

DREAM: Visual Decoding from REversing HumAn Visual SystemWeihaio Xia¹✉ Raoul de Charette² Cengiz Oztireli³ Jing-Hao Xue¹
¹University College London ²Inria ³University of Cambridge

{weihaio.xia.21, jinghao.xue}@ucl.ac.uk, raoul.de-charette@inria.fr, aco41@cam.ac.uk

Abstract

In this work we present *DREAM*, an fMRI-to-image method for reconstructing viewed images from brain activities, grounded on fundamental knowledge of the human visual system. We craft reverse pathways that emulate the hierarchical and parallel nature of how humans perceive the visual world. These tailored pathways are specialized to decipher semantics, color, and depth cues from fMRI data, mirroring the forward pathways from visual stimuli to fMRI recordings. To do so, two components mimic the inverse processes within the human visual system: the Reverse Visual Association Cortex (R-VAC) which reverses pathways of this brain region, extracting semantics from fMRI data; the Reverse Parallel PKM (R-PKM) component simultaneously predicting color and depth from fMRI signals. The experiments indicate that our method outperforms the current state-of-the-art models in terms of the consistency of appearance, structure, and semantics. Code will be available at <https://github.com/weihaiox/DREAM>.

1. Introduction

Exploring neural encoding unravel the intricacies of brain function. In last years, we have witnessed tremendous progress in visual decoding [14] which aims at decoding a Functional Magnetic Resonance Imaging (fMRI) to reconstruct the test image seen by a human subject during the fMRI recording. Visual decoding could significantly affect our society from how we interact with machines to helping paralyzed patients [36]. However, existing methods still suffer from missing concepts and limited quality in the image results. Recent studies turned to deep generative models for visual decoding due to their remarkable generation capabilities, particularly the text-to-image diffusion models [29, 40]. These methods heavily rely on aligning brain signals with the vision-language model [27]. This strategic utilization of CLIP helps mitigate the scarcity of annotated data and the complexities of underlying brain information.

Still, the inherent nature of CLIP, which fails to preserve the scene structural and positional information, limits vi-

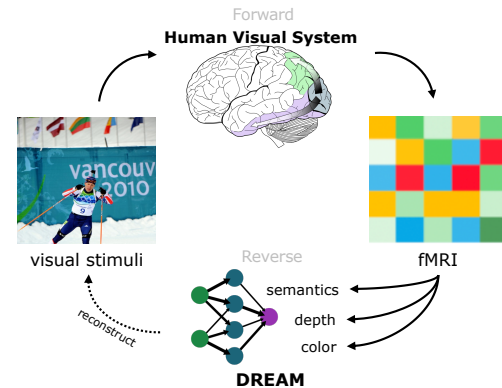


Figure 1. **Forward and Reverse Cycle.** Forward (HVS): visual stimuli \mapsto color, depth, semantics \mapsto fMRI; Reverse (**DREAM**): fMRI \mapsto color, depth, semantics \mapsto reconstructed images.

visual decoding. Hence, current methods have endeavored to incorporate structural and positional details, either through depth maps [7, 35] or by utilizing the decoded representation of an initial guessed image [25, 31]. However, these methods primarily focus on merging inputs that fit well within the pretrained generative model for visual decoding, lacking the insights from the human visual system.

We commence our study with the foundational principles [3] governing the Human Visual System (HVS) and dissect essential cues crucial for effective visual decoding. Our method draw insights from HVS — how humans perceive visual stimuli (forward route in Fig. 1) — to address the potential information loss during the transition from the fMRI to the visual domain (reverse route in Fig. 1). We do that by deciphering crucial cues from fMRI recordings, thereby contributing to enhanced consistency in terms of appearance, structure, and semantics. As cues, we investigate: *color* for accurate scene appearance [24], *depth* for scene structure [28], and the popular *semantics* for high-level comprehension [27]. Our study shows that current visual decoding methods often underscored and unnoticed color which in fact plays an indispensable role. Fig. 2 highlights color inconsistencies in a recent work [25]. The generated images, while accurate in semantics, deviate in struc-




Figure 2. **Appearance Inconsistency.** When decoding the fMRI data of a subject viewing a test image (top), recent visual decoding methods, here [25], reconstruct images (bottom) which are semantically close but still suffer from strong color inconsistencies.

ture and color from the original visual stimuli. This phenomenon arises due to the absence of proper color guidance.

Following above analysis, we propose **DREAM**, a visual Decoding method from REversing huMAn visual system. It aims to mirror the forward process from visual stimuli to fMRI recordings (Sec. 3). Specifically, we design two reverse pathways specialized in deciphering semantics, color, and depth from fMRI. Reverse-VAC (Sec. 4.1) replicates the reverse operations of the visual associated cortex, extracting semantics from fMRI. Reverse-PKM (Sec. 4.2) predicts color and depth simultaneously from fMRI signals. Deciphered cues are then fed into Stable Diffusion [29] with T2I-Adapter [24] to guide the image reconstruction (Sec. 4.3). Our contributions are summarized as follows:

- We scrutinize the limitations of recent diffusion-based visual decoding methods, shedding light on the potential loss of information, and introduce a novel formulation based on the principles of human perception.
- We mirror the forward process from visual stimuli to fMRI recordings within the visual system and devise two reverse pathways specialized in extracting semantics, color, and depth information from fMRI data.
- We show through experiments that our biologically interpretable method, DREAM, outperforms state-of-the-art methods while maintaining better consistency of appearance, structure, and semantics.

2. Related Work

2.1. Diffusion Probabilistic Models

Recently, diffusion models have risen to prominence as cutting-edge generative models. Denoising Diffusion Probabilistic Model is a parameterized bi-directional Markov chain that utilizes variational inference to produce matching samples. The forward diffusion process is designed to transform any data distribution into a basic prior distribution (e.g., isotropic Gaussian), and the reverse denoising

process learns to denoise by learning transition kernels parameterized by deep neural networks such as U-Net. In Latent Diffusion Models (LDMs) [29], diffusion process is applied within the latent space rather than in the pixel space, enabling faster inference and reducing training costs. The text-conditioned LDM, known as Stable Diffusion (SD), has gained widespread usage due to its versatile applications and capabilities. ControlNet [42] and T2I-Adapter [24] aim to enhance the control capabilities even more by training versatile modality-specific encoders. These encoders align external control (e.g., sketch, depth, and spatial palette) with internal knowledge in SD, thereby enabling more precise control over the generated output. Unlike SD, which solely employs the CLIP text encoder, Versatile Diffusion (VD) [40] incorporates both CLIP text and image encoders, thereby enabling the utilization of multimodal capabilities.

2.2. Image Decoding from fMRI

The advancements in visual decoding are closely intertwined with the evolution of various modeling frameworks. For instance, in [13], sparse linear regression was applied to preprocessed fMRI data to predict features extracted from early convolutional layers in a pretrained CNN. In the past few years, researchers have advanced visual decoding techniques by mapping the brain signals to the latent space of generative adversarial networks (GANs) [10] to reconstruct human faces [5] and natural scenes [25, 32]. More recently, visual decoding has reached an unprecedented level of quality [11, 26, 34] with the release of vision-language models [27], multimodal diffusion models [29, 40], and large-scale fMRI datasets [2]. Lin *et al.* [19] learned to project voxels to the CLIP space and then processed outcomes through a fine-tuned conditional StyleGAN2 [16] to reconstruct natural images. Takagi *et al.* [34] employed the ridge regression to associate fMRI signals with the CLIP text embedding and the latent space of Stable Diffusion, opting for varied voxels based on different components. Recent research [22, 26] explored the process of mapping fMRI signals to both CLIP text and image embeddings, subsequently utilizing the pre-trained Versatile Diffusion model [40] that accommodates multiple inputs for image reconstruction.

2.3. Multi-Level Modeling in Visual Decoding

Hierarchical visual feature representations are frequently utilized in visual decoding. Early studies [1, 13] have indicated that hierarchical features extracted through pretrained CNN models demonstrated strong correlation with neural activities of visual cortices. Recent research delved into combining both low-level visual cues and high-level semantics inferred from brain activity, using frozen diffusion models for reconstruction [22, 26, 31, 41]. The low-level visual cues are commonly incorporated in an implicit manner, such as utilizing intermediate features predicted by a

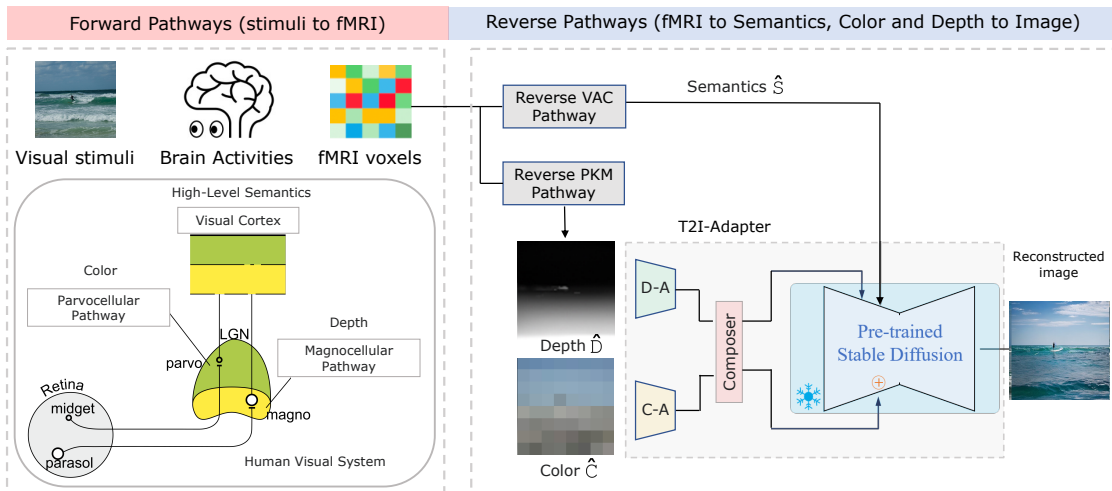


Figure 3. **Relation of the HVS and Our proposed DREAM.** Grounding on the Human Visual System (HVS), we devise reverse pathways aimed at deciphering semantics, depth, and color cues from fMRI to guide image reconstruction. **(Left)** Schematic view of HVS, detailed in Sec. 3. When perceiving visual stimuli, connections from the retina to the brain can be separated into two parallel pathways. The Parvocellular Pathway originates from midget cells in the retina and is responsible for transmitting color information, while the Magnocellular Pathway starts with parasol cells and is specialized in detecting depth and motion. The conveyed information is channeled into the visual cortex for undertaking intricate processing of high-level semantics from the visual image. **(Right)** DREAM mimics the corresponding inverse processes within the HVS: the Reverse VAC (Sec. 4.1) replicates the opposite operations of this brain region, analogously extracting semantics \hat{S} as a form of CLIP embedding from fMRI; and the Reverse PKM (Sec. 4.2) maps fMRI to color \hat{C} and depth \hat{D} in the form of spatial palettes and depth maps to facilitate subsequent processing by the Color Adapter (C-A) and the Depth Adapter (D-A) in T2I-Adapter [24] in conjunction with SD [29] for image reconstruction from deciphered semantics, color, and depth cues $\{\hat{S}, \hat{C}, \hat{D}\}$.

large-scale vision model [26] or encoded from an initial estimated image [31]. The high-level semantics is frequently represented as CLIP embedding. Recent research [7, 35] also suggested the explicit provision of both low-level and high-level information, predicting captions and depth maps from brain signals. Our method differs in terms of how to predict auxiliary information and the incorporation of color.

3. Preliminary on the Human Visual System

Human Visual System (HVS) endows us with the ability of visual perception. The visual information is *concurrently* relayed from various cell types in the retina, each capturing distinct facets of data, through the optic nerve to the brain. Connections from the retina to the brain, as shown in Fig. 3, can be separated into a parvocellular pathway and a magnocellular pathway¹. The parvocellular pathway originates from midget cells in the retina and is responsible for transmitting color information, while the magnocellular pathway starts with parasol cells and is specialized in detecting depth and motion. The visual information is first directed to a sensory relay station known as the lateral geniculate nucleus (LGN) of the thalamus, before being channeled to the visual cortex (V1) for the initial processing of visual stimuli. The visual association cortex (VAC) receives processed in-

¹An additional set of neurons, known as the koniocellular layers, are found ventral to each of the magnocellular and parvocellular layers [3].

formation from V1 and undertakes intricate processing of high-level semantic contents from the visual image.

The hierarchical and parallel manner where visual stimuli are broken down and passed forward as color, depth, and semantics guided our choices to reverse the HVS for decoding. For a detailed illustration of human perception and an analysis on the feasibility of extracting desired cues from fMRI recordings, please consult supplementary material.

4. DREAM

The task of visual decoding aims to recover the viewed image $I \in \mathbb{R}^{H \times W \times 3}$ from brain activity signals elicited by visual stimuli. Functional MRI (fMRI) is usually employed as a proxy of the brain activities, typically encoded as a set of voxels $\text{fMRI} \in \mathbb{R}^{1 \times N}$. Formally, the task optimizes $f(\cdot)$ so that $f(\text{fMRI}) = \hat{I}$, where \hat{I} best approximates I .

To address this task, we propose **DREAM**, a method grounded on fundamental principles of human perception. Following Sec. 3, our method relies on explicit design of reverse pathways to decipher Semantics, Color, and Depth intertwined in the fMRI data. These reverse pathways mirror the forward process from visual stimuli to brain activity. Considering that an fMRI captures changes in the brain regions during the forward process, it is feasible to derive the desired cues of the visual stimuli from such recording [2].

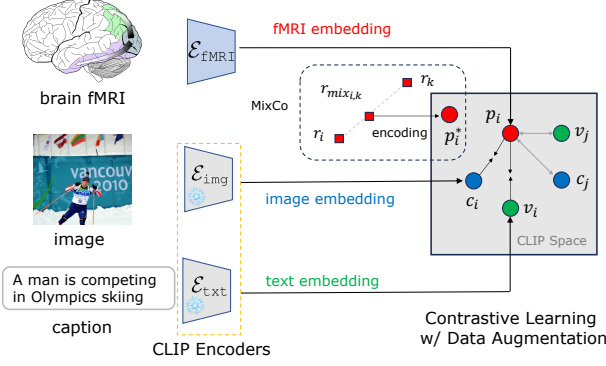


Figure 4. **R-VAC Training.** To decipher semantics from fMRI, we train an encoder $\mathcal{E}_{\text{fmri}}$ in a contrastive fashion which aligns fMRI data with the frozen CLIP space [27]. Data augmentation (represented by the dashed rectangle) [17] combats the data scarcity of the fMRI modality. See Sec. 4.1 for more details.

Overview. Fig. 3 illustrates an overview of DREAM. It is constructed on two consecutive phases, namely, Pathways Reversing and Guided Image Reconstruction. These phases break down the reverse mapping from fMRI to image into two subprocesses: $\text{fMRI} \mapsto \{\hat{S}, \hat{C}, \hat{D}\}$ and $\{\hat{S}, \hat{C}, \hat{D}\} \mapsto \hat{I}$. In the first phase, two Reverse Pathways decipher the cues of semantics, color and depth from fMRI with parallel components: Reverse Visual Association Cortex (R-VAC, Sec. 4.1) inverts operations of the VAC region to extract semantic details from the fMRI, encoded as CLIP embedding [27], and Reverse Parallel Parvo-, Konio- and Magno-Cellular (R-PKM, Sec. 4.2) is designed to predict color and depth simultaneously from fMRI signals. Given the lossy nature of fMRI data and the non-bijective transformation of image \mapsto fMRI, we then cast the decoding process as a generative task while using the extracted $\hat{S}, \hat{C}, \hat{D}$ cues as conditions for image reconstruction. Therefore, in the second phase Guided Image Reconstruction (GIR, Sec. 4.3), we follow recent visual decoding practices [34, 35] and employ a frozen SD with T2I-Adapter [24] to generate images through benefiting here the additional $\hat{S}, \hat{C}, \hat{D}$ guidance.

4.1. R-VAC (Semantics Decipher)

The Visual Association Cortex (VAC), as detailed in Sec. 3, is responsible for interpreting the high-level semantics of visual stimuli. We design R-VAC to reverse such process through analogous learning of the mapping from fMRI to semantics $\text{fMRI} \mapsto \hat{S}$. This is achieved by training an encoder $\mathcal{E}_{\text{fmri}}$ whose goal is to align the fMRI embedding with the shared CLIP space [27]. Though CLIP was initially trained with image-text pairs, prior works [21, 31] demonstrated ability to align new modalities. To fight the scarcity of fMRI data, we also carefully select ad-hoc data augmentation strategy [17].

Contrastive Learning. In practice, we train the fMRI encoder $\mathcal{E}_{\text{fmri}}$ with triplets of {fMRI, image, caption} to pull the fMRI embeddings closer to the rich shared semantic space of CLIP. Given that both text encoder (\mathcal{E}_{txt}) and image encoder (\mathcal{E}_{img}) of CLIP are frozen, we minimize the embedding distances of fMRI-image and fMRI-text which in turns forces alignment of the fMRI embedding with CLIP. The training is illustrated in Fig. 4. Formally, with the embeddings of fMRI, text, and image denoted by p, c, v , respectively, the initial contrastive loss writes

$$\mathcal{L}_p = -\log \frac{\exp(p_i \cdot c_i / \tau)}{\sum_{j=0}^K \exp(p_i \cdot c_j / \tau)} - \log \frac{\exp(p_i \cdot v_i / \tau)}{\sum_{j=0}^K \exp(p_i \cdot v_j / \tau)}, \quad (1)$$

where τ is a temperature hyperparameter. The sum for each term is over one positive and K negative samples. Each term represents the log loss of a $(K+1)$ -way softmax-based classifier [12], which aims to classify p_i as c_i (or v_i). The sum over samples of the batch size n is omitted for brevity.

Data Augmentation. An important issue to consider is that there are significant fewer fMRI samples ($\approx 10^4$) compared to the number of samples used to train CLIP (10^8), which may damage contrastive learning [15, 43]. To address this, we utilize a data augmentation loss based on MixCo [17], which generates mixed fMRI data $r_{\text{mix}_{i,k}}$ from convex combination of two fMRI data r_i and r_k :

$$r_{\text{mix}_{i,k}} = \lambda_i \cdot r_i + (1 - \lambda_i) \cdot r_k, \quad (2)$$

where k represents the arbitrary index of any data in the same batch, and its encoding writes $p_i^* = \mathcal{E}_{\text{fmri}}(r_{\text{mix}_{i,k}})$. The data augmentation loss, which excludes the image components for brevity, is formulated as

$$\mathcal{L}_{\text{MixCo}} = -\sum_{i=1}^n \left[\lambda_i \cdot \log \frac{\exp(p_i^* \cdot c_i / \tau)}{\sum_{j=0}^K \exp(p_i^* \cdot c_j / \tau)} + (1 - \lambda_i) \cdot \log \frac{\exp(p_i^* \cdot c_k / \tau)}{\sum_{j=0}^K \exp(p_i^* \cdot c_j / \tau)} \right]. \quad (3)$$

Finally, the total loss is a combination of \mathcal{L}_p and $\mathcal{L}_{\text{MixCo}}$ weighted with hyperparameter α :

$$\mathcal{L}_{\text{total}} = \mathcal{L}_p + \alpha \mathcal{L}_{\text{MixCo}}. \quad (4)$$

4.2. R-PKM (Depth & Color Decipher)

While R-VAC provides semantics knowledge, the latter is inherently bounded by the CLIP space capacity, unable to encode spatial colors and geometry. To address this issue, inspired by the human visual system, we craft the R-PKM component to reverse pathways of the Parvo-, Konio- and Magno-Cellular (PKM), subsequently predicting color and depth from the fMRI data denoted as $\text{fMRI} \mapsto \{\hat{C}, \hat{D}\}$.

While color and depth can be represented in various ways (e.g., histograms, graphs), we represent them as *spatial color palettes* and *depth maps* to facilitate the reconstruction guidance, as discussed in Sec. 4.3. Visuals are in Fig. 3.

In practice, we formulate the problem as RGBD estimation. The color palette is then derived from the RGB prediction by first $\times 64$ downscaling it and then upscaling back to its original size. There are readily available methods for $\text{fMRI} \mapsto \text{RGBD}$ mapping [7, 35] but they offer limited performance due to the scarcity fMRI data. Instead, we introduce a multi-stage encoder-decoder training [9], which benefits from both the scarcely available (fMRI, RGBD) pairs and the abundant RGBD data without fMRI. Fig. 5 shows the training procedure for R-PKM.

Stage 1. Given limited pairs $\{(r, d)\} = \{\text{fMRI}, \text{RGBD}\}$, we first train an encoder to map RGBD to their corresponding fMRI data. To compensate for the absence of depth in fMRI datasets, we use MiDaS-estimated depth maps [28] as surrogate ground-truth depth. The encoder is trained with a convex combination of mean square error and cosine proximity between the input r and its predicted counterpart \hat{r} :

$$\mathcal{L}_r(r, \hat{r}) = \beta \cdot \text{MSE}(r, \hat{r}) - (1 - \beta) \cos(\angle(r, \hat{r})), \quad (5)$$

where β is determined empirically as a hyperparameter.

Stage 2. Similar to stage 1, we now train the decoder with pairs $\{(r, d)\}$ in a supervised manner:

$$\mathcal{L}_s(d, \hat{d}) = \|d - \hat{d}\|_1 + \mathcal{J}(\hat{d}), \quad (6)$$

where $\hat{d} = \mathcal{D}(r)$ and the total variation regularization $\mathcal{J}(\hat{d})$ encourages spatial smoothness in the reconstructed \hat{d} .

Stage 3. To address the scarcity of fMRI data and improve the model generalization to unseen categories, we employ a self-supervised strategy to finetune the decoder while keeping the encoder frozen. This facilitates the usage of any natural images (e.g., from ImageNet [6] or LAION [30]) along with their estimated depth maps, without need of paired fMRI and image data. Hence, we train solely with the RGBD data by ensuring a cycle consistency through the Encoder-Decoder transformation, i.e. $\hat{d} = \mathcal{D}(\mathcal{E}(d))$, with the loss in Eq. (6). Given that this stage involves images for which fMRI data was never collected, the model greatly improves its generalization capability.

4.3. Guided Image Reconstruction (GIR)

Equipped with R-VAC (Sec. 4.1) and R-PKM (Sec. 4.2), our method can decipher semantics \hat{S} , color \hat{C} , and depth \hat{D} in the form of CLIP embedding, spatial color palette, and depth map. Finally, guided image reconstruction $\{\hat{S}, \hat{C}, \hat{D}\} \mapsto \hat{I}$ completes the reverse mapping of the forward process during the visual perception $I \mapsto \text{fMRI}$.




Figure 5. **R-PKM Training.** Our multi-stage training reverses the PKM pathway and decodes color and depth cues in fMRI data. Stages 1 and 2 employ (RGBD, fMRI) pairs to train an encoder \mathcal{E} that maps RGBD to fMRI and a decoder \mathcal{D} to decode RGBD from fMRI. Stage 3 benefits from additional RGBD images without fMRI to train \mathcal{D} in a cycle-consistent manner $\hat{d} = \mathcal{D}(\mathcal{E}(d))$, while keeping \mathcal{E} frozen. See Sec. 4.2 for details.

We utilize Stable Diffusion (SD) [29] to reconstruct the final image from the predicted CLIP embedding \hat{S} and the additional guidance from predicted color palette \hat{C} and depth map \hat{D} . Such guidance is produced using the color adapter \mathcal{R}_c and the depth adapter \mathcal{R}_d within T2I-adapter [24]. This process is formulated as follows:

$$\begin{aligned} F_{\mathcal{R}} &= \omega_c \mathcal{R}_c(\hat{C}) + \omega_d \mathcal{R}_d(\hat{D}), \\ \hat{I} &= \text{SD}(z, F_{\mathcal{R}}, \hat{S}), \end{aligned} \quad (7)$$

where z is a random noise, ω_c and ω_d are adjustable weights to control the relative significance of the adapters.

5. Experiments

Following common practice in the field, we evaluate DREAM with the largest public neuroimaging dataset, the Natural Scene Dataset [2]. We report our performance against five leading methods [11, 19, 26, 31, 34]. We detail our experimental methodology in Sec. 5.1 and report quantitative and qualitative evaluations in Sec. 5.2. Ablation studies are presented in Sec. 6.

5.1. Experimental Setting

Dataset. We use the Natural Scenes Dataset (NSD) [2] in all experiments, which follows the standard practices in the field [11, 19, 22, 26, 31, 34]. NSD, as the largest fMRI dataset, records brain responses from eight human subjects successively isolated in an MRI machine and passively observed a wide range of visual stimuli, namely, natural images sourced from MS-COCO [20], which allows retrieving the associated captions. In practice, because brain activity patterns highly vary across subjects [14], a separate model is trained per subject. The standardized splits contain 982 fMRI test samples and 24,980 fMRI training samples. Please refer to the supplementary material for more details.

Metrics for Visual Decoding. The same set of eight metrics is utilized for our evaluation in accordance with prior




Figure 6. Sample Visual Decoding Results from the SOTA Methods on NSD.

Table 1. **Quantitative Evaluation.** Following standard NSD metrics, DREAM performs on a par or better than the SOTA methods (we highlight **best** and **second**). We also report ablation of the two strategies fighting fMRI data scarcity: R-VAC without Data Augmentation (DA) and R-PKM without the third-stage decoder training (S3) that allows additional RGBD data without fMRI.

Method	Low-Level				High-Level			
	PixCorr \uparrow	SSIM \uparrow	AlexNet(2) \uparrow	AlexNet(5) \uparrow	Inception \uparrow	CLIP \uparrow	EffNet-B \downarrow	SwAV \downarrow
Mind-Reader [19]	—	—	—	—	78.2%	—	—	—
Takagi <i>et al.</i> [34]	—	—	83.0%	83.0%	76.0%	77.0%	—	—
Gu <i>et al.</i> [11]	.150	.325	—	—	—	—	.862	.465
Brain-Diffuser [26]	.254	.356	94.2%	96.2%	87.2%	91.5%	.775	.423
MindEye [31]	.309	.323	<u>94.7%</u>	97.8%	<u>93.8%</u>	<u>94.1%</u>	<u>.645</u>	.367
DREAM (Ours)	<u>.288</u>	<u>.338</u>	95.0%	<u>97.5%</u>	94.8%	95.2%	.638	<u>.413</u>
w/o DA (R-VAC)	.279	.340	86.8%	88.1%	87.2%	89.9%	.662	.517
w/o S3 (R-PKM)	.203	.295	92.7%	96.2%	92.1%	94.6%	.642	.463

research [26, 31]. To be specific, PixCorr is the pixel-level correlation between the reconstructed and ground-truth images. PixCorr is the pixel-level correlation between the reconstructed and ground-truth images. Structural Similarity Index (SSIM) [38] quantifies similarity between two images. It measures the structural and textural similarity rather than just pixel-wise differences. AlexNet(2) and AlexNet(5) are two-way comparisons of the second and fifth layers of AlexNet [18], respectively. Inception is the two-way comparison of the last pooling layer of InceptionV3 [33]. CLIP is the two-way comparison of the last layer of the CLIP-Vision [27] model. EffNet-B and SwAV are distances gathered from EfficientNet-B1 [37] and SwAV-ResNet50 [4], respectively. The first four metrics can be categorized as low-level measurements, whereas the remaining four capture higher-level characteristics.

Metrics for Depth and Color. We use metrics from depth estimation and color correction to assess depth and color consistencies in the final reconstructed images. The depth metrics, as elaborated in [23], include Abs Rel (absolute error), Sq Rel (squared error), RMSE (root mean squared error), and RMSE log (root mean squared logarithmic error). For color metrics, we use CD (Color Discrepancy) [39]

Table 2. **Consistency of the Decoded Images.** We evaluate the color and depth consistencies in decoded images by comparing the distances between test images and visual decoding results. DREAM significantly outperforms the other two methods [26, 31].

Method	Depth				Color	
	Abs Rel \downarrow	Sq Rel \downarrow	RMSE \downarrow	RMSE log \downarrow	CD \downarrow	STRESS \downarrow
Brain-Diffuser [26]	10.162	4.819	<u>9.871</u>	1.157	4.231	47.025
MindEye [31]	<u>8.391</u>	<u>4.176</u>	9.873	<u>1.075</u>	<u>4.172</u>	<u>45.380</u>
DREAM (Ours)	7.695	4.031	9.862	1.039	2.957	37.285

and STRESS (Standardized Residual Sum of Squares) [8]. Please consult the supplementary material for details.

Implementation Details. One NVIDIA A100-SXM-80GB GPU is used in all experiments, including the training of fMRI \mapsto Semantics encoder $\mathcal{E}_{\text{fMRI}}$ and fMRI \mapsto Depth & Color encoder \mathcal{E} and decoder \mathcal{D} . We use pretrained color and depth adapters from T2I-adapter [24] to extract guidance features from predicted spatial palettes and depth maps. These guidance features, along with predicted CLIP representations, are then input into the pretrained SD model for the purpose of image reconstruction. The hyperparameters $\alpha = 0.3$, $\beta = 0.9$, ω_c and ω_d are set as 1.0 unless otherwise mentioned. For further details regarding the network



Figure 7. **Visual Decoding with DREAM.** Sample outputs demonstrate DREAM’s ability to accurately decode the visual stimuli from fMRI. Our depth and color predictions from the R-PKM (Sec. 4.2) are in line with the pseudo ground-truth, despite the extreme complexity of the task. DREAM reconstructions closely match the test images, and the rightmost samples demonstrate the benefit of color guidance.

Table 3. **Effectiveness of R-VAC (Semantics Decipherer) and R-PKM (Depth & Color Decipherer).** We conducted two sets of experiments using ground-truth (GT) or predicted (Pred) cues to reconstruct the visual stimuli, respectively.

	Reconstruction (Sec. 4.3)	Low-Level				High-Level			
		PixCorr \uparrow	SSIM \uparrow	AlexNet(2) \uparrow	AlexNet(5) \uparrow	Inception \uparrow	CLIP \uparrow	EffNet-B \downarrow	SwAV \downarrow
GT	semantics {S}	.244	.272	96.68%	97.39%	87.82%	92.45%	1.00	.415
	+depth {S, D}	.186	.286	99.58%	99.78%	98.78%	98.09%	.723	.322
	+color {S, D, C}	.413	.366	99.99%	99.98%	99.19%	98.66%	.702	.278
Pred	semantics {S-hat}	.194	.278	91.82%	92.57%	93.11%	91.24%	.645	.369
	+depth {S-hat, D-hat}	.083	.282	88.07%	94.69%	94.13%	96.05%	.802	.429
	+color {S-hat, D-hat, C-hat}	.288	.338	94.99%	97.50%	94.80%	95.24%	.638	.413

architecture, please refer to the supplementary material.

5.2. Experimental Results and Analysis

Visual Decoding. Our method is compared with five state-of-the-art methods: Mind-Reader [19], Takagi *et al.* [34], Gu *et al.* [11], Brain-Diffuser [26], and MindEye [31]. The quantitative visual decoding results are presented in Tab. 1, indicating a competitive performance. Our method, with explicit deciphering mechanism, appears to be more proficient at discerning scene structure and semantics, as evidenced by the favorable high-level metrics. The qualitative results, depicted in Fig. 6, align with the numerical findings, indicating that DREAM produces more realistic outcomes that maintains consistency with the viewed images in terms of semantics, appearance, and structure, compared to the other methods. Striking DREAM outputs are the food plate (left, middle row) which accurately decodes the presence of vegetables and a tablespoon, and the baseball scene (right, middle row) which showcases the correct number of players (3) with poses similar to the test image.

Besides visual appearance, we wish to measure the consistency of depth and color in the decoded images with respect to the test images viewed by the subject. We achieve this by measuring the variance in the estimated depth (and color palettes) of the test image and the reconstructed results from Brain-Diffuser [26], MindEye [31], or DREAM. Results presented in Tab. 2 indicate that our method yields images that align more consistently in color and depth with the visual stimuli than the other two methods.

Cues Deciphering. Our method decodes three cues from fMRI data: semantics, depth, and color. To assess semantics deciphered from R-VAC (Sec. 4.1), we simply refer to the CLIP metric of Tab. 1 which quantifies CLIP embeddings distances with the test image. From the aforementioned table, DREAM is at least 1.1% better than others. Fig. 7 shows examples of the depth (\hat{D}) and color (\hat{C}) deciphered by R-PKM (Sec. 4.2). While accurate depth is beneficial for image reconstruction, faithfully recovering the original depth from fMRI is nearly impossible due to the information loss in capturing the brain activities [2]. Still, coarse depth is sufficient in most cases to guide the scene structure and object position such as determining the location of an airplane or the orientation of a bird standing on a branch. This is intuitively understood from the bottom row of Fig. 7, where our coarse depth (\hat{D}) leaves no doubt on the giraffe’s location and orientation. Interestingly, despite not precisely preserving the local color, the estimated color palettes (\hat{C}) provide a reliable constraint and guidance on the overall scene appearance. This is further demonstrated in last three columns of Fig. 7 by removing color guidance which, despite appealing visuals, proves to produce images drastically differing from the test image.

6. Ablation Study

Here we present ablation studies while discussing first the effect of using color, depth and semantics as guidance for image reconstruction thanks to our reversed pathways (*i.e.*, R-VAC and R-PKM).



Figure 8. **Effect of the Composition Weight in GIR (Sec. 4.3).** The two weights ω_c and ω_d control the relative importance of corresponding features from three deciphered cues: semantics, color, and depth. When predicted depth or color fail to provide reliable guidance, we can manually tweak the weights to achieve satisfactory reconstructed results.

Effect of Color Palettes. As highlighted earlier, the deciphered color guidance noticeably enhances the visual quality of the reconstructed images in Fig. 7. We further quantify the color’s significance through two additional sets of experiments detailed in Tab. 3, where the reconstruction uses: 1) ground-truth (GT) depth and caption with or without color, and 2) fMRI-predicted (Pred) depth and semantic embedding with or without the predicted color. The results using ground-truth cues serve as a proxy. The generated results exhibit improved color consistency and enhanced quantitative performance across the board, underscoring the importance of using color for visual decoding.

Effect of Depth and Semantics. Tab. 3 presents results where we ablate the use of depth and semantics. Comparison of GT and predicted semantics (*i.e.*, $\{S\}$ and $\{\hat{S}\}$) suggests that the fMRI embedding effectively incorporates high-level semantic cues into the final images. The overall image quality can be further improved by integrating either ground truth depth (represented as $\{S, D\}$) or predicted depth ($\{\hat{S}, \hat{D}\}$) combined with color (denoted as $\{S, D, C\}$ and $\{\hat{S}, \hat{D}, \hat{C}\}$). Introducing color cues not only bolsters the structural information but also strengthens the semantics, possibly because it compensates for the color information absent in the predicted fMRI embedding. Of note, all metrics (except for PixCorr) improve smoothly with more GT guidance. Yet, the impact of predicted cues varies across metrics, highlighting intriguing research avenues and emphasizing the need for more reliable measures.

Effect of Data Scarcity Strategies. We ablate the two strategies introduced to fight fMRI data scarcity: data augmentation (DA) in R-VAC and the third-stage decoder training (S3) which allows R-PKM to use additional RGBD data without fMRI. The results shown in the two bottom rows of Tab. 1 demonstrate that the two data augmentation strate-

gies address the limited availability of fMRI data and subsequently bolster the model’s generalization capability.

Effect of Weighted Guidance. The features inputted into SD are formulated as $S + \omega_c \mathcal{R}_c(C) + \omega_d \mathcal{R}_d(D)$, where two weights ω_c and ω_d from Eq. (7) control the relative importance of the deciphered cues and play a crucial role in the final image quality and alignment with the cues. In DREAM, ω_c and ω_d are set to 1.0 showing no preference of color guidance over depth guidance. Still, Fig. 8 shows that in some instances the predicted components fail to provide dependable guidance on structure and appearance, thus compromising results. There is also empirical evidence that indicates the T2I-adapter slightly underperforms when compared to ControlNet. The performance of the T2I-adapter further diminishes when multiple conditions are used, as opposed to just one. Taking both factors into account, there are instances where manual adjustments to the weighting parameters become necessary to achieve images of the desired quality, semantics, structure, and appearance.

7. Conclusion

This paper presents DREAM, a visual decoding method founded on principles of human perception. We design reverse pathways that mirror the forward pathways from visual stimuli to fMRI recordings. These pathways specialize in deciphering semantics, color, and depth cues from fMRI data and then use these predicted cues as guidance to reconstruct visual stimuli. Experiments demonstrate that our method surpasses current state-of-the-art models in terms of consistency in appearance, structure, and semantics.

Acknowledgements This work was supported by the Engineering and Physical Sciences Research Council [grant number EP/W523835/1] and a UKRI Future Leaders Fellowship [grant number G104084].

References

- [1] Pulkit Agrawal, Dustin Stansbury, Jitendra Malik, and Jack L Gallant. Pixels to voxels: modeling visual representation in the human brain. *arXiv preprint arXiv:1407.5104*, 2014. [2](#)
- [2] Emily J Allen, Ghislain St-Yves, Yihan Wu, Jesse L Breedlove, Jacob S Prince, Logan T Dowdle, Matthias Nau, Brad Caron, Franco Pestilli, Ian Charest, et al. A massive 7t fmri dataset to bridge cognitive neuroscience and artificial intelligence. *Nature neuroscience*, 25(1):116–126, 2022. [2](#), [3](#), [5](#), [7](#)
- [3] Per Brodal. *The central nervous system: structure and function*. oxford university Press, 2004. [1](#), [3](#)
- [4] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. *NeurIPS*, 33:9912–9924, 2020. [6](#)
- [5] Thirza Dado, Yağmur Güçlütürk, Luca Ambrogioni, Gabriëlle Ras, Sander Bosch, Marcel van Gerven, and Umüt Güçlü. Hyperrealistic neural decoding for reconstructing faces from fmri activations via the gan latent space. *Scientific reports*, 12(1):141, 2022. [2](#)
- [6] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, pages 248–255, 2009. [5](#)
- [7] Matteo Ferrante, Furkan Ozcelik, Tommaso Boccatto, Rufin VanRullen, and Nicola Toschi. Brain Captioning: Decoding human brain activity into images and text. *arXiv preprint arXiv:2305.11560*, 2023. [1](#), [3](#), [5](#)
- [8] Pedro A Garcia, Rafael Huertas, Manuel Melgosa, and Guihua Cui. Measurement of the relationship between perceived and computed color differences. *JOSA A*, 24(7):1823–1829, 2007. [6](#)
- [9] Guy Gaziv and Michal Irani. More than meets the eye: Self-supervised depth reconstruction from brain activity. *arXiv preprint arXiv:2106.05113*, 2021. [5](#)
- [10] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020. [2](#)
- [11] Zijin Gu, Keith Jamison, Amy Kuceyeski, and Mert Sabuncu. Decoding natural image stimuli from fmri data with a surface-based convolutional network. In *MIDL*, 2023. [2](#), [5](#), [6](#), [7](#)
- [12] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, pages 9729–9738, 2020. [4](#)
- [13] Tomoyasu Horikawa and Yukiyasu Kamitani. Generic decoding of seen and imagined objects using hierarchical visual features. *Nature communications*, 8(1):15037, 2017. [2](#)
- [14] Shuo Huang, Wei Shao, Mei-Ling Wang, and Dao-Qiang Zhang. fmri-based decoding of visual information from human brain activity: A brief review. *International Journal of Automation and Computing*, 18(2):170–184, 2021. [1](#), [5](#)
- [15] Ziyu Jiang, Tianlong Chen, Bobak J Mortazavi, and Zhangyang Wang. Self-damaging contrastive learning. In *ICLR*, pages 4927–4939. PMLR, 2021. [4](#)
- [16] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of StyleGAN. In *CVPR*, pages 8107–8116, 2020. [2](#)
- [17] Sungnyun Kim, Gihun Lee, Sangmin Bae, and Se-Young Yun. Mixco: Mix-up contrastive learning for visual representation. In *NeurIPS Workshop*, 2020. [4](#)
- [18] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6):84–90, 2017. [6](#)
- [19] Sikun Lin, Thomas Sprague, and Ambuj K Singh. Mind Reader: Reconstructing complex images from brain activities. *NeurIPS*, 35:29624–29636, 2022. [2](#), [5](#), [6](#), [7](#)
- [20] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: Common objects in context. In *ECCV*, pages 740–755, 2014. [5](#)
- [21] Yulong Liu, Yongqiang Ma, Wei Zhou, Guibo Zhu, and Nanning Zheng. Brainclip: Bridging brain and visual-linguistic representation via clip for generic natural visual stimulus decoding from fmri. *arXiv preprint arXiv:2302.12971*, 2023. [4](#)
- [22] Yizhuo Lu, Changde Du, Dianpeng Wang, and Huiguang He. Minddiffuser: Controlled image reconstruction from human brain activity with semantic and structural diffusion. *arXiv preprint arXiv:2303.14139*, 2023. [2](#), [5](#)
- [23] Yue Ming, Xuyang Meng, Chunxiao Fan, and Hui Yu. Deep learning for monocular depth estimation: A review. *Neurocomputing*, 438:14–33, 2021. [6](#)
- [24] Chong Mou, Xintao Wang, Liangbin Xie, Jian Zhang, Zhonggang Qi, Ying Shan, and Xiaohu Qie. T2i-adapter: Learning adapters to dig out more controllable ability for text-to-image diffusion models. *arXiv preprint arXiv:2302.08453*, 2023. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#)
- [25] Furkan Ozcelik, Bhavin Choksi, Milad Mozafari, Leila Reddy, and Rufin VanRullen. Reconstruction of Perceived Images from fMRI Patterns and Semantic Brain Exploration using Instance-Conditioned GANs. In *IJCNN*, pages 1–8. IEEE, 2022. [1](#), [2](#)
- [26] Furkan Ozcelik and Rufin VanRullen. Brain-Diffuser: Natural scene reconstruction from fMRI signals using generative latent diffusion. *arXiv preprint arXiv:2303.05334*, 2023. [2](#), [3](#), [5](#), [6](#), [7](#)
- [27] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, pages 8748–8763. PMLR, 2021. [1](#), [2](#), [4](#), [6](#)
- [28] René Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun. Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *TPAMI*, 44(3):1623–1637, 2020. [1](#), [5](#)
- [29] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, pages 10684–10695, 2022. [1](#), [2](#), [3](#), [5](#)

- [30] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. Laion-5b: An open large-scale dataset for training next generation image-text models. *NeurIPS*, 35:25278–25294, 2022. [5](#)
- [31] Paul S Scotti, Atmadeep Banerjee, Jimmie Goode, Stepan Shabalin, Alex Nguyen, Ethan Cohen, Aidan J Dempster, Nathalie Verlinde, Elad Yundler, David Weisberg, et al. Reconstructing the mind’s eye: fmri-to-image with contrastive learning and diffusion priors. *arXiv preprint arXiv:2305.18274*, 2023. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#)
- [32] Guohua Shen, Tomoyasu Horikawa, Kei Majima, and Yukiyasu Kamitani. Deep image reconstruction from human brain activity. *PLoS computational biology*, 15(1):e1006633, 2019. [2](#)
- [33] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *CVPR*, pages 2818–2826, 2016. [6](#)
- [34] Yu Takagi and Shinji Nishimoto. High-resolution image reconstruction with latent diffusion models from human brain activity. In *CVPR*, pages 14453–14463, 2023. [2](#), [4](#), [5](#), [6](#), [7](#)
- [35] Yu Takagi and Shinji Nishimoto. Improving visual image reconstruction from human brain activity using latent diffusion models via multiple decoded inputs. *arXiv preprint arXiv:2306.11536*, 2023. [1](#), [3](#), [4](#), [5](#)
- [36] Desney Tan and Anton Nijholt. *Brain-computer interfaces and human-computer interaction*. Springer, 2010. [1](#)
- [37] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *ICML*, pages 6105–6114. PMLR, 2019. [6](#)
- [38] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *TIP*, 13(4):600–612, 2004. [6](#)
- [39] Menghan Xia, Jian Yao, Renping Xie, Mi Zhang, and Jinsheng Xiao. Color consistency correction based on remapping optimization for image stitching. In *ICCV workshops*, pages 2977–2984, 2017. [6](#)
- [40] Xingqian Xu, Zhangyang Wang, Eric Zhang, Kai Wang, and Humphrey Shi. Versatile diffusion: Text, images and variations all in one diffusion model. *arXiv preprint arXiv:2211.08332*, 2022. [1](#), [2](#)
- [41] Bohan Zeng, Shanglin Li, Xuhui Liu, Sicheng Gao, Xiaolong Jiang, Xu Tang, Yao Hu, Jianzhuang Liu, and Baochang Zhang. Controllable mind visual diffusion model. *arXiv preprint arXiv:2305.10135*, 2023. [2](#)
- [42] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *ICCV*, pages 3836–3847, 2023. [2](#)
- [43] Jianggang Zhu, Zheng Wang, Jingjing Chen, Yi-Ping Phoebe Chen, and Yu-Gang Jiang. Balanced contrastive learning for long-tailed visual recognition. In *CVPR*, pages 6908–6917, 2022. [4](#)