
Second Sight: Using brain-optimized encoding models to align image distributions with human brain activity

Reese Kneeland

Department of Computer Science
University of Minnesota
Minneapolis MN, 55455
rek@umn.edu

Jordyn Ojeda

Department of Computer Science
University of Minnesota
Minneapolis MN, 55455
ojeda040@umn.edu

Ghislain St-Yves

Department of Neuroscience
University of Minnesota
Minneapolis MN, 55455
gstyves@umn.edu

Thomas Naselaris

Department of Neuroscience
University of Minnesota
Minneapolis MN, 55455
nase0005@umn.edu

Abstract

Two recent developments have accelerated progress in image reconstruction from human brain activity: large datasets that offer samples of brain activity in response to many thousands of natural scenes, and the open-sourcing of powerful stochastic image-generators that accept both low- and high-level guidance. Most work in this space has focused on obtaining point estimates of the target image, with the ultimate goal of approximating literal pixel-wise reconstructions of target images from the brain activity patterns they evoke. This emphasis belies the fact that there is always a family of images that are equally compatible with any evoked brain activity pattern, and the fact that many image-generators are inherently stochastic and do not by themselves offer a method for selecting the single best reconstruction from among the samples they generate. We introduce a novel reconstruction procedure (Second Sight) that iteratively refines an image distribution to explicitly maximize the alignment between the predictions of a voxel-wise encoding model and the brain activity patterns evoked by any target image. We use an ensemble of brain-optimized deep neural networks trained on the Natural Scenes Dataset (NSD) as our encoding model, and a latent diffusion model as our image generator. At each iteration, we generate a small library of images and select those that best approximate the measured brain activity when passed through our encoding model. We extract semantic and structural guidance from the selected images, used for generating the next library. We show that this process converges on a distribution of high-quality reconstructions by refining both semantic content and low-level image details across iterations. Images sampled from these converged image distributions are competitive with state-of-the-art reconstruction algorithms. Interestingly, the time-to-convergence varies systematically across visual cortex, with earlier visual areas generally taking longer and converging on narrower image distributions, relative to higher-level brain areas. Second Sight thus offers a succinct and novel method for exploring the diversity of representations across visual brain areas.

1 Introduction

Reconstructing visual stimuli from measured brain activity is a long-standing problem in computational neuroscience. Improvements bring us closer to understanding visual representation in the brain by providing interpretable tools for examining brain states [1, 2].

Two recent developments have accelerated progress in image reconstruction from human brain activity: large datasets that offer samples of brain activity in response to many thousands of natural scenes [3], and the open-sourcing of powerful image-generators that can invert abstract feature representations into distributions over natural scenes [4]. In many recent approaches, these data are used to map brain activity directly to feature representations that are fed into the generator.

Although this approach has been proven quite successful [5, 6, 7, 8, 9, 10], it has clear limitations. First, latent representations that are used to drive image generators may overlap with only a subset of the rich and diverse representations encoded in brain activity across visual cortex [1].

Second, even with the very large datasets that have recently been distributed, bottom-up approaches that seek to decode only low-level representations do not yield semantically meaningful reconstructions [11]. Thus, most successful pipelines extract some form of semantic representation in conjunction with low-level structural representations to guide reconstruction. This practice requires a policy for weighting the relative influence of high-level semantic representations and low-level structural representations on the reconstruction. Semantic representations are structurally ambiguous, so weighting them strongly will result in broad output distributions that capture semantics at the expense of uncertainty about structure; strongly weighting structural representations offers lower structural uncertainty at the expense of semantic accuracy. Because of this tension, some amount of stochasticity in the reconstructions may be inevitable [10]. The problem of how to reduce structural uncertainty without distorting semantic content is currently unsolved.

Finally, given that some amount of stochasticity in the reconstruction pipeline must be accepted, a method for scoring individual candidate reconstruction sampled from a distribution over images is also desirable.

We develop an approach to reconstructing natural images that addresses all three problems. Rather than training decoding models that minimize loss in the space of high- or low-level representations, we train “brain-optimized” encoding models that minimize loss in the space of brain activity. At test, we use the encoding model to select image distributions that are closely aligned to the brain activity pattern evoked by the target image. Specifically, we select images from small libraries generated by a diffusion model [4]. By using encoding models optimized to predict activity in most of the visual cortex, we leverage the rich diversity of representations in the human visual system. By incrementally lowering the breadth of reconstructed image distributions, and guiding it with selections from sequential search iterations, we are able to minimize uncertainty about structural detail while preserving and even improving the semantic aspects of the reconstruction. In this way, we are able to leverage the stochasticity inherent in diffusion models, progressively reducing the width of the generated image distribution until it aligns with the ambiguity measured in cortical representations of natural stimuli. We refer to our pipeline as “Second Sight”.

We show that selecting the most brain-aligned image distribution from each search yields semantically meaningful reconstructions that often correctly reflect low-level details of target images (e.g., pose). Our reconstructions are competitive with SOTA systems in terms of their proximity to ground truth images in a variety of feature spaces and, as expected, establish a new SOTA in terms of their alignment to activity patterns in the brain.

Furthermore, by analyzing search dynamics we show that, for most brain areas, samples from output distributions exceed ground truth alignment after several iterations, demonstrating that any single pattern of brain activity admits a distribution of highly aligned images. We show how this degeneracy increases across visual cortex, as evidenced by the rapid rate at which reconstructions align to activity in high-level areas, relative to low-level areas. We argue that “rate of convergence” is a succinct and attractive metric of representational invariance across brain areas and brain states.

2.2 Models

The pipeline’s model preparation phase comprises several integral components, some of which are pre-trained implementations and some we prepare for this work. For the Library Assembler, we load a pre-trained OpenClip ViT-H/14 [13, 14, 15] image embedding model for the $I \rightarrow c$ step, and a pre-trained VDVAE decoder [16] for the $I \rightarrow z$ step. The ensemble of encoding models includes a brain-optimized neural network (GNet; [2]), and an independent 3-layer MLP that maps the ViT-H/14 CLIP image embedding of an input image to a predicted brain activity pattern that includes all of nsdgeneral. Both encoding models were trained on the NSD training/validation sets. Unless noted otherwise, we averaged the outputs of this ensemble of encoding models to obtain a single predicted activity pattern β' . We refer to this average as the output of a “hybrid” encoding model. While more nuanced methods for combining encoding model outputs exist [17], this methodology produces an encoding model with sufficient sensitivity to both structural and semantic features, more details of which can be found in Appendix B.5. The encoded β' s corresponding to the ground truth images serve as targets in conjunction with the original fMRI scans to train a denoising autoencoder. This process enables more robust comparisons between the original brain data and the encoded predictions by removing some of the noise associated with fMRI data. This work uses the Stable Diffusion [4] Reimagine model (SD) as our image generator. The SD model is adapted in our implementation to receive a latent z vector and a strength parameter in addition to a CLIP embedding. More details of the individual model components and our pipeline can be found in Appendix B.

2.3 Stage 1: Library Assembler

Numbered items reference Figure 1, (top).

1. Three samples of measured brain activity β for a target image in the test set are passed through an autoencoder and concatenated to yield denoised activity patterns $\hat{\beta}$.
2. Each of the $n = 63,000$ images (I) in a large image library is passed through an encoding model (either a hybrid model or the GNet model), yielding a predicted activity pattern β' for each library image. The library is a subset of the COCO dataset [18]. None of the library images were included in our training, test, or validation sets.
3. Each image in the large library is scored by calculating Pearson correlation between $\hat{\beta}$ and β' . We refer to this form of scoring as “brain correlation”. Where noted, brain correlation scores may also be computed with respect to activity in a single ROI (e.g., V1).
4. Library images are sorted in decreasing order of their brain correlation scores, and the top k images are selected.
5. The top k images are embedded into either a high-level or a low-level representation. The high-level feature representation is a ViT-H/14 CLIP vector c , with $k = 100$. The low-level representation is the first 31 components of the VDVAE latent vector [16], with k parameter = 25. For high-level representation, the top k selections are made using brain correlation between the hybrid encoding model and all voxels in nsdgeneral. For low-level representations the top k selections are made using brain correlation between the GNet model and all voxels in early visual cortex. The VDVAE representations are converted to pixels using the pre-trained VDVAE decoder and then passed through SD’s AutoKL embedder to generate the low-level structural guidance z .
6. The top k high-level (low-level) representations are averaged to yield seed representations c_0 (z_0).
7. The top-scoring image from the library is also saved as an alternative encoding based reconstruction approach. We refer to this method as the “Best Library Image”.

2.4 Stage 2: Stochastic Search

Lettered items refer to Figure 1 (bottom).

- (a) As above, three samples of brain activity patterns evoked by the test target image are denoised and concatenated ($\hat{\beta}$).
- (b) Representations c_0 and z_0 (produced by the Library Assembler) seed the first search iteration (blue).

- (c) High-level (c_0) and low-level (z_0) representations are input to the image generator. Iterations $i > 0$ yield multiple high-level and low-level representations, denoted $c_{i,b}$ and $z_{i,b}$, $0 < i < 6$, $0 < b < 3$.
- (d) Image batches ($I_{i,b}$) are sampled from SD conditioned on $c_{i,b}$ and $z_{i,b}$. SD exposes a strength parameter λ that modulates the quantity of noise introduced to the latent z vector, with a value of 1 representing the addition of 100% noise and 0 signifying the addition of 0% noise. Furthermore, the strength parameter dictates the number of CLIP-guided denoising steps to be executed, as z vectors with reduced noise necessitate less denoising. This parameter effectively controls the breadth of the distribution of generated images originating from a given z vector.
- (e) Images in each batch are passed through the encoding model to yield predicted brain activity β' .
- (f) Images are scored and sorted as in Stage 1.
- (g) The four images with the best correlation scores are embedded into their latent c and z vectors. By using the top four images to seed parts of the following iteration, we ensure our search does not get stuck in a local maximum. The z vectors from each of the four batch seeds b act as the corresponding $z_{i,b}$ vector for the n_b images in the forthcoming batch $I_{i,b}$. We observed that using the embedded c vector directly in the next iteration led to semantic instability, which we correct by spherically interpolating between the current set of c vectors and the previous iteration's best CLIP vector c_{i-1} according to a momentum parameter γ . This produces our four $c_{i,b}$ vectors, each mixing a little of the semantics from the new batch seeds b into the next iteration of batches $I_{i,b}$. n_b , γ , and the strength parameter λ are progressively reduced across iterations in accordance with monotonically decreasing exponential or cubic schedules (described in Appendix B.7.)
- (h) The image distribution $I_{i,b}$ that yields the highest average brain correlation throughout all iterations is selected as the output distribution, from which the image with the highest brain correlation I is selected as the output image displayed in Figures 2 and 3.




Figure 2: Example reconstructions. The first row is the ground truth image, indicated by the red background, while the remaining rows are the reconstructions for each individual subject (1, 2, 5, and 7). Some additional samples can be found in Appendix E.

3 Results

3.1 Assessment of reconstruction quality

Figure 2 shows example reconstructions for four subjects. As shown, in the best case reconstructions recapitulate both the semantic content of images and their low-level structural details for all subjects.




Figure 3: Comparative assessment of reconstructions. The first row is the ground truth image, indicated by the red background, the second row is Second Sight reconstructions, the third row shows the "Best Library Image" approach, while the remaining rows represent subject 1 results from previous works in Ozcelik et al. [6], Tagaki et al. [8], Lu et al. [9], and Gu et al. [7]

Figure 3 presents a comparison of reconstructions from a number of recent works. At a glance, our reconstructions compare favorably.

To make explicit comparisons of reconstruction quality across different approaches, we calculated the proximity of our and others' reconstructions to ground truth in a variety of feature spaces (Table 1) and the space of brain activity (Table C.2). To calculate these metrics, we contacted the authors of all of the previous works we compare against to obtain the full set of their generated results. Wherever possible, we calculated the metrics as an average of 5 generated images sampled from the final image distribution of the method. The numbers for Second Sight are calculated on 5 images sampled from the distribution with the highest average brain correlation with the nsdgeneral mask of the visual cortex. The numbers for our "Best Library Image" were calculated on a single image per sample, as this method is deterministic and does not produce a distribution. Although we performed no explicit optimization in feature space, our reconstructions are competitive with current SOTA at approximating ground truth features. As expected, our reconstructions establish SOTA brain correlation scores for all visual ROIs we examined. For more details on how metrics were calculated as well as results in a multi-subject analysis, see Appendix C.

3.2 Analysis of search dynamics

Figure 4 shows examples of the best-aligned reconstructions across iterations of several searches. As seen in these examples, using an encoding model to iteratively align reconstructions to brain activity results in gradual refinement of both the semantic content (i.e., the category of objects depicted) and structural details (i.e., their pose and composition) of the reconstructions. For more details on this ablation study, see Appendix D.

Methods	Low Level Structural Metrics				
	PixCorr \uparrow	SSIM \uparrow	AlexNet(2) \uparrow	AlexNet(5) \uparrow	AlexNet(7) \uparrow
Second Sight (ours)	.185	.288	91.7%	95.0%	94.2%
Best Library Image (ours)	.138	.255	86.9%	93.1%	93.3%
Ozcelik et al. [6]	.304	.294	96.7%	97.6%	96.3%
Tagaki et al.[8]	.269	.217	75.1%	77.0%	76.0%
Lu et al.[9]	.246	.221	74.7%	78.9%	77.1%
Gu et al.[7]	.146	.279	79.4%	90.1%	93.7%

Methods	High Level Semantics Metrics				
	CLIP(2-way) \uparrow	CLIP(cos) \uparrow	Inception V3 \uparrow	EffNet-B \downarrow	SwAV \downarrow
Second Sight (ours)	88.3%	.664	84.7%	.776	.418
Best Library Image (ours)	85.9%	.657	82.5%	.798	.445
Ozcelik et al.[6]	92.6%	.657	88.5%	.764	.411
Tagaki et al.[8]	64.4%	.566	64.0%	.954	.668
Lu et al.[9]	70.8%	.578	69.6%	.913	.552
Gu et al.[7]	78.5%	.597	79.0%	.864	.451

Table 1: Quantitative comparison against past reconstruction methods on subject 1. For each measure, the best value is in bold. For EffNet-B and SwAV distances, lower is better. Higher is better for all other metrics. This is indicated by the arrow pointing up or down.

Methods	Brain Correlation Metrics					
	V1 \uparrow	V2 \uparrow	V3 \uparrow	V4 \uparrow	Higher Vis \uparrow	nsdgeneral \uparrow
Second Sight (ours)	.739	.774	.817	.854	.932	.905
Best Library Image (ours)	.697	.750	.800	.847	.927	.895
Ozcelik et al.[6]	.672	.676	.705	.757	.862	.828
Tagaki et al.[8]	.324	.271	.282	.307	.361	.356
Lu et al.[9]	.286	.276	.304	.364	.511	.467
Gu et al.[7]	.397	.425	.460	.550	.740	.676

Table 2: Evaluation of brain correlation scores against past reconstruction methods on subject 1. This metric evaluates how well the reconstructed images represent activity in various brain regions according to our hybrid encoding model.

Given the degeneracy of representations encoded in brain activity, any search through image space can in principle identify images that, when passed through our encoding model, yield predictions that are more highly correlated with brain activity than the target image that evoked it. We analyzed the evolution of the alignment between brain activity and our reconstructions to determine at what point the process exceeded the alignment of the ground truth target images (Figure 5, plots A and B). In this analysis, surpassing the brain correlation score of the ground truth image serves as functional confirmation that we have converged on a distribution of images that aligns closely with the brain’s representation of the original stimuli.

In high-level areas, alignment to brain activity achieves parity with ground truth by iteration 0 (on average), while in V1 average alignment is still below ground truth by iteration 5 (the highest iteration tested here). Interestingly, the "rate of convergence" to parity with ground truth orders brain areas by their level of representational invariance, with V1 converging most slowly, followed by V2, V3, V4, and higher-level visual cortex.

We also determined the relative increase in alignment over several "unrefined" reconstructions sampled at random from output distributions guided only by low-level representations ("VDVAE"), only by high-level representations ("CLIP"), or both ("CLIP+VDVAE"). In V1, the percent of reconstructions that achieve parity with ground truth is quite high for the "low-level" only reconstructions but is eventually exceeded at iteration 6 (see Figure 4 for visualizations). Analysis of the subset of reconstructions that achieved parity with ground truth for all brain areas (Figure 5 plot C) revealed a similar ordering of brain areas by their rate of convergence to parity with ground truth.




Figure 4: Examples of best reconstructions across iterations of the search process. “VDVAE” indicates reconstructions generated from a VDVAE with low-level input prepared by the Library Assembler. “CLIP” indicates reconstructions generated by SD using CLIP guidance (c_0) prepared by the Library Assembler. “CLIP+VDVAE” indicates reconstructions generated by SD using CLIP guidance (c_0) and structural guidance (z_0) prepared by the Library Assembler.

4 Limitations

One limitation of the current implementation is the rather small size of the fixed Large Library used in the Library Assembler and the smaller libraries generated throughout the search. Small library sizes make it difficult to break correlations between object categories. For example, trains are prevalent in the COCO dataset used as the large library for this work, and are often depicted near buildings. If brain activity encodes a picture of a building, it is likely that a picture of a train will be included in the subset of images with the highest brain correlation scores. Of course, this limitation becomes an advantage when reconstructing images that depict trains, but is indicative of larger biases in the prior enforced by the image library that might constrain this method’s ability to represent smaller image categories or finer semantic details within images. The prior is also entirely restricted to the image content and classes present within the library, as object classes that are out of distribution for the large library would never be represented in the reconstructions.

For the technical limitations of our implementation, the most immediate limitation lies in our hyperparameter values and schedules, which are likely suboptimal, as they were manually tuned on qualitative results obtained from only 20 samples withheld from the test set. A grid search over possible values is currently impractical due to the computational costs of the algorithm, but improvements are definitely possible. This computational cost also represents a general limitation of the algorithm during inference, as the algorithm requires far more GPU cycles than previous methods to generate the artificial libraries at the intermediate stages of the search process. This could be improved with faster image generation models, but will always represent an increase in GPU cycles of several orders of magnitude over direct decoding approaches, as it requires generating many intermediate images in addition to the final image distribution. For details on the computational requirements of our algorithm, see Appendix B.6 and B.7.

5 Conclusion

This work introduces a procedure for reducing the breadth of decoded image distributions while maximizing their alignment to brain activity. We find that this procedure generates high-quality

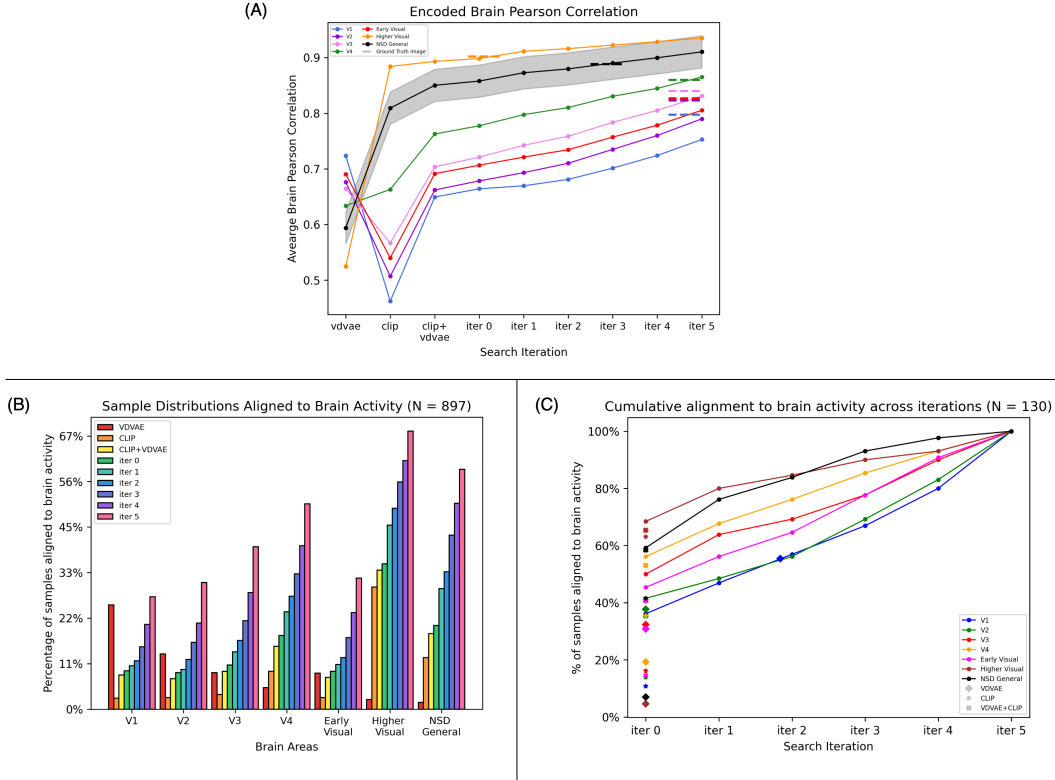


Figure 5: (A) Correlation of predicted and actual brain activity (brain correlation) for the top image at each iteration for different ROIs (curves). Dashed lines are the score of the target ground truth image in each ROI. (B) Percentage of samples in the test set for which reconstructions achieve a brain correlations score at or above parity with the score of the ground truth image (y-axis) for each of the indicated ROIs (x-axis) and across iterations (colors). (C) Cumulative percentage of reconstructions that achieve brain correlations score at or above parity with the score of the ground truth image (y-axis) across iterations (x-axis) and for “unrefined” reconstructions methods (shapes). This plot uses only samples that reach parity with the score of the ground truth image for all brain areas.

reconstructions that approximate ground truth images in a variety of feature spaces without explicitly minimizing loss in any one of them. A key component of our system is our use of brain-optimized encoding models that yield accurate predictions of brain activity in much of visual cortex [2], allowing us to leverage the diversity of representations for natural images across visual cortex. The potential of this new tool is demonstrated in quantifying the invariance in representations across the brain via each brain area’s rate of convergence to parity with the ground truth image, and finding that it sorts brain areas in agreement with complementary measures of invariance explored in other works, such as receptive field size [19]. We believe image reconstruction algorithms have a pivotal role to play in understanding and interpreting brain activity in visual cortex. By utilizing encoding models to align reconstructions directly with the measured brain activity, this work aims to refine work in this space to a more descriptive and scientific objective that allows these tools to give new insights into the structure and function of various cortical regions.

6 Acknowledgements

We would like to thank Stability AI for providing code and pre-trained models for Stable Diffusion. Funding is provided by R01EY023384.

Appendix

A Data Withholding

Previous examination in Tagaki et al. [8] has found that 65 of the images used in the NSD shared 1000 test data [3], approximately 7% of the samples, were present in LAION-5B, which Stable Diffusion and the ViT-H/14 CLIP model utilized for training. Consequently, we excluded these overlapping images from our generated samples, and they were not used in either our quantitative analysis or any of our qualitative examples.

In addition to the samples held out from LAION-5B[12], we also withhold the first 20 samples in the test set, which we used to manually tune various hyperparameters, including the number of images generated at each iteration, the number of batches, the number of iterations, the strength and momentum schedules, as well as the k parameter of the Library Assembler. We withhold these samples to avoid overfitting our selected hyperparameters to those data points.

The full list of the 897 samples our quantitative metrics were calculated on will be released along with our code and models in late June 2023.

B Architecture details

B.1 VDVAE

For the VDVAE model [16], we use the pre-trained VDVAE Encoder to extract and concatenate the latent variables for the first 31 layers of the hierarchy, creating a vector with 91,168 dimensions. This allows us to pass images from our large library through the model and collect the VDVAE latents as output, which we can then combine in the Library Assembler. Once the Library Assembler returns a set of averaged VDVAE latents, this set is then reconstructed into an image using the pre-trained VDVAE Decoder. This reconstruction serves as an initial guess for the low-level features of our reconstruction, seeding future stages of the architecture. It's important to note that all VDVAE layers, including both encoder and decoder blocks, are pre-trained and frozen, and we train no additional models for this stage of the pipeline. To convert the image into the z vector format accepted for structural guidance in Stable Diffusion, we pass it through the AutoKL encoder utilized in Stable Diffusion's `img2img` script.

B.2 CLIP

For the generation of CLIP vectors representing semantic content [13], we utilize an OpenClip implementation of the ViT-H/14 image embedding model [14, 15]. This generates a CLIPVision vector c with 1024 dimensions. Stable Diffusion 2.1 utilizes this CLIP embedding format, making it a natural choice for this work. It is notably larger than the ViT-L/14 model used in previous versions of Stable Diffusion and other diffusion models, which may contribute to an increase in the semantic fidelity of the reconstructions.

B.3 Stable Diffusion

For the diffusion model responsible for image generation, the pre-trained Stable Diffusion unCLIP model from the Diffusers library [4] is adopted. This model, released as Stable Diffusion Reimagine, is a modified version of Stable Diffusion 2.1, which has been trained to accept OpenClip ViT-H/14 image embeddings. In this work, the model pipeline is adapted to also receive a latent z vector and a strength parameter λ in addition to the c embeddings.

B.4 Denoising autoencoder

For the places of our pipeline where we want to compare encoded brain data against measured brain data to generate a brain correlation score, we train a denoising autoencoder to remove some of the additional noise present in the measured brain data. To do this, we encode the original brain data, whose layer size is the voxel count of the `nsdgeneral` mask, [15724,14278,13039,12682] for subjects

1, 2, 5, and 7, respectively, down to a set of compressed intermediary layers of 5000, 1000, and 500 dimensions. We then decode this layer back up through these layer sizes to the input layer size to produce a denoised representation of the original scan that contains only the relevant signal to our comparison. To train this model, we use the measured brain activity of a sample β as input and the encoded β' of the ground truth image computed by our hybrid encoding model detailed in section B.5 as the target. By training the model in this way, our autoencoder is trained to keep only the signal elements that our encoding model deems relevant, leading to the very high brain correlation scores seen in some metrics, often above 0.9. The fact that the images in this work are tuned to maximize the measured correlation between the outputs of the autoencoder and encoding model may limit the relevance of the brain correlation metric to other works that haven't been optimized for this measure of performance, but it still provides a useful insight into how various reconstructions represent measured brain data.

B.5 Hybrid encoder

For the hybrid encoder used in various places in our architecture, we combine the outputs of GNet, a pre-trained interpretable voxel-wise encoding model designed to extract and associate image features with the feature-weighted receptive fields of individual voxels in the visual cortex, and a 3-layer deep neural network that we train to map ViT-H/14 CLIP vectors to fMRI signals. The GNet model was initially trained on the same combined training-validation dataset employed in this research, allowing for the combination of their respective outputs. The CLIP encoding model is sensitive only to semantic features represented in the CLIP embedding, and so largely predicts voxels in areas outside the early visual cortex, while GNet provides predictions that are very accurate within the early visual cortex but less sensitive to higher regions corresponding to semantic representations of visual content. This discrepancy serves as the motivation for combining their outputs to get a more comprehensive estimate of brain activity. Different variations of combining the outputs of each method were briefly tested, including weighting individual voxel predictions by the relative prediction accuracies of each model; however, this did not lead to an increase in prediction accuracy, so a simple averaging procedure was used to combine outputs. More rigorous experiments for best combining the outputs of different encoding models could be done in the future to expand upon this implementation, such as the work done in [17].

B.6 Library Assembler

The library used for this work is a subset of the 73,000 COCO [18] images used as stimuli in the Natural Scenes Dataset [3]. Each subject of the experiment viewed up to 10,000 of these images, and so when computing the initial latent variables for a subject, the 63,000 images not seen by that subject at any point during the experiment are used. This means that the 63,000 images used in the Library Assembler stage will be different for each subject, but should contain a very similar distribution of images and semantic classes.

For the k parameter used in the Library Assembler, we performed a grid search over possible values of k on the 20 samples held out from the test set to find values that sufficiently maximized quantitative and qualitative comparisons against the latent variables extracted from the corresponding ground truth images. In this analysis, we determined that a value of ≈ 25 works well for the VDVAE latent configuration, while a value of ≈ 100 is optimal for the CLIP configuration. Notably, CLIP vectors lie in a Euclidean vector space, meaning you can average as many of them together as you like without significant loss in reconstruction quality. The optimal k value in this instance should correspond to the number of images in our large library ($n=63,000$) that fall within close proximity of the semantic class of the seen image. The VDVAE latents, however, do not fall within a neat Euclidean vector space. While they can still be averaged together, the larger the set of image latents averaged together, the greater the distortion of the produced image. The k value of 25, in this case, was selected to strike a qualitative balance between not overfitting our initial structural guess to a single library image, and not overly distorting our initial structural guess by averaging too many latents together.

The only GPU operations required for the Library Assembler are the inference stage of the encoding model used to create β' values for each image, and the inference stage of the denoising autoencoder to create $\hat{\beta}$ s for each of your measured samples. This makes the GPU overhead of the architecture quite reasonable, taking about 1 hour for both of these steps for our $n=63,000$ library on an NVIDIA A100 with 17GB of VRAM being utilized. The other major computation requirement for this

architecture is its memory requirement. The current implementation loads all 63,000 latent variables corresponding to library images into memory to provide fast access to them during the sorting stage. At 0.7mb for our largest latent variable, the VDVAE latents, that amounts to 46GB of memory usage. The inference stage of the Library Assembler, where it performs the search and averaging steps, takes about 10 minutes per subject to compute latent variables for the 897 test samples.

B.7 Stochastic Search

Within the SCS algorithm, the strength parameter λ at each iteration is calculated using the equation $\lambda = 0.92 - 0.3\left(\frac{i+1}{6}\right)^3$, where i denotes the current iteration number. The momentum parameter γ is calculated by the equation $\gamma = 0.2\left(\frac{i+1}{6}\right)^2$. These schedules, as well as the numerous hyperparameters of the algorithm, including i , the number of iterations, n the number of images generated at each iteration, and n_b the number of images in each batch, were all tuned to their respective values through qualitative testing on the first 20 held out test samples described in Section A.

The Stochastic Search algorithm presents a very large GPU computing requirement. At each iteration of the search, we use a diffusion model to generate hundreds of artificial library images and an encoding model to encode each of them into β 's to be scored against our denoised $\hat{\beta}$ s. By far, the largest cost comes in the diffusion process of image generation. Each image takes about 3 seconds to generate on an NVIDIA A100. At 100 images per iteration and 6 iterations, that is 30 minutes of computing time for just the image generations of a single test sample. To help address this, we implement an optimization to reduce the number of samples generated in each batch as the strength parameter decreases, as we need fewer samples to represent the same variance in our distribution. This optimization is calculated by taking the base number of branch samples (25) and multiplying it by our current strength parameter value λ , so the number of batch samples is calculated by $n_b = 25\lambda$. With these optimizations, we reduce the number of images generated by 37% and reduce the time requirement for the diffusion step to about 20 minutes per sample. The additional overhead of the encoding model predictions for those images as well as other file operation steps, brings our total compute time for a single sample to 31.5 minutes on an NVIDIA A100. Our total test set generations were computed across 4 A100s, one per subject, for a total of 19 days.

C Quantitative Metrics

C.1 Details for Table 1

The metrics for Ozcelik et al. [6] were calculated using 5 repetitions of the results generated using their open-source implementation. Ozcelik et al. also provided a single repetition of their results, which we verified had very similar metrics to our 5-fold reproduction. The results for Tagaki et al. [8] were calculated on 5 repetitions of their samples provided by the authors. The authors of Lu et al. [9] and Gu et al. [9] only provided one repetition of their reconstructions, and do not have open-source code, and so their metrics are calculated only across the single image for each sample that the authors provided us. This difference should not have a large effect on their metrics, besides making their results slightly less confident.

The CLIP model used for the quantitative values in both the CLIP(2-way) and CLIP(cos) metrics was the official ViT-L/14 model (768 dimensions) released by OpenAI. This is notably different than the OpenCLIP ViT-H/14 model (1024 dimensions) used for our pipeline but is a more commonly used model for calculating CLIP metrics in other papers such as Ozcelik et al[6]. and Tagaki et al.[8], and so was used for the sake of consistency with other works. It is different than the CLIP model used in Lu et al. [9], which appears to use a ViT-B/16 model (512 dimensions). Out of curiosity, we tested these other CLIP models to see if they would produce a difference in quantitative benchmarks and confirmed that the results vary drastically in scale between models, confirming the need for consistency in the calculation of this metric.

Another notable item is that for all data reported in this paper, the SSIM metric was calculated using version 0.20.0 of the Scikit-Image python package, which is the newest version at the time of this work. This version carries a significant change in functionality from previous versions of the package used to calculate the metric in previous works such as Ozcelik et al., Lu et al., and Gu et al. [6, 9, 7], so the SSIM metrics reported here may not align with the numbers reported in their papers.

C.2 Multi-subject results

The authors of Lu et al. [9] did not produce results on additional subjects beyond subject 1, and so their results are excluded from Table C.1. The remaining results are averaged across subjects 1, 2, 5, and 7, in accordance with the sample repetition methodology detailed in Section C.1.

Methods	Low Level Structural Metrics				
	PixCorr \uparrow	SSIM \uparrow	AlexNet(2) \uparrow	AlexNet(5) \uparrow	AlexNet(7) \uparrow
Second Sight (ours)	.156	.285	88.4%	93.5%	93.1%
Best Library Image (ours)	.124	.252	84.2%	91.6%	91.9%
Ozcelik et al.	.253	.284	94.1%	96.3%	95.6%
Tagaki et al.	.213	.213	71.0%	74.0%	73.9%
Gu et al.	.141	.276	78.2%	89.0%	92.5%

Methods	High Level Semantics Metrics				
	CLIP(2-way) \uparrow	CLIP(cos) \uparrow	Inception V3 \uparrow	EffNet-B \downarrow	SwAV \downarrow
Second Sight (ours)	87.0%	.655	82.0%	.792	.435
Best Library Image (ours)	84.4%	.645	80.5%	.818	.458
Ozcelik et al.	91.7%	.653	87.5%	.774	.423
Tagaki et al.	64.1%	.564	63.2%	.957	.671
Gu et al.	77.7%	.587	78.6%	.870	.459

Table C.1: Quantitative comparison against past reconstruction methods on subjects 1, 2, 5, and 7. For each measure, the best value is in bold. For EffNet-B and SwAV distances, lower is better. Higher is better for all other metrics. This is indicated by the arrow pointing up or down.

Methods	Brain Correlation Metrics					
	V1	V2	V3	V4	Higher Vis	nsdgeneral
Second Sight (ours)	.795	.814	.845	.866	.937	.92
Best Library Image (ours)	.573	.588	.606	.612	.656	.65
Ozcelik et al.	.716	.715	.739	.766	.866	.844
Tagaki et al.	.311	.260	.264	.306	.360	.353
Gu et al.	.482	.499	.524	.574	.729	.691

Table C.2: Evaluation of brain correlation scores against past reconstruction methods on subjects 1, 2, 5 and 7. Evaluating how well the reconstructed images represent activity in various brain regions according to our brain correlation metric.

D Ablation studies

Here we present an ablation study of how the various parts of our Second Sight architecture perform in isolation with regard to subject 1. In the qualitative analysis found in Figures D.1 and D.2, the Only VDVAE model demonstrated potential for reproducing low-level details, such as shapes and layouts, of the original ground truth images. However, these reconstructions often resembled hazy outlines, falling short of recreating realistic images. On the contrary, the Only CLIP model performed well in high-level properties, which are essentially the semantics of the ground truth images. However, it struggled with capturing the accurate placement and layout of objects within these images. The combination of both CLIP and VDVAE outperforms both standalone models by capturing elements of both high-level and low-level information. This combined approach, thus, serves as our initial benchmark for the reconstruction process. As the iterative search process begins, the accuracy of the reconstructions improves, aligning more closely with the reference images, thereby enhancing the accuracy of the reconstructions.

In our quantitative analysis found in Tables D.1 and D.2, we observe similar conclusions. The first variant, using only the low-level structural guess provided by the VDVAE latents, excelled in the low-level metrics, specifically PixCorr, SSIM, and AlexNet(2). However, it faltered in high-level measures and in terms of brain correlation outside of V1. The second variant, Only CLIP, yielded underwhelming results across all measures, including low-level, high-level, and brain correlation metrics. The hybrid approach of CLIP+VDVAE demonstrated intermediate results, approximating the

average performance of the two isolated methods. These outcomes provide a benchmark for image reconstruction in the future iterative stages, which aim to improve them. Observing these iterations, PixCorr and SSIM performance deteriorates as the search progresses, while AlexNet(2) maintains a stable performance. Most notably, the performances of AlexNet(5) and AlexNet(7), as well as the high-level metrics of CLIP(2-way), Inception V3, and SwAV, note an increase in performance as the search advances. However, CLIP cosine and EffNet-B performance initially improve before declining in later stages. At the fifth iteration, a significant drop in performance is noticeable across multiple metrics, including AlexNet(2), AlexNet(5), AlexNet(7), CLIP(2-way), Inception V3, EffNet-B, and SwAV. This decline is not observed in the brain correlation, as observed in the brain correlation measurements, reaffirming the conclusion that the brain’s representation of seen images is abstracted away from the original seen image, and can be adequately reproduced by many different images. Overall, the image reconstruction quality initially improves as the search progresses, reaches an optimal point, and then declines as the image begins to align more closely with the subject’s brain responses. This study provides insights into the nuanced effects of the search iterations on image reconstruction quality, as well as how the different parts of our pipeline contribute to our final results.




Figure D.1: Qualitative ablation study of our Second Sight architecture on subject 1. The samples are the same as in Figure 2 to allow for easy comparison. The first column is the ground truth image, indicated by the red background, while the remaining columns represent the output at various stages of the Second Sight algorithm.

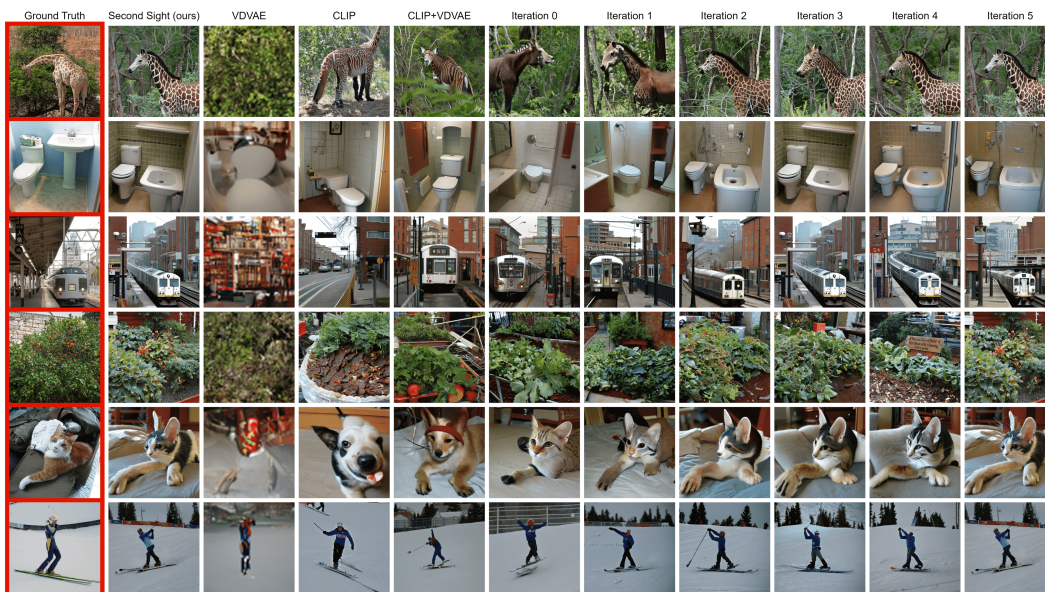


Figure D.2: Some additional ablation study samples with comparison across stages of our pipeline, ground truth images are indicated by a red background. These samples were picked as examples that specifically demonstrate the effects of the iterative refinement capabilities of our architecture.

Methods	Low Level Structural Metrics				
	PixCorr \uparrow	SSIM \uparrow	AlexNet(2) \uparrow	AlexNet(5) \uparrow	AlexNet(7) \uparrow
Only VDVAE	.252	.350	92.4%	91.7%	87.0%
Only CLIP	.083	.268	80.6%	89.8%	92.5%
CLIP+VDVAE	.220	.310	90.6%	93.7%	93.4%
Iteration 0	.219	.310	91.1%	93.8%	93.5%
Iteration 1	.206	.296	90.8%	94.3%	93.9%
Iteration 2	.197	.293	90.6%	94.5%	93.9%
Iteration 3	.186	.288	90.7%	94.5%	94.0%
Iteration 4	.181	.285	91.8%	94.7%	94.1%
Iteration 5	.136	.281	76.9%	87.1%	90.4%
Second Sight (ours)	.185	.288	91.7%	95.0%	94.2%

Methods	High Level Semantics Metrics				
	CLIP(2-way) \uparrow	CLIP(cos) \uparrow	Inception V3 \uparrow	EffNet-B \downarrow	SwAV \downarrow
Only VDVAE	67.2%	.571	70.1%	.933	.600
Only CLIP	87.8%	.656	82.2%	.792	.432
CLIP+VDVAE	87.7%	.661	83.0%	.788	.425
Iteration 0	87.7%	.665	82.8%	.786	.423
Iteration 1	88.1%	.665	83.1%	.783	.421
Iteration 2	88.1%	.666	83.1%	.780	.420
Iteration 3	88.2%	.663	83.3%	.779	.421
Iteration 4	88.0%	.662	83.8%	.780	.421
Iteration 5	75.3%	.660	75.8%	.893	.479
Second Sight (ours)	88.3%	.664	84.7%	.776	.418

Table D.1: Quantitative ablation study on subject 1. For each measure, the best value is in bold. For EffNet-B and SwAV distances, lower is better. Higher is better for all other metrics. This is indicated by the arrow pointing up or down.

Methods	Brain Correlation Metrics					
	V1	V2	V3	V4	Higher Vis	nsdgeneral
Only VDVAE	.724	.676	.665	.634	.525	.594
Only CLIP	.462	.507	.567	.663	.884	.809
CLIP+VDVAE	.649	.662	.704	.763	.893	.850
Iteration 0	.664	.678	.721	.777	.898	.858
Iteration 1	.670	.693	.742	.798	.911	.873
Iteration 2	.681	.710	.759	.810	.916	.880
Iteration 3	.701	.735	.784	.831	.922	.890
Iteration 4	.724	.760	.805	.845	.929	.900
Iteration 5	.753	.790	.831	.865	.935	.910
Second Sight (ours)	.739	.774	.817	.854	.932	.905

Table D.2: Ablation study of brain correlation scores on subject 1. Evaluating how well the reconstructed images represent activity in various brain regions according to our hybrid encoding model.

E Additional samples

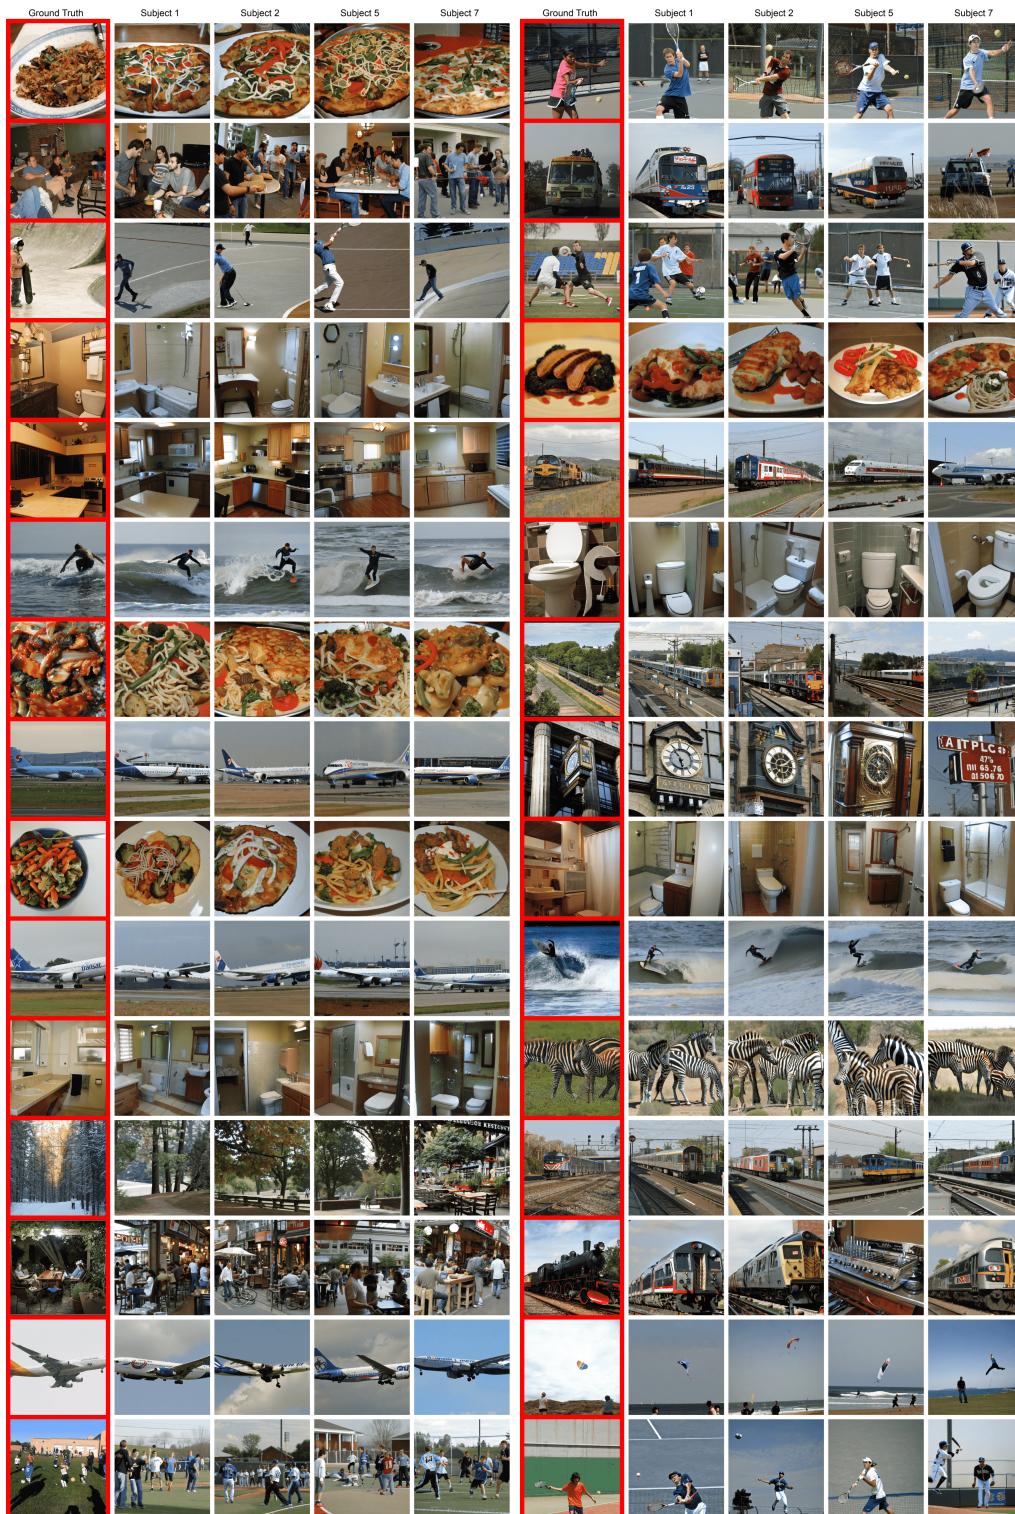


Figure E.1: Some additional samples with comparison across subjects, ground truth images are indicated by a red background.



Figure E.2: Some additional samples with comparison across subjects, ground truth images are indicated by a red background.

References

- [1] Thomas Naselaris, Kendrick N. Kay, Shinji Nishimoto, and Jack L. Gallant. Encoding and decoding in fmri. *NeuroImage*, 56(2), 2011.

- [2] Ghislain St-Yves, Emily J. Allen, Yihan Wu, Kendrick Kay, and Thomas Naselaris. Brain-optimized neural networks learn non-hierarchical models of representation in human visual cortex. *bioRxiv*, 2022.
- [3] E.J. Allen, G. St-Yves, and Y. et al. Wu. A massive 7t fmri dataset to bridge cognitive neuroscience and artificial intelligence. In *Nat Neurosci*.
- [4] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models, 2022.
- [5] Sikun Lin, Thomas Sprague, and Ambuj K Singh. Mind reader: Reconstructing complex images from brain activities, 2022.
- [6] Furkan Ozcelik and Rufin VanRullen. Brain-diffuser: Natural scene reconstruction from fmri signals using generative latent diffusion, 2023.
- [7] Zijin Gu, Keith Jamison, Amy Kuceyeski, and Mert Sabuncu. Decoding natural image stimuli from fmri data with a surface-based convolutional network, 2023.
- [8] Yu Takagi and Shinji Nishimoto. High-resolution image reconstruction with latent diffusion models from human brain activity. *bioRxiv*, 2023.
- [9] Yizhuo Lu, Changde Du, Dianpeng Wang, and Huiguang He. Minddiffuser: Controlled image reconstruction from human brain activity with semantic and structural diffusion, 2023.
- [10] Ghislain St-Yves and Thomas Naselaris. Generative adversarial networks conditioned on brain activity reconstruct seen images. pages 1054–1061. Institute of Electrical and Electronics Engineers Inc., January 2019.
- [11] Thomas Naselaris, Ryan J. Prenger, Kendrick N. Kay, Michael Oliver, and Jack L. Gallant. Bayesian reconstruction of natural images from human brain activity. *Neuron*, 63(6):902–915, 2009.
- [12] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, Patrick Schramowski, Srivatsa Kundurthy, Katherine Crowson, Ludwig Schmidt, Robert Kaczmarczyk, and Jenia Jitsev. Laion-5b: An open large-scale dataset for training next generation image-text models, 2022.
- [13] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision, 2021.
- [14] Gabriel Ilharco, Mitchell Wortsman, Ross Wightman, Cade Gordon, Nicholas Carlini, Rohan Taori, Achal Dave, Vaishal Shankar, Hongseok Namkoong, John Miller, Hannaneh Hajishirzi, Ali Farhadi, and Ludwig Schmidt. Openclip, July 2021.
- [15] Mehdi Cherti, Romain Beaumont, Ross Wightman, Mitchell Wortsman, Gabriel Ilharco, Cade Gordon, Christoph Schuhmann, Ludwig Schmidt, and Jenia Jitsev. Reproducible scaling laws for contrastive language-image learning, 2022.
- [16] Rewon Child. Very deep vaes generalize autoregressive models and can outperform them on images, 2021.
- [17] Ruogu Lin, Thomas Naselaris, Kendrick Kay, and Leila Wehbe. Stacked regressions and structured variance partitioning for interpretable brain maps. *bioRxiv*, 2023.
- [18] Tsung-Yi Lin, Michael Maire, Serge Belongie, Lubomir Bourdev, Ross Girshick, James Hays, Pietro Perona, Deva Ramanan, C. Lawrence Zitnick, and Piotr Dollár. Microsoft coco: Common objects in context, 2015.
- [19] Kendrick N. Kay, Jonathan Winawer, Aviv Mezer, and Brian A. Wandell. Compressive spatial summation in human visual cortex. *Journal of Neurophysiology*, 110(2):481–494, 2013. PMID: 23615546.