In Silico Mapping of Visual Categorical Selectivity Across the Whole Brain

Ethan Hwang

Zuckerman Mind Brain Behavior Institute Columbia University eh2976@columbia.edu

Wenxuan Guo

Zuckerman Mind Brain Behavior Institute Columbia University wg2361@columbia.edu

Hossein Adeli

Zuckerman Mind Brain Behavior Institute Columbia University ha2366@columbia.edu

Andrew Luo

University of Hong Kong aluo@hku.hk

Nikolaus Kriegeskorte

Zuckerman Mind Brain Behavior Institute Columbia University nk2765@columbia.edu

Abstract

A fine-grained account of functional selectivity in the cortex is essential for understanding how visual information is processed and represented in the brain. Classical studies using designed experiments have identified multiple category-selective regions; however, these approaches rely on preconceived hypotheses about categories. Subsequent data-driven discovery methods have sought to address this limitation but are often limited by simple, typically linear encoding models. We propose an in silico approach for data-driven discovery of novel category-selectivity hypotheses based on an encoder-decoder transformer model. The architecture incorporates a brain-region to image-feature cross-attention mechanism, enabling nonlinear mappings between high-dimensional deep network features and semantic patterns encoded in the brain activity. We further introduce a method to characterize the selectivity of individual parcels by leveraging diffusion-based image generative models and large-scale datasets to synthesize and select images that maximally activate each parcel. Our approach reveals regions with complex, compositional selectivity involving diverse semantic concepts, which we validate in silico both within and across subjects. Using a brain encoder as a "digital twin" offers a powerful, data-driven framework for generating and testing hypotheses about visual selectivity in the human brain—hypotheses that can guide future fMRI experiments. Our code is available at: https://kriegeskorte-lab.github.io/in-silico-mapping/.

1 Introduction

Over the past few decades, researchers have extensively studied the visual hierarchy in the brain, from early cortical areas that encode low-level features to higher-level regions that represent categorical information. Neuroimaging experiments, especially studies using functional magnetic resonance imaging (fMRI), have revealed specialized cortical regions for faces, places, words, bodies, and food [48, 35, 5, 63, 15, 8, 21, 20, 61, 65, 34, 38, 69, 54]. However, visual perception is more nuanced than this short list of categories. It remains unknown what additional visual concepts have

dedicated regions that enable humans to make sense of the complex world. Common mapping methods depend on experimenter-curated concepts, and empirically-driven alternatives require more data and expensive fMRI experiments. For example, one could present a subject with a large set of images and then label the selectivity of each cortical parcel by the images that elicit the strongest mean parcel response. As the stimulus set grows, however, the cost of data acquisition (operating the scanner, paying subjects, bonuses)—at least with current fMRI technology and experimental paradigms—may prove prohibitive. We propose methods that rely on the state-of-the-art encoding models to generalize to concepts beyond the stimuli for which fMRI responses have been measured.

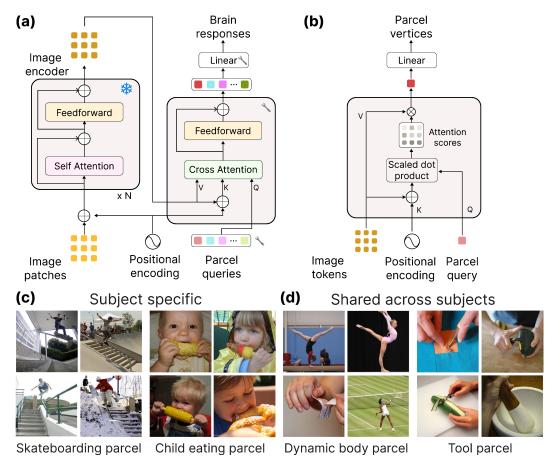


Figure 1: Brain encoder architecture and images predicted to maximally activate selected parcels outside the visual cortex. (a) Brain encoder architecture. (b) Cross-attention for parcel fMRI prediction. (c) A brain encoder ranks ImageNet [17] images by how much each image would activate a parcel. Images from two sample parcels, four images from a single subject are shown. (d) Images from two sample parcels, one image from each of the four subjects is shown for each parcel. All images are curated from the top 25, see Appendix A.7 for the parcel locations and full collage.

Our encoding model leverages recent advances in AI and large-scale neural datasets to serve as a "digital twin," upon which we perform extensive in silico experimentation [27, 29] to generate hypotheses of complex categorical selectivity beyond the visual cortex. The encoding model enables us to predict the neural activity of image sets far larger and more diverse than the limited number of images shown to the participant during fMRI, effectively expanding the search space for optimal visual stimuli. In addition, the encoding model is fully observable and differentiable, enabling interpretability queries using attention maps and gradient-based analyses, such as diffusion-based models that use gradients to find stimuli that elicit high activity [46].

Through experimentation with large image datasets on our model, we generated hypotheses for the selectivity of many parcels outside the visual area, as well as optimal sets of images that maximally activate the parcels. These hypotheses can be tested in targeted future fMRI experiments by showing only the sets of optimal stimuli, accelerating data collection and experiment iteration, and lowering

the cost of data acquisition. We demonstrate that our pipeline using brain encoders can test concepts that were not explicitly shown to the subject in the scanner, effectively enriching the diversity and size of the fMRI training set.

Key contributions of our work:

- Massive scale: applying in silico mapping on millions of images (ImageNet, BrainDIVE)
 with a transformer-based brain encoder, enabling discovery of parcel selectivity for concepts
 never shown in training. To the best of our knowledge, no other study has been done on this
 scale.
- Mapping of the whole brain: expanding beyond visual cortex and revealing human-specific semantic selectivity.
- 3. **In silico verification:** our pipeline verifies selectivity hypotheses in silico with rigorous tests that evaluate how well a label can predict ground-truth activation on a held-out set within and across subjects.
- 4. **New fMRI experimental paradigm:** as datasets grow and encoding models improve, our pipeline offers a way to leverage these advances to accelerate and improve the accuracy of whole-brain mapping.

2 Related Work

Semantic mapping. Our work builds upon a growing body of computational modeling and machine learning research that investigates how semantic information is represented in the higher visual cortex [42]. Some approaches leverage large image datasets to build decoders [19, 22, 66, 36, 62, 74, 14, 44] or models for generating optimal stimuli [13, 46, 45, 28, 58, 47], while others use cross-domain (e.g. vision-language) mapping [33, 32, 37, 9, 50, 51, 55, 64, 30]. These studies face a key challenge of dataset size, since the collection of neural data is often expensive and time-consuming. Our work seeks to address this constraint by using encoders trained on large-scale datasets to perform in silico mapping. This allows us to expand the set of concepts that can be probed, beyond those stimuli shown to the subjects.

Brain encoding models. Highlighting its importance, several community-driven efforts have sought to benchmark models predicting brain responses [60, 68, 26]. With the release of increasingly large neural datasets each year, researchers have introduced novel architectures and methodologies to improve the accuracy of brain encoding model, including leveraging multiple datasets and pretrained networks [2, 43, 73, 11, 72, 57, 24, 41, 67, 7]. While our paper uses a cutting-edge encoding model [3], our pipeline is ultimately encoder-agnostic, and can use any encoder that is image-computable. As researchers build better brain encoders, we expect that the space of hypotheses our pipeline could generate and their accuracy will only grow.

Brain-optimized stimuli. Previous studies have introduced encoding model-based stimulus selection and empirically validated the superstimuli in non-human primates and mice [10, 71, 56, 23]. In this work, we extend this general approach of stimulus optimization for studying neural populations to the fMRI domain, revealing high-level human-specific selectivity beyond the visual cortex.

3 Methods

Our goal is to map the visual selectivity of parcels across the whole brain. First, we train a brain encoder to predict fMRI responses from natural scene images. Then we select visually responsive and robust parcels for experimentation to determine categorical selectivity.

3.1 Parcellation Strategy

We partitioned the 327,684 cortical vertices across the whole brain into 1,000 functional parcels using the Schaefer resting-state functional connectivity parcellation (see Figure 2) [59].



Figure 2: Schaefer-1000 parcellation

3.2 Brain Encoder Architecture

Extending the work of Adeli et al. [2], our brain encoder predicts vertex-wise fMRI activity across the whole brain from an input image (see Figure 1a). An image encoder backbone (see Section 4.2) first extracts patch embeddings from the image. A transformer decoder uses parcel-specific, learnable queries to attend to relevant patch embeddings via cross-attention. The decoder consists of a single transformer layer with cross-attention followed by a feedforward projection (see Figure 1b for cross-attention). Each output decoder token is linearly mapped to predict the fMRI responses for vertices in the corresponding parcel, and predictions from all output tokens are aggregated to obtain a whole-brain prediction. Parcel queries, the transformer decoder, and linear mappings are optimized using Adam [39] to minimize the mean squared error between the predicted and actual fMRI responses. All other layers, including the backbone, are frozen. Separate models are trained for each subject.

To improve prediction accuracy, we ensemble multiple instances of the brain encoder. For each subject, we trained two random seeds with features from four different DINOv2 backbone layers (the 0th, 2nd, 4th, and 6th layers from the last). To predict a vertex, we take the weighted average across model predictions, scaled by softmax weights from validation set accuracy for each model on that vertex.

3.3 Parcel selection process for further experimentation

Not all regions in the brain are visually responsive, so we selected parcels for further experimentation that satisfy three criteria: (1) Location: Fewer than 10% of parcel vertices overlap with the labeled area, since we are most interested in parcels beyond the visual cortex. (2) Visual responsiveness: The average noise ceiling must be in the top 25% of parcels that satisfy (1). (3) Model prediction accuracy: The average prediction accuracy must be in the top 25% of parcels that satisfy (1).

Since the data quality varies slightly across subjects, we used percentiles rather than numerical cutoffs. The selection process for parcels in subject 1's left hemisphere is shown in Figure S1 in Appendix A.3. Cutoffs are determined separately for each hemisphere to maintain comparable parcel counts.

Among the 500 Schaefer parcels in each hemisphere for every subject, 409 ± 6 parcels satisfy condition (1). After filtering out parcels with low mean noise ceiling or low mean model prediction accuracy, 179 ± 12 parcels per subject are chosen for further experimentation.

3.4 Superstimulus Generation Process

We choose images that maximally activate (mean z-scored beta values) a parcel of interest using three different methods:

- 1. Natural Scenes Dataset (NSD) Ground Truth Images: Images from the held-out NSD [4] test split, ranked based on ground truth data.
- 2. **Diffusion-generated superstimuli:** BrainDIVE [46] uses a generative backbone guided by gradients from a brain encoder to generate images that can maximally activate specified brain parcels. We generated 400 images per parcel and reranked them with the brain encoder.
- 3. **Encoder-selected ImageNet superstimuli:** ImageNet [17] images that maximally activate the parcel, according to the encoder.

4 Experiments

4.1 Setup

We used the NSD [4], the largest fMRI dataset to date, with 7T fMRI responses from 8 subjects who each viewed up to 10,000 distinct natural scenes. Each image is presented up to three times, and our model is trained on the neural response averaged over the presentations. We report results for subjects who completed all NSD scan sessions (1, 2, 5, and 7), though we observed comparable results in all subjects. fMRI responses were preprocessed according to [4]. The resulting beta estimates were centered to zero mean and scaled to unit variance before training and experiments. ROI labels were obtained from NSD. V1–hV4 ROIs are derived from a pRF experiment; body-, face-, place-, word-selective ROIs are derived from a fLoc experiment [4].

4.2 Brain encoder results

Model prediction accuracy. Figure 3 shows the performance of the ensemble model with DINOv2 (ViT-B) backbone and the transformer cross-attention mapping function for subject 1 projected onto the cortical surface using Pycortex [25]. As expected, the model performs well on predicting the activity in the visual cortex (the area in the center of the flatmap), but also on several regions beyond the typical visual pathways.

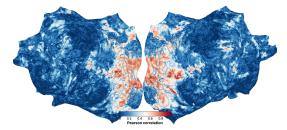


Figure 3: **Brain encoder prediction accuracy.** Pearson correlation between model predictions and ground truth data for subject 1 on the held-out test set.

Model comparison. We evaluate several encoding models with different backbones and mapping functions to find the most suitable model for whole-brain voxel-wise prediction. The Algonauts 2023 challenge [26] leaderboard showed that transformer-based backbones generally outperformed other model families such as convolutional networks (CNNs) at predicting fMRI activity. Therefore, we focused on evaluating features from several transformer backbones paired with either transformer attention-based mapping to neural data [2] or a parameter-matched model that linearly maps the CLS token to vertex values. For each architecture, we used an ensemble of 16 models for each subject (4 feature backbone layers x 2 hemispheres x 2 random seeds). We compared the encoding accuracy, which is Pearson's correlation on the held-out set corrected for noise ceilings (see [4] Methods, Noise ceiling estimation).

TC 1 1 1 TO	. 1	1.	•	1' CC	1 1
Table I. Rra	in encoder e	ancoding.	accuracy using	different	architectures
Table 1. Dia	iii ciicoaci t	circouning	accuracy using	different	architectures

Architecture	S 1	S2	S5	S 7	Backbone + Mapping fn size (M)
DINOv2 (ViT-B) + Linear	0.33	0.34	0.39	0.33	$\sim 87 + 252$
CLIP vision + Transf.	0.44	0.41	0.46	0.43	$\sim 304 + 258$
RADIOv2.5-H + Transf.	0.26	0.35	0.34	0.30	$\sim 652 + 258$
DINOv2 (ViT-G w/reg)					
+ Transf.	0.45	0.43	0.48	0.45	$\sim 1136 + 258$
DINOv2 (ViT-B) + Transf.	0.45	0.43	0.48	0.43	$\sim 87 + 258$

Across the subjects that completed all NSD scan sessions (Table 1), the transformer models with the DINOv2 backbone [53] outperform the other non-DINOv2 backbones. The two DINOv2 backbones perform similarly, which is consistent with past work [1] that showed diminishing or worse performance for DINOv2 models larger than the base (ViT-B). The transformer-based encoder

significantly outperformed the linear baseline, leveraging the attention mechanism to flexibly route information [2, 6, 56]. Figure 4a plots the difference in prediction accuracy (Pearson's correlation) between the two.

Figure 4b plots the difference in prediction accuracy between DINOv2 and CLIP vision. While DINOv2 performs far better in the visual cortex, they perform similarly in many regions outside the visual area. For the rest of the paper, we performed our experiments on the brain encoder with DINOv2 (ViT-B) + Transformer, since it offered the best overall speed-performance tradeoff.

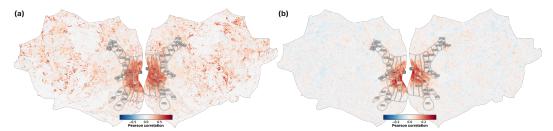


Figure 4: (a) DINOv2 Transformer vs. Linear head-to-head comparison. Difference in prediction accuracy for subject 1 between DINOv2 (ViT-B) with transformer vs. linear mapping functions (> 0 or red means transformer is better). (b) DINOv2 vs. CLIP head-to-head comparison. Difference in prediction accuracy for subject 1 between DINOv2 (ViT-B) and CLIP vision (> 0 or red means DINOv2 is better).

4.3 Sanity check on known regions

We first validate our paradigm by replicating the well-documented category selectivity of ventral pathway categorical areas. We show the results for a sample parcel from the labeled area, one which significantly overlaps with aTL-faces (47.6% of the vertices in the parcel overlap with aTL-faces). As shown in Figure 5, the images maximally activating the parcel overlapping with aTL-faces prominently feature faces, which agrees with previous work on the selectivity of this area [63]. In Appendix A.2, we reproduce the selectivity of body-, place-, and word-selective areas. In Appendix A.4, we show that stimuli generated by BrainDIVE reproduce the fine-to-coarse visual hierarchy progressing from V1–V4 to FFA.

4.4 Extension to unlabeled areas

We now extend our analyses to parcels beyond the visual cortex. We show images selected by the encoder to maximally activate each parcel for qualitative evaluation, then verify that our labels indeed explain parcel activations using formal statistical tests.

Labeling unlabeled parcels. We apply the same image-selection method to generate selectivity hypotheses for unlabeled parcels. Figure 1c displays ImageNet images predicted by the encoder to maximally activate two unlabeled subject-specific parcels: one appears to depict skateboarding, the other a child eating food. Full results, including parcel location and the corresponding NSD and BrainDIVE images, are shown in Figures S5 and S6 (Appendix A.7). Many parcels with selectivity that appeared to be consistent and complex were found across the whole brain, suggesting that there exist regions in the brain that respond to more complex concepts than the basic categories labeled in visual cortex.

Evaluating our labels. Our sanity checks demonstrated that our brain encoder retrieves images that align with the category selectivity of the labeled parcels. However, since the parcels we are interested in are outside the visual area and therefore unlabeled, we would like to quantitatively show that the images selected by our encoder faithfully reflect underlying neural selectivity.

We conducted two tests against the NSD test set—the only held-out fMRI data not used during training. If the selected images (hereafter "labels") genuinely reflect parcel selectivity, other images with greater semantic similarity to the label should result in higher parcel activation. Prior studies have shown that categorical selectivity is not binary but graded across diverse stimuli [12, 49]. For

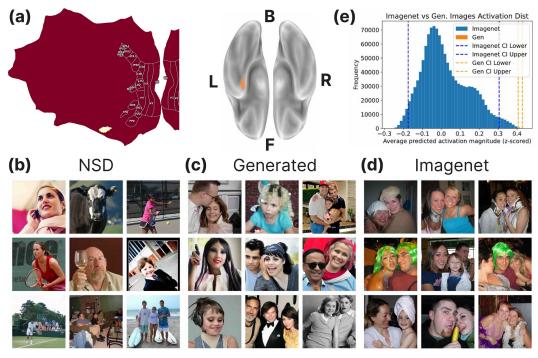


Figure 5: **Verifying the selectivity of aTL-faces.** (a) The location of the parcel. The inflated cortical surface (ventral view), with the left hemisphere on the left. (b) Held-out NSD images that maximally activate the parcel (based on ground-truth fMRI). (c) BrainDIVE generated and re-ranked images optimized for the parcel. (d) Maximally-activating images from ImageNet according to the encoder. (e) Distribution of predicted parcel activation for all of ImageNet images compared to images in (c).

example, fMRI responses to 1,705 object and action categories revealed that brain-response similarity strongly correlates with semantic similarity [33]. Accordingly, we expect a parcel's mean activity to scale with how semantically close an image is to that parcel's preferred concept, as defined by our encoder's maximally activating examples.

In each test, we compared two predictors: (1) labels derived from our encoder, and (2) a baseline that forms its hypothesis from each parcel's most activating NSD-training images, selected using their measured responses. This baseline mirrors conventional fMRI studies that rely solely on stimuli shown in the scanner. If our encoder performs better than the baseline training set, it shows we can use brain encoders to discover novel concepts not explicitly shown to the subject.

For Test 1, we asked whether semantic similarity to a label predicted the activation rank order of NSD test images better than chance. Our pipeline represents a label as the mean CLIP embedding of the top 32 ImageNet or BrainDIVE images (both shown separately); the baseline uses the top 32 NSD-training images. Appendix A.10 details the procedure. Table 2 reports the number of parcels for which each label significantly outperforms a random ranking. To ensure that the results reflect an expansion of the stimulus space and not merely dense sampling near peak activation, Appendix A.16 reports the same results with a varying number of top images used to create the label. Each parcel whose model-derived label outperforms the null is deemed successfully labeled. Across tests, the encoder-selected stimuli generally outperformed the baseline, labeling a greater number of parcels.

Table 2: Fraction of parcels whose model-derived label predicts parcel activation rankings significantly better than chance (p < 0.05, FDR corrected)

Model type	S 1	S2	S5	S7
NSD train Our encoder w/ImageNet Our encoder w/BrainDIVE	150/181 139/181 135/181	$163/192 \\ 167/192 \\ \mathbf{170/192}$	136/175 $130/175$ $139/175$	155/196 150/196 156/196

For Test 2, we compare Spearman's rank correlations between the ground-truth activation ordering and each model's predicted ordering, quantifying which model's selectivity hypothesis best explains parcel activations. Table 3 summarizes these coefficients across all parcels. A head-to-head comparison between our encoder and the baseline appears in Table 8 (Appendix A.13).

Table 3: Spearman's ρ (mean \pm std) between the model-predicted and ground-truth activation rankings on the NSD test set, averaged across parcels.

Model type	S 1	S2	S5	S7
Null	0.000 ± 0.045	0.000 ± 0.045	0.000 ± 0.045	0.000 ± 0.045
NSD train	0.162 ± 0.098	0.164 ± 0.091	0.150 ± 0.094	0.148 ± 0.086
Our encoder w/INet	0.168 ± 0.106	0.163 ± 0.082	0.142 ± 0.092	0.133 ± 0.075
Our encoder w/BD	0.163 ± 0.121	0.190 ± 0.099	0.154 ± 0.094	0.133 ± 0.083

Notes. INet = ImageNet. BD = BrainDIVE.

The rank ordering derived from the encoder-selected stimuli generally outperforms the baseline ordering based on the NSD test set. To contextualize the rank correlation magnitude, we report in Appendix A.15 the results for known parcels in the visual area as a benchmark for what constitutes meaningful selectivity. The average rank correlation is in the range of areas like PPA and RSC—both widely studied and accepted in the literature. Because these encoder-driven selectivity hypotheses better explain activation patterns in the NSD test images, our results indicate that the brain-encoder pipeline can generate finer-grained categorical hypotheses than those afforded by the stimuli actually shown to the subject in the scanner.

Choosing parcels for future fMRI experimentation. For an fMRI study aimed at uncovering the selectivity of parcels outside visual cortex, an experimenter may be interested in parcels whose activity can be well explained by a semantic label. To prioritize such parcels, we define a metric that ranks the quality of the hypotheses generated by our pipeline. In Figure 6, we plot the proportion of top 32 maximally-activating images successfully retrieved against the number of images retrieved by the pipeline. The top five parcels per subject (ranked by area under the curve) are highlighted.

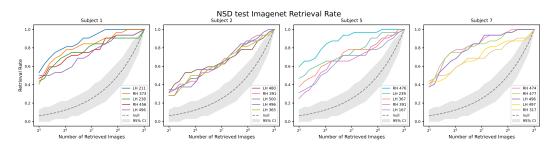


Figure 6: **Retrieval accuracy of parcel labels.** Using the concept vector as the label for each parcel, we retrieve a varying number of images from the test set based on cosine similarity, and calculate the fraction of overlap with 32 maximally-activating NSD test images (i.e. recall, on the y-axis). Curves that rise quickly and plateau high indicate concept vectors whose semantics closely match the parcel's true selectivity.

From an experimenter's perspective, the parcels shown in Figure 6 may be promising targets for follow-up fMRI studies, since semantic labels seem to capture the parcel selectivity exceptionally well. Our pipeline identifies images whose semantic representations appear to maximally activate these parcels—at least insofar as this can be evaluated with the modest NSD test set.

4.5 Mapping Cross-subject Selectivity

We now explore the selectivity of a parcel outside visual cortex (Figure 7a) by examining ImageNet images that the encoder predicts will maximally activate it. From the top 9 images per subject, we display 3 from each participant who completed all NSD scanning sessions. These maximally-activating images lie substantially outside the parcel's activation distribution for the remainder of

ImageNet (Figure 7c). The selected images consistently depict hands manipulating tools—such as writing instruments or cooking utensils (see full collages in Appendix A.9). When asked to identify a common theme among the top 25 ImageNet images, ChatGPT likewise highlights hands with objects. Recent work has reported tool-use representations in nearby cortical regions [16].

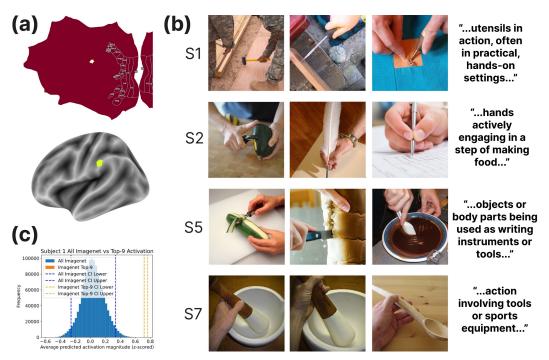


Figure 7: (a) Parcel location. (b) Selected from the 9 images in ImageNet that maximally activate the parcel of interest, across subjects 1, 2, 5, and 7. ChatGPT 40 generated labels for top-25 ImageNet images, shown on the right (see Appendix A.18 for details). (c) A comparison of the activation magnitude of the parcel from all ImageNet images and the top 9 images, with 95% confidence intervals.

Evaluating our labels. We qualitatively showed that our pipeline can uncover a parcel exhibiting consistent semantic selectivity across subjects. We now seek to quantitatively verify these hypotheses on ground-truth data. We ran statistical test similar to that in Section 4.4, with two key differences because we are evaluating selectivity across subjects.

- 1. **Combining across subjects:** For each parcel, we form the optimal image set by combining the top-32 maximally-activating images chosen by encoder models trained on the other three subjects.
- 2. **Retrieval set:** Because each subject's NSD training set contains distinct images, the ranking retrieval is performed on that subject's own NSD training set.

We again compare the selectivity hypotheses generated by our brain-encoder pipeline with the NSD-training-set baseline (described in the within-subject statistical tests). Table 4 reports the results of this modified test (cf. Table 2); additional details are provided in Appendix A.11.

Table 4: Fraction of parcels shared across subjects whose model-derived label predicts parcel activation rankings significantly better than chance (p < 0.05, FDR corrected)

Model type	S 1	S2	S5	S7
NSD train	37/49	44/49	40/49	40/49
Our encoder w/ImageNet	38/49	45/49	42/49	39/49
Our encoder w/BrainDIVE	37/49	40/49	38/49	38/49

Comparing the number of parcels for which each hypothesized label ranks NSD-training images better than chance, our pipeline generally outperforms the NSD-train baseline. These findings indicate that for parcels with high visual responsiveness shared across subjects, our pipeline can generate better hypotheses that describe the shared selectivity.

Table 5 presents results from the same statistical test used in Table 3, reporting the average Spearman rank correlations for each model across parcels. A head-to-head comparison appears in Table 9 (Appendix A.13), and Figure S8 (Appendix A.14) highlights the most promising cross-subject parcels for future fMRI experiments.

Table 5: Spearman's ρ (mean \pm std) between the model-predicted (from all subjects other than heldout) and ground-truth activation rankings on the NSD training set, averaged across parcels.

Model type	S 1	S2	S5	S7
Null	0.000 ± 0.011	0.000 ± 0.011	0.000 ± 0.011	0.000 ± 0.011
NSD train	0.068 ± 0.068	0.094 ± 0.065	0.090 ± 0.074	0.091 ± 0.070
Our encoder w/INet	0.072 ± 0.071	0.105 ± 0.073	0.105 ± 0.078	0.093 ± 0.075
Our encoder w/BD	0.078 ± 0.077	0.103 ± 0.087	0.106 ± 0.091	0.089 ± 0.079

Notes. INet = ImageNet. BD = BrainDIVE.

5 Discussion

Leveraging recent advances in AI and the availability of large-scale datasets, we introduce a datadriven paradigm for discovering parcel selectivity beyond visual cortex, paving the way toward systematic whole-brain labeling of higher-order visual representations. In particular, our transformerbased encoder, with its cross-attention mechanism and nonlinear mappings, routes visual information more effectively for parcels in and outside classical visual areas.

Limitations. Although we predict images that should maximally activate a parcel, we have not yet validated these "superstimuli" in new fMRI experiments. On a held-out NSD split we show that our semantic labels predict parcel activations well, but whether the synthesized superstimuli elicit even stronger responses remains unknown. Moreover, because the encoder was trained solely on NSD [4], it may inherit dataset biases. For example, one parcel in subject 2's left hemisphere appears zebra-selective (Figure S10, Appendix A.19), however, this could be due to the over-representation of giraffe and zebra images in the dataset. We did not quantify how such biases affect our results; future work should train on multiple datasets to address this concern.

More fundamentally, category selectivity is limited as a theoretical construct for understanding brain computation. We hope our approach will help the field go beyond the technical limitations of prior work and help reveal the fundamental limits of understanding brain computation through the lens of category selectivity.

Optimizing future fMRI experiments. Our framework paves the way for future fMRI studies that mitigate the effect of small datasets and experimenter bias in image selection by demonstrating the promise of in silico mapping to superstimuli. These superstimuli can then be tested in follow-up scans to verify the selectivity of newly discovered areas. By letting the encoder drive hypothesis generation, researchers can discover optimal stimuli empirically—even for concepts never presented in the scanner—maximizing data-collection efficiency. Because our encoder architecture is modality-agnostic, the same semantic-mapping approach could be extended to multi-modal backbones, enabling superstimuli generation across sensory domains.

Acknowledgments

Research reported in this publication was supported in part by the National Institute of Neurological Disorders and Stroke of the National Institutes of Health under award numbers 1RF1NS128897 and 4R01NS128897. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

References

- [1] Hossein Adeli, Seoyoung Ahn, Nikolaus Kriegeskorte, and Gregory Zelinsky. Affinity-based attention in self-supervised transformers predicts dynamics of object grouping in humans, 2023. URL https://arxiv.org/abs/2306.00294.
- [2] Hossein Adeli, Sun Minni, and Nikolaus Kriegeskorte. Predicting brain activity using Transformers, August 2023. URL http://biorxiv.org/lookup/doi/10.1101/2023.08.02.551743.
- [3] Hossein Adeli, Minni Sun, and Nikolaus Kriegeskorte. Transformer brain encoders explain human high-level visual responses, May 2025. URL http://arxiv.org/abs/2505.17329. arXiv:2505.17329 [q-bio].
- [4] Emily J. Allen, Ghislain St-Yves, Yihan Wu, Jesse L. Breedlove, Jacob S. Prince, Logan T. Dowdle, Matthias Nau, Brad Caron, Franco Pestilli, Ian Charest, J. Benjamin Hutchinson, Thomas Naselaris, and Kendrick Kay. A massive 7T fMRI dataset to bridge cognitive neuroscience and artificial intelligence. *Nature Neuroscience*, 25(1):116–126, January 2022. ISSN 1097-6256, 1546-1726. doi: 10.1038/s41593-021-00962-x. URL https://www.nature.com/articles/s41593-021-00962-x.
- [5] Truett Allison, Gregory McCarthy, Anna Nobre, Aina Puce, and Aysenil Belger. Human extrastriate visual cortex and the perception of faces, words, numbers, and colors. *Cerebral cortex*, 4(5):544–554, 1994.
- [6] Mehdi Azabou, Vinam Arora, Venkataramana Ganesh, Ximeng Mao, Santosh Nachimuthu, Michael Mendelson, Blake Richards, Matthew Perich, Guillaume Lajoie, and Eva L. Dyer. A unified, scalable framework for neural population decoding. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- [7] Mehdi Azabou, Vinam Arora, Venkataramana Ganesh, Ximeng Mao, Santosh Nachimuthu, Michael J. Mendelson, Blake Richards, Matthew G. Perich, Guillaume Lajoie, and Eva L. Dyer. A unified, scalable framework for neural population decoding, 2023. URL https://arxiv.org/abs/2310.16046.
- [8] Chris I. Baker, Jia Liu, Lawrence L. Wald, Kenneth K. Kwong, Thomas Benner, and Nancy Kanwisher. Visual word processing and experiential origins of functional selectivity in human extrastriate cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 104(21):9087–9092, May 2007. ISSN 0027-8424 1091-6490. doi: 10.1073/pnas. 0703300104. Place: United States.
- [9] Alfonso Barrós-Loscertales, Julio González, Friedemann Pulvermüller, Noelia Ventura-Campos, Juan Carlos Bustamante, Víctor Costumero, María Antonia Parcet, and César Ávila. Reading salt activates gustatory brain regions: fMRI evidence for semantic grounding in a novel sensory modality. *Cerebral cortex (New York, N.Y.: 1991)*, 22(11):2554–2563, November 2012. ISSN 1460-2199 1047-3211. doi: 10.1093/cercor/bhr324. Place: United States.
- [10] Pouya Bashivan, Kohitij Kar, and James J DiCarlo. Neural population control via deep image synthesis. *Science*, 364(6439):eaav9436, 2019.
- [11] Roman Beliy, Navve Wasserman, Amit Zalcher, and Michal Irani. The wisdom of a crowd of brains: A universal brain encoder, 2025. URL https://arxiv.org/abs/2406.12179.

- [12] Edwin J. Burns, Taylor Arnold, and Cindy M. Bukach. P-curving the fusiform face area: Meta-analyses support the expertise hypothesis. *Neuroscience & Biobehavioral Reviews*, 104: 209–221, 2019. ISSN 0149-7634. doi: https://doi.org/10.1016/j.neubiorev.2019.07.003. URL https://www.sciencedirect.com/science/article/pii/S014976341830798X.
- [13] Diego García Cerdas, Christina Sartzetaki, Magnus Petersen, Gemma Roig, Pascal Mettes, and Iris Groen. Brainactiv: Identifying visuo-semantic properties driving cortical selectivity using diffusion-based image manipulation. *bioRxiv*, 2024. doi: 10.1101/2024.10.29.620889. URL https://www.biorxiv.org/content/early/2024/10/31/2024.10.29.620889.
- [14] Zijiao Chen, Jiaxin Qing, Tiange Xiang, Wan Lin Yue, and Juan Helen Zhou. Seeing beyond the brain: Conditional diffusion model with sparse masked modeling for vision decoding. *arXiv* preprint arXiv:2211.06956, 1(2):4, 2022.
- [15] L. Cohen, S. Dehaene, L. Naccache, S. Lehéricy, G. Dehaene-Lambertz, M. A. Hénaff, and F. Michel. The visual word form area: spatial and temporal characterization of an initial stage of reading in normal subjects and posterior split-brain patients. *Brain: a journal of neurology*, 123 (Pt 2):291–307, February 2000. ISSN 0006-8950. doi: 10.1093/brain/123.2.291. Place: England.
- [16] Davide Cortinovis, Marius V. Peelen, and Stefania Bracci. Tool Representations in Human Visual Cortex. *Journal of Cognitive Neuroscience*, 37(3):515–531, March 2025. ISSN 0898-929X. doi: 10.1162/jocn_a_02281. URL https://doi.org/10.1162/jocn_a_02281.
- [17] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In 2009 IEEE Conference on Computer Vision and Pattern Recognition, pages 248–255, Miami, FL, June 2009. IEEE. ISBN 978-1-4244-3992-8. doi: 10. 1109/CVPR.2009.5206848. URL https://ieeexplore.ieee.org/document/5206848/.
- [18] R. Desimone and S. J. Schein. Visual properties of neurons in area V4 of the macaque: sensitivity to stimulus form. *Journal of neurophysiology*, 57(3):835–868, March 1987. ISSN 0022-3077. doi: 10.1152/jn.1987.57.3.835. Place: United States.
- [19] Adrien Doerig, Tim C Kietzmann, Emily Allen, Yihan Wu, Thomas Naselaris, Kendrick Kay, and Ian Charest. Semantic scene descriptions as an objective of human vision. *arXiv* preprint *arXiv*:2209.11737, 2022.
- [20] Paul E Downing, Yuhong Jiang, Miles Shuman, and Nancy Kanwisher. A cortical area selective for visual processing of the human body. *Science*, 293(5539):2470–2473, 2001.
- [21] Russell Epstein and Nancy Kanwisher. A cortical representation of the local visual environment. Nature, 392(6676):598–601, 1998.
- [22] Matteo Ferrante, Furkan Ozcelik, Tommaso Boccato, Rufin VanRullen, and Nicola Toschi. Brain captioning: Decoding human brain activity into images and text. *arXiv preprint arXiv:2305.11560*, 2023.
- [23] Katrin Franke, Nikos Karantzas, Konstantin Willeke, Maria Diamantaki, Kandan Ramakrishnan, Pavithra Elumalai, Kelli Restivo, Paul Fahey, Cate Nealley, Tori Shinn, et al. Dual-feature selectivity enables bidirectional coding in visual cortical neurons. *bioRxiv*, pages 2025–07, 2025.
- [24] Maëlle Freteault, Maximilien Le Clei, Loic Tetrel, Lune Bellec, and Nicolas Farrugia. Alignment of auditory artificial networks with massive individual fmri brain data leads to generalisable improvements in brain encoding and downstream tasks. *Imaging Neuroscience*, 3, 04 2025. ISSN 2837-6056. doi: 10.1162/imag\a_00525. URL https://doi.org/10.1162/imag\a_00525.
- [25] James S. Gao, Alexander G. Huth, Mark D. Lescroart, and Jack L. Gallant. Pycortex: an interactive surface visualizer for fMRI. *Frontiers in Neuroinformatics*, 9, September 2015. ISSN 1662-5196. doi: 10.3389/fninf.2015.00023. URL http://journal.frontiersin.org/Article/10.3389/fninf.2015.00023/abstract.

- [26] Alessandro T Gifford, Benjamin Lahner, Sari Saba-Sadiya, Martina G Vilas, Alex Lascelles, Aude Oliva, Kendrick Kay, Gemma Roig, and Radoslaw M Cichy. The algonauts project 2023 challenge: How the human brain makes sense of natural scenes. *arXiv* preprint *arXiv*:2301.03198, 2023.
- [27] Alessandro T. Gifford, Maya A. Jastrzębowska, Johannes J. D. Singer, and Radoslaw M. Cichy. In silico discovery of representational relationships across visual cortex, 2025. URL https://arxiv.org/abs/2411.10872.
- [28] Zijin Gu, Keith Wakefield Jamison, Meenakshi Khosla, Emily J. Allen, Yihan Wu, Ghislain St-Yves, Thomas Naselaris, Kendrick Kay, Mert R. Sabuncu, and Amy Kuceyeski. NeuroGen: Activation optimized image synthesis for discovery neuroscience. *NeuroImage*, 247:118812, February 2022. ISSN 10538119. doi: 10.1016/j.neuroimage.2021.118812. URL https://linkinghub.elsevier.com/retrieve/pii/S1053811921010831.
- [29] Umut Güçlü and Marcel A. J. van Gerven. Deep neural networks reveal a gradient in the complexity of neural representations across the ventral stream. *Journal of Neuroscience*, 35 (27):10005-10014, 2015. ISSN 0270-6474. doi: 10.1523/JNEUROSCI.5023-14.2015. URL https://www.jneurosci.org/content/35/27/10005.
- [30] Tomoyasu Horikawa, Alan S. Cowen, Dacher Keltner, and Yukiyasu Kamitani. The Neural Representation of Visually Evoked Emotion Is High-Dimensional, Categorical, and Distributed across Transmodal Brain Regions. *iScience*, 23(5), May 2020. ISSN 2589-0042. doi: 10.1016/j.isci.2020.101060. URL https://doi.org/10.1016/j.isci.2020.101060. Publisher: Elsevier.
- [31] D. H. Hubel and T. N. Wiesel. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *The Journal of physiology*, 160(1):106–154, January 1962. ISSN 0022-3751 1469-7793. doi: 10.1113/jphysiol.1962.sp006837. Place: England.
- [32] Alexander G. Huth, Wendy A. De Heer, Thomas L. Griffiths, Frédéric E. Theunissen, and Jack L. Gallant. Natural speech reveals the semantic maps that tile human cerebral cortex. *Nature*, 532 (7600):453–458, April 2016. ISSN 0028-0836, 1476-4687. doi: 10.1038/nature17637. URL https://www.nature.com/articles/nature17637.
- [33] Alexander G. Huth, Shinji Nishimoto, An T. Vu, and Jack L. Gallant. A Continuous Semantic Space Describes the Representation of Thousands of Object and Action Categories across the Human Brain. *Neuron*, 76(6):1210–1224, December 2012. ISSN 08966273. doi: 10. 1016/j.neuron.2012.10.014. URL https://linkinghub.elsevier.com/retrieve/pii/S0896627312009348.
- [34] Nidhi Jain, Aria Wang, Margaret M. Henderson, Ruogu Lin, Jacob S. Prince, Michael J. Tarr, and Leila Wehbe. Selectivity for food in human ventral visual cortex. *Communications Biology* 2023 6:1, 6:1–14, 2 2023. ISSN 2399-3642. doi: 10.1038/s42003-023-04546-2.
- [35] Nancy Kanwisher, Josh McDermott, and Marvin M Chun. The fusiform face area: a module in human extrastriate cortex specialized for face perception. *Journal of neuroscience*, 17(11): 4302–4311, 1997.
- [36] Kendrick N Kay, Thomas Naselaris, Ryan J Prenger, and Jack L Gallant. Identifying natural images from human brain activity. *Nature*, 452(7185):352–355, 2008.
- [37] Alexander J. E. Kell, Daniel L. K. Yamins, Erica N. Shook, Sam V. Norman-Haignere, and Josh H. McDermott. A Task-Optimized Neural Network Replicates Human Auditory Behavior, Predicts Brain Responses, and Reveals a Cortical Processing Hierarchy. *Neuron*, 98(3):630–644.e16, May 2018. ISSN 1097-4199 0896-6273. doi: 10.1016/j.neuron.2018.03.044. Place: United States.
- [38] Meenakshi Khosla, N. Apurva Ratan Murty, and Nancy Kanwisher. A highly selective response to food in human visual cortex revealed by hypothesis-free voxel decomposition. *Current Biology*, 32:1–13, 2022.

- [39] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980, 2014.
- [40] Ru Kong, Yan Rui Tan, Naren Wulan, Leon Qi Rong Ooi, Seyedeh-Rezvan Farahibozorg, Samuel Harrison, Janine D. Bijsterbosch, Boris C. Bernhardt, Simon Eickhoff, and B. T. Thomas Yeo. Comparison between gradients and parcellations for functional connectivity prediction of behavior. *NeuroImage*, 273:120044, June 2023. ISSN 1095-9572 1053-8119. doi: 10.1016/j.neuroimage.2023.120044. Place: United States.
- [41] Nikolaus Kriegeskorte. Deep neural networks: a new framework for modeling biological vision and brain information processing. *Annual review of vision science*, 1:417–446, 2015.
- [42] Nikolaus Kriegeskorte, Marieke Mur, Douglas A. Ruff, Roozbeh Kiani, Jerzy Bodurka, Hossein Esteky, Keiji Tanaka, and Peter A. Bandettini. Matching categorical object representations in inferior temporal cortex of man and monkey. *Neuron*, 60(6):1126–1141, Dec 2008. ISSN 0896-6273. doi: 10.1016/j.neuron.2008.10.043. URL https://doi.org/10.1016/j.neuron.2008.10.043.
- [43] Yuanning Li, Huzheng Yang, and Shi Gu. Upgrading voxel-wise encoding model via integrated integration over features and brain networks. *bioRxiv*, 2022. doi: 10.1101/2022.11.06.515387. URL https://www.biorxiv.org/content/early/2022/11/07/2022.11.06.515387.
- [44] Yulong Liu, Yongqiang Ma, Wei Zhou, Guibo Zhu, and Nanning Zheng. Brainclip: Bridging brain and visual-linguistic representation via clip for generic natural visual stimulus decoding from fmri. *arXiv preprint arXiv:2302.12971*, 2023.
- [45] Andrew Luo, Margaret Marie Henderson, Michael J. Tarr, and Leila Wehbe. Brainscuba: Fine-grained natural language captions of visual cortex selectivity. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=mQYHXUUTkU.
- [46] Andrew F. Luo, Margaret M. Henderson, Leila Wehbe, and Michael J. Tarr. Brain Diffusion for Visual Exploration: Cortical Discovery using Large Scale Generative Models, November 2023. URL http://arxiv.org/abs/2306.03089. arXiv:2306.03089 [cs].
- [47] Takuya Matsuyama, Shinji Nishimoto, and Yu Takagi. LaVCa: LLM-assisted Visual Cortex Captioning, February 2025. URL http://arxiv.org/abs/2502.13606. arXiv:2502.13606 [q-bio].
- [48] Gregory McCarthy, Aina Puce, John C Gore, and Truett Allison. Face-specific processing in the human fusiform gyrus. *Journal of cognitive neuroscience*, 9(5):605–610, 1997.
- [49] Marieke Mur, Douglas A. Ruff, Jerzy Bodurka, Peter De Weerd, Peter A. Bandettini, and Nikolaus Kriegeskorte. Categorical, yet graded – single-image activation profiles of human category-selective cortical regions. *Journal of Neuroscience*, 32(25):8649–8662, 2012. ISSN 0270-6474. doi: 10.1523/JNEUROSCI.2334-11.2012. URL https://www.jneurosci.org/ content/32/25/8649.
- [50] Till Nierhaus, Sara Wesolek, Daniel Pach, Claudia M. Witt, Felix Blankenburg, and Timo T. Schmidt. Content Representation of Tactile Mental Imagery in Primary Somatosensory Cortex. eNeuro, 10(6):ENEURO.0408–22.2023, June 2023. ISSN 2373-2822. doi: 10.1523/ENEURO.0408-22.2023. Place: United States.
- [51] Toshiki Okumura, Ikuhiro Kida, Atsushi Yokoi, Tomoya Nakai, Shinji Nishimoto, Kazushige Touhara, and Masako Okamoto. Semantic context-dependent neural representations of odors in the human piriform cortex revealed by 7T MRI. *Human brain mapping*, 45(6):e26681, April 2024. ISSN 1097-0193 1065-9471. doi: 10.1002/hbm.26681. Place: United States.
- [52] OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum, Irwan Bello, Jake Berdine, Gabriel Bernadett-Shapiro, Christopher Berner, Lenny Bogdonoff, Oleg Boiko, Madelaine Boyd, Anna-Luisa Brakman,

Greg Brockman, Tim Brooks, Miles Brundage, Kevin Button, Trevor Cai, Rosie Campbell, Andrew Cann, Brittany Carey, Chelsea Carlson, Rory Carmichael, Brooke Chan, Che Chang, Fotis Chantzis, Derek Chen, Sully Chen, Ruby Chen, Jason Chen, Mark Chen, Ben Chess, Chester Cho, Casey Chu, Hyung Won Chung, Dave Cummings, Jeremiah Currier, Yunxing Dai, Cory Decareaux, Thomas Degry, Noah Deutsch, Damien Deville, Arka Dhar, David Dohan, Steve Dowling, Sheila Dunning, Adrien Ecoffet, Atty Eleti, Tyna Eloundou, David Farhi, Liam Fedus, Niko Felix, Simón Posada Fishman, Juston Forte, Isabella Fulford, Leo Gao, Elie Georges, Christian Gibson, Vik Goel, Tarun Gogineni, Gabriel Goh, Rapha Gontijo-Lopes, Jonathan Gordon, Morgan Grafstein, Scott Gray, Ryan Greene, Joshua Gross, Shixiang Shane Gu, Yufei Guo, Chris Hallacy, Jesse Han, Jeff Harris, Yuchen He, Mike Heaton, Johannes Heidecke, Chris Hesse, Alan Hickey, Wade Hickey, Peter Hoeschele, Brandon Houghton, Kenny Hsu, Shengli Hu, Xin Hu, Joost Huizinga, Shantanu Jain, Shawn Jain, Joanne Jang, Angela Jiang, Roger Jiang, Haozhun Jin, Denny Jin, Shino Jomoto, Billie Jonn, Heewoo Jun, Tomer Kaftan, Łukasz Kaiser, Ali Kamali, Ingmar Kanitscheider, Nitish Shirish Keskar, Tabarak Khan, Logan Kilpatrick, Jong Wook Kim, Christina Kim, Yongjik Kim, Jan Hendrik Kirchner, Jamie Kiros, Matt Knight, Daniel Kokotajlo, Łukasz Kondraciuk, Andrew Kondrich, Aris Konstantinidis, Kyle Kosic, Gretchen Krueger, Vishal Kuo, Michael Lampe, Ikai Lan, Teddy Lee, Jan Leike, Jade Leung, Daniel Levy, Chak Ming Li, Rachel Lim, Molly Lin, Stephanie Lin, Mateusz Litwin, Theresa Lopez, Ryan Lowe, Patricia Lue, Anna Makanju, Kim Malfacini, Sam Manning, Todor Markov, Yaniv Markovski, Bianca Martin, Katie Mayer, Andrew Mayne, Bob McGrew, Scott Mayer McKinney, Christine McLeavey, Paul McMillan, Jake McNeil, David Medina, Aalok Mehta, Jacob Menick, Luke Metz, Andrey Mishchenko, Pamela Mishkin, Vinnie Monaco, Evan Morikawa, Daniel Mossing, Tong Mu, Mira Murati, Oleg Murk, David Mély, Ashvin Nair, Reiichiro Nakano, Rajeev Nayak, Arvind Neelakantan, Richard Ngo, Hyeonwoo Noh, Long Ouyang, Cullen O'Keefe, Jakub Pachocki, Alex Paino, Joe Palermo, Ashley Pantuliano, Giambattista Parascandolo, Joel Parish, Emy Parparita, Alex Passos, Mikhail Pavlov, Andrew Peng, Adam Perelman, Filipe de Avila Belbute Peres, Michael Petrov, Henrique Ponde de Oliveira Pinto, Michael, Pokorny, Michelle Pokrass, Vitchyr H. Pong, Tolly Powell, Alethea Power, Boris Power, Elizabeth Proehl, Raul Puri, Alec Radford, Jack Rae, Aditya Ramesh, Cameron Raymond, Francis Real, Kendra Rimbach, Carl Ross, Bob Rotsted, Henri Roussez, Nick Ryder, Mario Saltarelli, Ted Sanders, Shibani Santurkar, Girish Sastry, Heather Schmidt, David Schnurr, John Schulman, Daniel Selsam, Kyla Sheppard, Toki Sherbakov, Jessica Shieh, Sarah Shoker, Pranav Shyam, Szymon Sidor, Eric Sigler, Maddie Simens, Jordan Sitkin, Katarina Slama, Ian Sohl, Benjamin Sokolowsky, Yang Song, Natalie Staudacher, Felipe Petroski Such, Natalie Summers, Ilya Sutskever, Jie Tang, Nikolas Tezak, Madeleine B. Thompson, Phil Tillet, Amin Tootoonchian, Elizabeth Tseng, Preston Tuggle, Nick Turley, Jerry Tworek, Juan Felipe Cerón Uribe, Andrea Vallone, Arun Vijayvergiya, Chelsea Voss, Carroll Wainwright, Justin Jay Wang, Alvin Wang, Ben Wang, Jonathan Ward, Jason Wei, CJ Weinmann, Akila Welihinda, Peter Welinder, Jiayi Weng, Lilian Weng, Matt Wiethoff, Dave Willner, Clemens Winter, Samuel Wolrich, Hannah Wong, Lauren Workman, Sherwin Wu, Jeff Wu, Michael Wu, Kai Xiao, Tao Xu, Sarah Yoo, Kevin Yu, Qiming Yuan, Wojciech Zaremba, Rowan Zellers, Chong Zhang, Marvin Zhang, Shengjia Zhao, Tianhao Zheng, Juntang Zhuang, William Zhuk, and Barret Zoph. Gpt-4 technical report, 2024. URL https://arxiv.org/abs/2303.08774.

- [53] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, Mahmoud Assran, Nicolas Ballas, Wojciech Galuba, Russell Howes, Po-Yao Huang, Shang-Wen Li, Ishan Misra, Michael Rabbat, Vasu Sharma, Gabriel Synnaeve, Hu Xu, Hervé Jegou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. DINOv2: Learning Robust Visual Features without Supervision, February 2024. URL http://arxiv.org/abs/2304.07193. arXiv:2304.07193 [cs].
- [54] Marius V. Peelen and Paul E. Downing. Selectivity for the human body in the fusiform gyrus. *Journal of neurophysiology*, 93(1):603–608, January 2005. ISSN 0022-3077. doi: 10.1152/jn.00513.2004. Place: United States.
- [55] Francisco Pereira, Bin Lou, Brianna Pritchett, Samuel Ritter, Samuel J. Gershman, Nancy Kanwisher, Matthew Botvinick, and Evelina Fedorenko. Toward a universal decoder of linguistic meaning from brain activation. *Nature Communications*, 9(1):963, March 2018.

- ISSN 2041-1723. doi: 10.1038/s41467-018-03068-4. URL https://doi.org/10.1038/s41467-018-03068-4.
- [56] Paweł A Pierzchlewicz, Konstantin F Willeke, Arne F Nix, Pavithra Elumalai, Kelli Restivo, Tori Shinn, Cate Nealley, Gabrielle Rodriguez, Saumil Patel, Katrin Franke, et al. Energy guided diffusion for generating neurally exciting images. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*, pages 32574–32601, 2023.
- [57] Xuelin Qian, Yun Wang, Jingyang Huo, Xinwei Sun, Yanwei Fu, and Jianfeng Feng. fMRI-PTE: A large-scale fMRI pretrained transformer encoder for multi-subject brain activity decoding, 2024. URL https://openreview.net/forum?id=BZkKMQ25Z7.
- [58] N Apurva Ratan Murty, Pouya Bashivan, Alex Abate, James J DiCarlo, and Nancy Kanwisher. Computational models of category-selective brain regions enable high-throughput tests of selectivity. *Nature communications*, 12(1):5540, 2021.
- [59] Alexander Schaefer, Ru Kong, Evan M Gordon, Timothy O Laumann, Xi-Nian Zuo, Avram J Holmes, Simon B Eickhoff, and B T Thomas Yeo. Local-Global Parcellation of the Human Cerebral Cortex from Intrinsic Functional Connectivity MRI. Cerebral Cortex, 28(9):3095–3114, September 2018. ISSN 1047-3211, 1460-2199. doi: 10.1093/cercor/bhx179. URL https://academic.oup.com/cercor/article/28/9/3095/3978804.
- [60] Martin Schrimpf, Jonas Kubilius, Ha Hong, Najib J Majaj, Rishi Rajalingham, Elias B Issa, Kohitij Kar, Pouya Bashivan, Jonathan Prescott-Roy, Franziska Geiger, et al. Brain-score: Which artificial neural network for object recognition is most brain-like? *BioRxiv*, page 407007, 2018.
- [61] Rebecca F. Schwarzlose, Chris I. Baker, and Nancy Kanwisher. Separate face and body selectivity on the fusiform gyrus. *The Journal of neuroscience: the official journal of the Society for Neuroscience*, 25(47):11055–11059, November 2005. ISSN 1529-2401 0270-6474. doi: 10.1523/JNEUROSCI.2621-05.2005. Place: United States.
- [62] Paul S Scotti, Mihir Tripathy, Cesar Kadir Torrico Villanueva, Reese Kneeland, Tong Chen, Ashutosh Narang, Charan Santhirasegaran, Jonathan Xu, Thomas Naselaris, Kenneth A Norman, et al. Mindeye2: Shared-subject models enable fmri-to-image with 1 hour of data. arXiv preprint arXiv:2403.11207, 2024.
- [63] Justine Sergent, Shinsuke Ohta, and Brenna Macdonald. Functional Neuroanatomy of Face and Object Processing: A Positron Emission Tomography Study. *Brain*, 115(1):15–36, February 1992. ISSN 0006-8950. doi: 10.1093/brain/115.1.15. URL https://doi.org/10.1093/brain/115.1.15. _eprint: https://academic.oup.com/brain/article-pdf/115/1/15/836448/115-1-15.pdf.
- [64] Irina Simanova, Peter Hagoort, Robert Oostenveld, and Marcel A. J. van Gerven. Modality-Independent Decoding of Semantic Information from the Human Brain. *Cerebral Cortex*, 24 (2):426–434, February 2014. ISSN 1047-3211. doi: 10.1093/cercor/bhs324. URL https://doi.org/10.1093/cercor/bhs324.
- [65] W. Kyle Simmons, Alex Martin, and Lawrence W. Barsalou. Pictures of appetizing foods activate gustatory cortices for taste and reward. *Cerebral cortex (New York, N.Y.: 1991)*, 15(10): 1602–1608, October 2005. ISSN 1047-3211. doi: 10.1093/cercor/bhi038. Place: United States.
- [66] Yu Takagi and Shinji Nishimoto. High-resolution image reconstruction with latent diffusion models from human brain activity. bioRxiv, pages 2022–11, 2022.
- [67] Jerry Tang, Meng Du, Vy Vo, Vasudev Lal, and Alexander Huth. Brain encoding models based on multimodal transformers can transfer across language and vision. *Advances in neural information processing systems*, 36:29654–29666, 12 2023.
- [68] Polina Turishcheva, Paul G Fahey, Michaela Vystrčilová, Laura Hansel, Rachel Froebe, Kayla Ponder, Yongrong Qiu, Konstantin F Willeke, Mohammad Bashiri, Eric Wang, et al. The dynamic sensorium competition for predicting large-scale mouse visual cortex activity from videos. *ArXiv*, pages arXiv–2305, 2024.

- [69] L. N. van der Laan, D. T. D. de Ridder, M. A. Viergever, and P. A. M. Smeets. The first taste is always with the eyes: a meta-analysis on the neural correlates of processing visual food cues. *NeuroImage*, 55(1):296–303, March 2011. ISSN 1095-9572 1053-8119. doi: 10.1016/j.neuroimage.2010.11.055. Place: United States.
- [70] A. van der Schaaf and J.H. van Hateren. Modelling the power spectra of natural images: Statistics and information. Vision Research, 36(17):2759–2770, 1996. ISSN 0042-6989. doi: https://doi.org/10.1016/0042-6989(96)00002-8. URL https://www.sciencedirect.com/science/article/pii/0042698996000028.
- [71] Edgar Y. Walker, R. James Cotton, Wei Ji Ma, and Andreas S. Tolias. A neural basis of probabilistic computation in visual cortex. *Nature neuroscience*, 23(1):122–129, January 2020. ISSN 1546-1726 1097-6256. doi: 10.1038/s41593-019-0554-5. Place: United States.
- [72] Daniel L. K. Yamins and James J. DiCarlo. Using goal-driven deep learning models to understand sensory cortex. *Nature Neuroscience*, 19(3):356–365, Mar 2016.
- [73] Huzheng Yang, James Gee, and Jianbo Shi. Brain decodes deep nets. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 23030–23040, 2024.
- [74] Jacob Yeung, Andrew F Luo, Gabriel Sarch, Margaret M Henderson, Deva Ramanan, and Michael J Tarr. Neural representations of dynamic visual stimuli. *arXiv preprint* arXiv:2406.02659, 2024.

A Technical Appendices and Supplementary Material

A.1 Social impacts

In this work, we present a method for in silico mapping of semantic selectivity across cortical parcels in the whole brain. For cognitive neuroscience, our approach deepens the field's understanding of visual processing well beyond classical visual cortex. Clinically, deviations from normative whole-brain semantic maps may act as early biomarkers for neurological and neurodegenerative disorders. Furthermore, these detailed semantic atlases lay foundational groundwork for next-generation brain—computer interfaces, where fine-grained category-level decoding could translate a user's intended concepts—rather than low-level motor signals—into control commands for communication or assistive devices.

A.2 Reproducing selectivity for known areas

We show superstimuli from parcels overlapping with body-, place-, and word-selective areas in order to show that our pipeline can reproduce the known selectivity.



A.3 Parcel selection process

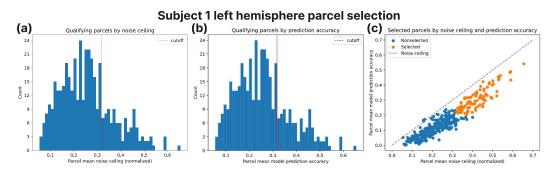


Figure S1: **Selecting parcels outside visual cortex for experimentation.** Parcels outside the visual area that have sufficiently high mean (a) noise ceiling and (b) model prediction accuracy are chosen for further experimentation. (c) Selected parcels by mean noise ceiling and model prediction accuracy.

A.4 Reproducing known visual hierarchy properties

We optimized 32 superstimuli (16 per hemisphere) that maximally activate parcels in V1, V2, and V4 using the BrainDIVE framework. We additionally include results on FFA as a representative IT-cortex ROI.

Generated superstimuli for early visual areas reproduce the classic fine-to-coarse hierarchy: V1 stimuli appear as cluttered scenes filled with dense, repetitive texture; V2 images add composite color patches and rudimentary objects; V4 stimuli reveal smoother, recognizable object forms; and FFA images almost exclusively depict close-up faces, often with clear emotional expressions.

To quantify hierarchical properties, we examined whether the spatial-frequency content of our stimuli mirrors classical physiological findings (e.g., [31, 18]). We computed the radial average power spectrum for each image [70] and calculated the proportion of spectral power above 10%, 20%, and 30% of the maximum spatial frequency. At every threshold the high-frequency energy ratio decreases monotonically from V1 > V2 > V4 > FFA, indicating that the images that best drive higher-level areas (FFA) contain proportionally less fine-scale texture and relatively more coarse, low-frequency structure.

Table 6: High-frequency energy ratio across ROIs at different thresholds, subject 1.

Threshold	ROI	High-frequency energy ratio
0.1	V1	0.006547
	V2	0.004160
	V4	0.002500
	FFA	0.001347
0.2	V1	0.001838
	V2	0.000944
	V4	0.000699
	FFA	0.000299
0.3	V1	0.000762
	V2	0.000339
	V4	0.000285
	FFA	0.000104

Below we include superstimuli optimized for each of the above ROIs to visualize the image frequency differences.

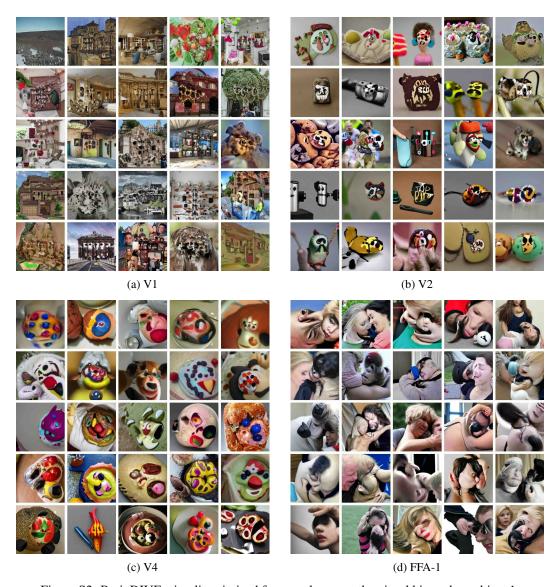


Figure S2: BrainDIVE stimuli optimized for parcels across the visual hierarchy, subject 1.

A.5 Analyzing ROI queries to map parcel connectivity

Taking advantage of our transformer-based approach, we map the connections between visual area parcels by examining the representational similarity of learned ROI queries, each of which corresponds to a single parcel. For both functional connectivity and ROI query similarity, we first take the full 1000 parcel x 1000 parcel matrix, then mask out the labeled area, and finally average across the top three parcels that overlap with each labeled area. (Thus, not every entry on the diagonal is 1.) High cosine similarity between ROI queries (queries averaged across ensemble models) suggests that two parcels are highly connected, as both attend to similar content in an image.

This similarity matrix for the visual areas closely replicates the functional connectivity matrix from the Schaefer parcellation (derived from resting-state correlated responses) [40]; the Pearson correlation between the upper triangle of the two matrices is ≈ 0.5 .

Below we report the functional connectivity and ROI query cosine similarity matrices for subject 1.

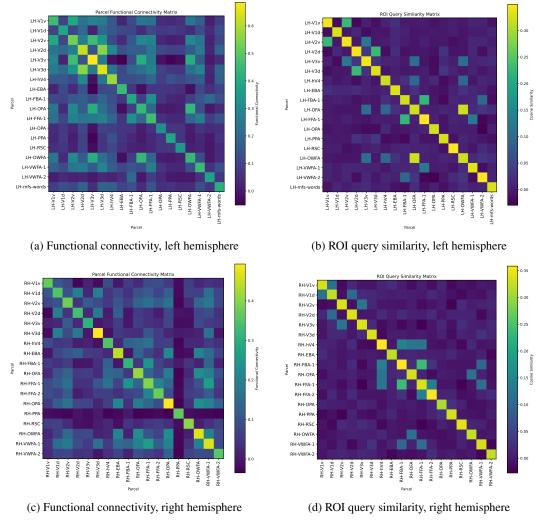


Figure S3: Functional connectivity matrix vs ROI query similarity, subject 1.

A.6 Attention maps

While one can easily make guesses about parcel selectivity simply by looking at the prominent feature of the image, we can examine the attention scores from cross-attention (Figure 1b) to interpret the selectivity of any parcel. The attention maps are visualized in Figure S4 for the images from Figure 5.



Figure S4: **Attention maps for aTL-faces and mfs-words.** Yellow areas represent the highest attention weights and purple represents the lowest. Held-out NSD images are selected based on ground-truth parcel activation, but attention maps are from the encoder since no ground-truth attention maps exist. See Figure 5 for the original images.

As expected, the areas with highest attention weights for the aTL-faces parcel primarily overlap with faces, while the attention weights for the mfs-words parcel overlap with words (when present). Compared to previous work using a similar encoder model [3], our attention maps appear to follow the expected stimuli considerably less closely, possibly because our parcel boundaries were determined a priori and don't align well with the boundaries of the actual labeled area.

A.7 Parcel locations and BrainDIVE/NSD images accompanying Figure 1

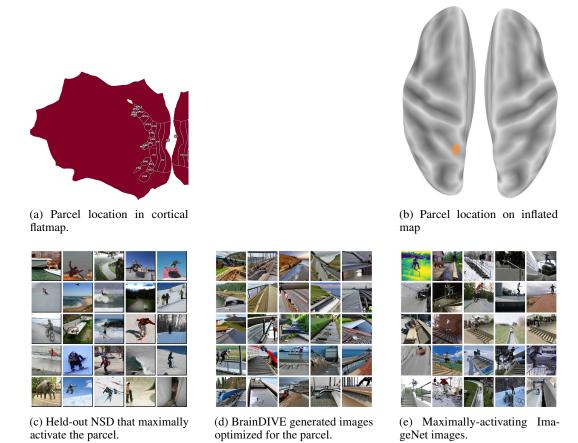
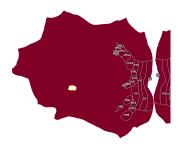


Figure S5: Skateboarding parcel location and full collages, subject 1.



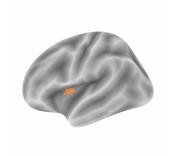
(a) Parcel location in cortical flatmap.



(c) Held-out NSD that maximally activate the parcel.



(d) BrainDIVE generated images optimized for the parcel.



(b) Parcel location on inflated map



(e) Maximally-activating ImageNet images.

Figure S6: Child eating parcel location and full collages, subject 1.

A.8 Labeling top subject-specific parcels

For each of the parcels shown in Figure 6, we show the top ImageNet images and the caption generated by GPT-40 (see Appendix A.18 for prompt details).

Parcel	ImageNet top 4	Caption	Parcel	ImageNet top 4	Caption
S1 LH 211		Sports, Tennis, Soccer	S5 RH 476		Sports, Team sports, Outdoor activities
S1 RH 373		Soccer, Tennis, Sports	S5 LH 239		Extreme sports, Skateboarding, BMX biking
S1 LH 238		Sports, Tennis, Soccer	S5 LH 367		Interior Design, Home Decor, Kitchen
S1 RH 456		Babies, Children, Family	S5 RH 391	1	Soccer, Ice hockey, Tennis
S1 LH 496		Family, Parenting, Children	S5 LH 167		Kitchen, Appliances, Cabinets
S2 LH 480		Food storage, Office supplies, Kitchen appliances	S7 RH 474		Tennis, Baseball, Sports
S2 RH 391	X X	Gymnastics, Dance, Tennis	S7 RH 477		Tennis, Baseball, Sports
S2 LH 500	多人	Tennis, Gymnastics, Baseball	S7 LH 496		Sports, Baseball, Tennis
S2 LH 496		Social Gatherings, Education and Learning, Sports and Recreation	S7 LH 497		Baseball, Pitching, Batting
S2 LH 365		Interior Design, Living Spaces, Furniture	S7 RH 317		Baseball, Pitching, Batting

Table 7: Maximally parcel-activating images (predicted) for unlabeled parcels.

A.9 Full collages for cross-subject tools parcel

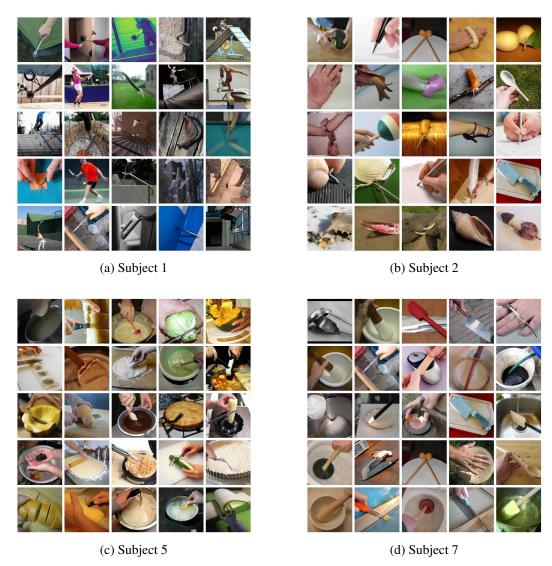


Figure S7: Top 25 maximally parcel-activating ImageNet images (based on encoder predictions) for the cross-subject parcel in Figure 7.

A.10 Statistical tests for subject-specific parcel evaluation

Here, we explain the details of the statistical tests performed in 4.4.

The underlying assumption in cognitive neuroscience is that for some parcels, there exists some category of visual stimuli that maximally activates the parcel.

Our goal is to assess whether the semantic content of the images our encoder identifies as highly activating aligns with the semantic content of NSD test images that actually elicit high ground-truth activations in that parcel. We ask whether a single "concept vector" derived from our encoder for a given parcel predicts parcel-activation rank order in the NSD test set more accurately than the baseline model.

We computed two concept vectors per parcel: one based on the top 32 maximally-activating ImageNet or BrainDIVE images predicted by our encoder, and one based on the top 32 maximally-activating NSD training images. We embedded each set of the images into CLIP space and then averaged the CLIP vectors to obtain a parcel-specific "concept vector" that reflects predicted parcel categorical

selectivity. We then embedded each NSD test image into CLIP space, and ranked images by cosine similarity to the concept vector (encoder-derived or baseline). Since we expect that higher semantic similarity to the "concept vector" will correspond to higher parcel activation, we quantified alignment with the empirical activation ordering using Spearman's rank correlation between concept-vector similarity rank and ground-truth fMRI activity rank.

In Table 2, for each parcel, we compared the rank correlation coefficient with a null distribution, bootstrapped with 10,000 random rankings of the NSD test images. We separate the results into most activating stimuli predicted by the model from BrainDIVE and ImageNet, as well as stimuli from the NSD training set (baseline).

A.11 Statistical tests for cross-subject parcel evaluation

For a given parcel and dataset, we generated a selectivity hypothesis by averaging the CLIP vectors corresponding to the top 32 images according to the brain encoder, then used that CLIP vector to rank the NSD training images using cosine similarity. We calculate Spearman rank correlation between that ranking and the ranking determined by ground truth parcel activation, and compare this to a null distribution of correlation coefficients, which consists of 10,000 draws of random rankings.

Note that there are a fewer number of total parcels measured here since each subject has different parcels selected for experimentation (see Section 3.3), so cross-subject selectivity can only be tested for parcels selected for experimentation with every subject.

A.12 Cosine similarity analysis of labels for cross-subject parcels

To demonstrate cross-subject consistency in semantic selectivity, we also compared directly the concept vectors generated by the encoder. Concept vectors generated for the same parcel, when analyzed across subjects, tend to be more similar (0.784 \pm 0.069 for BrainDIVE) than the parcels coming from around the target parcel (0.710 \pm 0.115 for BrainDIVE). Absolute cosine similarities are high because these parcels are selective for high-level visual concepts that share similar features, but our method can adjudicate between them using targeted stimulus generation.

A.13 Head-to-head encoder model versus NSD train baseline

Table 8 expands upon the results shown in table 2, reporting the fraction of parcels for which the encoder-selected stimuli outperforms the NSD train baseline in Spearman's rank correlation coefficient, evaluated within the same subject.

Table 8: Head to head comparison between our model and NSD train baseline: number of parcels where model rank correlation coefficient beats NSD train, evaluated within the same subject.

Comparison type	S 1	S2	S5	S7
Our encoder w/ImageNet > NSD train:	93/181	94/192 $123/192$	89/175	62/196
Our encoder w/BrainDIVE > NSD train:	95/181		100/175	76/196

Table 9 expands upon results in Table 4, reporting the number of parcels for which the rank correlation coefficient from the specified model is greater than that from the baseline model, when three subjects are used to predict a held-out subject.

Table 9: Head to head comparison between our model and NSD train baseline: number of parcels where model rank correlation coefficient beats NSD train, evaluated with all subjects used to predict one held out.

Comparison type	S 1	S2	S5	S7
Our encoder w/ImageNet > NSD train: Our encoder w/BrainDIVE > NSD train:		$\frac{35/49}{24/49}$	$\frac{36/49}{30/49}$	27/49 $22/49$

A.14 Choosing parcels shared across subjects for future fMRI experimentation

In Figure S8, we report similar results as in Figure 6, except for top parcels shared across subjects.

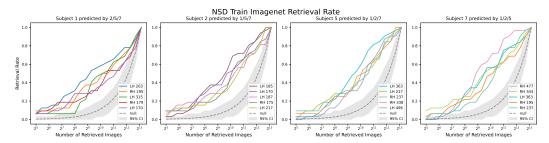


Figure S8: See caption from Figure 6 for description. Each parcel shown is shared across subjects and the encoder train on three subjects is used to predict the selectivity of that same parcel on a held-out subject.

A.15 Interpreting statistical test results by comparing to known visual areas

To quantify what rank correlation magnitude constitutes meaningful selectivity, we report the rank correlation the model achieves on known visual areas as a benchmark. Parcels with known high-level selectivity (EBA, FFA, FBA) tend to exhibit very high correlations, while parcels with known lower-level selectivity (V1–V4) tend to exhibit low correlations. The average Spearman's ρ from our concept vectors are in the range of areas like PPA and RSC.

Using CLIP as our metric allows us to capture high-level features that can explain semantic selectivity. Given that fMRI studies tend to rely on experimenter-curated topics, the CLIP space is a move toward a data-driven approach, but ultimately any metric can be substituted.

Figure S9: Spearman's ρ , actual vs. predicted activation ordering for subject 1. Top: high-level areas. Bottom: low-level areas.

Region	ImageNet	BrainDIVE
Unlabeled parcel mean (ours)	0.168 ± 0.106	0.163 ± 0.121
EBA	0.432	0.501
FFA-1	0.217	0.243
FFA-2	0.373	0.412
FBA-2	0.373	0.401
lateral	0.323	0.339
PPA	0.166	0.128
RSC	0.170	0.162
aTL-faces	0.120	0.138
mTL-words	0.164	0.199

(b) Low-level areas

Region	ImageNet	BrainDIVE
Unlabeled parcel mean (ours)	0.168 ± 0.106	0.163 ± 0.121
V1d	0.010	0.027
V1v	0.024	0.036
V2d	-0.032	-0.035
V2v	-0.030	-0.002
V3d	-0.004	-0.018
V3v	0.007	0.009
hV4	0.047	0.067
early	-0.025	-0.014

Note that since we're using the Schaefer-1000 parcellation, for each visual area, we average the rank correlation value across the top 3 Schaefer parcels with greatest overlap with that visual area.

A.16 Varying top images used for statistical test

To confirm that the superior performance of ImageNet and BrainDIVE superstimuli stems from a broader diversity of concepts present—and not sampling near peak activation—we varied the top-K image threshold used to define the concept vector. We report K=16,8,4,2,1 versions for tables 2, 3, 4, and 5.

Table 10: Fraction of parcels whose model-derived label predicts parcel activation rankings significantly better than chance (p < 0.05, FDR corrected), varying number of top images for label generation.

Model type	S 1	S2	S5	S7
k=64				
NSD train	134/181	156/192	134/175	154/196
Our encoder w/ImageNet	139/181	166/192	134/175	149/196
Our encoder w/BrainDIVE	128/181	168/192	140/175	152/196
k=32				
NSD train	150/181	163/192	136/175	155/196
Our encoder w/ImageNet	139/181	167/192	130/175	150/196
Our encoder w/BrainDIVE	135/181	170/192	139/175	156/196
k=16				
NSD train	154/181	169/192	137/175	157/196
Our encoder w/ImageNet	138/181	166/192	134/175	151/196
Our encoder w/BrainDIVE	139/181	171/192	136/175	152/196
k=8				
NSD train	144/181	160/192	131/175	149/196
Our encoder w/ImageNet	138/181	162/192	134/175	153/196
Our encoder w/BrainDIVE	135/181	168/192	136/175	153/196
k=4				
NSD train	134/181	146/192	122/175	140/196
Our encoder w/ImageNet	138/181	164/192	129/175	148/196
Our encoder w/BrainDIVE	140/181	166/192	142/175	151/196
k=2				
NSD train	131/181	125/192	109/175	110/196
Our encoder w/ImageNet	134/181	154/192	122/175	143/196
Our encoder w/BrainDIVE	141/181	171/192	136/175	136/196
k=1				
NSD train	114/181	120/192	92/175	86/196
Our encoder w/ImageNet	133/181	145/192	113/175	130/196
Our encoder w/BrainDIVE	137/181	162/192	119/175	131/196

Table 11: Average parcel-level correlations (mean \pm std) between model-derived labels and parcel activation rankings, varying number of top images for label generation.

Model type	S 1	S2	S5	S7
k=64				
NSD train	0.145 ± 0.098	0.153 ± 0.088	0.141 ± 0.096	0.142 ± 0.085
Our encoder w/ImageNet	0.167 ± 0.106	0.163 ± 0.082	0.143 ± 0.092	0.133 ± 0.075
Our encoder w/BrainDIVE	0.154 ± 0.119	0.187 ± 0.099	0.153 ± 0.095	0.131 ± 0.084
k=32				
NSD train	0.162 ± 0.098	0.164 ± 0.091	0.150 ± 0.094	0.148 ± 0.086
Our encoder w/ImageNet	0.168 ± 0.106	0.163 ± 0.082	0.142 ± 0.092	0.133 ± 0.075
Our encoder w/BrainDIVE	0.163 ± 0.121	0.190 ± 0.099	0.154 ± 0.094	0.133 ± 0.083
k=16		·	·	
NSD train	0.167 ± 0.096	0.170 ± 0.093	0.149 ± 0.096	0.148 ± 0.087
Our encoder w/ImageNet	0.169 ± 0.106	0.162 ± 0.082	0.142 ± 0.091	0.131 ± 0.075
Our encoder w/BrainDIVE	0.168 ± 0.120	0.190 ± 0.099	0.155 ± 0.092	0.134 ± 0.082
k=8				
NSD train	0.160 ± 0.103	0.164 ± 0.094	0.140 ± 0.100	0.139 ± 0.091
Our encoder w/ImageNet	0.168 ± 0.107	0.159 ± 0.081	0.142 ± 0.092	0.128 ± 0.073
Our encoder w/BrainDIVE	0.167 ± 0.121	0.188 ± 0.099	0.155 ± 0.092	0.133 ± 0.082
k=4				
NSD train	0.151 ± 0.112	0.147 ± 0.099	0.128 ± 0.097	0.124 ± 0.091
Our encoder w/ImageNet	0.165 ± 0.108	0.155 ± 0.082	0.140 ± 0.091	0.124 ± 0.073
Our encoder w/BrainDIVE	0.165 ± 0.120	0.184 ± 0.098	0.151 ± 0.092	0.131 ± 0.080
k=2				
NSD train	0.134 ± 0.114	0.127 ± 0.106	0.110 ± 0.104	0.106 ± 0.092
Our encoder w/ImageNet	0.158 ± 0.107	0.147 ± 0.082	0.135 ± 0.090	0.118 ± 0.071
Our encoder w/BrainDIVE	0.160 ± 0.117	0.176 ± 0.094	0.149 ± 0.090	0.127 ± 0.080
k=1				
NSD train	0.121 ± 0.120	0.109 ± 0.100	0.088 ± 0.105	0.087 ± 0.090
Our encoder w/ImageNet	0.148 ± 0.101	0.136 ± 0.084	0.125 ± 0.091	0.112 ± 0.075
Our encoder w/BrainDIVE	0.148 ± 0.115	0.165 ± 0.093	0.140 ± 0.092	0.119 ± 0.082

Table 12: Fraction of parcels shared across subjects whose model-derived label predicts parcel activation rankings significantly better than chance (p < 0.05, FDR corrected), varying number of top images for label generation.

Model type	S1	S2	S5	S7
k=64				
NSD train	37/49	44/49	40/49	40/49
Our encoder w/ImageNet	38/49	45/49	42/49	39/49
Our encoder w/BrainDIVE	34/49	40/49	37/49	38/49
k=32				
NSD train	37/49	44/49	40/49	40/49
Our encoder w/ImageNet	38/49	45/49	42/49	39/49
Our encoder w/BrainDIVE	37/49	40/49	38/49	38/49
k=16				
NSD train	37/49	42/49	40/49	40/49
Our encoder w/ImageNet	38/49	46/49	42/49	39/49
Our encoder w/BrainDIVE	37/49	40/49	38/49	38/49
k=8				
NSD train	36/49	40/49	41/49	40/49
Our encoder w/ImageNet	37/49	45/49	43/49	39/49
Our encoder w/BrainDIVE	37/49	40/49	38/49	39/49
k=4				
NSD train	40/49	41/49	40/49	38/49
Our encoder w/ImageNet	38/49	44/49	42/49	40/49
Our encoder w/BrainDIVE	37/49	40/49	39/49	41/49
k=2				
NSD train	40/49	40/49	40/49	40/49
Our encoder w/ImageNet	38/49	45/49	42/49	38/49
Our encoder w/BrainDIVE	37/49	43/49	39/49	40/49
k=1				
NSD train	37/49	35/49	38/49	40/49
Our encoder w/ImageNet	37/49	44/49	42/49	37/49
Our encoder w/BrainDIVE	39/49	40/49	40/49	40/49

Table 13: Average parcel-level correlations (mean \pm std) between model-derived labels and parcel activation rankings, varying number of top images for label generation.

Model type	S 1	S2	S5	S7
k=64				
NSD train	0.066 ± 0.068	0.092 ± 0.065	0.091 ± 0.073	0.090 ± 0.069
Our encoder w/ImageNet	0.072 ± 0.072	0.105 ± 0.074	0.105 ± 0.078	0.093 ± 0.076
Our encoder w/BrainDIVE	0.075 ± 0.076	0.102 ± 0.087	0.102 ± 0.090	0.087 ± 0.080
k=32				
NSD train	0.068 ± 0.068	0.094 ± 0.065	0.090 ± 0.074	0.091 ± 0.070
Our encoder w/ImageNet	0.072 ± 0.071	0.105 ± 0.073	0.105 ± 0.078	0.093 ± 0.075
Our encoder w/BrainDIVE	0.078 ± 0.077	0.103 ± 0.087	0.106 ± 0.091	0.089 ± 0.079
k=16				
NSD train	0.068 ± 0.067	0.093 ± 0.066	0.089 ± 0.072	0.090 ± 0.069
Our encoder w/ImageNet	0.073 ± 0.070	0.105 ± 0.073	0.107 ± 0.079	0.091 ± 0.076
Our encoder w/BrainDIVE	0.081 ± 0.078	0.103 ± 0.087	0.106 ± 0.093	0.090 ± 0.080
k=8				
NSD train	0.070 ± 0.067	0.095 ± 0.064	0.088 ± 0.072	0.088 ± 0.067
Our encoder w/ImageNet	0.071 ± 0.070	0.104 ± 0.074	0.105 ± 0.077	0.091 ± 0.075
Our encoder w/BrainDIVE	0.083 ± 0.077	0.102 ± 0.086	0.106 ± 0.093	0.093 ± 0.079
k=4				
NSD train	0.068 ± 0.065	0.093 ± 0.068	0.090 ± 0.072	0.086 ± 0.071
Our encoder w/ImageNet	0.071 ± 0.070	0.106 ± 0.075	0.104 ± 0.075	0.092 ± 0.076
Our encoder w/BrainDIVE	0.083 ± 0.078	0.102 ± 0.086	0.110 ± 0.093	0.093 ± 0.079
k=2				
NSD train	0.073 ± 0.068	0.093 ± 0.068	0.090 ± 0.076	0.082 ± 0.071
Our encoder w/ImageNet	0.070 ± 0.068	0.108 ± 0.073	0.102 ± 0.073	0.088 ± 0.073
Our encoder w/BrainDIVE	0.084 ± 0.079	0.103 ± 0.086	0.111 ± 0.087	0.094 ± 0.081
k=1				
NSD train	0.062 ± 0.069	0.079 ± 0.075	0.081 ± 0.076	0.078 ± 0.063
0 1 7 37	0.068 ± 0.070	0.105 ± 0.070	0.104 ± 0.076	0.086 ± 0.076
Our encoder w/ImageNet	0.000 ± 0.010	0.100 ± 0.010	0.101 - 0.010	0.000 0.0.0

A.17 Generating concept vectors using predicted activations

In the main text, we generated the baseline NSD train concept vector using the top maximally-activating images based on measured responses of the training set. Here, we generate concept vectors using top images based on predicted responses of the same training set—reproducing tables 2, 3, 4, and 5. Interestingly, the rankings generated by the model-predicted concept vector outperforms the concept vector generated by measured responses, and even our model in some cases. We suspect that using model-predicted responses may further reduce the noise compared to measured responses and be responsible for the increase in performance. We did not include the model-predicted concept vector performance in the main text since it's not quite a baseline (since it relies on the encoding model choice).

Table 14: Fraction of parcels whose model-derived label predicts parcel activation rankings significantly better than chance (p < 0.05, FDR corrected)

Model type	S1	S2	S5	S7
NSD train (ground truth)	150/181	163/192 176/192	136/175	155/196
NSD train (model-predicted)	169/181		156/175	168/196

Table 15: Spearman's ρ (mean \pm std) between the model-predicted and ground-truth activation rankings on the NSD test set, averaged across parcels.

Model type	S1	S2	S5	S7
NSD train (ground truth)	0.162 ± 0.098	0.164 ± 0.091	0.150 ± 0.094	0.148 ± 0.086
NSD train (model-predicted)	0.191 ± 0.092	0.187 ± 0.084	0.174 ± 0.087	0.155 ± 0.082

Notes. INet = ImageNet. BD = BrainDIVE.

Table 16: Fraction of parcels shared across subjects whose model-derived label predicts parcel activation rankings significantly better than chance (p < 0.05, FDR corrected)

Model type	S 1	S2	S5	S7
NSD train (ground truth) NSD train (model-predicted)	37/49 40/49	44/49 46/49	40/49 $43/49$	$40/49 \ 40/49$

Table 17: Spearman's ρ (mean \pm std) between the model-predicted (from all subjects other than heldout) and ground-truth activation rankings on the NSD training set, averaged across parcels.

Model type	S1	S2	S5	S7
NSD train (ground truth)	0.068 ± 0.068	0.094 ± 0.065	0.090 ± 0.074	0.091 ± 0.070
NSD train (model-predicted)	0.078 ± 0.074	0.110 \pm 0.068	0.113 \pm 0.081	0.101 ± 0.074

Notes. INet = ImageNet. BD = BrainDIVE.

A.18 LLM prompts used

We used OpenAI's GPT-4o [52] to caption the images selected by our model in order to obtain a linguistic representation of the selectivity of a parcel. The captions in Figure 7 were generated with ChatGPT 4o (04/06/2025) with a collage of top-25 encoder-ranked ImageNet images shown, followed by the instruction "Describe a theme present in most of the images presented." The captions in Table 7 were generated by GPT-4o (05/12/2025, via the API) with the prompt "Give keywords for the central concept or categories present in these images."

A.19 NSD dataset quirks

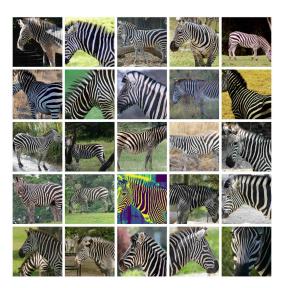


Figure S10: Top 25 maximally parcel-activating ImageNet images (based on encoder predictions) for a parcel in subject 2, left hemisphere.



Figure S11: Top 25 maximally parcel-activating NSD images (based on ground truth) for a parcel in subject 2, left hemisphere.

A.20 Compute used

We used CPU (AMD EPYC 7662), GPU (NVIDIA A40, L40), memory, and storage resources from an internal cluster. Storage for the entire project totals roughly 10TB. Training the model used roughly 3,000 GPU hours, 24,000 CPU core hours, and 32 GB per GPU hour. Running the remaining experiments used roughly 20,000 GPU hours, 160,000 CPU core hours, and 32 GB per GPU hour. The full project required more compute than these estimates due to failed experiments, experiments not included in the paper, and model iteration.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: In the abstract, we explain our approach and the results, include parcels of which we labeled the selectivity and our statistical tests that verify our labels. The Section 3 Methods explains the implementation details for our model. Section 4 Experiments show examples of parcels that we labeled and the results from the statistical tests that show our labels do drive the activation of the parcels.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals
 are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: Limitations are discussed in Section 5.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: The paper does not include theoretical results.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and crossreferenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: Information about the datasets, model architecture, and steps to perform statistical tests are all contained in the paper.

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
- (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
- (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
- (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: The supplementary zip file contains open access to the data and code with sufficient documentation.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (https://nips.cc/ public/guides/CodeSubmissionPolicy) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- · The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (https: //nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- · At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: Training and test details can be found in the supplementary zip file containing the code to reproduce the experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Justification: Error bars and p-values are reported in all tables. Details for implementation of the statistical tests are found in the main text and appendix.

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: See Appendix A.20 for estimates on computer resources used.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics https://neurips.cc/public/EthicsGuidelines?

Answer: [Yes]

Justification: The research conforms in every respect with the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a
 deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: The paper's societal impacts are discussed in Appendix A.1.

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.

- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: The paper poses no such risks.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with
 necessary safeguards to allow for controlled use of the model, for example by requiring
 that users adhere to usage guidelines or restrictions to access the model or implementing
 safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do
 not require this, but we encourage authors to take this into account and make a best
 faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: The paper cites any data and models used and respects the respective licenses.

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.

- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes]

Justification: Anonymized assets will be provided in the supplementary zip. Materials will include documentation.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects. Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.

• For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [No]

Justification: The core method in this research does not involve LLMs as any important, original, or non-standard components.

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (https://neurips.cc/Conferences/2025/LLM) for what should or should not be described.