
PLACE: Patch-Level Agnostic Concept Extraction

Gabriele Onorato^{1, 2*} Fabrizio Silvestri²
Nathaniel D. Bastian³ Francesco Restuccia¹

¹Department of Electrical and Computer Engineering, Northeastern University

²Department of Computer, Control and Management Engineering, Sapienza University of Rome

³Department of Electrical Engineering and Computer Science, United States Military Academy

Abstract

While Vision Foundation Models (VFM) are increasingly used in several computer vision tasks, their internal representations remain substantially opaque. This mostly stems from *polysemanticity*, i.e., individual neurons encode mixtures of unrelated visual patterns. In stark contrast, human reasoning is inherently *monosemantic*, i.e., it naturally relies on distinct, isolated features to process information. Therefore, monosemanticity needs to be implemented within the model’s activation space to obtain representations that are human-understandable *concepts*. Existing methods mostly operate on convolutional networks, so when applied to VFMs their "crop-and-resize" approach distorts the model’s representations by disrupting global self-attention mechanisms and discarding spatial geometry. Furthermore, prior work requires downstream classification labels and is based on KL-divergence, which requires to propagate the gradients back from the classifier head. Ultimately incurring in excessive concept extraction time, and making it hardly applicable to extract task-agnostic concepts. To overcome these issues, we propose PLACE, a *fully unsupervised* concept extraction framework that operates directly on the native patch-token activations of a VFM. By introducing a geometric Gram-matrix alignment loss, PLACE mechanistically controls downstream KL divergence, thus enforcing *faithfulness* – i.e., the notion that extracted concepts, although obtained in an unsupervised fashion, are relevant to downstream classification tasks. Evaluations on image classification tasks (ImageNet and COCO) on DINOv2, ViT-MAE and ViT-B/16 show that PLACE ❶ using Gram-alignment it extracts concepts about $20\times$ faster compared to KL-supervised extraction and $11\times$ faster than Sparse Autoencoders (SAEs). Moreover, PLACE yields highly sparse concepts that ❷ achieve up to 40 percentage points sparser compared to the state of the art (i.e., max 0.933 Gini sparsity score), and ❸ are up to **92.3%** similar to those obtained under classifier supervision (i.e., 0.923 cosine similarity).

1 Introduction

Vision Foundation Models (VFMs) [23] are currently seen as “black boxes”, i.e., mostly lacking human-understandable representations [14, 46, 61]. This opacity has become a barrier to the safe deployment of VFMs into real-world scenarios such as AI-guided surgeries or medical imaging [3, 4, 44], where relying on an unexplainable decision process poses unacceptable risks to human life [61]. A VFM is difficult for humans to understand because it compresses the vast amount of visual information seen during training into its activation space. This results in *polysemanticity* [41], i.e., a single neuron fires in response to a mixture of unrelated visual patterns, which ultimately obscures their decision-making process [57]. This behavior stands in stark contrast to the *monosemantic* nature of human reasoning; when perceiving and describing the visual world, humans naturally

*Corresponding author: onorato.g@northeastern.edu

rely on distinct, isolated features [8]. Therefore, to make VFMs more interpretable and human-understandable, it is necessary to enforce *monosemanticity* within their internal representations by aligning the model’s computations with human logic.

To achieve this alignment, the literature has defined the notion of *concept*, i.e., a *monosemantic* direction within the model’s latent activation space [7, 24–26, 35]. While a monosemantic direction might visually align with a human-understandable concept, *it does not guarantee that the model actively relies on it to execute its task*. For example, if the model is performing classification, a concept extraction method might identify a *rabbit ear* as a relevant concept, whereas the model could actually be using the *rabbit fur* to classify the image correctly. Without a mechanism to verify this alignment, extracted concepts risk being misleading, offering plausible but ultimately unreliable explanations of the model’s behavior. As such, *faithfulness* addresses this risk by ensuring the concept influences the model’s predictions (a formal definition of concept and faithfulness are provided in Section 3).

The key issue is that most prior methods were designed for convolutional architectures (CNNs) [25, 26, 35, 70]. Current attempts to interpret Vision Transformers (ViTs) rely on SAEs [33, 57, 60], architectural modifications [5, 9–11, 36, 50, 68], or external language models [15, 32, 40]. *These approaches share a critical issue: they require training auxiliary modules or fine-tuning the model’s weights*. However, imposing a retraining phase in VFMs is computationally prohibitive and may degrade their generalization capabilities [39]. Consequently, extracting faithful concepts without relying on retraining with task-specific labels is a compelling necessity.

There are several technical challenges to achieve this goal. First, VFMs do not possess a classifier head, which removes the explicit task-driven signal that *all existing methods* rely on to enforce faithfulness [7]. Second, existing extraction methods rely on cropping and resizing images [7, 25, 70]. While effective for convolutional architectures, this approach disrupts the global self-attention mechanisms of ViTs [58], biasing concept extraction (formalized in Section 3). Finally, since prior approaches extract concepts from globally aggregated representations, they inherently discard the image’s spatial geometry. Consequently, prior work is forced to rely on gradient-based attribution techniques to locate and identify concepts within the original input [25].



Figure 1: **Example of PLACE applied to a picture of a church.** While current literature works on CNNs and yields image crops as the extracted concepts, PLACE returns a fine-grained heatmap localizing the relevant concepts in the whole image.

To address these issues, we propose PLACE, a *fully unsupervised* concept extraction framework exemplified in Figure 1 and specifically designed for ViTs and VFMs [13, 30, 56, 63]. In opposition with state of the art work CRAFT [25] and FACE [7], PLACE operates directly on patch-token activations of a frozen backbone and replaces classifier-based faithfulness supervision with a Gram-matrix alignment loss. This approach allows PLACE to overcome the limitations of prior work and yield spatially resolved concept heatmaps without the need for architectural modifications or class labels while achieving highly stable concept extraction (up to 0.95 W-Stab, see Appendix E.7) without demanding downstream labels.

Summary of Novel Contributions:

- **Measurable Interpretability & Efficiency:** PLACE achieves strict concept sparsity – quantitatively measured via a high Gini sparsity of 0.933 – while accelerating the extraction pipeline by $11\times$ compared to SAEs and over $20\times$ compared to KL-supervised alternatives (Appendix E.9).
- **Label-Free Faithfulness via Geometric Alignment:** We introduce an unsupervised Gram-matrix alignment loss and show that (Section 3.1) that reconstruction error and Gram deviation are jointly sufficient to control downstream (D_{KL}) for any linear classifier, thus grounding faithfulness without labels.

- **Supervised Alignment without Labels:** We demonstrate via RSVC [37] that our extracted concepts tightly align (up to 0.923 cosine similarity) with those extracted under full classifier supervision.
- **Comprehensive Benchmarking:** We systematically evaluate PLACE against state of the art concept extraction baselines CRAFT [25] and FACE [7], which were adapted for sliding-window ViTs [37], and SAEs [33, 57, 60]. Our experimental evaluation indicates that PLACE avoids the representational distortion of sliding-window methods (Section 4.2) and thus reduces reconstruction error (MSE) by over 10x (0.092 vs. 1.015) while delivering high-quality concept heatmaps.

2 Related Work

Concept-based explanations. TCAV [35] pioneered concept-based interpretability by training linear classifiers on activation vectors of user-defined concepts and random counterexamples to measure a classifier’s sensitivity to specific concept directions. However, TCAV relies on manually curated concept datasets, limiting scalability and introducing human bias. ACE [26] addresses this by segmenting images into superpixels via SLIC [1] and clustering segment-level activations to discover concepts automatically, at the cost of potential segmentation artifacts. ICE [70] replaces clustering with Non-negative Matrix Factorization (NMF) [42], offering cleaner concept dictionaries but operating at the level of convolutional kernels only. CRAFT [25] improves on ICE with a recursive factorization pipeline and Sobol-based relevance scoring [64], but does not enforce predictive consistency through the classifier. FACE [7] introduced a KL-regularized NMF framework that explicitly aligns discovered concepts with the classifier’s predictions, at the cost of requiring class-label supervision. Concept Bottleneck Models (CBMs) [36] force predictions through an explicit concept layer, but require concept annotations at training time and architectural changes to the model. Concept Embedding Models [68] mitigate the predictive performance drop of standard CBMs by learning high-dimensional concept embeddings instead of scalar activations, though they still require architectural modifications and concept supervision during training.

Achtibat et al. [2] proposed Concept Relevance Propagation (CRP) and Relevance Maximization (RelMax). CRP propagates relevance scores to concept subspaces, identifying both *where* and *what* concepts contribute to individual predictions, while RelMax visualizes concepts by selecting representative training examples that maximize feature relevance during inference. However, CRP relies on layer-wise architecture-specific relevance propagation rules and RelMax requires access to the training data, which for VFMs like the DINO family [13, 56, 63] or Masked Autoencoders (ViT-MAE) [30] is typically either inaccessible or too large to navigate meaningfully. Sparse autoencoders (SAEs) [33, 57], instead, learn interpretable directions from model activations without label supervision. Recently, Gated SAEs [60] addressed an issue of SAEs, where the ℓ_1 penalty required to enforce sparsity inherently shrinks the magnitude of active features, degrading reconstruction fidelity. However, all SAE-based approaches require training a separate encoder-decoder architecture, thus adding significant computational overhead. A separate line of work leverages vision-language models (VLMs) to bypass manual concept curation entirely: LG-CAV [32] uses a VLM to score a shared pool of probe images according to textual concept descriptions, and uses those scores as supervision to learn concept directions in the target model without labeled concept images. CE-FAM [40] and SSG-CAM [15] learn to mimic VLM embeddings through a branched network with gradient-weighted activation-map fusion. While effective, these methods risk hallucinating explanations, directly inheriting the flaws of the external language models they rely on [43].

Concepts in Vision Transformers. Conversely from CNN feature maps, ViTs encode information through patch-token sequences and self-attention[23]. Dorszewski et al. [22] provide a layer-wise analysis of ViT concepts by labeling individual neurons via CLIP-Dissect [55], showing that ViTs develop a coarse-to-fine concept hierarchy despite lacking CNN inductive biases. B-cos ViTs [9–11] reformulate every ViT component as a dynamic linear mapping so that predictions can be faithfully summarized by a single linear transform, yet require retraining the model. Arya et al. [5] fine-tune existing pre-trained models into inherently interpretable networks. ProtoViT [50] extends prototype-based reasoning to ViT backbones with spatially deformable prototypes, but the prototype layer itself is an architectural addition trained jointly with the model. In contrast, *PLACE operates on the native patch-token activations of a frozen ViT*. This inherently preserves spatial structure, yielding high-quality heatmaps without introducing architectural overhead. Furthermore, we replace task-specific classifier supervision with an unsupervised Gram-matrix alignment. This allows us to guide the extraction toward faithful concepts without labels or external models.

3 The PLACE Framework

Prior concept extraction methods [7, 25, 70] assume a supervised classification setting in which a model $f : \mathcal{X} \rightarrow \mathcal{Y}$ decomposes into an encoder $g : \mathcal{X} \rightarrow \mathcal{G}$ and a classifier head $h : \mathcal{G} \rightarrow \mathcal{Y}$, so that $f(\mathbf{x}) = (h \circ g)(\mathbf{x})$. The classifier head (i) provides class-label supervision for concept discovery [7] and (ii) defines the output distribution against which faithfulness can be measured [25, 35]. However VFMs break this assumption; models such as DINO [13, 56, 63] or ViT-MAE [30] produce general-purpose embeddings $f : \mathcal{X} \rightarrow \mathcal{Z}$ with no classification head attached. As such, there is no h to supervise concept extraction, no class-label signal to enforce predictive consistency, and the representation space $\mathcal{Z}$ is not organized around decision boundaries. Consequently, the KL-regularized NMF objective of FACE [7], which aligns the factorized activations with the classifier’s output distribution, can be hardly applied to ViTs.

Effect of crop-and-resize on ViTs. CRAFT [25], FACE [7], and their ViT adaptations [37] extract concepts by cropping sub-regions from an image, resizing each crop to the model’s input resolution, feeding it to the model, and factorizing the resulting activation vectors. For CNNs, this is justified by translation equivariance [16, 69]. ViTs, however, process information in a different way [23, 27, 54, 58, 67]. This difference in input can be seen in Figure 2, and we discuss this mismatch formally in Appendix A. Due to this, our patch-token extraction and the sliding-window baseline are not directly comparable on classifier-dependent metrics (C-Ins, C-Del, D_{KL}), since they operate on activation spaces of different scale and address different input scenarios; we return to this in Section 4.2.



Figure 2: **The Crop-and-Resize Problem in ViTs.** While PLACE operates on the full-resolution image seen at training time (left), current crop-and-resize methods feed sub-resolution patches (right). We show in Section 4.2 this shift from the training distribution biases concept extraction.

Concept Extraction. We formulate concept extraction as a geometric problem on the backbone’s internal representations. Let f be a ViT with L layers. As illustrated in Figure 3, we split f at a chosen layer ℓ into a prefix $g_{\leq \ell} : \mathcal{X} \rightarrow \mathbb{R}^{N \times d}$, which maps an input image to a sequence of N patch-token activations of dimension d , and a suffix $g_{> \ell}$, comprising the remaining $L - \ell$ transformer blocks. Given a dataset of n images, the prefix produces an activation tensor $\mathcal{A} \in \mathbb{R}^{n \times N \times d}$. Unlike CNNs, where a global average pool reduces spatial dimensions before factorization [7, 25], the patch-token structure of ViTs preserves a regular spatial grid (e.g. 14×14 for a 224×224 input with patch size 16), which we exploit directly for concept localization.

To decompose the activation space into human-understandable components, we reshape the patch-token tensor $\mathcal{A}$ into an activation matrix $\mathbf{A} \in \mathbb{R}^{nN \times d}$, where each row corresponds to a single patch token. Since activations of transformer models (e.g., GELU, LayerNorm) contain negative values, we employ Semi-NMF [21]:

$$(\mathbf{U}, \mathbf{W}) = \arg \min_{\mathbf{U} \geq 0, \mathbf{W}} \frac{1}{2} \|\mathbf{A} - \mathbf{U}\mathbf{W}^\top\|_F^2 \quad (1)$$

where $\mathbf{W} \in \mathbb{R}^{d \times r}$ is the concept bank and $\mathbf{U} \in \mathbb{R}_{\geq 0}^{nN \times r}$ the coefficient matrix. This factorization naturally operationalizes our core interpretability primitives:

- **Concept:** A basis vector $\mathbf{w}_j \in \mathbb{R}^d$ (the j -th column of $\mathbf{W}$) representing a consistent, monosemantic direction in the latent representation space. The corresponding column $\mathbf{u}_j \in \mathbb{R}_{\geq 0}^{nN}$ encodes the magnitude of this concept across all tokens, i.e., *how intensely* the concept $\mathbf{w}_j$ appears in the scene.
- **Faithfulness:** A concept bank $\mathbf{W}$ is considered faithful if the reconstructed activations $\hat{\mathbf{A}} = \mathbf{U}\mathbf{W}^\top$ preserve the *predictive consistency* of the suffix $g_{> \ell}$ relative to the original activations $\mathbf{A}$.

The non-negativity constraint on $\mathbf{U}$ ensures that each patch is represented strictly as an *additive* combination of concepts, enabling the parts-based interpretation necessary for visual grounding.




Figure 3: **PLACE Framework.** An input image is passed through the transformer prefix $g_{\leq \ell}$ to extract patch-token activations $\mathbf{A}$. To control granularity and computational cost, adjacent tokens are spatially grouped. The grouped representations are then dimensionally flattened and decomposed via Semi-NMF to discover a shared Concept bank and non-negative Concept coefficients. These coefficients are then upsampled to the original resolution to yield fine-grained concept heatmaps.

Token Grouping. Dense token grids (e.g., 37×37 for DINOv2) produce a large activation matrix $\mathbf{A} \in \mathbb{R}^{nN \times d}$, making Semi-NMF optimization expensive. To control this cost, we introduce a spatial grouping factor s : before factorization, we apply non-overlapping average pooling over $s \times s$ blocks of adjacent patch tokens, reducing the token count per image from N to $\lfloor \sqrt{N}/s \rfloor^2$. After optimization, the learned coefficients are mapped back to the original grid via nearest-neighbor upsampling (repeat_interleave). Setting $s=2$ accelerates optimization by $\sim 3\times$ (Appendix E.5). Detailed spatial arithmetic and grid boundary handling are described in Appendix B.

3.1 Geometric Faithfulness

Without a classifier head h , faithfulness cannot be measured as the divergence between predictive distributions. Instead, we preserve the *geometric structure* of activations $\mathbf{U}\mathbf{W}^\top = \hat{\mathbf{A}}$ to ensure predictive consistency. We optimize $\mathbf{U}$ and $\mathbf{W}$ by minimizing the joint objective:

$$\mathcal{L} = \underbrace{\frac{1}{2} \|\mathbf{A} - \hat{\mathbf{A}}\|_F^2}_{\mathcal{L}_{\text{rec}}} + \lambda \underbrace{\frac{\|\mathbf{A}\mathbf{A}^\top - \hat{\mathbf{A}}\hat{\mathbf{A}}^\top\|_F^2}{\|\mathbf{A}\mathbf{A}^\top\|_F^2}}_{\mathcal{L}_{\text{gram}}} \quad (2)$$

where $\mathcal{L}_{\text{rec}}$ is the reconstruction error (MSE) and $\mathcal{L}_{\text{gram}}$ aligns the global correlation structure $\mathbf{A}\mathbf{A}^\top$, capturing the pairwise inner products among all tokens. This metric is made invariant to the scale of representations by normalizing by $\|\mathbf{A}\mathbf{A}^\top\|_F^2$. The following proposition formalizes why jointly minimizing these two terms suffices to control downstream predictive divergence.

Proposition 1 (Informal; proof in Appendix C). *Let $h(z) = \mathbf{W}_c z + b$ be any linear classifier mapping the CLS token z to class logits, and let $\delta_{\text{rec}} = \|\mathbf{A} - \hat{\mathbf{A}}\|_F$, $\delta_{\text{gram}} = \sqrt{\mathcal{L}_{\text{gram}}}$. Then:*

$$D_{\text{KL}}(\sigma(h \circ g_{>\ell}(\mathbf{A})) \parallel \sigma(h \circ g_{>\ell}(\hat{\mathbf{A}}))) \leq \frac{1}{2} \|\mathbf{W}_c\|_{\text{op}}^2 L_{\text{block}}^{2L_{\text{suf}}} \left(\sqrt{N} \delta_{\text{rec}} + \alpha \|\mathbf{A}\mathbf{A}^\top\|_F \cdot \delta_{\text{gram}} \right)^2 \quad (3)$$

where L_{block} is the per-block Lipschitz constant, $L_{\text{suf}} = L - \ell$ the number of suffix transformer blocks, α collects attention-projection constants, and $\|\cdot\|_{\text{op}}$, $\|\cdot\|_F$ denote the spectral and Frobenius norms, respectively.

The bound identifies $\mathcal{L}_{\text{rec}}$ and $\mathcal{L}_{\text{gram}}$ as the two sufficient control variables for downstream predictive divergence (the precise mechanism is detailed in Appendix C). Being worst-case, the bound compounds multiplicatively across the L_{suf} suffix blocks and is therefore loose as a numerical predictor; its value is structural rather than quantitative. Notably, the bound tightens as the extraction layer ℓ increases, since fewer suffix blocks reduce the compounding factor; at $\ell = L$ the composition term vanishes entirely (Appendix C). Empirical tightness is validated in Section 4.3. Moreover, optimizing $\mathcal{L}$ requires no forward pass and backpropagation through $g_{>\ell}$ or h , and no auxiliary network, requiring less computing than KL-supervised methods [7], which backpropagate through the full suffix at every step, and SAEs [33, 57, 60], which train a separate encoder-decoder (Appendix E.9).

3.2 Optimization and Visualization

We solve the objective in Equation (2) using AdamW [49], treating both $\mathbf{U}$ and $\mathbf{W}$ as free parameters updated via backpropagation. Unlike traditional Semi-NMF solvers [42], gradient-based optimization requires accommodating the non-convex geometric constraints introduced by $\mathcal{L}_{\text{gram}}$, therefore the Semi-NMF constraint $\mathbf{U} \geq 0$ is enforced by clamping negative entries to zero, and gradient clipping is performed to ensure stability [51]. To ensure convergence toward interpretable solutions, we apply an ℓ_1 regularization on $\mathbf{U}$ (quantified in Appendix D.2) to enhance concept sparsity and an ℓ_2 penalty on $\mathbf{W}$ to stabilize the dictionary. The initialization of $\mathbf{U}$ is partially performed via NNDSVD [12], which accelerates convergence compared to random starting points (ablation in Appendix E.1), while the concept bank is initialized via the Moore–Penrose pseudoinverse $\mathbf{W} \leftarrow \mathbf{U}^\dagger \mathbf{A}$.

Heatmaps. A key advantage of PLACE is that concept heatmaps are obtained with no additional computational cost. Since each entry of $\mathbf{u}_j$ is spatially anchored to a specific patch token, the distribution of coefficients provides a direct activation map. To visualize a concept, the coefficients are reshaped into the original $\sqrt{N} \times \sqrt{N}$ grid and upsampled via bicubic interpolation. This process delivers exact *spatial localization* without requiring gradient-based attribution or implicit differentiation [25], preserving the structural integrity of the model’s internal representations.

4 Experimental Results

Datasets. Following the evaluation protocol of FACE [7], we evaluate our method on matched subsets of ImageNet [20, 62] (10 classes: `cassette`, `chainsaw`, `church`, `english_springer`, `french_horn`, `garbage_truck`, `gas_pump`, `golf_ball`, `parachute`, `tench`; ≥ 800 images/class) and COCO [45] (5 classes: `hot_dog`, `teddy_bear`, `train`, `umbrella`, `zebra`).

Models and hyperparameters. We evaluate PLACE on DINOv2 [56] and ViT-MAE [30], alongside a supervised ViT-B/16 for parity. We compare against supervised baselines CRAFT [25] and FACE [7] on CNNs, and their ViT extensions [37], we train a surrogate linear probe to compute D_{KL} . For PLACE and SAE, we extract concepts from layer $\ell = 9$ and set the dictionary rank to $r = 25$ to align with FACE [7]. Hyperparameter configurations and detailed per-class ablations are provided in Appendices D and E. Crucially, evaluations confirm that PLACE passes the weight randomization sanity check (Appendix E.6), and maintains highly stable dictionaries (Appendix E.7) even under perturbations (Appendix E.13).

Metrics. Our primary evaluation relies on three metrics: KL-divergence for faithfulness, MSE for reconstruction quality, and the C-Gini index as a quantitative proxy for monosemanticity (i.e., concept sparsity). This choice aligns with our goal of task-agnostic evaluation, well suited for VFMs. Formal definitions for all metrics are deferred to Appendix F.

We additionally report C-Insertion and C-Deletion [7, 25], which measure confidence recovery as concept patches are progressively revealed or masked, respectively. These causal metrics require a classifier head and a downstream task (Appendix D.1), and we include them strictly for a more thorough comparison with prior work. We note, however, that they carry three well-known limitations: masking creates artifacts [28] (visualized in Figure 4), the sequential removal order confounds individual concept contributions [18], and the protocol implicitly assumes isolated single-concept reasoning, ignoring the multi-concept interactions typical of deep networks.

4.1 Benchmark Comparison

We focus on metrics aligned with our task-agnostic goal, organized by their dependency on a downstream classifier. First, two metrics require *no classifier*: MSE for reconstruction quality and C-

Table 1: Faithfulness, complexity, and reconstruction quality on IMAGENET and COCO. Prior art on ResNet-34 and sliding-window adaptations on ViTs are provided strictly as magnitude and methodological references. PLACE is evaluated on DINOv2 ViT-B/14, supervised ViT-B/16, and ViT-MAE B/16 ($r=25, \ell=9$). Higher is better: C-Ins, C-Del, C-Gini; lower: MSE, D_{KL} . Mean and standard deviation over 5 runs. **Bold values** indicate the best performance for a given backbone.

Dataset	Method	Backbone	C-Ins $\uparrow$	C-Del $\uparrow$	C-Gini $\uparrow$	MSE $\downarrow$	$D_{KL} \downarrow$
IMAGENET							
Context: Prior Art on Convolutional Networks							
	ICE [70]	ResNet-34	0.908 $\pm$ 0.034	0.484 $\pm$ 0.063	0.537 $\pm$ 0.071	0.296 $\pm$ 0.001	0.359 $\pm$ 0.004
	CRAFT [25]		0.932 $\pm$ 0.001	0.752 $\pm$ 0.031	0.835 $\pm$ 0.031	0.451 $\pm$ 0.004	0.240 $\pm$ 0.003
	FACE [7]		0.969 $\pm$ 0.010	0.891 $\pm$ 0.011	0.895 $\pm$ 0.001	0.497 $\pm$ 0.002	0.220 $\pm$ 0.002
Context: Sliding-Window Adaptations on ViTs							
	CRAFT adapt. [25, 37]	DINOv2	0.841 $\pm$ 0.009	0.710 $\pm$ 0.008	0.875 $\pm$ 0.004	0.938 $\pm$ 0.001	0.282 $\pm$ 0.001
	FACE adapt. [7, 37]		0.914 $\pm$ 0.004	0.829 $\pm$ 0.001	0.898 $\pm$ 0.006	1.015 $\pm$ 0.001	0.371 $\pm$ 0.002
	CRAFT adapt. [25, 37]	ViT-B/16	0.951 $\pm$ 0.003	0.812 $\pm$ 0.012	0.920 $\pm$ 0.002	3.489 $\pm$ 0.001	0.071 $\pm$ 0.001
	FACE adapt. [7, 37]		0.955 $\pm$ 0.005	0.825 $\pm$ 0.009	0.926 $\pm$ 0.001	3.510 $\pm$ 0.001	0.062 $\pm$ 0.001
Comparison: PLACE Framework vs. Sparse Autoencoders							
	Gated SAE [60]	DINOv2	0.472 $\pm$ 0.187	0.620 $\pm$ 0.163	0.724 $\pm$ 0.054	0.073$\pm$0.004	0.368$\pm$0.161
	SAE [33, 57]		0.652$\pm$0.212	0.537 $\pm$ 0.188	0.865 $\pm$ 0.012	0.079 $\pm$ 0.004	0.533 $\pm$ 0.235
	PLACE (reconstruction)		0.467 $\pm$ 0.004	0.895 $\pm$ 0.005	0.933$\pm$0.001	0.092 $\pm$ 0.001	1.205 $\pm$ 0.044
	PLACE (+Gram)		0.522 $\pm$ 0.014	0.914$\pm$0.003	0.921 $\pm$ 0.003	0.090 $\pm$ 0.001	1.162 $\pm$ 0.099
	PLACE (+KL)		0.509 $\pm$ 0.008	0.914$\pm$0.005	0.924 $\pm$ 0.001	0.092 $\pm$ 0.002	1.233 $\pm$ 0.046
	Gated SAE [60]	ViT-B/16	0.421 $\pm$ 0.262	0.411 $\pm$ 0.214	0.303 $\pm$ 0.051	0.273$\pm$0.017	0.868 $\pm$ 0.428
	SAE [33, 57]		0.528 $\pm$ 0.251	0.497$\pm$0.159	0.434 $\pm$ 0.051	0.284 $\pm$ 0.018	1.101 $\pm$ 0.538
	PLACE (reconstruction)		0.715 $\pm$ 0.005	0.392 $\pm$ 0.015	0.530 $\pm$ 0.003	0.292 $\pm$ 0.001	$< 10^{-3}$
	PLACE (+Gram)		0.701 $\pm$ 0.011	0.401 $\pm$ 0.009	0.614$\pm$0.003	0.296 $\pm$ 0.001	$< 10^{-3}$
	PLACE (+KL)		0.796$\pm$0.002	0.313 $\pm$ 0.011	0.437 $\pm$ 0.003	0.304 $\pm$ 0.001	0.016 $\pm$ 0.002
	Gated SAE [60]	ViT-MAE	0.734 $\pm$ 0.195	0.477 $\pm$ 0.104	0.276 $\pm$ 0.054	0.980 $\pm$ 0.093	0.285$\pm$0.135
	SAE [33, 57]		0.937 $\pm$ 0.090	0.267 $\pm$ 0.183	0.238 $\pm$ 0.042	0.978$\pm$0.093	0.306 $\pm$ 0.144
	PLACE (reconstruction)		0.911 $\pm$ 0.063	0.586$\pm$0.114	0.438 $\pm$ 0.060	0.995 $\pm$ 0.092	0.302 $\pm$ 0.139
	PLACE (+Gram)		0.989$\pm$0.017	0.161 $\pm$ 0.090	0.678$\pm$0.037	1.028 $\pm$ 0.090	0.339 $\pm$ 0.153
	PLACE (+KL)		0.989$\pm$0.016	0.189 $\pm$ 0.109	0.663 $\pm$ 0.050	1.023 $\pm$ 0.092	0.325 $\pm$ 0.146
COCO							
Context: Prior Art on Convolutional Networks							
	ICE [70]	ResNet-34	0.883 $\pm$ 0.029	0.632 $\pm$ 0.020	0.623 $\pm$ 0.086	0.308 $\pm$ 0.003	0.596 $\pm$ 0.002
	CRAFT [25]		0.861 $\pm$ 0.029	0.691 $\pm$ 0.029	0.874 $\pm$ 0.035	0.457 $\pm$ 0.002	0.600 $\pm$ 0.004
	FACE [7]		0.971 $\pm$ 0.013	0.894 $\pm$ 0.010	0.947 $\pm$ 0.002	0.462 $\pm$ 0.003	0.458 $\pm$ 0.001
Context: Sliding-Window Adaptations on ViTs							
	CRAFT adapt. [25, 37]	DINOv2	0.974 $\pm$ 0.015	0.487 $\pm$ 0.090	0.883 $\pm$ 0.076	1.237 $\pm$ 0.287	0.011 $\pm$ 0.007
	FACE adapt. [7, 37]		0.973 $\pm$ 0.016	0.604 $\pm$ 0.099	0.869 $\pm$ 0.083	1.242 $\pm$ 0.288	0.011 $\pm$ 0.006
	CRAFT adapt. [25, 37]	ViT-B/16	0.977 $\pm$ 0.009	0.596 $\pm$ 0.094	0.874 $\pm$ 0.033	1.023 $\pm$ 0.383	0.020 $\pm$ 0.013
	FACE adapt. [7, 37]		0.969 $\pm$ 0.014	0.608 $\pm$ 0.097	0.862 $\pm$ 0.040	1.024 $\pm$ 0.383	0.020 $\pm$ 0.013
Comparison: PLACE Framework vs. Sparse Autoencoders							
	Gated SAE [60]	DINOv2	0.202 $\pm$ 0.194	0.246 $\pm$ 0.137	0.729 $\pm$ 0.068	0.073$\pm$0.005	0.479$\pm$0.318
	SAE [33, 57]		0.143 $\pm$ 0.069	0.276 $\pm$ 0.234	0.859 $\pm$ 0.012	0.080 $\pm$ 0.005	0.589 $\pm$ 0.331
	PLACE (reconstruction)		0.205$\pm$0.022	0.732 $\pm$ 0.019	0.884 $\pm$ 0.001	0.076 $\pm$ 0.002	0.563 $\pm$ 0.006
	PLACE (+Gram)		0.177 $\pm$ 0.014	0.776$\pm$0.009	0.888$\pm$0.001	0.080 $\pm$ 0.001	0.584 $\pm$ 0.015
	PLACE (+KL)		0.173 $\pm$ 0.017	0.768 $\pm$ 0.006	0.887 $\pm$ 0.001	0.080 $\pm$ 0.001	0.576 $\pm$ 0.017
	Gated SAE [60]	ViT-B/16	0.949 $\pm$ 0.063	0.115 $\pm$ 0.103	0.277 $\pm$ 0.041	0.315$\pm$0.038	0.060 $\pm$ 0.028
	SAE [33, 57]		0.993$\pm$0.012	0.033 $\pm$ 0.060	0.415 $\pm$ 0.038	0.328 $\pm$ 0.038	0.068 $\pm$ 0.030
	PLACE (reconstruction)		0.988 $\pm$ 0.001	0.968 $\pm$ 0.004	0.638 $\pm$ 0.005	0.357 $\pm$ 0.004	$< 10^{-3}$
	PLACE (+Gram)		0.957 $\pm$ 0.004	0.979$\pm$0.010	0.655 $\pm$ 0.014	0.351 $\pm$ 0.001	$< 10^{-3}$
	PLACE (+KL)		0.949 $\pm$ 0.004	0.940 $\pm$ 0.004	0.677$\pm$0.005	0.359 $\pm$ 0.002	0.003 $\pm$ 0.001
	Gated SAE [60]	ViT-MAE	0.522 $\pm$ 0.239	0.416 $\pm$ 0.304	0.309 $\pm$ 0.068	0.609 $\pm$ 0.164	0.083 $\pm$ 0.043
	SAE [33, 57]		0.692$\pm$0.296	0.272 $\pm$ 0.190	0.200 $\pm$ 0.045	0.595 $\pm$ 0.189	0.076 $\pm$ 0.037
	PLACE (reconstruction)		0.435 $\pm$ 0.242	0.432$\pm$0.279	0.371 $\pm$ 0.062	0.585$\pm$0.191	0.075 $\pm$ 0.038
	PLACE (+Gram)		0.456 $\pm$ 0.231	0.432$\pm$0.286	0.429$\pm$0.126	0.608 $\pm$ 0.210	0.074$\pm$0.038
	PLACE (+KL)		0.454 $\pm$ 0.235	0.425 $\pm$ 0.284	0.424 $\pm$ 0.118	0.602 $\pm$ 0.203	0.074$\pm$0.038



Figure 4: *Visualizing the C-Deletion and C-Insertion protocols.* Reading from left to right illustrates C-Deletion: the original input image undergoes progressive masking of patch-tokens based on concept activations, culminating in a fully obscured state. Conversely, reading from right to left depicts C-Insertion, where patches are progressively revealed starting from the fully masked baseline.

Gini for concept sparsity (i.e., monosemanticity). Second, we report D_{KL} for predictive consistency; crucially, although D_{KL} requires a classification head, we use it purely for evaluation purposes (Appendix D.1). We also report C-Ins and C-Del to ensure comparison with prior work, even if they are incomparable across different architectures and classification heads. Consequently, the scores for ICE [70], CRAFT [25], and FACE [7] on ResNet-34 [29], as well as sliding-window adaptations on ViTs [37], are provided in Table 1 as methodological references rather than as direct comparisons. Focusing on the PLACE–SAE comparison, where both methods use matched dictionary rank $r=25$ for a controlled evaluation, Table 1 reveals three consistent findings across backbones and datasets:

- ❶ PLACE consistently outperforms both SAE variants in concept sparsity at matched dictionary rank (peak C-Gini 0.933 vs. 0.724 for Gated SAEs on DINOv2; 0.678 vs. 0.276 on ViT-MAE), confirming that strict non-negativity on U yields cleaner parts-based decompositions than ℓ_1 -penalized activations.
- ❷ PLACE (+Gram) closely matches the supervised reference (+KL) on D_{KL} and MSE across all configurations, and achieves lower D_{KL} on self-supervised backbones (e.g., DINOv2/ImageNet: 1.162 vs. 1.233). SAEs achieve lower absolute D_{KL} (0.368 for Gated SAE), but this reflects their learned encoder-decoder rather than superior concept quality; the relevant comparison for validating the Gram proxy is against +KL, which shares the same factorization structure.
- ❸ While SAEs achieve lower reconstruction error due to their encoder-decoder capacity, this does not translate to superior interpretability: their C-Gini scores are consistently lower. The causal metrics corroborate this: PLACE (+Gram) outperforms SAEs on C-Del on self-supervised backbones (e.g., 0.914 vs. 0.620 on DINOv2/ImageNet), while avoiding the costly training phase of SAEs (Section E.9).

4.2 Patch-Token Extraction vs. Crop-and-Resize

To address the input mismatch (Section 3), we compare PLACE against CRAFT [25] and FACE [7] adapted for ViTs by Kondapaneni et al. [37]. This variation factorizes the final CLS tokens of 64×64 sliding-window crops (stride 64, resized to 224×224), and optimizes via MSE and MSE+KL.

Discussion. The sliding-window baseline artificially inflates C-Ins metrics through two compounding artifacts. First, by factorizing the final CLS token immediately preceding the classifier head, the extracted concepts shortcut the predictive distribution, yielding high confidence scores that reflect architectural proximity to the logits. Second, as formally detailed in Appendix A, because each resized crop is processed independently, its CLS token aggregates context *as if the crop were a complete image*, violating the context of the original scene. The structural damage caused by the crop-and-resize pipeline is mathematically evident in the representation error gap: the sliding-window baseline yields an MSE of 3.489 on ViT-B/16, compared to 0.292 for PLACE, and highlights a representational mismatch. The final CLS token inherently aggregates highly entangled, global scene statistics, and forcing such a low-rank factorization leads to high reconstruction errors.

4.3 Semantic Alignment with Supervised Concepts

We validate the semantic quality of our unsupervised concepts by measuring their geometric alignment with the supervised “gold standard” (+KL). Because causal metrics inherently depend on the surrogate linear probe, we employ the RSVC framework [37] to compare the concept banks directly. We report

Maximum Cosine Similarity (MCS) for one-to-one vector alignment, and Cross-Model Cosine Similarity i.e., Pearson correlation of Lasso predictions (CMCS) for the functional equivalence of the overall subspace.

Table 2: Agreement of unsupervised concept banks with the supervised (+KL) reference on ImageNet, measured via RSVC. Mean $\pm$ std over 5 runs.

Backbone	Comparison	MCS $\uparrow$	CMCS $\uparrow$
DINOv2 (ViT-B/14)	Reconstruction vs. +KL	0.840 $\pm$ 0.043	0.732 $\pm$ 0.040
	+Gram vs. +KL	0.857 $\pm$ 0.038	0.722 $\pm$ 0.041
ViT-MAE (B/16)	Reconstruction vs. +KL	0.921 $\pm$ 0.033	0.863 $\pm$ 0.021
	+Gram vs. +KL	0.923 $\pm$ 0.026	0.856 $\pm$ 0.021

Discussion. Table 2 reveals that PLACE recovers the supervised concept space entirely without labels, though the mechanistic role of the geometric proxy varies by VFM architecture.

For DINOv2, pure reconstruction captures a functionally valid subspace (high CMCS) but leaves individual vectors loosely aligned. Introducing the geometric proxy (+Gram) forces the unsupervised concept vectors to lock into the same semantic directions as the supervised ones, boosting specific vector alignment (MCS 0.840 $\rightarrow$ 0.857). This validates our hypothesis: $\mathcal{L}_{\text{gram}}$ acts as a substitute for label-driven supervision to discover predictively relevant structures. Conversely, for the masked autoencoder ViT-MAE pure reconstruction inherently achieves high alignment (MCS 0.921), and the +Gram objective maintains fidelity. Ultimately, this confirms that PLACE bridges the gap between unsupervised geometry and task-specific semantics, adapting to the underlying structural priors of multiple foundation models.

5 Limitations and Future Work

While PLACE trivially extends to other ViT-based VFMs [13, 63], working on architectures with alternative tokenization schemes (Swin [47, 48], DeiT [66], JEPA [6, 52], and registers [19, 34]) is a natural next step. Our Gram-matrix alignment serves as a geometric proxy for faithfulness. We emphasize that our theoretical KL bound (Appendix C) is loose due to worst-case Lipschitz compounding and serves primarily as a mechanistic motivation rather than a strict numerical guarantee; tightening it via local Lipschitz estimates is an open direction. More broadly, complementing the Gram-matrix loss with alternative representational similarity measures [38, 53, 59] is a promising avenue. Nonetheless, the empirical agreement measured via RSVC (Section 4.3) remains robust, validating the efficacy of the proxy. Performance on COCO is lower than on ImageNet; as demonstrated by our dataset size ablation (Appendix E.8), this is a data-driven limitation of NMF-based factorization on small sample sets (e.g., < 50 images), leading to under-determined systems, rather than a failure of the extraction pipeline. PLACE assumes that concepts are linearly factorizable in the activation space and requires pre-defining the dictionary rank r , which may not perfectly reflect the dynamic conceptual complexity of varied classes; adaptive rank selection via MDL [65] could relax this constraint. Additionally, spatial token grouping ($s > 1$) introduces spatial quantization error, as tokens within an $s \times s$ block share a reconstructed coefficient, masking intra-block variation (Appendix B). Regarding evaluation, while we assess concept agreement via the RSVC framework [37], its reliance on linear predictive models (i.e., Lasso regression) to correlate subspaces may not fully capture the highly non-linear transformations inherent to deep transformer architectures; developing non-linear semantic alignment metrics remains an open avenue for future research. Finally, the obtained concept heatmaps are evaluated qualitatively (Appendix H) rather than via human studies, and PLACE does not automatically assign textual labels to the discovered concepts.

6 Conclusion

In this work, we have proposed Patch-Level Agnostic Concept Extraction (PLACE), a fully unsupervised framework for concept extraction on VFMs and ViTs. PLACE adapts concept discovery directly to the native patch-token representations of ViTs by ❶ bypassing the distortive crop-and-resize pipelines of prior methods to preserve the exact spatial geometry; ❷ enforcing sparsity via localized Semi-NMF; and ❸ introducing a Gram-matrix alignment loss that mathematically anchors the extracted concepts to the model’s predictive structure without requiring *any* downstream classifier supervision. This way, PLACE overcomes the computational overhead of training auxiliary modules

like Sparse Autoencoders (SAEs) and natively avoids resizing artifacts, delivering task-agnostic, spatially resolved, human-understandable concepts out of the box. Through comprehensive experiments on ImageNet and COCO using DINOv2, ViT-MAE, and ViT-B/16, we show that PLACE extracts concepts up to $20\times$ faster than KL-supervised alternatives and $11\times$ faster than SAEs. Furthermore, it yields highly robust concepts that achieve up to 40 percentage points higher sparsity (peaking at a 0.933 C-Gini score) and recover the semantic fidelity of supervised extraction (0.923 maximum cosine similarity).

Acknowledgments

This work was supported by the U.S. National Science Foundation Award CNS-2312875 and OAC-2530896, by the Office of Naval Research under grant N00014-23-1-2221 and by the Air Force Office of Scientific Research under grant FA9550-23-1-0261 and by DARPA under cooperative agreement D25AC00374-00.

References

- [1] Radhakrishna Achanta et al. “SLIC Superpixels Compared to State-of-the-Art Superpixel Methods”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 34.11 (2012), pp. 2274–2282. DOI: 10.1109/TPAMI.2012.120.
- [2] Reduan Achtibat et al. “From attribution maps to human-understandable explanations through concept relevance propagation”. In: *Nature Machine Intelligence* 5.9 (2023), pp. 1006–1019.
- [3] Suhyun Ahn, Donggyu Lee, and Jinah Park. “Adaptability of Vision Foundation Models for 3D Medical Image Segmentation”. In: *IEEE Open Journal of Signal Processing* 7 (2026), pp. 82–90. DOI: 10.1109/OJSP.2025.3650437.
- [4] Julia Amann et al. “Explainability for artificial intelligence in healthcare: a multidisciplinary perspective”. In: *BMC Medical Informatics and Decision Making* 20.1 (2020), p. 310. ISSN: 1472-6947. DOI: 10.1186/s12911-020-01332-6.
- [5] Shreyash Arya et al. “B-cosification: Transforming deep neural networks to be inherently interpretable”. In: *Advances in Neural Information Processing Systems* 37 (2024), pp. 62756–62786.
- [6] Mahmoud Assran et al. “Self-Supervised Learning from Images with a Joint-Embedding Predictive Architecture”. In: *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2023, pp. 15619–15629. DOI: 10.1109/CVPR52729.2023.01499.
- [7] Dipkamal Bhusal et al. “FACE: Faithful Automatic Concept Extraction”. In: *The Thirty-ninth Annual Conference on Neural Information Processing Systems*. 2025.
- [8] Irving Biederman. “Recognition-by-components: A theory of human image understanding”. In: *Psychological Review* 94.2 (Apr. 1987), pp. 115–147. DOI: 10.1037/0033-295X.94.2.115.
- [9] Moritz Böhle, Mario Fritz, and Bernt Schiele. “B-cos Networks: Alignment is All We Need for Interpretability”. In: *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2022, pp. 10319–10328. DOI: 10.1109/CVPR52688.2022.01008.
- [10] Moritz Böhle, Mario Fritz, and Bernt Schiele. “Holistically explainable vision transformers”. In: *arXiv preprint arXiv:2301.08669* (2023).
- [11] Moritz Böhle et al. “B-Cos Alignment for Inherently Interpretable CNNs and Vision Transformers”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 46.6 (2024), pp. 4504–4518. DOI: 10.1109/TPAMI.2024.3355155.
- [12] C. Boutsidis and E. Gallopoulos. “SVD based initialization: A head start for nonnegative matrix factorization”. In: *Pattern Recognition* 41.4 (2008), pp. 1350–1362. ISSN: 0031-3203. DOI: <https://doi.org/10.1016/j.patcog.2007.09.010>.
- [13] Mathilde Caron et al. “Emerging Properties in Self-Supervised Vision Transformers”. In: *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*. 2021, pp. 9630–9640. DOI: 10.1109/ICCV48922.2021.00951.
- [14] Hila Chefer, Shir Gur, and Lior Wolf. “Transformer Interpretability Beyond Attention Visualization”. In: *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2021, pp. 782–791. DOI: 10.1109/CVPR46437.2021.00084.

- [15] Zhiqing Chen et al. “SSG-CAM: enhancing visual interpretability through refined second-order gradients and evolutionary multi-layer fusion”. In: *Scientific Reports* 16 (2026), p. 6848. DOI: 10.1038/s41598-026-37278-4.
- [16] Taco Cohen and Max Welling. “Group Equivariant Convolutional Networks”. In: *Proceedings of The 33rd International Conference on Machine Learning*. Ed. by Maria Florina Balcan and Kilian Q. Weinberger. Vol. 48. Proceedings of Machine Learning Research. New York, New York, USA: PMLR, 2016, pp. 2990–2999.
- [17] Jean-Baptiste Cordonnier, Andreas Loukas, and Martin Jaggi. “On the Relationship between Self-Attention and Convolutional Layers”. In: *International Conference on Learning Representations*. 2020.
- [18] Ian C. Covert, Scott Lundberg, and Su-In Lee. “Explaining by removing: a unified framework for model explanation”. In: *J. Mach. Learn. Res.* 22.1 (Jan. 2021). ISSN: 1532-4435.
- [19] Timothée Darcet et al. “Vision Transformers Need Registers”. In: *The Twelfth International Conference on Learning Representations*. 2024.
- [20] Jia Deng et al. “ImageNet: A large-scale hierarchical image database”. In: *2009 IEEE Conference on Computer Vision and Pattern Recognition*. 2009, pp. 248–255. DOI: 10.1109/CVPR.2009.5206848.
- [21] Chris H.Q. Ding, Tao Li, and Michael I. Jordan. “Convex and Semi-Nonnegative Matrix Factorizations”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 32.1 (2010), pp. 45–55. DOI: 10.1109/TPAMI.2008.277.
- [22] Teresa Dorszewski et al. “From Colors to Classes: Emergence of Concepts in Vision Transformers”. In: *Explainable Artificial Intelligence*. Ed. by Riccardo Guidotti, Ute Schmid, and Luca Longo. Cham: Springer Nature Switzerland, 2026, pp. 28–47. ISBN: 978-3-032-08317-3.
- [23] Alexey Dosovitskiy et al. “An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale”. In: *International Conference on Learning Representations*. 2021.
- [24] Thomas Fel et al. “A Holistic Approach to Unifying Automatic Concept Extraction and Concept Importance Estimation”. In: *Thirty-seventh Conference on Neural Information Processing Systems*. 2023.
- [25] Thomas Fel et al. “CRAFT: Concept Recursive Activation FacTorization for Explainability”. In: *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2023, pp. 2711–2721. DOI: 10.1109/CVPR52729.2023.00266.
- [26] Amirata Ghorbani et al. “Towards automatic concept-based explanations”. In: *Proceedings of the 33rd International Conference on Neural Information Processing Systems*. Red Hook, NY, USA: Curran Associates Inc., 2019.
- [27] Kai Han et al. “A Survey on Vision Transformer”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 45.1 (2023), pp. 87–110. DOI: 10.1109/TPAMI.2022.3152247.
- [28] Peter Hase, Harry Xie, and Mohit Bansal. “The out-of-distribution problem in explainability and search methods for feature importance explanations”. In: *Proceedings of the 35th International Conference on Neural Information Processing Systems*. NIPS ’21. Red Hook, NY, USA: Curran Associates Inc., 2021. ISBN: 9781713845393.
- [29] Kaiming He et al. “Deep Residual Learning for Image Recognition”. In: *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2016, pp. 770–778. DOI: 10.1109/CVPR.2016.90.
- [30] Kaiming He et al. “Masked Autoencoders Are Scalable Vision Learners”. In: *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2022, pp. 15979–15988. DOI: 10.1109/CVPR52688.2022.01553.
- [31] Dan Hendrycks and Thomas Dietterich. “Benchmarking Neural Network Robustness to Common Corruptions and Perturbations”. In: *International Conference on Learning Representations*. 2019.
- [32] Qihan Huang et al. “Lg-cav: Train any concept activation vector with language guidance”. In: *Advances in Neural Information Processing Systems* 37 (2024), pp. 39522–39551.
- [33] Robert Huben et al. “Sparse Autoencoders Find Highly Interpretable Features in Language Models”. In: *The Twelfth International Conference on Learning Representations*. 2024.
- [34] Nicholas Jiang et al. “Vision Transformers Don’t Need Trained Registers”. In: *The Thirty-ninth Annual Conference on Neural Information Processing Systems*. 2026.

- [35] Been Kim et al. “Interpretability Beyond Feature Attribution: Quantitative Testing with Concept Activation Vectors (TCAV)”. In: *Proceedings of the 35th International Conference on Machine Learning*. PMLR, 2018, pp. 2673–2682.
- [36] Pang Wei Koh et al. “Concept bottleneck models”. In: *International conference on machine learning*. PMLR. 2020, pp. 5338–5348.
- [37] Neehar Kondapaneni, Oisín Mac Aodha, and Pietro Perona. “Representational Similarity via Interpretable Visual Concepts”. In: *The Thirteenth International Conference on Learning Representations*. 2025.
- [38] Simon Kornblith et al. “Similarity of Neural Network Representations Revisited”. In: *Proceedings of the 36th International Conference on Machine Learning*. Ed. by Kamalika Chaudhuri and Ruslan Salakhutdinov. Vol. 97. Proceedings of Machine Learning Research. PMLR, 2019, pp. 3519–3529.
- [39] Ananya Kumar et al. “Fine-Tuning can Distort Pretrained Features and Underperform Out-of-Distribution”. In: *International Conference on Learning Representations*. 2022.
- [40] Michihiro Kuroki and Toshihiko Yamasaki. “CE-FAM: Concept-Based Explanation via Fusion of Activation Maps”. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2025, pp. 1413–1422.
- [41] Victor Lecomte et al. “What Causes Polysemanticity? An Alternative Origin Story of Mixed Selectivity from Incidental Causes”. In: *ICLR 2024 Workshop on Representational Alignment*. 2024.
- [42] Daniel D Lee and H Sebastian Seung. “Learning the parts of objects by non-negative matrix factorization”. In: *Nature* 401.6755 (1999), pp. 788–791.
- [43] Yifan Li et al. “Evaluating Object Hallucination in Large Vision-Language Models”. In: *The 2023 Conference on Empirical Methods in Natural Language Processing*. 2023.
- [44] Pengchen Liang et al. “Vision Foundation Models in Medical Image Analysis: Advances and Challenges”. In: *Blockchain and Trustworthy Systems*. Ed. by Jianguo Chen, Xiaonan Luo, and Yuanlong Yu. Singapore: Springer Nature Singapore, 2026, pp. 170–181. ISBN: 978-981-95-3483-8.
- [45] Tsung-Yi Lin et al. “Microsoft COCO: Common Objects in Context”. In: *Computer Vision – ECCV 2014*. Ed. by David Fleet et al. Cham: Springer International Publishing, 2014, pp. 740–755. ISBN: 978-3-319-10602-1.
- [46] Zachary C. Lipton. “The mythos of model interpretability”. In: *Commun. ACM* 61.10 (Sept. 2018), pp. 36–43. ISSN: 0001-0782. DOI: 10.1145/3233231.
- [47] Ze Liu et al. “Swin Transformer V2: Scaling Up Capacity and Resolution”. In: *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2022, pp. 11999–12009. DOI: 10.1109/CVPR52688.2022.01170.
- [48] Ze Liu et al. “Swin Transformer: Hierarchical Vision Transformer using Shifted Windows”. In: *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*. 2021, pp. 9992–10002. DOI: 10.1109/ICCV48922.2021.00986.
- [49] Ilya Loshchilov and Frank Hutter. “Decoupled Weight Decay Regularization”. In: *International Conference on Learning Representations*. 2019.
- [50] Chiyu Ma et al. “Interpretable image classification with adaptive prototype-based vision transformers”. In: *Advances in Neural Information Processing Systems* 37 (2024), pp. 41447–41493.
- [51] Vien V. Mai and Mikael Johansson. “Stability and Convergence of Stochastic Gradient Clipping: Beyond Lipschitz Continuity and Smoothness”. In: *Proceedings of the 38th International Conference on Machine Learning*. Ed. by Marina Meila and Tong Zhang. Vol. 139. Proceedings of Machine Learning Research. PMLR, 2021, pp. 7325–7335.
- [52] Shentong Mo and Shengbang Tong. “Connecting Joint-Embedding Predictive Architecture with Contrastive Self-supervised Learning”. In: *The Thirty-eighth Annual Conference on Neural Information Processing Systems*. 2024.
- [53] Alex Graeme Murphy, Joel Zylberberg, and Alona Fyshe. “Correcting Biased Centered Kernel Alignment Measures in Biological and Artificial Neural Networks”. In: *ICLR 2024 Workshop on Representational Alignment*. 2024.

- [54] Muzammal Naseer et al. “Intriguing properties of vision transformers”. In: *Proceedings of the 35th International Conference on Neural Information Processing Systems*. NIPS ’21. Red Hook, NY, USA: Curran Associates Inc., 2021. ISBN: 9781713845393.
- [55] Tuomas Oikarinen and Tsui-Wei Weng. “CLIP-Dissect: Automatic Description of Neuron Representations in Deep Vision Networks”. In: *The Eleventh International Conference on Learning Representations*. 2023.
- [56] Maxime Oquab et al. “DINOv2: Learning Robust Visual Features without Supervision”. In: *Transactions on Machine Learning Research* (2024). Featured Certification. ISSN: 2835-8856.
- [57] Mateusz Pach et al. “Sparse Autoencoders Learn Monosemantic Features in Vision-Language Models”. In: *The Thirty-ninth Annual Conference on Neural Information Processing Systems*. 2025.
- [58] Maithra Raghu et al. “Do vision transformers see like convolutional neural networks?” In: *Advances in neural information processing systems* 34 (2021), pp. 12116–12128.
- [59] Maithra Raghu et al. “SVCCA: singular vector canonical correlation analysis for deep learning dynamics and interpretability”. In: *Proceedings of the 31st International Conference on Neural Information Processing Systems*. NIPS’17. Long Beach, California, USA: Curran Associates Inc., 2017, pp. 6078–6087. ISBN: 9781510860964.
- [60] Senthooran Rajamanoharan et al. “Improving Sparse Decomposition of Language Model Activations with Gated Sparse Autoencoders”. In: *The Thirty-eighth Annual Conference on Neural Information Processing Systems*. 2024.
- [61] Cynthia Rudin. “Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead”. In: *Nature Machine Intelligence* 1.5 (May 2019), pp. 206–215. ISSN: 2522-5839. DOI: 10.1038/s42256-019-0048-x.
- [62] Olga Russakovsky et al. “ImageNet Large Scale Visual Recognition Challenge”. In: *International Journal of Computer Vision* 115.3 (Dec. 2015), pp. 211–252. ISSN: 1573-1405. DOI: 10.1007/s11263-015-0816-y.
- [63] Oriane Siméoni et al. *DINOv3*. 2025. arXiv: 2508.10104 [cs.CV].
- [64] Ilya M Sobol. “Sensitivity estimates for nonlinear mathematical models”. In: *Mathematical Modelling and Computational Experiments* 1.4 (1993), pp. 407–414.
- [65] Vincent Y.F. Tan and Cédric Févotte. “Automatic Relevance Determination in Nonnegative Matrix Factorization with the β -Divergence”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 35.7 (2013), pp. 1592–1605. DOI: 10.1109/TPAMI.2012.240.
- [66] Hugo Touvron et al. “Training data-efficient image transformers; distillation through attention”. In: *Proceedings of the 38th International Conference on Machine Learning*. Ed. by Marina Meila and Tong Zhang. Vol. 139. Proceedings of Machine Learning Research. PMLR, 2021, pp. 10347–10357.
- [67] Shikhar Tuli et al. “Are Convolutional Neural Networks or Transformers more like human vision?” In: *43rd Annual Meeting of the Cognitive Science Society: Comparative Cognition: Animal Minds, CogSci 2021*. 2021.
- [68] Mateo Espinosa Zarlenga et al. “Concept Embedding Models”. In: *Advances in Neural Information Processing Systems*. Ed. by Alice H. Oh et al. 2022.
- [69] Richard Zhang. “Making Convolutional Networks Shift-Invariant Again”. In: *Proceedings of the 36th International Conference on Machine Learning*. Ed. by Kamalika Chaudhuri and Ruslan Salakhutdinov. Vol. 97. Proceedings of Machine Learning Research. PMLR, 2019, pp. 7324–7334.
- [70] Ruihan Zhang et al. “Invertible concept-based explanations for cnn models with non-negative concept activation vectors”. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 35. 13. 2021, pp. 11682–11690.

PLACE: Patch-Level Agnostic Concept Extraction

Supplementary Material

Contents

A	Why crop-and-resize concept extraction is sub-optimal for Vision Transformers	16
A.1	Translation equivariance in CNNs	16
A.2	Why crop-and-resize exploits equivariance	17
A.3	Vision Transformers are not translation equivariant	17
A.4	What goes wrong with crop-and-resize in ViTs	18
A.5	Empirical evidence	18
A.6	Our approach: factorizing patch tokens in place	19
B	Token grouping: detailed spatial arithmetic	19
B.1	From pixels to patch tokens	19
B.2	Token grouping and resolution	19
B.3	Per-architecture spatial resolution	20
C	Connecting Gram-matrix alignment to predictive consistency	20
C.1	Preliminary: Gram deviation under reconstruction error	20
C.2	Step 1: Attention logit deviation	21
C.3	Step 2: Single-block error propagation	22
C.4	Step 3: Composition through suffix blocks	23
C.5	Step 4: From CLS-token deviation to KL divergence	23
C.6	Discussion	24
D	Implementation Details	25
D.1	Linear probes for validation	25
D.2	PLACE Hyperparameters	25
D.3	PCA Visualization Details	26
E	Additional Analyses	26
E.1	NNDSVD subsampling ratio for Semi-NMF initialization	26
E.2	Faithfulness loss weight sensitivity	27
E.3	Extraction layer ablation	28
E.4	Decomposition rank ablation	29
E.5	Token group size ablation	30
E.6	Sanity check: weight randomization	30
E.7	Stability across seeds	31
E.8	Effect of Dataset Size on Concept Stability	32

E.9 Runtime and memory	32
E.10 Computational Complexity and Scalability to Larger Architectures	33
E.11 Per-class faithfulness trends	35
E.12 Per-class agreement scores	37
E.13 Robustness to Image Perturbations	38
F Evaluation Metrics Description	39
G Broader Impact	40
H Additional Qualitative Results	40

A Why crop-and-resize concept extraction is sub-optimal for Vision Transformers

In this appendix we formalise the argument, sketched in Section 3, that the crop-and-resize concept-proposal pipeline used by CRAFT [25] and RSVC [37] is well-motivated for CNNs but is sub-optimal for Vision Transformers. The root cause is a structural property of convolutional networks, translation equivariance, that ViTs do not possess.

A.1 Translation equivariance in CNNs

Let T_{τ} denote the operator that shifts an image $\mathbf{x} \in \mathbb{R}^{3 \times H \times W}$ by a displacement $\tau = (\tau_h, \tau_w)$:

$$[T_{\tau} \mathbf{x}](c, i, j) = \mathbf{x}(c, i - \tau_h, j - \tau_w). \quad (4)$$

A mapping $\phi : \mathbb{R}^{c_{\text{in}} \times H \times W} \rightarrow \mathbb{R}^{c_{\text{out}} \times H' \times W'}$ is *translation equivariant* if there exists a (possibly stride-adjusted) translation operator T'_{τ} on the output space such that

$$\phi(T_{\tau} \mathbf{x}) = T'_{\tau} \phi(\mathbf{x}) \quad \forall \mathbf{x}, \tau. \quad (5)$$

Convolution is equivariant. A convolutional layer with kernel $\mathbf{w} \in \mathbb{R}^{c_{\text{out}} \times c_{\text{in}} \times k \times k}$ computes

$$[\phi(\mathbf{x})](m, i, j) = \sum_{c=1}^{c_{\text{in}}} \sum_{p,q=0}^{k-1} w_{m,c,p,q} \mathbf{x}(c, i+p, j+q). \quad (6)$$

Evaluating on the translated input gives

$$\begin{aligned} [\phi(T_{\tau} \mathbf{x})](m, i, j) &= \sum_{c,p,q} w_{m,c,p,q} [T_{\tau} \mathbf{x}](c, i+p, j+q) \\ &= \sum_{c,p,q} w_{m,c,p,q} \mathbf{x}(c, (i-\tau_h)+p, (j-\tau_w)+q) \\ &= [\phi(\mathbf{x})](m, i-\tau_h, j-\tau_w) \\ &= [T'_{\tau} \phi(\mathbf{x})](m, i, j), \end{aligned} \quad (7)$$

where the third equality follows because the kernel weights $w_{m,c,p,q}$ are *shared across all spatial positions* and do not depend on (i, j) . This is the key structural property: the same filter is applied everywhere, so shifting the input shifts the output by exactly the same amount. Pointwise nonlinearities (ReLU, GELU, etc.) act independently at each spatial position and therefore commute with translation trivially.

Spatial pooling: equivariance and local invariance. Average or max pooling with a fixed spatial grid of size $s \times s$ maps each non-overlapping $s \times s$ block to a single value. This operation is equivariant only to translations that are multiples of the stride s : shifting the input by s pixels shifts the pooled output by exactly one position. For shifts smaller than s , the pooled output remains unchanged, introducing *local translation invariance*: content that moves within the same pooling window produces identical activations. Crucially, even with pooling, each output position still depends on a bounded local neighborhood of the input.

Consequence: local receptive fields. By composing equivariant convolutions, pointwise nonlinearities, and local pooling, the feature map at any layer ℓ of a CNN satisfies

$$f_{\text{cnn}}^{(\ell)}(T_{\tau} \mathbf{x}) = T'_{\tau} f_{\text{cnn}}^{(\ell)}(\mathbf{x}) \quad (8)$$

up to boundary padding artifacts. This means that each spatial position (i, j) in $\mathbf{A} = f_{\text{cnn}}^{(\ell)}(\mathbf{x}) \in \mathbb{R}^{c \times h \times w}$ depends only on a bounded neighborhood of (i, j) in the input, the *receptive field* $\mathcal{R}_{i,j}^{(\ell)}$:

$$\mathbf{A}[:, i, j] = g(\mathbf{x}|_{\mathcal{R}_{i,j}^{(\ell)}}). \quad (9)$$

A.2 Why crop-and-resize exploits equivariance

In CRAFT [25], FACE [7] and RSVC [37], a crop $\mathbf{x}_p$ from the spatial window $[i_0 : i_1, j_0 : j_1]$ of $\mathbf{x}$ is *resized* to the full input resolution $H \times W$ and passed through the network independently. CRAFT explicitly defines a “crop and resize function $\pi(\cdot)$ ” and justifies the resize by arguing that modern CNNs are trained with aggressive crop augmentation (RandomResizedCrop, CutMix, etc.), so resized crops lie within the training distribution.

For CNNs, this argument is reinforced by equivariance: thanks to the local receptive field property (Eq. (9)), the feature map of the resized crop satisfies

$$f_{\text{cnn}}^{(\ell)}(\text{resize}(\mathbf{x}_p)) \approx \mathbf{A}[:, i'_0 : i'_1, j'_0 : j'_1], \quad (10)$$

where $[i'_0 : i'_1, j'_0 : j'_1]$ is the corresponding region in the feature map. The approximation is imperfect because resizing changes the effective scale and boundary padding introduces edge effects, but the core property holds: the activation of the crop describes the *local visual content* of that image region. After global average pooling, each crop yields a vector $\bar{\mathbf{a}}_p \in \mathbb{R}^c$ that summarizes local content. Collecting these vectors into $\mathbf{A} \in \mathbb{R}^{n \times c}$ and factorizing via NMF discovers spatially localized concepts, which is the intended behavior.

A.3 Vision Transformers are not translation equivariant

A ViT [23] tokenizes the input into $N = (H/P)^2$ non-overlapping patches of size $P \times P$, linearly projects each patch into $\mathbb{R}^d$, prepends a CLS token, adds learnable positional embeddings, and processes the sequence through L transformer blocks. While Cordonnier et al. [17] prove that self-attention layers *can* implement convolutional filters, they are not constrained to do so: the mapping is strictly more general and imposes no spatial locality by default. We show below that the ViT architecture violates translation equivariance at multiple stages.

1. Positional embeddings break equivariance. Let $\mathbf{e}_i^{(0)} = \mathbf{W}_{\text{proj}} \mathbf{x}_i + \mathbf{p}_i$ be the initial embedding of patch i , where $\mathbf{p}_i \in \mathbb{R}^d$ is the positional embedding. Translating the image by one patch width permutes the patch content ($\mathbf{x}_i \rightarrow \mathbf{x}_{\sigma(i)}$ for some permutation σ), but the positional embeddings remain tied to their grid positions:

$$\mathbf{e}_i^{(0)}(T_\tau \mathbf{x}) = \mathbf{W}_{\text{proj}} \mathbf{x}_{\sigma(i)} + \mathbf{p}_i \neq \mathbf{W}_{\text{proj}} \mathbf{x}_{\sigma(i)} + \mathbf{p}_{\sigma(i)}. \quad (11)$$

The mismatch between content (which has shifted) and position (which has not) breaks equivariance before the first transformer block is even applied. This stands in direct contrast to CNNs, where spatial weight sharing guarantees that the same filter response is produced regardless of position [16].

2. Self-attention is global and content-dependent. Multi-head self-attention computes, for each token i , a weighted sum over *all* tokens:

$$\text{Attn}(\mathbf{Q}, \mathbf{K}, \mathbf{V})_i = \sum_{j=0}^N \frac{\exp(\mathbf{q}_i^\top \mathbf{k}_j / \sqrt{d_h})}{\sum_{m=0}^N \exp(\mathbf{q}_i^\top \mathbf{k}_m / \sqrt{d_h})} \mathbf{v}_j. \quad (12)$$

The attention weights a_{ij} depend on the *content* of all tokens through the queries and keys, whereas convolutional kernel weights are fixed parameters independent of the input (Eq. (6)). Two consequences follow. First, every token’s representation after layer 1 is a function of all input patches: there is no bounded receptive field and no analogue of Eq. (9). Second, the attention pattern itself changes when the input is translated, because the queries and keys change. Even without positional embeddings, the softmax normalization couples all tokens, making the mapping globally contextual.

3. Empirical confirmation. The theoretical absence of equivariance has measurable consequences on how ViTs process visual information [27]. Raghu et al. [58] provide a systematic comparison of ViT and CNN internal representations, finding that (i) ViTs aggregate global spatial information from the earliest layers, whereas CNNs build local-to-global hierarchies gradually, (ii) ViT representations are more uniform across layers, with lower-layer features already encoding high-level structure, and (iii) spatial information is preserved throughout ViT layers but in a qualitatively different manner from CNNs, with self-attention enabling long-range feature propagation via residual connections. Complementing these representational findings, Tuli et al. [67] analyze the behavioral consequences of these architectural differences. They show that ViTs, precisely because they relax the translation-invariance constraint of CNNs, develop recognition strategies whose error patterns are more consistent

with human vision. This suggests that ViTs learn fundamentally different visual features, relying on global structure and shape rather than local texture. Naseer et al. [54] further corroborate this by demonstrating that ViTs exhibit markedly different robustness properties from CNNs: they are more resilient to severe occlusions, texture corruption, and adversarial perturbations, which is consistent with a recognition strategy that does not depend on local, spatially equivariant feature detection.

Taken together, these findings confirm that ViTs do not merely “relax” translation equivariance as a minor architectural choice: they leverage its absence to develop qualitatively different representations. Any concept extraction method that assumes CNN-like locality, whether through crop-and-resize proposals or spatial-slice approximations, is therefore methodologically mismatched with the ViT computation.

A.4 What goes wrong with crop-and-resize in ViTs

Let $\mathbf{x}$ be a full image ($H \times W$ pixels) and $\mathbf{x}_p$ a crop from region $[i_0 : i_1, j_0 : j_1]$, resized to $H \times W$ before being fed to the ViT. We identify four distinct failure modes.

1. The patch grid tiles different content at a different scale. When the full image is tokenised with patch size $P=16$, each of the $N=196$ patches covers a 16×16 pixel region of the scene. When a 64×64 crop is resized to 224×224 and tokenised with the same $P=16$, each patch now covers a $\sim 4.6 \times 4.6$ pixel region of the original crop. The visual content within each patch is at a completely different scale and resolution: fine texture that would fall within a single patch of the full image is now spread across multiple patches of the resized crop. The linear patch projection $\mathbf{W}_{\text{proj}}$ was not trained to process content at this magnification. While CRAFT argues that CNNs are trained with crop augmentation so resized crops are in-distribution, this argument is weaker for ViTs: the *patch-level* statistics of a heavily upsampled small crop differ from those of the random crops seen during training (which are typically $\geq 60\%$ of the image area for standard RandomResizedCrop).

2. Positional embeddings encode wrong locations. Patch i in the resized crop receives positional embedding $\mathbf{p}_i$, which the network learned to associate with a particular absolute position in the 14×14 grid. But the visual content now at position i of the crop occupied a completely different position in the original image. A patch of sky that sat in the top-left corner of the full image might land at the center of the resized crop, receiving a positional embedding the network associates with a different semantic context. This inconsistency has no CNN analogue: convolutional filters are spatially shared and carry no positional information.

3. Attention context is entirely different. In the full image, a patch token at position i attends to tokens covering the entire scene (background, foreground, other objects). In the resized crop, the same positional slot attends only to tokens originating from the crop region. The keys and values are functions of different input content, so

$$a_{ij}^{\text{crop}} \neq a_{ij}^{\text{full}} \quad \forall l \geq 1, \quad \forall i, j. \quad (13)$$

The CLS token of the resized crop therefore summarizes the crop as if it were the entire world, incorporating attention patterns, LayerNorm statistics, and residual stream content determined exclusively by the crop.

4. No spatial correspondence to the full-image representation. In a CNN, Eq. (10) establishes an approximate correspondence between the crop activation and a spatial slice of the full-image feature map, grounded in equivariance. For the ViT CLS token, no such correspondence exists. The CLS token of a resized crop does not approximate any sub-vector, spatial slice, or linear projection of the full-image CLS token $\mathbf{z}_0^{(L)}(\mathbf{x})$. It is a fundamentally different representation of a fundamentally different input.

A.5 Empirical evidence

Kondapaneni et al. [37] report (their Table A2) reconstruction errors after Semi-NMF concept extraction that are 5–8 $\times$ higher for ViTs than for ResNets on the same ImageNet classes (e.g. 926.9 for ViT-S vs. 176.2 for ResNet-18 and 205.5 for ResNet-50, with $k=10$ concepts). While part of this gap reflects higher activation dimensionality ($d=384$ for ViT-S vs. $d=512$ for ResNet-18), the analysis above identifies a systematic, irreducible source of additional variance: two crops from the same image that share identical local content but differ in their remaining visible area produce

different CLS vectors, because global attention contextualizes each token by the entire input. This context-dependent variation inflates the effective rank of the activation matrix $\mathbf{A}$, making low-rank factorization inherently less faithful. RSVC does not investigate whether the extraction strategy itself (CLS tokens from resized crops) is appropriate for ViTs, treating the higher reconstruction error as a limitation of the decomposition method rather than a consequence of the extraction pipeline.

A.6 Our approach: factorizing patch tokens in place

Rather than extracting crops, resizing them, and factorizing CLS tokens, our method operates on the *patch-token activations* of the full image at an intermediate layer ℓ . For a single image, the prefix $g_{\leq \ell}$ produces N patch tokens $\mathbf{z}_1^{(\ell)}, \dots, \mathbf{z}_N^{(\ell)} \in \mathbb{R}^d$, each anchored to a fixed position in the $\sqrt{N} \times \sqrt{N}$ grid. Collecting patch tokens across n images gives $\mathbf{A} \in \mathbb{R}^{nN \times d}$, which we factorize via Semi-NMF (Eq. 1).

This avoids all four failure modes:

1. **No artificial re-tokenization.** Each patch token is computed by the network in the context of the full image, at the scale and resolution the model was trained on. The patch projection $\mathbf{W}_{\text{proj}}$ operates on content at its native scale.
2. **Correct positional encoding.** Each token receives the positional embedding corresponding to its actual location in the image, consistent with training.
3. **Consistent attention context.** All patch tokens within an image share the same global attention context (the full image), so differences between their concept coefficients reflect genuine spatial variation in how concepts are activated across the scene, rather than artifacts of crop composition.
4. **Spatial anchoring is preserved.** Although self-attention mixes information globally, each token maintains its spatial identity. Each row of $\mathbf{U}$ corresponds to the token anchored at a specific grid position. Concept heatmaps are thus obtained natively by reshaping the N coefficients per image into a $\sqrt{N} \times \sqrt{N}$ grid and upscaling via bicubic interpolation, without requiring gradient-based attribution or implicit differentiation [25].

Each patch token at layer $\ell \geq 1$ is no longer a purely local descriptor due to self-attention. Consequently, the discovered concepts reflect a mixture of local visual evidence and global context. However, this is a reflection of how the ViT actually represents information, organically grounded in the token grid, rather than an artifact of an extraction pipeline designed for a different architecture.

B Token grouping: detailed spatial arithmetic

The spatial resolution of discovered concepts depends on the interaction between three quantities: the model’s patch size P (in pixels), the input image resolution $H \times W$, and the token group size s used during factorisation. In this appendix, we make these relationships explicit for each architecture considered in our experiments.

B.1 From pixels to patch tokens

A ViT with patch size P and input resolution $H \times W$ produces a grid of

$$N = \frac{H}{P} \times \frac{W}{P} \quad (14)$$

non-overlapping patch tokens. Each token $\mathbf{z}_i \in \mathbb{R}^d$ corresponds to a $P \times P$ pixel region of the input. For a square input ($H = W$), the token grid has spatial dimensions $G_0 \times G_0$ where $G_0 = \sqrt{N}$. As detailed in Table 3, a standard ViT-B/16 at 224×224 resolution yields a 14×14 grid, whereas DINOv2 at 518×518 yields a 37×37 grid.

B.2 Token grouping and resolution

Given a token grid of size $G_0 \times G_0$, we apply non-overlapping average pooling with kernel and stride equal to s , producing a grouped grid of size $G_s = \lfloor G_0/s \rfloor$. Each grouped token is the average of

$s \times s$ adjacent patch tokens and therefore corresponds to a spatial region of $sP \times sP$ pixels in the input image.

After factorization, the grouped coefficient matrix $\mathbf{U}$ is mapped back to the original token resolution via `repeat_interleave`, which replicates each grouped coefficient across the $s \times s$ tokens it represents. This effectively quantizes the spatial resolution of the concept heatmap to $sP \times sP$ pixels. We note that when G_0 is not perfectly divisible by s (as occurs with DINOv2 where $G_0 = 37$ is prime), the floor operation in G_s implies that a small portion of the original tokens at the right and bottom boundaries is excluded from the grouping. Consequently, when applying `repeat_interleave`, the resulting $sG_s \times sG_s$ tensor is strictly smaller than the original $G_0 \times G_0$ activation map. To resolve this shape mismatch during optimization, we compute all losses ($\mathcal{L}_{\text{rec}}$, $\mathcal{L}_{\text{gram}}$) by cropping the original activation grid $\mathbf{A}$ to match the truncated $sG_s \times sG_s$ dimensions, explicitly discarding these marginal boundary tokens.

B.3 Per-architecture spatial resolution

Table 3 reports the effective spatial resolution for each architecture. While for $P = 16$ models we choose s to be divisors of $G_0 = 14$ to ensure perfect coverage, for DINOv2 a marginal approximation error is introduced for $s > 1$, which is quantified in Appendix E.5.

Table 3: Effective spatial resolution of concept heatmaps. P : patch size; H : input resolution; G_0 : native grid dimension (H/P); s : grouping factor; G_s : grouped grid dimension; sP : pixel resolution.

Architecture	P	H	G_0	s	G_s	sP (px)
ViT-B/16 (ImageNet-1k / MAE)	16	224	14	1	14	16
				2	7	32
				7	2	112
				14	1	224
DINOv2 ViT-B/14	14	518	37	1	37	14
				2	18	28
				7	5	98
				37	1	518

C Connecting Gram-matrix alignment to predictive consistency

We formalise the claim, stated informally in Section 3.1, that jointly minimising the reconstruction loss $\mathcal{L}_{\text{rec}}$ and the Gram-matrix alignment loss $\mathcal{L}_{\text{gram}}$ controls the KL divergence through a downstream linear classifier without ever computing it. The argument proceeds in five steps: we first establish a preliminary relationship between reconstruction error and Gram deviation (Preliminary), then bound the attention logit deviation (Step 1), propagate the error through a single transformer block (Step 2), compose the bound across the $L_{\text{suf}} = L - \ell$ suffix blocks (Step 3), and convert the final CLS-token deviation into a KL divergence bound (Step 4). The resulting bound identifies $\mathcal{L}_{\text{rec}}$ and $\mathcal{L}_{\text{gram}}$ as the two quantities whose joint minimization is sufficient, in principle, to control predictive consistency without access to the classifier. The empirical correspondence is stronger than the worst-case composition below would predict, and is verified in Section 4.3.

Notation. Let $\mathbf{A} \in \mathbb{R}^{N \times d}$ denote the patch-token activations at layer ℓ and $\hat{\mathbf{A}} = \mathbf{U}\mathbf{W}^\top$ their Semi-NMF reconstruction. Define the residual $\mathbf{E} = \mathbf{A} - \hat{\mathbf{A}}$, so that $\|\mathbf{E}\|_F^2 = 2\mathcal{L}_{\text{rec}}$. Let $\mathbf{G} = \mathbf{A}\mathbf{A}^\top$ and $\hat{\mathbf{G}} = \hat{\mathbf{A}}\hat{\mathbf{A}}^\top$ be the original and reconstructed Gram matrices. We write $\delta_{\text{rec}} = \|\mathbf{E}\|_F$ and $\delta_{\text{gram}} = (\mathcal{L}_{\text{gram}})^{1/2} = \|\mathbf{G} - \hat{\mathbf{G}}\|_F / \|\mathbf{G}\|_F$ for the two error terms controlled by our objective.

C.1 Preliminary: Gram deviation under reconstruction error

Before analyzing the attention mechanism, we establish a direct relationship between the reconstruction error and the Gram-matrix deviation. This clarifies why $\mathcal{L}_{\text{rec}}$ alone provides *indirect* control over the Gram structure, and motivates the need for the explicit $\mathcal{L}_{\text{gram}}$ term.

Lemma 1. Let $\mathbf{A} \in \mathbb{R}^{N \times d}$ and $\hat{\mathbf{A}} = \mathbf{A} - \mathbf{E}$. Then

$$\|\mathbf{A}\mathbf{A}^\top - \hat{\mathbf{A}}\hat{\mathbf{A}}^\top\|_F \leq 2\|\mathbf{A}\|_F \cdot \delta_{\text{rec}} + \delta_{\text{rec}}^2. \quad (15)$$

Proof. Substituting $\hat{\mathbf{A}} = \mathbf{A} - \mathbf{E}$ and expanding:

$$\mathbf{A}\mathbf{A}^\top - \hat{\mathbf{A}}\hat{\mathbf{A}}^\top = \mathbf{A}\mathbf{E}^\top + \mathbf{E}\mathbf{A}^\top - \mathbf{E}\mathbf{E}^\top.$$

Applying the triangle inequality and sub-multiplicativity ($\|\mathbf{X}\mathbf{Y}^\top\|_F \leq \|\mathbf{X}\|_F \|\mathbf{Y}\|_F$):

$$\|\mathbf{A}\mathbf{A}^\top - \hat{\mathbf{A}}\hat{\mathbf{A}}^\top\|_F \leq \|\mathbf{A}\|_F \|\mathbf{E}\|_F + \|\mathbf{E}\|_F \|\mathbf{A}\|_F + \|\mathbf{E}\|_F^2 = 2\|\mathbf{A}\|_F \cdot \delta_{\text{rec}} + \delta_{\text{rec}}^2. \quad \square$$

Interpretation. Lemma 1 shows that reconstruction fidelity alone exerts some control over the Gram structure. However, this indirect control scales with $\|\mathbf{A}\|_F$, which can be large for high-dimensional transformer activations. The explicit $\mathcal{L}_{\text{gram}}$ term provides a *direct* constraint on $\|\mathbf{G} - \hat{\mathbf{G}}\|_F / \|\mathbf{G}\|_F$ that is normalized and independent of the activation scale. This distinction explains why reconstruction alone nearly suffices when $\|\mathbf{A}\|_F \cdot \delta_{\text{rec}}$ is already small, while the Gram loss adds measurable value precisely when the indirect control is loose.

C.2 Step 1: Attention logit deviation

Consider a single self-attention head with query, key, and value projections $W_Q, W_K, W_V \in \mathbb{R}^{d \times d_h}$. The pre-softmax attention logits are:

$$\mathbf{S} = \frac{1}{\sqrt{d_h}} \mathbf{A} W_Q W_K^\top \mathbf{A}^\top \in \mathbb{R}^{N \times N}. \quad (16)$$

Define $M = W_Q W_K^\top \in \mathbb{R}^{d \times d}$. The reconstructed logits are $\hat{\mathbf{S}} = \hat{\mathbf{A}} M \hat{\mathbf{A}}^\top / \sqrt{d_h}$. We decompose the deviation:

$$\begin{aligned} \mathbf{S} - \hat{\mathbf{S}} &= \frac{1}{\sqrt{d_h}} (\mathbf{A} M \mathbf{A}^\top - \hat{\mathbf{A}} M \hat{\mathbf{A}}^\top) \\ &= \frac{1}{\sqrt{d_h}} (\mathbf{E} M \mathbf{A}^\top + \mathbf{A} M \mathbf{E}^\top - \mathbf{E} M \mathbf{E}^\top), \end{aligned} \quad (17)$$

where the identity follows from substituting $\hat{\mathbf{A}} = \mathbf{A} - \mathbf{E}$ and expanding. Taking the Frobenius norm and applying the triangle inequality yields:

$$\sqrt{d_h} \|\mathbf{S} - \hat{\mathbf{S}}\|_F \leq \|\mathbf{E} M \mathbf{A}^\top\|_F + \|\mathbf{A} M \mathbf{E}^\top\|_F + \|\mathbf{E} M \mathbf{E}^\top\|_F.$$

We bound each term using the sub-multiplicativity of matrix norms (specifically the mixed bound $\|XYZ\|_F \leq \|X\|_F \|Y\|_{\text{op}} \|Z\|_F$) and the property that transposition preserves the Frobenius norm ($\|\mathbf{A}^\top\|_F = \|\mathbf{A}\|_F$). Recalling our definition $\delta_{\text{rec}} = \|\mathbf{E}\|_F$, the three terms are bounded respectively by:

$$\begin{aligned} \|\mathbf{E} M \mathbf{A}^\top\|_F &\leq \|\mathbf{E}\|_F \|M\|_{\text{op}} \|\mathbf{A}\|_F = \|M\|_{\text{op}} \|\mathbf{A}\|_F \delta_{\text{rec}}, \\ \|\mathbf{A} M \mathbf{E}^\top\|_F &\leq \|\mathbf{A}\|_F \|M\|_{\text{op}} \|\mathbf{E}\|_F = \|M\|_{\text{op}} \|\mathbf{A}\|_F \delta_{\text{rec}}, \\ \|\mathbf{E} M \mathbf{E}^\top\|_F &\leq \|\mathbf{E}\|_F \|M\|_{\text{op}} \|\mathbf{E}\|_F = \|M\|_{\text{op}} \delta_{\text{rec}}^2. \end{aligned}$$

Summing these upper bounds and factoring out $\|M\|_{\text{op}}$ gives the final deviation bound:

$$\|\mathbf{S} - \hat{\mathbf{S}}\|_F \leq \frac{\|M\|_{\text{op}}}{\sqrt{d_h}} (2\|\mathbf{A}\|_F \cdot \delta_{\text{rec}} + \delta_{\text{rec}}^2). \quad (18)$$

To connect the Gram-matrix alignment loss to this deviation, we decompose the projection matrix into an identity component and a residual: $M = I + (M - I)$. When $M = I$, the bilinear deviation reduces to $\|\mathbf{G} - \hat{\mathbf{G}}\|_F = \|\mathbf{G}\|_F \cdot \delta_{\text{gram}}$ by definition of δ_{gram} , which is directly controlled by $\mathcal{L}_{\text{gram}}$. Note that by Lemma 1, reconstruction alone would only bound this quantity indirectly as $2\|\mathbf{A}\|_F \cdot \delta_{\text{rec}} + \delta_{\text{rec}}^2$, underscoring the value of the explicit Gram constraint. Splitting the full deviation accordingly:

$$\mathbf{A} M \mathbf{A}^\top - \hat{\mathbf{A}} M \hat{\mathbf{A}}^\top = \underbrace{(\mathbf{G} - \hat{\mathbf{G}})}_{\text{Gram error}} + \underbrace{\mathbf{A}(M - I)\mathbf{A}^\top - \hat{\mathbf{A}}(M - I)\hat{\mathbf{A}}^\top}_{\text{residual}}. \quad (19)$$

The first term is bounded by $\|\mathbf{G}\|_F \cdot \delta_{\text{gram}}$. The second term has the same structure as Eq. (17) with M replaced by $(M - I)$, yielding:

$$\|\mathbf{S} - \hat{\mathbf{S}}\|_F \leq \frac{1}{\sqrt{d_h}} \left(\|\mathbf{G}\|_F \cdot \delta_{\text{gram}} + \|M - I\|_{\text{op}} (2\|\mathbf{A}\|_F \cdot \delta_{\text{rec}} + \delta_{\text{rec}}^2) \right). \quad (20)$$

Equation (20) reveals the complementary roles of the two losses. $\mathcal{L}_{\text{gram}}$ directly controls the dominant identity component of the attention deviation, while $\mathcal{L}_{\text{rec}}$ handles the projection-specific residual. When $M \approx I$ (i.e. the head computes near-identity queries and keys), the Gram term dominates.

Crucially, this decomposition also explains why purely reconstruction-based methods (like SAEs or standard NMF) are insufficient for predictive faithfulness in ViTs. By Lemma 1, reconstruction alone (δ_{rec}) provides only an *indirect* bound on the Gram deviation that scales with $\|\mathbf{A}\|_F$, leaving the identity component of the attention logits loosely constrained. Our explicit Gram-matrix alignment plugs this gap with a scale-normalized, direct constraint, explaining why PLACE (+Gram) matches supervised performance.

C.3 Step 2: Single-block error propagation

A standard transformer block applies multi-head self-attention (MHSA) followed by a feed-forward network (FFN), each wrapped in a residual connection and layer normalization (LN):

$$\mathbf{A}' = \text{LN}(\mathbf{A} + \text{MHSA}(\mathbf{A})), \quad (21)$$

$$\mathbf{A}'' = \text{LN}(\mathbf{A}' + \text{FFN}(\mathbf{A}')). \quad (22)$$

We analyze the error propagation through each sub-layer.

MHSA output deviation. For a single head, the output is $\mathbf{O} = \text{softmax}(\mathbf{S}) \mathbf{A} W_V \in \mathbb{R}^{N \times d_h}$, where $\mathbf{S}$ is defined in Eq. (16). Define the attention matrices $\mathbf{P} = \text{softmax}(\mathbf{S})$ and $\hat{\mathbf{P}} = \text{softmax}(\hat{\mathbf{S}})$. The output deviation decomposes as:

$$\begin{aligned} \mathbf{O} - \hat{\mathbf{O}} &= \mathbf{P} \mathbf{A} W_V - \hat{\mathbf{P}} \hat{\mathbf{A}} W_V \\ &= \mathbf{P} \mathbf{A} W_V - \hat{\mathbf{P}} (\mathbf{A} - \mathbf{E}) W_V \\ &= (\mathbf{P} - \hat{\mathbf{P}}) \mathbf{A} W_V + \hat{\mathbf{P}} \mathbf{E} W_V. \end{aligned} \quad (23)$$

Taking the Frobenius norm:

$$\|\mathbf{O} - \hat{\mathbf{O}}\|_F \leq \|\mathbf{P} - \hat{\mathbf{P}}\|_F \cdot \|\mathbf{A}\|_F \cdot \|W_V\|_{\text{op}} + \sqrt{N} \cdot \|W_V\|_{\text{op}} \cdot \delta_{\text{rec}}, \quad (24)$$

where we used $\|\hat{\mathbf{P}}\|_{\text{op}} \leq \sqrt{N}$ since the spectral norm of an $N \times N$ row-stochastic matrix is bounded by $\sqrt{N}$.

The first term in Eq. (24) depends on the deviation of the attention probability matrix $\mathbf{P}$. Because the attention operation applies the softmax function independently to each row of the logit matrix $\mathbf{S}$, we first bound the error for a single row. Let $\mathbf{s}, \hat{\mathbf{s}} \in \mathbb{R}^N$ be corresponding row vectors extracted from $\mathbf{S}$ and $\hat{\mathbf{S}}$. The softmax function is 1-Lipschitz continuous:

$$\|\text{softmax}(\mathbf{s}) - \text{softmax}(\hat{\mathbf{s}})\|_2 \leq \|\mathbf{s} - \hat{\mathbf{s}}\|_2, \quad (25)$$

because its Jacobian matrix, $J = \text{diag}(\mathbf{p}) - \mathbf{p}\mathbf{p}^\top$, has a spectral norm strictly bounded by 1. Applying this 1-Lipschitz property independently across all N rows and summing the squared errors yields the matrix-level bound in the Frobenius norm:

$$\|\mathbf{P} - \hat{\mathbf{P}}\|_F \leq \|\mathbf{S} - \hat{\mathbf{S}}\|_F. \quad (26)$$

Residual connection and LayerNorm. The residual connection produces $\mathbf{A} + \mathbf{O} = \mathbf{A} + \text{MHSA}(\mathbf{A})$, so its deviation is bounded by $\delta_{\text{rec}} + \|\mathbf{O} - \hat{\mathbf{O}}\|_F$. LayerNorm is locally Lipschitz: for bounded inputs there exists a constant L_{LN} such that

$$\|\text{LN}(\mathbf{x}) - \text{LN}(\hat{\mathbf{x}})\| \leq L_{\text{LN}} \|\mathbf{x} - \hat{\mathbf{x}}\|. \quad (27)$$

The FFN is a two-layer MLP with GELU activation, which is $\sqrt{2/\pi}$ -Lipschitz, so the FFN sub-layer is Lipschitz with constant $L_{\text{FFN}} = \|W_1\|_{\text{op}} \cdot \|W_2\|_{\text{op}} \cdot \sqrt{2/\pi}$.

Composing all sub-layers of a single block, we obtain a per-block Lipschitz constant:

$$L_{\text{block}} = L_{\text{LN}}(1 + L_{\text{FFN}}) \cdot L_{\text{LN}}(1 + L_{\text{MHSA}}), \quad (28)$$

where L_{MHSA} absorbs the per-head bounds from Eq. (24) summed across all H heads.

C.4 Step 3: Composition through suffix blocks

Let $\mathbf{A}^{(0)} = \mathbf{A}$, $\hat{\mathbf{A}}^{(0)} = \hat{\mathbf{A}}$, and define $\mathbf{A}^{(t)}, \hat{\mathbf{A}}^{(t)}$ as the token sequences after t suffix blocks. Applying the per-block bound recursively:

$$\|\mathbf{A}^{(t)} - \hat{\mathbf{A}}^{(t)}\|_F \leq L_{\text{block}}^t \|\mathbf{A}^{(0)} - \hat{\mathbf{A}}^{(0)}\|_F + \sum_{k=0}^{t-1} L_{\text{block}}^{t-1-k} \eta_{\text{gram}}^{(k)}, \quad (29)$$

where $\eta_{\text{gram}}^{(k)}$ is the additional error injected at block k due to the Gram-matrix deviation of $\mathbf{A}^{(k)}$. In the worst case, $\eta_{\text{gram}}^{(k)} \leq c_k \cdot \|\mathbf{G}^{(k)}\|_F \cdot \delta_{\text{gram}}^{(k)}$ for constants c_k depending on the attention projections. At $k = 0$ this is directly controlled by our loss; at subsequent blocks it is controlled indirectly through the accumulated deviation.

After all $L_{\text{suf}} = L - \ell$ suffix blocks, the CLS token deviation satisfies:

$$\|z_{\text{cls}} - \hat{z}_{\text{cls}}\|_2 \leq \|\mathbf{A}^{(L_{\text{suf}})} - \hat{\mathbf{A}}^{(L_{\text{suf}})}\|_F \leq L_{\text{block}}^{L_{\text{suf}}} (\sqrt{N} \delta_{\text{rec}} + \alpha \|\mathbf{G}\|_F \cdot \delta_{\text{gram}}), \quad (30)$$

where α collects the constants from the attention projections and the geometric series, and $\sqrt{N}$ carries forward the sequence-length scaling introduced by the stochastic attention matrix in Eq. (24). This makes explicit that both δ_{rec} and δ_{gram} contribute to controlling the final representation.

C.5 Step 4: From CLS-token deviation to KL divergence

Let $h(z) = W_c z + b$ be a linear classifier mapping the CLS token to C -class logits, which we denote as $\mathbf{u}$. The logit deviation is bounded by:

$$\|\mathbf{u} - \hat{\mathbf{u}}\|_2 = \|W_c(z_{\text{cls}} - \hat{z}_{\text{cls}})\|_2 \leq \|W_c\|_{\text{op}} \|z_{\text{cls}} - \hat{z}_{\text{cls}}\|_2. \quad (31)$$

To rigorously convert this logit deviation into a KL divergence bound without introducing dependencies on logit magnitudes, we exploit the connection between KL divergence and Bregman divergences.

Let $A(\mathbf{u}) = \log \sum_{c=1}^C \exp(\mathbf{u}_c)$ be the log-partition function. Its gradient yields the softmax probabilities, $\nabla A(\mathbf{u}) = p$, and its Hessian is the covariance matrix of the categorical distribution, $\nabla^2 A(\mathbf{u}) = \text{diag}(p) - pp^\top$. Expressing the KL divergence $D_{\text{KL}}(p \parallel \hat{p})$ in terms of the logits gives:

$$\begin{aligned} D_{\text{KL}}(p \parallel \hat{p}) &= \sum_{c=1}^C p_c \log \frac{\exp(\mathbf{u}_c) / \exp(A(\mathbf{u}))}{\exp(\hat{\mathbf{u}}_c) / \exp(A(\hat{\mathbf{u}}))} \\ &= \sum_{c=1}^C p_c (\mathbf{u}_c - \hat{\mathbf{u}}_c) - A(\mathbf{u}) + A(\hat{\mathbf{u}}) \\ &= A(\hat{\mathbf{u}}) - A(\mathbf{u}) - \nabla A(\mathbf{u})^\top (\hat{\mathbf{u}} - \mathbf{u}). \end{aligned} \quad (32)$$

Equation (32) reveals that the KL divergence is exactly the Bregman divergence $D_A(\hat{\mathbf{u}}, \mathbf{u})$ generated by the log-partition function. By the multivariate Taylor's theorem, this error is precisely the second-order remainder:

$$D_{\text{KL}}(p \parallel \hat{p}) = \frac{1}{2} (\mathbf{u} - \hat{\mathbf{u}})^\top H(\xi) (\mathbf{u} - \hat{\mathbf{u}}), \quad (33)$$

where $H(\xi) = \text{diag}(q) - qq^\top$ is the Hessian evaluated at an intermediate point $\xi = t\mathbf{u} + (1-t)\hat{\mathbf{u}}$ (for some $t \in [0, 1]$) lying on the line segment connecting the two logit vectors, with corresponding softmax probabilities $q = \sigma(\xi)$.

Because $qq^\top$ is positive semi-definite, we have $H(\xi) \preceq \text{diag}(q)$. Since all probabilities $q_c \in (0, 1)$, the maximum eigenvalue of $\text{diag}(q)$ is strictly bounded by 1. Consequently, the spectral norm of the Hessian satisfies $\|H(\xi)\|_{\text{op}} \leq 1$ globally, regardless of the magnitude of the intermediate logits. Applying this to Equation (33) yields a strict, magnitude-independent quadratic bound:

$$D_{\text{KL}}(p \parallel \hat{p}) \leq \frac{1}{2} \|\mathbf{u} - \hat{\mathbf{u}}\|_2^2. \quad (34)$$

Substituting Eq. (31) and Eq. (30) into Eq. (34):

$$D_{\text{KL}}(\sigma(h \circ g_{>\ell}(\mathbf{A})) \parallel \sigma(h \circ g_{>\ell}(\hat{\mathbf{A}}))) \leq \frac{1}{2} \|W_c\|_{\text{op}}^2 L_{\text{block}}^{2L_{\text{suf}}} (\sqrt{N} \delta_{\text{rec}} + \alpha \|\mathbf{G}\|_F \cdot \delta_{\text{gram}})^2. \quad (35)$$

C.6 Discussion

What the bound identifies. Equation (35) attributes predictive consistency jointly to δ_{rec} and δ_{gram} . The decomposition $M = I + (M - I)$ in Section C.2 makes the roles of the two terms explicit: the identity component of every query-key product is controlled directly by $\mathcal{L}_{\text{gram}}$, since preserving $\mathbf{A}\mathbf{A}^\top$ is exactly preserving the bilinear form at $M = I$, while the residual $(M - I)$ is controlled by $\mathcal{L}_{\text{rec}}$ through Eq. (20). The two losses are therefore not redundant in the worst case: $\mathcal{L}_{\text{gram}}$ supplies a direct constraint on the attention-logit deviation that $\mathcal{L}_{\text{rec}}$ bounds only indirectly via $\|\mathbf{A}\|_F \cdot \delta_{\text{rec}}$ (Lemma 1), while $\mathcal{L}_{\text{rec}}$ supplies the pointwise activation fidelity that the Gram loss, being invariant under any orthogonal transformation acting on the right of $\mathbf{A}$, cannot constrain at all. This complementarity is the formal content of the claim that geometric preservation plus reconstruction fidelity substitute for classifier supervision.

Reconciling backbone-dependent effects. Lemma 1 also clarifies why the empirical benefit of $\mathcal{L}_{\text{gram}}$ varies across architectures (Tables 1 and 2). The indirect Gram control provided by reconstruction scales as $2\|\mathbf{A}\|_F \cdot \delta_{\text{rec}} + \delta_{\text{rec}}^2$. For ViT-MAE, the masked autoencoder pretraining objective produces activations whose pairwise geometry is already well-captured by reconstruction alone — the product $\|\mathbf{A}\|_F \cdot \delta_{\text{rec}}$ is small enough that the indirect bound is tight, explaining why Table 2 shows minimal MCS improvement from +Gram ($0.921 \rightarrow 0.923$). Conversely, DINOv2’s contrastive pretraining yields activations with richer pairwise structure and larger $\|\mathbf{A}\|_F$, so the indirect bound is looser and the explicit, scale-normalized Gram constraint provides a measurably tighter handle ($\text{MCS } 0.840 \rightarrow 0.857$).

What the bound does not characterise. The operator norm $\|W_Q W_K^\top - I\|_{\text{op}}$ and its per-head distribution determine how much of the logit deviation is absorbed by the Gram term versus the reconstruction term in practice; we do not measure it. We also do not track how $\mathcal{L}_{\text{gram}}$ at layer ℓ propagates to the Gram matrices of subsequent layers: the bound preserves only the ℓ -th Gram matrix and invokes Lipschitz composition (Eq. (28)) for everything beyond.

Worst-case bounds vs. empirical tightness. As is standard in the theoretical analysis of deep neural networks, our worst-case bound compounds multiplicatively across layers ($L_{\text{block}}^{2L_{\text{suf}}}$). Because it assumes adversarial alignment of sublayer errors, the theoretical bound is looser than the empirical divergences observed in Table 1. LayerNorm, in particular, acts as a strong empirical contractor for low-rank signals, a dynamic not captured by worst-case Lipschitz constants.

Therefore, Equation (35) should be interpreted not as a strict numerical predictor, but as a structural identification of the sufficient control variables: it formally shows that $\mathcal{L}_{\text{gram}}$ and $\mathcal{L}_{\text{rec}}$ are the quantities whose joint minimization controls downstream predictive divergence in a task-agnostic manner.

Dependence on extraction depth (ℓ). The looseness of the bound is strictly governed by the number of suffix blocks, $L_{\text{suf}} = L - \ell$. Extracting concepts closer to the classifier (increasing ℓ) tightens the theoretical bound by reducing the exponent in $L_{\text{block}}^{2L_{\text{suf}}}$. In the extreme case of extracting at the final layer ($\ell = L$, so $L_{\text{suf}} = 0$), the transformer composition term vanishes entirely ($L_{\text{block}}^0 = 1$), and the KL divergence bound collapses to:

$$D_{\text{KL}}(\sigma(h(\mathbf{A})) \parallel \sigma(h(\hat{\mathbf{A}}))) \leq \frac{1}{2} \|W_c\|_{\text{op}}^2 (\sqrt{N} \delta_{\text{rec}} + \alpha \|\mathbf{G}\|_F \cdot \delta_{\text{gram}})^2. \quad (36)$$

However, this introduces a representation trade-off: while the predictive consistency bound becomes remarkably tight, representations at deeper layers become increasingly task-specific, potentially degrading the part-level interpretability of the Semi-NMF concepts compared to intermediate layers.

What the bound does, and does not, justify about the method. Eq. (35) is a sufficiency statement: it identifies $\mathcal{L}_{\text{rec}}$ and $\mathcal{L}_{\text{gram}}$ as two quantities whose joint minimization controls predictive consistency through any linear classifier attached downstream, without requiring that classifier during concept extraction. It is not a necessity statement, and does not exclude alternative activation-space losses (for example, higher-order moment matching or layer-wise CKA) from achieving comparable control. The empirical question of whether $\mathcal{L}_{\text{gram}}$ recovers the same concept subspace as full classifier supervision, and whether it does so more faithfully than reconstruction alone, is answered in Section 4.3 via the RSVC framework. The empirical agreement is much stronger than this worst-case bound predicts. However, as expected from our analysis, the bound’s true value lies in identifying the correct mechanism rather than providing tight numerical guarantees.

D Implementation Details

D.1 Linear probes for validation

To evaluate our unsupervised concept banks against supervised ones via the RSVC framework, we attach a lightweight linear probe on top of each frozen backbone. The probe consists of a single linear layer ($d \rightarrow C$) applied to the CLS token. We train for 10 epochs using the Adam optimizer with a learning rate of 10^{-4} , no weight decay, and a batch size of 64. Table 4 reports the probe accuracy and macro-F1 scores on COCO for each backbone. The probes for DINOv2, ViT, and ViT-MAE on ImageNet are not finetuned; rather, we utilize the classification heads distributed with the respective checkpoints: for ViT we leverage timm (i.e. for vit_base_patch16_224 we use augreg_in21k_ft_in1k), we get dinov2_vitb14_lc from torch.hub, and for ViT-MAE we use the Meta AI version available from https://dl.fbaipublicfiles.com/mae/finetune/mae_finetuned_vit_base.pth that uses global average pooling + fully connected norm instead of the CLS token as the classifier-head input (this because the CLS token is not trained on ViT-MAE).

Table 4: Linear probe performance per backbone on COCO. All probes share the same architecture (CLS pooling, $d=768$) and training schedule.

Backbone	Accuracy (%)	Macro F1 (%)
DINOv2 ViT-B/14	95.5	95.5
ViT-B/16	93.8	93.8
ViT-MAE B/16	90.0	90.0

D.2 PLACE Hyperparameters

Table 5 lists all hyperparameters used by PLACE for concept extraction. These are shared across all three loss variants (reconstruction, +Gram, +KL) and across all backbones and datasets, unless otherwise noted. We use five seeds ($s \in \{42, 43, 44, 45, 46\}$) on all reported results; tables report mean $\pm$ standard deviation across seeds.

Table 5: PLACE concept extraction hyperparameters.

Parameter	Value	Description
extract_layer	9	Transformer block from which patch tokens are extracted
group_size s	2	Spatial grouping factor ($s \times s$ patches $\rightarrow$ 1 token)
rank r	25	Number of concepts per class
steps	22,000	Maximum optimization steps (early stopping on plateau)
learning_rate	5×10^{-4}	AdamW learning rate
batch_size	4	Mini-batch size for optimization
λ	1.0	Weight of the faithfulness loss ($\mathcal{L}_{\text{gram}}$ or $\mathcal{L}_{\text{KL}}$)
ϵ	10^{-6}	Early stopping criterion
$\ell_1(U)$	10^{-6}	ℓ_1 regularization on the coefficient matrix U
$\ell_2(W)$	10^{-6}	ℓ_2 regularization on the basis matrix W
NNDSVD %	80	Fraction of rows subsampled for NNDSVD initialization

The regularization terms $\ell_1(U)$ and $\ell_2(W)$ are set to small values (10^{-6}) to mildly encourage sparsity in the concept activations and prevent unbounded growth of the basis vectors, respectively, without materially affecting reconstruction quality. The faithfulness weight $\lambda=1.0$ gives equal importance to the reconstruction and geometric/KL loss terms; a sweep over λ is reported in Appendix E.2. It is worth noting that the batch_size is used for every loss (Gram, KL and MSE).

Sparse Autoencoders. SAEs baseline comparisons follow the same parameters where applicable, with a learning rate of 10^{-3} and a ℓ_1 penalty of 10^{-3} .

D.3 PCA Visualization Details

To visualize the raw patch-token activations in Appendix H, we project the high-dimensional token representations into RGB space. We follow a two-step principal component analysis (PCA) approach to separate the foreground object from the background and highlight intra-object feature variance. This process is applied independently to each image.

First, we extract the raw patch tokens $\mathbf{A} \in \mathbb{R}^{N \times d}$ and fit a 1D PCA over all N tokens. The first principal component typically captures the highest variance direction in the representation space, effectively acting as a foreground-background discriminator. Tokens with a positive projection score are designated as foreground. If this produces a foreground mask covering less than 30% of the patches, we assume the polarity is inverted and flip the sign of the scores.

Second, to capture the feature structure without interference from background variance, we fit a 3D PCA strictly on the foreground tokens. We then project all N tokens through these three components. To generate the final RGB values, we normalize each of the three component channels to $[0, 1]$ using the minimum and maximum values computed exclusively from the foreground tokens. Finally, the background patches are dimmed by a factor of 0.3, and the resulting $\sqrt{N} \times \sqrt{N}$ grid is bilinearly upsampled to the original input resolution $H \times W$.

E Additional Analyses

In this section, we present comprehensive analyses to justify our architectural choices and hyperparameter settings. Unless otherwise specified, all ablations are performed on the ImageNet dataset using the *church* class. We examine the impact of the NNDSVD subsampling ratio (Appendix E.1), the faithfulness loss weight λ (Appendix E.2), the extraction layer ℓ (Appendix E.3), and the decomposition rank r (Appendix E.4). Furthermore, we include a weight randomization sanity check (Appendix E.6), an analysis of the stability of our factorization across different random seeds (Appendix E.7) and runtime measurements (Appendix E.9). Finally we discuss the computational complexity and scalability of PLACE (Appendix E.10), the robustness to image perturbations (Appendix E.13), and the complete, per-class results (Appendix E.12). The default configuration used throughout the main experiments (Section 4) is listed in Appendix D.2.

E.1 NNDSVD subsampling ratio for Semi-NMF initialization

Computing the full SVD of the activation matrix $\mathbf{A}$ is expensive for large datasets; a natural question is whether computing NNDSVD on a random subset of rows degrades the final solution.

We run concept extraction on the ImageNet *church* class (878 correctly-classified images) using ViT-B/16 (timm, augreg2_in21k_ft_in1k) at layer $\ell=9$ with rank $r=25$ and group size $s=2$, varying the fraction of rows subsampled for NNDSVD initialization: 100%, 80%, 60%, 40%, and 20%. All runs use the MSE+KL loss with a budget of 22,000 steps and early stopping.

Table 6: Effect of NNDSVD subsampling ratio on Semi-NMF concept extraction (ViT-B/16, ImageNet *church*, $r=25$, $s=2$). Mean $\pm$ std over 5 seeds.

NNDSVD %	MSE $\downarrow$	$D_{\text{KL}} \downarrow$	Steps $\downarrow$	C-Ins $\uparrow$	C-Del $\uparrow$	C-Gini $\uparrow$
100	0.306 $\pm$ 0.002	0.011 $\pm$ 0.001	12100 $\pm$ 7477	0.903 $\pm$ 0.032	0.162 $\pm$ 0.014	0.458 $\pm$ 0.014
80	0.302 $\pm$ 0.002	0.015 $\pm$ 0.003	4847 $\pm$ 2265	0.911 $\pm$ 0.003	0.147 $\pm$ 0.023	0.439 $\pm$ 0.014
60	0.309 $\pm$ 0.002	0.013 $\pm$ 0.002	6012 $\pm$ 4814	0.906 $\pm$ 0.048	0.138 $\pm$ 0.022	0.460 $\pm$ 0.016
40	0.303 $\pm$ 0.004	0.015 $\pm$ 0.007	3896 $\pm$ 2252	0.913 $\pm$ 0.031	0.123 $\pm$ 0.063	0.462 $\pm$ 0.013
20	0.303 $\pm$ 0.003	0.013 $\pm$ 0.003	5587 $\pm$ 4162	0.912 $\pm$ 0.008	0.130 $\pm$ 0.055	0.452 $\pm$ 0.010
0	0.304 $\pm$ 0.003	0.014 $\pm$ 0.002	6992 $\pm$ 141	0.764 $\pm$ 0.073	0.127 $\pm$ 0.026	0.312 $\pm$ 0.039

Discussion. The results show that Semi-NMF is remarkably robust to the NNDSVD subsampling ratio in terms of final reconstruction accuracy. Reconstruction error (MSE) varies marginally across all settings, hovering between 0.302 and 0.309, while the KL divergence remains tightly bounded in the range $[0.011, 0.015]$. Even when entirely abandoning NNDSVD in favor of random initialization (0%), the model achieves comparable reconstruction metrics.

However, logging the optimization steps required for convergence reveals a high degree of variance in the optimization trajectory. Full 100% NNDSVD initialization required the highest average number of steps (12100 ± 7477) before triggering early stopping. Subsampling the initialization matrix reduced the mean convergence time: 80% took 4847 ± 2265 steps, and 40% took 3896 ± 2252 steps. Random initialization (0%) required a moderate 6992 ± 141 steps, notably exhibiting the lowest variance in convergence time across seeds.

Crucially, while reconstruction quality is agnostic to the initialization strategy, the interpretability and faithfulness of the extracted concepts are highly dependent on NNDSVD. The 0% setting suffers a degradation in concept quality: C-Ins drops precipitously to 0.764 ± 0.073 (compared to > 0.900 for all NNDSVD runs), and C-Gini falls to 0.312 ± 0.039 , indicating less sparse and interpretable concepts. As long as NNDSVD is used, even at just 20% subsampling, faithfulness metrics remain strong and stable (e.g., C-Ins of 0.912 ± 0.008). C-Del is highest (best) at 100% (0.162), with a monotonic but mild degradation as subsampling decreases, though random initialization performs comparably to the lower NNDSVD ratios.

In practice, subsampling to 40%–80% offers a superior trade-off. It avoids the decrease in interpretability observed with random initialization, while reducing both the initial NNDSVD compute requirements and the number of subsequent optimization steps needed to converge, all while retaining full-data solution quality.

E.2 Faithfulness loss weight sensitivity

We sweep the faithfulness loss weight λ in $\mathcal{L} = \mathcal{L}_{\text{MSE}} + \lambda \mathcal{L}_{\text{faith}}$ across seven values from 0.01 to 10 for both the KL and Gram variants, and report the effect on reconstruction quality, faithfulness, and concept sparsity. All experiments use ViT-B/16 (timm, augreg2_in21k_ft_in1k), ImageNet *church* class, layer $\ell=9$, $r=25$, $s=2$; values are mean $\pm$ std over five seeds.

Table 7: Effect of λ on concept extraction with the KL loss (ViT-B/16, ImageNet *church*, $r=25$, $s=2$). Mean $\pm$ std over 5 seeds.

λ	MSE $\downarrow$	$D_{\text{KL}} \downarrow$	C-Ins $\uparrow$	C-Del $\uparrow$	C-Gini $\uparrow$
0.01	0.296 ± 0.002	0.106 ± 0.015	0.882 ± 0.019	0.088 ± 0.023	0.402 ± 0.005
0.1	0.298 ± 0.003	0.033 ± 0.010	0.879 ± 0.026	0.110 ± 0.028	0.463 ± 0.011
0.5	0.300 ± 0.002	0.017 ± 0.003	0.887 ± 0.031	0.136 ± 0.026	0.445 ± 0.009
1.0	0.302 ± 0.002	0.015 ± 0.003	0.911 ± 0.003	0.147 ± 0.023	0.439 ± 0.014
2.0	0.302 ± 0.002	0.011 ± 0.002	0.917 ± 0.037	0.156 ± 0.031	0.433 ± 0.013
5.0	0.306 ± 0.001	0.008 ± 0.001	0.933 ± 0.009	0.149 ± 0.029	0.425 ± 0.007
10.0	0.312 ± 0.003	0.007 ± 0.002	0.932 ± 0.011	0.143 ± 0.044	0.423 ± 0.012

Table 8: Effect of λ on concept extraction with the Gram loss (ViT-B/16, ImageNet *church*, $r=25$, $s=2$). Mean $\pm$ std over 5 seeds.

λ	MSE $\downarrow$	$\mathcal{L}_{\text{gram}} \downarrow$	C-Ins $\uparrow$	C-Del $\uparrow$	C-Gini $\uparrow$
0.01	0.287 ± 0.001	0.018 ± 0.001	0.710 ± 0.020	0.309 ± 0.087	0.600 ± 0.010
0.1	0.292 ± 0.002	0.023 ± 0.003	0.710 ± 0.020	0.300 ± 0.077	0.580 ± 0.010
0.5	0.292 ± 0.010	0.024 ± 0.020	0.710 ± 0.020	0.319 ± 0.087	0.580 ± 0.010
1.0	0.293 ± 0.003	0.022 ± 0.010	0.707 ± 0.028	0.300 ± 0.011	0.571 ± 0.017
2.0	0.281 ± 0.006	0.010 ± 0.002	0.710 ± 0.020	0.300 ± 0.087	0.560 ± 0.020
5.0	0.282 ± 0.007	0.010 ± 0.002	0.700 ± 0.030	0.309 ± 0.058	0.560 ± 0.020
10.0	0.282 ± 0.014	0.012 ± 0.007	0.690 ± 0.020	0.338 ± 0.077	0.520 ± 0.060

Discussion. The two losses respond to λ in qualitatively different ways, reflecting the distinction between a supervised and an unsupervised faithfulness objective.

For the KL variant (Table 7), increasing λ produces a clear MSE–faithfulness trade-off: MSE rises monotonically from 0.296 to 0.312, while D_{KL} drops by more than an order of magnitude ($0.106 \rightarrow 0.007$). C-Ins improves steadily from 0.882 to 0.933, with the sharpest gain between

$\lambda=0.5$ and $\lambda=5$. C-Del improves rapidly from 0.088 at $\lambda=0.01$ to ~ 0.147 at $\lambda=1$ and then plateaus, indicating that moderate KL pressure suffices for faithful concept ordering.

The Gram variant (Table 8) tells a different story. Because $\mathcal{L}_{\text{gram}}$ constrains the pairwise geometry of the reconstruction without reference to the downstream classifier, its influence on C-Ins and C-Del is second-order: C-Ins varies only between 0.69 and 0.71, and C-Del remains in the 0.30–0.34 range across all λ values. MSE is similarly stable (0.28–0.29), and C-Gini decreases mildly from 0.60 to 0.52. The Gram loss itself saturates around ~ 0.010 for $\lambda \geq 2$, with diminishing returns beyond. The main practical difference is convergence speed: at $\lambda=0.01$ the optimizer converges in ~ 300 steps, while $\lambda \geq 2$ requires 1,000–3,000 steps.

Comparing the two, the KL loss is the more responsive lever for controlling faithfulness but introduces a genuine trade-off with reconstruction quality. The Gram loss is robust to λ across a wide range, making it the safer default for unsupervised settings where the practitioner cannot tune λ against a classifier. We select $\lambda=1$ for both variants: for the KL loss this balances low D_{KL} , competitive C-Ins, and minimal MSE penalty; for the Gram loss the choice is largely inconsequential but ensures consistency across the experimental setup.

E.3 Extraction layer ablation

We extract concepts at layers $\ell \in \{4, 6, 8, 9, 10\}$ of a 12-block ViT and report the effect on concept quality metrics. For ViT-B/16 (timm), which has a classification head, we use the MSE+KL loss and report C-Ins and C-Del alongside reconstruction error. For DINOv2 ViT-B/14, which lacks a classification head, we use the MSE+Gram loss and report $\mathcal{L}_{\text{gram}}$ and Gini sparsity. All runs use $r=25$, $s=2$, $\lambda=1$; values are mean $\pm$ std over five seeds.

Table 9: Layer ablation on ViT-B/16 (timm, MSE+KL, ImageNet *church*, $r=25$, $s=2$). † Layer 10 is degenerate: with only two transformer blocks remaining, representations are highly aggregated, and the factorization heavily overfits the nearly-formed class prediction rather than discovering part-level concepts.

ℓ	MSE $\downarrow$	D_{KL} $\downarrow$	C-Ins $\uparrow$	C-Del $\uparrow$	C-Gini $\uparrow$
4	0.053 ± 0.001	0.040 ± 0.013	0.403 ± 0.035	0.848 ± 0.004	0.361 ± 0.008
6	0.086 ± 0.000	0.022 ± 0.002	0.702 ± 0.018	0.889 ± 0.003	0.370 ± 0.002
8	0.222 ± 0.001	0.014 ± 0.002	0.696 ± 0.022	0.641 ± 0.016	0.413 ± 0.015
9	0.302 ± 0.002	0.015 ± 0.003	0.911 ± 0.003	0.147 ± 0.023	0.439 ± 0.014
10 [†]	0.621 ± 0.004	0.010 ± 0.001	0.948 ± 0.006	0.161 ± 0.042	0.434 ± 0.011

Table 10: Layer ablation on DINOv2 ViT-B/14 (MSE+Gram, ImageNet *church*, $r=25$, $s=2$).

ℓ	MSE $\downarrow$	$\mathcal{L}_{\text{gram}}$ $\downarrow$	C-Gini $\uparrow$
4	0.013 ± 0.001	0.042 ± 0.029	0.459 ± 0.053
6	0.012 ± 0.001	0.019 ± 0.003	0.605 ± 0.005
8	0.017 ± 0.001	0.020 ± 0.002	0.803 ± 0.008
9	0.093 ± 0.001	0.017 ± 0.004	0.923 ± 0.004
10	0.104 ± 0.001	0.019 ± 0.000	0.588 ± 0.009

Discussion. The two models exhibit complementary trends across depth. For ViT-B/16 with the KL loss, C-Ins increases sharply with depth, jumping from 0.403 at $\ell=4$ to 0.911 at $\ell=9$. This is expected: deeper layers encode increasingly class-discriminative features, so concepts extracted there are more faithful to the downstream prediction. Layer 10 achieves the highest raw C-Ins (0.948) but is functionally degenerate (marked †): at this depth (with only two transformer blocks remaining before the classifier), the representations are highly aggregated and task-specific. The CLS token has already accumulated global evidence, and reconstructing the patch tokens at this stage allows the Semi-NMF to trivially overfit the nearly-formed class logits. This inflates both C-Ins and classification accuracy (0.987 vs. a 0.984 ungrouped baseline) without reflecting genuine, spatially-grounded concept structure.

Conversely, C-Del exhibits a decreasing pattern: it is exceptionally high at layers 4-6 (0.889 at $\ell=6$), but drops to 0.147 at layer 9. High C-Del at early layers indicates that these concepts capture fundamental, low-level visual features; removing them completely breaks the model’s ability to process the image. At layer 9, concepts are more semantically distinct (higher C-Gini, 0.439) but the model exhibits higher redundancy, meaning it can maintain its prediction even when individual high-level concepts are deleted. We select $\ell=9$ as the default extraction layer for ViT-B/16 to prioritize the recovery of semantic, class-specific concepts (C-Ins) and sparsity (C-Gini), while acknowledging the trade-off in deletion sensitivity.

For DINOv2 ViT-B/14 with the Gram loss, MSE increases with depth (from 0.013 at $\ell=4$ to 0.104 at $\ell=10$), reflecting the growing complexity of later-layer representations. Gini sparsity increases significantly with depth, peaking at $\ell=9$ (0.923). Notably, at $\ell=10$, the factorization begins to collapse: sparsity drops sharply to 0.588. Similar to the supervised ViT, this indicates that representations at $\ell \geq 10$ are too globally aggregated to be meaningfully factorized into sparse, patch-level parts. Layer 9, used as the default throughout the paper, sits at the optimal point where representations are most distinct and specialized before this architectural collapse.

E.4 Decomposition rank ablation

We analyze the effect of the decomposition rank $r \in \{5, 10, 15, 20, 25, 30\}$ on faithfulness and concept quality. Table 11 and Table 12 report metrics averaged over five seeds on the church class.

Faithfulness (Table 11). C-Del decreases from 0.230 at $r=10$ to 0.125 at $r=30$, while C-Ins remains stable at ~ 0.90 across all ranks, indicating that even a small number of concepts suffices to recover most of the target-class logit under progressive insertion. C-Gini increases monotonically, reflecting greater per-concept specialization at higher ranks. We note that at $r=5$ two of five seeds converge to a degenerate solution (high C-Del, low C-Ins), inflating the variance; this instability disappears at $r \geq 10$, suggesting that very low rank under-determines the KL-constrained factorization. Gains plateau between $r=25$ and $r=30$, motivating our choice of $r=25$ for the main experiments.

Reconstruction and stability (Table 12). MSE decreases monotonically from 0.099 to 0.092 as rank increases, confirming that additional concepts capture progressively finer structure. This trade-off between reconstruction fidelity and interpretability is expected: at higher rank, the decomposition captures more nuanced features, but risks fragmenting cohesive visual concepts. Gains plateau between $r=25$ and $r=30$, validating our choice of $r=25$ to balance granularity and conciseness.

Table 11: Rank ablation for PLACE (+KL) on the church class (ViT-B/16, $s=2$). Mean $\pm$ std over 5 seeds. $\downarrow$ = lower is better, $\uparrow$ = higher is better.

r	C-Ins $\uparrow$	C-Del $\uparrow$	C-Gini $\uparrow$
5	0.88 ± 0.13	0.419 ± 0.258	0.34 ± 0.03
10	0.90 ± 0.02	0.230 ± 0.036	0.45 ± 0.01
15	0.90 ± 0.02	0.131 ± 0.056	0.48 ± 0.01
20	0.87 ± 0.05	0.174 ± 0.038	0.52 ± 0.01
25	0.91 ± 0.01	0.147 ± 0.023	0.44 ± 0.01
30	0.91 ± 0.04	0.125 ± 0.058	0.58 ± 0.01

Table 12: Rank ablation for PLACE (+Gram) on the church class (DINOv2 ViT-B/14, $s=2$). Mean $\pm$ std over 5 seeds.

r	MSE $\downarrow$	$\mathcal{L}_{\text{gram}}$ $\downarrow$	C-Gini $\uparrow$
5	0.099 ± 0.002	0.048 ± 0.002	0.64 ± 0.04
10	0.096 ± 0.001	0.031 ± 0.002	0.74 ± 0.03
15	0.094 ± 0.002	0.028 ± 0.006	0.83 ± 0.03
20	0.094 ± 0.000	0.034 ± 0.013	0.91 ± 0.00
25	0.093 ± 0.001	0.017 ± 0.004	0.923 ± 0.004
30	0.092 ± 0.001	0.017 ± 0.011	0.93 ± 0.02

E.5 Token group size ablation

We evaluate the effect of the token group size $s \in \{1, 2, 4, 7\}$ on faithfulness, reconstruction quality, and runtime. As discussed in Section 3 and Appendix B.3, larger s reduces computational cost but coarsens the spatial resolution of concept heatmaps and introduces quantization error through the `repeat_interleave` operation.

Table 13: Effect of token group size s on ViT-B/16 (ImageNet, $r=25$, PLACE (+KL)). Grouped and ungrouped MSE disentangle quantization error from factorization error. Values are reported as mean $\pm$ standard deviation.

s	C-Ins $\uparrow$	C-Del $\uparrow$	MSE (grouped) $\downarrow$	MSE (ungrouped) $\downarrow$	Acc (Rec) $\uparrow$	Acc (Ungrouped) $\uparrow$	Time (s) $\downarrow$
1	0.913 $\pm$ 0.063	0.177 $\pm$ 0.044	0.193 $\pm$ 0.001	-	0.980 $\pm$ 0.003	-	348.538 $\pm$ 53.164
2	0.962 $\pm$ 0.013	0.132 $\pm$ 0.008	0.093 $\pm$ 0.002	0.089 $\pm$ 0.002	0.990 $\pm$ 0.002	0.945 $\pm$ 0.002	125.419 $\pm$ 6.041
4	0.918 $\pm$ 0.025	0.112 $\pm$ 0.015	0.126 $\pm$ 0.005	0.093 $\pm$ 0.004	0.801 $\pm$ 0.024	0.731 $\pm$ 0.001	118.969 $\pm$ 2.400
7	0.869 $\pm$ 0.059	0.165 $\pm$ 0.020	0.077 $\pm$ 0.001	0.077 $\pm$ 0.001	0.875 $\pm$ 0.006	0.733 $\pm$ 0.003	117.602 $\pm$ 2.102
14	0.917 $\pm$ 0.057	0.101 $\pm$ 0.020	0.036 $\pm$ 0.000	0.036 $\pm$ 0.000	0.709 $\pm$ 0.003	0.615 $\pm$ 0.004	115.301 $\pm$ 2.059

Discussion. The results in Table 13 reveal a clear trade-off between computational efficiency and spatial resolution.

Dense factorization ($s = 1$) preserves the native token grid but incurs a computational bottleneck (~ 348 s). Setting $s = 2$ emerges as the optimal configuration: it reduces the optimization runtime by nearly $3 \times$ (~ 125 s) while simultaneously achieving the highest concept insertion score (C-Ins 0.962) and the highest post-reconstruction target accuracy (0.990). This indicates that a moderate grouping of 2×2 tokens effectively regularises the factorization, mitigating local variance among adjacent tokens and improving the extraction of predictively relevant concepts. While C-Del sees a slight drop compared to $s = 1$, the massive gains in C-Ins, accuracy, and speed make $s = 2$ the strictly superior choice.

For larger group sizes ($s \geq 4$), spatial quantization error becomes detrimental. Grouping 4×4 or larger blocks of tokens aggressively smooths out fine-grained visual features, causing a sharp collapse in target-class accuracy (dropping to 0.801 at $s = 4$ and further deteriorating at larger sizes).

Notably, the reconstruction error (MSE) tends to decrease at extreme grouping levels (e.g., 0.036 at $s = 14$). We emphasize that this does not indicate superior factorization quality; rather, heavy spatial pooling creates a significantly smaller, lower-variance activation matrix $\mathbf{A}$, which is trivially easier to fit with a low-rank approximation but devoid of meaningful spatial semantics. Based on these findings, we adopt $s = 2$ as the default hyperparameter, providing the ideal balance between fine-grained localization, predictive faithfulness, and computational tractability.

E.6 Sanity check: weight randomization

Following Fel et al. [25], we verify that our method’s explanations degrade when the model’s weights are randomized. We progressively randomize transformer blocks from the first to the last ($n \in \{0, 2, 4, 6, 8, 10\}$ out of 12 blocks) and re-extract concepts at layer $\ell=9$. If the discovered concepts genuinely reflect learned representations rather than artifacts of the factorization, concept quality should degrade as more blocks are randomized.

Discussion. MSE decreases monotonically with randomization ($0.302 \rightarrow 0.065$), which is expected: random features have lower effective dimensionality and are more compressible by a rank-25 factorization; low reconstruction error under a randomized model should not be confused with explanation quality. The sanity check confirms that our method’s explanations reflect learned representations. Randomizing even two blocks reduces target-class accuracy to 0%, collapsing C-Ins and C-Del to zero: the concepts extracted from a corrupted model carry no predictive information. C-Gini drops from 0.439 (intact model) to 0.271 at $n=6$, indicating that the coefficient matrix becomes uniformly diffuse rather than selectively activating distinct concepts. At $n=8$ and $n=10$, Gini partially recovers (0.371, 0.502) as one or two concepts absorb the dominant random direction, yielding a degenerate form of sparsity that does not correspond to meaningful visual concepts.

Table 14: Weight randomization sanity check. n denotes the number of randomized blocks (from input side). C-Ins and C-Del collapse to 0 for $n \geq 2$ because the randomized model has 0% target-class accuracy. We report results on ViT-B/16 (timm, MSE+KL, ImageNet *church*, $r=25$, $s=2$), mean $\pm$ std over five seeds.

n	MSE $\downarrow$	D_{KL} $\downarrow$	C-Gini $\uparrow$	C-Ins $\uparrow$	C-Del $\uparrow$
0	0.302 ± 0.002	0.015 ± 0.003	0.439 ± 0.014	0.911 ± 0.003	0.147 ± 0.023
2	0.181 ± 0.001	0.022 ± 0.003	0.332 ± 0.021	0.0	0.0
4	0.150 ± 0.002	0.020 ± 0.005	0.310 ± 0.014	0.0	0.0
6	0.136 ± 0.002	0.048 ± 0.015	0.271 ± 0.021	0.0	0.0
8	0.101 ± 0.003	0.029 ± 0.012	0.371 ± 0.031	0.0	0.0
10	0.065 ± 0.001	0.004 ± 0.000	0.502 ± 0.011	0.0	0.0

E.7 Stability across seeds

Since Semi-NMF, like NMF, yields non-unique solutions due to permutation and scaling ambiguities, we verify that the discovered concepts are stable across random seeds. We run concept extraction five times with different seeds on the same data. To account for permutation invariance, we measure stability by computing the mean cosine similarity between the columns of the concept dictionaries W across runs, after determining the optimal column alignment via the Hungarian algorithm (W -Stab).

Table 15: W -Stab (Mean matched cosine similarity of W across seeds) per class and model variant (DINOv2 ViT-B/14).

Dataset	Class	Reconstruction	+Gram	+KL	Gated SAE	SAE
ImageNet	cassette	0.8904	0.9310	0.9309	0.7667	0.9561
	chainsaw	0.9013	0.9043	0.9052	0.7452	0.9260
	church	0.8969	0.9535	0.9531	0.7224	0.9038
	English springer	0.9330	0.9874	0.9919	0.7477	0.8906
	French horn	0.9426	0.9437	0.9550	0.7609	0.9403
	garbage truck	0.9367	0.9359	0.9365	0.7913	0.9459
	gas pump	0.9829	0.9926	0.9944	0.7791	0.9038
	golf ball	0.9895	0.9905	0.9923	0.8040	0.9046
	parachute	0.9436	0.9420	0.9449	0.8154	0.9223
	tench	0.9589	0.9953	0.9959	0.7041	0.9227
Mean		0.9376	0.9576	0.9600	0.7637	0.9216
COCO	hot dog	0.8317	0.8242	0.8246	0.6459	0.8913
	teddy bear	0.8282	0.8515	0.8502	0.6612	0.8919
	train	0.8706	0.9053	0.9055	0.7329	0.9518
	umbrella	0.6871	0.6900	0.6877	0.5859	0.9088
	zebra	0.7805	0.8249	0.8232	0.6950	0.9121
Mean		0.7996	0.8192	0.8183	0.6642	0.9112

Discussion. All three Semi-NMF variants achieve high W -Stab on ImageNet (mean ≥ 0.94), confirming that the method converges to consistent concept dictionaries despite the non-convexity of the factorization. Adding a faithfulness loss further improves stability: both +Gram and +KL outperform the reconstruction-only baseline on most classes, with the largest gains on *cassette* ($0.890 \rightarrow 0.931$), *church* ($0.897 \rightarrow 0.953$), and *tench* ($0.959 \rightarrow 0.995$). This is consistent with the geometric constraints reducing the effective solution space and biasing Semi-NMF toward more determined factorizations. +KL and +Gram reach nearly identical mean stability on both datasets (0.960 vs. 0.958 on ImageNet, 0.818 vs. 0.819 on COCO), suggesting that faithfulness metrics remain the primary basis for differentiating the two.

The inclusion of Sparse Autoencoders (SAEs) reveals a divergence in architectural stability. **Vanilla SAEs** [33, 57] exhibit high stability, particularly on the low-data COCO regime (mean 0.911), where they significantly outperform all Semi-NMF variants. This suggests that the SAE’s architectural

priors and fixed initialization strategies may provide a more constrained optimization landscape than Semi-NMF. Conversely, **Gated SAEs** [60] show the lowest stability across all tested models (ImageNet mean 0.764, COCO mean 0.664). The additional gating parameters likely introduce a more complex, non-convex objective that is more sensitive to seed initialization, leading to more varied dictionary solutions across runs.

Stability remains noticeably lower for most methods on COCO (mean ≤ 0.82 for Semi-NMF), driven by *umbrella* (22 images) and *zebra* (43 images). With very few images, the patch-token activation matrix is lower-rank and noisier, leaving more freedom for the factorization to settle in different local minima across seeds. While Vanilla SAEs seem more robust to this data-regime limitation, for Semi-NMF, this remains a challenge, we thoroughly isolate and quantify this limitation in Appendix E.8.

E.8 Effect of Dataset Size on Concept Stability

Because Semi-NMF relies on statistical co-occurrences of features across the dataset, its stability depends heavily on the sample size. To quantify this, we extract concepts ($r = 25, s = 2$) from the ImageNet church class using PLACE (+Gram) on the ViT-B/16 backbone at layer $\ell = 9$, progressively restricting the dataset size to 1%, 5%, 10%, 25%, 50%, 75%, and 100% of the available images. We report the matched cosine similarity of the dictionary $\mathbf{W}$ computed across 5 random seeds (W-Stab).

Table 16: Effect of dataset size on dictionary stability (W-Stab). Evaluated on the ImageNet church class across 5 random seeds using PLACE (+Gram). Images refer to the correctly classified subset used for extraction.

Data Fraction	Images	W-Stab $\uparrow$
1%	9	0.587
5%	47	0.799
10%	91	0.767
25%	217	0.877
50%	439	0.842
75%	659	0.873
100%	878	0.906

Discussion. As shown in Table 16, stability is highly sensitive to the sample regime. In the extreme low-data regime (1%, corresponding to just 9 images), W-Stab collapses to 0.587. In such settings, the activation matrix $\mathbf{A}$ is severely under-determined, allowing the factorization to settle into different local minima across seeds. As the number of images increases, stability climbs rapidly, surpassing 0.87 with approximately 200 images, and steadily approaching 0.906 when the full 878-image subset is utilized.

This ablation provides crucial empirical context for the COCO results observed in Appendix E.7: classes like *umbrella* (22 images) and *zebra* (43 images) naturally exhibit lower stability (ranging ~ 0.78 to 0.85). These datasets fall exactly into the extreme low-data regime demonstrated by our 1% and 5% splits, confirming that the performance drop is a fundamental structural limitation of NMF-based factorization on small sample sets, rather than a shortcoming of the geometric loss formulation.

E.9 Runtime and memory

We report wall-clock time and peak GPU memory for concept extraction on ImageNet using DINOv2 ViT-B/14 ($r=25, s=2$) measured on a single NVIDIA RTX 5060 Ti (16 GB VRAM, CUDA 12.8). All timings exclude data loading and preprocessing.

Discussion. Activation extraction and NNDSVD initialization are shared across all variants and complete in under 5 s per class. Both the reconstruction-only and the Gram variants converge in comparable time (~ 5 s), indicating that the Gram-matrix alignment loss adds negligible overhead to the optimization loop. The full concept extraction pipeline thus completes in under 10 s per class for these two variants. The KL variant requires substantially longer optimization due to the

Table 17: Runtime and memory for concept extraction (DINOv2 ViT-B/14, ImageNet). Median across 10 classes.

Phase	Time (s)	Peak VRAM (GB)
Activation extraction + NNDSVD init	3.9	10.9
Optimization (PLACE, reconstruction)	5.3	10.9
Optimization (PLACE, +Gram)	5.6	10.9
Optimization (PLACE, +KL)	118.0	10.6
Training (SAE)	62.5	8.0
Training (Gated SAE)	71.4	8.0

per-step cost of computing the divergence over the full patch-token distribution and its gradients. Training for Sparse Autoencoders (SAEs) and Gated SAEs also requires notably more time than the baseline reconstruction variants (62.5 s and 71.4 s, respectively), though they remain faster than the KL approach. The gating mechanism in Gated SAEs introduces only a modest runtime overhead (~ 9 s) over standard SAEs. Notably, both SAE variants benefit from a lower peak VRAM footprint (8.0 GB). For the other variants, peak VRAM is reported for batch extraction over the entire class subset ($\sim 1,000$ images), where the dominant cost is storing the activation matrix $\mathbf{A}$ in memory. At inference time, when projecting a single image onto a precomputed concept bank, the memory footprint reduces to a single forward pass through the backbone plus a matrix multiplication with W , requiring less than 1 GB of additional VRAM beyond the model weights.

E.10 Computational Complexity and Scalability to Larger Architectures

We analyze the computational cost and memory footprint of PLACE as a function of the model and extraction parameters, and empirically validate its scalability on DINOv2 ViT-L/14.

Notation. Let P denote the patch size, H the input resolution, $G_0 = H/P$ the native grid dimension, s the token grouping factor, $G_s = \lfloor G_0/s \rfloor$ the grouped grid dimension, $N_s = G_s^2$ the number of grouped tokens per image, n the number of images, d the embedding dimension, and r the dictionary rank.

Activation extraction. A single forward pass through the prefix $g_{\leq \ell}$ costs $\mathcal{O}(n \cdot G_0^2 \cdot d^2)$, dominated by the self-attention computation. This cost is shared across all methods (PLACE, SAEs, CRAFT, FACE) and is *not* part of the concept extraction itself. Token grouping via non-overlapping average pooling is $\mathcal{O}(n \cdot G_0^2 \cdot d)$ and negligible in practice.

NNDSVD initialization. The initialization requires a truncated SVD of a subsampled activation matrix $\tilde{\mathbf{A}} \in \mathbb{R}^{m \times d}$, where $m = \rho \cdot n N_s$ and $\rho \in (0, 1]$ is the subsampling ratio. Using randomized SVD, this costs $\mathcal{O}(m \cdot d \cdot r)$. Since $r \ll d \ll m$ in all configurations, this step is dominated by the optimization phase only when ρ is large.

Per-step optimization cost. Each AdamW step involves:

1. **Reconstruction:** Computing $\hat{\mathbf{A}} = \mathbf{U}\mathbf{W}^\top$ for a mini-batch of b images costs $\mathcal{O}(b \cdot N_s \cdot d \cdot r)$, and $\mathcal{L}_{\text{rec}} = \frac{1}{2} \|\mathbf{A}_b - \hat{\mathbf{A}}_b\|_F^2$ is $\mathcal{O}(b \cdot N_s \cdot d)$.
2. **Gram loss:** Computing the Gram matrices $\mathbf{A}_b \mathbf{A}_b^\top$ and $\hat{\mathbf{A}}_b \hat{\mathbf{A}}_b^\top \in \mathbb{R}^{b N_s \times b N_s}$ costs $\mathcal{O}((b N_s)^2 \cdot d)$, and the Frobenius norm of their difference is $\mathcal{O}((b N_s)^2)$. This is the dominant cost per step.
3. **KL loss:** Requires a full forward pass through the suffix $g_{> \ell}$ and classifier h , costing $\mathcal{O}(b \cdot G_0^2 \cdot d^2 \cdot L_{\text{suf}})$ where $L_{\text{suf}} = L - \ell$. This is one reason for the $\sim 20\times$ slowdown of the +KL variant.

The total per-step cost of PLACE (+Gram) is therefore:

$$\mathcal{O}(b \cdot N_s \cdot d \cdot r + (b \cdot N_s)^2 \cdot d), \quad (37)$$

which is *independent* of the suffix depth L_{suf} and requires no forward pass through $g_{> \ell}$ or h .

Memory footprint. Peak GPU memory is dominated by storing the full activation matrix $\mathbf{A} \in \mathbb{R}^{n N_s \times d}$ and the two Gram matrices per mini-batch (each $b N_s \times b N_s$). In float32:

$$\text{VRAM} \approx 4 \cdot (n \cdot N_s \cdot d + 2 \cdot (b \cdot N_s)^2) \text{ bytes}, \quad (38)$$

plus model weights for the prefix (loaded once for activation extraction, then discarded during optimization).

Scaling properties across ViT scales. Table 18 reports the projected computational parameters for architectures from ViT-B to ViT-G.

Table 18: Projected computational parameters for PLACE (+Gram) across ViT scales. N_s : grouped tokens per image; $|\mathbf{A}|$: activation matrix size. All estimates assume $n=800$, $b=4$, $r=25$, $s=2$.

Architecture	P	H	d	N_s	$ \mathbf{A} $ (M floats)	Gram size
ViT-B/16	16	224	768	49	30.0	196×196
DINOv2 ViT-B/14	14	518	768	324	199.1	1296×1296
DINOv2 ViT-L/14	14	518	1024	324	265.4	1296×1296
DINOv2 ViT-G/14	14	518	1536	324	398.1	1296×1296

A crucial observation emerges: since DINOv2 ViT-B/14, ViT-L/14, and ViT-G/14 share the same patch size ($P=14$) and input resolution ($H=518$), they produce *identical* token grids (37×37 , grouped to 18×18). Consequently, the Gram matrix computation – the dominant per-step cost – is identical across all three scales. Only the linear reconstruction term $\mathcal{O}(bN_s \cdot d \cdot r)$ and the activation storage $\mathcal{O}(nN_s \cdot d)$ grow with the embedding dimension d . Furthermore, unlike the +KL variant, whose per-step cost scales as $\mathcal{O}(b \cdot G_0^2 \cdot d^2 \cdot L_{\text{sur}})$, the +Gram variant is entirely independent of the number of remaining transformer blocks. This advantage compounds with model depth: for ViT-G ($L=40$), extracting at $\ell=30$ would require backpropagating through 10 blocks per step under +KL, while +Gram incurs zero additional cost.

Empirical validation on DINOv2 ViT-L/14. To confirm these projections, we run the full PLACE pipeline on DINOv2 ViT-L/14 ($d=1024$, $L=24$) for two representative ImageNet classes (church and tench), using ($r=25$, $s=2$, $\ell=21$, $\lambda=1.0$).

Table 19: Wall-clock runtime breakdown on DINOv2 ViT-L/14 (median over runs). The +Gram variant completes in 23–38 s total, achieving a 5–25 $\times$ speedup over +KL.

Class	Method	Extract (s)	Init (s)	Optimize (s)	Total (s)
church	Reconstruction	9.6	2.2	13.6	25.4
	+Gram	9.4	2.1	11.1	22.7
	+KL	9.5	2.1	566.6	578.5
tench	Reconstruction	11.7	2.5	15.4	29.9
	+Gram	11.6	2.5	23.9	38.0
	+KL	11.6	2.5	188.4	202.5

The runtime confirms the complexity analysis: the +Gram optimization cost increases only modestly from ViT-B (5.6 s, Table 17) to ViT-L (11–24 s), consistent with the linear scaling in d , while the Gram bottleneck remains unchanged. The +KL variant, by contrast, suffers an increase (118 s $\rightarrow$ 202–578 s) due to the larger suffix with $d=1024$.

Table 20 reports the full suite of faithfulness, sparsity, reconstruction, and stability metrics on ViT-L/14.

Table 20: Faithfulness, reconstruction, sparsity, and dictionary stability on DINOv2 ViT-L/14 (ImageNet, $r=25$, $\ell=21$). Higher is better: C-Ins, C-Del, C-Gini, W-Stab; lower: MSE, D_{KL} . Mean $\pm$ std over 5 runs. Bold indicates best per class.

Class	Method	C-Ins $\uparrow$	C-Del $\uparrow$	C-Gini $\uparrow$	MSE $\downarrow$	D_{KL} $\downarrow$	W-Stab $\uparrow$
church	Reconstruction	0.590 $\pm$ 0.023	0.784 $\pm$ 0.022	0.667 $\pm$ 0.004	0.372 $\pm$ 0.000	0.356 $\pm$ 0.005	0.894
	+Gram	0.721 $\pm$ 0.068	0.783 $\pm$ 0.032	0.704 $\pm$ 0.023	0.373 $\pm$ 0.003	0.361 $\pm$ 0.009	0.938
	+KL	0.775 $\pm$ 0.118	0.831 $\pm$ 0.025	0.720 $\pm$ 0.026	0.377 $\pm$ 0.002	0.374 $\pm$ 0.011	0.929
tench	Reconstruction	0.857 $\pm$ 0.028	0.660 $\pm$ 0.030	0.632 $\pm$ 0.002	0.328 $\pm$ 0.001	0.123 $\pm$ 0.002	0.941
	+Gram	0.894 $\pm$ 0.045	0.763 $\pm$ 0.110	0.654 $\pm$ 0.033	0.327 $\pm$ 0.004	0.118 $\pm$ 0.002	0.957
	+KL	0.937 $\pm$ 0.006	0.889 $\pm$ 0.005	0.666 $\pm$ 0.013	0.335 $\pm$ 0.001	0.123 $\pm$ 0.003	0.941

The results confirm that all metrics transfer consistently from ViT-B to ViT-L. The +Gram variant achieves the highest dictionary stability on both classes (W-Stab 0.938 and 0.957), confirming that the geometric constraint continues to regularize the factorization at larger scale. Crucially, the +Gram variant achieves lower D_{KL} than the supervised +KL reference on `tench` (0.118 vs. 0.123), replicating the pattern observed on ViT-B/14 (Table 1). The MSE–faithfulness trade-off is consistent: reconstruction-only achieves the lowest MSE while +KL achieves the best causal metrics, with +Gram positioned between the two – matching its unsupervised character.

Table 21 reports the semantic alignment between unsupervised and supervised concept banks, measured via RSVC.

Table 21: Agreement of unsupervised concept banks with the supervised (+KL) reference on DINOv2 ViT-L/14 (ImageNet, $r=25$, $\ell=21$). Mean $\pm$ std over 5 runs.

Class	Comparison	MCS $\uparrow$	CMCS $\uparrow$
church	Reconstruction vs. +KL	0.921 $\pm$ 0.018	0.740 $\pm$ 0.006
	+Gram vs. +KL	0.947 $\pm$ 0.015	0.736 $\pm$ 0.010
tench	Reconstruction vs. +KL	0.918 $\pm$ 0.012	0.695 $\pm$ 0.010
	+Gram vs. +KL	0.935 $\pm$ 0.018	0.697 $\pm$ 0.008

The +Gram variant achieves substantially higher individual concept alignment on ViT-L than on ViT-B (MCS 0.947 vs. 0.866 for `church`; 0.935 vs. 0.829 for `tench`, cf. Table 28). This improvement of +0.08 to +0.10 MCS suggests that the Gram-matrix proxy becomes a *tighter* surrogate for classifier supervision at larger model scales, likely because the richer representations of ViT-L produce pairwise token geometries that are more predictive of downstream task structure.

Summary. PLACE’s computational cost is governed primarily by the *spatial* resolution of the token grid (controlled by P , H , and s) rather than the *model* scale (d , L). The empirical validation on ViT-L/14 confirms that the method scales gracefully: extraction completes in under 40s on consumer hardware, the Gram-matrix proxy for faithfulness becomes *more* effective at larger scale (MCS improvement of +0.08 to +0.10 over ViT-B), and the +Gram variant achieves the highest dictionary stability while closely matching the supervised reference across all metrics.

E.11 Per-class faithfulness trends

We report detailed, per-class faithfulness and reconstruction metrics for the three PLACE variants across both ImageNet and COCO datasets. Table 22 presents the results for the self-supervised DINOv2 ViT-B/14 backbone, while Table 23 reports the metrics for the supervised ViT-B/16 architecture.

Table 22: Per-class faithfulness and reconstruction quality for all three PLACE variants on ImageNet and COCO (DINOv2 ViT-B/14, $\ell = 9$, $r=25$). Mean $\pm$ std across 5 seeds.

Class	PLACE (rec.)					PLACE (+Gram)					PLACE (+KL)				
	C-Ins $\uparrow$	C-Del $\uparrow$	C-Gini $\uparrow$	MSE $\downarrow$	$D_{KL} \downarrow$	C-Ins $\uparrow$	C-Del $\uparrow$	C-Gini $\uparrow$	MSE $\downarrow$	$D_{KL} \downarrow$	C-Ins $\uparrow$	C-Del $\uparrow$	C-Gini $\uparrow$	MSE $\downarrow$	$D_{KL} \downarrow$
<i>ImageNet</i>															
cassette	0.422 $\pm$ 0.021	0.910 $\pm$ 0.015	0.917 $\pm$ 0.001	0.082 $\pm$ 0.000	0.501 $\pm$ 0.012	0.701 $\pm$ 0.025	0.940 $\pm$ 0.004	0.910 $\pm$ 0.003	0.083 $\pm$ 0.000	0.651 $\pm$ 0.010	0.682 $\pm$ 0.029	0.936 $\pm$ 0.005	0.913 $\pm$ 0.003	0.083 $\pm$ 0.000	0.658 $\pm$ 0.006
chaise	0.380 $\pm$ 0.038	0.784 $\pm$ 0.026	0.939 $\pm$ 0.001	0.096 $\pm$ 0.000	2.205 $\pm$ 0.407	0.401 $\pm$ 0.055	0.790 $\pm$ 0.039	0.926 $\pm$ 0.008	0.093 $\pm$ 0.002	1.703 $\pm$ 0.477	0.404 $\pm$ 0.061	0.805 $\pm$ 0.036	0.930 $\pm$ 0.005	0.095 $\pm$ 0.002	1.776 $\pm$ 0.460
church	0.113 $\pm$ 0.020	0.952 $\pm$ 0.004	0.930 $\pm$ 0.001	0.093 $\pm$ 0.000	1.277 $\pm$ 0.031	0.153 $\pm$ 0.041	0.962 $\pm$ 0.004	0.923 $\pm$ 0.004	0.093 $\pm$ 0.001	1.275 $\pm$ 0.094	0.130 $\pm$ 0.036	0.962 $\pm$ 0.000	0.928 $\pm$ 0.004	0.095 $\pm$ 0.001	1.326 $\pm$ 0.082
english springer	0.690 $\pm$ 0.018	0.957 $\pm$ 0.004	0.931 $\pm$ 0.001	0.101 $\pm$ 0.001	0.813 $\pm$ 0.012	0.746 $\pm$ 0.022	0.958 $\pm$ 0.002	0.918 $\pm$ 0.011	0.100 $\pm$ 0.003	0.765 $\pm$ 0.073	0.745 $\pm$ 0.023	0.958 $\pm$ 0.002	0.925 $\pm$ 0.006	0.102 $\pm$ 0.002	0.800 $\pm$ 0.052
french horn	0.076 $\pm$ 0.003	0.905 $\pm$ 0.011	0.936 $\pm$ 0.001	0.095 $\pm$ 0.000	1.196 $\pm$ 0.028	0.067 $\pm$ 0.025	0.927 $\pm$ 0.006	0.926 $\pm$ 0.011	0.094 $\pm$ 0.002	1.388 $\pm$ 0.425	0.052 $\pm$ 0.010	0.928 $\pm$ 0.004	0.928 $\pm$ 0.004	0.096 $\pm$ 0.001	2.264 $\pm$ 0.157
garbage truck	0.438 $\pm$ 0.010	0.960 $\pm$ 0.001	0.938 $\pm$ 0.000	0.091 $\pm$ 0.000	1.065 $\pm$ 0.038	0.704 $\pm$ 0.083	0.956 $\pm$ 0.001	0.923 $\pm$ 0.011	0.088 $\pm$ 0.003	1.067 $\pm$ 0.381	0.658 $\pm$ 0.021	0.955 $\pm$ 0.001	0.928 $\pm$ 0.004	0.090 $\pm$ 0.001	1.185 $\pm$ 0.183
gas pump	0.275 $\pm$ 0.017	0.712 $\pm$ 0.019	0.948 $\pm$ 0.000	0.094 $\pm$ 0.000	3.112 $\pm$ 0.122	0.286 $\pm$ 0.110	0.780 $\pm$ 0.022	0.932 $\pm$ 0.009	0.090 $\pm$ 0.003	2.540 $\pm$ 0.828	0.293 $\pm$ 0.036	0.760 $\pm$ 0.031	0.934 $\pm$ 0.003	0.091 $\pm$ 0.001	2.491 $\pm$ 0.294
gold bull	0.925 $\pm$ 0.001	0.960 $\pm$ 0.000	0.938 $\pm$ 0.000	0.092 $\pm$ 0.000	0.216 $\pm$ 0.009	0.915 $\pm$ 0.005	0.957 $\pm$ 0.002	0.927 $\pm$ 0.006	0.089 $\pm$ 0.002	0.221 $\pm$ 0.026	0.911 $\pm$ 0.003	0.957 $\pm$ 0.002	0.926 $\pm$ 0.006	0.089 $\pm$ 0.002	0.219 $\pm$ 0.030
parachute	0.831 $\pm$ 0.003	0.950 $\pm$ 0.001	0.934 $\pm$ 0.000	0.088 $\pm$ 0.000	0.539 $\pm$ 0.017	0.838 $\pm$ 0.003	0.960 $\pm$ 0.001	0.925 $\pm$ 0.005	0.086 $\pm$ 0.002	0.563 $\pm$ 0.037	0.839 $\pm$ 0.012	0.960 $\pm$ 0.000	0.921 $\pm$ 0.010	0.085 $\pm$ 0.003	0.551 $\pm$ 0.116
tench	0.521 $\pm$ 0.016	0.856 $\pm$ 0.006	0.924 $\pm$ 0.001	0.090 $\pm$ 0.001	0.435 $\pm$ 0.017	0.405 $\pm$ 0.038	0.910 $\pm$ 0.005	0.898 $\pm$ 0.014	0.087 $\pm$ 0.002	0.785 $\pm$ 0.332	0.376 $\pm$ 0.026	0.922 $\pm$ 0.004	0.911 $\pm$ 0.007	0.090 $\pm$ 0.002	1.055 $\pm$ 0.199
<i>COCO</i>															
hot dog	0.132 $\pm$ 0.024	0.894 $\pm$ 0.025	0.869 $\pm$ 0.001	0.077 $\pm$ 0.000	0.212 $\pm$ 0.007	0.096 $\pm$ 0.009	0.913 $\pm$ 0.018	0.899 $\pm$ 0.002	0.081 $\pm$ 0.001	0.300 $\pm$ 0.011	0.089 $\pm$ 0.013	0.915 $\pm$ 0.017	0.901 $\pm$ 0.002	0.083 $\pm$ 0.002	0.317 $\pm$ 0.023
ukky bear	0.474 $\pm$ 0.029	0.601 $\pm$ 0.044	0.908 $\pm$ 0.001	0.084 $\pm$ 0.000	1.215 $\pm$ 0.027	0.274 $\pm$ 0.028	0.776 $\pm$ 0.034	0.906 $\pm$ 0.003	0.087 $\pm$ 0.001	1.045 $\pm$ 0.063	0.265 $\pm$ 0.040	0.772 $\pm$ 0.039	0.905 $\pm$ 0.003	0.087 $\pm$ 0.002	1.032 $\pm$ 0.094
train	0.020 $\pm$ 0.007	0.852 $\pm$ 0.015	0.897 $\pm$ 0.001	0.073 $\pm$ 0.000	0.332 $\pm$ 0.014	0.041 $\pm$ 0.004	0.946 $\pm$ 0.009	0.899 $\pm$ 0.002	0.077 $\pm$ 0.001	0.354 $\pm$ 0.008	0.045 $\pm$ 0.003	0.948 $\pm$ 0.012	0.899 $\pm$ 0.002	0.077 $\pm$ 0.001	0.348 $\pm$ 0.006
umbrella	0.149 $\pm$ 0.027	0.898 $\pm$ 0.015	0.853 $\pm$ 0.001	0.077 $\pm$ 0.000	0.648 $\pm$ 0.026	0.156 $\pm$ 0.042	0.927 $\pm$ 0.029	0.853 $\pm$ 0.001	0.077 $\pm$ 0.009	0.666 $\pm$ 0.016	0.144 $\pm$ 0.032	0.910 $\pm$ 0.020	0.852 $\pm$ 0.001	0.077 $\pm$ 0.000	0.646 $\pm$ 0.020
zebra	0.251 $\pm$ 0.108	0.327 $\pm$ 0.026	0.873 $\pm$ 0.002	0.071 $\pm$ 0.000	0.408 $\pm$ 0.019	0.316 $\pm$ 0.082	0.317 $\pm$ 0.043	0.882 $\pm$ 0.002	0.076 $\pm$ 0.000	0.554 $\pm$ 0.017	0.303 $\pm$ 0.099	0.296 $\pm$ 0.051	0.879 $\pm$ 0.006	0.075 $\pm$ 0.002	0.539 $\pm$ 0.050

Discussion. The results demonstrate distinct behavioral trends across the three transformer backbones, particularly when comparing the PLACE framework against Sparse Autoencoder (SAE) baselines.

For the self-supervised **DINOv2** model (Table 25), the introduction of the unsupervised geometric loss (PLACE +Gram) improves both concept insertion (C-Ins) and deletion (C-Del) metrics compared to the reconstruction-only baseline across the majority of ImageNet and COCO classes. Furthermore, the KL divergence (D_{KL}) achieved by the +Gram variant closely matches the supervised reference upper bound (+KL variant). This confirms that enforcing Gram-matrix alignment successfully guides

Table 23: Per-class faithfulness and reconstruction quality for all three PLACE variants on ImageNet and COCO (ViT-B/16, $\ell = 9, r=25$). Mean $\pm$ std across 5 seeds.

Class	PLACE (rec.)					PLACE (+Gram)					PLACE (+KL)				
	C-Ins $\uparrow$	C-Del $\uparrow$	C-Gini $\uparrow$	MSE $\downarrow$	$D_{KL} \downarrow$	C-Ins $\uparrow$	C-Del $\uparrow$	C-Gini $\uparrow$	MSE $\downarrow$	$D_{KL} \downarrow$	C-Ins $\uparrow$	C-Del $\uparrow$	C-Gini $\uparrow$	MSE $\downarrow$	$D_{KL} \downarrow$
<i>ImageNet</i>															
cassette	0.614 $\pm$ 0.024	0.133 $\pm$ 0.015	0.514 $\pm$ 0.008	0.265 $\pm$ 0.000	0.000 $\pm$ 0.000	0.533 $\pm$ 0.034	0.268 $\pm$ 0.024	0.579 $\pm$ 0.014	0.271 $\pm$ 0.003	0.000 $\pm$ 0.000	0.730 $\pm$ 0.021	0.234 $\pm$ 0.032	0.448 $\pm$ 0.028	0.281 $\pm$ 0.006	0.012 $\pm$ 0.007
chainsaw	0.685 $\pm$ 0.014	0.362 $\pm$ 0.079	0.618 $\pm$ 0.006	0.314 $\pm$ 0.000	0.000 $\pm$ 0.000	0.654 $\pm$ 0.026	0.432 $\pm$ 0.047	0.456 $\pm$ 0.009	0.317 $\pm$ 0.004	0.000 $\pm$ 0.000	0.765 $\pm$ 0.014	0.116 $\pm$ 0.018	0.459 $\pm$ 0.012	0.327 $\pm$ 0.001	0.027 $\pm$ 0.004
church	0.606 $\pm$ 0.032	0.434 $\pm$ 0.084	0.474 $\pm$ 0.003	0.289 $\pm$ 0.000	0.000 $\pm$ 0.000	0.707 $\pm$ 0.028	0.300 $\pm$ 0.011	0.571 $\pm$ 0.017	0.293 $\pm$ 0.003	0.000 $\pm$ 0.000	0.911 $\pm$ 0.003	0.147 $\pm$ 0.023	0.439 $\pm$ 0.014	0.392 $\pm$ 0.002	0.015 $\pm$ 0.003
english springer	0.747 $\pm$ 0.012	0.233 $\pm$ 0.045	0.536 $\pm$ 0.009	0.304 $\pm$ 0.001	0.000 $\pm$ 0.000	0.762 $\pm$ 0.007	0.104 $\pm$ 0.021	0.587 $\pm$ 0.006	0.306 $\pm$ 0.002	0.000 $\pm$ 0.000	0.736 $\pm$ 0.019	0.211 $\pm$ 0.040	0.328 $\pm$ 0.011	0.314 $\pm$ 0.001	0.016 $\pm$ 0.004
french horn	0.746 $\pm$ 0.030	0.451 $\pm$ 0.099	0.510 $\pm$ 0.013	0.306 $\pm$ 0.001	0.000 $\pm$ 0.000	0.694 $\pm$ 0.027	0.406 $\pm$ 0.073	0.605 $\pm$ 0.013	0.309 $\pm$ 0.003	0.000 $\pm$ 0.000	0.822 $\pm$ 0.056	0.363 $\pm$ 0.088	0.452 $\pm$ 0.035	0.316 $\pm$ 0.002	0.019 $\pm$ 0.009
garbage truck	0.463 $\pm$ 0.015	0.139 $\pm$ 0.047	0.294 $\pm$ 0.006	0.278 $\pm$ 0.000	0.000 $\pm$ 0.000	0.506 $\pm$ 0.007	0.090 $\pm$ 0.003	0.559 $\pm$ 0.004	0.280 $\pm$ 0.001	0.000 $\pm$ 0.000	0.503 $\pm$ 0.015	0.314 $\pm$ 0.069	0.303 $\pm$ 0.021	0.280 $\pm$ 0.003	0.020 $\pm$ 0.009
gas pump	0.601 $\pm$ 0.006	0.536 $\pm$ 0.061	0.571 $\pm$ 0.008	0.298 $\pm$ 0.000	0.000 $\pm$ 0.000	0.708 $\pm$ 0.037	0.452 $\pm$ 0.037	0.615 $\pm$ 0.017	0.301 $\pm$ 0.003	0.000 $\pm$ 0.000	0.928 $\pm$ 0.010	0.195 $\pm$ 0.036	0.479 $\pm$ 0.015	0.310 $\pm$ 0.001	0.023 $\pm$ 0.008
golf ball	0.900 $\pm$ 0.008	0.441 $\pm$ 0.035	0.620 $\pm$ 0.011	0.316 $\pm$ 0.000	0.000 $\pm$ 0.000	0.693 $\pm$ 0.012	0.812 $\pm$ 0.017	0.689 $\pm$ 0.005	0.320 $\pm$ 0.003	0.000 $\pm$ 0.000	0.764 $\pm$ 0.039	0.702 $\pm$ 0.067	0.497 $\pm$ 0.005	0.325 $\pm$ 0.001	0.012 $\pm$ 0.003
parachute	0.915 $\pm$ 0.009	0.476 $\pm$ 0.028	0.633 $\pm$ 0.008	0.258 $\pm$ 0.000	0.000 $\pm$ 0.000	0.900 $\pm$ 0.009	0.314 $\pm$ 0.016	0.698 $\pm$ 0.008	0.263 $\pm$ 0.003	0.000 $\pm$ 0.000	0.964 $\pm$ 0.006	0.165 $\pm$ 0.009	0.588 $\pm$ 0.012	0.267 $\pm$ 0.001	0.012 $\pm$ 0.003
tench	0.787 $\pm$ 0.033	0.716 $\pm$ 0.100	0.431 $\pm$ 0.014	0.296 $\pm$ 0.000	0.000 $\pm$ 0.000	0.856 $\pm$ 0.016	0.828 $\pm$ 0.020	0.587 $\pm$ 0.005	0.300 $\pm$ 0.003	0.000 $\pm$ 0.000	0.830 $\pm$ 0.013	0.682 $\pm$ 0.060	0.379 $\pm$ 0.036	0.304 $\pm$ 0.002	0.007 $\pm$ 0.003
<i>COCO</i>															
hot dog	0.983 $\pm$ 0.002	0.990 $\pm$ 0.005	0.577 $\pm$ 0.005	0.351 $\pm$ 0.000	0.000 $\pm$ 0.000	0.948 $\pm$ 0.005	0.990 $\pm$ 0.002	0.628 $\pm$ 0.013	0.351 $\pm$ 0.003	0.000 $\pm$ 0.000	0.925 $\pm$ 0.001	0.934 $\pm$ 0.015	0.646 $\pm$ 0.005	0.356 $\pm$ 0.001	0.003 $\pm$ 0.002
teddy bear	0.970 $\pm$ 0.001	0.965 $\pm$ 0.018	0.657 $\pm$ 0.006	0.369 $\pm$ 0.000	0.000 $\pm$ 0.000	0.959 $\pm$ 0.006	0.950 $\pm$ 0.017	0.664 $\pm$ 0.038	0.362 $\pm$ 0.002	0.000 $\pm$ 0.000	0.961 $\pm$ 0.003	0.893 $\pm$ 0.012	0.692 $\pm$ 0.027	0.369 $\pm$ 0.002	0.002 $\pm$ 0.001
train	0.996 $\pm$ 0.000	0.958 $\pm$ 0.003	0.576 $\pm$ 0.002	0.368 $\pm$ 0.000	0.000 $\pm$ 0.000	0.965 $\pm$ 0.010	0.986 $\pm$ 0.002	0.606 $\pm$ 0.047	0.359 $\pm$ 0.002	0.000 $\pm$ 0.000	0.956 $\pm$ 0.004	0.946 $\pm$ 0.010	0.684 $\pm$ 0.036	0.373 $\pm$ 0.003	0.003 $\pm$ 0.000
umbrella	0.991 $\pm$ 0.003	0.947 $\pm$ 0.005	0.704 $\pm$ 0.005	0.415 $\pm$ 0.001	0.000 $\pm$ 0.000	0.927 $\pm$ 0.017	0.985 $\pm$ 0.002	0.664 $\pm$ 0.038	0.403 $\pm$ 0.003	0.000 $\pm$ 0.000	0.933 $\pm$ 0.019	0.945 $\pm$ 0.022	0.646 $\pm$ 0.017	0.416 $\pm$ 0.003	0.006 $\pm$ 0.001
zebra	0.999 $\pm$ 0.000	0.979 $\pm$ 0.003	0.673 $\pm$ 0.021	0.282 $\pm$ 0.001	0.000 $\pm$ 0.000	0.985 $\pm$ 0.001	0.986 $\pm$ 0.001	0.711 $\pm$ 0.008	0.279 $\pm$ 0.000	0.000 $\pm$ 0.000	0.969 $\pm$ 0.004	0.981 $\pm$ 0.003	0.719 $\pm$ 0.005	0.283 $\pm$ 0.000	0.001 $\pm$ 0.000

Table 24: Per-class faithfulness and reconstruction quality for all three PLACE variants on ImageNet and COCO (ViT-MAE B/16, $\ell = 9, r=25$). Mean $\pm$ std across 5 seeds.

Class	PLACE (rec.)					PLACE (+Gram)					PLACE (+KL)				
	C-Ins $\uparrow$	C-Del $\uparrow$	C-Gini $\uparrow$	MSE $\downarrow$	$D_{KL} \downarrow$	C-Ins $\uparrow$	C-Del $\uparrow$	C-Gini $\uparrow$	MSE $\downarrow$	$D_{KL} \downarrow$	C-Ins $\uparrow$	C-Del $\uparrow$	C-Gini $\uparrow$	MSE $\downarrow$	$D_{KL} \downarrow$
<i>ImageNet</i>															
cassette	0.930 $\pm$ 0.018	0.465 $\pm$ 0.142	0.396 $\pm$ 0.005	1.021 $\pm$ 0.000	0.374 $\pm$ 0.002	0.953 $\pm$ 0.007	0.151 $\pm$ 0.039	0.645 $\pm$ 0.005	1.064 $\pm$ 0.003	0.434 $\pm$ 0.007	0.955 $\pm$ 0.006	0.206 $\pm$ 0.045	0.583 $\pm$ 0.022	1.035 $\pm$ 0.008	0.383 $\pm$ 0.020
chainsaw	0.942 $\pm$ 0.007	0.408 $\pm$ 0.071	0.434 $\pm$ 0.023	1.148 $\pm$ 0.001	0.401 $\pm$ 0.004	1.000 $\pm$ 0.000	0.115 $\pm$ 0.006	0.631 $\pm$ 0.009	1.168 $\pm$ 0.005	0.477 $\pm$ 0.007	1.000 $\pm$ 0.000	0.115 $\pm$ 0.009	0.642 $\pm$ 0.028	1.185 $\pm$ 0.025	0.491 $\pm$ 0.020
church	0.806 $\pm$ 0.020	0.599 $\pm$ 0.056	0.398 $\pm$ 0.007	0.936 $\pm$ 0.000	0.471 $\pm$ 0.014	0.996 $\pm$ 0.004	0.119 $\pm$ 0.020	0.652 $\pm$ 0.015	0.975 $\pm$ 0.012	0.529 $\pm$ 0.015	0.996 $\pm$ 0.002	0.168 $\pm$ 0.106	0.635 $\pm$ 0.032	0.965 $\pm$ 0.018	0.501 $\pm$ 0.028
english springer	0.813 $\pm$ 0.009	0.484 $\pm$ 0.037	0.420 $\pm$ 0.033	1.003 $\pm$ 0.001	0.387 $\pm$ 0.006	0.997 $\pm$ 0.006	0.063 $\pm$ 0.003	0.674 $\pm$ 0.010	1.042 $\pm$ 0.007	0.412 $\pm$ 0.003	0.996 $\pm$ 0.006	0.075 $\pm$ 0.007	0.635 $\pm$ 0.009	1.023 $\pm$ 0.003	0.396 $\pm$ 0.003
french horn	0.962 $\pm$ 0.004	0.639 $\pm$ 0.032	0.410 $\pm$ 0.007	1.088 $\pm$ 0.000	0.247 $\pm$ 0.003	1.000 $\pm$ 0.000	0.165 $\pm$ 0.032	0.686 $\pm$ 0.006	1.117 $\pm$ 0.004	0.286 $\pm$ 0.007	1.000 $\pm$ 0.000	0.220 $\pm$ 0.092	0.675 $\pm$ 0.023	1.116 $\pm$ 0.015	0.278 $\pm$ 0.021
garbage truck	0.862 $\pm$ 0.036	0.685 $\pm$ 0.033	0.386 $\pm$ 0.017	0.987 $\pm$ 0.001	0.314 $\pm$ 0.004	0.978 $\pm$ 0.021	0.250 $\pm$ 0.048	0.666 $\pm$ 0.017	1.018 $\pm$ 0.014	0.363 $\pm$ 0.021	0.979 $\pm$ 0.020	0.265 $\pm$ 0.049	0.663 $\pm$ 0.015	1.021 $\pm$ 0.009	0.359 $\pm$ 0.014
gas pump	0.963 $\pm$ 0.009	0.645 $\pm$ 0.014	0.395 $\pm$ 0.052	1.082 $\pm$ 0.001	0.416 $\pm$ 0.003	0.970 $\pm$ 0.002	0.234 $\pm$ 0.017	0.646 $\pm$ 0.007	1.117 $\pm$ 0.007	0.486 $\pm$ 0.013	0.971 $\pm$ 0.003	0.278 $\pm$ 0.055	0.628 $\pm$ 0.017	1.103 $\pm$ 0.017	0.451 $\pm$ 0.009
golf ball	0.959 $\pm$ 0.007	0.706 $\pm$ 0.015	0.533 $\pm$ 0.016	0.952 $\pm$ 0.001	0.066 $\pm$ 0.000	0.997 $\pm$ 0.001	0.109 $\pm$ 0.034	0.737 $\pm$ 0.005	0.983 $\pm$ 0.007	0.068 $\pm$ 0.001	0.995 $\pm$ 0.000	0.115 $\pm$ 0.041	0.737 $\pm$ 0.009	0.986 $\pm$ 0.014	0.069 $\pm$ 0.001
parachute	0.931 $\pm$ 0.010	0.686 $\pm$ 0.025	0.543 $\pm$ 0.017	0.820 $\pm$ 0.000	0.158 $\pm$ 0.003	1.000 $\pm$ 0.000	0.339 $\pm$ 0.066	0.732 $\pm$ 0.007	0.858 $\pm$ 0.012	0.176 $\pm$ 0.004	0.997 $\pm$ 0.003	0.373 $\pm$ 0.090	0.721 $\pm$ 0.025	0.857 $\pm$ 0.032	0.171 $\pm$ 0.013
tench	0.964 $\pm$ 0.023	0.544 $\pm$ 0.044	0.465 $\pm$ 0.034	0.911 $\pm$ 0.000	0.126 $\pm$ 0.003	1.000 $\pm$ 0.000	0.062 $\pm$ 0.018	0.712 $\pm$ 0.011	0.940 $\pm$ 0.012	0.133 $\pm$ 0.010	1.000 $\pm$ 0.000	0.070 $\pm$ 0.023	0.708 $\pm$ 0.017	0.940 $\pm$ 0.017	0.151 $\pm$ 0.016
<i>COCO</i>															
hot dog	0.201 $\pm$ 0.082	0.176 $\pm$ 0.113	0.322 $\pm$ 0.019	0.657 $\pm$ 0.000	0.070 $\pm$ 0.002	0.202 $\pm$ 0.087	0.173 $\pm$ 0.113	0.325 $\pm$ 0.015	0.657 $\pm$ 0.000	0.070 $\pm$ 0.002	0.193 $\pm$ 0.081	0.164 $\pm$ 0.115	0.325 $\pm$ 0.016	0.657 $\pm$ 0.000	0.070 $\pm$ 0.002
teddy bear	0.328 $\pm$ 0.055	0.844 $\pm$ 0.073	0.455 $\pm$ 0.028	0.575 $\pm$ 0.001	0.107 $\pm$ 0.001	0.447 $\pm$ 0.057	0.860 $\pm$ 0.054	0.605 $\pm$ 0.037	0.629 $\pm$ 0.029	0.106 $\pm$ 0.002	0.436 $\pm$ 0.017	0.864 $\pm$ 0.052	0.590 $\pm$ 0.044	0.618 $\pm$ 0.027	0.106 $\pm$ 0.001
train	0.429 $\pm$ 0.084	0.369 $\pm$ 0.175	0.315 $\pm$ 0.024	0.295 $\pm$ 0.000	0.005 $\pm$ 0.000	0.421 $\pm$ 0.101	0.371 $\pm$ 0.190	0.313 $\pm$ 0.024	0.295 $\pm$ 0.000	0.005 $\pm$ 0.000	0.426 $\pm$ 0.084	0.378 $\pm$ 0.176	0.319 $\pm$ 0.024	0.295 $\pm$ 0.000	0.004 $\pm$ 0.000
umbrella	0.352 $\pm$ 0.095	0.224 $\pm$ 0.195	0.358 $\pm$ 0.040	0.529 $\pm$ 0.000	0.086 $\pm$ 0.002	0.360 $\pm$ 0.096	0.296 $\pm$ 0.206	0.360 $\pm$ 0.040	0.529 $\pm$ 0.000	0.086 $\pm$ 0.002	0.353 $\pm$ 0.101	0.240 $\pm$ 0.209	0.362 $\pm$ 0.041	0.530 $\pm$ 0.000	0.088 $\pm$ 0.002
zebra	0.864 $\pm$ 0.063	0.548 $\pm$ 0.101	0.405 $\pm$ 0.052	0.871 $\pm$ 0.001	0.105 $\pm$ 0.003	0.851 $\pm$ 0.050	0.520 $\pm$ 0.147	0.540 $\pm$ 0.031	0.928 $\pm$ 0.035	0.104 $\pm$ 0.002	0.859 $\pm$ 0.054	0.477 $\pm$ 0.130	0.521 $\pm$ 0.034	0.908 $\pm$ 0.023	0.102 $\pm$ 0.

Table 27: Per-class faithfulness and reconstruction quality for SAEs and Gated SAEs on ImageNet and COCO (ViT-MAE B/16, $\ell = 9$). Mean $\pm$ std across 5 seeds.

Class	Vanilla SAE					Gated SAE				
	C-Ins $\uparrow$	C-Del $\uparrow$	C-Gini $\uparrow$	MSE $\downarrow$	D_{KL} $\downarrow$	C-Ins $\uparrow$	C-Del $\uparrow$	C-Gini $\uparrow$	MSE $\downarrow$	D_{KL} $\downarrow$
<i>ImageNet</i>										
cassette	0.902 $\pm$ 0.018	0.441 $\pm$ 0.028	0.190 $\pm$ 0.014	1.021 $\pm$ 0.000	0.391 $\pm$ 0.006	0.557 $\pm$ 0.058	0.423 $\pm$ 0.091	0.265 $\pm$ 0.024	1.021 $\pm$ 0.009	0.373 $\pm$ 0.010
chainsaw	0.993 $\pm$ 0.007	0.274 $\pm$ 0.114	0.228 $\pm$ 0.007	1.127 $\pm$ 0.000	0.438 $\pm$ 0.004	0.515 $\pm$ 0.048	0.387 $\pm$ 0.033	0.220 $\pm$ 0.047	1.124 $\pm$ 0.009	0.413 $\pm$ 0.014
church	0.714 $\pm$ 0.069	0.212 $\pm$ 0.054	0.196 $\pm$ 0.002	0.921 $\pm$ 0.000	0.488 $\pm$ 0.001	0.437 $\pm$ 0.047	0.468 $\pm$ 0.038	0.268 $\pm$ 0.065	0.927 $\pm$ 0.019	0.465 $\pm$ 0.021
springer	0.967 $\pm$ 0.012	0.137 $\pm$ 0.086	0.221 $\pm$ 0.001	0.984 $\pm$ 0.000	0.428 $\pm$ 0.004	0.681 $\pm$ 0.081	0.442 $\pm$ 0.078	0.276 $\pm$ 0.048	0.993 $\pm$ 0.015	0.386 $\pm$ 0.009
french horn	1.000 $\pm$ 0.000	0.004 $\pm$ 0.001	0.234 $\pm$ 0.003	1.067 $\pm$ 0.000	0.243 $\pm$ 0.001	0.920 $\pm$ 0.055	0.525 $\pm$ 0.066	0.284 $\pm$ 0.029	1.075 $\pm$ 0.011	0.229 $\pm$ 0.005
garbage truck	0.958 $\pm$ 0.029	0.436 $\pm$ 0.059	0.210 $\pm$ 0.010	0.970 $\pm$ 0.000	0.307 $\pm$ 0.002	0.668 $\pm$ 0.083	0.442 $\pm$ 0.083	0.234 $\pm$ 0.024	0.965 $\pm$ 0.009	0.282 $\pm$ 0.006
gas pump	0.864 $\pm$ 0.002	0.528 $\pm$ 0.008	0.209 $\pm$ 0.005	1.059 $\pm$ 0.000	0.413 $\pm$ 0.005	0.795 $\pm$ 0.121	0.441 $\pm$ 0.062	0.254 $\pm$ 0.038	1.065 $\pm$ 0.009	0.371 $\pm$ 0.004
golf ball	0.972 $\pm$ 0.004	0.377 $\pm$ 0.052	0.319 $\pm$ 0.009	0.932 $\pm$ 0.000	0.062 $\pm$ 0.001	0.967 $\pm$ 0.010	0.601 $\pm$ 0.063	0.332 $\pm$ 0.021	0.924 $\pm$ 0.009	0.058 $\pm$ 0.001
parachute	0.991 $\pm$ 0.006	0.254 $\pm$ 0.013	0.290 $\pm$ 0.006	0.799 $\pm$ 0.000	0.152 $\pm$ 0.002	0.849 $\pm$ 0.066	0.414 $\pm$ 0.078	0.336 $\pm$ 0.052	0.799 $\pm$ 0.011	0.147 $\pm$ 0.002
tench	1.000 $\pm$ 0.000	0.003 $\pm$ 0.003	0.278 $\pm$ 0.006	0.891 $\pm$ 0.000	0.126 $\pm$ 0.001	0.945 $\pm$ 0.051	0.613 $\pm$ 0.086	0.283 $\pm$ 0.013	0.901 $\pm$ 0.004	0.115 $\pm$ 0.002
<i>COCO</i>										
hot dog	0.491 $\pm$ 0.158	0.443 $\pm$ 0.084	0.200 $\pm$ 0.009	0.667 $\pm$ 0.000	0.076 $\pm$ 0.000	0.390 $\pm$ 0.108	0.150 $\pm$ 0.070	0.246 $\pm$ 0.045	0.658 $\pm$ 0.009	0.070 $\pm$ 0.003
teddy bear	0.250 $\pm$ 0.026	0.480 $\pm$ 0.057	0.274 $\pm$ 0.002	0.582 $\pm$ 0.000	0.104 $\pm$ 0.001	0.498 $\pm$ 0.165	0.165 $\pm$ 0.210	0.349 $\pm$ 0.052	0.580 $\pm$ 0.009	0.104 $\pm$ 0.006
train	1.000 $\pm$ 0.000	0.012 $\pm$ 0.025	0.164 $\pm$ 0.023	0.305 $\pm$ 0.000	0.006 $\pm$ 0.000	0.887 $\pm$ 0.108	0.532 $\pm$ 0.337	0.378 $\pm$ 0.029	0.367 $\pm$ 0.012	0.006 $\pm$ 0.001
umbrella	0.813 $\pm$ 0.056	0.146 $\pm$ 0.067	0.202 $\pm$ 0.012	0.541 $\pm$ 0.000	0.092 $\pm$ 0.002	0.336 $\pm$ 0.080	0.499 $\pm$ 0.146	0.320 $\pm$ 0.020	0.572 $\pm$ 0.013	0.121 $\pm$ 0.007
zebra	0.902 $\pm$ 0.043	0.273 $\pm$ 0.044	0.157 $\pm$ 0.007	0.875 $\pm$ 0.000	0.098 $\pm$ 0.002	0.495 $\pm$ 0.169	0.729 $\pm$ 0.090	0.246 $\pm$ 0.041	0.864 $\pm$ 0.014	0.108 $\pm$ 0.006

the factorization towards predictively faithful concept subspaces without requiring a downstream classification head. When compared to SAE baselines, PLACE demonstrates a superior ability to enforce sparsity without losing representational integrity. While Gated and Vanilla SAEs achieve competitive reconstruction errors (MSE ≈ 0.07) and moderate D_{KL} , they struggle to maintain tight concept sparsity (C-Gini dropping to ≈ 0.70 – 0.86), whereas PLACE consistently achieves C-Gini scores exceeding 0.90. Furthermore, on complex COCO scenes, SAEs exhibit a severe drop in concept insertion (e.g., C-Ins ≈ 0.16 for hot dog), highlighting PLACE’s enhanced robustness in isolating semantically meaningful features in multi-object environments.

Conversely, for the supervised **ViT-B/16** architecture (Table 26), the supervised +KL variant yields the best overall faithfulness alongside minimal D_{KL} degradation. The +Gram variant remains a robust, fully unsupervised alternative, matching the near-zero divergence of the +KL reference while preserving strong C-Gini sparsity. We explicitly note that for this specific supervised backbone, the reconstruction-only PLACE baseline achieves a KL divergence less than 10^{-3} (reported as 0.000 due to rounding), indicating that the subspace captured by pure MSE already aligns with the classifier’s decision boundaries. The contrast with SAE baselines here is striking. Both Vanilla and Gated SAEs fundamentally fail to preserve the original predictive distribution, suffering from KL divergence degradation ($D_{KL} \approx 0.5$ – 1.7 on ImageNet) and a collapse in sparsity (C-Gini falling to 0.25 – 0.45). This indicates that standard L_1 -penalized SAE objectives disrupt the highly specific, discriminative manifolds learned by supervised ViTs—a limitation effectively bypassed by PLACE. Notably, we observe a systemic architectural difference: supervised ViT-B/16 tends to exhibit remarkably high C-Ins values coupled with comparatively lower C-Del scores, suggesting it encodes class-specific features with a higher degree of redundancy than self-supervised models.

Finally, the self-supervised **ViT-MAE B/16** architecture (Table 27) presents a different behavioral profile. Unlike DINOv2, which benefits from the +Gram regularization, ViT-MAE exhibits less sensitivity to the introduced geometric and divergence losses, with metrics remaining relatively static across the PLACE variants. Across ImageNet, ViT-MAE yields high C-Ins scores alongside moderate-to-low C-Del values—a redundancy pattern somewhat echoing the supervised ViT-B/16, but accompanied by higher reconstruction errors (MSE ≈ 1.0) and elevated KL divergences ($D_{KL} \approx 0.25$ – 0.50). This entanglement poses an even greater challenge for traditional SAEs. Vanilla and Gated SAEs evaluated on ViT-MAE exhibit an almost complete collapse in concept sparsity across both datasets, with C-Gini frequently falling to ~ 0.20 . While PLACE also encounters difficulties with C-Del volatility on MAE (e.g., COCO ‘hot dog’ and ‘teddy bear’), it prevents this sparsity collapse, maintaining C-Gini scores that are higher than the SAE baselines. This reinforces the hypothesis that while reconstructive masked autoencoding yields dense, highly distributed representations that resist linear disentanglement, the PLACE framework provides a more constrained and structurally sound factorization than unconstrained sparse dictionary learning.

E.12 Per-class agreement scores

This section provides the comprehensive, class-by-class breakdown of the alignment metrics summarized in Section 4.3 (Table 2). We evaluate how closely the unsupervised concept banks extracted via

pure reconstruction and PLACE (+Gram) recover the supervised representation subspace (+KL) on both DINOv2 ViT-B/14 and ViT-MAE.

Alignment is measured via RSVC [37] using Maximum Cosine Similarity (MCS) for individual vector alignment and CMCS (Pearson correlation of Lasso predictions) for broader subspace span. Values are reported as mean $\pm$ std over 5 seeds.

As shown in Table 28, the impact of geometric regularization varies significantly between the backbones. For DINOv2, the Gram-matrix alignment is remarkably consistent: PLACE (+Gram) improves or perfectly matches the Maximum Cosine Similarity (MCS) of the pure reconstruction baseline in 9 out of 10 evaluated classes. This confirms its robustness in aligning individual concepts with predictively-relevant directions.

Conversely, for ViT-MAE, pure reconstruction already achieves exceptionally high baseline alignment scores (often reaching an MCS > 0.90), leaving much less headroom for geometric refinement. Consequently, the +Gram regularization improves individual concept alignment in only a minority of classes (e.g., church, garbage_truck, cassette). This further supports our broader observation that the masked autoencoder’s representation space is deeply entangled by the reconstruction objective, making it less responsive to the explicit geometric factorizations that successfully structure the contrastive DINOv2 space.

Table 28: Per-class agreement of unsupervised concept banks with the supervised (+KL) reference on ImageNet for both DINOv2 and ViT-MAE. Higher values indicate stronger alignment. Best MCS per class for each backbone is bolded.

Class	DINOv2 (ViT-B/14)				ViT-MAE (ViT-B/16)			
	Reconstruction vs. +KL		+Gram vs. +KL		Reconstruction vs. +KL		+Gram vs. +KL	
	MCS $\uparrow$	CMCS $\uparrow$	MCS $\uparrow$	CMCS $\uparrow$	MCS $\uparrow$	CMCS $\uparrow$	MCS $\uparrow$	CMCS $\uparrow$
cassette	0.859 $\pm$ 0.018	0.754 $\pm$ 0.016	0.875 $\pm$ 0.010	0.747 $\pm$ 0.014	0.900 $\pm$ 0.014	0.867 $\pm$ 0.011	0.905 $\pm$ 0.019	0.862 $\pm$ 0.011
chainsaw	0.841 $\pm$ 0.019	0.707 $\pm$ 0.007	0.828 $\pm$ 0.011	0.687 $\pm$ 0.007	0.905 $\pm$ 0.025	0.827 $\pm$ 0.014	0.903 $\pm$ 0.027	0.816 $\pm$ 0.017
church	0.841 $\pm$ 0.014	0.712 $\pm$ 0.019	0.866 $\pm$ 0.007	0.703 $\pm$ 0.018	0.907 $\pm$ 0.047	0.838 $\pm$ 0.017	0.942 $\pm$ 0.014	0.838 $\pm$ 0.018
english_springer	0.822 $\pm$ 0.013	0.760 $\pm$ 0.009	0.863 $\pm$ 0.011	0.750 $\pm$ 0.009	0.911 $\pm$ 0.033	0.862 $\pm$ 0.011	0.902 $\pm$ 0.027	0.858 $\pm$ 0.011
french_horn	0.830 $\pm$ 0.017	0.719 $\pm$ 0.012	0.849 $\pm$ 0.011	0.705 $\pm$ 0.014	0.925 $\pm$ 0.016	0.871 $\pm$ 0.007	0.922 $\pm$ 0.021	0.863 $\pm$ 0.008
garbage_truck	0.819 $\pm$ 0.007	0.697 $\pm$ 0.007	0.843 $\pm$ 0.007	0.687 $\pm$ 0.011	0.879 $\pm$ 0.033	0.847 $\pm$ 0.005	0.912 $\pm$ 0.012	0.843 $\pm$ 0.005
gas_pump	0.775 $\pm$ 0.036	0.663 $\pm$ 0.011	0.798 $\pm$ 0.038	0.658 $\pm$ 0.009	0.953 $\pm$ 0.014	0.869 $\pm$ 0.012	0.940 $\pm$ 0.027	0.850 $\pm$ 0.013
golf_ball	0.909 $\pm$ 0.017	0.794 $\pm$ 0.008	0.912 $\pm$ 0.016	0.782 $\pm$ 0.011	0.954 $\pm$ 0.013	0.878 $\pm$ 0.008	0.948 $\pm$ 0.019	0.874 $\pm$ 0.008
parachute	0.902 $\pm$ 0.006	0.777 $\pm$ 0.007	0.908 $\pm$ 0.005	0.774 $\pm$ 0.009	0.927 $\pm$ 0.021	0.889 $\pm$ 0.004	0.918 $\pm$ 0.026	0.881 $\pm$ 0.004
tench	0.805 $\pm$ 0.022	0.734 $\pm$ 0.016	0.829 $\pm$ 0.036	0.727 $\pm$ 0.014	0.946 $\pm$ 0.014	0.878 $\pm$ 0.005	0.940 $\pm$ 0.026	0.870 $\pm$ 0.006

E.13 Robustness to Image Perturbations

In this section, we analyze the structural resilience of PLACE against input signal degradations. We verify whether the extracted semantic directions and their predictive faithfulness remain stable even in non-ideal, real-world scenarios.

Experimental Setup. Following the ImageNet-C protocol [31], we apply five orthogonal perturbations (Gaussian Noise, Gaussian Blur, JPEG Compression, Contrast, and Pixelate) at five increasing severity levels ($L \in \{1, \dots, 5\}$). We do not filter the images to execute the algorithm only on correctly classified ones; avoiding this filtering prevents dataset reduction and ensures we strictly isolate the effect of the noise perturbations on the representation geometry.

Discussion. The results in Table 29 highlight a strength of PLACE: the stability of the representation and the uncoupling of semantic faithfulness from structural noise. Across almost all perturbations, dictionary stability (W_{stab}) remains strong (> 0.88), and activation sparsity (C -Gini) is practically unfazed. Interestingly, rather than deteriorating, the predictive faithfulness (D_{KL}) generally tightens (decreases) or remains stable relative to the clean baseline as perturbation severity increases, particularly in noise, blur, and pixelation scenarios. Meanwhile, reconstruction metrics (MSE and $\mathcal{L}_{gram}$) exhibit only constrained variance. This demonstrates that our localized geometric constraints successfully anchor the extraction process to the model’s intrinsic abstractions, preventing the concepts from collapsing into visual noise, and confirming that the extracted concepts remain functionally robust in downstream predictive influence.

Table 29: Comprehensive Robustness Analysis with stability and faithfulness metrics across five independent perturbations. Mean baseline values for clean images: $MSE = 0.0944$, $W_{stab} = 0.9402$, $C\text{-Gini} = 0.9269$, $D_{KL} = 1.3829$, $\mathcal{L}_{gram} = 0.0005$. We report results on DINOv2 ViT-B/14 (MSE+Gram, ImageNet *church*, $r=25$, $s=2$), mean $\pm$ std over five seeds.

Perturbation	Metric	L1	L2	L3	L4	L5
Gaussian Noise	MSE	0.0935 ± 0.0022	0.0930 ± 0.0010	0.0903 ± 0.0012	0.0862 ± 0.0014	0.0780 ± 0.0007
	W_{stab}	0.9097	0.9480	0.9288	0.8836	0.8889
	$C\text{-Gini}$	0.9270 ± 0.0092	0.9323 ± 0.0032	0.9288 ± 0.0038	0.9249 ± 0.0053	0.9273 ± 0.0035
	D_{KL}	1.2902 ± 0.1526	1.2518 ± 0.0485	1.0652 ± 0.0776	1.0010 ± 0.0694	0.9120 ± 0.0388
	$\mathcal{L}_{gram}$	0.0004 ± 0.0001	0.0005 ± 0.0000	0.0004 ± 0.0000	0.0004 ± 0.0001	0.0008 ± 0.0002
Gaussian Blur	MSE	0.0952 ± 0.0003	0.0931 ± 0.0017	0.0930 ± 0.0017	0.0935 ± 0.0019	0.0990 ± 0.0017
	W_{stab}	0.9543	0.9483	0.9686	0.9312	0.9491
	$C\text{-Gini}$	0.9306 ± 0.0013	0.9244 ± 0.0060	0.9280 ± 0.0057	0.9315 ± 0.0085	0.9268 ± 0.0066
	D_{KL}	1.4464 ± 0.0323	1.3124 ± 0.0885	1.2385 ± 0.1189	0.9192 ± 0.0939	0.7228 ± 0.0448
	$\mathcal{L}_{gram}$	0.0005 ± 0.0000	0.0005 ± 0.0001	0.0004 ± 0.0000	0.0004 ± 0.0000	0.0004 ± 0.0001
JPEG Comp.	MSE	0.0959 ± 0.0002	0.0944 ± 0.0016	0.0948 ± 0.0022	0.0947 ± 0.0014	0.0895 ± 0.0010
	W_{stab}	0.9437	0.9607	0.9465	0.9432	0.9445
	$C\text{-Gini}$	0.9315 ± 0.0010	0.9259 ± 0.0062	0.9264 ± 0.0098	0.9288 ± 0.0055	0.9139 ± 0.0069
	D_{KL}	1.4505 ± 0.0300	1.3003 ± 0.0943	1.2801 ± 0.1347	1.2554 ± 0.0604	1.0057 ± 0.0621
	$\mathcal{L}_{gram}$	0.0004 ± 0.0001	0.0004 ± 0.0000	0.0005 ± 0.0001	0.0005 ± 0.0001	0.0006 ± 0.0001
Contrast	MSE	0.0934 ± 0.0016	0.0937 ± 0.0013	0.0926 ± 0.0025	0.0935 ± 0.0021	0.0922 ± 0.0004
	W_{stab}	0.9543	0.9672	0.9529	0.9459	0.9689
	$C\text{-Gini}$	0.9234 ± 0.0063	0.9257 ± 0.0051	0.9196 ± 0.0135	0.9282 ± 0.0078	0.9303 ± 0.0017
	D_{KL}	1.3018 ± 0.1076	1.3533 ± 0.0848	1.2708 ± 0.1883	1.3685 ± 0.1358	1.3094 ± 0.0375
	$\mathcal{L}_{gram}$	0.0005 ± 0.0001	0.0005 ± 0.0000	0.0006 ± 0.0001	0.0006 ± 0.0001	0.0006 ± 0.0001
Pixelate	MSE	0.0939 ± 0.0017	0.0948 ± 0.0017	0.0924 ± 0.0015	0.0801 ± 0.0016	0.0787 ± 0.0012
	W_{stab}	0.9457	0.9890	0.9926	0.9242	0.9233
	$C\text{-Gini}$	0.9275 ± 0.0066	0.9285 ± 0.0062	0.9306 ± 0.0045	0.9243 ± 0.0076	0.9203 ± 0.0072
	D_{KL}	1.1853 ± 0.0891	1.2319 ± 0.1203	1.3160 ± 0.0763	0.6125 ± 0.0504	0.4480 ± 0.0470
	$\mathcal{L}_{gram}$	0.0004 ± 0.0001	0.0004 ± 0.0000	0.0004 ± 0.0000	0.0007 ± 0.0002	0.0007 ± 0.0001

F Evaluation Metrics Description

To rigorously evaluate the quality, faithfulness, and interpretability of the extracted concepts, we employ the following standard metrics. Let $f(\cdot)$ denote the classification model, $A \in \mathbb{R}^{N \times D}$ the original patch-token activations, and $U \in \mathbb{R}^{N \times r}$, $W \in \mathbb{R}^{D \times r}$ the extracted concept activation and dictionary matrices, respectively, such that $A \approx UW^\top$.

Concept Insertion (C-Ins): Measures the ability of the extracted concepts to independently drive the model’s prediction. Starting from a completely masked image, we progressively reveal the input patches corresponding to the highest values in U . For K insertion steps, let $x_{ins}^{(k)}$ be the image with the top k -th fraction of patches revealed, and y the target class. C-Ins is computed as the Area Under the Curve (AUC) of the target-class probability:

$$C\text{-Ins} = \frac{1}{K} \sum_{k=1}^K p(y \mid x_{ins}^{(k)}) \quad (39)$$

A higher score indicates that the concepts capture the core features necessary for classification.

Concept Deletion (C-Del): Measures the model’s causal reliance on the extracted concepts. Conversely to C-Ins, we progressively mask the top-activating patches. Let $x_{del}^{(k)}$ be the image with the top k -th fraction of patches masked. Following FACE [7], we report $1 - \text{AUC}$, so a higher score indicates a sharper drop in confidence:

$$C\text{-Del} = 1 - \left(\frac{1}{K} \sum_{k=1}^K p(y \mid x_{del}^{(k)}) \right) \quad (40)$$

Concept Gini Sparsity (C-Gini): Quantifies the monosemanticity and distinctiveness of the extracted concepts. For a given concept activation vector $\mathbf{u}_j \in \mathbb{R}^N$ (the j -th column of the coefficient matrix U , representing the magnitude of concept c_j across the N spatial patch tokens), the spatial Gini coefficient is computed as:

$$G(\mathbf{u}_j) = \frac{\sum_{i=1}^N \sum_{m=1}^N |u_{i,j} - u_{m,j}|}{2N \sum_{i=1}^N u_{i,j}} \quad (41)$$

We report the mean C-Gini across all r concepts. The Gini index operates as a direct quantitative proxy for monosemanticity within the spatial geometry of Vision Transformers. Because human-understandable visual features – such as isolated object part – are inherently localized, a rigorously monosemantic direction in the latent space naturally yields a spatially sparse activation map. Therefore, a high C-Gini score (approaching 1.0) mathematically guarantees that the extracted concept does not collapse into a polysemantic mixture of global statistics (e.g., activating densely across the entire scene or distributed textures), but instead fires selectively on a highly specific and isolated subset of visual features.

Reconstruction Error (MSE): Evaluates the geometric fidelity of the low-rank factorization, computed as the Mean Squared Error between the original activations and the reconstructed space:

$$\text{MSE} = \frac{1}{N \cdot D} \|A - UW^\top\|_F^2 \quad (42)$$

Predictive Divergence (D_{KL}): Measures the downstream predictive consistency. It is the Kullback-Leibler divergence between the class probability distribution produced by the original frozen model and the distribution produced when intermediate activations are replaced by their reconstruction $\hat{A} = UW^\top$:

$$D_{\text{KL}} = \sum_{c=1}^C p_c(f(A)) \log \left(\frac{p_c(f(A))}{p_c(f(UW^\top))} \right) \quad (43)$$

Maximum Cosine Similarity (MCS): Evaluates representation subspace agreement between our unsupervised dictionary W and a supervised reference dictionary $\hat{W}$ as used in RSVC [37]. We find the optimal assignment via the Hungarian algorithm (permutation π):

$$\text{MCS}(W, \hat{W}) = \frac{1}{r} \max_{\pi \in \Pi} \sum_{i=1}^r \frac{\mathbf{w}_i^\top \hat{\mathbf{w}}_{\pi(i)}}{\|\mathbf{w}_i\|_2 \|\hat{\mathbf{w}}_{\pi(i)}\|_2} \quad (44)$$

Subspace Correlation (CMCS): Assesses the broader representational subspace agreement. It is the Pearson correlation coefficient ρ of the predictions made by Lasso regression models trained on the respective concept activation spaces, used in the RSVC framework [37].

Dictionary Stability (W-Stab): Evaluates the reproducibility of the factorization across different random initialization seeds. For two dictionaries $W^{(1)}$ and $W^{(2)}$ extracted from independent runs, it is computed exactly as the MCS equation above, verifying that the extraction robustly converges to consistent concept directions.

G Broader Impact

PLACE is a diagnostic tool for understanding the internal representations of Vision Foundation Models. On the positive side, faithful concept-based explanations can support safety auditing and trust calibration when deploying VFMs in high-stakes settings such as medical imaging or autonomous systems, by revealing which visual patterns the model relies on. Making such tools applicable without classifier supervision broadens their reach to the growing class of foundation models that are deployed as general-purpose feature extractors with no fixed downstream task.

On the negative side, detailed knowledge of a model’s concept structure could be exploited to craft targeted adversarial perturbations that manipulate specific concept activations while remaining imperceptible, or to engineer inputs that suppress safety-relevant concepts. More broadly, concept-based explanations may convey a false sense of transparency: users may over-trust a model because its explanations appear interpretable, even when faithfulness metrics indicate imperfect alignment. We emphasise that PLACE provides a decomposition of intermediate activations, not a causal account of the model’s decision process, and should be used as one component of a broader evaluation pipeline rather than as a standalone guarantee of model reliability.

H Additional Qualitative Results

In this section, we present extended visualizations for all 10 ImageNet classes used in our evaluation to demonstrate the consistency of our approach across diverse semantic categories. For each class,

three concepts extracted from the supervised ViT-B/16 model using the PLACE (+Gram) loss at layer $\ell = 9$.

Each class panel displays the concepts with the following views: the original image (highlighting the highest-activating region), the spatial heatmap produced by the coefficient matrix $\mathbf{U}$, and the PCA projection of the original patch-token activations as described in Appendix D.3.

Object and Context Disentanglement. Crucially, our qualitative analysis reveals that PLACE does not merely act as a foreground segmenter. Because the extraction is fully unsupervised and task-agnostic, the Semi-NMF factorization naturally captures the entire semantic composition of the scene, cleanly disentangling the primary object from its contextual environment. As highlighted in the figures below, the top concepts frequently isolate interacting subjects (e.g., humans holding objects), specific backgrounds, skies, and terrain, demonstrating the richness of the discovered geometric subspace.




Figure 5: Top-3 highest-activating concepts for evaluated ImageNet classes (Part 1/5). Notably, for the `cassette` class, the first concept cleanly captures the background. For the `chainsaw`, the factorization isolates the interacting environment (a log/tree in the first concept) separately from the tool itself.




Figure 6: Top-3 highest-activating concepts (Part 2/5). Here, PLACE clearly disentangles the main subject from its environment: the second concept of church perfectly isolates the sky, while the third concept of english_springer isolates the lawn.




Figure 7: Top-3 highest-activating concepts (Part 3/5). The decomposition successfully captures complex interactions and backgrounds: the first concept of french_horn focuses on the person playing the instrument, while a top concept in garbage_truck isolates the background trees rather than the vehicle itself.




Figure 8: Top-3 highest-activating concepts (Part 4/5). Contextual separation is prominent: `gas_pump` features background foliage in its components, while `golf_ball` cleanly separates the core object (the ball) from the surrounding environmental elements like trees and the lawn.




Figure 9: Top-3 highest-activating concepts (Part 5/5). The semantic disentanglement is remarkably coherent here: for `parachute`, the components strictly isolate the sky (first), the skydiver (second), and the canopy (third). Similarly, for `tench`, the decomposition distinctly separates the person holding the catch (first), the fish itself (second), and the background trees (third).