding space through the integration of visual and textual supervision. The CLIP embeddings of the ground-truth image and its corresponding caption text are denoted by  $\mathbf{I}_{s,n} \in \mathbb{R}^{v \times c}$  and  $\mathbf{T}_{s,n} \in \mathbb{R}^{t \times c}$ , respectively. Here,  $v$  and  $t$  represent the number of tokens in the CLIP image and text embeddings, while  $c$  indicates their dimensions. During inference, the encoded fMRI embeddings,  $\tilde{\mathbf{I}}_{s,n}$  and  $\tilde{\mathbf{T}}_{s,n}$ , serve as conditional inputs to a versatile diffusion model (Xu et al. 2023), which facilitates image reconstruction.

### Method Overview

BRAINGUARD requires only a single training session using multisubject data by introducing a collaborative training framework, consisting of individual models and a global model (§3.1). These models are trained in a collaborative approach, where individual and global models engage in bidirectional parameter fusion during their training process. Specifically, the individual models are optimized not only through their respective subject’s fMRI training objectives but also via integration with the global model. Conversely, the update of the global model’s parameters is informed by the amalgamation of multiple individual models. To achieve the dynamic parameter fusion process, we devise a hybrid synchronization strategy (§3.2). An overview of our framework is illustrated in Fig. 3, and the detailed network architecture and learning objectives is presented in §3.3.

### 3.1 Privacy-Preserving Collaborative Training

BRAINGUARD implements a privacy-preserving collaborative training approach designed to protect the privacy of subjects’ fMRI while simultaneously leveraging commonalities across different subjects. The training process is cyclical and consists of three essential steps: **1) Individual Training:** At the beginning of each training cycle, BRAINGUARD trains a distinct model for each subject using their specific fMRI. This approach ensures that sensitive neural data remains localized, thereby eliminating the need for centralization. The individual model, denoted as  $f_s$ , is optimized to predict the CLIP image embedding  $\tilde{\mathbf{I}}_{s,n}$  and the text embedding  $\tilde{\mathbf{T}}_{s,n}$  based on the subject’s data. **2) Parameters Aggregating:** After each training round, the individual models transmit their updated parameters to a global model for aggregation, which aggregates these parameters to capture shared patterns

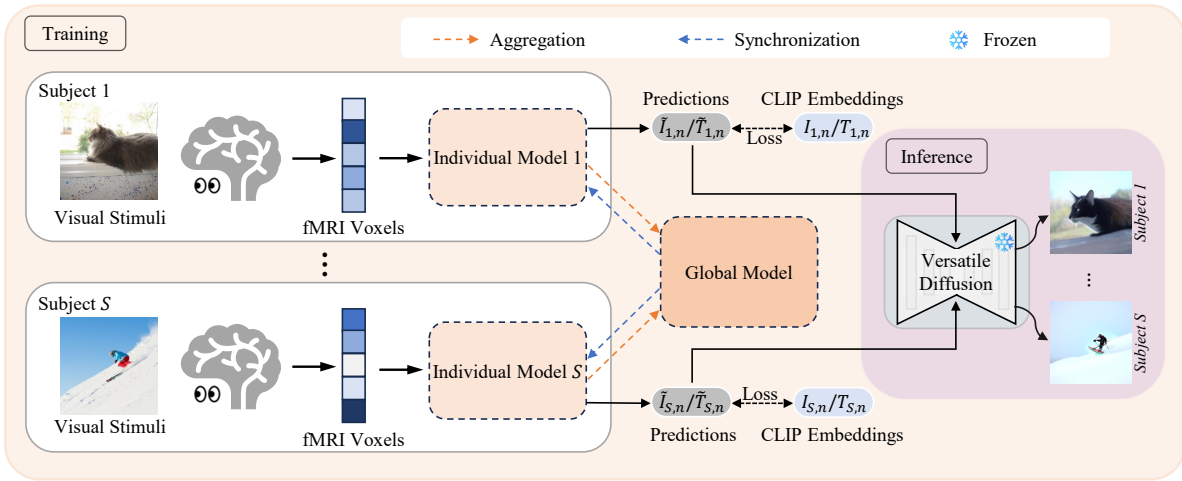


Figure 3: An overview of the BRAINGUARD training and inference framework (§3.1).

across subjects. Before this aggregation, individual models update their parameters using the Exponential Moving Average (EMA):  $\theta'_s = \alpha\theta'_s + (1 - \alpha)\theta_s$ , where  $\alpha$  is the EMA factor, typically set to 0.999. This step aims to smooth parameter changes and enhance individual model stability, ensuring that the subsequent aggregation reflects both refined parameters and shared information. The global model’s parameters  $\theta_g$  is updated using  $\theta_g = \sum_s k_s \theta'_s$ , where  $k_s$  is a weight coefficient representing each subject’s data proportion in the overall dataset, calculated as  $k_s = |\mathbf{X}_{s,n}| / \sum_{i=1}^S |\mathbf{X}_{i,n}|$  where  $|\mathbf{X}_{s,n}|$  denotes the total number of fMRI for subject  $s$ . This aggregation process enables the global model to effectively capture and integrate shared patterns across subjects. **3) Model Synchronizing:** After the global model is updated, individual models are refined by incorporating the aggregated knowledge from the global through a hybrid synchronization strategy (§3.2). This process ensures that each individual benefits from the shared knowledge while preserving its unique adaptations to the specific subject data, thereby achieving a balance between generalization and personalization.

### 3.2 Hybrid Synchronization Strategy

Unlike typical image and text data, fMRI reflects complex neural activity and is typically preprocessed into a one-dimensional vector comprising high-dimensional voxels. This inherent complexity makes standard data processing methods less effective. To address this challenge, inspired by FedALA (Zhang et al. 2023), we employ a hybrid synchronization strategy that balances individual uniqueness with shared insights through the integration of individual and global models. It is implemented through three distinct approaches, each targeting specific layers within the model to optimize different aspects of the learning process based on the unique characteristics of fMRI.

**Retention for foundational layers.** The fMRI is characterized by highly individualized brain activity patterns, in the lower layers that capture foundational neural signals. These patterns are unique to each subject and vary significantly depending on experimental conditions. To preserve these individualized features, the lower layers of the local models

retain their original parameters without updates from the global model. This approach ensures that each local model remains finely tuned to its specific data distribution, improving the accuracy of personalized interpretations of brain activity.

**Global alignment for intermediate layers.** Despite the individualized nature of foundational layers, intermediate layers of fMRI processing often reveal commonalities across different subjects and experimental conditions, such as shared patterns in brain connectivity. These layers benefit from global insights, which help in aligning the models across different datasets. In this strategy, the intermediate layers of the local models are updated by directly adopting parameters from the corresponding layers in the global model. This promotes a unified feature extraction process that captures the shared characteristics of fMRI, enhancing the model’s ability to generalize across diverse subjects and conditions.

**Adaptive tuning for advanced layers.** The advanced layers of the model are tasked with interpreting complex, high-level features related to cognitive processes and brain network interactions. These layers must balance the nuanced variations in fMRI with broader patterns. To achieve this, we introduce a Dynamic Fusion Learner (DFL) module (Zhang et al. 2023), which adaptively tunes the advanced layers’ parameters. This tuning is crucial for integrating both global patterns and local variations, ensuring the model can accurately reflect the cognitive and neural processes underlying the fMRI, thereby improving overall performance and generalization. Together, these strategies form a cohesive hybrid synchronization approach that effectively balances the unique and shared aspects of individual fMRI, optimizing both individual model performance and global model robustness.

**Dynamic Fusion Learner Module.** Since the advanced layers of a network generally handle the most abstract high-level features, including complex patterns that are highly task-specific (Yosinski et al. 2014; LeCun, Bengio, and Hinton 2015), they are critical for fine-tuning in the context of BRAINGUARD. The distinct characteristics of fMRI further emphasize the importance of these advanced layers for accurately interpreting and responding to individual brain activity. In BRAINGUARD, we employ a Dynamic Fusion Learner

(DFL) module, which performs layer-wise dynamic aggregation of the global and individual model parameters in the advanced layers rather than completely overwriting them. DFL is strategically applied to the advanced  $p$  layers, fine-tuning them to reflect individual specificities. This approach seeks to balance the general patterns with the personalized aspects of the individual model through targeted adjustments.

Formally, the parameters of the individual model,  $\theta_s$ , at iteration  $t$ , are updated through DFL as follows:

$$\theta_s^t := \theta_s^{t-1} + (\theta_g^{t-1} - \theta_s^{t-1}) \odot W_s, \quad (1)$$

where  $\odot$  denotes the Hadamard product.  $(\theta_g^{t-1} - \theta_s^{t-1})$  accounts for the difference between the global and the individual, while  $W_s$  represents the aggregation weights between the global and the individual for specific layer parameters. We apply element-wise weight clipping (Courbariaux et al. 2016; Zhang et al. 2023), defined as  $\tau(w) = \max(0, \min(1, w))$ , ensuring that  $w \in [0, 1]$  for all  $w \in W_s$ . When  $\forall w_s \in W_s, w_s \equiv 1$ , it indicates a scenario where the global completely overwrites the individual models' parameters.

Since we focus exclusively on the advanced  $m$  layers in the model, Eq. (1) is given by:

$$\theta_s^t := \theta_s^{t-1} + (\theta_g^{t-1} - \theta_s^{t-1}) \odot [\mathbf{1}^{|\theta_s|-m}; W_s^m], \quad (2)$$

where  $|\theta_s|$  denotes the total number of layers in the model, and  $\mathbf{1}^{|\theta_s|-m}$  is an array of ones corresponding to the foundational and intermediate layers.  $W_s^m$  represents the aggregation weights for the advanced  $m$  layers, allowing the model to adapt to subject-specific features of higher level. Initially, each element in  $W_s^m$  is set to one and is updated iteratively. The  $W_s^m$  is optimized through a gradient-based learning method:

$$W_s^m \leftarrow W_s^m - \eta \frac{\partial \mathcal{L}(\theta_s; \theta_g)}{\partial W_s^m}, \quad (3)$$

where  $\eta$  is the learning rate for weight learning.

By employing DFL, BRAINGUARD adeptly balances the exploration of intersubject commonalities with the delineation of individual specificities. This balance is crucial for neuroscientific studies and clinical applications where understanding shared brain functions is as important as recognizing subject-specific patterns. This methodology highlights the framework's potential to foster a new paradigm for collaborative yet privacy-preserving neuroimaging studies, wherein collective insights are obtained without compromising individual data confidentiality.

### 3.3 Network Architecture and Training Objectives

**Network Architecture.** The network architecture for both the individual and global models is uniform, comprising an initial linear layer, several residual blocks, and two final linear layers (corresponding to the three different synchronization parts). The model projects the fMRI voxels to a shared CLIP latent space, enabling semantic interpretation. During inference, the predicted fMRI embeddings are input into a SOTA diffusion model (Xu et al. 2023), which guides the image reconstruction process. Both the individual and global models can be used to produce the resultant fMRI embeddings.

**Training Objectives.** In practice, BRAINGUARD is optimized by incorporating fMRI as an additional modality,

aiming to align the fMRI-derived embeddings more closely with the CLIP space. Specifically, the individual models employ two types of loss functions to learn CLIP image and text embeddings. The first loss function is the mean squared error, which ensures the accurate prediction of CLIP embeddings:

$$\mathcal{L}_{\text{MSE}}(\mathbf{P}, \mathbf{Y}) = \frac{1}{B} \sum_{i=1}^B (\mathbf{p}_i - \mathbf{y}_i)^2, \quad (4)$$

where  $\mathbf{p}$  is the predicted CLIP embedding by the individual model, and  $\mathbf{y}$  is the target embedding in a batch size of  $B$ . The second loss function, SoftCLIP loss (Scotti et al. 2023), facilitates the alignment of fMRI and CLIP embeddings by leveraging contrastive learning to increase positive pair similarity and decrease negative pair similarity. Positive pairs are identified using soft labels derived from the dot product of embeddings within a batch.

$$\mathcal{L}_{\text{SoftCLIP}}(\mathbf{P}, \mathbf{Y}) = - \sum_{i=1}^B \sum_{j=1}^B \frac{\exp\left(\frac{\mathbf{y}_i \cdot \mathbf{y}_j}{\tau}\right)}{\sum_{k=1}^B \exp\left(\frac{\mathbf{y}_i \cdot \mathbf{y}_k}{\tau}\right)} \cdot \log\left(\frac{\exp\left(\frac{\mathbf{p}_i \cdot \mathbf{y}_j}{\tau}\right)}{\sum_{k=1}^B \exp\left(\frac{\mathbf{p}_i \cdot \mathbf{y}_k}{\tau}\right)}\right), \quad (5)$$

where  $\tau$  denotes the temperature hyperparameter.

The complete set of losses  $\mathcal{L}_s^V$  and  $\mathcal{L}_s^T$  for predicting image and text CLIP embeddings includes:

$$\mathcal{L}_s^V = \mathcal{L}_{\text{MSE}}(\tilde{\mathbf{I}}_s, \mathbf{I}_s) + \mathcal{L}_{\text{SoftCLIP}}(\tilde{\mathbf{I}}_s, \mathbf{I}_s), \quad (6)$$

$$\mathcal{L}_s^T = \mathcal{L}_{\text{MSE}}(\tilde{\mathbf{T}}_s, \mathbf{T}_s) + \mathcal{L}_{\text{SoftCLIP}}(\tilde{\mathbf{T}}_s, \mathbf{T}_s). \quad (7)$$

The main objective of the entire framework is to minimize the global objective function, balancing the performance of individual models with the integration of commonalities. This is achieved by optimizing both image and text prediction models across all subjects as follows:

$$\{(\theta_s^*)\}_{s=1}^S = \arg \min_{\theta_s} \sum_{s=1}^S (\mathcal{L}_s^V + \mathcal{L}_s^T). \quad (8)$$

## 4 Experiments

### Experimental Setup

**Datasets.** The Natural Scenes Dataset (NSD) (Allen et al. 2022) encompasses fMRI obtained from eight participants, who were exposed to a total of 73,000 RGB images. This dataset has been extensively employed in numerous studies (Lin, Sprague, and Singh 2022; Chen et al. 2023; Takagi and Nishimoto 2023; Gu et al. 2023; Scotti et al. 2023) for the purpose of reconstructing images perceived during fMRI. Following (Scotti et al. 2023), our study utilizes data from subjects 1, 2, 5, and 7, who successfully completed all predetermined trials, and these subjects viewed a set of 10,000 natural scene images, each presented three times. The training set for each subject comprises 8,859 image stimuli and 24,980 fMRI trials, while the test set includes 982 image stimuli and 2,770 fMRI trials. All images and their accompanying captions were derived from MS-COCO (Lin et al. 2014). It is different from previous methods which employed all subject data to train an omnifit model, our BRAINGUARD maintains data privacy and simultaneously enables each network to leverage insights from others.

Methods	Low-Level				High-Level			
	PixCorr $\uparrow$	SSIM $\uparrow$	Alex(2) $\uparrow$	Alex(5) $\uparrow$	Incept. $\uparrow$	CLIP $\uparrow$	EffNet-B $\downarrow$	SwAV $\downarrow$
Mind-Reader <small>NeurIPS 2022</small>	—	—	—	—	78.2%	—	—	—
Mind-Vis <small>CVPR 2023</small>	.080	.220	72.1%	83.2%	78.8%	76.2%	.854	.491
Takagi et al. <small>CVPR 2023</small>	—	—	83.0%	83.0%	76.0%	77.0%	—	—
Gu et al. <small>MIDL 2023</small>	.150	.325	—	—	—	—	.862	.465
Brain-Diffuser <small>arXiv 2023</small>	.254	<b>.356</b>	94.2%	96.2%	87.2%	91.5%	.775	.423
MindEye <small>NeurIPS 2023</small>	.309	.323	<b>94.7%</b>	<b>97.8%</b>	93.8%	94.1%	.645	.367
MindBridge <small>CVPR 2024</small>	.148	.259	86.9%	95.3%	92.2%	94.3%	.713	.413
Psychometry <small>CVPR 2024</small>	.295	.328	<u>94.5%</u>	<u>96.8%</u>	<u>94.9%</u>	<u>95.3%</u>	<u>.632</u>	<u>.361</u>
<b>BRAINGUARD (ours)</b>	<b>.313</b>	<b>.330</b>	<b>94.7%</b>	<b>97.8%</b>	<b>96.1%</b>	<b>96.4%</b>	<b>.624</b>	<b>.353</b>

Table 1: **Quantitative comparison results on NSD test dataset between BRAINGUARD and previous SOTA methods.** The best result is highlighted in bold, and the second-best result is underlined. Note that all methods employ a *per-subject-per-model* approach for the final inference process, and all metrics are calculated as the average across four subjects.



Figure 4: **Qualitative comparisons on the NSD test dataset.** BRAINGUARD performs a single training session on multisubject fMRI data, demonstrates superior reconstruction accuracy compared to four recent state-of-the-art methods (Quan et al. 2024; Wang et al. 2024; Scotti et al. 2023; Ozcelik and VanRullen 2023), while effectively preserving data privacy.

#### 4.1 Comparison

**Evaluation Metrics.** For qualitative evaluation, we visually compare our reconstructed images with the ground-truth images and the results of state-of-the-art methods, as shown in Fig. 4. For quantitative evaluation, we present eight different image quality metrics for both low-level and high-level evaluations, following (Scotti et al. 2023). Low-level image features provide basic information about the visual content and structure of the image, while high-level features capture semantic information, object relationships, and contextual understanding. For low-level evaluation, *PixCorr* represents the pixel-level correlation, and *SSIM* (Wang et al. 2004) is the structural similarity. *Alex(2)* and *Alex(5)* denote comparisons of the second and fifth layers of AlexNet (Krizhevsky, Sutskever, and Hinton 2012), respectively. For high-level evaluation, *Incept.* refers to a comparison of the last pooling layer of Inception-v3 (Szegedy et al. 2016), and *CLIP* is a comparison of the output layer of the CLIP-Vision model (Radford et al. 2021). *EffNet-B* and *SwAV* are distance metrics is from the EfficientNet-B1 (Tan and Le 2019) and SwAV-ResNet50 (Caron et al. 2020) models, respectively.

**Quantitative Results.** We compare BRAINGUARD with seven state-of-the-art methods, namely Mind-Reader (Lin, Sprague, and Singh 2022), Mind-Vis (Chen et al. 2023), Takagi et al. (Takagi and Nishimoto 2023), Gu et al. (Gu et al. 2023), Brain-Diffuser (Ozcelik and VanRullen 2023), Mind-

Eye (Scotti et al. 2023), MindBridge (Wang et al. 2024) and Psychometry (Quan et al. 2024). BRAINGUARD achieves 31.3%, 33.0%, 94.7%, and 97.8% in the low-level metrics of *PixCorr*, *SSIM*, *Alex(2)*, and *Alex(5)*, respectively. For high-level metrics, BRAINGUARD secures 96.1%, 96.4%, 62.4%, and 35.3% in *Incept.*, *CLIP*, *EffNet-B*, and *SwAV*, respectively. It is important to note that the results presented in Table 1 are the averages from four subject-specific models, each trained with data from the respective subject. BRAINGUARD significantly outperforms previous SOTA methods across both low-level and high-level evaluation metrics. For low-level metrics (*PixCorr*, *SSIM*, *Alex(2)*, and *Alex(5)*), it achieves improvements over Mind-Vis by 23.3%, 11.0%, 22.6%, and 14.6%, respectively; and over MindBridge by 16.5%, 7.1%, 7.8%, and 2.5%. Similarly for high-level metrics (*Incept.*, *CLIP*, *EffNet-B*, and *SwAV*), BRAINGUARD surpasses Mind-Vis by 17.3%, 20.2%, 23%, and 13.8%, Brain-diffuser by 8.9%, 4.9%, 15.1%, and 7%, and MindEye by 4.1%, 2.1%, 8.9%, and 6%. These advancements are particularly notable given that our method requires training in a single session.

**Qualitative Results.** As illustrated in Fig. 4, the qualitative outcomes align with the quantitative data, indicating that our methodology yields reconstructions of higher quality and greater realism in comparison to alternative approaches. We compare the reconstructions image with MindEye, Psychom-

<i>Found.</i>	<i>Inter.</i>	<i>Advan.</i>	Low-Level				High-Level			
			PixCorr $\uparrow$	SSIM $\uparrow$	Alex(2) $\uparrow$	Alex(5) $\uparrow$	Incept. $\uparrow$	CLIP $\uparrow$	EffNet-B $\downarrow$	SwAV $\downarrow$
			.163	.238	88.7%	90.9%	83.2%	86.1%	.856	.471
$\checkmark$	$\checkmark$		.237	.287	90.3%	93.8%	90.5%	89.3%	.794	.423
$\checkmark$		$\checkmark$	.261	.298	93.3%	95.9%	93.4%	92.5%	.685	.387
$\checkmark$	$\checkmark$	$\checkmark$	<b>.313</b>	<b>.330</b>	<b>94.7%</b>	<b>97.8%</b>	<b>96.1%</b>	<b>96.4%</b>	<b>.624</b>	<b>.353</b>

Table 2: **Synchronization Strategies Ablation:** foundational (*Found.*), intermediate (*Inter.*), and advanced layers (*Advan.*). The results of the first row are from BRAINGUARD without the hybrid synchronization strategy.




Figure 5: (a) We evaluate our trained global model and compare it with the SOTA (MindBridge). (b) Performance changes with different numbers of participating subjects.

etry, MindBridge, and Brain-Diffuser, our BRAINGUARD maintains a high level of consistency with the visual stimuli in terms of semantics, appearance, and structure. This indicates that our method effectively captures intersubject commonality and individual specificity across subjects, resulting in high-quality image reconstructions from fMRI, and preserving the privacy of individual fMRI.

## 4.2 Diagnostic Experiment

To systematically demonstrate the influence of each component in BRAINGUARD on its performance, a series of ablation studies were conducted using the NSD test set.

**Effective of Hybrid Synchronization Strategy.** We examined the effectiveness of the hybrid synchronization strategy across foundational, intermediate, and upper layers between the global model and individual models. Table 2 displays the experimental results when applying global and individual model synchronization strategies at different levels. For low-level features (such as PixCorr and SSIM), a significant improvement in model performance is observed as we transition to higher layers, particularly when the hybrid synchronization strategy is utilized. For instance, in low-level features, Alex(5) achieves a performance of 97.8%, whereas in high-level features, Inception and CLIP attain performance levels of 96.1% and 96.4%, respectively. Additionally, models like EffNet-B and SwAV exhibit varying degrees of improvement when the hybrid synchronization strategy is applied. These findings suggest that the hybrid synchronization strategy effectively enhances the performance of individual models, especially in the intermediate and upper layers.

**Number of Subjects Participating in Training.** The augmentation in the number of subjects participating in the model training leads to enhanced performance, as each subject can mutually benefit from the shared insights. To empirically ver-

ify this premise, we conducted experiments by incrementally increasing the number of subjects involved in the learning process. As shown in Fig. 5(b), our findings highlight a positive correlation between the number of participating subjects and the improvement in results. In particular, performance peaks when all four subjects are trained collaboratively, substantiating the efficacy of our approach. This phenomenon elucidates that with a greater number of subjects engaged in the training process, while each maintains its distinctiveness, the models are adept at identifying and leveraging the shared patterns across different subjects’ brain activities, thereby increasing their individual performance.

**Intersubject Commonality.** As discussed earlier, BRAINGUARD utilizes a unified training approach on multisubject fMRI, which not only streamlines training costs but also leverages intersubject commonalities. This dual benefit is substantiated through both quantitative and qualitative analyses. Fig. 5(a) presents results comparing our BRAINGUARD, where the foundational layer in the individual model is followed by the global model’s intermediate and advanced layers, with those obtained using MindBridge. These results demonstrate that BRAINGUARD has a superior ability to capture intersubject commonalities compared to MindBridge. Additionally, Fig. 4 reveals the semantic coherence of reconstructions across various subjects when exposed to the same visual stimuli. This consistency highlights BRAINGUARD’s proficiency in capturing shared patterns among different subjects, further establishing its effectiveness in fMRI-based image reconstruction.

## 5 Conclusion and Discussion

We introduce BRAINGUARD, an innovative framework designed to address the challenges of f2I reconstruction, particularly focusing on individual variability and data privacy. By ingeniously integrating individual models for each subject with a global model that captures common and unique patterns of brain activity, BRAINGUARD ensures a collaborative approach to exploit the intersubject commonalities from the fMRI of multiple subjects while preserving their privacy. A pivotal feature of this framework is the hybrid synchronization strategy, which effectively incorporates insights from the commonality model into personalized models, providing a robust approach for refining image reconstructions. The deployment of BRAINGUARD represents a significant advance in brain-computer interfaces, offering a novel and secure method to derive meaningful interpretations from brain.

## Acknowledgments

This work was supported by the Key R&D Program of China under Grant No. 2021ZD0112801, the National Natural Science Foundation of China under Grant Nos. 62176108 and U2336212, the Natural Science Foundation of Qinghai Province of China under No. 2022-ZJ-929, the Science Foundation of National Archives Administration of China under No. 2024-B-006, and the Supercomputing Center of Lanzhou University.

## 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.
- Beliy, R.; Gaziv, G.; Hoogi, A.; Strappini, F.; Golan, T.; and Irani, M. 2019. From voxels to pixels and back: Self-supervision in natural-image reconstruction from fmri. *NeurIPS*, 32.
- Briggs, C.; Fan, Z.; and Andras, P. 2020. Federated learning with hierarchical clustering of local updates to improve training on non-IID data. In *IJCNN*, 1–9.
- Bullmore, E.; and Sporns, O. 2009. Complex brain networks: graph theoretical analysis of structural and functional systems. *Nature Reviews Neuroscience*, 10(3): 186–198.
- Caron, M.; Misra, I.; Mairal, J.; Goyal, P.; Bojanowski, P.; and Joulin, A. 2020. Unsupervised learning of visual features by contrasting cluster assignments. In *NeurIPS*, volume 33, 9912–9924.
- Chang, N.; Pyles, J. A.; Marcus, A.; Gupta, A.; Tarr, M. J.; and Aminoff, E. M. 2019. BOLD5000, a public fMRI dataset while viewing 5000 visual images. *Scientific data*, 6(1): 49.
- Chen, Y.; Qin, X.; Wang, J.; Yu, C.; and Gao, W. 2020. Fed-health: A federated transfer learning framework for wearable healthcare. *IEEE Intelligent Systems*, 35(4): 83–93.
- 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 *CVPR*.
- Courbariaux, M.; Hubara, I.; Soudry, D.; El-Yaniv, R.; and Bengio, Y. 2016. Binarized neural networks: Training deep neural networks with weights and activations constrained to +1 or -1. *arXiv preprint arXiv:1602.02830*.
- Cox, D. D.; and Savoy, R. L. 2003. Functional magnetic resonance imaging (fMRI) “brain reading”: Detecting and classifying distributed patterns of fMRI activity in human visual cortex. *Neuroimage*, 19(2): 261–270.
- Fallah, A.; Mokhtari, A.; and Ozdaglar, A. 2020. Personalized federated learning with theoretical guarantees: A model-agnostic meta-learning approach. In *NeurIPS*, volume 33, 3557–3568.
- Ghosh, A.; Chung, J.; Yin, D.; and Ramchandran, K. 2020. An efficient framework for clustered federated learning. In *NeurIPS*, volume 33, 19586–19597.
- Gong, Z.; Zhang, Q.; Bao, G.; Zhu, L.; Liu, K.; Hu, L.; and Miao, D. 2024. MindTuner: Cross-Subject Visual Decoding with Visual Fingerprint and Semantic Correction. *arXiv preprint arXiv:2404.12630*.
- Gu, Z.; Jamison, K.; Kuceyeski, A.; and Sabuncu, M. 2023. Decoding natural image stimuli from fMRI data with a surface-based convolutional network. In *MIDL*.
- He, C.; Annavaram, M.; and Avestimehr, S. 2020. Group knowledge transfer: Federated learning of large cnns at the edge. In *NeurIPS*, volume 33, 14068–14080.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. In *NeurIPS*, volume 33, 6840–6851.
- Horikawa, T.; and Kamitani, Y. 2017. Generic decoding of seen and imagined objects using hierarchical visual features. *Nature Communications*, 8(1): 15037.
- Kamitani, Y.; and Tong, F. 2005. Decoding the visual and subjective contents of the human brain. *Nature Neuroscience*, 8(5): 679–685.
- Karras, T.; Laine, S.; Aittala, M.; Hellsten, J.; Lehtinen, J.; and Aila, T. 2020. Analyzing and improving the image quality of stylegan. In *CVPR*, 8110–8119.
- Kay, K. N.; Naselaris, T.; Prenger, R. J.; and Gallant, J. L. 2008. Identifying natural images from human brain activity. *Nature*, 452(7185): 352–355.
- Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. Imagenet classification with deep convolutional neural networks. In *NeurIPS*, volume 25.
- Kwong, K. K.; Belliveau, J. W.; Chesler, D. A.; Goldberg, I. E.; Weisskoff, R. M.; Poncelet, B. P.; Kennedy, D. N.; Hoppel, B. E.; Cohen, M. S.; and Turner, R. 1992. Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation. *PNAS*, 89(12): 5675–5679.
- LeCun, Y.; Bengio, Y.; and Hinton, G. 2015. Deep learning. *Nature*, 521(7553): 436–444.
- Li, D.; and Wang, J. 2019. FedMD: Heterogenous federated learning via model distillation. In *NeurIPS Workshop*.
- Li, Q.; He, B.; and Song, D. 2021. Model-contrastive federated learning. In *CVPR*, 10713–10722.
- Li, T.; Sahu, A. K.; Zaheer, M.; Sanjabi, M.; Talwalkar, A.; and Smith, V. 2020. Federated optimization in heterogeneous networks. *Proceedings of Machine Learning and Systems*, 2: 429–450.
- Li, Y.; Quan, R.; Zhu, L.; and Yang, Y. 2023. Efficient multimodal fusion via interactive prompting. In *CVPR*, 2604–2613.
- Lin, S.; Sprague, T.; and Singh, A. K. 2022. Mind reader: Reconstructing complex images from brain activities. *NeurIPS*, 35: 29624–29636.
- Lin, T.; Kong, L.; Stich, S. U.; and Jaggi, M. 2020. Ensemble distillation for robust model fusion in federated learning. In *NeurIPS*, volume 33, 2351–2363.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *ECCV*, 740–755.

- Lu, J.; Batra, D.; Parikh, D.; and Lee, S. 2019. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *NeurIPS*, volume 32.
- Mai, W.; Zhang, J.; Fang, P.; and Zhang, Z. 2023. Brain-Conditional Multimodal Synthesis: A Survey and Taxonomy. *arXiv preprint arXiv:2401.00430*.
- McMahan, B.; Moore, E.; Ramage, D.; Hampson, S.; and y Arcas, B. A. 2017. Communication-efficient learning of deep networks from decentralized data. In *AISTATS*, 1273–1282.
- Miyawaki, Y.; Uchida, H.; Yamashita, O.; Sato, M.-a.; Morito, Y.; Tanabe, H. C.; Sadato, N.; and Kamitani, Y. 2008. Visual image reconstruction from human brain activity using a combination of multiscale local image decoders. *Neuron*, 60(5): 915–929.
- Nishimoto, S.; Vu, A. T.; Naselaris, T.; Benjamini, Y.; Yu, B.; and Gallant, J. L. 2011. Reconstructing visual experiences from brain activity evoked by natural movies. *Current biology*, 21(19): 1641–1646.
- Ozcelik, F.; and VanRullen, R. 2023. Natural scene reconstruction from fMRI signals using generative latent diffusion. *Scientific Reports*, 13(1): 15666.
- Parthasarathy, N.; Batty, E.; Falcon, W.; Rutten, T.; Rajpal, M.; Chichilnisky, E.; and Paninski, L. 2017. Neural networks for efficient bayesian decoding of natural images from retinal neurons. In *NeurIPS*, volume 30.
- Quan, R.; Wang, W.; Tian, Z.; Ma, F.; and Yang, Y. 2024. Psychometry: An omnifit model for image reconstruction from human brain activity. In *CVPR*, 233–243.
- 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 *ICML*, 8748–8763.
- Ramesh, A.; Pavlov, M.; Goh, G.; Gray, S.; Voss, C.; Radford, A.; Chen, M.; and Sutskever, I. 2021. Zero-shot text-to-image generation. In *ICML*, 8821–8831.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *CVPR*, 10684–10695.
- Sattler, F.; Müller, K.-R.; and Samek, W. 2020. Clustered federated learning: Model-agnostic distributed multitask optimization under privacy constraints. *IEEE transactions on neural networks and learning systems*, 32(8): 3710–3722.
- 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. In *NeurIPS*, volume 36.
- 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. 2024. MindEye2: Shared-Subject Models Enable fMRI-To-Image With 1 Hour of Data. *arXiv preprint arXiv:2403.11207*.
- Shen, G.; Horikawa, T.; Majima, K.; and Kamitani, Y. 2019. Deep image reconstruction from human brain activity. *PLoS Computational Biology*, 15(1): e1006633.
- Sohl-Dickstein, J.; Weiss, E.; Maheswaranathan, N.; and Ganguli, S. 2015. Deep unsupervised learning using nonequilibrium thermodynamics. In *ICML*, 2256–2265.
- Song, Y.; and Ermon, S. 2019. Generative modeling by estimating gradients of the data distribution. In *NeurIPS*, volume 32.
- Song, Y.; Sohl-Dickstein, J.; Kingma, D. P.; Kumar, A.; Ermon, S.; and Poole, B. 2020. Score-based generative modeling through stochastic differential equations. In *ICLR*.
- Szegedy, C.; Vanhoucke, V.; Ioffe, S.; Shlens, J.; and Wojna, Z. 2016. Rethinking the inception architecture for computer vision. In *CVPR*, 2818–2826.
- Takagi, Y.; and Nishimoto, S. 2023. High-resolution image reconstruction with latent diffusion models from human brain activity. In *CVPR*, 14453–14463.
- Tan, M.; and Le, Q. 2019. Efficientnet: Rethinking model scaling for convolutional neural networks. In *ICML*, 6105–6114.
- Van Essen, D. C.; Smith, S. M.; Barch, D. M.; Behrens, T. E.; Yacoub, E.; Ugurbil, K.; Consortium, W.-M. H.; et al. 2013. The WU-Minn human connectome project: an overview. *Neuroimage*, 80: 62–79.
- Wang, S.; Liu, S.; Tan, Z.; and Wang, X. 2024. Mindbridge: A cross-subject brain decoding framework. In *CVPR*, 11333–11342.
- Wang, Z.; Bovik, A. C.; Sheikh, H. R.; and Simoncelli, E. P. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on Image Processing*, 13(4): 600–612.
- Xu, X.; Wang, Z.; Zhang, G.; Wang, K.; and Shi, H. 2023. Versatile diffusion: Text, images and variations all in one diffusion model. In *CVPR*.
- Yang, H.; He, H.; Zhang, W.; and Cao, X. 2020. FedSteg: A federated transfer learning framework for secure image steganalysis. *IEEE Transactions on Network Science and Engineering*, 8(2): 1084–1094.
- Yang, Z.; Chen, G.; Li, X.; Wang, W.; and Yang, Y. 2024. Doraemongpt: Toward understanding dynamic scenes with large language models (exemplified as a video agent). In *ICML*.
- Yao, X.; and Sun, L. 2020. Continual local training for better initialization of federated models. In *ICIP*, 1736–1740.
- Yosinski, J.; Clune, J.; Bengio, Y.; and Lipson, H. 2014. How transferable are features in deep neural networks? In *NeurIPS*, volume 27.
- Zhang, J.; Hua, Y.; Wang, H.; Song, T.; Xue, Z.; Ma, R.; and Guan, H. 2023. FedALA: Adaptive local aggregation for personalized federated learning. In *AAAI*, volume 37, 11237–11244.
- Zhao, S.; Quan, R.; Zhu, L.; and Yang, Y. 2024. CLIP4STR: A simple baseline for scene text recognition with pre-trained vision-language model. *IEEE transactions on Image Processing*.
- Zhu, Z.; Hong, J.; and Zhou, J. 2021. Data-free knowledge distillation for heterogeneous federated learning. In *ICML*, 12878–12889.