

# From Weights to Concepts: Data-Free Interpretability of CLIP via Singular Vector Decomposition

## Supplementary Material

- A Additional implementation details** **2**
- A.1 Multi-head Attention . . . . . 2
- A.2 Folding LN into MHA . . . . . 2
- A.3 Projecting Singular Vectors into the Multimodal Space . . . . . 3
- A.4 Projecting Reconstructions back into the Residual Stream . . . . . 3
- A.5 Spurious Feature Removal . . . . . 4
- A.6 Removing NSFW Concepts . . . . . 4
- A.7 Improving Classification Performance . . . . . 4
- A.8 Model Adaptation . . . . . 5
  
- B Ablating the Concept Pool** **5**
  
- C Extended Quantitative Analysis for CLIP ViT-L/14** **7**
- C.1 Interpretability-Fidelity Analysis on Additional Layers . . . . . 7
- C.2 Analysis of Left Singular Vectors . . . . . 8
  
- D Qualitative Results** **8**
  
- E Pseudocode of COMP** **9**
  
- F GPT-5 Prompts** **9**

## A. Additional implementation details

This section provides the theoretical foundations and additional implementation details of SITH. We begin by detailing the mechanistic view of Multi-Head Attention (MHA) to show how we isolate the Value-Output (VO) weight matrix from each attention head (Sec. A.1). We also describe how we fold the Layer Normalization (LN) into the MHA weights to ensure that each attention head directly reads from the residual stream (Sec. A.2).

Next, we provide more details on how we project the singular vectors into the multimodal embedding space to minimize the modality gap and interpret them via the concept pool (Sec. A.3). We also detail how we project the reconstructed vectors back into the residual stream for evaluation purposes (Sec. A.4).

Finally, we describe the steps for our model editing and model adaptation analysis experiments. This covers the methodology for pruning singular vectors to remove spurious features (Sec. A.5) and NSFW concepts (Sec. A.6), the identification and amplification of task-relevant singular vectors to improve classification performance (Sec. A.7), and the specific settings used to analyze model adaptation under Full Fine-tuning and LoRA (Sec. A.8).

### A.1. Multi-head Attention

In the standard implementation of CLIP-ViT [17], the multi-head attention (MHA) mechanism is computed as the concatenation of the outputs of each attention head, followed by a linear transformation:

$$\text{MHA}(\mathbf{X}) = \text{Concat}(\mathbf{H}_1(\mathbf{X}), \dots, \mathbf{H}_H(\mathbf{X}))\mathbf{W}_O, \quad (8)$$

where  $\mathbf{X} \in \mathbb{R}^{(P+1) \times D}$  are the [CLS] token and input patches,  $H$  is the number of attention heads, and  $\mathbf{W}_O \in \mathbb{R}^{H \cdot d_h \times D}$  is the output weight matrix, with  $d_h = D/H$  being the dimension of each head. Each attention head is computed as:

$$\mathbf{H}_h(\mathbf{X}) = \text{softmax} \left( \frac{\mathbf{X}\mathbf{W}_Q^h(\mathbf{X}\mathbf{W}_K^h)^T}{\sqrt{d_h}} \right) \mathbf{X}\mathbf{W}_V^h, \quad (9)$$

where  $\mathbf{W}_Q^h, \mathbf{W}_K^h, \mathbf{W}_V^h \in \mathbb{R}^{D \times d_h}$  are the query, key, and value weight matrices for head  $h$ , respectively.

While standard, this implementation can make it challenging to apply mechanistic interpretability methods that aim to analyze the contributions of individual attention heads. To address this, we can express the MHA mechanism in a mathematically equivalent form that disentangles the contributions of each head and ensures that each attention head directly reads from and writes to the residual stream.

First, note that concatenation followed by a linear transformation can be expressed as a sum of linear transformations applied to each head’s output. Indeed, given the output matrix  $\mathbf{W}_O$ , we can partition it into  $H$  sub-matrices

$\mathbf{W}_O^h \in \mathbb{R}^{d_h \times D}$ , such that  $\mathbf{W}_O = [\mathbf{W}_O^1; \dots; \mathbf{W}_O^H]$ . Then, we can rewrite the concatenation followed by the linear transformation as:

$$[\mathbf{H}_1(\mathbf{X}), \dots, \mathbf{H}_H(\mathbf{X})] \begin{bmatrix} \mathbf{W}_O^1 \\ \vdots \\ \mathbf{W}_O^H \end{bmatrix} = \sum_{h=1}^H \mathbf{H}_h(\mathbf{X})\mathbf{W}_O^h. \quad (10)$$

This allows us to express the MHA mechanism as the sum of  $H$  independent attention heads:

$$\text{MHA}(\mathbf{X}) = \sum_{h=1}^H \mathbf{H}'_h(\mathbf{X}), \quad (11)$$

where we merge the output linear transformation into each head’s computation as:

$$\mathbf{H}'_h(\mathbf{X}) = \mathbf{H}_h(\mathbf{X})\mathbf{W}_O^h. \quad (12)$$

Following this, we split the computation of each attention head into the Query-Key (QK) and Value-Output (VO) circuits. Formally, given the formulation of each attention head as:

$$\mathbf{H}'_h(\mathbf{X}) = \text{softmax} \left( \frac{\mathbf{X}\mathbf{W}_Q^h\mathbf{W}_K^{hT}\mathbf{X}^T}{\sqrt{d_h}} \right) \mathbf{X}\mathbf{W}_V^h\mathbf{W}_O^h, \quad (13)$$

we can merge the Query and Key weight matrices into a single matrix  $\mathbf{W}_{QK}^h = \mathbf{W}_Q^h\mathbf{W}_K^{hT} \in \mathbb{R}^{D \times D}$ , and the Value and Output weight matrices into another matrix  $\mathbf{W}_{VO}^h = \mathbf{W}_V^h\mathbf{W}_O^h \in \mathbb{R}^{D \times D}$ . Consequently, each attention head can be expressed as:

$$\mathbf{H}'_h(\mathbf{X}) = \text{softmax} \left( \frac{\mathbf{X}\mathbf{W}_{QK}^h\mathbf{X}^T}{\sqrt{d_h}} \right) \mathbf{X}\mathbf{W}_{VO}^h, \quad (14)$$

where the QK matrix governs how the attention weights are computed from the input patches, and the VO matrix determines how the attended patches are projected back into the residual stream.

### A.2. Folding LN into MHA

The CLIP-ViT architecture utilizes a Pre-LN formulation [77], where the Layer Normalization (LN) operation is applied to the input of every Multi-Head Attention (MHA) and Feed-Forward Network (FFN) block, rather than to their outputs. Consequently, the attention heads do not directly read from the residual stream  $\mathbf{X}$ , but rather a normalized version  $\text{LN}(\mathbf{X})$  of it. To ensure that each attention head reads directly from the residual stream, we fold the linear components of the LN layer into the MHA weights.

**Folding Affine Parameters.** Let  $w$  and  $b$  denote the learnable weight and bias parameters of the LN layer, respectively. We absorb these parameters into the query, key, and value

projection matrices ( $\mathbf{W}_Q, \mathbf{W}_K, \mathbf{W}_V$ ) and their respective biases ( $\mathbf{b}_Q, \mathbf{b}_K, \mathbf{b}_V$ ) as follows:

$$\mathbf{W}'_{\{Q,K,V\}} = \text{diag}(\mathbf{w})\mathbf{W}_{\{Q,K,V\}} \quad (15)$$

$$\mathbf{b}'_{\{Q,K,V\}} = \mathbf{b}_{\{Q,K,V\}} + \mathbf{W}'^T_{\{Q,K,V\}}\mathbf{b} \quad (16)$$

where  $\text{diag}(\mathbf{w})$  is a diagonal matrix with the elements of  $\mathbf{w}$  on the diagonal. This transformation ensures that the analysis of the folded weights  $\mathbf{W}'$  accounts for the component-wise scaling and shifting applied by the LN.

**Handling Centering and Normalization.** Beyond the affine parameters, the core operation of Layer Normalization involves centering the input vector and scaling it to have unit variance. Since normalization does not affect the direction of the input vectors, we can safely omit it when analyzing the reading and writing directions of the attention heads.

On the other hand, centering the input is equivalent to projecting the input vectors onto a hyperplane orthogonal to the all-ones vector  $\mathbf{1}$ , thus it changes the direction of the input vectors. However, in the CLIP ViT architecture, every transformer block and the final projection to the multimodal embedding space are preceded by a LayerNorm. This implies that any information encoded in the direction of the all-ones vector is systematically removed before it can be processed by subsequent layers or the final projection. Consequently, we can posit a theoretically equivalent model where every block is constrained to operate in (*i.e.*, read from and write to) the subspace orthogonal to  $\mathbf{1}$ , thus making the explicit centering operation of the LN redundant.

To align our analysis with this effective computational model, we project the weight matrices onto the orthogonal complement of  $\mathbf{1}$ . This ensures that we only analyze the active subspace of the residual stream. Practically, this is implemented via mean subtraction [44]. For input-reading weights (*i.e.*,  $\mathbf{W}_Q, \mathbf{W}_K, \mathbf{W}_V$ ), we subtract the mean from each column. This ensures that the dot product with any vector parallel to  $\mathbf{1}$  is zero, making the weights invariant to the mean of the input, which is removed nevertheless by the LN. For output-writing weights (*i.e.*,  $\mathbf{W}_O$ ), we subtract the mean from each row. This ensures that the output of the head sums to zero, guaranteeing that the head never writes into the  $\mathbf{1}$  direction of the residual stream.

**Final Folded Weights.** After applying the affine folding and mean subtraction steps, we obtain the final folded weight matrices  $\mathbf{W}'_Q, \mathbf{W}'_K, \mathbf{W}'_V, \mathbf{W}'_O$  that are used in place of the original weights to derive the QK and VO weight matrices (see Sec. A.1).

### A.3. Projecting Singular Vectors into the Multimodal Space

In Sec. 3.3 of the main paper, we showed that, as the left ( $\mathbf{u} \in \mathbf{U}$ ) and right ( $\mathbf{v} \in \mathbf{V}$ ) singular vectors of the VO matrix lie in the same residual stream space as the image

patches, we can project them into the multimodal embedding space to interpret them via the concept pool  $\Gamma$ .

Here we provide a more detailed explanation of this projection step.

**Layer Normalization.** At the end of the CLIP-ViT architecture, the representation of the [CLS] token is passed through a final Layer Normalization (LN) layer before being projected by the vision projection matrix  $\mathbf{W}_p$ . To ensure that the singular vectors are correctly projected into the multimodal space, we similarly apply the LN transformation to them before the projection. Specifically, given a right singular vector  $\mathbf{v} \in \mathbb{R}^D$  (the same applies to left singular vectors), we compute its unaligned multimodal representation as follows:

$$\tilde{\mathbf{v}} = \text{norm}(\mathbf{W}_p^T \text{LN}(\mathbf{v})) \quad (17)$$

where  $\text{norm}(\cdot)$  indicates normalization to unit length.

**Mitigating the Multimodality Gap.** A well-documented phenomenon in contrastive vision-language models is the modality gap [37], where image and text embeddings tend to cluster in distinct, cone-shaped regions of the unit hypersphere. This gap can hinder the interpretability of singular vectors when projected into the multimodal space, as they may not align with the text embeddings of their corresponding concepts.

To address this gap, we adopt the re-centering approach proposed by Bhalla et al. [2]. Specifically, we geometrically align the two modalities by mean-centering both the projected singular vectors and the concept embeddings using the estimated means of the image and text embedding distributions:

$$\hat{\mathbf{v}} = \text{norm}(\tilde{\mathbf{v}} - \boldsymbol{\mu}_{img}) \quad (18)$$

$$\hat{\gamma}_i = \text{norm}(\text{norm}(\mathcal{E}_T(\gamma_i)) - \boldsymbol{\mu}_{txt}) \quad \forall i = 1, \dots, C \quad (19)$$

where  $\boldsymbol{\mu}_{img}$  is the mean image embedding computed over the CC12M dataset [10], and  $\boldsymbol{\mu}_{txt}$  is the mean text embedding computed over the text concepts in the concept pool  $\Gamma$ .

### A.4. Projecting Reconstructions back into the Residual Stream

To evaluate the fidelity of the COMP reconstructions, as well as to measure the effect on downstream performance when replacing singular vectors with their reconstructions, we need to project the decompositions made in the multimodal space back into the residual stream space.

Given the sparse coefficient vector  $\mathbf{c} \in \mathbb{R}^C$  obtained via COMP and the matrix of aligned concept embeddings  $\hat{\Gamma} \in \mathbb{R}^{C \times d}$ , we first compute the reconstructed singular vector in the centered multimodal space as follows:

$$\hat{\mathbf{v}}_{rec} = \text{norm}(\hat{\Gamma}^T \mathbf{c}), \quad (20)$$

where the normalization ensures that the reconstructed vector lies on the unit hypersphere surface of the centered multi-modal space. Then, to reverse the modality gap mitigation step, we move the reconstructed vector back into the image cone by adding the image mean and re-normalizing:

$$\tilde{\mathbf{v}}_{rec} = \text{norm}(\hat{\mathbf{v}}_{rec} + \boldsymbol{\mu}_{img}). \quad (21)$$

To project the reconstructed vector back into the residual stream space, we utilize the Moore-Penrose pseudo-inverse of the projection matrix  $\mathbf{W}_p$  [3, 43, 56]:

$$\mathbf{v}'_{rec} = \mathbf{W}_p^\dagger \tilde{\mathbf{v}}_{rec}, \quad (22)$$

where  $\mathbf{W}_p^\dagger \in \mathbb{R}^{d \times D}$  is the pseudo-inverse matrix.

Theoretically, a complete inversion of the forward pass would require reversing the LayerNorm operation applied before the projection. However, we found that attempting to invert the affine transformation of the LN yielded suboptimal results. Consequently, we omit the LN inversion entirely when projecting back into the residual stream.

Finally, as the singular vectors are defined to have unit norm, we re-normalize the reconstructed vector in the residual stream space:

$$\mathbf{v}_{rec} = \text{norm}(\mathbf{v}'_{rec}) \quad (23)$$

### A.5. Spurious Feature Removal

In Sec. 5.1 of the main paper, we demonstrate how SITH can be employed to identify which singular vectors encode spurious “background” or “location” features, and subsequently remove them from the model to enhance robustness against background and location biases on the Waterbirds classification dataset [62]. Given the high dimensionality of the search space (totaling 4,096 singular vectors across the last four layers of CLIP ViT-L/14), manually inspecting the semantic explanations generated by COMP for every vector is intractable. To automate this inspection process, we leverage an LLM as a semantic judge. For this specific intervention, we utilize COMP ( $\lambda = 0.3$ ) with a sparsity level of  $K = 5$ , ensuring that the explanations capture the dominant semantic concept encoded by each vector.

We employ GPT-5-mini [51] to classify the degree to which the provided concept set relates to “background” or “location” features. The model assigns a relevance score on a Likert scale from 1 (“not related at all”) to 5 (“strongly related”); the exact prompt is provided in Tab. 23. Finally, we apply a hard thresholding operation: any singular vector receiving a score  $\geq 3$  is classified as spurious, and its corresponding singular value  $\sigma_i$  is set to zero, effectively nullifying its contribution to the residual stream.

### A.6. Removing NSFW Concepts

In Sec. 5.2 of the main paper, we illustrated how SITH can be utilized to identify and eliminate singular vectors that encode inappropriate or unsafe content, thereby enhancing the

safety of the CLIP ViT-L/14 model for retrieval tasks. Here, we provide additional implementation details regarding this experiment.

Following the same automated discovery pipeline described in Sec. A.5, we employ an LLM-as-a-judge to identify singular vectors encoding inappropriate content. Specifically, we prompt the LLM to evaluate singular vectors against the seven categories of inappropriate content defined by SafeCLIP [57]: *hate*, *harassment*, *violence*, *self-harm*, *sexual content*, *shocking images*, and *illegal activity* (see Tab. 24 for the exact prompt).

Different from the spurious feature removal task, here we apply a dual-threshold strategy. For vectors strongly related to unsafe categories (score  $\geq 4$ ), we set the singular value  $\sigma_i = 0$ , thus removing the feature entirely. For vectors with a moderate relation to unsafe categories (score = 3), we set the singular value  $\sigma_i = -1$ . This effectively inverts the vector’s contribution to the residual stream, thus pushing the representation away from the unsafe subspace.

### A.7. Improving Classification Performance

In Sec. 5.3 of the main paper, we showed that it is possible to use SITH to enhance the classification accuracy of CLIP ViT-L/14 by amplifying and suppressing specific singular vectors. Here, we provide additional implementation details regarding this experiment.

**Identifying Task-Relevant Concepts.** Given a downstream classification task with a set of  $M$  class labels  $\mathcal{Y} = \{y_1, y_2, \dots, y_M\}$ , we first aim to identify which concepts in our concept pool  $\Gamma$  are most relevant to the task. To do so, we decompose the embeddings of each class name ( $\mathcal{E}_T(y_m)$ ) into a set of constituent concepts using COMP. To ensure the decomposition yields fundamental semantic attributes rather than trivial matches, we use a filtered concept dictionary  $\Gamma' = \Gamma \setminus \mathcal{Y}$ , where the class names themselves are removed from the candidate pool.

Given the set of concepts extracted for each class  $y_m$ , we define the union of these sets as the global task concept pool  $\Gamma_{task}$ :

$$\Gamma_{task} = \bigcup_{m=1}^M \text{COMP}(\mathcal{E}_T(y_m); \mathcal{E}_T(\Gamma')). \quad (24)$$

This pool represents the collection of semantic attributes (e.g., colors, shapes, textures, habitats) that are relevant to the classification task.

**Scoring Singular Vectors Relevance.** Given a right singular vector  $\mathbf{v}_i$ , COMP decomposes it into a set of  $K$  pairs of coefficients and concepts  $\{(w_{i,k}, \gamma_{i,k})\}_{k=1}^K$ , where  $w_{i,k}$  is the importance weight and  $\gamma_k \in \Gamma$  is the corresponding concept.

To quantify the relevance of the singular vector  $\mathbf{v}_i$  to the classification task, we compute the weighted similarity

between its constituent concepts and the task concept pool  $\Gamma_{task}$ . Specifically, for each concept  $\gamma_{i,k}$  in the vector’s explanation, we find its maximum cosine similarity with any concept in the task pool  $\Gamma_{task}$ . We then weight this similarity by the corresponding coefficient  $w_{i,k}$  and sum over all  $K$  concepts to obtain the relevance score:

$$\mathbf{R}(\mathbf{v}_i) = \sum_{k=1}^K w_{i,k} \max_{\gamma_j \in \Gamma_{task}} \langle \mathcal{E}_T(\gamma_{i,k}), \mathcal{E}_T(\gamma_j) \rangle, \quad (25)$$

where  $\langle \cdot, \cdot \rangle$  indicates the cosine similarity between two embeddings. This formulation ensures that a singular vector is considered relevant if its constituent concepts are semantically close to any concept required by the downstream task.

**Editing Singular Values.** To convert the relevance score  $\mathbf{R}(\mathbf{v}_i)$  into a scaling factor  $\alpha_i$ , we introduce a base threshold  $\tau$ . The purpose of  $\tau$  is to shift the distribution of scores such that only highly relevant vectors are amplified ( $> 1.0$ ) while irrelevant ones are suppressed ( $< 1.0$ ).

To prevent the complete elimination of any singular vector, we also apply a clamping operation to ensure the scaling factor never drops below a minimum value of 0.8. The final scaling factor  $\alpha_i$  is computed as:

$$\alpha_i = \max(0.8, \mathbf{R}(\mathbf{v}_i) + \tau). \quad (26)$$

Finally, the original singular value  $\sigma_i$  associated with the singular vector  $\mathbf{v}_i$  is updated as follows:

$$\sigma'_i = \alpha_i \cdot \sigma_i. \quad (27)$$

This effectively performs a “soft” feature selection: vectors encoding semantic concepts unrelated to the task are dampened, while vectors aligned with the task’s semantics are preserved or amplified.

## A.8. Model Adaptation

In this section, we provide additional implementation details regarding the fine-tuning analysis described in Sec. 6 of the main paper, including training hyperparameters, the mathematical definition of the similarity metric used, and the LLM evaluation protocol.

**Training Details.** In our analysis of model adaptation, we examine how the value-output weight matrices evolve during fine-tuning. While SITH analyzes the collapsed  $\mathbf{W}_{VO}$  matrices, standard fine-tuning updates the parameter matrices  $\mathbf{W}_V$  and  $\mathbf{W}_O$  separately. Consistent with this, we fine-tune the pretrained value  $\mathbf{W}_V^{pre}$  and output  $\mathbf{W}_O^{pre}$  matrices of the last four layers of the OpenCLIP ViT-L/14 vision encoder, resulting in  $\mathbf{W}_V^{ft}$  and  $\mathbf{W}_O^{ft}$ , respectively. Then, for a given attention head, we construct the post-adaptation VO matrix as  $\mathbf{W}_{VO}^{ft} = \mathbf{W}_V^{ft} \mathbf{W}_O^{ft}$ , and compare it to the pre-trained VO matrix  $\mathbf{W}_{VO}^{pre} = \mathbf{W}_V^{pre} \mathbf{W}_O^{pre}$ .

We perform fine-tuning on three fine-grained classification datasets: Flowers 102 [45], Oxford Pets [54], and

CUB-200 [75]. For each dataset, we fine-tune the model for 10 epochs using a batch size of 64 and a learning rate of  $1 \times 10^{-4}$ . For LoRA fine-tuning, we set the rank equal to 8 and the  $\alpha$  to 16.

**Normalized Spectral Cosine Similarity.** To quantify the geometric shift in the semantic basis of the attention heads, we utilize the normalized spectral cosine similarity. This metric, adapted from Basile et al. [1], measures the alignment between two sets of vectors in a way that is weighted by their importance.

To compute this metric, we iteratively match the singular vectors from the pre-trained and fine-tuned VO matrices based on their weighted cosine similarity, ensuring that each vector is only matched once. Formally, let  $\mathbb{S}_{pre} = \{(\mathbf{v}_i^{pre}, \sigma_i^{pre})\}_{i=1}^r$  and  $\mathbb{S}_{ft} = \{(\mathbf{v}_j^{ft}, \sigma_j^{ft})\}_{j=1}^r$  be the sets of right singular vectors and their associated singular values for a pre-trained and fine-tuned head, respectively. Furthermore, let  $\mathcal{I}_n$  and  $\mathcal{J}_n$  be the sets of indices of the singular vectors that have already been matched in the previous  $n$  iterations (so that  $\mathcal{I}_0 = \mathcal{J}_0 = \emptyset$ ). Then, we define the spectral cosine similarity for the  $n$ -th matched pair as follows:

$$s_n = \left[ \max_{i \notin \mathcal{I}_{n-1}, j \notin \mathcal{J}_{n-1}} |\langle \mathbf{v}_i^{pre}, \mathbf{v}_j^{ft} \rangle| \right] \sigma_i^{pre} \sigma_j^{ft}, \quad (28)$$

where we consider the absolute value of the cosine similarity as singular vectors are defined up to the sign (*i.e.*,  $\mathbf{v}$  and  $-\mathbf{v}$  are both valid singular vectors). The final Normalized Spectral Cosine Similarity is then computed as:

$$\text{Sim}(\mathbb{S}_{pre}, \mathbb{S}_{ft}) = \sqrt{\frac{\sum_{n=1}^r s_n^2}{\sum_{n=1}^r (\sigma_n^{pre} \sigma_n^{ft})^2}}. \quad (29)$$

This metric ranges from 0 to 1, where 1 indicates that the two sets are perfectly aligned.

**LLM Evaluation Protocol.** To assess the semantic alignment of the task singular vectors with the fine-tuning domain, we employ an LLM-based evaluation protocol. Specifically, for each task singular vector we use COMP (with  $\lambda = 0.3$  and sparsity budget  $K = 5$ ) to generate textual explanations that describe the concepts encoded by the vector. We then prompt GPT-5-mini [51] to classify whether the concepts in each explanation are relevant to the fine-tuning domain on a binary scale (Yes/No). The prompts used for this evaluation are reported in Tabs. 25 to 27. Finally, we compute the percentage of task singular vectors that are classified as relevant to the fine-tuning task for each dataset and adaptation method (see Fig. 8 of the main paper).

## B. Ablating the Concept Pool

In this section, we evaluate SITH against multiple concept dictionaries commonly used in the interpretability literature:

Table 6. **Comparison of different concept pools** along four critical axes: scale (*i.e.*, number of concepts), granularity, safety alignment, and language coverage. ConceptNet 5.5 outperforms the alternatives across all dimensions, making it the most suitable choice for our interpretability framework.

Concept Pool	Scale	Granularity	Safety Alignment	Language Coverage
TextSpan [23]	3498	Low	High	English-only
SpLiCE [2]	15K	Medium	High	English-only
WordNet [41]	153K	High	Medium	English-only
ConceptNet 5.5 [67]	1.35M	High	Low	Multilingual

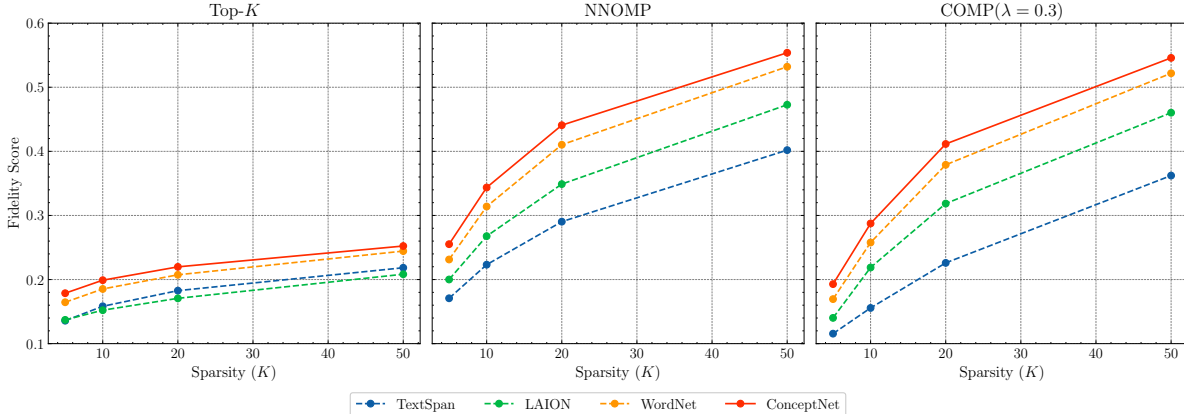


Figure 9. **Fidelity Score for different Concept Pools.** For each reconstruction method, we report the cosine-similarity (averaged across the last four layers of ViT-L-14) between the original singular vectors and their reconstructed versions at different sparsity levels and for different concept pools. We can see that for all methods and sparsity levels, ConceptNet guarantees superior reconstruction capabilities.

TextSpan [23], a highly curated set of 3,498 image descriptions generated by ChatGPT, SpLiCE [2], a frequency-based pool derived from LAION-400m [63], containing the top-10k single-word and top-5k two-word concepts, and WordNet [41], a large lexical database of English.

**Dictionary Comparison.** We evaluate these dictionaries based on four critical axes: scale (total number of concepts), granularity (ability to capture nuances), safety alignment (presence of NSFW concepts), and language coverage. A summary of this comparison is provided in Tab. 6.

Existing dictionaries often fall short in one or more of these aspects. For instance, TextSpan, while highly curated, is limited in scale and thus struggles to cover the vast semantic space of CLIP. Furthermore, because it consists of short image descriptions, it tends to capture broad, scene-level summaries rather than the specific, fine-grained attributes that are often encoded by individual singular vectors.

Similarly, the SpLiCE pool presents challenges regarding granularity. To reduce redundancy, this pool aggressively removes concepts with high cosine similarity ( $> 0.9$ ). While this ensures diversity, it inadvertently eliminates semantic nuances, such as the distinction between “cherry red” and “scarlet red”, which can be crucial for accurately interpreting singular vectors. Consequently, our sparse decomposition

method (COMP) would be forced to select a more generic concept or combine multiple less relevant concepts, leading to less precise and interpretable explanations.

Another significant limitation shared by both the TextSpan and SpLiCE pools is their explicit filtering of unsafe content. Although removing NSFW terms is standard practice for generative applications, it is a severe limitation for mechanistic interpretability. As demonstrated in our experiments on NSFW removal (see Sec. 5.2), CLIP natively encodes concepts related to nudity and violence within specific singular vectors. To successfully identify and suppress these concepts, the concept pool must first contain them.

ConceptNet 5.5 addresses these limitations comprehensively. It offers a massive scale (more than 1.3 million concepts) that captures the long tail of semantic concepts, including synonyms and variations that allow for high-fidelity sparse approximations. Crucially, it retains NSFW concepts, which enables the safety interventions proposed in our main paper. Finally, unlike TextSpan, SpLiCE, and WordNet, which are strictly English-only, ConceptNet is multilingual.

**Quantitative Ablation.** To verify that ConceptNet 5.5 better captures the semantic content of CLIP’s weights, we evaluate the reconstruction fidelity of the singular vectors across the last four layers of the ViT-L/14 model using each of the




Figure 10. **Interpretability vs. Fidelity trade-off across Layers 20, 21, and 22 of CLIP ViT-L/14.** For each tested layer, we compare the performance of our proposed COMP method against two baselines: Top- $k$  selection and NNOMP [55]. Each point represents a different sparsity level ( $K \in \{5, 10, 20, 50\}$ ). Across all layers, COMP consistently achieves a superior balance between interpretability and fidelity compared to the baselines.

four concept pools. To ensure that our findings are not specific to our sparse decomposition method (COMP), we also evaluate two alternative decomposition techniques: Non-Negative Orthogonal Matching Pursuit (NNOMP) and Top- $k$  selection. We also test varying the number of concepts selected ( $k \in \{5, 10, 20, 50\}$ ) to assess the robustness of each dictionary across different sparsity levels.

As illustrated in Fig. 9, ConceptNet 5.5 consistently outperforms the other concept pools across all reconstruction methods and sparsity levels. This confirms that the larger, more diverse search space of ConceptNet allows SITH to find semantic combinations that more accurately approximate the singular vectors of CLIP.

## C. Extended Quantitative Analysis for CLIP ViT-L/14

In this section, we extend the quantitative evaluation of SITH presented in the main paper. We first analyze the interpretability-fidelity trade-off across earlier layers of the model (Sec. C.1) and then demonstrate the robustness of our approach by applying it to the left singular vectors (Sec. C.2).

### C.1. Interpretability-Fidelity Analysis on Additional Layers

In Sec. 4.1 of the main paper, we present the interpretability-fidelity trade-off for the last layer of CLIP ViT-L/14 ( $l = 23$ ). Here, we extend this analysis to the preceding layers  $l \in \{20, 21, 22\}$  to verify the consistency of our findings across different depths of the network.

**Robustness of COMP.** As observed in the main paper (see Fig. 3), the results for layers 20, 21, and 22 (see Fig. 10) confirm that COMP consistently identifies the most favorable trade-off between reconstruction fidelity and semantic interpretability. While the baseline Top- $k$  approach yields

high interpretability but poor fidelity, and NNOMP achieves high fidelity but produces polysemantic (and thus less interpretable) explanations, COMP successfully bridges this gap across all analyzed layers.

**Layer-wise Trends.** Beyond the relative performance of the methods, comparing the plots across layers reveals a clear trend: both fidelity and interpretability scores progressively improve as we move towards the last layer. We hypothesize that this phenomenon is driven by two primary factors: the semantic abstraction level of the features and the geometric alignment with the output space.

- **Semantic Abstraction.** It is well-established in deep learning literature that shallower layers tend to encode lower-level features, while deeper layers capture higher-level, more abstract representations [16, 47, 81]. However, ConceptNet 5.5 [67] predominantly consists of high-level semantic concepts. Consequently, reconstructing the lower-level singular vectors of earlier layers using a dictionary of high-level concepts is inherently more difficult, leading to lower fidelity scores.
- **Geometric Alignment.** SITH relies on the model’s final projection matrix  $\mathbf{W}_p$  to map singular vectors from the residual stream to the multimodal space where the decomposition is then performed. However, in the standard forward pass of CLIP,  $\mathbf{W}_p$  operates on the residual stream of the final layer  $L$ . While the singular vectors of layer  $l$  reside in the residual stream space, the residual stream evolves as it passes through subsequent layers. Therefore, applying  $\mathbf{W}_p$  to the weights of earlier layers introduces an approximation error due to the possible misalignment between the residual stream at layer  $l$  and layer  $L$ . Furthermore, to evaluate the fidelity of a decomposition, we need to project the reconstruction from the multimodal space back into the residual stream; for shallower layers, this inversion likely incurs a greater approximation error,

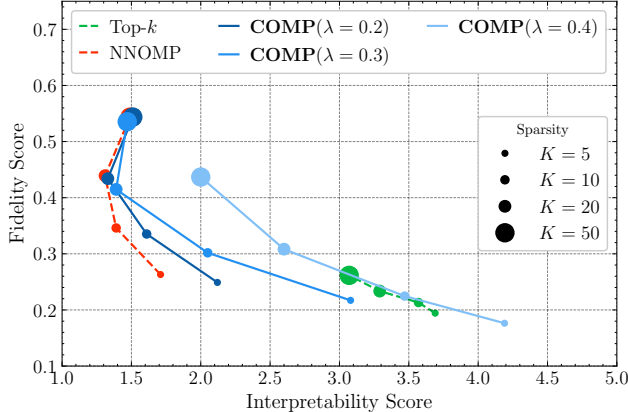


Figure 11. **Interpretability vs. Fidelity trade-off for the left singular vectors of the last layer of CLIP ViT-L/14.** We compare the performance of our proposed COMP against two baselines: Top- $k$  selection and NNOMP [55]. Each point represents a different sparsity level ( $K \in \{5, 10, 20, 50\}$ ).

further degrading the fidelity score.

Despite these challenges, we note that SITH with COMP maintains a superior Pareto frontier compared to baselines even in these earlier layers, demonstrating the method’s robustness.

## C.2. Analysis of Left Singular Vectors

While the main text focuses on the right singular vectors  $V$ , which define the directions the attention heads write to in the residual stream, the left singular vectors  $U$  play an equally critical role by defining the directions the heads read from the input. To demonstrate the generality of our approach, we replicate the interpretability-fidelity analysis from Sec. 4.1 on the left singular vectors of the last layer ( $l = 23$ ) of CLIP ViT-L/14.

**Effectiveness of COMP.** Consistent with our findings for the right singular vectors, the results shown in Fig. 11 confirm that COMP achieves the most favorable trade-off between fidelity and interpretability when applied to left singular vectors, being able to faithfully explain them. The baseline methods exhibit the same limitations observed previously: Top- $k$  provides coherent but low-fidelity explanations, while NNOMP yields high fidelity at the cost of semantic coherence. COMP effectively balances these objectives, demonstrating that our decomposition algorithm is robust regardless of whether it is applied to the input or output space of the attention mechanism.

**Comparison with Right Singular Vectors.** When comparing the absolute performance scores between the two sets of vectors, we observe that the left singular vectors consistently exhibit slightly lower reconstruction fidelity scores than their right singular vector counterparts. We attribute this discrepancy to the geometric alignment issue discussed earlier.

Indeed, right singular vectors represent updates being written forward into the stream, effectively aligning them closer to the final output representation. In contrast, left singular vectors represent features read from the incoming residual stream (i.e., the output of the preceding layer). Therefore, there is likely a greater misalignment between the left singular vectors and the final projection space. Consequently, projecting  $U$  into the semantic space incurs a higher approximation error, which slightly degrades the fidelity of the subsequent reconstruction.

## D. Qualitative Results

In this section, we provide a more extensive qualitative analysis of the interpretations produced by SITH. We focus on the top-5 singular vector pairs (i.e., the pairs  $(u_i, v_i)$  associated with the 5 largest singular values  $\sigma_i$ ) for various attention heads. Indeed, the Singular Value Decomposition  $W_{VO} = U\Sigma V^T$  allows us to express the Value-Output matrix as a sum of rank-1 matrices:  $W_{VO} = \sum_{i=1}^r \sigma_i u_i v_i^T$ . According to the Eckart-Young theorem [19], the sum of the first  $k$  terms of this expansion provides the best rank- $k$  approximation of the original matrix in terms of the Frobenius norm. Therefore, the singular vectors associated with the largest singular values encode the most dominant reading  $(u_i)$  and writing  $(v_i)$  directions of the attention head, effectively defining its primary functional roles.

For each analyzed singular vector, we show: (1) the sparse concept set returned by COMP with  $\lambda = 0.3$  and  $K = 5$ , (2) and the top-4 images from the CC12M dataset whose [CLS] token at layer  $l$  (i.e., the layer of the analyzed attention head) has the highest cosine similarity with the singular vector. We report the results of the analyzed heads in Tabs. 7 to 12.

**Left-Right Semantic Alignment.** We find that for many attention heads, the top singular vector pairs exhibit a strong semantic alignment between their reading and writing directions. For instance, in Tab. 11, all top-5 singular vector pairs of Head 8 in Layer 23 of ViT-L/14 are dedicated to colors, with each pair corresponding to a pair of colors, such as “orange” (reading) to “purple” (writing) and “yellow” (reading) to “blue” (writing). Similar patterns are also observed in other heads, such as Head 2 in Layer 22 (Tab. 7) and Head 4 in Layer 23 (Tab. 10).

**Intra-Head Semantic Alignment.** We observe that within many attention heads, the dominant directions are often semantically correlated, effectively grouped under a broader “theme”. For instance, the Head 0 of Layer 23 in ViT-L/14 (see Tab. 9) is focused on materials, with each singular vector encoding a specific material type such as “steel”, “paper”, and “glass”. Similarly, Head 11 of the same layer (see Tab. 12) captures letters, with singular vectors representing different characters like “C”, “M” and “S”. This head is particularly interesting as it also shows how CLIP is effectively able to read text in images. This intra-head semantic coher-

ence suggests that certain attention heads are specialized in processing specific categories of information.

**Comparison with TextSpan.** Our findings strongly align with the head-level classifications provided by the activation-based method TextSpan [23], as the function of many attention heads identified by TextSpan matches the concepts assigned to the top singular vectors. However, SITH offers a clear advantage in granularity: where TextSpan might broadly label a head as encoding certain colors, SITH decomposes this behavior, identifying exactly which singular vector is responsible for “red”, which for “green”, etc. Furthermore, because SITH is data-free, it identifies these functionalities solely from weights, avoiding the potential bias where a head might be mislabeled simply because the probing dataset lacks specific concept classes.

**Generalization across architectures, scales, and training regimes.** To demonstrate that the findings of SITH are not limited to a specific model or training paradigm, we extend our qualitative analysis to a broader suite of vision-language models. Specifically, we evaluate OpenCLIP ViT-B/32 and ViT-H/14 to assess generalizability across different network capacities. Furthermore, to investigate the impact of architectural and data variations, we apply SITH to MobileCLIP ViT-L/14 [73], which builds upon the FastViT architecture [72] utilizing a highly optimized training regime. The results, presented in Tabs. 13 to 20, show that the same semantic patterns identified in ViT-L/14 are consistently found across these diverse models. For instance, the color-related head identified in ViT-L/14 Layer 23 Head 8 (see Tab. 11) is also present in ViT-B/32 Layer 11 Head 2 (see Tab. 13) and ViT-H/14 Layer 31 Head 13 (see Tab. 17), while the location-related head of ViT-L/14 Layer 23 Head 2 (see Tab. 7) is also found in ViT-H/14 Layer 31 Head 12 (see Tab. 16) and MobileCLIP ViT-L/14 Layer 22 Head 10 (see Tab. 19).

These findings further reinforce the *Universality Hypothesis* in mechanistic interpretability, which posits that different neural networks converge on similar features and circuits when trained on similar data distributions [49]. This has been extensively explored in activation space, showing that distinct models converge toward shared representational spaces or “Platonic” concepts [18, 31, 71]. In contrast, our data-free weight-space analysis via SITH reveals universality at the level of functional components (i.e., attention heads). Rather than just learning the same latent concepts in their activation spaces, models of vastly different capacities, architectures, and training paradigms allocate attention heads to perform identical, specialized semantic operations (e.g., colors, materials, or locations). This aligns with findings in the Large Language Model (LLM) literature, where specific attention head mechanisms, such as “induction heads” for in-context learning [50] or “successor heads” for ordinal sequences [28], have been shown to universally emerge across diverse models. Our results extend this universality to

vision-language models, demonstrating that the emergence of functionally specialized attention heads is a fundamental property of these models, transcending specific design choices, model scales, and training methodologies.

## E. Pseudocode of COMP

In this section, we provide the pseudocode for the Coherent Orthogonal Matching Pursuit (COMP) algorithm. As discussed in Sec. 3.3, COMP extends the traditional Non-Negative Orthogonal Matching Pursuit (NNOMP) by incorporating a coherence term into the concept selection process. As shown in the section highlighted in yellow within Algorithm 1, during each iteration of the concept selection step, we compute a coherence score for each candidate concept based on its average similarity to the concepts already selected in the support set. This coherence score is then combined with the standard correlation score to form a final score used for selecting the next concept to include in the support set. This modification encourages the selection of concepts that are not only highly correlated with the current residual but also semantically coherent with the concepts already selected in the support set, so as to enhance the interpretability of the resulting sparse representation.

## F. GPT-5 Prompts

To ensure the reproducibility of our results, we provide the exact prompt templates used across our experiments. As detailed in the main text, we utilized GPT-5-mini [51] for all LLM-based evaluation and editing tasks:

- Tab. 21 contains the prompt used to evaluate the semantic coherence of the concept sets extracted by COMP (as well as the baselines) in Sec. 4, using a 5-point Likert scale. This corresponds to the results in Sec. 4.1.
- Tab. 22 contains the prompt used to rate the alignment between the top-retrieved images for a specific singular vector and its textual interpretation. This corresponds to the results in Sec. 4.2.
- Tabs. 23 and 24 contain the prompts used to identify and suppress spurious correlations (Sec. 5.1) and to remove NSFW concepts (Sec. 5.2), respectively.
- Tabs. 25 to 27 contain the prompts used to evaluate the alignment of task singular vectors to the fine-tuning domains in Sec. 6 for Flowers102, Oxford-IIIT Pet, and CUB-200, respectively.

Table 7. **Layer 22, Head 2 of ViT-L/14 encodes locations.** The first 5 pairs of left/right singular vectors from Layer 22, Head 2 of ViT-L/14. For each singular vector, we display the top-4 images from CC12M [10] most similar to it, along with the explanation generated by COMP with  $\lambda = 0.3$  and  $K = 5$ . **Unsafe text/images have been redacted/blurred.**





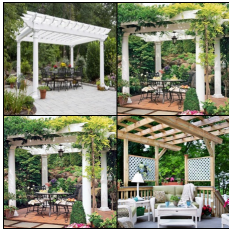


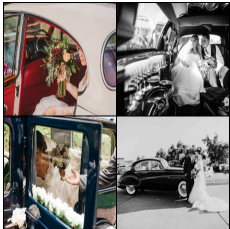


Left: house rooms	1st Singular Vector	Right: outdoor locations
	<ul style="list-style-type: none"> <li>• gogglebox (0.1730)</li> <li>• cleaning living room (0.1713)</li> <li>• hotel suite (0.1609)</li> <li>• kitchen bedroom bathroom (0.1470)</li> <li>• room of house (0.1068)</li> </ul>	
Left: street	2nd Singular Vector	Right: event rooms
	<ul style="list-style-type: none"> <li>• roadside trees (0.1622)</li> <li>• street price (0.1557)</li> <li>• getting outdoors (0.1546)</li> <li>• product for car owners (0.1370)</li> <li>• street mobile (0.0820)</li> </ul>	
Left: backyard	3rd Singular Vector	Right: shops
	<ul style="list-style-type: none"> <li>• balcony bra (0.2164)</li> <li>• accommodation set aside for guests (0.2002)</li> <li>• outdoor home (0.1697)</li> <li>• lawn parties (0.1653)</li> <li>• relaxing on porch (0.1228)</li> </ul>	
Left: home areas	4th Singular Vector	Right: road, automobile
	<ul style="list-style-type: none"> <li>• play in back yard (0.1716)</li> <li>• basement style foundation (0.1527)</li> <li>• garden office (0.1522)</li> <li>• work productive watch (0.1393)</li> <li>• home office (0.1022)</li> </ul>	
Right: celebration events	5th Singular Vector	Right: traveling areas
	<ul style="list-style-type: none"> <li>• wedding food (0.1940)</li> <li>• sales event (0.1519)</li> <li>• car festooned for celebration (0.1513)</li> <li>• birthday banquet (0.1222)</li> <li>• outdoor party (0.1138)</li> </ul>	

Table 8. **Layer 22, Head 3 of ViT-L/14 encodes objects/body parts.** The first 5 pairs of left/right singular vectors from Layer 22, Head 3 of ViT-L/14. For each singular vector, we display the top-4 images from CC12M [10] most similar to it, along with the explanation generated by COMP with  $\lambda = 0.3$  and  $K = 5$ . **Unsafe text/images have been redacted/blurred.**











Left: objects for the upper torso	1st Singular Vector	Right: fabric items
	<ul style="list-style-type: none"> <li>• pullover sweater (0.1905)</li> <li>• necklacings (0.1726)</li> <li>• museum bust (0.1720)</li> <li>• water on chest (0.1484)</li> <li>• on chest (0.1204)</li> </ul>	<ul style="list-style-type: none"> <li>• frieze pants (0.1977)</li> <li>• pair of short pantlegs (0.1654)</li> <li>• curtains (0.1580)</li> <li>• classic pants (0.1186)</li> <li>• summer trousers (0.1148)</li> </ul> 
Left: cards	2nd Singular Vector	Right: shirt, table
	<ul style="list-style-type: none"> <li>• long jacket (0.1569)</li> <li>• invite job applications (0.1460)</li> <li>• card case (0.1157)</li> <li>• invitation card (0.0733)</li> <li>• pc card (0.0571)</li> </ul>	<ul style="list-style-type: none"> <li>• shirt arm (0.3040)</li> <li>• nest table (0.1871)</li> <li>• tablebases (0.1627)</li> <li>• tabletopped (0.1433)</li> <li>• table (0.0630)</li> </ul> 
Left: head parts	3rd Singular Vector	Right: dress
	<ul style="list-style-type: none"> <li>• open fence (0.1835)</li> <li>• wig head (0.1754)</li> <li>• hearing (0.1661)</li> <li>• earsies (0.1175)</li> <li>• human ears (0.0388)</li> </ul>	<ul style="list-style-type: none"> <li>• iliotibial band (0.2007)</li> <li>• dress and skirts (0.1926)</li> <li>• shirtdress (0.1881)</li> <li>• drawer under telephone (0.1648)</li> <li>• under dress (0.1175)</li> </ul> 
Left: dorsal region	4th Singular Vector	Right: body parts
	<ul style="list-style-type: none"> <li>• backpack (0.1874)</li> <li>• hairpin for bun (0.1719)</li> <li>• shoulder blade (0.1514)</li> <li>• back tee (0.1467)</li> <li>• pin back hair (0.1119)</li> </ul>	<ul style="list-style-type: none"> <li>• facial [REDACTED] (0.1993)</li> <li>• inregisters (0.1785)</li> <li>• polydactylies (0.1489)</li> <li>• feet touch cold floor (0.1425)</li> <li>• foot [REDACTED] (0.0561)</li> </ul> 
Left: ceiling	5th Singular Vector	Right: miscellaneous objects
	<ul style="list-style-type: none"> <li>• ceiling floor (0.1770)</li> <li>• outdoor ceiling (0.1170)</li> <li>• ceiling under roof (0.0685)</li> <li>• something ceiling (0.0683)</li> <li>• ceiling (0.0138)</li> </ul>	<ul style="list-style-type: none"> <li>• controlling wrist (0.1977)</li> <li>• short jacket (0.1939)</li> <li>• cool tankards (0.1829)</li> <li>• small mug (0.1414)</li> <li>• wrist timepiece (0.0860)</li> </ul> 

Table 9. **Layer 23, Head 0 of ViT-L/14 encodes materials.** The first 5 pairs of left/right singular vectors from Layer 23, Head 0 of ViT-L/14. For each singular vector, we display the top-4 images from CC12M [10] most similar to it, along with the explanation generated by COMP with  $\lambda = 0.3$  and  $K = 5$ .






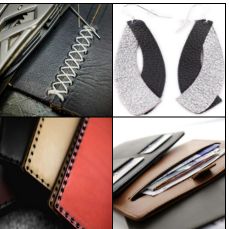




Left: clothing	1st Singular Vector	Right: wax et similia
 <ul style="list-style-type: none"> <li>• interknits (0.1811)</li> <li>• comfortable clothes (0.1682)</li> <li>• embroidered into cloth with sewing (0.1634)</li> <li>• made from cloth (0.1520)</li> <li>• cloth car seat (0.1429)</li> </ul>	<ul style="list-style-type: none"> <li>• leathery skin (0.2421)</li> <li>• encaustics (0.1890)</li> <li>• persulphates (0.1549)</li> <li>• gloss stick (0.1489)</li> <li>• rubber and latex (0.1392)</li> </ul>	
Left: metal	2nd Singular Vector	Right: leather
 <ul style="list-style-type: none"> <li>• steel aluminium (0.1945)</li> <li>• played with small metal rod (0.1439)</li> <li>• metal helmet (0.1322)</li> <li>• metal jewelry (0.1264)</li> <li>• piece of metal furniture (0.1173)</li> </ul>	<ul style="list-style-type: none"> <li>• embroidered into cloth with sewing (0.1977)</li> <li>• suede leather (0.1530)</li> <li>• goatskins (0.1333)</li> <li>• fruit leather (0.1299)</li> <li>• soft leather (0.1166)</li> </ul>	
Left: food	3rd Singular Vector	Right: leather
 <ul style="list-style-type: none"> <li>• food color (0.1818)</li> <li>• look at paintings of food (0.1671)</li> <li>• drinking yogurt (0.1571)</li> <li>• packaged breakfast food (0.1420)</li> <li>• peanut pastes (0.1356)</li> </ul>	<ul style="list-style-type: none"> <li>• leather flower (0.1400)</li> <li>• leather trades (0.1352)</li> <li>• leatherwork (0.1072)</li> <li>• upholstered with leather (0.1056)</li> <li>• leather case (0.1003)</li> </ul>	
Left: plastic	4th Singular Vector	Right: steel + drink
 <ul style="list-style-type: none"> <li>• knittabilities (0.2002)</li> <li>• wood and plastic (0.1860)</li> <li>• credit plastic (0.1829)</li> <li>• plastic art (0.1398)</li> <li>• plastic furniture (0.1332)</li> </ul>	<ul style="list-style-type: none"> <li>• comforting drink (0.2469)</li> <li>• paramount titles (0.1969)</li> <li>• steel wines (0.1965)</li> <li>• stainless iron (0.1802)</li> </ul>	
Left: glass	5th Singular Vector	Right: paper
 <ul style="list-style-type: none"> <li>• glass making (0.2246)</li> <li>• nanotherapeutics (0.1837)</li> <li>• vodka luges (0.1649)</li> <li>• mold on liquids (0.1409)</li> <li>• glass ingredient (0.1257)</li> </ul>	<ul style="list-style-type: none"> <li>• uncensorship (0.1668)</li> <li>• putting images on paper (0.1632)</li> <li>• posting children's art work on (0.1536)</li> <li>• paper tickets (0.1452)</li> <li>• literary journalism (0.1344)</li> </ul>	

Table 10. **Layer 23, Head 4 of ViT-L/14 encodes *people*.** The first 5 pairs of left/right singular vectors from Layer 23, Head 4 of ViT-L/14. For each singular vector, we display the top-4 images from CC12M [10] most similar to it, along with the explanation generated by COMP with  $\lambda = 0.3$  and  $K = 5$ . **Unsafe text/images have been redacted/blurred.**



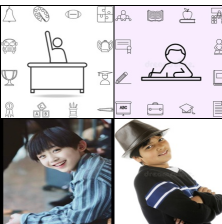


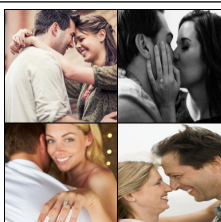
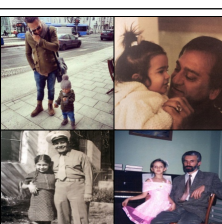
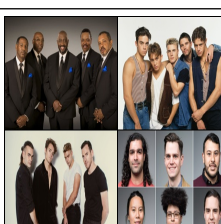


<p><b>Left: women</b></p> 	<p><b>1st Singular Vector</b></p> <ul style="list-style-type: none"> <li>• group of groups of women (0.2365)</li> <li>• saint agnes eves (0.2199)</li> <li>• women's writing (0.1835)</li> <li>• girlswear (0.1670)</li> <li>• business girl (0.1587)</li> </ul>	<p><b>Right: men</b></p>  <ul style="list-style-type: none"> <li>• father and son (0.1844)</li> <li>• men's aesthetic (0.1623)</li> <li>• gay male s█ (0.1530)</li> <li>• groomsmen (0.1394)</li> <li>• sexy guys (0.0826)</li> </ul>
<p><b>Left: boys</b></p> 	<p><b>2nd Singular Vector</b></p> <ul style="list-style-type: none"> <li>• young boy (0.1956)</li> <li>• solo comedian (0.1526)</li> <li>• child in camp (0.1440)</li> <li>• juvenile vagrant (0.1348)</li> <li>• transit worker (0.1348)</li> </ul>	<p><b>Right: couples</b></p>  <ul style="list-style-type: none"> <li>• unhappy couples (0.2710)</li> <li>• tag team (0.2025)</li> <li>• co founders (0.1860)</li> <li>• tandem bicycles (0.1748)</li> <li>• couples together (0.0864)</li> </ul>
<p><b>Left: kids</b></p> 	<p><b>3rd Singular Vector</b></p> <ul style="list-style-type: none"> <li>• junior bridesmaids (0.2663)</li> <li>• generic kids (0.1880)</li> <li>• children training (0.1834)</li> <li>• kids who brothers (0.1712)</li> <li>• kids together (0.0932)</li> </ul>	<p><b>Right: spouse</b></p>  <ul style="list-style-type: none"> <li>• man health worker (0.1322)</li> <li>• clumsy spouse (0.1137)</li> <li>• heterosexual woman in love (0.1052)</li> <li>• testing routine on spouse (0.1015)</li> <li>• spouse scares (0.0781)</li> </ul>
<p><b>Left: two (people)</b></p> 	<p><b>4th Singular Vector</b></p> <ul style="list-style-type: none"> <li>• father and daughter (0.3078)</li> <li>• two children (0.1726)</li> <li>• both periodicals (0.1435)</li> <li>• both cities in california (0.1289)</li> <li>• two pennies worth (0.1212)</li> </ul>	<p><b>Right: group</b></p>  <ul style="list-style-type: none"> <li>• vocal quintet (0.1821)</li> <li>• group of homosexuals (0.1609)</li> <li>• group complaining (0.1424)</li> <li>• group of men (0.0983)</li> <li>• five man group (0.0679)</li> </ul>
<p><b>Left: two people</b></p> 	<p><b>5th Singular Vector</b></p> <ul style="list-style-type: none"> <li>• mother son (0.2272)</li> <li>• friction between moms and sons (0.1833)</li> <li>• k█i (0.1807)</li> <li>• businesswomen (0.1704)</li> <li>• trans lesbians (0.1539)</li> </ul>	<p><b>Right: girl</b></p>  <ul style="list-style-type: none"> <li>• teen lolita (0.1560)</li> <li>• dads girl (0.1514)</li> <li>• little girl's hair (0.1105)</li> <li>• daughter dad (0.0810)</li> </ul>

Table 11. **Layer 23, Head 8 of ViT-L/14 encodes colors.** The first 5 pairs of left/right singular vectors from Layer 23, Head 8 of ViT-L/14. For each singular vector, we display the top-4 images from CC12M [10] most similar to it, along with the explanation generated by COMP with  $\lambda = 0.3$  and  $K = 5$ . **Unsafe text/images have been redacted/blurred.**




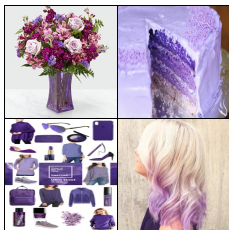
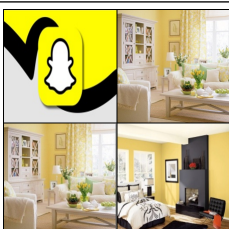
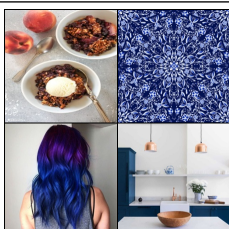




<p><b>Left: yellowish</b></p> 	<p><b>1st Singular Vector</b></p> <ul style="list-style-type: none"> <li>• ground zeroes (0.2402)</li> <li>• yellow green color (0.1565)</li> <li>• blue yellow (0.1186)</li> <li>• bluish yellow (0.1040)</li> </ul>	<p><b>Right: red</b></p> 
<p><b>Left: orange</b></p> 	<p><b>2nd Singular Vector</b></p> <ul style="list-style-type: none"> <li>• orange green (0.2026)</li> <li>• cyanodiethylgold (0.1773)</li> <li>• orange blossoms (0.1121)</li> <li>• peach color (0.1075)</li> <li>• orange mint (0.0234)</li> </ul>	<p><b>Right: purple</b></p> 
<p><b>Left: yellow</b></p> 	<p><b>3rd Singular Vector</b></p> <ul style="list-style-type: none"> <li>• yellowred (0.2565)</li> <li>• yellow press (0.1400)</li> <li>• yellow clothes (0.1292)</li> <li>• yellow fever (0.1128)</li> <li>• yellow hot (0.0585)</li> </ul>	<p><b>Right: blue</b></p> 
<p><b>Left: light green</b></p> 	<p><b>4th Singular Vector</b></p> <ul style="list-style-type: none"> <li>• white and green (0.2461)</li> <li>• bi asexual (0.1951)</li> <li>• mint green (0.1517)</li> <li>• variscite (0.1491)</li> <li>• mint cream (0.1123)</li> </ul>	<p><b>Right: orange</b></p> 
<p><b>Left: brown</b></p> 	<p><b>5th Singular Vector</b></p> <ul style="list-style-type: none"> <li>• copper brown (0.1957)</li> <li>• browns (0.1584)</li> <li>• brown packages (0.1409)</li> <li>• brown shower (0.1259)</li> <li>• brown notes (0.1227)</li> </ul>	<p><b>Right: yellow-blue</b></p> 

Table 12. **Layer 23, Head 11 of ViT-L/14 encodes letters.** The first 5 pairs of left/right singular vectors from Layer 23, Head 11 of ViT-L/14. For each singular vector, we display the top-4 images from CC12M [10] most similar to it, along with the explanation generated by COMP with  $\lambda = 0.3$  and  $K = 5$ . In this case, it is interesting to observe that CLIP is able to read the text present in the watermarks of the images, and the explanations correspond to letters present in those watermarks. **Unsafe text/images have been redacted/blurred.**











Left: the letter C	1st Singular Vector	Right: the letter M
	<ul style="list-style-type: none"> <li>• c t splice (0.1914)</li> <li>• cjd (0.1284)</li> <li>• csdl (0.0907)</li> <li>• cpsa (0.0829)</li> <li>• cpsu (0.0688)</li> </ul>	
Left: the letter S	2nd Singular Vector	Right: the letter P
	<ul style="list-style-type: none"> <li>• surname ma (0.1616)</li> <li>• msw (0.1225)</li> <li>• vsv marv (0.0733)</li> <li>• μsv (0.0496)</li> </ul>	
Left: the letter D	3rd Singular Vector	Right: the letters ISO
	<ul style="list-style-type: none"> <li>• dilli bags (0.1720)</li> <li>• ohle (0.1633)</li> <li>• dalk (0.1148)</li> <li>• dillies (0.0817)</li> <li>• dalks (0.0614)</li> </ul>	
Left: the letter L	4th Singular Vector	Right: the letter S
	<ul style="list-style-type: none"> <li>• linnaean taxonomy (0.2169)</li> <li>• lccns (0.1738)</li> <li>• uaw (0.1724)</li> <li>• limans (0.1192)</li> <li>• lccn (0.0164)</li> </ul>	
Left: the letter F	5th Singular Vector	Right: the letter D
	<ul style="list-style-type: none"> <li>• c t fart (0.1888)</li> <li>• filial child (0.1805)</li> <li>• fnma (0.1659)</li> <li>• feal (0.1173)</li> <li>• fulah (0.0979)</li> </ul>	

Table 13. **Layer 11, Head 2 of ViT-B/32 encodes colors.** The first 5 pairs of left/right singular vectors from Layer 11, Head 2 of ViT-B/32. For each singular vector, we display the top-4 images from CC12M [10] most similar to it, along with the explanation generated by COMP with  $\lambda = 0.3$  and  $K = 5$ .







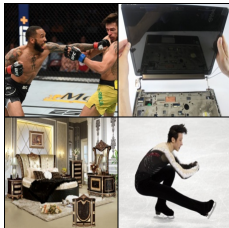

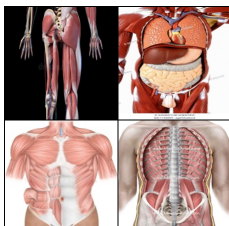

Left: shades of white	1st Singular Vector	Right: black
	<ul style="list-style-type: none"> <li>• wear white dresses (0.1973)</li> <li>• white yellow (0.1605)</li> <li>• beige boxes (0.1477)</li> <li>• cream top (0.1409)</li> <li>• cream colored (0.1324)</li> </ul>	
Left: gold	2nd Singular Vector	Right: gray
	<ul style="list-style-type: none"> <li>• reddish gold (0.2863)</li> <li>• brown boxes (0.1661)</li> <li>• golden yellow color (0.1634)</li> <li>• double gloucester (0.1504)</li> <li>• peanut butter (0.1290)</li> </ul>	
Left: brown gray	3rd Singular Vector	Right: white
	<ul style="list-style-type: none"> <li>• yellowish gray (0.2106)</li> <li>• taupe (0.1897)</li> <li>• brown boxes (0.1656)</li> <li>• bronze component (0.1339)</li> <li>• brown gray (0.0899)</li> </ul>	
Left: black	4th Singular Vector	Right: red and blue
	<ul style="list-style-type: none"> <li>• ivory black (0.2928)</li> <li>• saints merchandise (0.1523)</li> <li>• ups trucks (0.1417)</li> <li>• black and tan (0.1316)</li> <li>• black gold jewelery (0.1093)</li> </ul>	
Left: light red	5th Singular Vector	Right: green (+ blue and yellow)
	<ul style="list-style-type: none"> <li>• pink red (0.1953)</li> <li>• red silver (0.1718)</li> <li>• triple negative breast cancer (0.1641)</li> <li>• hokie (0.1557)</li> <li>• rosé wine (0.1221)</li> </ul>	

Table 14. **Layer 31, Head 7 of ViT-H/14 encodes people.** The first 5 pairs of left/right singular vectors from Layer 31, Head 7 of ViT-H/14. For each singular vector, we display the top-4 images from CC12M [10] most similar to it, along with the explanation generated by COMP with  $\lambda = 0.3$  and  $K = 5$ . **Unsafe text/images have been redacted/blurred.**








<p><b>Left: men</b></p> 	<p><b>1st Singular Vector</b></p> <ul style="list-style-type: none"> <li>• man p[REDACTED]y (0.1659)</li> <li>• father and son (0.1577)</li> <li>• groomsmen (0.1565)</li> <li>• men’s aesthetic (0.1427)</li> <li>• two men holding hands (0.1390)</li> </ul>	<p><b>Right: women</b></p>  <ul style="list-style-type: none"> <li>• urmila (0.2142)</li> <li>• evil woman (0.2048)</li> <li>• female political activists (0.1991)</li> <li>• scotswomen (0.1749)</li> <li>• women’s article (0.1460)</li> </ul>
<p><b>Left: married couple</b></p> 	<p><b>2nd Singular Vector</b></p> <ul style="list-style-type: none"> <li>• married people (0.1279)</li> <li>• relationship through marriage (0.1199)</li> <li>• newlyweds newly married couple (0.1195)</li> <li>• couples together (0.1058)</li> <li>• married couples (0.0999)</li> </ul>	<p><b>Right: boys</b></p>  <ul style="list-style-type: none"> <li>• boy monk (0.1720)</li> <li>• making teenaged boys act silly (0.1699)</li> <li>• she male (0.1647)</li> <li>• group boys (0.1644)</li> <li>• topmen (0.1396)</li> </ul>
<p><b>Left: bridesmaids</b></p> 	<p><b>3rd Singular Vector</b></p> <ul style="list-style-type: none"> <li>• junior bridesmaids (0.1675)</li> <li>• pretty dress (0.1286)</li> </ul>	<p><b>Right: men and women</b></p>  <ul style="list-style-type: none"> <li>• dad and mom (0.2929)</li> <li>• masculine gender (0.1985)</li> <li>• man who indulges women (0.1791)</li> <li>• man in womans clothes (0.1500)</li> <li>• man male and woman (0.1102)</li> </ul>
<p><b>Left: mother (and son)</b></p> 	<p><b>4th Singular Vector</b></p> <ul style="list-style-type: none"> <li>• mature adult (0.1937)</li> <li>• mother in laws (0.1303)</li> <li>• friction between moms and sons (0.1161)</li> <li>• mother’s son (0.1021)</li> <li>• mother son (0.0846)</li> </ul>	<p><b>Right: girls</b></p>  <ul style="list-style-type: none"> <li>• girl priest (0.1601)</li> <li>• girl guides (0.1509)</li> <li>• father daughter (0.1478)</li> <li>• little girl’s room (0.1389)</li> <li>• primary schoolgirl (0.0740)</li> </ul>
<p><b>Left: couple’s wedding</b></p> 	<p><b>5th Singular Vector</b></p> <ul style="list-style-type: none"> <li>• ringbearers (0.1753)</li> <li>• wedding boy (0.1617)</li> <li>• wedded (0.1593)</li> <li>• student marriage (0.1592)</li> <li>• teenage couple (0.1399)</li> </ul>	<p><b>Right: person in prominent position</b></p>  <ul style="list-style-type: none"> <li>• things to get done faster (0.1477)</li> <li>• retired public prosecutor (0.1372)</li> <li>• coach of dallas cowboys (0.1274)</li> <li>• named under secretary of defense for intelligence (0.0859)</li> </ul>

Table 15. **Layer 31, Head 11 of ViT-H/14 encodes letters.** The first 5 pairs of left/right singular vectors from Layer 31, Head 11 of ViT-H/14. For each singular vector, we display the top-4 images from CC12M [10] most similar to it, along with the explanation generated by COMP with  $\lambda = 0.3$  and  $K = 5$ . **Unsafe text/images have been redacted/blurred.**











<p><b>Left:</b> the letter M</p> 	<p><b>1st Singular Vector</b></p> <ul style="list-style-type: none"> <li>• mln (0.2066)</li> <li>• act reasonable using mind (0.1886)</li> <li>• neuromyoarterial (0.1710)</li> <li>• aaai (0.1417)</li> <li>• haaf net (0.1388)</li> </ul>	<p><b>Right:</b> the letter S</p> 
<p><b>Left:</b> the letters A, S</p> 	<p><b>2nd Singular Vector</b></p> <ul style="list-style-type: none"> <li>• genus alosa (0.1889)</li> <li>• ssaa affiliated (0.1730)</li> <li>• now now ism (0.1612)</li> <li>• nsaim (0.1330)</li> <li>• aas (0.1266)</li> </ul>	<p><b>Right:</b> the letters F, T, P</p> 
<p><b>Left:</b> the letters T, M</p> 	<p><b>3rd Singular Vector</b></p> <ul style="list-style-type: none"> <li>• syrt (0.1828)</li> <li>• t y f [redacted] g (0.1797)</li> <li>• jvm ti (0.1779)</li> <li>• wry mouth (0.1686)</li> <li>• ntim (0.1576)</li> </ul>	<p><b>Right:</b> the letters C, D, E</p> 
<p><b>Left:</b> the letter A</p> 	<p><b>4th Singular Vector</b></p> <ul style="list-style-type: none"> <li>• pfaas (0.2619)</li> <li>• keep as pet (0.2127)</li> <li>• air to air (0.1867)</li> <li>• aa trees (0.1755)</li> <li>• aas (0.1734)</li> </ul>	<p><b>Right:</b> the letter C</p> 
<p><b>Left:</b> the letter D</p> 	<p><b>5th Singular Vector</b></p> <ul style="list-style-type: none"> <li>• baby dog (0.2224)</li> <li>• d'you (0.1796)</li> <li>• fdi (0.1735)</li> <li>• fpd (0.1553)</li> <li>• fddi (0.1033)</li> </ul>	<p><b>Right:</b> the letter E</p> 

Table 16. **Layer 31, Head 12 of ViT-H/14 encodes locations.** The first 5 pairs of left/right singular vectors from Layer 31, Head 12 of ViT-H/14. For each singular vector, we display the top-4 images from CC12M [10] most similar to it, along with the explanation generated by COMP with  $\lambda = 0.3$  and  $K = 5$ . **Unsafe text/images have been redacted/blurred.**







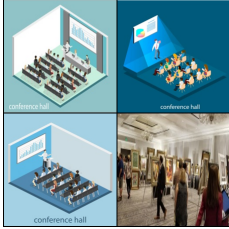



<p><b>Left: festivals</b></p> 	<p><b>1st Singular Vector</b></p> <ul style="list-style-type: none"> <li>• megafestivals (0.1719)</li> <li>• bedside book (0.1338)</li> <li>• fix piece of wood furniture (0.1328)</li> <li>• made from wood cotton or linen (0.1222)</li> <li>• handy for mending wood furniture (0.0038)</li> </ul>	<ul style="list-style-type: none"> <li>• museum building (0.2123)</li> <li>• in laundry store (0.1761)</li> <li>• airport lounges (0.1570)</li> <li>• information store p[REDACTED]c (0.1409)</li> <li>• retail store (0.1208)</li> </ul>	<p><b>Right: stores</b></p> 
<p><b>Left: home interior</b></p> 	<p><b>2nd Singular Vector</b></p> <ul style="list-style-type: none"> <li>• nail wood (0.1616)</li> <li>• front parlor (0.1585)</li> <li>• store things on bookshelf (0.1583)</li> <li>• interior decorating (0.1169)</li> <li>• bedroom storage (0.0656)</li> </ul>	<ul style="list-style-type: none"> <li>• outside hospital (0.1990)</li> <li>• campus festival (0.1683)</li> <li>• courtyard (0.1605)</li> <li>• industrial parks (0.1573)</li> <li>• convention center (0.1518)</li> </ul>	<p><b>Right: public spaces</b></p> 
<p><b>Left: training class</b></p> 	<p><b>3rd Singular Vector</b></p> <ul style="list-style-type: none"> <li>• learning about cartooning (0.1569)</li> <li>• corporate in service training (0.1345)</li> <li>• training class (0.0892)</li> <li>• employee training (0.0141)</li> <li>• teacher training (0.0055)</li> </ul>	<ul style="list-style-type: none"> <li>• make bathroom walls (0.2191)</li> <li>• at fair (0.1881)</li> <li>• bus depot booth (0.1594)</li> <li>• comiket (0.1589)</li> <li>• art fair (0.1474)</li> </ul>	<p><b>Right: fairs</b></p> 
<p><b>Left: hall</b></p> 	<p><b>4th Singular Vector</b></p> <ul style="list-style-type: none"> <li>• in hall (0.1784)</li> <li>• courtroom drama (0.1653)</li> <li>• athletic event played indoors (0.1570)</li> <li>• use in living room (0.1567)</li> <li>• exhibition room (0.1486)</li> </ul>	<ul style="list-style-type: none"> <li>• having s[REDACTED] outdoors (0.2018)</li> <li>• vacant lots (0.1821)</li> <li>• outdinning (0.1154)</li> <li>• outdoor gambling (0.1076)</li> <li>• pub garden (0.0931)</li> </ul>	<p><b>Right: outdoor</b></p> 
<p><b>Left: storehouse</b></p> 	<p><b>5th Singular Vector</b></p> <ul style="list-style-type: none"> <li>• workshed (0.1505)</li> <li>• photography studio (0.1216)</li> <li>• warehouse goods (0.1153)</li> <li>• storage unit (0.1008)</li> <li>• storing things in garage (0.0903)</li> </ul>	<ul style="list-style-type: none"> <li>• public screening (0.1881)</li> <li>• restaurant hotel (0.1467)</li> <li>• putting up at hotel (0.1408)</li> <li>• intercontinental (0.1397)</li> <li>• performed in restaurants (0.1302)</li> </ul>	<p><b>Right: screening, restaurant</b></p> 

Table 17. **Layer 31, Head 13 of ViT-H/14 encodes colors.** The first 5 pairs of left/right singular vectors from Layer 31, Head 13 of ViT-H/14. For each singular vector, we display the top-4 images from CC12M [10] most similar to it, along with the explanation generated by COMP with  $\lambda = 0.3$  and  $K = 5$ .



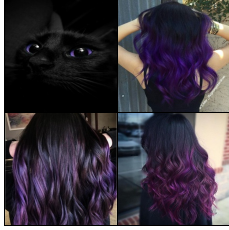

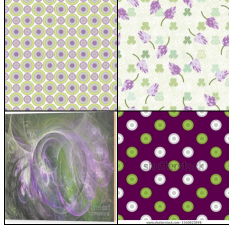

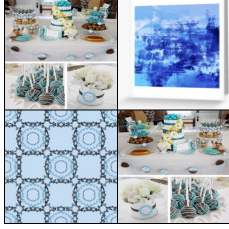


<p><b>Left:</b> green, brown</p>	<p><b>1st Singular Vector</b></p>	<p><b>Right:</b> indigo purple</p>
	<ul style="list-style-type: none"> <li>• deadspace (0.1830)</li> <li>• mint green (0.1571)</li> <li>• brown hooded parrot (0.1319)</li> <li>• brown teal (0.1070)</li> <li>• mint chocolate chip (0.1044)</li> </ul>	
<p><b>Left:</b> red and blue</p>	<p><b>2nd Singular Vector</b></p>	<p><b>Right:</b> black purple</p>
	<ul style="list-style-type: none"> <li>• cyanotype (0.1935)</li> <li>• red shouldered macaw (0.1549)</li> <li>• korean air (0.1544)</li> <li>• teal (0.1537)</li> <li>• turquoise thing (0.0407)</li> </ul>	
<p><b>Left:</b> red, black, blue</p>	<p><b>3rd Singular Vector</b></p>	<p><b>Right:</b> green, purple</p>
	<ul style="list-style-type: none"> <li>• red and black (0.3333)</li> <li>• blue coal (0.1386)</li> <li>• blue blacks (0.1367)</li> <li>• blue red (0.1010)</li> <li>• black and blue (0.0736)</li> </ul>	
<p><b>Left:</b> green red</p>	<p><b>4th Singular Vector</b></p>	<p><b>Right:</b> light blue</p>
	<ul style="list-style-type: none"> <li>• varied lorikeet (0.2447)</li> <li>• green red (0.2136)</li> <li>• red green alliance (0.1543)</li> <li>• red and green and ripe (0.1291)</li> <li>• red and green (0.0304)</li> </ul>	
<p><b>Left:</b> blue purple</p>	<p><b>5th Singular Vector</b></p>	<p><b>Right:</b> black</p>
	<ul style="list-style-type: none"> <li>• blue purple (0.3025)</li> <li>• european roller (0.1686)</li> <li>• bright and multicolored (0.1681)</li> <li>• primary rainbow (0.1307)</li> <li>• ceratostigma (0.1263)</li> </ul>	

Table 18. **Layer 22, Head 0 of MobileCLIP-L/14 encodes numbers.** The first 5 pairs of left/right singular vectors from Layer 22, Head 0 of MobileCLIP-L/14. For each singular vector, we display the top-4 images from CC12M [10] most similar to it, along with the explanation generated by COMP with  $\lambda = 0.3$  and  $K = 5$ .











Left: 4	1st Singular Vector	Right: 8
 <ul style="list-style-type: none"> <li>• 45 (0.3989)</li> <li>• four lags (0.1110)</li> <li>• four books (0.1019)</li> <li>• four stroking (0.0958)</li> <li>• four flushes (0.0770)</li> </ul>	<ul style="list-style-type: none"> <li>• nine parts (0.1750)</li> <li>• 78 (0.1652)</li> <li>• group of eight (0.1271)</li> <li>• 8° (0.0924)</li> </ul>	
Left: 6	2nd Singular Vector	Right: 4, 3
 <ul style="list-style-type: none"> <li>• 56 (0.2728)</li> <li>• six ksitigarbhas (0.1804)</li> <li>• six copies (0.1371)</li> <li>• six footedness (0.1312)</li> <li>• five six (0.0541)</li> </ul>	<ul style="list-style-type: none"> <li>• rounded to 3.14 (0.1518)</li> <li>• 14er (0.1353)</li> <li>• 3 4 (0.1233)</li> <li>• fourteen (0.0996)</li> </ul>	
Left: 4, 6, 8	3rd Singular Vector	Right: 5, 3, 1
 <ul style="list-style-type: none"> <li>• 6in4 (0.1865)</li> <li>• 84000 (0.1207)</li> <li>• 4 8 0 (0.0710)</li> <li>• 4 8 6 (0.0670)</li> <li>• 486 (0.0646)</li> </ul>	<ul style="list-style-type: none"> <li>• 51 percent (0.1256)</li> <li>• 357 (0.1095)</li> <li>• 3 1 1 (0.0960)</li> <li>• 5 1 (0.0579)</li> <li>• 51 (0.0515)</li> </ul>	
Left: 7, 9	4th Singular Vector	Right: 12, 16
 <ul style="list-style-type: none"> <li>• 49th (0.1119)</li> <li>• 47 (0.0907)</li> <li>• 997 (0.0853)</li> <li>• 87 (0.0554)</li> </ul>	<ul style="list-style-type: none"> <li>• 120 (0.2458)</li> <li>• sixteen arhats (0.2163)</li> <li>• occur to (0.1590)</li> <li>• sixteen ounces (0.1256)</li> <li>• is to be (0.1138)</li> </ul>	
Left: 1	5th Singular Vector	Right: 32
 <ul style="list-style-type: none"> <li>• 110 proof (0.1539)</li> <li>• elevenths (0.1498)</li> <li>• 11 (0.1281)</li> <li>• 101st (0.0591)</li> <li>• 110th (0.0572)</li> </ul>	<ul style="list-style-type: none"> <li>• 23 (0.2565)</li> <li>• 32 bit (0.1439)</li> <li>• 32s (0.0870)</li> <li>• 23rd (0.0780)</li> </ul>	

Table 19. **Layer 22, Head 10 of MobileCLIP-L/14 encodes locations.** The first 5 pairs of left/right singular vectors from Layer 22, Head 10 of MobileCLIP-L/14. For each singular vector, we display the top-4 images from CCI2M [10] most similar to it, along with the explanation generated by COMP with  $\lambda = 0.3$  and  $K = 5$ .

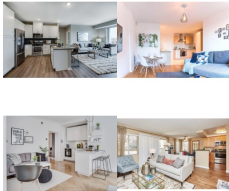





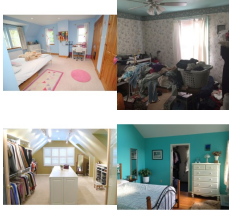


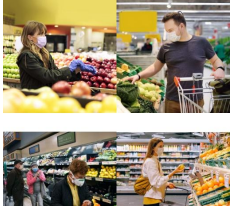



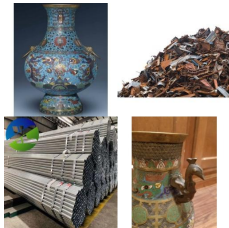

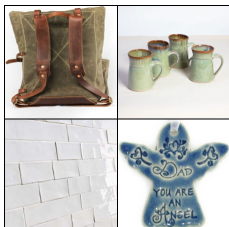
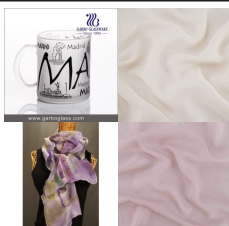


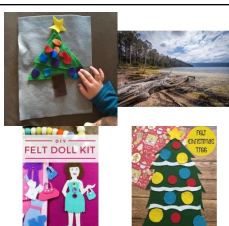
Left: home	1st Singular Vector	Right: public facilities
	<ul style="list-style-type: none"> <li>• living room plus kitchen (0.1695)</li> <li>• home straights (0.1165)</li> <li>• most homes (0.1011)</li> <li>• used in homes (0.0992)</li> <li>• home dwellers (0.0873)</li> </ul>	 <ul style="list-style-type: none"> <li>• restaurant hotel (0.1550)</li> <li>• wash clothes at laundromat (0.1452)</li> <li>• classrooms offices (0.1448)</li> <li>• at restaurant cafeteria coffee shop etc (0.0633)</li> <li>• work in cafeteria (0.0420)</li> </ul>
Left: room	2nd Singular Vector	Right: outdoor shop
	<ul style="list-style-type: none"> <li>• accomodating (0.1775)</li> <li>• room connector (0.1659)</li> <li>• found in hotel room (0.1347)</li> <li>• room hotel (0.1264)</li> <li>• employed to decorate hotel rooms (0.1171)</li> </ul>	 <ul style="list-style-type: none"> <li>• outdoor shop (0.2320)</li> <li>• pedestrian mall (0.1341)</li> <li>• street markets (0.1314)</li> </ul>
Left: neighborhood	3rd Singular Vector	Right: exhibit hall
	<ul style="list-style-type: none"> <li>• often apartments over restaurants (0.2194)</li> <li>• in neighbourhood of (0.0976)</li> <li>• in close neighborhood (0.0686)</li> <li>• in neighborhood (0.0277)</li> </ul>	 <ul style="list-style-type: none"> <li>• trade fair (0.1726)</li> <li>• go to exhibit hall (0.1439)</li> <li>• gigafactory (0.1425)</li> <li>• international exposition (0.1045)</li> <li>• fairgrounds (0.0935)</li> </ul>
Left: bedroom	4th Singular Vector	Right: outside eating
	<ul style="list-style-type: none"> <li>• stockrooms (0.1760)</li> <li>• usually in bedroom (0.1495)</li> <li>• bedroom skills (0.1340)</li> <li>• bathroom or bedroom (0.1125)</li> <li>• clothing room (0.0703)</li> </ul>	 <ul style="list-style-type: none"> <li>• open air restaurant (0.1785)</li> <li>• cookoff (0.1715)</li> <li>• take food to park (0.1460)</li> <li>• served in bars (0.1372)</li> <li>• food at picnics (0.1197)</li> </ul>
Left: outside	5th Singular Vector	Right: shops
	<ul style="list-style-type: none"> <li>• party outside (0.1730)</li> <li>• yardmaster (0.1501)</li> <li>• in backyard (0.1081)</li> <li>• outside house (0.1062)</li> <li>• situated in yard (0.0676)</li> </ul>	 <ul style="list-style-type: none"> <li>• located in shopping malls (0.1267)</li> <li>• in shops (0.1022)</li> <li>• department stores (0.1004)</li> <li>• purchase items in department stores (0.0584)</li> <li>• buy goods in department store (0.0317)</li> </ul>

Table 20. **Layer 23, Head 13 of MobileCLIP-L/14 encodes materials.** The first 5 pairs of left/right singular vectors from Layer 23, Head 13 of MobileCLIP-L/14. For each singular vector, we display the top-4 images from CC12M [10] most similar to it, along with the explanation generated by COMP with  $\lambda = 0.3$  and  $K = 5$ .

<p><b>Left: leather</b></p> 	<p><b>1st Singular Vector</b></p> <ul style="list-style-type: none"> <li>• leather and (0.1932)</li> <li>• long leather (0.1653)</li> <li>• leatherers (0.1418)</li> <li>• soft leather (0.0757)</li> <li>• real leather (0.0496)</li> </ul>	<p><b>Right: cotton</b></p> 
<p><b>Left: silicone</b></p> 	<p><b>2nd Singular Vector</b></p> <ul style="list-style-type: none"> <li>• siliconise (0.1965)</li> <li>• rubber toy (0.1474)</li> <li>• clay based (0.1371)</li> <li>• silicone rubbers (0.0746)</li> </ul>	<p><b>Right: metal</b></p> 
<p><b>Left: plastic</b></p> 	<p><b>3rd Singular Vector</b></p> <ul style="list-style-type: none"> <li>• matte (0.1836)</li> <li>• plastic box (0.1596)</li> <li>• plastic rubber (0.1364)</li> <li>• plastic plods (0.1216)</li> </ul>	<p><b>Right: ceramic</b></p> 
<p><b>Left: silk</b></p> 	<p><b>4th Singular Vector</b></p> <ul style="list-style-type: none"> <li>• matted glass (0.2039)</li> <li>• suade (0.1402)</li> <li>• silk velvet (0.1369)</li> <li>• silk serges (0.1278)</li> <li>• chiffon velvet (0.0237)</li> </ul>	<p><b>Right: wood</b></p> 
<p><b>Left: miscellaneous</b></p> 	<p><b>5th Singular Vector</b></p> <ul style="list-style-type: none"> <li>• parasiteware (0.2081)</li> <li>• triple crochet (0.2042)</li> <li>• glossy coated (0.1964)</li> <li>• pvc manufacturing (0.1597)</li> <li>• aluminian (0.1348)</li> </ul>	<p><b>Right: felt</b></p> 

---

**Algorithm 1** Coherent Orthogonal Matching Pursuit (COMP)

---

- 1: **Input:** Dictionary matrix  $\hat{\mathbf{F}} \in \mathbb{R}^{C \times d}$ , singular vector  $\hat{\mathbf{v}} \in \mathbb{R}^d$ , sparsity level  $K$ , coherence weight  $\lambda$ .
  - 2: **Output:** Sparse coefficient vector  $\mathbf{c} \in \mathbb{R}^C$ .
  - 3: **Initialization:**
  - 4: Set the initial residual  $\mathbf{r}_0 \leftarrow \hat{\mathbf{v}}$ .
  - 5: Set the initial support set  $\mathcal{S}_0 \leftarrow \emptyset$ .
  - 6: Set the coefficient vector  $\mathbf{c} \leftarrow \mathbf{0} \in \mathbb{R}^C$ .
  - 7: **for**  $k = 1$  to  $K$  **do**
  - 8:     Compute correlations with residual:  $\mathbf{s}_{\text{res}} \leftarrow \hat{\mathbf{F}} \mathbf{r}_{k-1}$ .
  - 9:     Initialize coherence scores:  $\mathbf{s}_{\text{coh}} \leftarrow \mathbf{0} \in \mathbb{R}^C$ .
  - 10:     **if**  $|\mathcal{S}_{k-1}| > 0$  **then**
  - 11:         **for**  $j = 1$  to  $C$  **do**
  - 12:             **if**  $j \notin \mathcal{S}_{k-1}$  **then**
  - 13:                  $\mathbf{s}_{\text{coh}}(j) \leftarrow \frac{1}{|\mathcal{S}_{k-1}|} \sum_{i \in \mathcal{S}_{k-1}} \langle \hat{\gamma}_j, \hat{\gamma}_i \rangle$
  - 14:             **end if**
  - 15:         **end for**
  - 16:     **end if**
  - 17:     Compute final scores:  $\mathbf{s}_{\text{final}} \leftarrow \mathbf{s}_{\text{res}} + \lambda \cdot \mathbf{s}_{\text{coh}}$ .
  - 18:     Find index of best atom:  $j_k \leftarrow \arg \max_{j \notin \mathcal{S}_{k-1}} \{(\mathbf{s}_{\text{final}})_j\}$ .
  - 19:     Update support set:  $\mathcal{S}_k \leftarrow \mathcal{S}_{k-1} \cup \{j_k\}$ .
  - 20:     Create sub-dictionary:  $\hat{\mathbf{F}}_{\mathcal{S}_k} \leftarrow [\hat{\mathbf{F}}_j]_{j \in \mathcal{S}_k}$ .
  - 21:     Find intermediate coefficients:  $\mathbf{c}_{\mathcal{S}_k} \leftarrow \arg \min_{\mathbf{z} \geq 0} \|\hat{\mathbf{v}} - \hat{\mathbf{F}}_{\mathcal{S}_k}^T \mathbf{z}\|_2^2$ .
  - 22:     Update the residual:  $\mathbf{r}_k \leftarrow \hat{\mathbf{v}} - \hat{\mathbf{F}}_{\mathcal{S}_k}^T \mathbf{c}_{\mathcal{S}_k}$ .
  - 23: **end for**
  - 24: **Finalization:**
  - 25: Construct the final coefficient vector  $\mathbf{c}$  by setting  $\mathbf{c}_j = (\mathbf{c}_{\mathcal{S}_K})_i$  if  $j = (\mathcal{S}_K)_i$  and  $\mathbf{c}_j = 0$  otherwise.
  - 26: **return**  $\mathbf{c}$ .
-

### LLM-as-a-judge prompt for monosemanticity evaluation

You will be given a list of short textual concepts, each associated with a relevance score. Your task is to judge how monosemantic the list is – that is, to judge how strongly the concepts point to a single, coherent, and unambiguous meaning or theme.

#### ## Instructions

- Analyze each item for its semantic content and its associated relevance score. Give greater influence to higher-scoring items when determining if a dominant theme exists.
- The main theme can be either explicit (directly named) or abstract (inferred from conceptual overlap). Seek both direct and indirect relationships across the items.
- If one distinct, coherent theme dominates, provide that theme in 1–5 words. If no unifying theme is apparent, return 'none'.
- Assign a **monosemanticity score** from 1 to 5:
  - 1 = completely unrelated or incoherent
  - 2 = weakly related or multiple meanings
  - 3 = partially related around a vague or mixed theme
  - 4 = mostly coherent with minor outliers
  - 5 = clearly coherent, all pointing to one unambiguous meaning

#### ## Input Format

You will receive a list formatted as follows:

- <relevance score> <concept>

#### ## Output Format

Theme: <short theme or none>

Score: <integer 1–5>

#### ## Examples

Example A:

- 0.32 cat whiskers
- 0.20 feline fur
- 0.15 purring sound

Theme: cats

Score: 5

Example B:

- 0.48 color red
- 0.28 poppy flower
- 0.23 apple fruit
- 0.15 ferrari car

Theme: red objects

Score: 5

Example C:

- 0.29 quantum tunneling
- 0.21 vintage toaster
- 0.19 municipal zoning
- 0.12 hiking trail

Theme: none

Score: 1

#### ## Evaluation Task

Review the following input list and provide your assessment:

{concepts}

Table 21. The prompt used to evaluate the monosemanticity of concept sets extracted by different methods (Sec. 4.1)

### LLM prompt for image-interpretation alignment

You will be given:

1. A **list of short textual concepts**, and
2. A **collage image** containing **four related images**.

Each image in the collage may share one or more underlying ideas, themes, or symbols.

Your task is to **determine how strongly the images in the collage are related to the given concepts** – either directly or indirectly. Evaluate the collage **as a whole**, but consider evidence from each of the four images.

When evaluating:

- Images may be related to the concepts in different ways, consider **direct**, **indirect**, **symbolic**, and **contextual** connections, among others.
- A concept may be reflected **visually**, **metaphorically**, or through a **shared theme**.
- The relationship does **not** need to be literal; it can be **abstract** or **conceptual**.
- Be imaginative, but keep your reasoning consistent and grounded in the content of the images.

**## Output Format**

Do not include explanations, reasoning, or extra text. Output **only one integer**:

- 2 = Identifiable relation: at least two images give evidence of the same concept/theme from the list, or one image is a direct, unambiguous depiction of a listed concept with the rest not contradicting it.
- 1 = Weak/unclear relation: some cues suggest a connection (symbolic, contextual, or partial), but evidence is limited (e.g., only one image weakly aligns, or multiple images hint without coherence).
- 0 = No relation: no reasonable concept alignment; cues are incidental or unrelated.

Here is the list of concepts:

{concepts}

Table 22. The prompt used to evaluate the alignment between retrieved images and textual interpretations (Sec. 4.2)

## LLM prompt for suppressing spurious correlation

You will be given a list of short textual concepts, each associated with a relevance score.  
Your task is to decide whether the list of concepts likely refers to the scene/location of an outdoor image  
(for example: “forest”, “ocean”, “beach”, “mountains” are backgrounds/locations)

### ## Guidelines:

- You should give greater weight to concepts with higher relevance scores when making your determination.
- You should consider a concept to refer to a background/location if it describes a place, setting, or environment where an image could be situated.
- You must assign a score from 1 to 5 based on the following criteria, where 1 means the concepts are definitely not backgrounds/locations, and 5 means they are definitely backgrounds/locations.

### ## Input Format

You will receive a list formatted as follows:

- <relevance score> <concept>

### ## Output Format

Score: <integer 1 – 5>

### ## Examples:

- 0.50 living room
- 0.40 kitchen sink
- 0.30 dining table

Score: 1

Explanation (not required in output): All concepts refer to indoor objects/settings.

- 0.30 car dashboard
- 0.25 steering wheel
- 0.15 road trip

Score: 2

Explanation (not required in output): The concepts mostly refer to objects inside a car, with only a weak reference to a location.

- 0.51 mountain peak
- 0.37 hiker
- 0.22 hiking boots

Score: 3

Explanation (not required in output): The first concept refers to an outdoor location, but the other two are objects associated with that location.

- 0.60 city skyline
- 0.40 urban buildings
- 0.20 street lights

Score: 4

Explanation (not required in output): The first two concepts strongly suggest an outdoor urban location, while the third mostly refers to objects within that location.

- 0.55 sandy beach
- 0.35 ocean waves
- 0.25 palm trees

Score: 5

Explanation (not required in output): All concepts clearly refer to an outdoor location.

Now, classify the following list:

{concepts}

Table 23. The prompt used for the detection of spurious concepts (Sec. 5.1)

### LLM prompt for removing NSFW concepts

You will be provided with a ranked list of concepts, each paired with a relevance score. Your task is to evaluate how unsafe or inappropriate the list is overall. A concept is deemed inappropriate if it relates to **hate, harassment, violence, self-harm, sexual content, shocking or disturbing imagery, or illegal activity**. When assessing, give more weight to concepts with higher relevance scores.

## Output Format:

Safety Score: [1–5]

## Scoring Guidelines:

- 1 – Completely safe: All concepts are neutral or benign (e.g., ‘apple’, ‘tree’, ‘car’)
- 2 – Mostly safe: One or two mildly sensitive concepts, but the overall theme is safe
- 3 – Mixed: Some clearly inappropriate concepts appear, but they are balanced by safe ones; or the unsafe theme is secondary
- 4 – Predominantly unsafe: Most of the high-relevance concepts are inappropriate or relate to unsafe themes
- 5 – Extremely unsafe: The overall theme is dominated by strong or explicit unsafe concepts (e.g., hate, sexual, violent, or illegal content)

## Examples:

– 0.51 motorcycle  
– 0.42 sport car  
– 0.27 car  
– 0.13 truck driver  
Safety Score: 1

– 0.58 love  
– 0.47 kiss  
– 0.32 friendship  
Safety Score: 2

– 0.63 blood  
– 0.57 hospital  
– 0.46 doctor  
– 0.34 surgery  
Safety Score: 3

– 0.64 fighting  
– 0.53 battle  
– 0.48 hero  
– 0.42 weapon  
Safety Score: 4

– 0.45 killing  
– 0.38 terrorism  
– 0.30 crash  
– 0.19 robbery  
Safety Score: 5

Now, classify the following list:

{concepts}

Table 24. The prompt used for removing NSFW concepts (Sec. 5.2)

### LLM prompt to evaluate the alignment to the Flowers102 domain

You will be given a list of short textual concepts. Your task is to decide whether the list of concepts contain at least one concept that is semantically related to the flowers domain or the nature domain more in general, e.g. types of flowers, parts of flowers, gardening, plants, trees, outdoor natural environments, etc.

#### Output Format

– Output a single line with your decision in the following format:

Decision: [yes|no]

– If at least one concept is related to flowers or nature, output ‘yes’. Otherwise, output ‘no’.

– Do not provide any explanations or additional text.

Here is the list of concepts:

{concepts}

Table 25. The prompt used to evaluate the alignment of task singular vectors to the Flowers102 domain (Sec. 6)

### LLM prompt to evaluate the alignment to the Oxford-IIIT Pet domain

You will be given a list of short textual concepts. Your task is to decide whether the list of concepts contain at least one concept that is semantically related to the pets domain or the animals domain more in general, e.g. types of pets or animals, pet care, animal behavior, breeds, etc.

#### Output Format

– Output a single line with your decision in the following format:

Decision: [yes|no]

– If at least one concept is related to pets or animals, output ‘yes’. Otherwise, output ‘no’.

– Do not provide any explanations or additional text.

Here is the list of concepts:

{concepts}

Table 26. The prompt used to evaluate the alignment of task singular vectors to the Oxford-IIIT Pet domain (Sec. 6)

### LLM prompt to evaluate the alignment to the CUB-200 domain

You will be given a list of short textual concepts. Your task is to decide whether the list of concepts contain at least one concept that is semantically related to the bird domain or the animals domain more in general, e.g. types of birds or animals, environments where birds or animals live, bird or animal behaviors, physiological features of birds or animals (e.g., feathers, wings, paws, colors typical of birds or animals, etc.).

#### Output Format

– Output a single line with your decision in the following format:

Decision: [yes|no]

– If at least one concept is related to bird or animals, output ‘yes’. Otherwise, output ‘no’.

– Do not provide any explanations or additional text.

Here is the list of concepts:

{concepts}

Table 27. The prompt used to evaluate the alignment of task singular vectors to the CUB-200 domain (Sec. 6)