



Reconstruction of Perceived Images from fMRI Patterns and Semantic Brain Exploration using Instance-Conditioned GANs

Furkan Ozcelik, Bhavin Choksi, Milad Mozafari, Leila Reddy, Rufin VanRullen

► To cite this version:

Furkan Ozcelik, Bhavin Choksi, Milad Mozafari, Leila Reddy, Rufin VanRullen. Reconstruction of Perceived Images from fMRI Patterns and Semantic Brain Exploration using Instance-Conditioned GANs. International Joint Conference on Neural Networks (IJCNN 2022), Jul 2022, Padua, Italy. pp.1-9, 10.48550/arXiv.2202.12692 . hal-03859829

HAL Id: hal-03859829

<https://hal.science/hal-03859829>

Submitted on 28 Nov 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Reconstruction of Perceived Images from fMRI Patterns and Semantic Brain Exploration using Instance-Conditioned GANs

Furkan Ozcelik
CerCo, CNRS UMR5549
Université de Toulouse
Toulouse, France
 furkan.ozcelik@univ-tlse3.fr

Bhavin Choksi
CerCo, CNRS UMR5549
Université de Toulouse
Toulouse, France
 bhavin.choksi@cnrs.fr

Milad Mozafari
CerCo, CNRS UMR5549
IRIT, CNRS UMR5505
Toulouse, France
 milad.mozafari@cnrs.fr

Leila Reddy
CerCo, CNRS UMR5549 and
ANITI, Université de Toulouse
Toulouse, France
 leila.reddy@cnrs.fr

Rufin VanRullen
CerCo, CNRS UMR5549 and
ANITI, Université de Toulouse
Toulouse, France
 rufin.vanrullen@cnrs.fr

Abstract—Reconstructing perceived natural images from fMRI signals is one of the most engaging topics of neural decoding research. Prior studies had success in reconstructing either the low-level image features or the semantic/high-level aspects, but rarely both. In this study, we utilized an Instance-Conditioned GAN (IC-GAN) model to reconstruct images from fMRI patterns with both accurate semantic attributes and preserved low-level details. The IC-GAN model takes as input a 119-dim noise vector and a 2048-dim instance feature vector extracted from a target image via a self-supervised learning model (SwAV ResNet-50); these instance features act as a conditioning for IC-GAN image generation, while the noise vector introduces variability between samples. We trained ridge regression models to predict instance features, noise vectors, and dense vectors (the output of the first dense layer of the IC-GAN generator) of stimuli from corresponding fMRI patterns. Then, we used the IC-GAN generator to reconstruct novel test images based on these fMRI-predicted variables. The generated images presented state-of-the-art results in terms of capturing the semantic attributes of the original test images while remaining relatively faithful to low-level image details. Finally, we use the learned regression model and the IC-GAN generator to systematically explore and visualize the semantic features that maximally drive each of several regions-of-interest in the human brain.

Index Terms—Natural Image Reconstruction, fMRI Decoding, IC-GAN, Brain-Computer Interface

I. INTRODUCTION

Understanding the brain and cognition has always been one of the fundamental research areas of science. One of the ways researchers approach this task is by establishing neural encoding and decoding methods. New ways to decode information from brain signals have emerged with recent developments in modeling and computation.

In vision research, many studies have used statistical methods and machine learning to decode specific information like position [1] or orientation [2], [3], to classify image categories [4], [5], to retrieve the closest images from a candidate set [6], or even to reconstruct images with low-complexity like basic shapes and structures [7].

With the emergence of deep learning, and in particular advanced deep generative models, reconstructing more complex images like handwritten digits [8], faces [9], and natural scenes [10] has become possible. These deep generative models include variational auto-encoders (VAEs), generative adversarial networks (GANs), and many variants and hybrids of both. Although many studies have used these models, they typically managed to reconstruct either low-level or high-level features of the images, but rarely both at the same time.

Here, we propose a method to reconstruct natural images from fMRI activation patterns with both accurate semantic attributes and relatively preserved low-level details, using an Instance-Conditioned GAN (IC-GAN) – a recent generative model [11] inspired by the success of self-supervised feature learning [12]. In our framework, we first extract latent representations for a set of training images (see Figure 1): high-level attributes of the images, called “instance features” in IC-GAN, are computed with a single forward pass through the SwAV ResNet-50 model; lower-level aspects of the image (e.g. reflecting the size, position or orientation of an object, details of the background, etc.) are obtained via a two-stage optimization of the IC-GAN “noise” and “dense” latent vectors (inspired by the method of Pividori et al. [13]). Next, we train three ridge regression models to predict these latent image representations from the corresponding fMRI patterns, recorded while human subjects viewed the same training images (Figure 2, Step 1). Finally, for each image in the test

set, we predict the instance feature, noise vector, and dense vector from fMRI data (using the previously learned regression models), and then reconstruct an estimate of the image using IC-GAN’s generator (Figure 2, Step 2 and 3). The code of this paper can be found in the official GitHub repository¹.

Our method establishes a new state-of-the-art performance for capturing the semantic attributes of the images, while preserving a reasonable amount of low-level details. We present both qualitative and quantitative results, and a comparison with previous methods to support our claims. We also take advantage of our brain-based image reconstruction system to explore and visualize the semantic image attributes encoded in various brain regions-of-interest (ROIs), and discuss how these findings align with neuroscientific evidence.

II. PREVIOUS WORKS

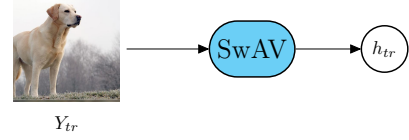
Many methods have been proposed in the literature for reconstructing natural images from fMRI patterns using deep learning models. Shen et al. [10] optimized input images with a deep generator network using the loss provided by fMRI decoded CNN features. Seeliger et al. [14] trained a linear regression between fMRI signals and DCGAN’s latent space using a feature loss from a CNN model and a pixel-space MSE loss. In addition to supervised training with {fMRI, stimulus} pairs, Belyi et al. [15] used a consistency loss for unsupervised training with test fMRI data (without corresponding stimuli) and additional image data. Later, Gaziv et al. [16] improved on this method by imposing a perceptual loss on reconstructed images, resulting in sharper reconstructions. Mozafari et al. [17] proposed the first semantically-oriented reconstruction model using BigBiGAN [18]. Finally, Ren et al. [19] developed a dual VAE-GAN model that uses a three-stage learning strategy incorporating adversarial learning and knowledge distillation. We qualitatively and quantitatively compare the results of these studies with our proposed method.

III. MATERIALS AND METHODS

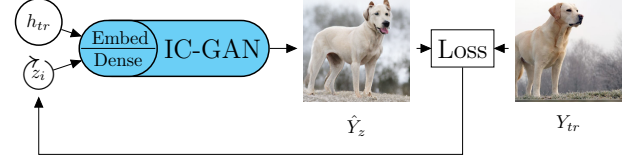
A. Instance-Conditioned GAN

We utilized an Instance-Conditioned GAN (IC-GAN) model, pretrained for natural image generation on the ImageNet dataset [20]. IC-GAN can be considered as a generic framework rather than a single model, because it can be applied to different GAN backbones, e.g. StyleGAN [21] or BigGAN [22]. In the usual conditional GAN setting [23], class labels are provided along with noise vectors sampled from a normal distribution to generate images. Images belonging to that specific class are labeled as “real”, and generated images from the generator are labeled as “fake”. Both the generator and discriminator are trained with these images and labels in an adversarial learning framework. In the instance-conditional setting, instead of giving a class label, instance features that capture the semantic attributes of a given image are extracted (via a pre-trained feature extractor) and provided to the generator as conditioning, alongside a sampled noise vector. For

Step 1 - Compute instance features h for training images



Step 2 - Optimize noise vector z for training images



Step 3 - Optimize dense vector d for training images

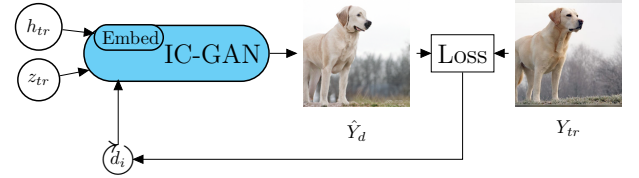


Fig. 1. Extraction of the latent variables (h_{tr} , z_{tr} and d_{tr}) for each training image (Y_{tr}). Step 1: Instance features of training images (h_{tr}) are extracted using SwAV ResNet-50. This 2048-dim instance feature vector (h_{tr}) captures the semantic attributes of the image. Step 2: In addition to the instance feature vector, the IC-GAN also requires a noise vector (z_i) as input, which encodes lower-level properties of the image (e.g., pose, orientation, background etc.). While providing h_{tr} obtained from Step 1 to the IC-GAN’s generator, we optimize the noise vector (z_i) to generate the closest image (\hat{Y}_z) to the groundtruth image (Y_{tr}). The resulting optimized noise vector is z_{tr} . Step 3: To further improve image reconstruction so as to better match the more detailed spatial structure of the training image, we apply another optimization stage, in which we optimize the dense layer vectors of IC-GAN itself. To achieve, this, we pass the first 17 dimensions of z_{tr} to the dense layer of the IC-GAN’s generator and obtain initial dense vectors (d_0). While keeping both h_{tr} and the remaining 102 dimensions of z_{tr} fixed, we optimize the dense vector d_i to generate the closest image (\hat{Y}_d) to the groundtruth image (Y_{tr}). d_{tr} is the resulting optimized dense vector.

training, IC-GAN selects k images in the neighborhood of the conditioning image (according to the feature extractor); these images are labeled as real, while generated images are considered as fake images to train both the generator and discriminator.

For the instance feature extraction, IC-GAN models use the SwAV (Swapping Assignments between Views) architecture [24] with a ResNet-50 [25] backbone. SwAV is a self-supervised learning model which means that it does not require handcrafted labels from humans. Similar to contrastive learning methods [12], SwAV minimizes the distance in feature space between representations of two transformed images (coming from the same original image).

It is possible to train the IC-GAN framework with different feature extractors, as long as they provide rich feature representations. However, using features from self-supervised learning models (e.g., SwAV) is better suited to the problem of neural decoding and natural image reconstruction. Indeed, many recent studies show that representations gathered from

¹https://github.com/ozcelikfu/IC-GAN_fmri_Reconstruction

self-supervised learning models present more similarity to brain representations than other learning methods [26], [27].

The specific IC-GAN model we used here relies on a BigGAN [22] architecture with 7 layers. It generates $256 \times 256 \times 3$ images from a 2048-dim (dimensional) instance feature vector extracted from SwAV ResNet-50 and a 119-dim noise vector sampled from a normal distribution. The 2048-dim instance features are given to an embedding layer and thus reduced to 512-dim embedded vectors. The 119-dim noise vector, which encodes lower-level properties of the image (e.g. pose, size, orientation of the object), is split into seven hierarchical levels, each with 17 dimensions. The first 17-dim level is directly given to the first dense layer of the IC-GAN generator. The remaining six hierarchical levels are concatenated with the embedded instance vector to be fed to the generator in each of the six BigGAN residual blocks.

Overall, the purpose of IC-GAN is to generate, from one conditioning image, new and diverse image instances that share semantic attributes (as captured by SwAV instance features), but differ in low-level properties (e.g. object position, size, orientation, background details). The diversity of low-level properties is determined by randomly sampled “noise” vectors (and by the “dense” vectors directly derived from them). However, for the purpose of fMRI-based image reconstruction, both high-level and low-level properties must be specified. Therefore, rather than randomly sampling noise vectors, we computed a specific noise vector (and the associated dense vector) for each training image in the dataset, as detailed below.

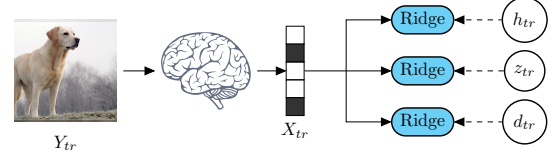
B. Extracting Latent Variables from Training Stimuli

We illustrate the computation of latent variables in Figure 1. We first extracted a 2048-dim instance feature vector for each training image in our dataset (see dataset details below) by presenting it to a SwAV ResNet-50 feature extractor. We then provided these instance features to the IC-GAN generator, and optimized the 119-dim noise vector for the same image using the covariance matrix adaptation evolution strategy (CMA-ES) [28]. We used this method because we empirically observed that global optimization strategies worked better than local optimization strategies (like gradient-based methods) for the noise vector. The loss function for this optimization was the distance between the generated image and the original training image in Layer-4 of SwAV ResNet-50; this representation level, hierarchically lower than the instance feature level, encodes more spatially structured information.

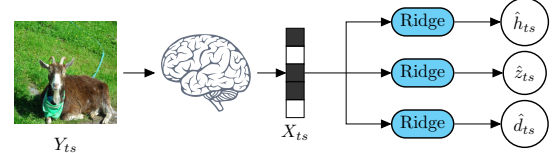
Finally, to further match the more detailed spatial structure of the original image, we applied one more optimization stage. Inspired by the two-stage inversion method of Pividori et al. [13], we provided the first 17 dimensions of the previously optimized noise vector to the first dense layer of the IC-GAN, resulting in a $1536 \times 4 \times 4$ -dim dense vector. While the instance features and the remaining 102 dimensions of the noise vector were kept fixed, we optimized these dense vectors with the RMSProp optimizer. For this second-stage optimization, the previous loss (SwAV ResNet-50 Layer-4

feature distance) was combined with a Learned Perceptual Image Patch Similarity (LPIPS) [29] loss gathered from a pretrained VGG16 model [30] and a pixel (MSE) loss from 64×64 resized images.

Step 1 - Train regression models with training fMRI data



Step 2 - Decode latent variables using test fMRI data



Step 3 - Reconstruct images from decoded latent variables

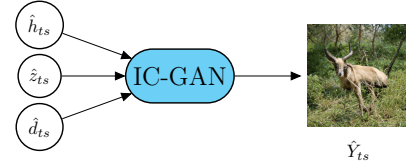


Fig. 2. Decoding latent variables from fMRI patterns and reconstructing images from decoded variables. Step 1: Having obtained the instance features (h_{tr}), noise vectors (z_{tr}) and dense vectors (d_{tr}) of training images (Y_{tr}) as described in Figure 1, we train three ridge regression models to map fMRI patterns of the training set (X_{tr}) to these latent variables. Step 2: Using these trained regression models, we decode latent variables of the test set ($\hat{h}_{ts}, \hat{z}_{ts}, \hat{d}_{ts}$) from test fMRI patterns (X_{ts}). Step 3: We pass the decoded latent variables to the IC-GAN Generator to obtain reconstructed images (\hat{Y}_{ts})

C. Generic Object Decoding Dataset

In this study, we used previously published fMRI recordings of five human subjects presented with images from the ImageNet dataset [31]. The dataset contains training and testing image perception sessions where subjects looked at 1200 training samples drawn from 150 categories (8 samples each) and 50 testing samples chosen from 50 categories (1 sample each), respectively. Training and testing categories were chosen independently and were non-overlapping. Each training image was presented only once, while testing images were repeated 35 times during the whole experiment. All fMRI runs followed a similar design: fixation (33s), 50 image presentations (9s per image flashing at 2Hz), fixation (6s). Moreover, subjects were also asked to perform a one-back task by pressing a button whenever the same image was presented two times in a row (five such events occurred per run).

The fMRI data were pre-processed for each subject by three-dimensional motion correction followed by coregistration to the high-resolution anatomical image. Then, the brain representation of each image was calculated by averaging the percent signal change values of each voxel over the 9-s presentation window. Additionally, the dataset provides

functional regions of interest (ROIs) that cover the entire visual cortex, including V1-V4, the fusiform face area (FFA), parahippocampal place area (PPA), and lateral occipital complex (LOC). The pre-processed data is available to download at brainliner.jp².

D. fMRI Decoding and Image Reconstruction

Details of fMRI decoding and image reconstruction are depicted in Figure 2. The procedure involves two separate stages for training and testing the brain decoding system of each subject.

First, we trained three separate ridge regression models to predict the latent variables (instance features; noise vectors; dense vectors) for each of the 1200 training images based on the corresponding fMRI patterns. Since both the fMRI data and the latent variables are high-dimensional, we applied L_2 regularization on the regression weights during training.

At test time, we averaged the 35 repetitions of fMRI signals corresponding to each test stimulus. Next, we used the previously trained regression models to predict the instance features, noise vectors, and dense vectors from these averaged fMRI signals. Finally, we used these predicted latent variables to generate image reconstructions using the IC-GAN generator.

IV. RESULTS AND ANALYSES

A. Image Reconstruction Results

Examples of image reconstructions produced by our method are displayed in Figure 3. First of all, it is important to examine IC-GAN reconstructions (second column) based on the optimized “ground-truth” latent vectors (derived as detailed in Figure 1): we can see that IC-GAN can successfully reconstruct the semantic attributes of the test images; however, it often misses some visual details, like parts of the vehicle (third row), liquid in the glass (fifth row), or the precise text in the gravestone (eighth row). These reconstructions help us understand how the IC-GAN generator would behave if we perfectly decoded latent variables from fMRI patterns, i.e. they serve as an upper bound on the expected reconstruction quality.

When we inspect actual fMRI reconstructions for the five subjects (third to seventh columns), our first observation is that reconstructions look like natural images. Furthermore, they are consistent across subjects. Again, these reconstructions capture some of the semantic attributes, while also missing specific aspects of the test images. For example, the system generates images of horned animals for the goat image (second row), but their species are not clearly identifiable. For the token image (fifth row), round objects are reconstructed, but not with the right texture. For the gravestone image (last row), similar square-shaped objects with text and symbols are generated, but most of them would not qualify as a gravestone. Overall, our method appears to reconstruct semantic attributes with slight but significant variations in details.

How does it compare to previously proposed methods? In Figure 4, we present image reconstructions using alternative

methods proposed in five other studies, together with our results for comparison³. From these reconstructions, we can see that many methods capture low-level details rather than high-level ones; as a result, many of the reconstructions do not look natural. A notable exception is the study of Mozafari et al., based on the BigBiGAN architecture [17], in which reconstructions often capture high-level properties and are more naturalistic. Even this method, however, does not correctly reconstruct semantic details for some of the images; furthermore, it misses many of the low-level details. Among the other studies, Ren et al. [19] succeed in reconstructing colors and textures better than other methods, while Gaziv et al. [16] give sharper object edges. Our method generates realistic-looking image reconstructions with appropriate semantic features, while preserving the low-level aspects to a certain degree.

TABLE I
QUANTITATIVE COMPARISON OF IMAGE RECONSTRUCTIONS. FOR EACH MEASURE, THE BEST VALUE IS IN BOLD. (FOR PIX-COMP/SSIM, HIGHER IS BETTER; FOR INCEPTION/CLIP DISTANCE, LOWER IS BETTER)

Method	Similarity Measure			
	Low-Level		High-Level	
	Pix-Comp \uparrow	SSIM \uparrow	Inception \downarrow	CLIP \downarrow
Shen et al. [10]	79.7%	0.582	0.829	0.358
Beliy et al. [15]	85.3%	0.597	0.865	0.424
Gaziv et al. [16]	91.5%	0.601	0.841	0.387
Ren et al. [19]	87.3%	0.588	0.847	0.383
Mozafari et al. [17]	54.3%	0.450	0.818	0.352
IC-GAN (Random)	64.1%	0.467	0.761	0.328
IC-GAN (Noise)	66.5%	0.489	0.744	0.320
IC-GAN (Dense)	67.2%	0.491	0.742	0.330

These qualitative observations are supported by the quantitative comparison of methods in Table I, according to both low-level measures of image quality (Pix-Comp, SSIM) and higher-level “semantic” measures (Inception or CLIP distance). Pix-Comp is a 2-way comparison of pixel-wise correlation measures computed over the whole test set. We used the results reported by authors in their respective papers, except for Gaziv et al. [16], who did not report Pix-Comp: we re-computed it over the reconstructed images provided in their supplementary material. All other metrics (SSIM [32], Inception-V3 [33] distance, and CLIP ViT-B/32 [34] distance), were computed over the seven common image reconstructions presented in Figure 4. Our own results are presented for three different versions of IC-GAN decoding, using different combinations of the three brain regression models in Figure 2, to evaluate the effects of each regressor on performance. First, the IC-GAN (Random) version uses brain-decoded instance features together with randomly sampled noise vectors from a normal distribution. Second, the IC-GAN (Noise) version combines brain-decoded instance features with brain-decoded noise vectors, without using the brain-decoded dense vectors (instead, the output of the first dense layer is used directly). Finally, IC-GAN (Dense) is the complete framework described

²http://brainliner.jp/data/brainliner/Generic_Object_Decoding

³We selected these seven images because it was the only common set of reconstruction exemplars presented across all of the considered studies.

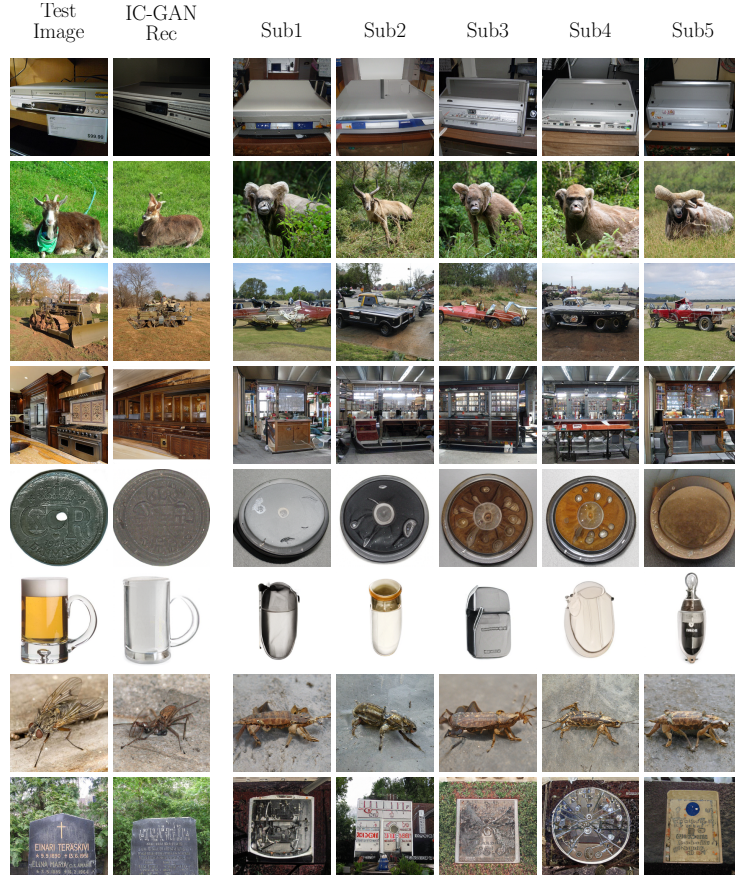


Fig. 3. fMRI Reconstructions by the IC-GAN model for all subjects. The first column is the groundtruth test image, whereas the second column is the reconstructed image by IC-GAN using true extracted latent variables. The following five columns demonstrate the equivalent reconstructions using fMRI-decoded latent variables for each subject. fMRI reconstructions are generally consistent with the groundtruth images in terms of semantic attributes, while they preserve the low-level details to a certain degree.

in Figure 2, which uses all the brain-decoded latent variables (thus overriding the dense vector with its brain-decoded version). The table indicates that most other methods yield better results than IC-GAN on the low-level measures (Pix-Comp, SSIM), except for Mozafari et al [17]; like ours, that study was aimed at matching higher-level “semantic” aspects of the input images. Importantly, IC-GAN outperforms the Mozafari et al method for both low-level measures. For the high-level measures (Inception and CLIP Distances), IC-GAN demonstrates state-of-the-art performance, surpassing all methods—including Mozafari et al.—by a significant margin.

The comparison of the 3 versions of our IC-GAN method reveals that the inclusion of both the brain-decoded noise vector (IC-GAN Noise) and the brain-decoded dense vector (IC-GAN Dense) helps improve the model’s ability to capture low-level details. Still, the full model remains inferior to many previous methods in this respect. Regarding high-level semantic attributes, while the full method IC-GAN (Dense) is superior to IC-GAN (Noise) for the Inception distance, the opposite is true for the CLIP distance. This could be because Inception features include more spatially structured information than CLIP features; indeed, the function of dense

vectors in our method is precisely to capture the image spatial structure that is less explicitly encoded in the noise vectors.

B. Semantic Analysis of Visual Encoding in Brain ROIs

Our brain decoding model, relying on the latent space(s) of the IC-GAN network, can reconstruct the high-level content of perceived images better than all prior methods, while retaining more low-level details than at least some of these methods. From a neuroscience viewpoint, can this brain decoding model also help us understand the neural coding of visual information in the brain? Here, we use our model to explore and directly visualize the types of information that are preferentially represented in various brain regions-of-interest (ROIs).

The fMRI dataset counts seven distinct ROIs across visual cortex for each subject—in hierarchical order: V1, V2, V3, V4, LOC (Lateral Occipital Complex), FFA (Fusiform Face Area) and PPA (Parahippocampal Place Area). First, we ask whether each region carries more information about high-level latent features—as captured by the model’s instance features—or about low-level properties—as captured by the model’s dense vector (note that similar results, not shown here, were obtained for the noise vector instead of the dense vector). To answer this

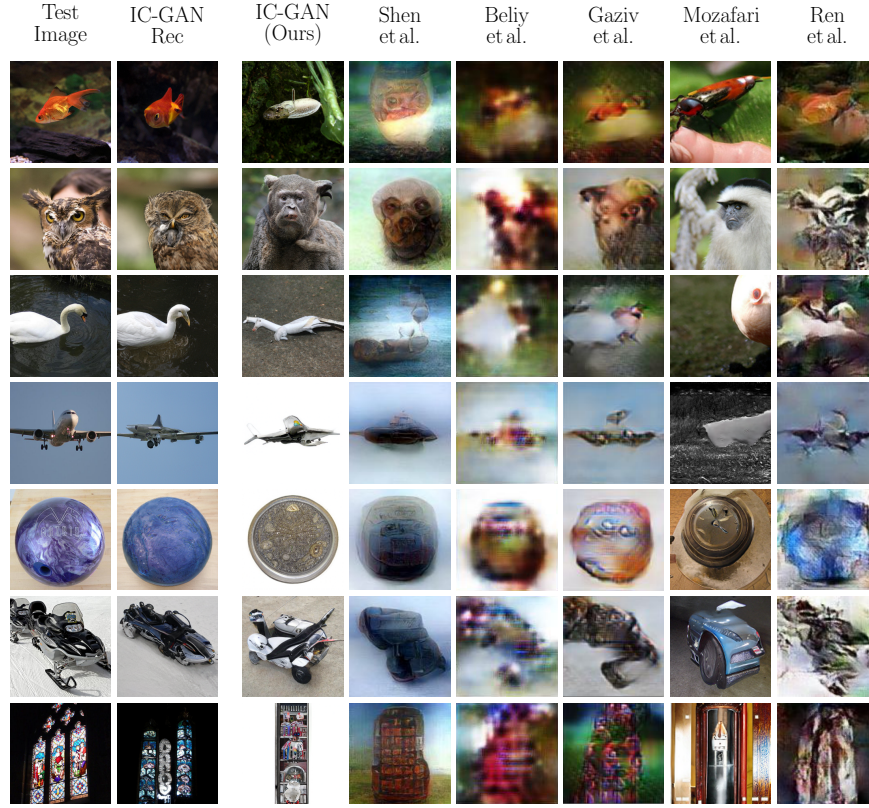


Fig. 4. Comparison of fMRI reconstructions for several methods. The first column is the groundtruth image, the second column is the reconstructed image with IC-GAN using true extracted latent variables. Columns three to eight present fMRI reconstructions from IC-GAN (Ours), Shen et al. [10], Beliy et al. [15], Gaziv et al. [16], Mozafari et al. [17], and Ren et al. [19], respectively. fMRI reconstructions by the IC-GAN method demonstrate more naturalistic-looking images with accurate semantic attributes, while preserving some low-level details (e.g. object position, size or orientation).

question, for each brain voxel we compared the L_1 norm of the model’s ridge regression weights for the instance features vs. dense vectors (Figure 5). As expected, lower brain regions (V1-V3) were more informative about the dense vector, while higher brain regions (V4, LOC, FFA, PPA) carried more information about instance features.

Next, we use our brain decoding model to visualize the “optimal” stimulus for each brain region. Instead of using fMRI patterns recorded from subjects, we synthesized seven patterns, one for each ROI, with a value of 1 for all voxels inside the ROI and 0 outside. We provided these synthetic patterns to the three trained ridge regression models to obtain predicted latent variables (as described in Figure 2). To mitigate the scaling problem, we normalized instance features to have unit norms. We then passed the predicted latent variables through the IC-GAN generator to generate images.

Previously, Gu et al. [35] synthesized optimal images for different ROIs using a BigGAN generator and a feature extractor. They iteratively optimized the latent variables of the generator in such a way that predicted fMRI patterns (obtained via the feature extractor) maximized activation in a specific ROI. In contrast, our method involves a single pass through our image reconstruction pipeline, and does not require iterative optimization of the latent variables. Figure 6 presents the

generated images from each subject (second to sixth columns), together with reconstructions using averaged latent variables across all five subjects (first column).

In lower visual cortex (V1-V2), basic textures (foliage, trees, stones) are produced rather than (or in addition to) identifiable objects. The textures emphasize the periphery of the visual field, in line with the fact that V1-V2 have small receptive fields that can be positioned at high visual eccentricity. For V3 and V4, the generated textures present more regularity than V1 and V2, and we begin to see visuals close to objects with multiple parts, including text-like symbols, notably in V4. LOC is known for its selectivity to object shapes; when maximizing this region’s response, IC-GAN generates complete objects at the center of the image, rather than extended textures. At this stage, the visual periphery appears empty or blurry, in contrast with the crisp peripheral textures produced for V1-V2. In FFA, a high-level region known for its selectivity to face images, IC-GAN generates human and animal faces. The presence of animal faces is not unexpected, since the ImageNet dataset (on which IC-GAN was trained) contains many more animal images than human images. Some previous experimental and computational work [35], [36] also suggests that fusiform regions may show a preferential response to animals, and particularly dogs. Nonetheless, the model still

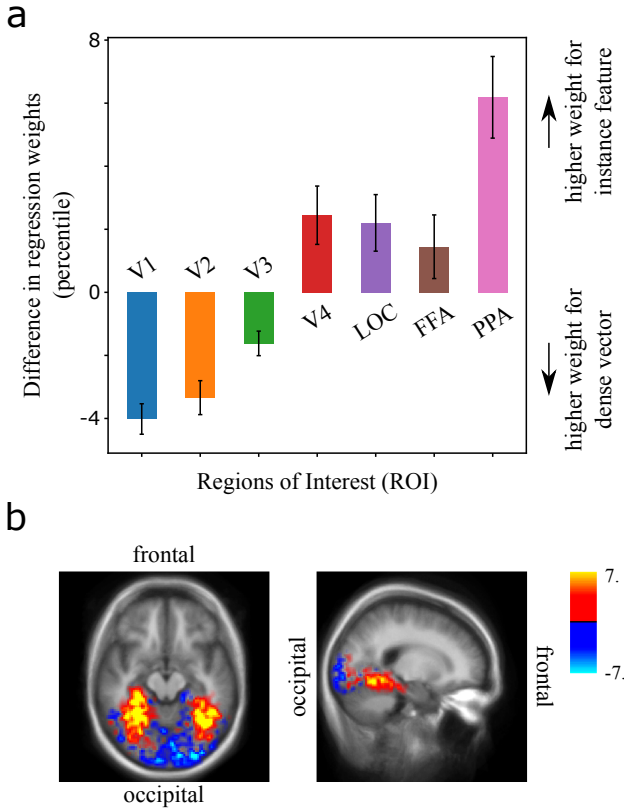


Fig. 5. Mapping of instance features vs. dense vectors over brain regions. (a) Difference between the percentiles of the regression weights (L_1 norm) for the instance features vs. the dense vector, averaged over voxels in each ROI. Positive values indicate relatively higher weight for instance features compared to the dense vector, and vice versa. Error bars represent standard error of the mean across 5 subjects. (b) Voxel-by-voxel maps (left: axial; right: sagittal) of the difference between the percentiles of the regression weights (L_1 norm) for the instance features (red) vs. the dense vector (blue), averaged over the 5 subjects. Dense vector weights are higher in early visual cortex (occipital regions), while instance feature weights are larger in higher visual cortex (temporal regions).

generates human face images for two of the subjects. The last ROI from the higher visual cortex is PPA, known for its selectivity to environmental scenes like indoor and outdoor places. IC-GAN also generates indoor and outdoor places when the voxels of this region are activated. Some of the images have an object in the center of the scene; this might be caused by the training of the IC-GAN model on ImageNet—an object-centered dataset. It is worth noting that PPA-optimized images produce more details in the visual periphery than FFA-optimized images; this is compatible with the known difference in preferential eccentricity between the two regions [37]. Overall, the outcomes of this analysis are consistent with findings from the neuroscience literature, indicating that our IC-GAN-based model learned to appropriately decode visual feature selectivity in the brain. Most importantly, the method allows us to *directly visualize* this selectivity, rather than inferring it from extended experiments.



Fig. 6. Generated images from synthetic fMRI patterns constructed by activating all voxels in a specific brain region-of-interest (ROI), and none outside of the ROI. The rows represent various brain regions: V1, V2, V3, V4, LOC, FFA, and PPA. The first column is generated after averaging the brain-predicted latent variables for all five subjects. The following columns are for individual subjects.

V. DISCUSSION

In this paper, we presented a framework for natural image reconstruction from fMRI patterns using the IC-GAN model, pretrained on ImageNet. First, we extracted instance features, noise vectors, and dense vectors from training images, and trained ridge regression models from fMRI patterns to these latent variables. With these regression models, we decoded latent variables from the test fMRI patterns, and finally reconstructed images with the IC-GAN generator.

Many previous studies implemented fMRI reconstruction frameworks with deep generative models. However, these models were able to reconstruct either low-level or high-level features of the images. Our method demonstrated state-of-the-art performance on reconstructing semantic (high-level) attributes of the images, both qualitatively and quantitatively, while generating naturalistic-looking images. Meanwhile, compared to other semantically oriented models (e.g. Mozafari et al. [17], an approach based on BigBiGAN), it was able to maintain more low-level details. Furthermore, we could use our fMRI-based image reconstruction model to visualize images decoded from synthetic fMRI patterns, designed to maximize activations in specific brain ROIs. The results of this analysis were aligned with the existing neuroscience literature, opening a range of possibilities for future brain exploration and

visualization techniques.

We acknowledge that there is still room for improving our model, especially in terms of better reproducing low-level details. This may be achieved in future work by improving our optimization of the noise and dense vectors, or by pairing IC-GAN with other generative networks more focused on low-level image properties.

REFERENCES

- [1] B. Thirion, E. Duchesnay, E. Hubbard, J. Dubois, J.-B. Poline, D. Lebihan, and S. Dehaene, "Inverse retinotopy: inferring the visual content of images from brain activation patterns," *Neuroimage*, vol. 33, no. 4, pp. 1104–1116, 2006.
- [2] Y. Kamitani and F. Tong, "Decoding the visual and subjective contents of the human brain," *Nature neuroscience*, vol. 8, no. 5, pp. 679–685, 2005.
- [3] J.-D. Haynes and G. Rees, "Predicting the orientation of invisible stimuli from activity in human primary visual cortex," *Nature neuroscience*, vol. 8, no. 5, pp. 686–691, 2005.
- [4] J. V. Haxby, M. I. Gobbini, M. L. Furey, A. Ishai, J. L. Schouten, and P. Pietrini, "Distributed and overlapping representations of faces and objects in ventral temporal cortex," *Science*, vol. 293, no. 5539, pp. 2425–2430, 2001.
- [5] D. D. Cox and R. L. Savoy, "Functional magnetic resonance imaging (fmri) "brain reading": detecting and classifying distributed patterns of fmri activity in human visual cortex," *Neuroimage*, vol. 19, no. 2, pp. 261–270, 2003.
- [6] K. N. Kay, T. Naselaris, R. J. Prenger, and J. L. Gallant, "Identifying natural images from human brain activity," *Nature*, vol. 452, no. 7185, pp. 352–355, 2008.
- [7] Y. Miyawaki, H. Uchida, O. Yamashita, M.-a. Sato, Y. Morito, H. C. Tanabe, N. Sadato, and Y. Kamitani, "Visual image reconstruction from human brain activity using a combination of multiscale local image decoders," *Neuron*, vol. 60, no. 5, pp. 915–929, 2008.
- [8] S. Schoenmakers, M. Barth, T. Heskes, and M. Van Gerven, "Linear reconstruction of perceived images from human brain activity," *NeuroImage*, vol. 83, pp. 951–961, 2013.
- [9] R. VanRullen and L. Reddy, "Reconstructing faces from fmri patterns using deep generative neural networks," *Communications biology*, vol. 2, no. 1, pp. 1–10, 2019.
- [10] G. Shen, T. Horikawa, K. Majima, and Y. Kamitani, "Deep image reconstruction from human brain activity," *PLoS computational biology*, vol. 15, no. 1, p. e1006633, 2019.
- [11] A. Casanova, M. Careil, J. Verbeek, M. Drozdal, and A. Romero, "Instance-conditioned GAN," in *Advances in Neural Information Processing Systems*, A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan, Eds., 2021. [Online]. Available: https://openreview.net/forum?id=aUuTEEcY_
- [12] A. Jaiswal, A. R. Babu, M. Z. Zadeh, D. Banerjee, and F. Makedon, "A survey on contrastive self-supervised learning," *Technologies*, vol. 9, no. 1, p. 2, 2021.
- [13] M. Pividori, G. L. Grinblat, and L. C. Uzal, "Exploiting gan internal capacity for high-quality reconstruction of natural images," *arXiv preprint arXiv:1911.05630*, 2019.
- [14] K. Seeliger, U. Güçlü, L. Ambrogioni, Y. Güçlütürk, and M. van Gerven, "Generative adversarial networks for reconstructing natural images from brain activity," *NeuroImage*, vol. 181, pp. 775–785, 2018. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S105381191830658X>
- [15] R. Beliy, G. Gaziv, A. Hoogi, F. Strappini, T. Golan, and M. Irani, "From voxels to pixels and back: Self-supervision in natural-image reconstruction from fmri," *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [16] G. Gaziv, R. Beliy, N. Granot, A. Hoogi, F. Strappini, T. Golan, and M. Irani, "Self-supervised natural image reconstruction and rich semantic classification from brain activity," *bioRxiv*, 2020.
- [17] M. Mozafari, L. Reddy, and R. VanRullen, "Reconstructing natural scenes from fmri patterns using biggan," in *2020 International joint conference on neural networks (IJCNN)*. IEEE, 2020, pp. 1–8.
- [18] J. Donahue and K. Simonyan, "Large scale adversarial representation learning," *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [19] Z. Ren, J. Li, X. Xue, X. Li, F. Yang, Z. Jiao, and X. Gao, "Reconstructing seen image from brain activity by visually-guided cognitive representation and adversarial learning," *NeuroImage*, vol. 228, p. 117602, 2021.
- [20] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *2009 IEEE conference on computer vision and pattern recognition*. Ieee, 2009, pp. 248–255.
- [21] T. Karras, S. Laine, M. Aittala, J. Hellsten, J. Lehtinen, and T. Aila, "Analyzing and improving the image quality of stylegan," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 8110–8119.
- [22] A. Brock, J. Donahue, and K. Simonyan, "Large scale GAN training for high fidelity natural image synthesis," in *International Conference on Learning Representations*, 2019. [Online]. Available: <https://openreview.net/forum?id=B1xsqj09Fm>
- [23] M. Mirza and S. Osindero, "Conditional generative adversarial nets," *arXiv preprint arXiv:1411.1784*, 2014.
- [24] M. Caron, I. Misra, J. Mairal, P. Goyal, P. Bojanowski, and A. Joulin, "Unsupervised learning of visual features by contrasting cluster assignments," *Advances in Neural Information Processing Systems*, vol. 33, pp. 9912–9924, 2020.
- [25] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [26] T. Konkle and G. A. Alvarez, "A self-supervised domain-general learning framework for human ventral stream representation," *Nature Communications*, vol. 13, no. 1, pp. 1–12, 2022.
- [27] C. Zhuang, S. Yan, A. Nayebi, M. Schrimpf, M. C. Frank, J. J. DiCarlo, and D. L. Yamins, "Unsupervised neural network models of the ventral visual stream," *Proceedings of the National Academy of Sciences*, vol. 118, no. 3, 2021.
- [28] N. Hansen and A. Ostermeier, "Completely derandomized self-adaptation in evolution strategies," *Evolutionary Computation*, vol. 9, no. 2, pp. 159–195, 2001.
- [29] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, "The unreasonable effectiveness of deep features as a perceptual metric," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 586–595.
- [30] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.
- [31] T. Horikawa and Y. Kamitani, "Generic decoding of seen and imagined objects using hierarchical visual features," *Nature communications*, vol. 8, no. 1, pp. 1–15, 2017.
- [32] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE transactions on image processing*, vol. 13, no. 4, pp. 600–612, 2004.
- [33] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, "Rethinking the inception architecture for computer vision," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 2818–2826.
- [34] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark *et al.*, "Learning transferable visual models from natural language supervision," in *International Conference on Machine Learning*. PMLR, 2021, pp. 8748–8763.
- [35] Z. Gu, K. W. Jamison, M. Khosla, E. J. Allen, Y. Wu, T. Naselaris, K. Kay, M. R. Sabuncu, and A. Kuceyeski, "Neurogen: activation optimized image synthesis for discovery neuroscience," *NeuroImage*, vol. 247, p. 118812, 2022.
- [36] L. X. Blonder, C. D. Smith, C. E. Davis, M. L. Kesler, T. F. Garrity, M. J. Avison, A. H. Andersen *et al.*, "Regional brain response to faces of humans and dogs," *Cognitive Brain Research*, vol. 20, no. 3, pp. 384–394, 2004.
- [37] I. Levy, U. Hasson, G. Avidan, T. Hendler, and R. Malach, "Center-periphery organization of human object areas," *Nature neuroscience*, vol. 4, no. 5, pp. 533–539, 2001.