

Eccentricity-Constrained CNN Training Reveals Adaptive Information Coding Around the Visual Field

Dylan M. Diaz

Departments of Computer Science and Psychology, California State University, San Bernardino
dylan.diaz4811@coyote.csusb.edu

Margaret M. Henderson

Department of Psychology and Neuroscience Institute, Carnegie Mellon University
mmhender@andrew.cmu.edu

Abstract

Within topographic eccentricity maps in the primate visual system, center-preferring populations have higher spatial resolution and overlap face- and word-selective regions while periphery-preferring populations have lower spatial resolution and overlap scene-selective regions. Prior behavioral and neuroimaging evidence suggests that this “eccentricity bias” may reflect the relevance of visual field portions for different tasks: the central visual field may be more informative for fine-grained tasks like face recognition and reading, while the periphery may be more informative for large-scale scene understanding tasks. To examine whether such eccentricity-dependent coding can emerge from natural experience, we leveraged egocentric video and eye-tracking data from the Visual Experience Dataset (VEDB). We trained ResNet-18 models using contrastive learning (SimCLR) on video frames modified to isolate information available at different eccentricities (gaze-contingent fovea-only crops, periphery-only crops, and periphery-only crops with a NeuroFovea transform applied). We then evaluated downstream task performance and model alignment with human fMRI data (Natural Scenes Dataset; encoding model framework). When examining in-domain classification of VEDB frame categories, we observed systematic variability in the performance of fovea-only and periphery-only models across categories, suggesting differential informativeness of visual field eccentricities across tasks. On downstream classification without fine-tuning, VEDB-pretrained models generalized more strongly to scene recognition (Places365) than to face recognition (VGGFace2), with fovea-only models showing an advantage on both tasks. Across visual cortex, VEDB-pretrained models achieved similar neural predictivity to models trained on mid-sized non-egocentric datasets (ImageNet-100), suggesting experience-sampled egocentric data, despite its low diversity and constrained semantic content, supports emergence of cortically-aligned representations. In scene-selective sub-regions (PPA, RSC), periphery-only models held a small but consistent advantage in explained variance over fovea-only models, suggesting scene-selective cortex may be adapted to peripheral-field statistics. Together, these results suggest that naturalistic egocentric experience provides an

organizing constraint on perception, leading to adaptive, task-aligned information processing.

Introduction

Visual field eccentricity is a key organizing dimension of the primate visual system, with topographic maps of preferred eccentricity found across early and higher visual cortex. These spatial eccentricity maps also align with topographic representations of mid-level features (curvature, spatial frequency) and semantic categories (faces, words, scenes), such that regions coding the center of visual space tend to have higher resolution and be selective for curved contours, as well as overlapping with face-selective and word-selective regions, while regions preferring the periphery tend to overlap with scene-selective visual regions (Arcaro & Livingstone, 2017; Hasson et al., 2002; Henderson et al., 2023; Levy et al., 2001; Ponce et al., 2017; Srihasam et al., 2014; Yue et al., 2014, 2020).

According to one hypothesis, this organization may reflect the differential task-relevance of the central and peripheral visual field portions for different behavioral tasks. The central visual field, with its small spatial receptive fields and high acuity, may be useful for fine-grained tasks like face recognition and reading, while more global tasks like scene recognition may benefit from the larger receptive fields of neurons in peripheral-coding regions (Hasson et al., 2002; Henderson et al., 2023; Larson & Loschky, 2009; Wang & Cottrell, 2017). Relatedly, the eccentricity biases found in higher visual cortex in the adult visual system may emerge in part due to the portion of retinotopic space in which face, word, and scene content typically occur (e.g., faces are often foveated, buildings often land in the peripheral field).

This account suggests that in natural egocentric inputs, image statistics may differ systematically between the central and peripheral portions of the visual field. These differences may include mid-level properties like curvature and spatial frequency, and/or semantic content (e.g., faces or scene components). If the visual system






Figure 1: Method overview. VEDB frames and gaze metadata are processed to create eccentricity-constrained inputs (Baseline, Fovea-Gaze, Periph-NF, Periph). Fovea-Gaze uses 112×112 crops centered on per-frame gaze position. Condition-specific SimCLR pretraining uses a Siamese architecture with tied (shared) ResNet-18 weights across the two augmented views; see *Methods*. Downstream evaluation uses frozen-backbone linear probes (in-domain, VGGFace2, Places365) and voxel-wise encoding on NSD.

is adapted to these statistics, this eccentricity-specific coding could have functional benefits for downstream tasks like face and scene recognition.

Critically, testing whether such eccentricity-dependent coding emerges from natural experience requires large-scale naturalistic egocentric training data that incorporate eye movement behavior. To achieve this here, we leverage the Visual Experience Dataset (VEDB; Greene et al., 2024) to train eccentricity-constrained convolutional neural network (CNN) models. We extract VEDB frames and modify them to isolate either a fixed gaze-centered crop around the participant’s per-frame gaze location (Fovea-Gaze) or the complementary peripheral region (Periph); we additionally use the NeuroFovea transform (Deza et al., 2019) to simulate information loss in peripheral vision (Periph-NF). We then pretrain CNNs using self-supervised contrastive learning (SimCLR; Chen et al., 2020) on each dataset version, and test performance on downstream tasks probing various aspects of scene (Places365; Zhou et al., 2017) and face (VGGFace2; Cao et al., 2017) perception. Finally, we evaluate neural predictivity of each model using a large fMRI dataset (Natural Scenes Dataset; NSD; Allen et al., 2022). Our key contributions are as follows:

- We show that pretraining with Fovea-Gaze vs. Periph or Periph-NF inputs leads to systematic variability in downstream task performance across categories, providing preliminary evidence that the central and pe-

ripheral field are differentially relevant for classifying object-centric versus large-scale scene content.

- We demonstrate VEDB-pretrained models generalize more effectively to scene classification (Places365) than face classification (VGGFace2), and that Fovea-Gaze outperforms peripheral models on both tasks.
- We show that self-supervised VEDB pretraining leads to comparable fMRI encoding model prediction accuracy relative to mid-sized non-egocentric datasets (STL-10, ImageNet-100), approaching accuracy of models trained on a large dataset (ImageNet-1K).
- We find small but statistically reliable differences in unique variance, with the Periph model favored in scene-selective cortex (PPA, RSC) and the Fovea-Gaze model favored in primary visual cortex.

Related Work

Behavioral evidence suggests that object and face recognition are best in the central visual field (Henderson & Hollingworth, 1999; Nelson & Loftus, 1980; Pelli & Tillman, 2008; Rosenholtz, 2016), while scene recognition may rely more on peripheral features (Larson & Loschky, 2009; McCotter et al., 2005). The role of peripheral vision in scene recognition may reflect large receptive fields and gist-like structural information, with some models proposing peripheral vision represents information in a

texture-like code with large pooling fields (Freeman & Simoncelli, 2011; Oliva & Torralba, 2006; Rosenholtz, 2016; Rosenholtz et al., 2012). In the higher visual system, information from the central and peripheral field is partially anatomically segregated, with central-preferring regions overlapping face-selective and word-selective areas, and peripheral-preferring regions overlapping with scene-selective areas (Grill-Spector & Weiner, 2014; Hasson et al., 2002; Levy et al., 2001).

Computational work has also examined the functional task-relevance of central and peripheral vision. CNN modeling has shown that peripheral features are more useful than central features for scene classification (Wang & Cottrell, 2017), and that modeling separate central and peripheral streams benefits downstream task performance (Zhang et al., 2019). Other work has demonstrated simulating cortical magnification and texture-like peripheral encoding can improve model robustness and performance (Deza & Konkle, 2020; Wang et al., 2021), and that modeling peripheral processing can boost classification performance in a transformer model (Min et al., 2022). Recent work has further connected these ideas to neural encoding and cortical organization, showing that biologically motivated spatial sampling of DNN feature maps improves encoding models of human neural responses (Müller et al., 2024), and that retinotopically organized inputs can scaffold consistent category-selective topography in models of high-level visual cortex (Blauch et al., 2025).

Much of this previous computational work has used non-egocentric computer vision datasets, whose curation and content coverage may differ strongly from the statistics of naturalistic first-person visual experience (Grauman et al., 2022; Greene et al., 2024). Moreover, this past work does not account for eye gaze position within scenes, meaning the retinotopic location of features may not be modeled accurately. Here we fill these gaps by training eccentricity-constrained CNN models on gaze-aligned frames of egocentric video data. Importantly, it is not trivial *a priori* that self-supervised models trained on such data, which is temporally correlated and less diverse than curated image datasets, would learn useful or performant visual representations, relative to curated image datasets. We also provide the novel contribution of directly comparing our models with human fMRI data recorded across early and category-selective cortical regions, enabling a test of the link between egocentric image statistics and cortical representations.

Methods

VEDB Processing and Metamer Generation

We analyzed egocentric video from the Visual Experience Database (VEDB; Greene et al. 2024), retaining sessions with synchronized gaze. Frames were sampled uniformly every 2 seconds within task-relevant session windows,

decoded to RGB, and center-cropped to 224×224 . All splits were performed at the session level to prevent leakage from temporally adjacent frames; SimCLR pre-training used an 80/10/10 train/val/test session split, and in-domain classification was restricted to a core subset of 17 tasks. To approximate peripheral texture-tiling representations, we generated gaze-contingent peripheral metamers using NeuroFovea (NF; Deza et al. 2019), a feed-forward foveated style-transfer model (Freeman & Simoncelli, 2011; Rosenholtz, 2019), applied per frame with pooling centered at the per-frame fixation. See Supp. Figure 2 for frame examples and *Supp. Methods* for full details.

Experimental Setup and Augmentations

The Baseline, Fovea-Gaze, Periph, and Periph-NF conditions followed an identical two-stage pipeline: (1) SimCLR pre-training on condition-specific augmented inputs, followed by (2) supervised linear-probe evaluation on in-domain classification and transfer tasks (VGGFace2 and Places365 classification).

1) Baseline. Models were trained on the original, unedited 224×224 frames.

2) Fovea-Gaze. We generated Fovea-Gaze views by extracting a fixed 112×112 crop from each 224×224 frame, centered on the participant’s per-frame gaze location (see *Supp. Methods*). Crops were clamped to image bounds and upsampled to 224×224 (bilinear interpolation) to maintain a constant input size. To suppress non-crop information, we applied a feathered circular aperture; surrounding region masked with uniform gray (gray=128; Gaussian blur kernel=15); frames without high-confidence gaze used a centered fallback crop. This condition should be interpreted as gaze-centered central-only input, not a simulation of biological fovea.

3) Periph-NF. Periph-NF images were constructed by applying the NeuroFovea transform to each frame prior to occluding the foveal region with a gray circular scotoma (pixel value = 128, Gaussian-blurred mask edges, kernel = 15) centered at the per-frame gaze-defined location. This differs from the Periph condition only in the application of NeuroFovea prior to masking, which introduces eccentricity-dependent texture pooling approximating the reduced acuity of peripheral vision.

4) Periph. Control condition of Periph-NF, constructed identically but without the NeuroFovea transform.

SimCLR Pre-Training

SimCLR pre-training was run once for each condition using a ResNet-18 encoder for 120 epochs (batch size = 64). The augmentation pipeline comprised random resized cropping, horizontal flipping, color jitter, grayscale




Figure 2: Performance of VEDB models during pretraining and linear probing. **(A)** SimCLR pre-training curves across conditions (120 epochs). **(B)** In-domain linear probe learning curves across conditions (20 probe epochs). **(C)** VGGFace2 transfer linear probe learning curves across conditions (20 probe epochs). **(D)** Places365 transfer linear probe learning curves across conditions (20 probe epochs).

conversion, and Gaussian blur, and was identical across all conditions, applied after condition-specific frame construction. Full training hyperparameters and projection head architecture are provided in Supp. Table 2.

Non-VEDB Baseline Models

As comparison baselines, we included ResNet-18 models pretrained with SimCLR on three non-egocentric datasets: STL-10 (Coates et al., 2011), using a pretrained checkpoint (<https://github.com/Spijkervet/SimCLR>; 40 epochs); ImageNet-100 (Tian et al., 2019; a subset of ImageNet including 100 randomly sampled classes), trained using the LightlySSL codebase, following a standard SimCLR augmentation pipeline for 120 epochs; and ImageNet-1K (Russakovsky et al., 2015; images resized to 256 px to reduce computational demand), trained using LightlySSL for 100 epochs.

Downstream Classifiers

For in-domain classification and transfer to VGGFace-2 and Places-365, we trained linear probes on a frozen ResNet-18 backbone initialized from each SimCLR checkpoint. For in-domain evaluation, frames were pre-processed identically to pretraining; no such modifications were applied for transfer. All probes were trained for 20 epochs (AdamW); checkpoints were selected by validation Macro-F1 (in-domain) or Top-1 accuracy (transfer). Hyperparameters and dataset splitting details are summarized in Supp. Table 1 and *Supp. Methods*.

fMRI Encoding Models

To compare the ability of eccentricity-constrained pre-trained models to predict brain activity across the human visual system, we used a large 7T fMRI dataset consisting of whole-brain responses to complex natural scene

images in 8 human participants (NSD; Allen et al., 2022). We constructed voxelwise encoding models (Naselaris & Kay, 2015) which map from viewed stimuli to single-voxel image-wise activation. Voxel beta values were extracted using GLMSingle (Prince et al., 2022) and averaged over repetitions of each image.

To build encoding models for each participant, we first extracted the model activations from a pretrained ResNet-18 model in response to each stimulus image viewed by the fMRI participant. This was done for each of the Baseline, Fovea-Gaze, Periph, and Periph-NF models separately (following the SimCLR pretraining step). Features were extracted from the original, intact NSD images for all model conditions; condition-specific transforms were applied only during SimCLR pretraining and were not reapplied at the encoding stage. As control comparisons, we also included the non-VEDB models described above (see *Methods* "Non-VEDB Baseline Models"). From each model, we used the output of each of the layers: Conv1, Layer1.1, Layer2.1, Layer3.1, Layer4.1, Avgpool. To reduce dimensionality of features, we used average pooling to spatially downsample convolutional layers, flattened the features, then applied principal components analysis (PCA) and retained the top 200 components. We concatenated features across all layers, z-scored the features across images, and used regularized (L2) linear regression to learn a voxel-specific linear mapping from the model features to the voxel's response, including an intercept term. Encoding model fit was evaluated based on the model's ability to predict responses to held-out images, quantified using R^2 on the held-out test set. For each participant, we used the 1000 images in the NSD that were shared across all participants (see Allen et al., 2022 for details) as our test set and fit the model on the remaining 9000 images viewed by the participant.

Table 1: Summary Metrics for Pretraining and Linear Probes. SimCLR Top-1 is computed from the self-supervised contrastive objective and is not directly comparable to downstream classification. For downstream tasks, classifier checkpoints were selected by best validation Macro-F1. In-domain Top-1 accuracy is omitted because label imbalance across frames can make accuracy misleading; Macro-F1 is reported as the primary class-balanced metric. STL-10, ImageNet-100, and ImageNet-1K are treated as out-of-domain baselines because they were not pretrained on VEDB.

Task	Condition	Val Loss	Top-1 (%)	Top-5 (%)	Best Macro-F1 (%)
SimCLR	Baseline	0.4331	87.60	—	—
	Fovea-Gaze	0.3749	90.43	—	—
	Periph-NF	0.4548	90.04	—	—
	Periph	0.4545	89.26	—	—
In-Domain	Baseline	0.9811	—	—	42.17
	Fovea-Gaze	1.2031	—	—	43.64
	Periph-NF	1.3090	—	—	30.93
	Periph	1.0623	—	—	36.56
	STL-10	1.6666	—	—	25.41
	ImageNet-100	1.2342	—	—	41.23
	ImageNet-1K	0.9713	—	—	43.33
VGGFace2	Baseline	7.8101	5.21	11.73	3.26
	Fovea-Gaze	7.9104	4.58	10.76	2.70
	Periph-NF	8.0232	3.39	8.17	1.90
	Periph	8.1681	2.54	6.39	1.35
	STL-10	6.9973	9.55	18.96	7.43
	ImageNet-100	6.7985	10.77	21.07	8.71
	ImageNet-1K	6.7964	10.74	21.08	8.77
Places365	Baseline	3.9690	25.63	51.90	23.16
	Fovea-Gaze	4.2347	21.86	46.21	19.14
	Periph-NF	4.2621	20.51	44.58	17.86
	Periph	4.2671	20.26	44.10	17.65
	STL-10	3.8281	26.57	53.47	24.82
	ImageNet-100	3.9207	24.99	51.21	23.32
	ImageNet-1K	3.6264	30.17	58.46	28.36

Functional regions of interest (ROIs) were defined using data from independent localizers. Retinotopic early visual ROIs (V1, V2, V3, hV4) were defined using a receptive field mapping task, and category-selective ROIs (scene-selective: OPA, PPA, RSC; face-selective: FFA, OFA; word-selective: VWFA, OWFA) were defined using a category localizer task (see Allen et al., 2022). We performed additional thresholding of the category-selective ROI definitions provided by Allen et al., 2022, retaining only voxels with a localizer test statistic of $t > 2$ for the region’s preferred category. Voxels were further thresholded by their estimated noise ceiling (threshold 0.10).

To perform variance partitioning between a pair of models A and B , we created a combined feature space that concatenated all the features across both models, fit the weights and computed cross-validated R^2 for this combined model. Based on this combined fit and the fits of each model alone, we computed: $R^2_{\text{unique:A}} = R^2_{\text{combined}} - R^2_{B \text{ only}}$, $R^2_{\text{unique:B}} = R^2_{\text{combined}} - R^2_{A \text{ only}}$.

Results

Downstream task performance

After pretraining each model, we used linear probing to evaluate performance on in-domain frame label classification, face classification and scene classification, revealing differences in performance across conditions (Figure 2 and Table 1). In the SimCLR pretraining phase, all four VEDB conditions performed well and within a similar range, but Fovea-Gaze yielded the strongest performance (Fovea-Gaze Top-1 = 90.43%), slightly exceeding both peripheral variants (Periph-NF = 90.04%; Periph = 89.26%), suggesting a small foveal advantage during representation learning. In the in-domain supervised probe, final validation loss was lowest for the Baseline condition, which is expected given that the Baseline images have more total information retained compared to the modified versions. However, class-balanced Macro-F1 accuracy showed that the Fovea-Gaze model outperformed both the Baseline and the two Periph models (Macro-F1: Fovea-Gaze = 43.64%, Periph-NF = 30.93%, Periph = 36.56%, Baseline = 42.17%). The advantage for the Fovea-Gaze model in this task may suggest that foveally-cropped frames isolate the most diagnostic features for supporting frame category classification. The advantage for the Periph condition over the Periph-NF condition suggests that fine spatial detail in the periphery, which is removed from images via the NeuroFovea transform, provides another source of information for scene classification.

Comparing against non-VEDB baselines on the in-domain task, the VEDB Baseline and Fovea-Gaze models (Macro-F1: 42.17%, 43.64%) performed comparably to ImageNet-1K (43.33%) and ImageNet-100 (41.23%), and outperformed STL-10 (25.41%). This result suggests that egocentric pretraining led to representations competitive with more curated datasets for classifying naturalistic scene content.

To better understand these condition differences, we next examined differences across individual categories (Table 2). When comparing Macro-F1 across conditions, we observed the Fovea-Gaze model outperformed the Periph-NF and Periph models for a subset of tasks, such as “playing video game”, “computer work”, and “socializing”, while Periph-NF and Periph outperformed the Fovea-Gaze model for other tasks like “skateboarding”, “playing frisbee”, “playing pool”. These results suggest a pattern in which fovea-constrained models learn more accurate representations of scene categories that include fine-scale detail or social information (screens, faces), while peripheral-constrained models learn more accurate representations of scene categories that include outdoor context or involve large-scale motor actions. However, there are exceptions to this suggestive pattern (e.g., “washing dishes” shows a peripheral advantage that does not intu-

Table 2: Top classes for in-domain classification, ranked by $\max_j |\Delta_j|$, where Δ_j is the F1 difference between two conditions. Positive values indicate condition A > B. Asterisks indicate comparisons whose 95% bootstrap CI does not overlap zero (1000-replicate parametric multinomial bootstrap; see Supp. Figure 4 for full CIs). See Supp. Figure 2 for examples of each label and Supp. Figure 1 for bar-plot visualization of $\Delta F1$.

Rank	Label	Fovea > Base	Fovea > Periph-NF	Fovea > Periph
1	playing video game	0.433 *	0.433 *	0.433 *
2	solving rubik's cube	-0.116 *	0.428 *	0.018
3	socializing	0.365 *	0.318 *	0.336 *
4	computer work	0.079 *	0.300 *	0.281 *
5	skateboarding	-0.292 *	-0.170 *	-0.282 *
6	playing frisbee	-0.251 *	-0.214 *	-0.263 *
7	washing dishes	-0.263 *	-0.100 *	-0.075 *
8	cooking	-0.221 *	0.029 *	-0.078 *
9	playing pool	-0.200 *	-0.178 *	-0.194 *
10	playing cards	0.178 *	0.147 *	0.178 *
11	exploring	0.023	0.124 *	0.077 *
12	playing poker	-0.116 *	0.053 *	-0.055 *
13	preparing	-0.038	-0.041 *	-0.094 *
14	playing sudoku	-0.071 *	-0.048 *	0.010
15	walking	-0.040 *	-0.028 *	-0.030 *
16	playing piano	0.002	0.002	0.002
17	standing	0.000	0.000	0.000

itively map onto properties of that task), indicating further work is needed to build on these preliminary results.

To further explore representational differences across conditions, we performed MDS on the confusion matrices from each classifier (Figure 3; Supp. Figure 3), where categories positioned closer together are those more frequently confused. These plots show coarse categorical organization, such that tasks like “cooking”, “washing dishes”, and “playing sudoku” are represented similarly, while tasks like “standing”, “walking”, and “exploring” are also represented similarly. These groupings suggest the models capture semantic or functional structure, such as stationary indoor activities versus mobile or navigational tasks. Notably, this organization emerged from self-supervised training without explicit category supervision, suggesting that the natural statistics of first-person visual experience contain sufficient structure to support categorical distinctions.

Differences in this structure are evident across conditions. For example, the “socializing” category in the Baseline, Periph and Periph-NF model is tightly clustered with categories like “standing”, “walking” and “preparing”, while in the Fovea-Gaze condition it is more separated from these neighboring categories. This may indicate that fovea-restricted inputs led to more effective differentiation of images containing social information (e.g., faces) from those not containing social information, which could underlie the difference in macro-F1 across conditions for the “socializing” label.

To provide a stronger test of these preliminary observations, we tested transfer to out-of-domain categoriza-




Figure 3: Multidimensional scaling (MDS) was performed using confusion patterns across the 17 task labels (see Supp. Figure 3 for confusion matrices). Categories positioned closer together are more often confused by the classifier. Embeddings were aligned across conditions by a Procrustes transform.

tion of face (VGGFace2) and scene (Places365) images using a linear probe. Across all four VEDB conditions, we observed better overall transfer to Places365 than VGGFace2, which likely reflects the greater input distribution shift between VEDB and VGGFace2. Similarly, the non-VEDB baseline models (STL-10, ImageNet-100, ImageNet-1K) also performed better on Places365 than VGGFace2 (note also that models were not fine-tuned). The non-VEDB baselines performed comparably to the VEDB-trained models on Places365 (Macro-F1: STL-10 = 24.82%, ImageNet-100 = 23.32%, ImageNet-1K = 28.36% vs. VEDB Baseline = 23.16%), but all non-VEDB baselines outperformed all VEDB-pretrained models on the VGGFace2 task, potentially reflecting limited face content in VEDB (see *Discussion*).

Across the VEDB eccentricity-constrained conditions, both VGGFace2 and Places365 showed a clear fovea>periphery ordering: on VGGFace2, Fovea-Gaze outperformed both peripheral variants (Macro-F1: 2.70% vs. 1.90% and 1.35%). Surprisingly, Fovea-Gaze similarly exceeded peripheral variants on Places365 (Macro-F1: 19.14% vs. 17.86% and 17.65%). This finding may be related to the semantic nature of this scene categorization task; see *Discussion*. Comparing the two peripheral conditions, Periph-NF outperformed Periph, suggesting periphery-inspired metamerization provided a small robustness benefit for out-of-domain transfer. This held for both types of transfer (VGGFace2 Macro-F1: 1.90% vs. 1.35%; Places365 Macro-F1: 17.86% vs. 17.65%). This is in contrast to the in-domain probe, in

which Periph achieved higher Macro-F1 than Periph-NF (36.56% vs. 30.93%).

fMRI encoding model performance

Next, we used a voxelwise encoding model framework (see *Methods*) to examine alignment of each model's learned representations with the representations in target visual cortex regions. We hypothesized that training on foveal-only inputs would lead to models more predictive of face- and word-selective cortical regions, while training on peripheral-only inputs would lead to models more predictive of scene-selective cortical regions. We first assessed overall model performance using the cross-validated predictive accuracy (R^2) of each VEDB-pretrained model and each non-VEDB comparison model (Figure 4A-B). We observed moderately high values of R^2 across voxels in early and higher visual cortex regions for all of our VEDB-pretrained models, with voxelwise R^2 close to the noise ceiling in many voxels (Supp. Figure 4). Across ROIs, R^2 for the Baseline VEDB model was comparable to that of the ImageNet-100 model, consistently exceeded the STL-10 model, and was lower but approaching the accuracy of the ImageNet-1K model. This difference was largest in FFA, where Baseline-VEDB R^2 was 0.11 ± 0.01 (mean $\pm$ SEM across 8 participants) and ImageNet-1K R^2 was 0.15 ± 0.01 . This suggests that pretraining with egocentric frames from VEDB, despite the more constrained and temporally correlated nature of experience-sampled data, led to comparably cortically-aligned models as pretraining on mid-sized non-egocentric datasets.

In addition, we observed differences across the different VEDB conditions, which varied across ROIs. This pattern was quantified using a two-way repeated-measures ANOVA which revealed significant main effects of ROI and model condition, as well as an interaction between ROI and model condition (ROI: $F(10, 70) = 43.68$, $p < 0.0001$; ModelCondition: $F(3, 21) = 70.17$, $p < 0.0001$; ROI:ModelCondition: $F(30, 210) = 5.95$, $p < 0.0001$; all p -values obtained using a permutation test; see *Supp. Methods*). Follow-up t -tests revealed that across ROIs, all eccentricity-constrained conditions (Fovea-Gaze, Periph, and Periph-NF) explained less variance than the Baseline VEDB model (see Supp. Table 3 for test statistics), consistent with the idea that cropping images reduces their total information content leading to worse neural alignment. The main exception to this was V1, in which the Fovea-Gaze and Periph-NF models performed as well as the Baseline model, though the Periph model was worse than Baseline.

Most importantly, consistent with our initial hypothesis, the Periph model out-performed the Fovea-Gaze model in R^2 for scene-selective cortical regions PPA and RSC (PPA: $t(7) = -4.86$, $p = 0.0074$; RSC $t(7) = -4.91$, $p = 0.0070$; paired t -test with permutation; see Supp. Table 3




Figure 4: fMRI encoding model performance. **(A)** Voxelwise prediction accuracy (cross-validated R^2) for Baseline VEDB and non-VEDB comparison models, averaged across voxels in functional ROIs for each participant. **(B)** As in (A), for the four VEDB conditions. See Supp. Figure 4 for comparison of R^2 vs. voxelwise noise ceiling, and see Supp. Table 3 for significance testing between conditions. **(C)** Unique variance for Fovea-Gaze model and Periph model. **(D)** Unique variance for Periph model and Periph-NF model. Brackets and * indicate significant difference between unique variance across models (paired t -test with permutation, all $p < 0.01$; see *Supp. Methods*). In all panels: gray dots and lines indicate individual fMRI participants, error bars show the SEM across 8 participants. Sel. = selective.

for all test statistics). In face-selective and word-selective cortical regions, R^2 did not significantly differ between the

Fovea-Gaze and the Periph model, but in primary visual cortex (V1), R^2 for the Fovea-Gaze model was significantly higher than for the Periph model ($t(7) = 7.04$, $p = 0.0062$). Further confirming these differences, a variance partition analysis between the Fovea-Gaze and Periph model (Figure 4C) demonstrated that the Periph model explained significantly more unique variance in PPA and RSC than the Fovea-Gaze model, while the Fovea-Gaze model explained significantly more unique variance than the Periph model in V1. While numerically small, this difference between Periph and Fovea-Gaze was found for all 8 individual NSD participants for both PPA and RSC, reinforcing the reliability of this small effect. We also observed a difference between the Periph and Periph-NF conditions, such that the Periph-NF model outperformed the Periph model in V1 ($t(7) = -5.62$, $p = 0.0074$), while the Periph model outperformed the Periph-NF model in scene-selective regions OPA and PPA (OPA: $t(7) = 3.76$, $p = 0.0066$; PPA: $t(7) = 2.97$, $p = 0.0084$). This pattern was confirmed with a variance partition analysis (Figure 4D), which showed significantly more unique variance in OPA and PPA was explained by the Periph model, while more unique variance in V1 was explained by the Periph-NF model. The performance of the Periph-NF model was similar to the Fovea-Gaze model across all ROIs, with no significant differences observed between these models.

Discussion

Natural visual inputs are highly structured, and the visual system utilizes this structure to enable adaptive behavior. In this work, we have demonstrated that CNN models trained on foveal or peripheral portions of an egocentric visual field have differential patterns of performance across categorization tasks, implying a functional difference in the statistics of these visual field portions. We also showed that self-supervised training with egocentric video frames led to comparable fMRI encoding model accuracy as training on mid-sized non-egocentric datasets, suggesting that experience-sampled egocentric data, despite its more constrained distribution, captures sufficient image statistics to support learning of cortically-aligned representations. Finally, we demonstrated that periphery-trained models had higher predictive accuracy in scene-selective cortical regions compared to fovea-trained models. This indicates scene-selective cortical regions may selectively represent the features found in the peripheral visual field, potentially supporting behaviors such as navigation.

Across all tasks, our Fovea-Gaze model outperformed the Periph and Periph-NF models. This foveal advantage was most pronounced for face classification (VGGFace2), which is directionally consistent with our hypothesis that the central field is more informative for face processing. However, absolute performance on VGGFace2 was

low across all conditions, potentially because VEDB pre-training has limited face content. The fovea > periphery ordering on this task should therefore be interpreted as suggestive rather than definitive.

Surprisingly, a foveal advantage was also observed for scene classification (Places365). One explanation for this is that the Places365 task is primarily a semantic task, which may be solvable based on diagnostic objects or materials, rather than a depth or navigational task that relies on scene geometry. This also fits with our finding that the Periph model did explain more unique variance in scene-selective cortical regions (PPA, RSC) than the Fovea-Gaze model. These scene-selective regions are thought to play key roles in encoding scene layout and other navigational information (Epstein, 2008; Epstein & Baker, 2019; Vann et al., 2009), so our results may suggest that peripheral training led to learning of more navigationally-relevant spatial information, leading to improved alignment with these areas. However, the Periph advantage in PPA and RSC was small, indicating further room for improvement in modeling scene-selective cortex.

We also observed a modest advantage for the Periph-NF model over the Periph model on out-of-distribution images, for both face and scene classification. This advantage may suggest that including the NeuroFovea transform as a precursor to self-supervised learning led to more robust model generalization performance. This is consistent with past work showing fovea-like transforms can improve robustness (Deza & Konkle, 2020; Wang et al., 2021), and also with work suggesting that blurring images during model training can benefit face recognition (Jang & Tong, 2024) and object categorization (Jinsi et al., 2023). In the brain, the Periph-NF model explained more unique variance in V1 than the Periph model, consistent with some past work suggesting a link between predictivity in early visual regions and model robustness (Dapello et al., 2020; Li et al., 2019; Piper et al., 2025), although V1 also showed an advantage in R^2 for the Fovea-Gaze over the Periph model. One explanation for this is that our analysis pools across all voxels regardless of retinotopic preference, so the Fovea-Gaze and Periph-NF advantages could reflect improvements in distinct voxel subpopulations. In contrast to V1, scene-selective cortical regions (OPA, PPA) were better predicted by the Periph model than the Periph-NF model, potentially indicating that high-resolution, non-texturized peripheral images are required in order to learn a good model of scene-selective cortex.

Our encoding model analysis did not reveal an advantage for the Fovea-Gaze model in predicting face-selective or word-selective cortical regions, as we had predicted. One potential explanation for this may be the large size of the cropping window we used in the Fovea-Gaze condition (112×112 px); future work can examine

the effect of a smaller cropping window (e.g., 10×10 px) to better approximate biological foveal vision. Pretraining models on foveal inputs during a supervised face or character recognition task may additionally lead to gains in prediction for face and word-selective regions.

Limitations of this current work include our use of static frames; future studies could use self-supervised video models. Future work should also investigate how foveally-trained and peripherally-trained models perform on a wider set of tasks, including more spatially-demanding scene tasks. Additionally, some downstream performance differences were modest in absolute magnitude, although we provide confidence intervals for in-domain comparisons, a full subsampling-based robustness analysis would better assess their stability; these results should be interpreted as suggestive rather than definitive. Furthermore, although we have shown differences between central and peripheral inputs, we have not yet determined which image properties underly these effects. Central and peripheral vision may differ in myriad features including spatial frequency, curvature, and semantic content; future work should aim to disentangle these factors.

Broadly, our work has provided a step forward in understanding the link between egocentric natural image statistics, behavioral task performance, and neural representations. Understanding this link will lead to new theories on the computational and biological constraints that govern visual cortex organization.

Acknowledgments and Disclosure

D.M.D. was supported by the National Institute of General Medical Sciences through the CSUSB U-RISE Program [grant number T34GM136467], the National Science Foundation through the Cal-Bridge Program [grant number 2425075], and the NSF ExCELS Program [grant number 2322436]. This work was initiated as part of D.M.D.'s participation in the CMU/Pitt Undergraduate Program in Neural Computation at the Center for the Neural Basis of Cognition in Pittsburgh, PA, with support from the National Institute on Drug Abuse through the Interdisciplinary Training in Computational Neuroscience program [grant number 5R90DA060340-03]. D.M.D. thanks Jason F. Reimer for helpful comments and discussions throughout the study. D.M.D. and M.M.H. also thank David Pane for technical support with Carnegie Mellon University's MiND high-performance computing cluster, on which the computational experiments were performed. AI and large language models were not extensively used in the creation of this work. The authors take full responsibility for all content and declare no conflicts of interest.

References

- Allen, E. J., St-Yves, G., Wu, Y., Breedlove, J. L., Prince, J. S., Dowdle, L. T., Nau, M., Caron, B., Pestilli, F., Charest, I., Hutchinson, J. B., Naselaris, T., & Kay, K. (2022). A massive 7t fmri dataset to bridge cognitive neuroscience and artificial intelligence. *Nature Neuroscience*, 25, 116–126. <https://doi.org/10.1038/s41593-021-00962-x>.
- Arcaro, M. J., & Livingstone, M. S. (2017). A hierarchical, retinotopic proto-organization of the primate visual system at birth. *eLife*, 6, e26196. <https://doi.org/10.7554/eLife.26196>.
- Blauch, N. M., Behrmann, M., & Plaut, D. C. (2025). Retinotopic scaffolding of high-level vision. *PsyArXiv*. https://doi.org/10.31234/osf.io/ryn timer _v2.
- Cao, Q., Shen, L., Xie, W., Parkhi, O. M., & Zisserman, A. (2017). Vggface2: A dataset for recognising faces across pose and age. *CoRR*, abs/1710.08092. <http://arxiv.org/abs/1710.08092>
- Chen, T., Kornblith, S., Norouzi, M., & Hinton, G. (2020). A simple framework for contrastive learning of visual representations. *37th International Conference on Machine Learning, ICML 2020, Part F168147-3*, 1575–1585.
- Coates, A., Lee, H., & Ng, A. Y. (2011). An analysis of single-layer networks in unsupervised feature learning. *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics (AISTATS)*.
- Dapello, J., Marques, T., Schrimpf, M., Geiger, F., Cox, D. D., & DiCarlo, J. J. (2020). Simulating a primary visual cortex at the front of cnns improves robustness to image perturbations. *Advances in Neural Information Processing Systems (NeurIPS)*, 33.
- Deza, A., Jonnalagadda, A., & Eckstein, M. P. (2019). Towards metamerism via foveated style transfer. *International Conference on Learning Representations (ICLR)*.
- Deza, A., & Konkle, T. (2020). Emergent properties of foveated perceptual systems. *arXiv preprint arXiv:2006.07991*. <https://arxiv.org/abs/2006.07991>
- Epstein, R. A. (2008). Parahippocampal and retrosplenial contributions to human spatial navigation. *Trends in Cognitive Sciences*, 12(10), 388–396. <https://doi.org/10.1016/j.tics.2008.07.004>.
- Epstein, R. A., & Baker, C. I. (2019). Scene perception in the human brain. *Annual Review of Vision Science*, 5, 373–397. <https://doi.org/10.1146/annurev-vision-091718-014809>.
- Freeman, J., & Simoncelli, E. P. (2011). Metamers of the ventral stream. *Nature Neuroscience*, 14, 1195–1204. <https://doi.org/10.1038/nn.2889>.
- Grauman, K., Westbury, A., Byrne, E., Chavis, Z., Furnari, A., Girdhar, R., Hamburger, J., Jiang, H., Liu, M., Liu, X., Martin, M., Nagarajan, T., Radosavovic, I., Ramakrishnan, S. K., Ryan, F., Sharma, J., Wray, M., Xu, M., Xu, E. Z., ... Malik, J. (2022). Ego4d: Around the world in 3,000 hours of egocentric video. <https://arxiv.org/abs/2110.07058>
- Greene, M. R., Balas, B. J., Lescroart, M. D., MacNeilage, P. R., Hart, J. A., Binaee, K., Hausamann, P. A., Mezile, R., Shankar, B., Sinnott, C. B., et al. (2024). The visual experience dataset: Over 200 recorded hours of integrated eye movement, odometry, and egocentric video. *Journal of Vision*, 24(11), 6–6.
- Grill-Spector, K., & Weiner, K. S. (2014). The functional architecture of the ventral temporal cortex and its role in categorization. *Nature Reviews Neuroscience*, 15, 536–548. <https://doi.org/10.1038/nrn3747>.
- Hasson, U., Levy, I., Behrmann, M., Hendler, T., & Malach, R. (2002). Eccentricity bias as an organizing principle for human high-order object areas. *Neuron*, 34, 479–490. [https://doi.org/10.1016/S0896-6273\(02\)00662-1](https://doi.org/10.1016/S0896-6273(02)00662-1).
- Henderson, J. M., & Hollingworth, A. (1999). The role of fixation position in detecting scene changes across saccades. *Psychological Science*, 10(5), 438–443. <https://doi.org/10.1111/1467-9280.00183>.
- Henderson, M. M., Rademaker, R. L., & Serences, J. T. (2022). Flexible utilization of spatial-and motor-based codes for the storage of visuo-spatial information. *eLife*, 11. <https://doi.org/10.7554/ELIFE.75688>.
- Henderson, M. M., Serences, J. T., & Rungratsameetaweemana, N. (2025). Dynamic categorization rules alter representations in human visual cortex. *Nature Communications*, 16, 3459.
- Henderson, M. M., Tarr, M. J., & Wehbe, L. (2023). Low-level tuning biases in higher visual cortex reflect the semantic informativeness of visual features. *Journal of Vision*, 23, 8–8. <https://doi.org/10.1167/JOV.23.4.8>.
- Jang, H., & Tong, F. (2024). Improved modeling of human vision by incorporating robustness to blur in convolutional neural networks. *Nature Communications*, 15, 1989. <https://doi.org/10.1038/s41467-024-46097-1>.
- Jinsi, O., Henderson, M. M., & Tarr, M. J. (2023). Early experience with low-pass filtered images facilitates visual category learning in a neural network model. *PLOS ONE*, 18(1), e0280145. <https://doi.org/10.1371/journal.pone.0280145>.
- Larson, A. M., & Loschky, L. C. (2009). The contributions of central versus peripheral vision to scene gist recognition. *Journal of Vision*, 9(10), 6. <https://doi.org/10.1167/9.10.6>.
- Levy, I., Hasson, U., Avidan, G., Hendler, T., & Malach, R. (2001). Center-periphery organization of human object areas. *Nature Neuroscience*, 4, 533–539. <https://doi.org/10.1038/87490>.
- Li, Z., Brendel, W., Walker, E. Y., Cobos, E., Muhammad, T., Reimer, J., Bethge, M., Sinz, F. H., Pitkow, X., & Tolias, A. S. (2019). Learning from brains how to

- regularize machines. *Advances in Neural Information Processing Systems (NeurIPS)*, 32.
- McCotter, M., Gosselin, F., Sowden, P., & Schyns, P. (2005). The use of visual information in natural scenes. *Visual Cognition*, 12(6), 938–953. <https://doi.org/10.1080/13506280444000645>.
- Min, J., Zhao, Y., Luo, C., & Cho, M. (2022). Peripheral vision transformer. *Advances in Neural Information Processing Systems (NeurIPS)*. <https://arxiv.org/abs/2110.12271>
- Müller, N., Scholte, H. S., & Groen, I. I. A. (2024). Spatial sampling of deep neural network features improves encoding models of foveal and peripheral visual processing in humans. *bioRxiv*. <https://doi.org/10.1101/2024.08.05.606515>.
- Naselaris, T., & Kay, K. N. (2015). Resolving ambiguities of mvpa using explicit models of representation. *Trends in Cognitive Sciences*, 19, 551–554. <https://doi.org/10.1016/j.tics.2015.07.005>.
- Nelson, W. W., & Loftus, G. R. (1980). The functional visual field during picture viewing. *Journal of Experimental Psychology: Human Learning and Memory*, 6(4), 391–399.
- Oliva, A., & Torralba, A. (2006). Building the gist of a scene: The role of global image features in recognition. In *Progress in brain research* (Vol. 155). Elsevier. [https://doi.org/10.1016/S0079-6123\(06\)55002-2](https://doi.org/10.1016/S0079-6123(06)55002-2).
- Pelli, D. G., & Tillman, K. A. (2008). The uncrowded window of object recognition. *Nature Neuroscience*, 11, 1129–1135. <https://doi.org/10.1038/nn.2187>.
- Piper, L., Oliveira, A. L., & Marques, T. (2025). Explicitly modeling subcortical vision with a neuro-inspired front-end improves cnn robustness. *Advances in Neural Information Processing Systems (NeurIPS)*.
- Ponce, C. R., Hartmann, T. S., & Livingstone, M. S. (2017). End-stopping predicts curvature tuning along the ventral stream. *Journal of Neuroscience*, 37, 648–659. <https://doi.org/10.1523/JNEUROSCI.2507-16.2016>.
- Prince, J. S., Charest, I., Kurzwaski, J. W., Pyles, J. A., Tarr, M. J., & Kay, K. N. (2022). Improving the accuracy of single-trial fmri response estimates using glmsingle. *eLife*, 11. <https://doi.org/10.7554/ELIFE.77599>.
- Rademaker, R. L., Chunharas, C., & Serences, J. T. (2019). Coexisting representations of sensory and mnemonic information in human visual cortex. *Nature Neuroscience*, 22, 1336–1344.
- Rosenholtz, R. (2016). Capabilities and limitations of peripheral vision. *Annual Review of Vision Science*, 2, 437–457. <https://doi.org/10.1146/annurev-vision-082114-035733>.
- Rosenholtz, R. (2019). Challenges to pooling models of peripheral vision: Implications for texture tiling models. *Journal of Vision*, 19(7), 15. <https://doi.org/10.1167/19.7.15>.
- Rosenholtz, R., Huang, J., Raj, A., Balas, B. J., & Ilie, L. (2012). A summary statistic representation in peripheral vision explains visual search. *Journal of Vision*, 12(4), 14. <https://doi.org/10.1167/12.4.14>.
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A. C., & Fei-Fei, L. (2015). ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3), 211–252. <https://doi.org/10.1007/s11263-015-0816-y>.
- Sechidis, K., Tsoumakas, G., & Vlahavas, I. (2011). On the stratification of multi-label data. *Machine Learning and Knowledge Discovery in Databases (ECML PKDD 2011)*, 6913, 145–158.
- Sprague, T. C., & Serences, J. T. (2013). Attention modulates spatial priority maps in the human occipital, parietal and frontal cortices. *Nature Neuroscience*, 16, 1879–1887.
- Srihasam, K., Vincent, J. L., & Livingstone, M. S. (2014). Novel domain formation reveals proto-architecture in inferotemporal cortex. *Nature Neuroscience*, 17, 1776–1783. <https://doi.org/10.1038/nn.3855>.
- Tian, Y., Krishnan, D., & Isola, P. (2019). Contrastive multiview coding. *CoRR*, abs/1906.05849. <http://arxiv.org/abs/1906.05849>
- Vann, S. D., Aggleton, J. P., & Maguire, E. A. (2009). What does the retrosplenial cortex do? *Nature Reviews Neuroscience*, 10, 792–802. <https://doi.org/10.1038/nrn2733>.
- Wang, B., Mayo, D., Deza, A., Barbu, A., & Cadena, C. (2021). On the use of cortical magnification and saccades as biological proxies for data augmentation. *3rd Workshop on Shared Visual Representations in Human and Machine Intelligence (SVRHM) at NeurIPS*. <https://arxiv.org/abs/2112.07173>
- Wang, P., & Cottrell, G. W. (2017). Central and peripheral vision for scene recognition: A neurocomputational modeling exploration. *Journal of Vision*, 17(4), 9. <https://doi.org/10.1167/17.4.9>.
- Yue, X., Pourixadian, I. S., Tootell, R. B., & Ungerleider, L. G. (2014). Curvature-processing network in macaque visual cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 111, E3467. <https://doi.org/10.1073/pnas.1412616111>.
- Yue, X., Robert, S., & Ungerleider, L. G. (2020). Curvature processing in human visual cortical areas. *NeuroImage*, 222, 117295. <https://doi.org/10.1016/j.neuroimage.2020.117295>.
- Zhang, X., Gao, X., He, L., & Lu, W. (2019). A gated peripheral-foveal convolutional neural network for unified image aesthetic prediction. *IEEE Transactions on Multimedia*, 21(9), 2815–2826. <https://doi.org/10.1109/TMM.2019.2895346>.

Zhou, B., Lapedriza, A., Khosla, A., Oliva, A., & Torralba, A. (2017). Places: A 10 million image database for scene recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.

Supp. Material

Supp. Methods

Frame Sampling and Processing. A total of 717 sessions were retrieved from the VEDB via Databrary; 203 sessions without synchronized gaze data were excluded, yielding 514 sessions retained for processing and analysis. For each retained session, a task-relevant analysis window was defined from session logs by excluding non-task periods (e.g., setup, calibration, validation) via keyword filtering and using logged frame indices to identify task onset and offset. The window began at the final calibration segment immediately preceding task onset when available, and otherwise at the first task frame, and ended at the final task frame. Within this window, frames were sampled uniformly every 2 seconds (step size = $2 \times \text{FPS}$ at the native 25 FPS), with a maximum of 1000 frames per session. Sampled frames were decoded, converted to RGB, and processed using a deterministic pipeline consisting of a bicubic resize to 256 px followed by a 224×224 center crop. For downstream metamer generation, we additionally retained a bicubic-upsampled 512×512 version of each frame.

Gaze Retrieval and Processing. Per-frame gaze positions were derived from the VEDB’s eye-tracking recordings. First, gaze timestamps were aligned to video-frame indices. For each session, we loaded raw binocular gaze data (normalized x/y positions and per-sample confidence values) from the VEDB’s processed gaze files, along with temporal sync metadata (calibration and validation timestamp pairs) stored in the session’s marker files. A first-order polynomial was fit from eye-tracker timestamps to video timestamps using these sync points, and sessions with fewer than 4 sync points or a fit $R^2 < 0.9$ were excluded. Each gaze sample was then assigned to the nearest video frame index based on the session’s native frame rate. The resulting per-sample records (frame index, eye identity, normalized gaze position, and confidence) were stored as one Parquet file per session.

Second, using the frame-aligned Parquet files from the previous step, per-frame fixation locations were computed for the foveal crop construction. For each sampled video frame, all raw gaze samples from both eyes assigned to that frame were pooled and filtered by a confidence threshold (≥ 0.30). The per-frame gaze position was then computed as the confidence-weighted centroid over the surviving samples (e.g., $\bar{x} = \sum_i w_i x_i / \sum_i w_i$, where w_i is the sample confidence). The resulting fixation coordinates, crop radius, and a flag indicating whether the gaze-based centroid or a center fallback was used, were saved as per-session CSV files at both the 224×224 and 512×512 image sizes. These fixation metadata were subsequently used for constructing the Fovea-Gaze, Periph, and Periph-NF conditions (see *Methods*).

Dataset Splitting Procedure. All dataset splits were performed at the video session level to prevent information leakage from highly correlated frames within the same recording (e.g., shared environment and temporally adjacent frames). Each sampled frame inherits a task label. For SimCLR self-supervised pretraining (Chen et al., 2020), we used the full set of sessions with synchronized gaze (514 sessions; 433,564 frames) and created an approximately 80/10/10 session split (455/28/31 train/val/test sessions) using stratified splitting to ensure balanced representation of task labels across splits. This yielded 377,462 training frames (87.06%), 26,026 validation frames (6.00%), and 30,076 test frames (6.94%). Validation and test sessions were held out from all representation learning and hyperparameter selection. For downstream in-domain evaluation, we restricted the task label space to a core subset of 17 tasks meeting inclusion criteria of at least 500 frames per task and at least 7 sessions per task, yielding 298 eligible sessions (167,437 frames). These sessions were partitioned into 239/28/31 train/val/test sessions using multilabel-stratified shuffle splitting (iterative stratification) (Sechidis et al., 2011) to preserve marginal label frequencies while maintaining strict session disjointness. This core-17 split was nested within the original session split, such that sessions retained their original train/val/test assignment after subsetting. The resulting frame counts were 116,540 train frames (69.60%), 23,626 validation frames (14.11%), and 27,271 test frames (16.29%), with no sessions spanning multiple splits. For VGGFace2, the official test split contains identities disjoint from the training identities and is therefore incompatible with our closed-set identification setup. Therefore, using the official train split we created a stratified 80/20 train/validation split within each identity using a fixed random seed, enforcing at least one image in each split for identities with ≥ 2 samples; split indices were cached to disk and reused across reruns to ensure identical partitions. Places365 underwent the same splitting protocol with a stratified 80/20 train/val split within each scene class using a fixed random seed, enforcing at least one image per split for classes with ≥ 2 samples.

NeuroFovea Metamer Generation. To approximate peripheral texture-tiling style representations, we generated gaze-contingent peripheral metamers using the NeuroFovea (NF) foveated style-transfer framework (Deza et al., 2019). NeuroFovea implements a feed-forward pipeline that (i) encodes an image with a VGG-based feature extractor, (ii) applies eccentricity-dependent pooling in feature space, and (iii) decodes the pooled representation back to pixel space, producing a texture-preserving metamer-like rendering. Starting from the 512×512 session-frame crops, we applied NF independently to each frame using scale parameter $s = 0.4$ (within the model’s standard scaling regime). Pooling was centered at the per-frame gaze-defined fixation location; when fix-

ation metadata were unavailable, we defaulted to the image center (see *Supp. Methods* "Gaze Retrieval and Processing"). NF outputs were saved as PNGs in session-wise directories and subsequently resized to 224×224 for peripheral-view construction and SimCLR pretraining.

Statistical testing of encoding model performance.

To evaluate the significance of differences between conditions in fMRI encoding model performance, we used a repeated measures ANOVA test (2-way), with p -values obtained using a non-parametric shuffling method. This was done by shuffling the values within each participant individually, which preserves the repeated measures structure of the data, and computing F -statistics for each effect using the shuffled data, over 10,000 iterations. To obtain a p -value for each effect, we computed the proportion of iterations where the true F -statistic was less than or equal to the shuffled F -statistic. This non-parametric method is appropriate for relatively small sample sizes because it is more conservative than standard parametric p -values, and is similar to previous work (Henderson et al., 2022, 2025; Rademaker et al., 2019; Sprague & Serences, 2013).

We also performed paired comparisons for differences between conditions, within each ROI, using a paired t -test with permutation. To shuffle the data, we randomly swapped (i.e., swapped with 50% probability) the values corresponding to the two conditions for each participant, and then computed a t -statistic for this shuffled data. This was done over each of 10,000 permutations. We obtained a two-tailed p -value calculated based on the minimum of the proportions of permuted t -statistics greater or less than the true t -statistic, across iterations, multiplied by 2.

Supp. Table 1: Linear probe hyperparameters by task. CE = cross-entropy; LS = label smoothing; pat = early-stopping patience (epochs); RLP = ReduceLROnPlateau (monitor val loss; factor = 0.1; patience = 5). In-domain used cosine LR decay with linear warmup (5 epochs). Transfer probes used early stopping (pat=10) and RLP.

Hyperparam.	In-Domain	VGGFace2	Places365
Backbone (frozen)	ResNet-18	ResNet-18	ResNet-18
Input size	224×224	224×224	224×224
Optimizer	AdamW	AdamW	AdamW
Batch size	256	60	60
Learning rate	1×10^{-3}	1×10^{-3}	1×10^{-3}
Weight decay	1×10^{-4}	1×10^{-2}	1×10^{-2}
Loss	CE	CE (+ LS)	CE (+ LS)
Label smoothing	0.0	0.1	0.1
Probe epochs	20	20	20
BN during probe	eval	train	train
Early stopping	off	pat=10	pat=10
Checkpoint select	Macro-F1 (val)	Top-1 (val)	Top-1 (val)
LR schedule	cosine + warmup	RLP	RLP
Warmup epochs	5	–	–
FP16	True	True	True
Seed	1337	1337	1337

Supp. Table 2: SimCLR pretraining hyperparameters. Adam = Adam optimizer; NT-Xent = normalized temperature-scaled cross-entropy loss; τ = softmax temperature; p = probability of applying augmentation; σ = Gaussian blur strength sampled uniformly from $U(0.1, 2.0)$; blur kernel size = $0.1 \times$ image size; h-flip = random horizontal flip; FP16 = mixed-precision training.

Hyperparam.	SimCLR
Backbone	ResNet-18
Input size	224×224
Optimizer	Adam
Batch size	64
Learning rate	3×10^{-4}
Weight decay	1×10^{-4}
Loss	NT-Xent
Temperature τ	0.07
Epochs	120
LR schedule	fixed
Proj. head	Linear(512,512) $\rightarrow$ ReLU $\rightarrow$ Linear(512,128)
Augmentations	Random crop, h-flip, color jitter ($p=0.8$), grayscale ($p=0.2$), Gaussian blur ($\sigma \sim U(0.1, 2.0)$)
FP16	True


Supp. Table 3: Comparing encoding model R^2 values across each pair of conditions, for each ROI. Average R^2 values for each ROI (across 8 participants) were compared across conditions using a paired t-test (df = 7) with permutation, see *Supp. Methods*. P-values < 0.01 are shown in bold.

	Baseline vs. Fovea-Gaze		Baseline vs. Periph		Baseline vs. Periph-NF		Fovea-Gaze vs. Periph	
	t	p	t	p	t	p	t	p
V1	1.63	0.1452	5.72	0.0088	2.16	0.0538	7.04	0.0062
V2	5.25	0.0082	7.15	0.0088	4.38	0.0084	3.39	0.0124
V3	6.89	0.0112	10.57	0.0080	4.78	0.0090	2.80	0.0340
hV4	5.71	0.0080	6.97	0.0066	5.87	0.0086	1.59	0.1638
OFA	6.67	0.0090	6.45	0.0074	6.14	0.0108	1.31	0.2054
FFA	6.45	0.0064	6.90	0.0096	6.32	0.0086	1.05	0.3040
OWFA	11.18	0.0050	4.94	0.0096	4.86	0.0078	0.75	0.4480
VWFA	5.45	0.0098	4.04	0.0158	4.72	0.0082	-0.12	0.9344
OPA	7.86	0.0088	6.03	0.0086	8.25	0.0090	-0.99	0.3130
PPA	9.40	0.0084	5.18	0.0090	8.39	0.0060	-4.86	0.0074
RSC	6.80	0.0090	0.72	0.5054	2.39	0.0550	-4.91	0.0070


	Fovea-Gaze vs. Periph-NF		Periph vs. Periph-NF	
	t	p	t	p
V1	0.78	0.4580	-5.62	0.0074
V2	0.61	0.5154	-3.04	0.0232
V3	0.53	0.6186	-1.73	0.1284
hV4	-0.95	0.3968	-2.68	0.0260
OFA	0.71	0.4862	-0.25	0.8086
FFA	0.98	0.3424	-0.07	0.9492
OWFA	0.87	0.4350	0.15	0.8876
VWFA	2.31	0.0344	1.98	0.0970
OPA	0.77	0.4432	3.76	0.0066
PPA	-0.46	0.6086	2.97	0.0084
RSC	-3.37	0.0210	1.12	0.2682

Supp. Table 4: Full confidence intervals for in-domain classification $\Delta F1$ (corresponding to Table 2 in the main text). Classes are ranked by $\max_j |\Delta_j|$, where Δ_j is the F1 difference between two conditions. Positive values indicate condition A > B. 95% confidence intervals were estimated using a 1000-replicate parametric (multinomial) bootstrap over the aggregate confusion matrices, with row-wise resampling conditional on observed class counts. Asterisks indicate comparisons whose 95% CI does not overlap zero. See Supp. Figure 2 for examples of each label and Supp. Figure 1 for bar-plot visualization of $\Delta F1$.


Rank	Label	Fovea > Base	95% CI		Fovea > Periph-NF	95% CI		Fovea > Periph	95% CI	
1	playing video game	0.433	[0.399, 0.466]	*	0.433	[0.401, 0.464]	*	0.433	[0.399, 0.467]	*
2	solving rubik's cube	-0.116	[-0.185, -0.051]	*	0.428	[0.362, 0.491]	*	0.018	[-0.049, 0.088]	*
3	socializing	0.365	[0.328, 0.402]	*	0.318	[0.281, 0.358]	*	0.336	[0.301, 0.370]	*
4	computer work	0.079	[0.046, 0.116]	*	0.300	[0.274, 0.326]	*	0.281	[0.253, 0.309]	*
5	skateboarding	-0.292	[-0.323, -0.261]	*	-0.170	[-0.203, -0.139]	*	-0.282	[-0.314, -0.253]	*
6	playing frisbee	-0.251	[-0.286, -0.215]	*	-0.214	[-0.248, -0.181]	*	-0.263	[-0.295, -0.229]	*
7	washing dishes	-0.263	[-0.299, -0.224]	*	-0.100	[-0.144, -0.062]	*	-0.075	[-0.114, -0.035]	*
8	cooking	-0.221	[-0.244, -0.200]	*	0.029	[0.008, 0.051]	*	-0.078	[-0.098, -0.054]	*
9	playing pool	-0.200	[-0.226, -0.178]	*	-0.178	[-0.201, -0.154]	*	-0.194	[-0.218, -0.170]	*
10	playing cards	0.178	[0.151, 0.206]	*	0.147	[0.117, 0.178]	*	0.178	[0.151, 0.207]	*
11	exploring	0.023	[-0.040, 0.086]		0.124	[0.072, 0.176]	*	0.077	[0.023, 0.129]	*
12	playing poker	-0.116	[-0.138, -0.094]	*	0.053	[0.027, 0.083]	*	-0.055	[-0.078, -0.032]	*
13	preparing	-0.038	[-0.074, 0.001]		-0.041	[-0.075, -0.002]	*	-0.094	[-0.140, -0.048]	*
14	playing sudoku	-0.071	[-0.094, -0.049]	*	-0.048	[-0.072, -0.025]	*	0.010	[-0.009, 0.029]	
15	walking	-0.040	[-0.046, -0.035]	*	-0.028	[-0.034, -0.022]	*	-0.030	[-0.036, -0.024]	*
16	playing piano	0.002	[0.000, 0.006]		0.002	[0.000, 0.006]		0.002	[0.000, 0.006]	
17	standing	0.000	[0.000, 0.000]		0.000	[0.000, 0.000]		0.000	[0.000, 0.000]	




Supp. Figure 1: Class-wise F1 differences ($\Delta F1$) between the Fovea condition and each comparison condition, where each panel reports $\Delta F1 = F1_{\text{Fovea}} - F1_{\text{comparison}}$ (top: Fovea–Base; middle: Fovea–Periph-NF; bottom: Fovea–Periph). Error bars indicate 95% confidence intervals from a 1000-replicate parametric (multinomial) bootstrap computed from the aggregate confusion matrices (row-wise resampling conditional on observed class counts). Positive values indicate higher F1 for Fovea; negative values indicate higher F1 for the comparison condition.



Supp. Figure 2: Frame examples for the 17 task categories.



Supp. Figure 3: Row-normalized confusion matrices (rows sum to 1) for the in-domain 17-class task across conditions (Baseline, Fovea-Gaze, Periph-NF, Periph). Axes show true labels (y) and predicted labels (x).



Supp. Figure 4: fMRI encoding model performance for the Baseline VEDB model condition, plotted as a function of voxelwise noise ceiling. Each panel plots voxels from a different functionally-defined ROI, each point is a voxel, colors indicate voxels from different participants. Voxelwise noise ceiling is computed with respect to beta weights averaged over 3 repetitions of each image in the NSD dataset, see Allen et al., 2022 for details.