



OPEN Natural scene reconstruction from fMRI signals using generative latent diffusion

Furkan Ozcelik^{1,2}✉ & Rufin VanRullen^{1,2,3}

In neural decoding research, one of the most intriguing topics is the reconstruction of perceived natural images based on fMRI signals. Previous studies have succeeded in re-creating different aspects of the visuals, such as low-level properties (shape, texture, layout) or high-level features (category of objects, descriptive semantics of scenes) but have typically failed to reconstruct these properties together for complex scene images. Generative AI has recently made a leap forward with latent diffusion models capable of generating high-complexity images. Here, we investigate how to take advantage of this innovative technology for brain decoding. We present a two-stage scene reconstruction framework called “Brain-Diffuser”. In the first stage, starting from fMRI signals, we reconstruct images that capture low-level properties and overall layout using a VDVAE (Very Deep Variational Autoencoder) model. In the second stage, we use the image-to-image framework of a latent diffusion model (Versatile Diffusion) conditioned on predicted multimodal (text and visual) features, to generate final reconstructed images. On the publicly available Natural Scenes Dataset benchmark, our method outperforms previous models both qualitatively and quantitatively. When applied to synthetic fMRI patterns generated from individual ROI (region-of-interest) masks, our trained model creates compelling “ROI-optimal” scenes consistent with neuroscientific knowledge. Thus, the proposed methodology can have an impact on both applied (e.g. brain–computer interface) and fundamental neuroscience.

Establishing neural encoding and decoding techniques is one way for researchers to discover how the brain and cognition work. Recent developments in modeling and computation have opened up new ways of decoding information from brain signals. Numerous studies in the field of vision research have employed statistical techniques and machine learning to decode specific information from fMRI (functional magnetic resonance imaging) neural activity, such as position¹ or orientation^{2,3}, to predict categories of images^{4,5}, to match exemplar images from a candidate set⁶, and to reconstruct images with low levels of complexity, such as simple shapes and structures⁷.

In recent years, following the success in the development of deep learning models, many studies utilized deep generative models to reconstruct entire images. These deep generative models included Variational Autoencoders (VAE), Generative Adversarial Networks (GAN), and recently Latent Diffusion Models (LDM). Most of these studies used existing deep generative models, pretrained on large-scale data, and then learned a mapping (with simple regression or more advanced neural network architectures) to reconstruct the corresponding latent variables from the brain signals. This general method was used to reconstruct images with different levels of complexity such as faces^{8,9}, single-object-centered images¹⁰, and more complex scenes^{11,12}.

Most of the earlier works on natural scene reconstruction studied either the Generic Object Decoding¹³ or the Deep Image Reconstruction¹⁰ datasets curated by the Kamitani Lab. These datasets consist of 1200 training and 50 testing images from the ImageNet¹⁴ dataset and they differ in the number of fMRI repetitions for training and testing images. One of the pioneer studies in this area is by Shen et al.¹⁰ who optimized input images using a deep generator network with a loss function provided by fMRI-decoded CNN features. Belyi et al.¹⁵ utilized supervised training with {fMRI, stimulus} pairs, alongside an additional consistency loss for unsupervised training with test fMRI data and additional image data. Building on this, Gaziv et al.¹⁶ further improved the method by incorporating a perceptual loss on reconstructed images, resulting in sharper reconstructions. Mozafari et al.¹⁷ introduced a reconstruction model based on BigBiGAN that focused on semantics. Ren et al.¹⁸ devised a dual VAE-GAN model with a three-stage learning strategy that incorporates adversarial learning and knowledge distillation. Ozcelik et al.¹⁹ employed the Instance-Conditioned GAN model to generate reconstructions focused on accurate semantics (by extracting semantic information with the SwAV model) and pose information (with

¹CerCo, CNRS UMR5549, Toulouse, France. ²Universite de Toulouse, Toulouse, France. ³ANITI, Toulouse, France. ✉email: ozcelikfu@gmail.com

latent optimization). Chen et al.²⁰ utilized a sparse masked brain modeling on large-scale fMRI data and then trained a double-conditioned diffusion model for visual decoding.

Recently, Allen et al. curated another dataset for visual encoding and decoding studies called Natural Scenes Dataset¹¹. For this dataset, 8 subjects viewed thousands of images from the COCO²¹ dataset. COCO images contain multiple objects and they are more complex in nature compared to ImageNet images. Because of the number, diversity, and complexity of images included, the NSD dataset—although very recent—is becoming the de facto benchmark for fMRI-based natural scene reconstruction. Thus, it is the dataset that we chose for the present work. There are already three studies that reconstructed images from this dataset, and we can use them as baselines against which to compare our model's performance. The first one is by Lin et al.¹², who utilized the Lafite framework that adapts the StyleGAN2 model for text-to-image generation. Takagi et al.²² devised a method based on Stable Diffusion, using captions for the semantic information and latent variables from images for the low-level information. Gu et al.²³ improved upon Ozelik et al.'s¹⁹ IC-GAN framework, by establishing a surface-based convolutional network to process fMRI data instead of using vectorized data in the regression models; they also trained an encoder network to predict pose information, instead of performing latent optimization.

The above studies have fostered advances in reconstructing images with high fidelity, especially in the case of object-centered images (i.e., ImageNet images from the Kamitani dataset). Yet, reconstructing scenes with multiple objects and complex semantic descriptions (i.e., COCO images from the NSD dataset) remains a challenge. Given the remarkable recent success of latent diffusion models²⁴ in generative AI applications such as text-to-image generation^{24–28}, we reasoned that brain decoding studies could also take advantage of such models. Thus, we propose here a visual reconstruction framework called “Brain-Diffuser”, relying on the powerful generation capabilities of Versatile Diffusion²⁸, a model conditioned on both vision and language representations acquired from the pretrained CLIP²⁹ model.

Our framework consists of two stages. The first stage, illustrated in Fig. 1, generates a low-level reconstruction of images (akin to an “initial guess”) using a Very Deep Variational Autoencoder (VDVAE)³⁰. We generate these reconstructions by training a regression model to associate fMRI signals to the corresponding latent variables of VDVAE for the same training images. In the second stage, illustrated in Fig. 2, we train two additional regression models: one from fMRI patterns to CLIP-Vision features (extracted by feeding the corresponding images to the CLIP model); and the other one from fMRI patterns to CLIP-Text features (collected by providing to the CLIP model the captions of the corresponding images). Finally, we use the multimodal dual-guidance as well as the image-to-image abilities of the pretrained Versatile Diffusion (VD) model to generate the final reconstructions for test images. Using our trained regression models, for each test fMRI pattern we obtain an “initial guess” image (stage 1, VDVAE reconstruction) used by VD's image-to-image pipeline, as well as predicted CLIP-Vision and CLIP-Text feature vectors (stage 2), jointly used for conditioning VD's diffusion process. We used VDVAE, CLIP, and Versatile Diffusion with their pretrained weights, and did not apply any finetuning. We only trained regression models that transform fMRI patterns to latent variables of the models.

We demonstrate below that the resulting scene images reconstructed by the Brain-Diffuser model are highly naturalistic and retain the overall layout and semantic information of the groundtruth images while showing only minor variations in finer details. Compared to earlier models that exhibited proficiency in capturing certain features of groundtruth images, Brain-Diffuser demonstrates qualitatively and quantitatively superior performance in terms of both high-level and low-level metrics, thus establishing itself as state-of-the-art.

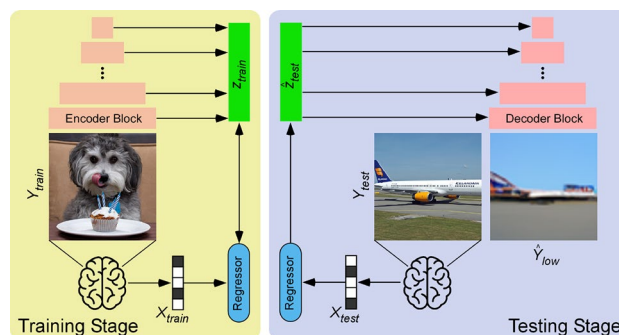


Figure 1. Reconstruction of Images via VDVAE (first stage). Training stage (left). Latent variables (z_{train}) are extracted and concatenated for the first 31 layers of the hierarchy by passing training images (Y_{train}) into the pretrained VDVAE Encoder. A ridge regression model (Regressor) is trained between fMRI patterns (X_{train}) and corresponding latent variables (z_{train}). Testing Stage (right). Test fMRI data (X_{test}) are passed through the trained Regressor to obtain predicted latent variables (\hat{z}_{test}). These predicted latent variables are fed to the pretrained VDVAE Decoder to get the low-level reconstruction (\hat{Y}_{low}) of the test images (Y_{test}), which will serve as a sort of “initial guess” for the second stage. Note that all VDVAE layers (encoder and decoder blocks) are pretrained and frozen, only the brain-to-latent regression layer (blue box) is trained.

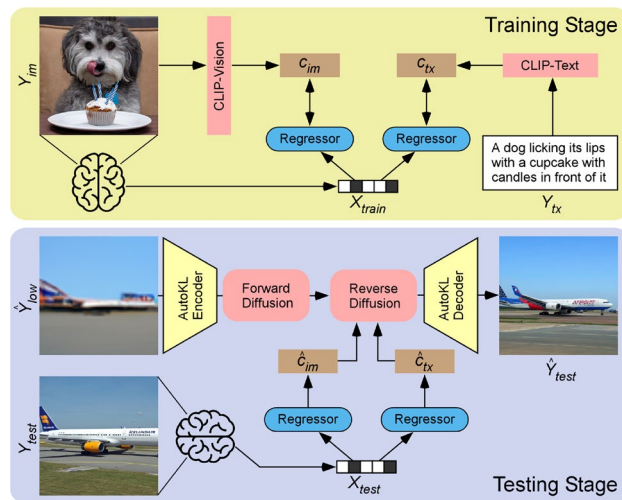


Figure 2. Final reconstruction of images via versatile diffusion (second stage). Training stage: CLIP-Vision features (c_{im}) are extracted by feeding training images (Y_{im}) to the pretrained CLIP model. CLIP-text features (c_{tx}) are extracted by providing the corresponding captions (Y_{tx}) to the pretrained CLIP Model. Two different ridge regression models (Regressors) are trained to learn the mapping between these features and fMRI patterns (X_{train}). Testing stage: predicted CLIP-Vision (\hat{c}_{im}) and CLIP-text (\hat{c}_{tx}) features are computed by giving test fMRI patterns (X_{test}) to the trained regression models. In the image-to-image pipeline of the latent diffusion model, VDVAE reconstructions of test images (the “initial guess” \hat{Y}_{low} from the first stage) are passed through the AutoKL Encoder of the pretrained Versatile Diffusion model, and the obtained latent vectors undergo 37 steps of the forward diffusion process (noise addition). The resulting noisy latent vectors are used to initialize the reverse diffusion process, which is also guided by predicted CLIP-Vision (\hat{c}_{im}) and CLIP-text (\hat{c}_{tx}) features jointly in a dual-guided framework. At last, the resulting denoised latent vector is passed through the AutoKL Decoder to generate the final reconstructed image (\hat{Y}_{test}). Note that all CLIP (vision and text encoders) and Versatile Diffusion layers (AutoKL encoder and decoder, forward and reverse diffusion blocks) are pretrained and frozen, only the brain-to-latent regression layers (blue boxes) are trained.

Materials and methods

Dataset. We used the publicly available Natural Scenes Dataset (NSD), a large-scale 7T fMRI dataset¹¹. The NSD was collected from 8 subjects viewing images from the COCO²¹ dataset. Each image was viewed for 3 seconds, while subjects were engaged in a continuous recognition task (reporting whether they had seen the image at any previous point in the experiment). For our study, we used the 4 subjects (sub1, sub2, sub5, sub7) who completed all trials. The training set that we used thus contained 8859 images and 24,980 fMRI trials (up to 3 repetitions for each image), and 982 images and 2770 fMRI trials for the test set. We averaged fMRI trials for the images that had multiple repetitions. We also used the corresponding captions from the COCO dataset. Test images are common for all subjects, while training images are different. We used the provided single-trial beta weights, obtained using generalized linear models with fitted hemodynamic response functions and additional GLMdenoise and ridge regression procedures (‘betas_fitrf_GLMdenoise_RR’). We masked preprocessed fMRI signals using the provided NSDGeneral ROI (Region-of-Interest) mask in 1.8 mm resolution. The ROI consists of [15,724, 14,278, 13,039, 12,682] voxels for the 4 subjects respectively, and includes many visual areas from the early visual cortex to higher visual areas. For further details on this dataset and the corresponding fMRI preprocessing steps, we refer the reader to the initial paper describing the Natural Scenes Dataset¹¹.

Low-level reconstruction of images using VDVAE (first stage). A Variational Auto-Encoder (VAE)³¹ is a generative model trained to capture an input distribution (such as an image dataset) via a low-dimensional latent space, constrained to follow a particular prior distribution (e.g. Gaussian). When the input dataset takes on a more complex distribution, training a Variational Autoencoder (VAE) can be challenging. Indeed, prior work has found that datasets consisting of natural scene images require many latent variables with complex distributions for which a simple VAE would not suffice; this is why the Very Deep Variational Autoencoder (VDVAE) was introduced³⁰. The VDVAE is a hierarchical VAE model, with several layers of conditionally dependent latent variables, each layer adding different details from coarse to fine when transitioning from top to bottom. The hierarchical dependence can be seen in Eqs. (1) and (2), where z indicates latent representations, x is the input variable, q_ϕ represents the approximate posterior distribution that is learned when training the encoder, and p_θ represents the prior distribution that is learned when training the decoder. The latent variable z_0 is at the top of the hierarchy with the smallest dimension (low resolution, with coarse details) and z_N is at the bottom of the hierarchy with the largest dimension (high resolution, with fine details). Equation (1) shows that the latent variables at the bottom of the hierarchy are dependent on those who are at the top (and on the input x). When there is no input (x), it is still possible to generate samples using the prior distribution described in Eq. (2).

This hierarchical structure helps the VDVAE learn sufficiently expressive latent variables to represent complex distributions like natural scene images.

$$q_{\phi}(\mathbf{z} | \mathbf{x}) = q_{\phi}(\mathbf{z}_0 | \mathbf{x})q_{\phi}(\mathbf{z}_1 | \mathbf{z}_0, \mathbf{x}) \dots q_{\phi}(\mathbf{z}_N | \mathbf{z}_{<N}, \mathbf{x}) \quad (1)$$

$$p_{\theta}(\mathbf{z}) = p_{\theta}(\mathbf{z}_0)p_{\theta}(\mathbf{z}_1 | \mathbf{z}_0) \dots p_{\theta}(\mathbf{z}_N | \mathbf{z}_{<N}) \quad (2)$$

For our study, we used the model provided in³⁰, trained on a 64×64 resolution ImageNet dataset, and consisting of 75 layers; we only utilized the latent variables from the first 31 layers for the sake of size in regression, since we observed that adding further layers did not make much difference in the reconstruction results (at test time, the latent variables from the remaining layers are sampled according to the prior distribution given in Eq. (2)).

In the training stage, we fed images to the encoder part of the VDVAE to extract latent variables for each training image (as described in Fig. 1). We concatenated the latent variables from the 31 layers, which resulted in 91,168-dim vectors. Then, we trained a ridge regression model between fMRI training patterns and these concatenated vectors. In the testing stage, we provided test fMRI patterns to the trained regression model and thus predicted latent values for each test image. Then, we fed those latent values to the decoder part of the VDVAE and obtained reconstructed images (64×64 pixels) from the VDVAE. These low-level reconstructions served as an “initial guess” for the diffusion model (second stage).

Final reconstruction of images using versatile diffusion (second stage). Although the VDVAE was helpful to reconstruct the layout of the image, it is not sufficient for the high-level features, nor does it produce fully naturalistic pictures. For that, we use the Versatile Diffusion²⁸ model in the second stage of our reconstruction framework. Versatile Diffusion is a recently proposed latent diffusion model (LDM)²⁴.

LDMs have become highly popular after their success in high-resolution text-to-image generation. In order to train an LDM, first an autoencoder (with encoder $E(\cdot)$ and decoder $D(\cdot)$) is trained on a large-scale image dataset to learn a compressed representation of images x_0 , i.e. a latent space $z_0 = E(x_0)$. Then, the forward diffusion process is applied to these latent variables z_0 by adding Gaussian noise in successive timesteps (described in Eq. (3), where t represents the timestep, α_t indicates a coefficient derived from the standard deviation of the Gaussian noise, and ε represents the Gaussian noise). The reverse diffusion process is learned via a neural network (Denoising U-Net in the original paper) to predict and remove noise from the noisy latent so as to retrieve the original latent variables. This is done by minimizing the loss function in Eq. (4), where ε is the true Gaussian noise, $\varepsilon_{\theta}(\cdot)$ represents the neural network being trained to predict the noise, z_t is the latent variable, t is the timestep, and $\tau_{\theta}(y)$ is the conditioning input for the U-Net. After the reverse diffusion process, the denoised latent variables are passed through the trained decoder $D(\cdot)$ to generate the images. The critical part of this process is that it is possible to condition this reverse diffusion process on different representations (e.g text captions, images, semantic maps). This conditioning process is done by merging conditions ($\tau_{\theta}(y)$) in the cross-attention block of the Denoising U-Net.

$$z_t = \sqrt{\alpha_t}z_0 + \sqrt{1 - \alpha_t}\varepsilon \quad (3)$$

$$L_{\text{LDM}} = \mathbb{E}_{t, z_0, \varepsilon, y} \left[\|\varepsilon - \varepsilon_{\theta}(z_t, t, \tau_{\theta}(y))\|^2 \right] \quad (4)$$

The Versatile Diffusion model (see Fig. 2) is a latent diffusion model with different pathways which allow us to condition the generation process on both text and image features to guide the reverse diffusion process. It is possible to provide CLIP-Vision, CLIP-Text, or both features as conditions in the reverse diffusion process. It is also possible to initialize the reverse diffusion with latent variables obtained from a particular image, rather than from a purely random distribution—this is the image-to-image pipeline that we will use to take advantage of our “initial guess” image from stage 1. The Versatile Diffusion model that we utilized in our framework was trained on the Laion2B-en³² dataset with 512×512 resolution images and corresponding captions. CLIP (Contrastive Language-Image Pre-training)²⁹ is a multimodal model designed to assist in different tasks that involve natural language processing and computer vision. It is trained in a contrastive learning approach, where features gathered from images vs. text captions are projected onto separate latent spaces of identical dimensions: CLIP-V refers to the latent space for images and CLIP-T for captions. Similarity scores (e.g. cosine similarity) of the latent space projections for matching images and captions are optimized throughout training. CLIP is widely used as a feature extractor, due to its high representational capabilities. The CLIP network used in Versatile Diffusion is based on the transformer architecture (ViT-L/14) and pretrained on a large-scale contrastive task.

In stage 2, we thus train two regression models, the first one between fMRI patterns and CLIP-Vision features (with 257×768 -dim extracted from the corresponding images where the first vector with 768-dim represents the category-related embedding and the remaining 256 embeddings represent the patches acquired from the images) and the second one between fMRI patterns and CLIP-Text features (77×768 -dim extracted from the COCO captions associated with the corresponding images where the 77 embeddings correspond to the number of tokens given to the model as inputs). At testing time, we use the image-to-image pipeline of the latent diffusion model. First, we encode the image reconstructed with the VDVAE model (stage 1) with the AutoKL Encoder (after upsampling the image from 64×64 to 512×512) and add noise to the latent vector for 37 steps of forward diffusion (corresponding to 75% of the 50 steps of full diffusion, which is a commonly used value in the image-to-image pipeline of LDMs). In this image-to-image pipeline, it is necessary to first add some amount of noise to the latent values using forward diffusion, since LDMs generate images via denoising using reverse diffusion (without noise on the image, the reverse diffusion step would end up with no change). Then, we feed

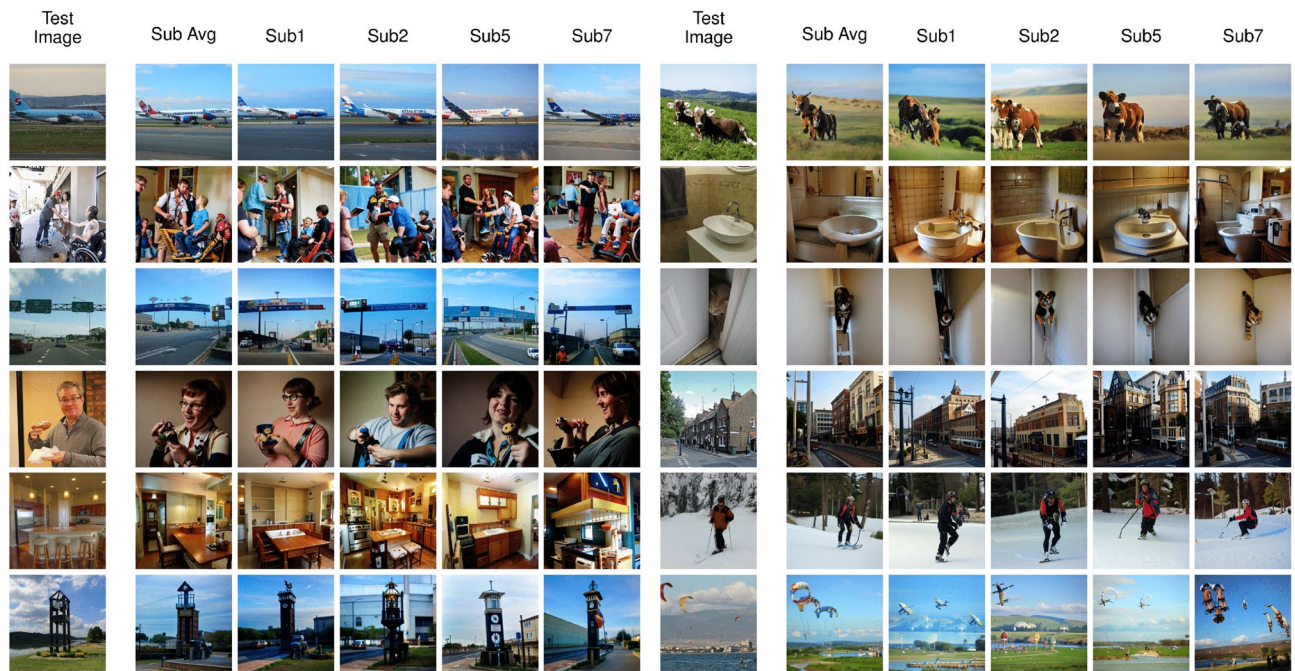


Figure 3. Examples of fMRI reconstructions from our Brain-Diffuser model. The first column is the groundtruth image (Test Image). The second column is generated by averaging the predicted latent variables over all 4 subjects seeing the same picture (Sub Avg). The remaining columns are for each individual subject (Sub1, Sub2, Sub5, Sub7).

this noisy latent as initialization to the diffusion model and denoise it for 37 steps while conditioning with the predicted CLIP-Vision and CLIP-Text features (stage 2). In every step of reverse diffusion, we use CLIP-Vision and CLIP-Text jointly in the double-guided diffusion pipeline of Versatile Diffusion, where the cross-attention matrices for both conditions are mixed through linear interpolation (with CLIP-Vision having a relative strength of 0.6 and CLIP-Text of 0.4). The diffusion result is passed through the AutoKL Decoder to produce our final 512×512 pixel reconstruction.

Results and analyses

Image reconstruction examples. We present examples of reconstructions from our model in Fig. 3. While we present the results of each individual subject in different columns, we also added results gathered by averaging the latent variables predicted by all subjects. In general, we see that reconstructed images capture most of the layout and semantics of the groundtruth images, while there remain differences in pixel-level details. For instance, looking specifically at the first four images on the left, we see that the reconstructed pose (3D orientation) of the plane (first image) is correct for every subject although there are some differences in the details of the plane and also in the texture of the background. Nonetheless, the fact that a commercial plane on a runway, facing to the right on a blue sky background was reconstructed in all instances is not a trivial feat. For the second example, all reconstructed images display a group of people, although layouts tend to differ. Still, a person in a wheelchair is visible in the bottom right corner for three of the four subjects. For the third image, the model reconstructed a highway with road signs correctly, although the orientation of the road is different for some of the subjects, and the details of the signs are not entirely matched. On the fourth sample, all reconstructed images show a single person facing left and holding an object in their hand, as in the groundtruth image. The person's details (gender, age, clothing) are different across subjects, e.g. with glasses only reconstructed for subject 1 and in the average across subjects. Reconstructed image contrast also differs from the ground truth. Similar conclusions can be generalized to most images of the test set: while never passing for a picture-perfect copy, with visible differences in especially color and contrast (due to inherent limitations of the Versatile Diffusion model in this respect), the reconstructed images are always naturalistic (that is, as much as diffusion models can generate) and plausible alternate renditions of the ground truth. Some of the remaining errors and differences may be caused by inherent limitations of the LDMs instead of unsuccessful predictions made by the fMRI-latent mapping model, as it is known that (current) diffusion models can generate unrealistic images in some occasions (e.g. unusual numbers of eyes on faces, fingers in hands).

We also present some examples of reconstruction failures from our model in Fig. 4. In these examples, we see that our model can fail due to different reasons. In the first example, although Brain-diffuser reconstructs oval objects around the center, the complex texture of the background seems to interfere with the object, which is not reconstructed as a clock. For the second example, the reconstructions show sea in the background, although there is no sea in the ground-truth image. On the fourth sample, the teddy bear occluding the kid's face seems to confuse the model, as it generates human faces in the reconstructions. For the sixth example, Brain-Diffuser

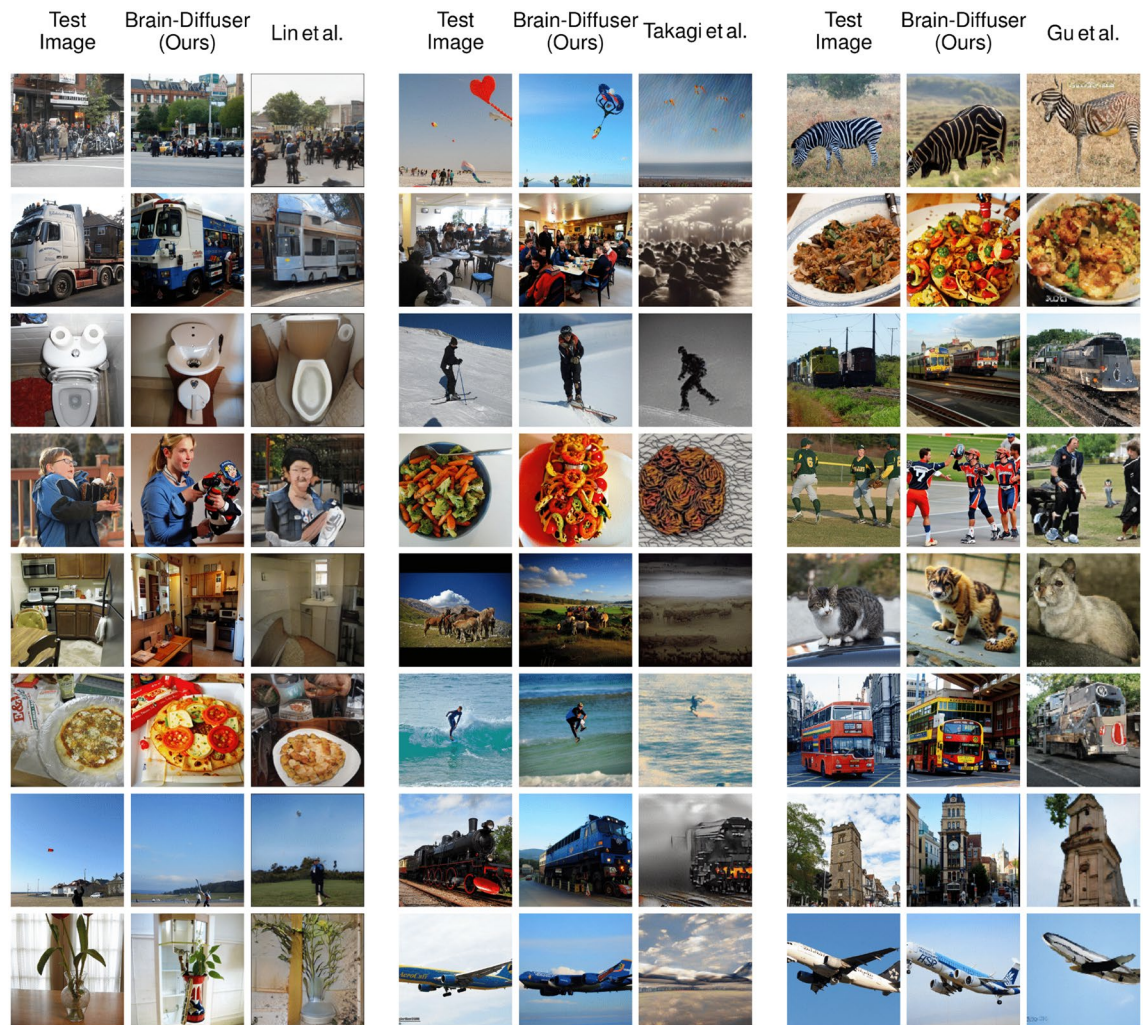


Figure 6. Comparison of fMRI reconstructions for different models on images presented in the papers of the previous methods. Since the presented test images in all methods were different, we did comparisons separately for each model. On the left (first 3 columns), we present the comparison of our model with Lin et al. together with groundtruth test images. On the center (columns 4–6), we present the comparison of our model with Takagi et al. together with groundtruth test images. On the right (last 3 columns), we present the comparison of our model with Gu et al. together with groundtruth test images.

Finally, Gu et al.²³ used an Instance-Conditioned GAN model trained on ImageNet. In Fig. 5, we compared our results with previous studies for the same set of images as in Fig. 3. Since Lin et al. used a different train-test split for their model, we used a replication of their model on the same train-test split as ours. Takagi et al. and Gu et al. shared generated images from their models with us for comparison. From these reconstructions, we can see that all methods capture high-level information to a degree, but not all of them are equally good at utilizing this information for image reconstruction. For instance, in the first image with a plane, Brain-Diffuser reconstructed a plane image that looks more similar to the ground-truth image and has a more realistic structure compared to Lin et al. and Gu et al. (and there is no recognizable plane in the reconstruction of Takagi et al.). In the fourth image with a man with glasses, the face is barely recognizable for Lin et al., the reconstruction by Gu et al. does not contain arms or glasses, while Takagi et al. reconstruct an unnatural rendition of a face and arms; in contrast, Brain-Diffuser exhibits a more natural-looking reconstruction and also manages to reconstruct the glasses.

In Fig. 6, we compare our method against the same three baselines, but using image reconstructions that were reported by the original authors in their papers, which might be more representative of each method's performance. Although Lin et al. seems to be performing better than the other two prior models, in some instances the quality of their reconstructions still lags behind ours. For instance, in the second image, the details of the truck are better represented in our model, while for the third image, the shape of the toilet is better represented in Lin et al. In the fourth image, the color of the clothes is presented more accurately in our model, as well as the fact that the person is holding an item; the person's face also looks more realistic compared to Lin et al. On the other hand, the color and location of the pizza in the sixth image appear more aligned with the ground-truth image for Lin et al. Although Takagi et al. generates easily recognizable silhouettes, they do not seem to perform as well as our model in any qualitative aspect including low-level details, semantics, or naturalness. Finally, when

Method	Quantitative measures							
	Low-level				High-level			
	PixCorr \uparrow	SSIM \uparrow	AlexNet(2) \uparrow	AlexNet(5) \uparrow	Inception \uparrow	CLIP \uparrow	EffNet-B \downarrow	SwAV \downarrow
Lin et al. ¹²	–	–	–	–	78.2%	–	–	–
Takagi et al. ²²	–	–	83.0%	83.0%	76.0%	77.0%	–	–
Gu et al. ²³	0.150	0.325	–	–	–	–	0.862	0.465
Brain-Diffuser (Ours)	0.254	0.356	94.2%	96.2%	87.2%	91.5%	0.775	0.423

Table 1. Quantitative analysis of fMRI reconstructions. For each measure, the best value is in bold (for PixCorr, SSIM, AlexNet(2), AlexNet(5), Inception and CLIP metrics, higher is better. For EffNet-B and SwAV distances, lower is better. This is indicated by the arrow pointing up or down, respectively).

Method	Quantitative measures							
	Low-level				High-level			
	PixCorr \uparrow	SSIM \uparrow	AlexNet(2) \uparrow	AlexNet(5) \uparrow	Inception \uparrow	CLIP \uparrow	EffNet-B \downarrow	SwAV \downarrow
Only-VDVAE	0.358	0.437	97.7%	97.6%	77.0%	71.1%	0.906	0.581
Brain-Diffuser w/o VDVAE	0.143	0.302	85.6%	93.0%	87.3%	92.6%	0.775	0.414
Brain-Diffuser w/o CLIP-Text	0.279	0.333	95.6%	97.0%	87.9%	91.2%	0.796	0.436
Brain-Diffuser w/o CLIP-Vision	0.327	0.433	93.9%	94.1%	84.7%	84.5%	0.821	0.509
Brain-Diffuser	0.305	0.367	96.7%	97.4%	87.8%	92.5%	0.768	0.415

Table 2. Quantitative comparisons of test fMRI reconstructions of Sub1 with various ablations of the full model. For each measure, the best value is in bold (for PixCorr, SSIM, AlexNet(2), AlexNet(5), Inception, and CLIP metrics, higher is better. For EffNet-B and SwAV distances, lower is better. This is indicated by the arrow pointing up or down, respectively).

we compare our results to Gu et al., we can see that, although both appear good at reconstructing images with similar semantics, structural aspects are less well represented in their reconstructed images (e.g. unrealistic warped shapes for the train, bus, and building). In contrast, the shape and texture details of our model are more realistic. Since their model has a BigGAN backbone, with few parameters to encode the entire layout of the image (including the object's class, its pose, size, and location), and since it is trained on a single-object-centric dataset (ImageNet), the model seems to be limited in reconstructing complex scenes with multiple objects. On the other hand, since LDMs include a spatially organized map of features, it is more convenient for them to represent multiple objects; as an example, we see one train in the third image although there are two trains in the groundtruth image, and in the reconstructed image from our model.

Quantitative results. To make the comparison with other models more quantitative, we present the results of 8 different image quality metrics in Table 1. PixCorr is the pixel-level correlation of reconstructed and groundtruth images. SSIM³³ is the structural similarity index metric. AlexNet(2) and AlexNet(5) are the 2-way comparisons of the second and fifth layers of AlexNet³⁴, respectively. Inception is the 2-way comparison of the last pooling layer of InceptionV3³⁵. CLIP is the 2-way comparison of the output layer of the CLIP-Vision²⁹ model. EffNet-B and SwAV are distance metrics gathered from EfficientNet-B1³⁶ and SwAV-ResNet50³⁷ models, respectively. The first four can be considered as low-level metrics, while the last four reflect higher-level properties. For PixCorr and SSIM metrics, we downsampled generated images from 512×512 resolution to 425×425 resolution (i.e. the resolution of groundtruth images in NSD dataset). For the rest of the measures, generated images are preprocessed according to the input properties of each network. Note that not all measures are available for each previous model (depending on what they chose to report). However, each model has at least one point of comparison with ours. Our quantitative comparisons with Takagi et al. and Gu et al. are made according to the exact same test set, i.e., the 982 images that are common for all 4 subjects. Lin et al., on the other hand, reported their results on only Subject 1 and with a custom train-test set split. However, when measuring our model's image quality on the same train-test split as Lin et al, we observed nearly identical results (Inception Score of 87.0%, compared to 78.2% for Lin et al). Our model is the best-performing model by a decent margin for all of the quantitative metrics. Overall, these results show that our model can be considered state-of-the-art for both low-level and high-level quantitative measures.

Ablation studies.

In order to reveal the contribution of each component of Brain-Diffuser, we performed an ablation study (with fMRI data of Sub1), and report both quantitative (Table 2) and qualitative (Fig. 7) results. The quantitative results are given in Table 2. Our first ablation (Only-VDVAE) considers the results from stage-1 reconstruction only (Fig. 1) without stage-2 reconstruction (Fig. 2). This Only-VDVAE model provides the best results for all low-level measures, but the worst (by a large margin) for all high-level measures. This pattern of results is expected since the VDVAE reconstruction relies on low-level information without a contribution of semantic information

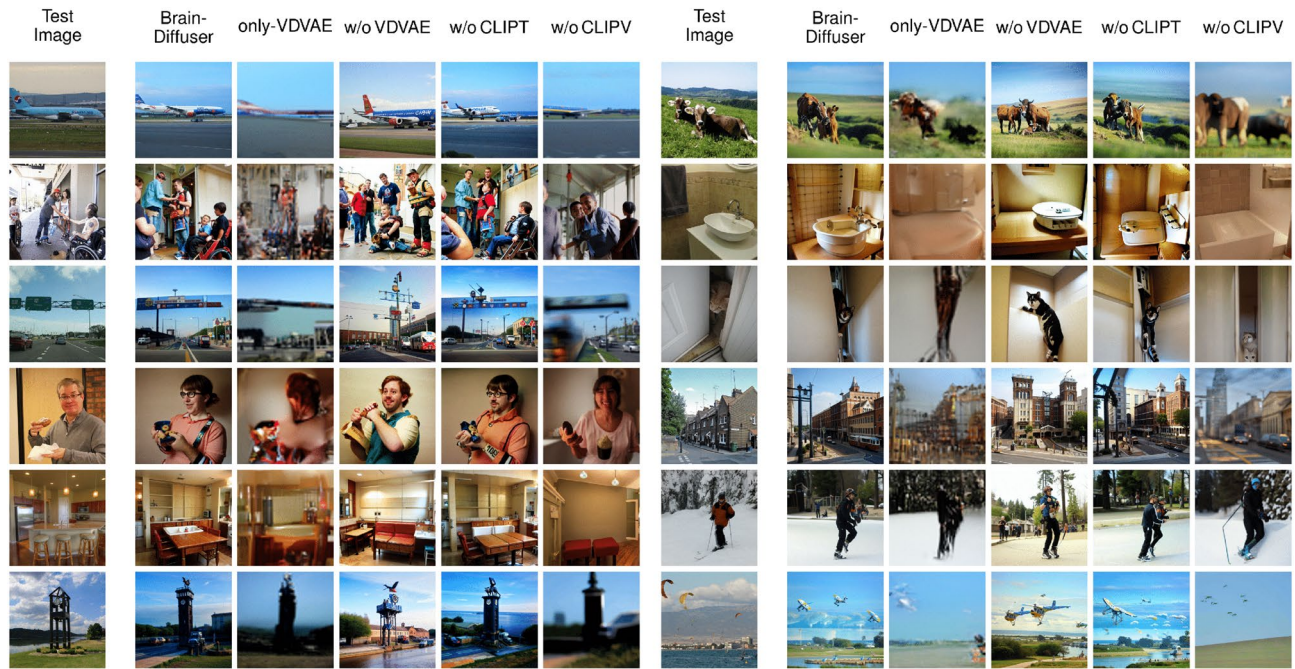


Figure 7. Examples of fMRI test reconstructions from Sub1 with various ablations of the full model. The first column is the groundtruth image (Test Image). The second column shows reconstructions from the full Brain-Diffuser model with all of its components. The third column is for reconstructions of the Only-VDVAE model. The remaining columns are for Brain-Diffuser with one of its components excluded, in order: without VDVAE, without CLIP-Text, and without CLIP-Vision.

from stage-2. By contrast, Brain-Diffuser without the VDVAE component (i.e., stage-2 reconstruction but with random initialization of the autoKL latent vector) performs worst on low-level measures (by a large margin), while it is among the best in high-level measures. This is also reasonable since this ablated model generates the reconstructions from only high-level features obtained from CLIP-Text and CLIP-Vision models and does not have much information about low-level information such as layout. Together, these results indicate that the VDVAE “initial guess” (stage-1) is necessary but not sufficient for optimal reconstruction. This is evident in the results from the full Brain-Diffuser model (last row in Table 2), where the contribution from VDVAE (stage 1) brings strong improvements in low-level measures, with near-optimal high-level features. In another ablation, we evaluate Brain-Diffuser without CLIP-Text. Compared to the full model, there is a sizeable decrement in both low-level and high-level measures, except Inception. While the contribution of CLIP-Text to the reconstruction of high-level semantic features is expected, its improvement of low-level measures is more surprising but could be explained by semantic information related to the image layout itself, such as the number or orientation of objects (see examples in Fig. 7). Finally, Brain-Diffuser without CLIP-Vision, surprisingly, retains high performance on the low-level PixCorr and SSIM measures (lower than Only-VDVAE, but higher than the full model); we assume that this could be due to insufficient diffusion steps (as discussed further below), preventing the reconstruction from deviating from the VDVAE initial guess. For all other measures (including low-level AlexNet measures), removing CLIP-Vision guidance severely impairs the performance of Brain-Diffuser. Overall, when jointly considering low-level and high-level measures, these quantitative results show that the full Brain-Diffuser model is better than any other variation or ablation.

We also present qualitative results in Fig. 7 with the same set of images presented in Fig. 3 of the main manuscript. Reconstructions from the Only-VDVAE model (i.e., stage-1 without stage-2) match the low-level details (e.g. shapes, layouts) of the groundtruth images, but they look like vague silhouettes rather than natural images. In contrast, Brain-Diffuser without VDVAE generates images that match high-level properties (semantics) of groundtruth images but lack positional information about the objects and their layout. This is particularly clear for the fourth image in the right part of the figure, where the layout of the street and buildings is properly captured by VDVAE (and thus, also by the full model), but is lacking in the VDVAE ablation. The images generated by Brain-Diffuser without CLIP-Text appear very close to those from the full model but with some notable differences. One example is the ski image (Row 5 on the right part of the figure), where the full model generates a single person (as in the groundtruth) while the model without CLIP-Text generates two people. Another example is the plane image (Row 1 on the left part of the figure) where the model without CLIP-Text does not produce an image with the correctly positioned plane. Finally, reconstructions from Brain-Diffuser without CLIP-Vision appear quite blurry, and somehow in between the Only-VDVAE and the full model reconstructions. This could be an indication that forward and reverse diffusion steps were not sufficient for this model. Still, increasing the number of diffusion steps may not be a good solution since that would cause the model to lose low-level information provided by VDVAE. Overall, these qualitative examples corroborate the quantitative findings in

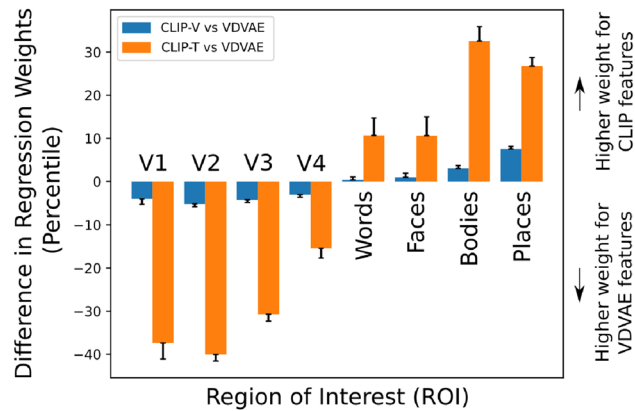


Figure 8. Difference between the percentiles of the regression weights (L_1 norm) for the CLIP features (CLIP-V and CLIP-T) vs. the VDVAE features, averaged over voxels in each ROI and normalized by the average percentile of VDVAE features for the same ROI. Positive values indicate relatively higher regression weight for CLIP features compared to the VDVAE features, and vice versa. Error bars represent the standard error of the mean across 4 subjects.

Table 2 and make it clear that the Brain-Diffuser model represents the optimal compromise for both low-level details and high-level semantic features.

Which brain regions are used? In order to understand the relationship between brain regions and the various components of our model (VDVAE, CLIP-Vision, CLIP-Text), we performed a region-of-interest (ROI) analysis of the regression weights. We used 4 visual ROIs derived from population receptive field (pRF) experiments, and 4 ROIs derived from functional localization (fLoc) experiments. All experiments were provided along with the NSD dataset by the original authors. These ROIs are as follows (region names following the terminology adopted in Allen et al.¹¹): V1 is the concatenation of V1 ventral (V1v) and V1 dorsal (V1d), and similarly for V2 and V3; V4 is the human V4 (hV4); the Face-ROI consists of the union of OFA, FFA-1, FFA-2, mTL-faces, and aTL-faces; Word-ROI consists of OWFA, VWFA-1, VWFA-2, mfs-words, and mTL-words; Place-ROI consists of OPA, PPA, and RSC; Body-ROI consists of EBA, FBA-1, FBA-2, and mTL-bodies. For each voxel in these regions, we computed the strength of the regression weights (L_1 norm) for the CLIP features (CLIP-V and CLIP-T) and the VDVAE features, expressed as a percentile. Because the absolute regression weights can be affected by the number of voxels in each region, as well as the overall activity level and the noise level, we report our results as a *difference* in regression weights between CLIP features and VDVAE features. The results in Fig. 8 show that early regions (V1–V4) are more informative about the VDVAE features, while category-specific higher brain regions (Words, Faces, Bodies, Places) carry more information about CLIP features. Another important observation is that the differences between CLIP-V and VDVAE are in the same direction, but much weaker than the differences between CLIP-T and VDVAE. This may indicate that although the Versatile Diffusion model uses CLIP-V features for high-level guidance, these features still contain more information about low-level properties than CLIP-T features.

ROI-optimal stimuli. Beyond brain decoding, we show here that our method can also be used to help understand the functional properties of specific regions-of-interest (ROIs) in the brain. Although we know from early studies in the neuroscience literature^{4,38–45} what sort of visual properties would best activate neurons in each brain region, there are only a few studies^{19,46–49} which attempted to directly visualize an “optimal” stimulus for a given brain region. Our method can easily be adapted for this purpose. We define “ROI-optimal” as images that would activate a certain ROI maximally while not activating other ROIs (or just activating them minimally). We analyzed the same 8 ROIs (V1, V2, V3, V4, Face-ROI, Word-ROI, Place-ROI, and Body-ROI) that we discussed in the previous section. We used the intersection of these regions with NSDGeneral (the one we used for training our decoding system), each time creating a synthetic fMRI pattern where the ROI was active (signal set to 1) and the rest of the brain inactive (signal set to 0). From this synthetic pattern, our system could then generate predicted latent variables, and directly reconstruct an equivalent visual scene, corresponding to the “ROI-optimal” image. Surprisingly, this simple and deterministic approach, inspired by the analysis in Ozeçelik et al.¹⁹, still gives plausible results. Since the synthetic fMRI patterns can be considered out-of-distribution (because there are no similar patterns in the training set), we re-normalized the generated latent variables to give them a similar euclidean norm to the training samples. This procedure helped the diffusion model to generate meaningful images that are shown in Fig. 9.

Upon inspecting the generated “ROI-optimal” images for visual ROIs, we see that V1 produces high-contrast scenes with very detailed textures extending to the visual periphery, such as trees and foliage in a park with numerous small human or animal figures. V2 is similar (especially for Sub1 and Sub5, which also display humans in a luxuriant garden environment), but with slightly broader elements and less peripheral detail (e.g. trays filled with various foods in Sub2, Sub7, and in the subject-average). Continuing along the same trend, V3 and V4

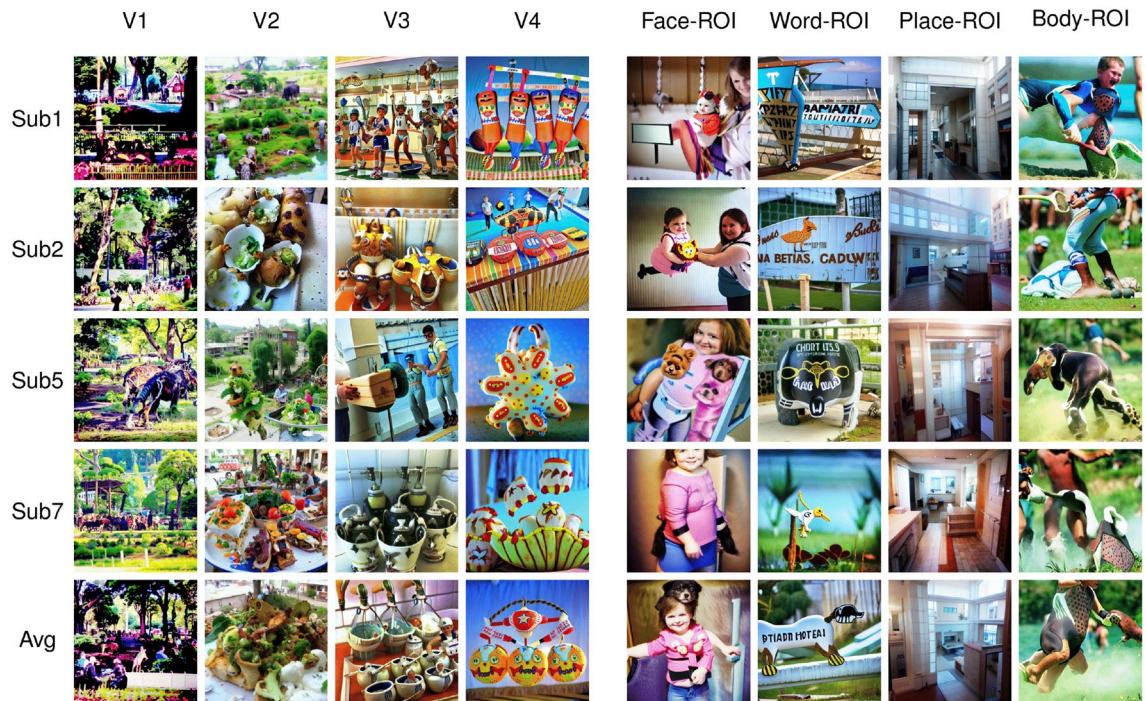


Figure 9. Images reconstructed from synthetic fMRI patterns created by activating specific regions-of-interests (ROIs). The first 4 rows present individual subjects: Sub1, Sub2, Sub5, and Sub7. The last row is generated by averaging the latent vectors predicted from all 4 subjects. The columns present ROIs: First four are ROIs from the visual cortex (V1–V4) gathered by population receptive field experiments, and the last four are ROIs that are specified with functional localization experiments (Face-ROI, Word-ROI, Place-ROI, Body-ROI). Since our synthetic fMRI patterns produce distribution shifts in the latent variables, which in turn can affect the contrast of the reconstructed images, histogram stretching and equalization are applied on color histograms of generated images for visualization purposes.

produce larger objects compared to the earlier regions, with repeating patterns and global structure. V4 especially generates colorful, high-contrast objects resembling toys on a bright background.

The ROI-optimal images for functionally defined high-level ROIs are even easier to interpret, as they tend to coincide with each region's known category preference. For instance, the model generated multiple face images for the Face-ROI, including humans and sometimes even animal faces (e.g. dogs in Sub5 and in the subject-average). For the Word-ROI, the model generated characters and pseudo-words on objects or signs (except for Sub7). Architecturally plausible indoor scene layouts were produced for the Place-ROI. Finally, for the Body-ROI, the reconstructed images show both human and animal body parts like arms and legs engaged in active behavior like sports or running.

In another exploratory experiment aiming to understand the effects of combining ROIs, we repeated the analysis of Fig. 9, this time using combinations of activations for different ROIs (Fig. 10). In the first column where we activated all the low-level visual regions together (V1–V4), we see scrambled and regular patterns in different parts of the images as well as some identifiable objects for some of the subjects—but there are no apparent objects that are commonly identifiable across all subjects. In the next columns, we combined activations across all low-level visual regions (V1–V4) and one of the functionally defined high-level ROIs. In the second column where we combined the visual regions (V1–V4) with Face-ROI, we see human and animal faces in all images, although some scrambled high-contrast patterns also continue to exist in different parts of images. In the third column where we combined the visual regions with Word-ROI, letter-like patterns or pseudo-words can be seen in the upper part of the image for subject 1 and middle and lower regions for subject 5, but they are less visible than in the analysis of Fig. 9. In the fourth column where we combined the visual regions with Place-ROI, the model generates architectural interior and exterior parts, and the scrambled patterns cease to exist for these images. Finally, in the fifth and last column where we combined the visual regions with Body-ROI, we see vaguely identifiable human and animal body parts like arms and legs. This proof-of-concept experiment reveals what happens when we combine the activations of different regions instead of activating one ROI in isolation. Although there are visual differences between generated images from Figs. 9 and 10, we continue to observe similar semantic relationships between the functional ROIs and the corresponding images.

While these results mainly confirm decades of converging knowledge from the neuroscience literature on neuronal selectivity in the ventral visual pathway, this method allowed us to directly visualize functional properties in vivid detail and high-resolution images. Furthermore, the technique introduced here could easily be extended to study retinotopic or eccentricity-based cortical organization. As a proof of concept, we also applied our ROI analysis to visual regions defined by different eccentricity preferences. Similar to hierarchical regions



Figure 10. Images reconstructed from synthetic fMRI patterns created by activating combinations of different regions-of-interests (ROIs). The first 4 rows present individual subjects: Sub1, Sub2, Sub5, and Sub7. The last row is generated by averaging the latent vectors predicted from all 4 subjects. The columns present different combinations of ROIs: The first column is where all four regions in the visual cortex are activated at once (V1, V2, V3, and V4). The remaining columns are combinations of activations of these visual ROIs with one of the functional ROIs: Face-ROI, Word-ROI, Place-ROI, and Body-ROI, respectively.

in the visual cortex (V1, V2, V3, and hV4) these eccentricity-based regions ($0^\circ < e < 0.5^\circ$, $0.5^\circ < e < 1^\circ$, $1^\circ < e < 2^\circ$, $2^\circ < e < 4^\circ$ and $4^\circ < e$, where e stands for “eccentricity”) were also extracted by population receptive field (pRF) experiments. These regions thus reflect the eccentricity preference of the retinotopic cortex, where degrees close to 0° indicate central vision (closer to the fovea) and higher degrees indicate peripheral vision. The corresponding results are shown in Fig. 11. It is difficult to see a clear pattern for eccentricities between 0° and 0.5° ($0^\circ < e < 0.5^\circ$), as the corresponding portion of the image might be too small to be considered meaningful for the model. A noticeable aspect, however, is that all images for that ROI have detailed and high-contrast objects in the center (though there are also objects in the periphery). For eccentricities between 0.5° and 1° ($0.5^\circ < e < 1^\circ$), and between 1° and 2° ($1^\circ < e < 2^\circ$), we begin to see larger objects (e.g. humans, animals, blobs) at the center of the images. When we reach eccentricities between 2° and 4° ($2^\circ < e < 4^\circ$) and beyond ($4^\circ < e$), we start to see these objects (or animals, humans, and blobs) move towards the periphery, while the center of the images is mostly empty. These results highlight two important findings: first, the latent representations used by Brain-Diffuser (combining latent features from VDVAE, CLIP-Vision, and CLIP-Text) are precise enough to convey information about the spatial localization of objects in the image; second, we see that Brain-Diffuser managed to learn the eccentricity-based retinotopic organization of the cortex from these representations.

Discussion

In this study, we designed a two-stage framework (Brain-Diffuser) that reconstructs images from fMRI patterns using generative models based on latent diffusion. In the first stage, we used the VDVAE model to generate “initial guess” reconstructions focusing on low-level details. Then in the second stage, we used the image-to-image pipeline of the Versatile Diffusion model, starting from this initial guess, to generate final reconstructions via diffusion, guided by both predicted CLIP-Vision and CLIP-Text features. As we relied on pre-trained and publicly available models for image generation (VDVAE, Versatile Diffusion) and multimodal feature extraction (CLIP), our method only required training ridge regression models from multivoxel brain activity to the relevant model latent spaces (Figs. 1, 2).

We analyzed the results both qualitatively (Fig. 3) and quantitatively (Table 1) We observed that reconstructed scene images generated by Brain-Diffuser, although not perfectly identical to groundtruth images, preserve most of the layout and semantic information. They also appear more naturalistic compared to reconstructions from earlier studies (Fig. 6). When evaluated quantitatively, we saw that Brain-Diffuser outperforms previous models in both high-level and low-level metrics. After advancing the state-of-the-art in image generation applications^{24–28}, it appears that latent diffusion models can also be used to improve the state-of-the-art in fMRI-based image reconstruction.

Although latent diffusion models are very recent²⁴, we noted at least two competing studies that used LDMs for fMRI-based image reconstruction. Chen et al.²⁰ proposed MinD-Vis, a method based on an LDM conditioned on image category labels (rather than text captions) to reconstruct images from the Kamitani dataset. As



Figure 11. Images reconstructed from synthetic fMRI patterns created by activating regions-of-interests (ROIs) in the visual cortex with different eccentricities. The first 4 rows present individual subjects: Sub1, Sub2, Sub5, and Sub7. The last row is generated by averaging the latent vectors predicted from all 4 subjects. The columns present concentric regions with increasing eccentricity coverage ($0^\circ < e < 0.5^\circ$, $0.5^\circ < e < 1^\circ$, $1^\circ < e < 2^\circ$, $2^\circ < e < 4^\circ$, and $4^\circ < e$, where “e” stands for eccentricity). Histogram stretching and equalization is applied for visualization purposes.

mentioned above, this is a less challenging, single-object-centered dataset; thus, their results cannot be directly compared with ours, obtained using the richer and more complex NSD dataset. Takagi et al.²², on the other hand, used NSD and were thus included in our quantitative comparisons. There are multiple possible reasons why our model performed better than theirs, on both low-level and high-level metrics. Beyond the use of distinct pretrained LDMs (Stable Diffusion²⁴ for Takagi et al. vs. Versatile Diffusion²⁸ in our study), our framework contains several improvements such as the use of VDVAE reconstructions for low-level details (Fig. 1) and the dual conditioning on CLIP-Vision and CLIP-Text features (Fig. 2), which together resulted in better qualitative and quantitative results.

There are several ways in which this work may be pursued in the future. First of all, it will be important to test and validate our method on other image-fMRI datasets. As deep generative models will likely continue to improve at a breakneck pace, it is probable that there will soon come models better suited for complex scene reconstruction from fMRI signals. Of course, among a pool of many generative models, it may not be a trivial task to select the most appropriate ones and to experiment on them, and adapt them for brain decoding and image reconstruction. If future generative models reach a ceiling in their ability to linearly explain brain activity, we may need to look for better alternatives than just doing ridge regression between fMRI patterns and latent variables. These alternatives (non-linear regressions, deep hierarchical networks), however, may require larger training datasets to learn the correspondence between fMRI patterns and visual features than ridge regression. It is possible that our reconstructions would benefit from including larger brain regions (or even the whole-brain) in our analysis. However, this is not guaranteed due to the possibility of overfitting in the presence of high-dimensional inputs. Also, expanding the region of fMRI inputs would dramatically raise the computational cost (in time and/or memory) of the training process of regression models. This is why the NSDGeneral ROI appears as an ideal compromise used in most studies (including ours). Using a common ROI also facilitates comparisons between studies. In the future, we may also see more accurate movie reconstruction studies that process temporal patterns together with spatial ones on movie-fMRI datasets^{50,51}. Besides improving the reconstruction quality, future work could also design novel experiments and analyses on the NSD dataset using generative models. For instance, in this study, we have shown that we can use generative models to reveal the “optimal” stimulus for anatomically, functionally, or retinotopically-defined ROIs, by analyzing the reconstructions of synthetic fMRI patterns created from the corresponding ROI masks. This approach could easily be extended to probe less well-known regions of the visual cortex, to help settle theoretical arguments about distinct sub-regions of (e.g.) the face processing network, or to render images for arbitrary combinations of ROIs (e.g., what image would optimally activate V1, V4, and the face-ROI, but not V2 or the Body-ROI; see also Fig. 10). Important advances in this direction were made using an iterative optimization method by Gu et al.⁴⁸ Directly passing synthetic fMRI patterns to the image reconstruction pipeline, however, is computationally advantageous, which may prove important when there are numerous combinations of sub-regions to be tried. Similar “virtual experiments” in this framework could help us address outstanding questions in neuroscience, and understand the organization of sensory and semantic knowledge in the brain.

Data availability

The information on accessing the NSD dataset that is analyzed during the current study is available in the Natural Scenes Dataset repository, <http://naturalscenesdataset.org/>.

Code availability

The code for our project, including scripts to train regression models, pretrained weights, and scripts to produce reconstructions for test images and for ROI-based synthetic patterns, is publicly available at <http://github.com/ozcelikfu/brain-diffuser>.

Received: 23 March 2023; Accepted: 15 September 2023

Published online: 20 September 2023

References

- Thirion, B. *et al.* Inverse retinotopy: Inferring the visual content of images from brain activation patterns. *Neuroimage* **33**, 1104–1116 (2006).
- Kamitani, Y. & Tong, F. Decoding the visual and subjective contents of the human brain. *Nat. Neurosci.* **8**, 679–685 (2005).
- Haynes, J.-D. & Rees, G. Predicting the orientation of invisible stimuli from activity in human primary visual cortex. *Nat. Neurosci.* **8**, 686–691 (2005).
- Haxby, J. V. *et al.* Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science* **293**, 2425–2430 (2001).
- Cox, D. D. & Savoy, R. L. Functional magnetic resonance imaging (fMRI) “brain reading”: Detecting and classifying distributed patterns of fMRI activity in human visual cortex. *Neuroimage* **19**, 261–270 (2003).
- Kay, K. N., Naselaris, T., Prenger, R. J. & Gallant, J. L. Identifying natural images from human brain activity. *Nature* **452**, 352–355 (2008).
- Miyawaki, Y. *et al.* Visual image reconstruction from human brain activity using a combination of multiscale local image decoders. *Neuron* **60**, 915–929 (2008).
- VanRullen, R. & Reddy, L. Reconstructing faces from fMRI patterns using deep generative neural networks. *Commun. Biol.* **2**, 1–10 (2019).
- Dado, T. *et al.* Hyperrealistic neural decoding for reconstructing faces from fMRI activations via the gan latent space. *Sci. Rep.* **12**, 141 (2022).
- Shen, G., Horikawa, T., Majima, K. & Kamitani, Y. Deep image reconstruction from human brain activity. *PLoS Comput. Biol.* **15**, e1006633 (2019).
- Allen, E. J. *et al.* A massive 7t fmri dataset to bridge cognitive neuroscience and artificial intelligence. *Nat. Neurosci.* **25**, 116–126 (2022).
- Lin, S., Sprague, T. C. & Singh, A. Mind reader: Reconstructing complex images from brain activities. In *Advances in Neural Information Processing Systems* (Oh, A. H., Agarwal, A., Belgrave, D. & Cho, K. Eds.) (2022).
- Horikawa, T. & Kamitani, Y. Generic decoding of seen and imagined objects using hierarchical visual features. *Nat. Commun.* **8**, 1–15 (2017).
- Deng, J. *et al.* Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, 248–255 (IEEE, 2009).
- Belyi, R. *et al.* From voxels to pixels and back: Self-supervision in natural-image reconstruction from fMRI. *Adv. Neural Inf. Process. Syst.* **32**, 25 (2019).
- Gaziv, G. *et al.* Self-supervised natural image reconstruction and large-scale semantic classification from brain activity. *Neuroimage* **254**, 119121 (2022).
- Mozafari, M., Reddy, L. & VanRullen, R. Reconstructing natural scenes from fmri patterns using bigbigan. In *2020 International Joint Conference on Neural Networks (IJCNN)*, 1–8 (IEEE, 2020).
- Ren, Z. *et al.* Reconstructing seen image from brain activity by visually-guided cognitive representation and adversarial learning. *Neuroimage* **228**, 117602 (2021).
- Ozcelik, F., Choksi, B., Mozafari, M., Reddy, L. & VanRullen, R. Reconstruction of perceived images from fmri patterns and semantic brain exploration using instance-conditioned gans. In *2022 International Joint Conference on Neural Networks (IJCNN)*, 1–8 (IEEE, 2022).
- Chen, Z., Qing, J., Xiang, T., Yue, W. L. & Zhou, J. H. Seeing beyond the brain: Conditional diffusion model with sparse masked modeling for vision decoding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 22710–22720 (2023).
- Lin, T.-Y. *et al.* Microsoft coco: Common objects in context. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V* 13, 740–755 (Springer, 2014).
- Takagi, Y. & Nishimoto, S. High-resolution image reconstruction with latent diffusion models from human brain activity. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14453–14463 (2023).
- Gu, Z., Jamison, K., Kuceyeski, A. & Sabuncu, M. R. Decoding natural image stimuli from fMRI data with a surface-based convolutional network. In *Medical Imaging with Deep Learning* (2023).
- Rombach, R., Blattmann, A., Lorenz, D., Esser, P. & Ommer, B. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10684–10695 (2022).
- Ramesh, A., Dhariwal, P., Nichol, A., Chu, C. & Chen, M. Hierarchical text-conditional image generation with clip latents. [arXiv:2204.06125](https://arxiv.org/abs/2204.06125) (arXiv preprint) (2022).
- Nichol, A. Q. *et al.* Glide: Towards photorealistic image generation and editing with text-guided diffusion models. In *International Conference on Machine Learning*, 16784–16804 (PMLR, 2022).
- Saharia, C. *et al.* Photorealistic text-to-image diffusion models with deep language understanding. In *Advances in Neural Information Processing Systems* (Oh, A. H., Agarwal, A., Belgrave, D. & Cho, K. eds.) (2022).
- Xu, X., Wang, Z., Zhang, E., Wang, K. & Shi, H. Versatile diffusion: Text, images and variations all in one diffusion model. [arXiv:2211.08332](https://arxiv.org/abs/2211.08332) (arXiv preprint) (2022).
- Radford, A. *et al.* Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, 8748–8763 (PMLR, 2021).
- Child, R. Very deep {vae}s generalize autoregressive models and can outperform them on images. In *International Conference on Learning Representations* (2021).
- Kingma, D. P. & Welling, M. Auto-encoding variational bayes. [arXiv:1312.6114](https://arxiv.org/abs/1312.6114) (arXiv preprint) (2013).
- Schuhmann, C. *et al.* Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. [arXiv:2111.02114](https://arxiv.org/abs/2111.02114) (arXiv preprint) (2021).

33. Wang, Z., Bovik, A. C., Sheikh, H. R. & Simoncelli, E. P. Image quality assessment: From error visibility to structural similarity. *IEEE Trans. Image Process.* **13**, 600–612 (2004).
34. Krizhevsky, A., Sutskever, I. & Hinton, G. E. Imagenet classification with deep convolutional neural networks. *Commun. ACM* **60**, 84–90 (2017).
35. Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J. & Wojna, Z. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2818–2826 (2016).
36. Tan, M. & Le, Q. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International Conference on Machine Learning*, 6105–6114 (PMLR, 2019).
37. Caron, M. *et al.* Unsupervised learning of visual features by contrasting cluster assignments. *Adv. Neural. Inf. Process. Syst.* **33**, 9912–9924 (2020).
38. Hubel, D. H. & Wiesel, T. N. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J. Physiol.* **160**, 106 (1962).
39. Gross, C. G., Rocha-Miranda, C. D. & Bender, D. Visual properties of neurons in inferotemporal cortex of the macaque. *J. Neurophysiol.* **35**, 96–111 (1972).
40. Perrett, D., Rolls, E. & Caan, W. Visual neurones responsive to faces in the monkey temporal cortex. *Exp. Brain Res.* **47**, 329–342 (1982).
41. Gallant, J. L., Braun, J. & Van Essen, D. C. Selectivity for polar, hyperbolic, and cartesian gratings in macaque visual cortex. *Science* **259**, 100–103 (1993).
42. Van Essen, D. C. & Gallant, J. L. Neural mechanisms of form and motion processing in the primate visual system. *Neuron* **13**, 1–10 (1994).
43. Kanwisher, N., McDermott, J. & Chun, M. M. The fusiform face area: A module in human extrastriate cortex specialized for face perception. *J. Neurosci.* **17**, 4302–4311 (1997).
44. Epstein, R. & Kanwisher, N. A cortical representation of the local visual environment. *Nature* **392**, 598–601 (1998).
45. Orban, G. A., Van Essen, D. & Vanduffel, W. Comparative mapping of higher visual areas in monkeys and humans. *Trends Cogn. Sci.* **8**, 315–324 (2004).
46. Bashivan, P., Kar, K. & DiCarlo, J. J. Neural population control via deep image synthesis. *Science* **364**, eaav9436 (2019).
47. Ratan Murty, N. A., Bashivan, P., Abate, A., DiCarlo, J. J. & Kanwisher, N. Computational models of category-selective brain regions enable high-throughput tests of selectivity. *Nat. Commun.* **12**, 5540 (2021).
48. Gu, Z. *et al.* Neurogen: Activation optimized image synthesis for discovery neuroscience. *Neuroimage* **247**, 118812 (2022).
49. Mueller, K. N., Carter, M. C., Kansupada, J. A. & Ponce, C. R. Macaques recognize features in synthetic images derived from ventral stream neurons. *Proc. Natl. Acad. Sci.* **120**, e2213034120 (2023).
50. Wang, C. *et al.* Reconstructing rapid natural vision with fmri-conditional video generative adversarial network. *Cereb. Cortex* **32**, 4502–4511 (2022).
51. Kupersmidt, G., Belyi, R., Gaziv, G. & Irani, M. A penny for your (visual) thoughts: Self-supervised reconstruction of natural movies from brain activity. [arXiv:2206.03544](https://arxiv.org/abs/2206.03544) (arXiv preprint) (2022).

Acknowledgements

This work was funded by the Agence Nationale de la Recherche ANR Grants AI-REPS ANR-18-CE37-0007-01 and ANTI ANR-19-PI3A-0004, and the European Union (Grant ERC Advanced GLoW 101096017). Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or ERCEA. Neither the European Union nor the granting authority can be held responsible for them. We thank Yu Takagi and Zijin Gu for sharing their test results with us and also we thank Alex Nguyen, Paul Scotti, and MedARC team for helping acquire the replicated results of the Mind-Reader study.

Author contributions

F.O. and R.V. conceptualized the idea. F.O. analyzed the data, designed the framework, and generated the results. F.O. and R.V. analyzed and interpreted the results. F.O. wrote the original draft. F.O. and R.V. reviewed and edited the draft. F.O. and R.V. prepared the figures. All authors reviewed the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to F.O.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2023