

Supplementary Material

A. Reproducibility

All the source code, figures, and models will be available at <https://github.com/jyhong0304/bi-ICE>.

B. Comparison to Other Intrinsically Interpretable Methods

We emphasize the novelties of our Bi-ICE framework compared to Concept Bottleneck Models (CBMs) and prototype-based approaches. In general, both CBMs and prototype-based models are restricted to providing interpretability only at the implementation level.

B.1. Limitations of CBMs and Prototype-based Approaches

CBMs suffer from a strong dependence on annotations, as they require either ground-truth concept labels or highly accurate concept predictors. They also face scalability challenges, often showing degraded performance when applied to large-scale datasets, such as ImageNet or Places365. In contrast, Bi-ICE achieves strong performance across both small- and large-scale settings, and operates effectively with or without concept supervision. From the standpoint of concept computation, Bi-ICE adopts a more general formulation than CBMs. Whereas concepts in CBMs are often assumed to be independently predictive in a relatively simple way by representing each concept as a scalar value [66], our framework models concepts as high-dimensional vector embeddings. This richer representation allows the model to (i) capture semantic meaning in a distributed and relational manner through patch-wise alignments, (ii) preserve spatial grounding of concepts (e.g., the concept *beak* aligning with head patches), and (iii) support bi-directional interactions, where image patches refine concepts and concepts in turn broadcast back to refine image embeddings.

Prototype-based approaches face significant scalability challenges. Since each prototype corresponds to a stored exemplar, the number of prototypes must increase as dataset size and class diversity grow. This inevitably leads to *prototype explosion*, where hundreds or even thousands of prototypes are required for large-scale datasets, such as ImageNet. In practice, these prototype-based approaches are rarely applied with backbones larger than ViT-S, as the computational burden becomes prohibitive and the risk of overfitting to individual prototypes increases with larger-capacity models. As a result, these models incur substantial memory and computational costs and lose interpretability due to redundant or overlapping prototypes. In contrast, Bi-ICE leverages *shared concept embeddings* that are reused across all training examples. Each concept is represented as a high-dimensional vector that is refined during train-

ing. This design offers: i) efficiency—far fewer parameters compared to storing per-class prototypes, ii) generalization—concepts capture semantic features (e.g., *beak*, *wing*) that transfer across the dataset rather than being tied to specific instances, and iii) robust scalability—Bi-ICE scales effectively to datasets like ImageNet and Places365 with small number of concepts while maintaining accuracy, where prototype-based approaches become impractical. Moreover, prototype-based methods typically provide only qualitative faithfulness validation, relying on nearest-example visual matches. They lack systematic causal testing to ensure that prototypes are truly necessary for predictions. Bi-ICE overcomes this limitation by: i) performing concept insertion and deletion to verify causal necessity, ii) applying counterfactual interventions (e.g., recoloring a bird’s body) to induce systematic and interpretable shifts in concept activations and predictions and iii) offering patch-level grounding, where concepts localize to specific image regions, ensuring that explanations align with the actual input evidence.

Model	CIFAR	ImageNet	CUB	Places
LaBo [68]	69.1	70.4	71.8	39.4
PCBM [69]	57.2	62.6	63.9	39.7
Lf-CBM [48]	65.3	72.0	74.5	43.7
Sparse CBM [55]	74.9	71.6	80.0	41.3
Bi-ICE (Ours)	85.1\pm0.1	76.4\pm0.1	90.3\pm0.1	50.5\pm0.2

ing. This design offers: i) efficiency—far fewer parameters compared to storing per-class prototypes, ii) generalization—concepts capture semantic features (e.g., *beak*, *wing*) that transfer across the dataset rather than being tied to specific instances, and iii) robust scalability—Bi-ICE scales effectively to datasets like ImageNet and Places365 with small number of concepts while maintaining accuracy, where prototype-based approaches become impractical. Moreover, prototype-based methods typically provide only qualitative faithfulness validation, relying on nearest-example visual matches. They lack systematic causal testing to ensure that prototypes are truly necessary for predictions. Bi-ICE overcomes this limitation by: i) performing concept insertion and deletion to verify causal necessity, ii) applying counterfactual interventions (e.g., recoloring a bird’s body) to induce systematic and interpretable shifts in concept activations and predictions and iii) offering patch-level grounding, where concepts localize to specific image regions, ensuring that explanations align with the actual input evidence.

B.2. Additional Performance Comparison

While we have outlined the distinctions between our Bi-ICE and CBM/prototype-based methods above, we also recognize that these architectures are widely considered important milestones in image-based XAI for their explicit emphasis on interpretability. Accordingly, we include some representative CBMs and prototype-based models as baselines and provide a detailed comparison of their performance against our proposed approach.

Table A-1 compares the performance of our model, Bi-ICE, with various CBM architectures across four datasets: CIFAR100, ImageNet, CUB-200-2011, and Places365. Bi-ICE achieves 85.1% on CIFAR100, 76.4% on ImageNet,

Table A-2. Performance comparison between prototype-based approaches and ours on CUB-200-2011. For the prototype-based approaches, we report their best results from the original papers. The backbones and parameter sizes are as follows: PIP-Net uses ConvNeXt-Tiny (28.6M), ProtoPFormer uses DeiT-S (22M), ProtoPool uses ResNet-50 (25.6M), ProtoPNet uses an ensemble of VGG19+ResNet34+DenseNet121 (173.5M), Deform-ProtoPNet uses DenseNet-161 (28.7M), and ProtoConcept uses DenseNet-161 (28.7M). Note that in prototype-based models, the number of prototypes is heuristically fixed for each class and cannot be shared across classes. This design often results in redundant concepts (e.g., separate *head* prototypes for different bird species) and limits scalability to larger backbones such as ViT-L. In contrast, Bi-ICE allows concept vectors to be shared across classes, leading to more compact and efficient representations. Even with a lightweight backbone (ViT-S, 22M), Bi-ICE achieves competitive performance while offering better scalability and efficiency.

Model	CUB-200-2011
ViT-S : 22M	81.3 \pm 0.4
PIP-Net [47]	84.3
ProtoPFormer [67]	84.9
ProtoPool [52]	87.6
ProtoPNet [6]	84.8
Deform-ProtoPNet [13]	86.5
ProtoConcept [40]	85.2
Bi-ICE (Ours): 25.4M	85.3 \pm 0.2

90.3% on CUB-200-2011, and 50.5% on Places365, surpassing the strongest CBM competitor by clear margins on every dataset. Importantly, Bi-ICE obtains these results without relying on CLIP-like vision–language pretraining, which other CBMs used to compensate for the absence of ground-truth concept annotations. This demonstrates that Bi-ICE delivers state-of-the-art accuracy while maintaining interpretability, even under settings where other CBMs depend on external pretrained supervision.

For comparison with prototype-based approaches, we chose CUB-200-2011 only because for larger datasets such as ImageNet and Places365, the computational and memory requirements of prototype-based models grow significantly with the number of classes and prototypes, making them impractical for efficient evaluation.

Table A-2 shows the performance comparison with the prototype-based methods of CUB-200-2011. To ensure a fairer comparison with them, we configured Bi-ICE in an unsupervised concept learning setup with 50 number of concepts and a ViT-S backbone (22M parameters), closer in scale to ResNet- or DeiT-based backbones commonly used for the baseline. Even under this lightweight setup, Bi-ICE delivers 85.3% accuracy on CUB-200-2011. This performance is notable given that

prototype-based methods often depend on large and redundant ensembles (e.g., VGG19+ResNet34+DenseNet121) or heavy backbones (e.g., DenseNet-161). Moreover, while prototype-based architectures offer strong interpretability by grounding predictions in visual exemplars, they usually allocate a fixed set of prototypes per class, which limits their scalability to larger or more complex settings. For example, in the case of CUB-200-2011, most prototype-based methods allocate 10 prototypes per class, resulting in a total of 2,000 prototypes across the 200 classes. This design not only inflates the number of parameters but also produces redundant concepts (e.g., separate all different body parts prototypes for different bird species), thereby limiting scalability. In contrast, Our Bi-ICE adopts a concept-sharing mechanism, where a compact set of concept vectors (e.g., 50 concept vectors only) is shared across all classes. This yields far greater efficiency while preserving interpretability. Importantly, if prototype-based approaches were to increase the number of prototypes per class to match our concept capacity (e.g., 50 prototypes per class), the resulting computational and memory costs would become prohibitive—especially when scaling to larger backbones or more complex datasets. Bi-ICE thus achieves competitive accuracy with significantly fewer resources, highlighting its scalability advantage over prototype-based models.

C. Experimental Details

C.1. Dataset Statistics

Tab. A-3 summarizes the statistics of all benchmark datasets used in our experiments. Because none of the datasets includes a predefined validation split, we manually selected a portion of the training data to serve as a validation set for hyperparameter tuning.

The CIFAR-100 dataset contains 100 fine-grained image classes, but there are no ground-truth concept annotations that we can leverage. Thus, we arbitrarily set the number of latent spatial concepts to 20 to train our model.

For the CUB-200-2011, we followed the pre-processing steps outlined in [50]. Initially, the original dataset has 312 binary concepts, but we filtered them to retain only those that occur in at least 45% of all samples within a given class and occur in at least 8 classes. Thus, we got a total of 108 concepts, which we grouped into two types of concepts: spatial and global concepts. Specifically, we identified 13 global and 95 spatial concepts by looking at each concept. For example, `has_shape::perching-like` and `has_primary_color::black` are global concepts, and `has_eye_color::black` and `has_leg_color::grey` are spatial concepts.

The ImageNet and the Places365 [73] datasets do not contain the ground-truth concept annotations the same as CIFAR100. Therefore, we arbitrarily set the number of la-

Table A-3. Benchmark Dataset Statistics. † indicates the rescaled size of inputs for the ViT backbone, which is different from the original sizes of the datasets. ‡ indicates the number of concepts corresponding to latent spatial concepts.

Dataset	CIFAR100	CUB-200-2011	ImageNet	Places365
Input size	$3 \times 28 \times 28$	$3 \times 224 \times 224^\dagger$	$3 \times 224 \times 224^\dagger$	$3 \times 224 \times 224^\dagger$
# Classes	100 (class)	200 (bird species)	1000 (objects)	365 (scenes)
# Concepts	20 [‡]	13 (global), 95 (spatial)	150 [‡]	200 [‡]
# Training samples	55,000	5,994	255,224	1,803,460
# Validation samples	5,000	1,000	10,000	30,000
# Test samples	10,000	4,794	20,000	36,500

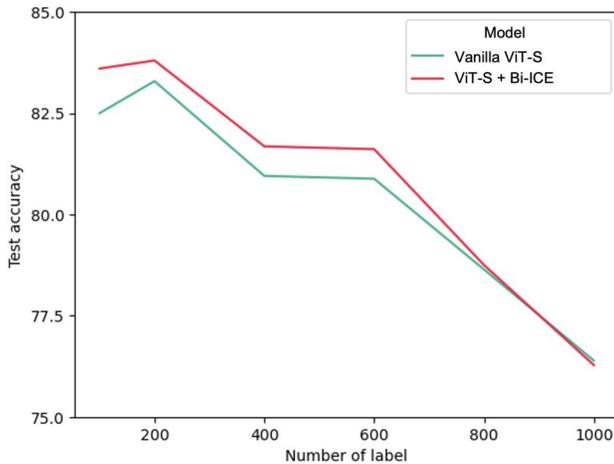


Figure A-1. Performance comparison on ImageNet with different numbers of labels.

tent concepts to 150 for ImageNet and 200 for Places365, respectively.

C.2. Model Architectures

We provide additional details of the ResNet50 [27] and Vision Transformers (ViT) [14] used as backbones. We employ `timm` Python library supported by Hugging FaceTM. In our experiments, we leverage ResNet50 (`resnet50.tv.in1k`), and three variants of ViT (`vit_large_patch16_224`, `vit_small_patch16_224`, `vit_tiny_patch16_224` in `timm`). These variants are defined by their number of encoder blocks, the number of attention heads on each block, and the dimension of the hidden layer. Specifically, ViT-L has 24 encoder blocks with 16 heads, and the dimension of the hidden layer is 1024. ViT-S has 12 encoder blocks with 6 heads, and the dimension of the hidden layer is 384. Finally, ViT-T has 12 encoder blocks with 3 heads, and the dimension of the hidden layer is 192, which is much more lightweight.

Table A-4. Ablation study of the number of concepts on ImageNet.

# of concept vectors	50	100	150	200
Test Acc. (%)	76.28 \pm 0.2	76.31 \pm 0.2	76.43 \pm 0.1	76.35 \pm 0.1

Table A-5. Hyperparameter setups for Vanilla ResNet50, Vanilla ViT-T, and all configurations of Bi-ICE on CIFAR100 in the main text.

Name	Value
Batch size	64
Epochs	20
Warmup Iters.	10
Learning rate	$1e - 4$
Explanation lambda λ_{expl}	1.0
Weight decay	$1e - 3$
Sparsity lambda λ_{sparse}	0.5

Table A-6. Hyperparameter setups for Vanilla ResNet50, Vanilla ViT-L, and all configurations of Bi-ICE on CUB-200-2011 in the main text. The learning rate for vanilla ResNet50 and ResNet50 + Bi-ICE is $1e - 4$, and the learning rate for vanilla ViT-L and ViT-L + Bi-ICE is $1e - 5$.

Name	Value
Batch size	32
Epochs	50
Warmup Iters.	10
Explanation lambda λ_{expl}	1.0
Weight decay	$1e - 3$
Sparsity lambda λ_{sparse}	0.5

Training Setups in the Main Text. For the vanilla ResNet50 and ViT models, we replaced the last linear layer with a new one and fine-tuned it to compare their performance with our Bi-ICE module. This is because our module can be seen as a drop-in replacement for the classifier head of any backbone network. All configurations, including the vanilla models and our proposed module, were

Table A-7. Hyperparameter setups for Vanilla ResNet50, Vanilla ViT-S, and all configurations of Bi-ICE on ImageNet in the main text. The learning rate for vanilla ResNet50 and ResNet50 + Bi-ICE is $5e-5$. The learning rate for vanilla ViT-S is $1e-5$ whereas the one for ViT-S + Bi-ICE is $1e-4$.

Name	Value
Batch size	64
Epochs	10
Warmup Iters.	10
Explanation lambda λ_{expl}	0.
Weight decay	$1e-3$
Sparsity lambda λ_{sparse}	10.

Table A-8. Hyperparameter setup for ViT-S + Bi-ICE on Places365

Name	Value
Batch size	64
Epochs	10
Warmup Iters.	10
Learning rate	$1e-4$
Explanation lambda λ_{expl}	0.
Weight decay	$1e-3$
Sparsity lambda λ_{sparse}	10.

trained using the AdamW optimizer with a Linear Warmup Cosine Annealing scheduler. The hyperparameter settings for CIFAR100, CUB-200-2011, and ImageNet are detailed in Tables A-5, A-6 and A-7. All models are trained on a NVIDIA GeForce Titan Xp GP102 (Pascal architecture, 3840 CUDA Cores @ 1.6 GHz, 384-bit bus width, 12 GB GDDR G5X memory).

D. Further Experimental Results

D.1. Performance Evaluation

Ablation Study. Fig. A-1 shows the experimental results on ImageNet with different numbers of labels to train vanilla ViT-S and ours. Interestingly, our Bi-ICE outperforms vanilla ViT-S when the number of labels is relatively small (< 800), while the performance of Bi-ICE becomes comparable or slightly worse than that of the vanilla model as the number of labels is enormous (≥ 800).

Table A-4 shows the experimental results on ImageNet with different numbers of concepts. We selected the number of concepts leading to the best performance, which is 150 for the performance evaluation on ImageNet.

Quantitative Comparison of Clustering. As shown in Fig. 4 in the main text, we qualitatively and quantitatively compared the clustering quality of the embeddings between the vanilla ViT (\mathcal{z}) and our Bi-ICE ($\tilde{\mathcal{z}}$). Table A-9 presents

Table A-9. ECE, Brier Score, NLL, and V-Score comparison of vanilla ViT and Bi-ICE embedding clusters on CUB-200-2011. The smaller the value, the better clustering for ECE, Brier Score, and NLL, while the higher the value, the better for the V-score.

Model	ECE	Brier Score	NLL	V-Score
Vanilla ViT	0.005	0.1737	0.4915	0.9124
Bi-ICE	0.002	0.1496	0.3982	0.9366

the additional comparison of quantitative metrics using Expected Calibration Error (ECE) and Negative Log Likelihood (NLL) [25], Brier Score [5], and V-score [51]. ECE measures calibration error, and NLL assesses probabilistic quality. Brier Score quantifies the accuracy of probabilistic predictions. Overall, Bi-ICE embedding is better clustered quantitatively than vanilla ViT, leading to better classification performance.

D.2. Hierarchical Classification Analysis on CIFAR100Super-class

Setup. The CIFAR100 Super-class dataset is a modified version of the CIFAR100 image dataset introduced by [38]. It comprises 100 detailed (fine-grained) image classes, which are grouped into 20 broader super-classes. For example, the fine-grained classes *baby*, *boy*, *girl*, *man*, and *woman* are all categorized under the super-class *people*.

Our Bi-ICE model can be trained using ground-truth concept labels by leveraging the Explanation loss detailed in Section 3.3 of the main text. This allows us to effectively evaluate the model’s ability to align with these annotated concepts. Since the advent of deep learning models with embedded logical constraints [18], the CIFAR100 Super-class dataset has served as a benchmark to test how well such constraints can be incorporated into neural networks (for extensive reviews, see [10, 21]). Our study demonstrates that the concepts learned by our model—and those used in Concept Transformer (CT) [50]—can be interpreted as constraints. In our experiments, we treated the fine-grained classes as global concept-level explanations for multi-class prediction across the 20 super-classes, utilizing a ViT-T backbone architecture.

In line with [30], we compare the performance of Bi-ICE against three baseline model groups. The first includes standard backbone architectures such as *Wide ResNet 28-10* [70] and *ViT-T*. The second group consists of neural models that embed logical constraints, such as the *Hierarchical Model* [30], *DL2* [18], and *MultiplexNet* [30]. The third group includes concept-based explainable models; however, with the exception of CT, these models are typically unable to manage both fine-grained and super-class predictions simultaneously. In our evaluation, we assess fine-grained class accuracy by comparing the top-1 pre-

Table A-10. Test accuracy on fine-grained class (F.C.) and super-class (S.C.) label prediction on CIFAR100. Notice that the classification of S.C. and F.C. are performed simultaneously and that this kind of experiment can be done in deep learning with constraints, but our method and CT are only among concept-based approaches. † indicates results from [30]. MultiplexNet used WideResNet 28-10 (36.5M) as a backbone. ProtoPNet and Deform-ProtoPNet used DenseNet-121 (8.0M) as a backbone. As a backbone for CT and our Bi-ICE, ViT-T was used (5.7M).

Model	F.C. Acc. (%)	S.C. Acc. (%)
Vanilla WideResNet†	NA	83.2±0.2
Vanilla ViT-T	NA	86.2±0.3
Hierarchical Model†	71.2±0.2	84.7±0.1
DL2† [18]	75.3±0.1	84.3±0.1
MultiplexNet† [30]	74.4±0.2	85.4±0.3
ProtoPNet [6]	NA	82.3±0.1
Deform-ProtoPNet [13]	NA	83.7±1.0
CT [50]	73.3±2.9	92.1±0.2
Bi-ICE	83.4±0.1	93.0±0.1

dicted concept with the ground-truth annotation, based on the resulting attention scores.

Results. As illustrated in Table A-10, our Bi-ICE method outperforms all competing baselines. Figure A-2 provides two illustrative examples where Bi-ICE demonstrates superior performance in both quantitative and qualitative terms compared to CT. One reason for CT’s weaker performance on fine-grained classification may be its tendency to generate hallucinated concepts in an attempt to optimize super-class accuracy without forming meaningful latent representations. While both models achieve correct super-class predictions in these examples, their reasoning processes diverge significantly—highlighting Bi-ICE’s superior explanatory power.

D.3. Concept-conditioned Pointing Game at Patch Resolution on CUB-200-2011

In addition to the CIFAR100Super-class experiments, we evaluate our model on CUB-200-2011 using a pointing-game protocol tailored to concept localization.

Setup. We introduce *Concept-Conditioned Pointing Game at Patch Resolution (CPGP)*, which specifies (i) a patch-level concept localization task, (ii) well-defined metrics, and (iii) reporting and interpretation guidelines. Conceptually, CPGP is a concept-level counterpart to the grid-pointing game [4]: rather than asking whether the most salient *location* falls on the ground-truth object, we ask whether the most salient *patch* for a *given concept* lands on any ground-truth occurrence of that concept. This

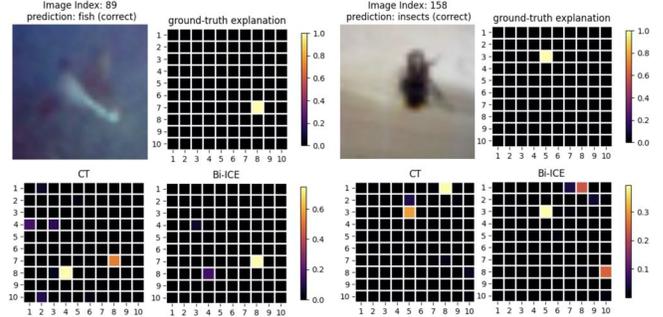


Figure A-2. Comparison of class predictions for CT [50] and Bi-ICE in examples where both make correct CIFAR100 super-class predictions. The 100 classes are indexed from 1 (top left in 10x10 grid) to 100 (bottom right in 10x10 grid). (Left) Ground-truth class label is *ray* (68), but CT mispredicted as *shark* (74), whereas Bi-ICE’s prediction is correct. (Right) Ground-truth label is *cockroach* (25), but CT incorrectly selects *beetle* (8) whereas our Bi-ICE again makes a correct class prediction.

patch-level formulation avoids threshold tuning and makes the evaluation robust to score scaling across images and concepts.

Given a trained model, each test image is partitioned into a patch lattice of size $H_p \times W_p$ with $L = H_p \cdot W_p$ patches. As described in Sec. 4.2, the ground-truth (GT) spatial concept annotations are represented by $Q^{\text{spatial}}(x) \in \{0, 1\}^{L \times K_{\text{spatial}}}$, where K_{spatial} is the number of spatial concepts and $Q_{i,k}^{\text{spatial}}(x) = 1$ indicates that patch i of image x contains concept k (otherwise 0). Our framework produces matching-dimension concept importance scores $\Phi^{\text{spatial}}(x) \in \mathbb{R}^{L \times K_{\text{spatial}}}$ where larger values indicate stronger evidence for the concept at that patch. We consider an image–concept pair (x, k) *valid* if the image contains at least one positive GT patch for k , i.e., $\sum_{i=1}^L \mathbf{1}[Q_{i,k}^{\text{spatial}}(x) = 1] > 0$. Let \mathcal{V} denote the set of all valid pairs.

Metric 1: Pointing-game accuracy (hit@1). For each $(x, k) \in \mathcal{V}$, we locate the most activated patch:

$$i^*(x, k) \triangleq \arg \max_{i \in \{1, \dots, L\}} \Phi_{i,k}^{\text{spatial}}(x).$$

A *hit* occurs if $Q_{i^*(x,k),k}^{\text{spatial}}(x) = 1$. The overall CPGP pointing accuracy is then:

$$\text{CPGP-Acc} \triangleq \frac{\sum_{(x,k) \in \mathcal{V}} \mathbf{1}[Q_{i^*(x,k),k}^{\text{spatial}}(x) = 1]}{|\mathcal{V}|}.$$

Metric 2: Per-concept ROC–AUC (threshold-free ranking). For each concept k , we pool all patches across the test set into labeled score pairs $\{(q_i^{(k)}, s_i^{(k)})\}_i$, where $q_i^{(k)} \in$

Table A-11. Hyperparameter setting for additional computational and implementation-level analyses.

Name	Value
Batch size	64
Epochs	50
Warmup Iters.	10
Learning rate	$5e - 5$
Explanation lambda λ_{expl}	10.0
Weight decay	$1e - 3$
Sparsity lambda λ_{sparse}	0.5

$\{0, 1\}$ comes from Q^{spatial} and $s_i^{(k)} = \Phi_{i,k}^{\text{sp}}$. We compute the ROC-AUC for k (defined only when both positive and negative patches are present):

$$\text{AUC}(k) \triangleq \text{ROC-AUC}\left(\{(q_i^{(k)}, s_i^{(k)})\}_i\right).$$

Let $\mathcal{D}_{\text{valid}}$ be the set of concepts with well-defined AUC; we report the macro mean:

$$\text{Mean-AUC} \triangleq \frac{1}{|\mathcal{D}_{\text{valid}}|} \sum_{k \in \mathcal{D}_{\text{valid}}} \text{AUC}(k).$$

Results. Our model attains CPGP-Acc of 0.593, indicating that in 59% of valid (image, concept) cases the single highest concept-attention peak lies on a GT-positive patch, demonstrating sharp concept-level localization. The Mean per-concept ROC-AUC is 0.719 (averaged over concepts with both classes present), evidencing strong ranking fidelity: GT-positive patches receive higher scores than negatives even when the top patch narrowly misses.

Novelty relative to grid-pointing games. (i) *Concept-conditioned evaluation:* CPGP scores localization per semantic concept (e.g., *wing*, *color*), directly auditing concept faithfulness rather than generic class focus. (ii) *Patch-resolution localization:* Hits are computed on the native patch lattice (e.g., ViT tokens), revealing fine-grained successes and failures that coarse 2×2 or 3×3 grid pointing games may obscure.

D.4. Analysis on Inner Interpretability

D.4.1. Computational Level

Training Setups. For the analysis, we trained an additional simple model as follows. We combined the frozen ViT-S as a feature extractor with our Bi-ICE module. The model leveraged the ground-truth global and spatial annotations during the training process to retrieve the concept importance scores for both global and spatial concepts. Table A-11 shows the hyperparameter setting of the model.

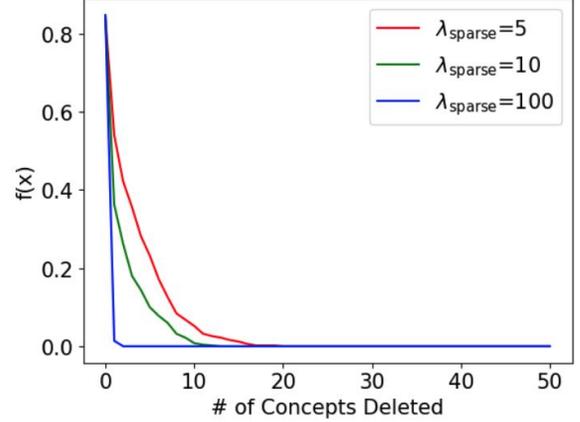


Figure A-3. C-deletion on ImageNet with the different λ_{sparse} . The higher the value, the sparser the concepts.

Additional Figures and Results. Fig A-3 shows the C-deletion experiment on ImageNet with varying sparsity lambda λ_{sparse} . Figs. A-4, A-5, and A-6 present additional examples of concept importance scores with varying numbers of activated patches, generated using our trained model with the aforementioned setting. Detailed explanations are included in the captions of each figure.

D.4.2. Algorithmic Level

Training Setups in the main text. For this experiment, we trained a simplified model with 15 concept vectors and 50 class labels using a ViT-S backbone without ground-truth concept annotations. This configuration was specifically designed to ensure a focused and interpretable analysis, enabling clear observation of the convergence of concept vectors both within each class and across different instances.

Additional Figures on CUB. In Fig. A-8a, we present an enlarged view of the concept vector convergence and refinement visualization initially shown in the main text (Fig. 6), offering a clearer perspective on the developmental process. This enlarged figure demonstrates how each concept vector focuses on specific semantic regions of bird images, such as the beak, wings, and lower body. As shown in Fig. A-8b, the framework consistently converges across different instances to identify the same meaningful regions, underscoring its reliability in capturing interpretable concepts.

In the t-SNE analysis shown in Fig. A-8c, we used the experimental setup to include 50 class labels and 15 concepts, utilizing a ViT-small backbone. This offers a comprehensive perspective on the clustering behavior of concept vectors. Additionally, the t-SNE visualization at the final epoch (epoch 50), presented in Fig. A-8d, shows that the model groups semantically similar features together. Notably, we observed well-defined clusters for features such as body underparts, head, legs, eyes, and beak regions.

Additional Figures on ImageNet. The algorithmic-level experiments on ImageNet are designed to analyze the behavior of concept vectors to demonstrate the applicability and versatility of our proposed method on a more challenging dataset. In the first set of ImageNet results shown in Fig. A-9, we focus on fish-like classes, training the model with 20 class labels and five concepts using ViT-S backbone. Through out our t-SNE graph (Figs. A-9b), we observed that only three concept vectors are sufficient to classify fish-like classes: Concept Vector 1 emphasizes the head of the fish, Concept Vector 3 captures the general body, and Concept Vector 4 focuses on the fins and tail.

Across other coupled ImageNet examples (Fig. A-10, Fig. A-11, Fig. A-12, and Fig. A-13), a consistent pattern emerges: concept vectors consistently capture semantically similar regions of the images, even when the labels vary. This consistency highlights the ability of the model to disentangle meaningful components of the images. Detailed explanations for each pair are provided in their respective captions.

Experimental Details and Submodel Differentiation on ImageNet. It is important to clarify that each of the experiments on ImageNet (Fig. A-9 - Fig. A-13) corresponds to a distinct submodel trained independently. These submodels were specifically designed to operate on subsets of ImageNet classes with closely related semantic meanings, ensuring that the concept vectors are tailored to capture features relevant to a coherent subset of labels.

For instance, the submodel in Fig. A-9 was trained on fish-like classes (e.g., Tench and Goldfish) leveraging five concept vectors to capture semantically meaningful components such as the fish’s body, fins, and broader scene elements. Similarly, in Fig. A-10, the submodel focused on amphibians, such as Salamander and Triturus vulgaris, with concept vectors isolating parts like the head, body, and background contours.

This approach of using semantically related subclasses for training was consistent across all figures. In Fig. A-11, the submodel was trained on antelope-like classes (e.g., Impala and Gazelle), which allowed the concept vectors to disentangle features such as horns, the general body, and the background. Similarly, in Fig. A-12, the focus was on human-centered image classes (e.g., Bikini and Two-Piece), allowing the model to identify specific body parts and their surrounding context. In Fig. A-13, the submodel targeted panda-like classes (e.g., Red Panda and Giant Panda), capturing distinctive facial features and contours.

In total, we trained 30 submodels, each focusing on a distinct subset of semantically related ImageNet classes. Among these, we selected and presented five representative submodels. By focusing on these specific cases, we demonstrated the robustness and adaptability of concept vectors

across various semantic contexts. This methodology underscores the scalability of our approach when applied to a complex and diverse datasets like ImageNet.

D.4.3. Implementation Level

Training Setups. For the analysis, we trained an additional simple model, which follows. We used the frozen ViT-S as a feature extractor and combined it with our Bi-ICE module. The model leveraged the ground-truth spatial annotation only during the training, which is the key difference from the setup in Sec. D.4.1. This is because the concept localization is mainly determined by the ground-truth spatial annotations. For the hyperparameter setting, we used the same setup in Table A-11.

Sensitivity Analysis of Thresholding the Most Activated Patches. Figure A-14 illustrates how different threshold values affect the visualization of concept activations. The top row shows whether a concept is activated or not, while the bottom row highlights which attributes are activated and the corresponding strength of their activations (with red indicating stronger activations and blue/transparent indicating weaker ones). At a high threshold of 0.8, for example, the eye region (Attribute 48: *has_eye_color::black*) is technically activated, but the intensity is so low that it is barely visible in the heatmap. At lower thresholds (e.g., 0.5, 0.3 or 0.1), the activations spread excessively, leading to blurred and ambiguous visualizations that are less interpretable. Empirically, we found that a threshold of 0.6 provides the clearest and most human-interpretable balance: Relevant regions become visible without the excessive diffusion observed at lower thresholds. Therefore, we chose 0.6 as the default threshold for our main experiments.

Additional Figures and Results. Firstly, Fig. A-15 depicts the detail of the example in Fig. 7 in the main text.

Furthermore, Fig. A-16 shows the additional concept localization results for three bird classes (Carolina Wren, Pigeon Guillemot, and Rock Wren), highlighting specific regions associated with various concepts. Concept 48, *has_eye_color::black*, consistently aligns with eye regions across all samples, demonstrating accurate identification of black eyes. Similarly, Concepts 79 (*has_leg_color::grey*) and 80 (*has_leg_color::black*) are localized to leg regions, emphasizing the module’s ability to localize and differentiate leg colors. These results demonstrate our module’s capacity for localized feature detection, aiding interpretability and classification accuracy.

E. Use cases

E.1. Use Case 1: Counterfactual Intervention on CUB.

Fig. A-7 demonstrates counterfactual intervention on an exemplary CUB sample. The caption of the figure includes the details.

E.2. Use Case 2: Empirical Evidence for the Number of Concepts.

The t-SNE visualizations from the ImageNet experiments (Fig. A-9b, Fig. A-10b, Fig. A-11b, Fig. A-12b, and Fig. A-13b) provide empirical evidence for determining the optimal number of concepts. By examining the results across different instance types—such as (i) fish-like classes (Fig. A-9), (ii) amphibians (Fig. A-10), (iii) animals with horns (Fig. A-11), (iv) human-centered images (Fig. A-12), and (v) panda-like instances (Fig. A-13)—we observed a recurring pattern: out of the five concepts used in training (except amphibian case in Fig. A-10), typically two concepts were redundant. These redundant concepts either focused on irrelevant areas of the image or were non-interpretable (e.g., capturing only a few random dots on the image). Meanwhile, the remaining three concepts consistently captured relevant and human-interpretable features.

To further validate this, we extended this analysis by varying the number of concepts used in the training. When we reduced the number of concepts to four, the same pattern persisted: two concepts remained redundant, while two concepts were useful. Similarly, with three concepts, one was redundant while two remained meaningful. In contrast, when we increased the number of concepts to ten, the number of redundant concepts grew to four, leaving only six as relevant.

These observations suggest that our framework can serve as a tool for determining the “sweet spot” for the number of concepts. For most experiments, setting the number of concepts to five provided a balanced trade-off between interpretability and representational capacity, avoiding unnecessary redundancy in the concept space.

By integrating these insights, our analysis underscores the utility of concept vectors in understanding not only the behavior of the model but also the fine-tuning of the model configurations for optimal performance and interpretability.

F. User Study

F.1. User Study Design

To evaluate the interpretability of unsupervised concepts on the CUB-200-2011 dataset, we conducted a user study under unsupervised setting. Specifically, we set the number of concept vectors to $K = 20$ and did not provide any hu-

man expert knowledge during concept discovery. All other model configurations followed the same setup described in the main experiments. For each concept vector, we then extracted the top-5 most highly activated samples, which serve as the clearest evidence of what the concept vector has captured in the absence of supervision.

A natural idea would be to ask participants to freely describe each concept by examining these top-5 activated samples (see Figure A-17). While such visualizations can reveal recurring patterns, open-ended descriptions are difficult to compare and aggregate among participants. To address this and following the approach in [64], we provide participants with a controlled vocabulary derived from the ground-truth annotations of the CUB dataset, allowing them to select relevant terms.

In each question, participants are shown bird images with highlighted regions, where red indicates stronger model attention. They are asked to assign the most appropriate labels in three categories: *body parts*, *action*, and *color*. The body part set includes head, beak, wing, leg, breast, tail, back, belly, eye, whole body shape, background, and a “none” option. The action set consists of flying, swimming, climbing, sitting, and “none,” while the color set covers red, gray, beige, black, yellow, brown, white, blue, green, colorful and “none.” Participants must select exactly one option per category, or “none” if no clear match is visible. Each participant answers twenty such questions. Figure A-18 illustrates the survey interface.

This design provides a structured and consistent protocol for evaluating concept vectors in an unsupervised setting. Inclusion of “none” options avoids bias from forced choices. More importantly, aggregating responses across annotators enables us to assess whether the discovered concepts support consistent human interpretation. High agreement would suggest that the concepts are coherent and semantically meaningful, while low agreement would reveal ambiguity and highlight the limitations of interpretability in the unsupervised setting.

F.2. Metrics

We designed the following two metrics to evaluate learned concepts based on the user study.

F.2.1. χ^2 tests of independence (per category).

To test whether participants’ vocabulary choices were systematically dependent on the concept being annotated, we used a χ^2 test of independence [8] for each category (body part, color, action).

For each category, we created a contingency table with rows corresponding to the 20 concepts and columns corresponding to the vocabulary terms (e.g., head, wing, tail for body part). Each entry in the table represents the number of participants who chose a given term for a given example.

The test evaluates whether the observed distribution of responses $O_{i,j}$ for cell (i, j) differs from the distribution expected under independence $E_{i,j}$. The χ^2 statistics is:

$$\chi^2 = \sum_{i,j} \frac{(O_{i,j} - E_{i,j})^2}{E_{i,j}},$$

with degrees of freedom $(R - 1)(C - 1)$, where R and C are the number of rows and columns in the contingency table. A large χ^2 statistic with a small p -value indicates that the choice of vocabulary term is not independent of the concept; in other words, different concepts elicit different label distributions.

Furthermore, to quantify the strength of association, we used Cramer’s V [8]:

$$V = \sqrt{\frac{\chi^2/n}{\min(R - 1, C - 1)}},$$

where n is the total number of responses. Values closer to 0 indicate weak associations, while values closer to 1 indicate strong associations.

Finally, beyond global significance, we inspected standardized residuals:

$$z = \frac{O_{i,j} - E_{i,j}}{\sqrt{E_{i,j}}}$$

to identify which specific (concept \times vocabulary) cells contributed most. Large positive residuals indicate terms chosen disproportionately often for that example, while negative residuals indicate terms chosen less often than expected.

F.2.2. Mutual Information between Concepts

While χ^2 tests assess global dependencies between concepts and responses within each category, they do not capture the pairwise similarity or redundancy between examples. To address this, we adapted the Mutual Information between Concepts (MIC) metric from [64].

In our configuration, each of the 20 examples is treated as a concept. For each concept, we collected the distribution of participant responses in a given category. Here, we consider body part category only. This gives us a categorical random variable L_i for each concept i .

For each pair of concept examples i, j , we computed the mutual information between L_i and L_j :

$$\text{MIC}(i, j) \triangleq \text{MI}(L_i; L_j) = \sum_{u,v} p(u, v) \log \frac{p(u, v)}{p(u)p(v)},$$

where $p(u, v)$ is the joint probability of participant assigning vocabulary u to concept i and vocabulary v to another concept j .

Table A-12. χ^2 test results per category. df indicates degree of freedom.

Category	χ^2	df	p (asymptotic)	Cramer’s V
Body part	4171.64	190	$< 1 \times 10^{-300}$	0.82
Color	1308.05	209	2.8×10^{-158}	0.41
Action	777.16	76	9.1×10^{-117}	0.54

We report MIC in bits, which reflects the amount of information one example’s labels provide about another. For example, 1 bit corresponds to a binary distinction (yes/no). To determine whether observed MIC values were greater than expected by chance, we used a permutation test: participant labels were randomly shuffled across concepts 10,000 times, and MIC was recomputed for each shuffle. The permutation p -value is the proportion of shuffled MIC values greater than or equal to the observed MIC.

Low MIC indicates concepts are distinct, with participants assigning very different labels. While high MIC indicates redundancy, with participants using similar vocabularies for both concepts.

F.3. Results

F.3.1. χ^2 Analysis

Figures A-19, A-20 and A-21 depict the contingency heatmap for categories: *Body Parts*, *Color* and *Action*. Body part terms show the strongest discriminability, with most concepts dominated by a single label (e.g., beak, tail), producing near one-to-one mappings. Action terms also show structured patterns, with many examples dominated by sitting on a branch or flying, though a few cases exhibit ambiguity. Color terms are comparatively more diffuse, with some examples clearly tied to a single color (e.g., black, blue) but others displaying overlap across multiple terms. Together, these maps illustrate that body part cues provide the most robust separation, action terms offer strong but somewhat less distinct signals, and color terms are moderately discriminative.

Table A-12 summarizes the results, including χ^2 statistic, degrees of freedom, p -values (asymptotic), and effect sizes (Cramer’s V). The χ^2 tests revealed strong and highly significant dependencies between concept identity and vocabulary choices across all categories.

- Body part yielded the largest χ^2 statistic with Cramer’s $V = 0.82$, indicating an extremely strong association. This means participants’ body part choices were almost perfectly aligned with specific examples, reflecting the high diagnostic power of morphological cues.
- Action also showed a strong association, demonstrating that behavioral descriptors consistently differentiated examples.
- Color produced a smaller but still statistically robust ef-

fect. Although color was less discriminative than body part or action, it still contributed significantly to distinguishing among examples.

Taken together, these results confirm that all three vocabularies carried meaningful information for distinguishing examples, with body part as the most discriminative category, followed by action and then color. The effect sizes show that participants were not guessing, but instead made systematic and interpretable choices aligned with human-recognizable features.

F.3.2. Mutual Information between Concepts

The heatmap (Fig. A-22) presents the pairwise mutual information (MI) between the 20 concepts using only Body part annotations.

Most cells in the matrix are dark (low NMI < 0.3), indicating that body part responses for different examples are generally distinct. This aligns with the contingency map, where body part terms (e.g., beak, tail, wing) mapped sharply to specific examples. A cluster in the matrix shows significantly elevated NMI values (0.7–1.0, yellow cells with significance markers). On inspection, these examples were consistently labeled as *background* or *whole body parts*. This redundancy is desirable rather than problematic: it reflects that participants reserved these labels for examples where no finer morphological feature stood out. In other words, overlap here does not reflect confusion, but rather systematic use of "catch-all" terms.

Thus, this pattern strengthens the interpretability of the body part vocabulary. It shows not only that most examples are distinct but also that redundancy, where it exists, is aligned with intuitive and desirable fallback labels.

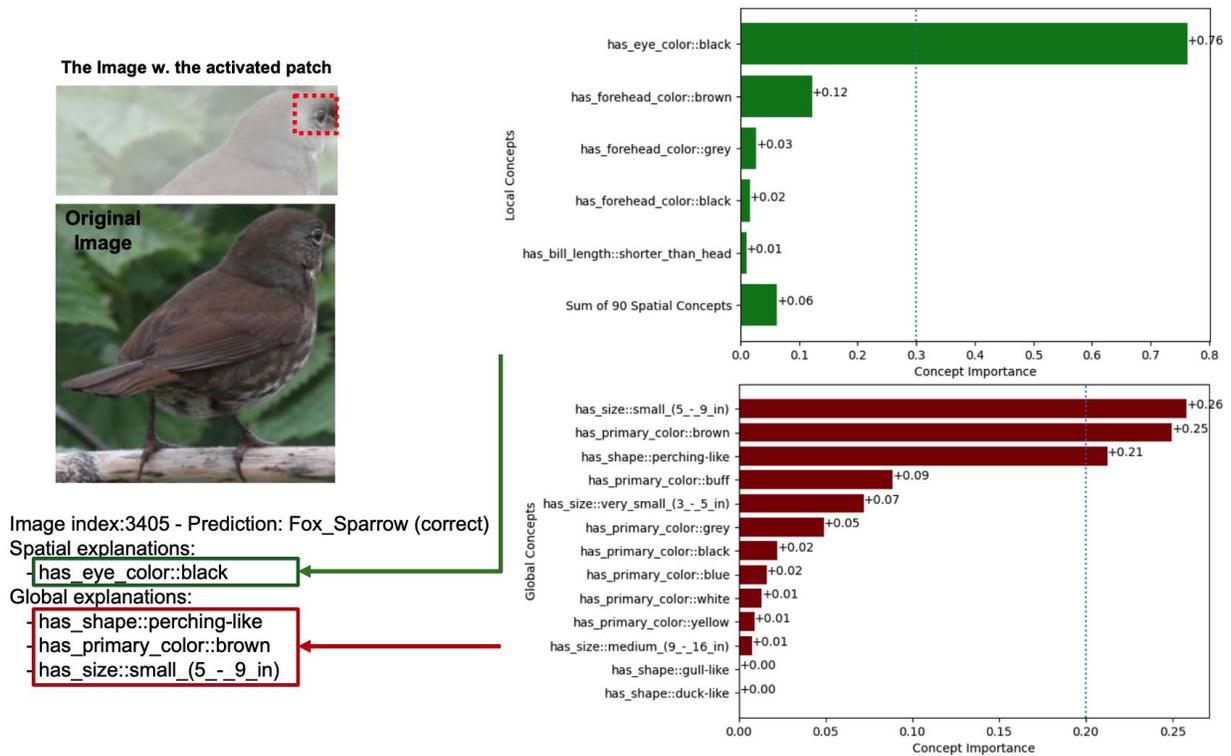


Figure A-4. The additional example for the computational-level interpretability analysis. The concept contributions of the class 'Fox Sparrow' illustrate how the concepts influence the model's decision. The dashed red box indicates the activated patch (> 0.6). Green for spatial, and red for global. The activated patch represents the region of the bird's eye and the forehead area. Accordingly, the spatial concept "has_eye_color::black" holds the highest contribution score (0.76), reflecting the clear presence of black in the bird's eye. Similarly, "has_forehead_color::brown" also shows the second highest importance, aligning with the brown coloration of the bird's forehead in the red box. In the global explanations, the bird's size and primary color were key concepts influencing the prediction. Concepts like "has_size::small_(5_-9_in)" and "has_primary_color::brown" rank the highest, with scores of 0.26 and 0.25, respectively, matching the Fox Sparrow's small size and dominant brown plumage.

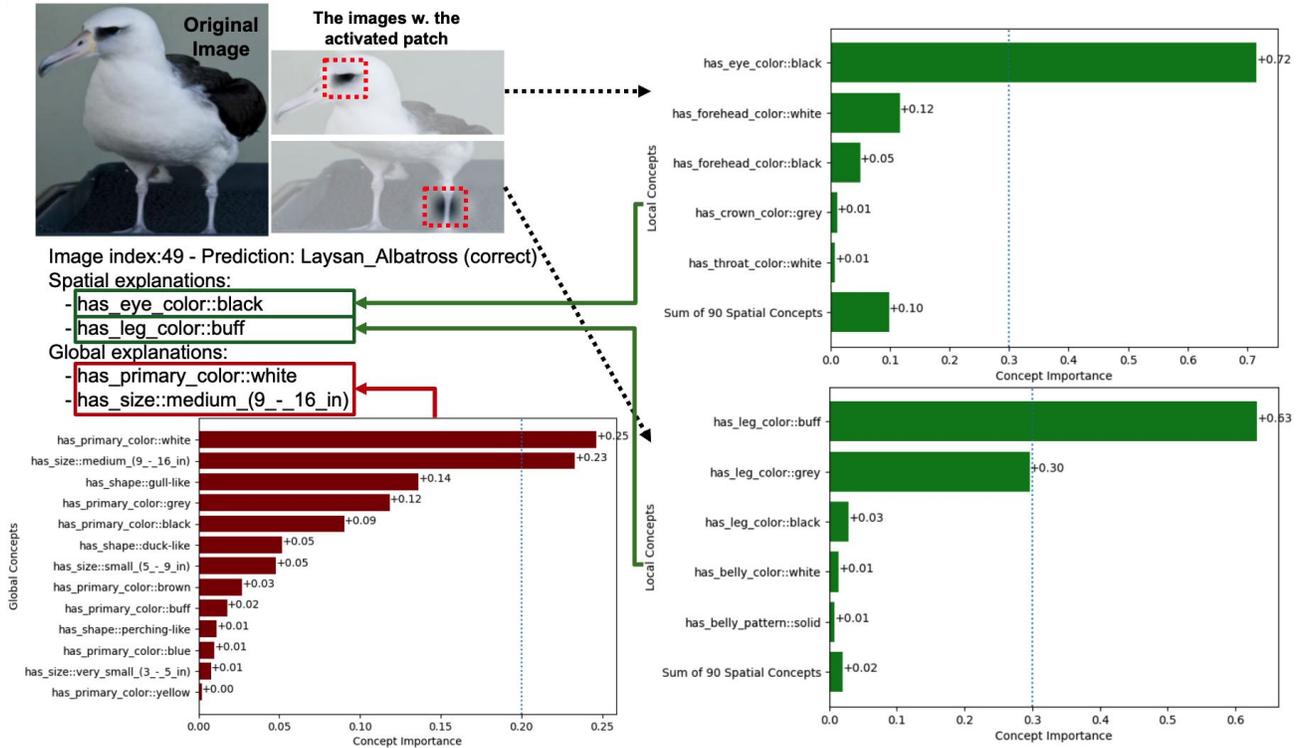


Figure A-5. The additional example for the computational-level interpretability analysis. The concept Contributions of the class ‘Laysan Albatross’ illustrate how the concepts influence the model’s decision. The dashed red boxes indicate the activated patches (> 0.6). Green for spatial, and red for global. The activated patches represent the regions of the bird’s eye and the leg areas, respectively. Accordingly, the spatial concept “has_eye_color::black” holds the highest contribution score (0.72) with the first patch, reflecting the dark eyes visible in the first patch. Similarly, “has_leg_color::buff” contributes significantly with a score of 0.63, aligning with the distinct buff-colored legs captured in the second patch. In the global explanations, broader concepts like “has_primary_color::white” and “has_size::medium_(9_-16_in)” are the most influential, with scores of 0.25 and 0.23, respectively, matching the Laysan Albatross’s overall appearance and size.



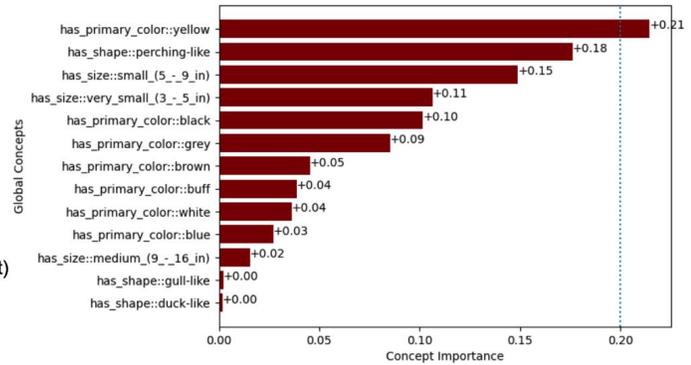
Image index:4678 - Prediction: Cape_May_Warbler (correct)

Spatial explanations:

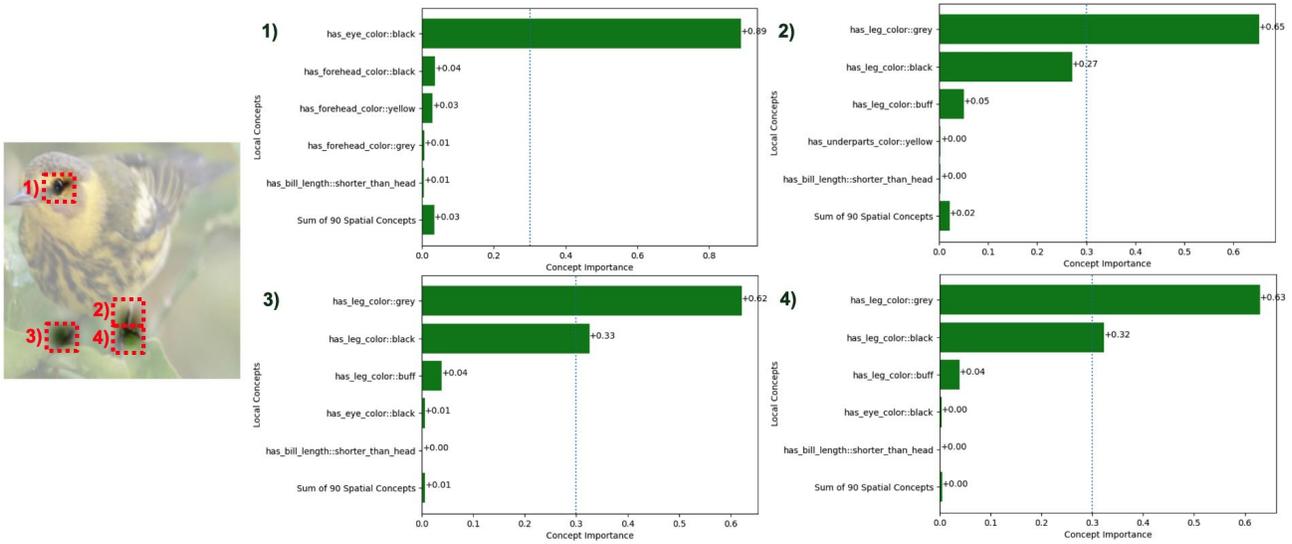
- has_eye_color::black
- has_leg_color::grey
- has_leg_color::black

Global explanations:

- has_primary_color::yellow

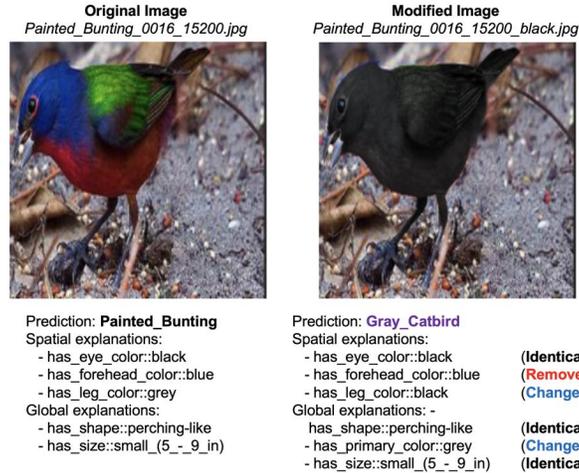


(a) Concept importance score of global explanation. Among the global concepts, “has_primary_color::yellow” holds the highest contribution score (0.21), reflecting the bird’s dominant yellow plumage. This is followed by “has_shape::perching-like” (0.18) and “has_size::small_(5_-9_in)” (0.15), which align with the bird’s distinctive perching posture and small size. Other color-related concepts, such as “has_primary_color::black” (0.10) and “has_primary_color::grey” (0.09), further highlight the importance of color features in the model’s prediction.

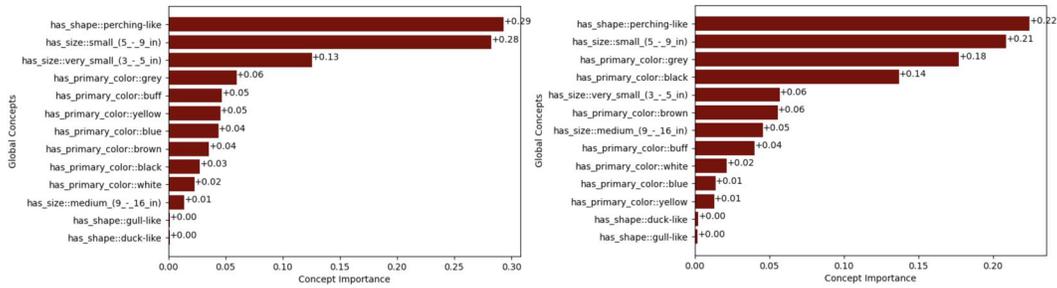


(b) Concept importance score of spatial explanations with their activated patches (red boxes). **1)** Patch 1 (Head region): The spatial concept “has_eye_color::black” contributes the highest score (0.89), accurately reflecting the bird’s black eye, a prominent feature in this region. **2)** Patch 2 (Upper leg region): The concept “has_leg_color::grey” holds the highest importance with a score of 0.65, emphasizing the grey coloration observed on the bird’s upper leg. “has_leg_color::black” also contributes with a score of 0.27, highlighting additional black tones in the same area. **3)** Patch 3 (Lower region): Similar to Patch 2, “has_leg_color::grey” dominates with a score of 0.62, representing the grey-colored lower leg. “has_leg_color::black” follows with a contribution of 0.33. **4)** Patch 4 (Feet region): Again, “has_leg_color::grey” holds the highest contribution (0.63), accurately reflecting the grey feet of the bird. “has_leg_color::black” is the second most important feature, with a score of 0.32.

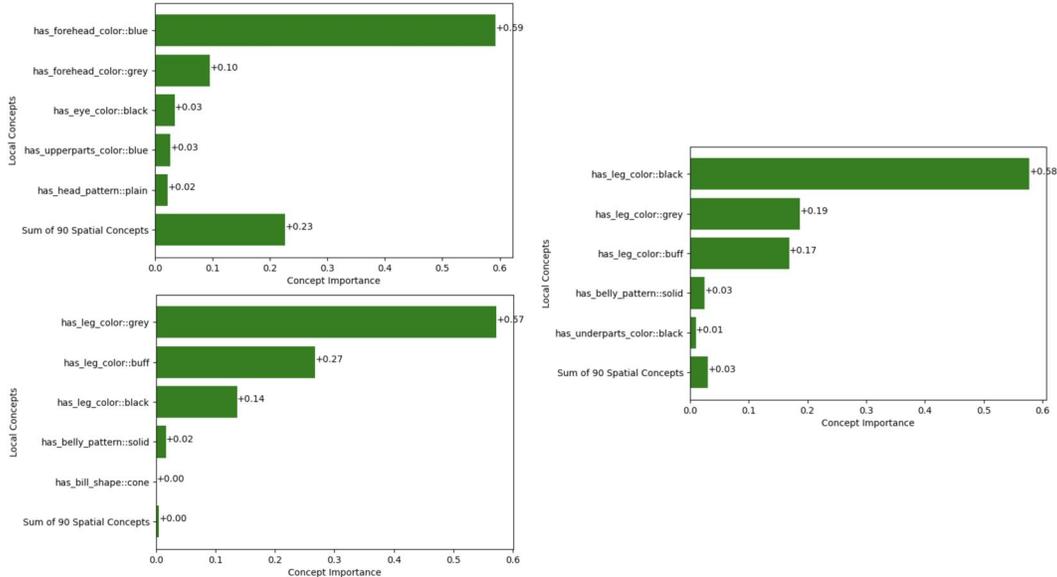
Figure A-6. The additional example for the computational-level interpretability analysis. The concept Contributions of the class ‘Cape May Warbler’ illustrate how the concepts influence the model’s decision. The dashed red boxes indicate the activated patches (> 0.6). Green for spatial, and red for global. The global and spatial explanations reveal how the model integrates both overarching and localized features-such as color, size, and posture-to accurately classify the class ‘Cape May Warbler’.



(a) (Left) The original image of *Painted Bunting* and its result, including prediction and explanations (Right) The modified image by coloring the body to black and its altered results. This artificially changing the color of the image input resulted in dramatic shifts by recognizing the concept `has_primary_color::grey` and updating others. Correspondingly, the bird originally classified as a *Painted Bunting* is then classified as a *Gray Catbird*.

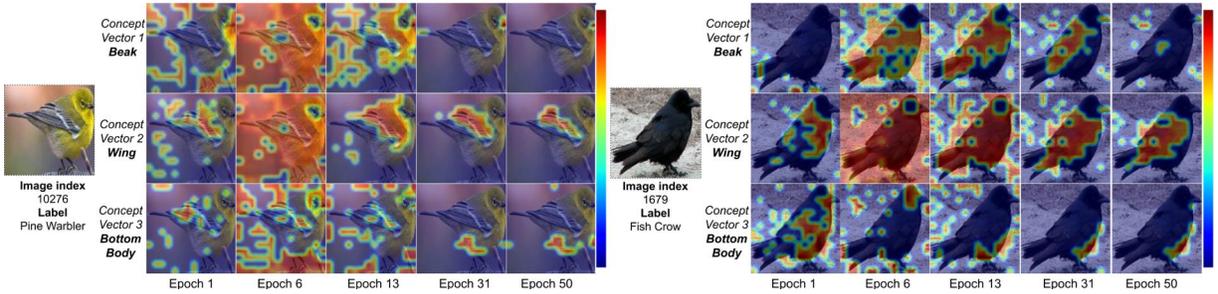


(b) The comparison of global explanations (Left) from the original image (Right) from the modified image

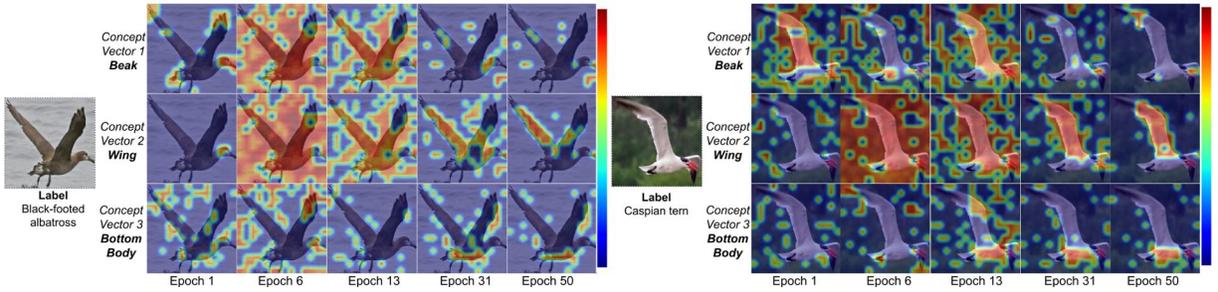


(c) The comparison of spatial explanations (Left) from the original image (Right) from the modified image

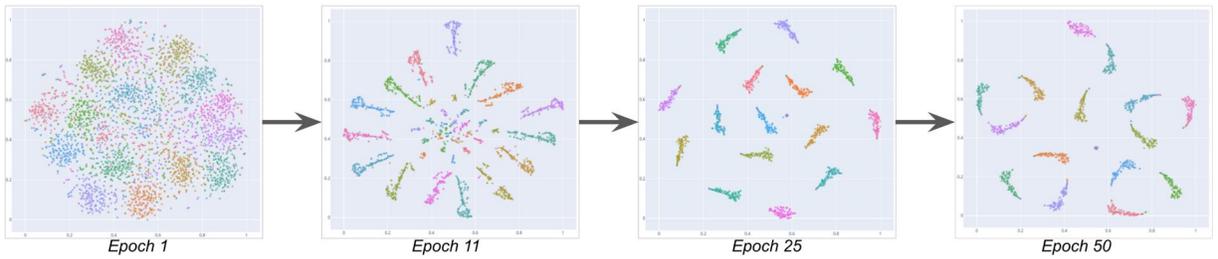
Figure A-7. Counterfactual intervention on an exemplary CUB sample. Artificially changing the body color of a *Painted Bunting* to black causes global and spatial concept shifts.



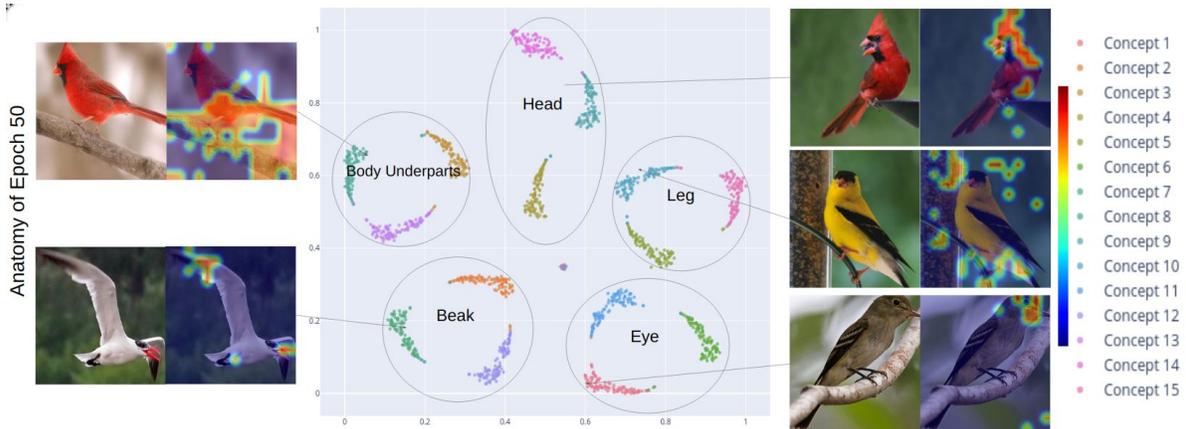
(a) Enlarged view of concept vector convergence and refinement (epoch 1 to 50). In the main text, a smaller version of this figure shows how concept vectors converge and refine from epoch 1 to 50. Here, we provide an enlarged version to better capture the details, highlighting the developmental process.



(b) New examples showing concept vector convergence and refinement. This panel introduces additional examples not shown in the main text. The model exhibits a similar developmental trend with different instances, where each concept vector gradually converges and refines into specific regions, as seen in the earlier example.

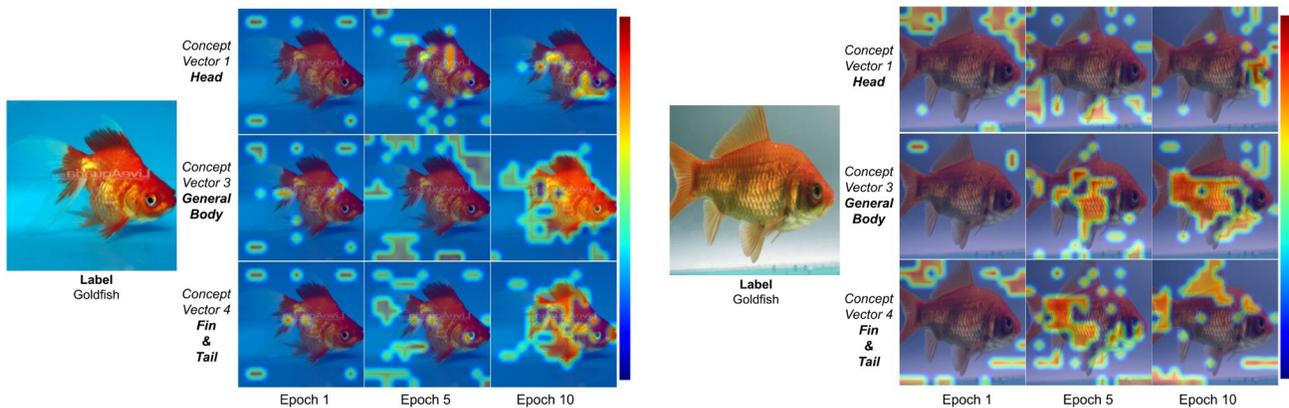


(c) t-SNE of concept vector convergence over epochs. To analyze the internal structure further, t-SNE visualizations show how concept vectors converge and refine across epochs. The clusters form as epochs advance, and by later epochs, the concept vectors align closely within their respective clusters.

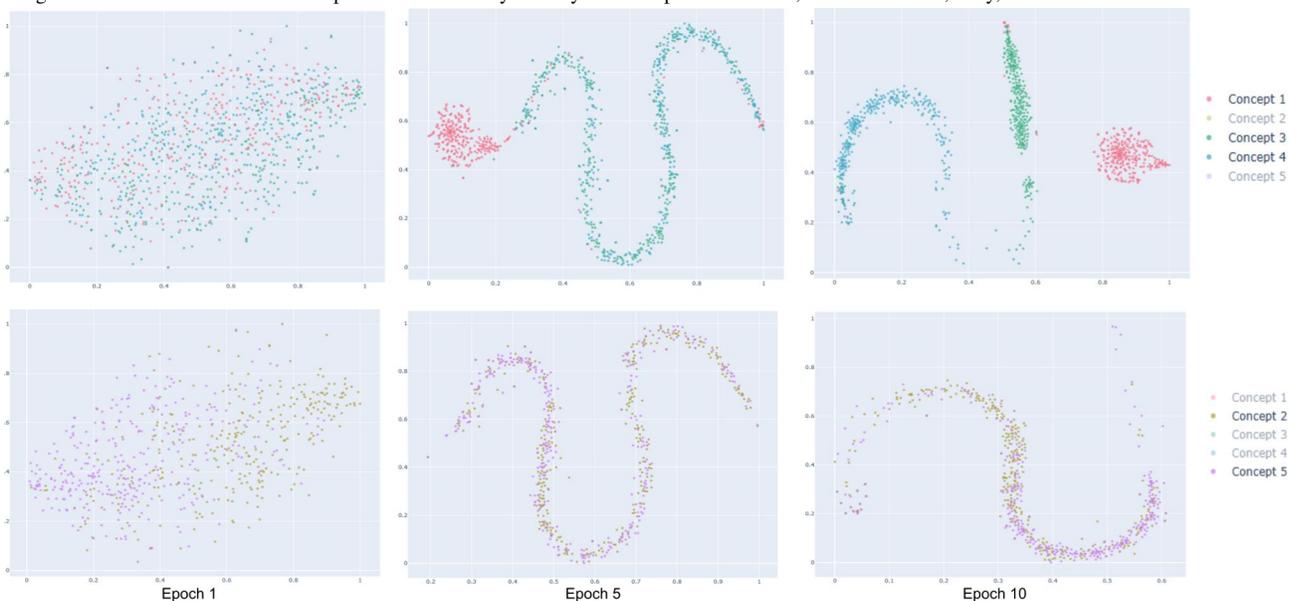


(d) t-SNE at epoch 50 showing final clustering. t-SNE visualization for epoch 50 (an enlarged version of the last panel in Fig.A-8c) demonstrates the fully developed state. At this stage, the clusters are well-formed, with each concept vector distinctly grouped.

Figure A-8. Convergence and Refinement of Concept Vectors Across Epochs on CUB instances. This figure illustrates the developmental process of the model as it refines its concept vectors over training epochs, showing both individual instance-level convergence and the overall clustering behavior of internal representations.

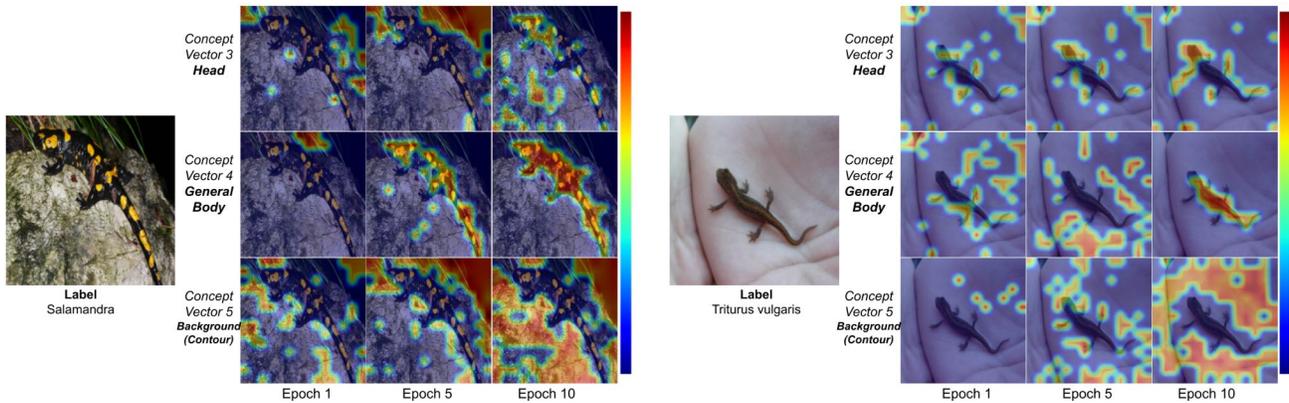


(a) Instance : Goldfish – Concept Vector Convergence and Refinement. This panel shows the evolution of concept vectors from epoch 1 to 10 for an image labeled as “Goldfish.” The concept vectors effectively identify different parts of the fish, such as the head, body, and fin with tail.

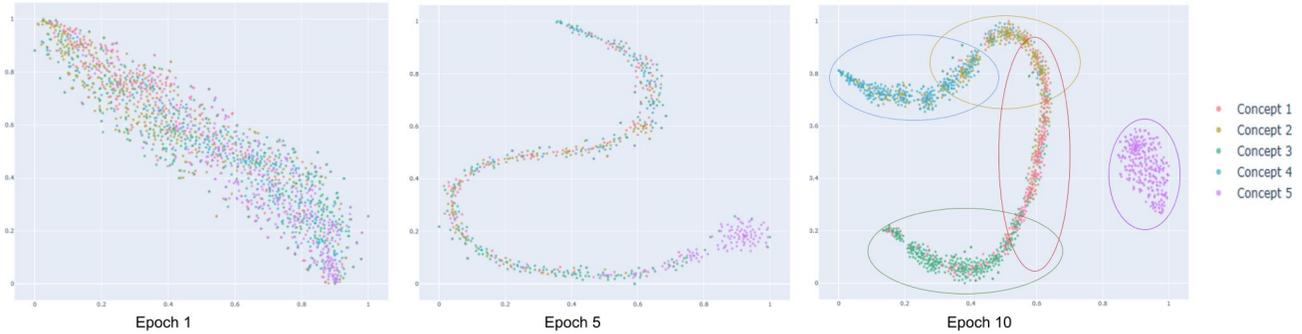


(b) t-SNE Visualization of Concept Vector Clustering. This panel shows the t-SNE projections of five internal concept vectors across epochs. Clusters 1, 3, and 4 group well, indicating meaningful internal representations, while clusters 2 and 5 overlap significantly, highlighting redundancy or focus on irrelevant image regions. Over epochs, convergence and refinement of these vectors can be observed.

Figure A-9. Convergence and Refinement of Concept Vectors Across Epochs on ImageNet instances. This figure illustrates the developmental process of the model as it refines its concept vectors over training epochs, showing both individual instance-level convergence and the overall clustering behavior of internal representations. The submodel presented here was trained on “Goldfish” with other fish-like classes to capture semantically meaningful features, such as the fish’s body, fins, and broader scene elements.

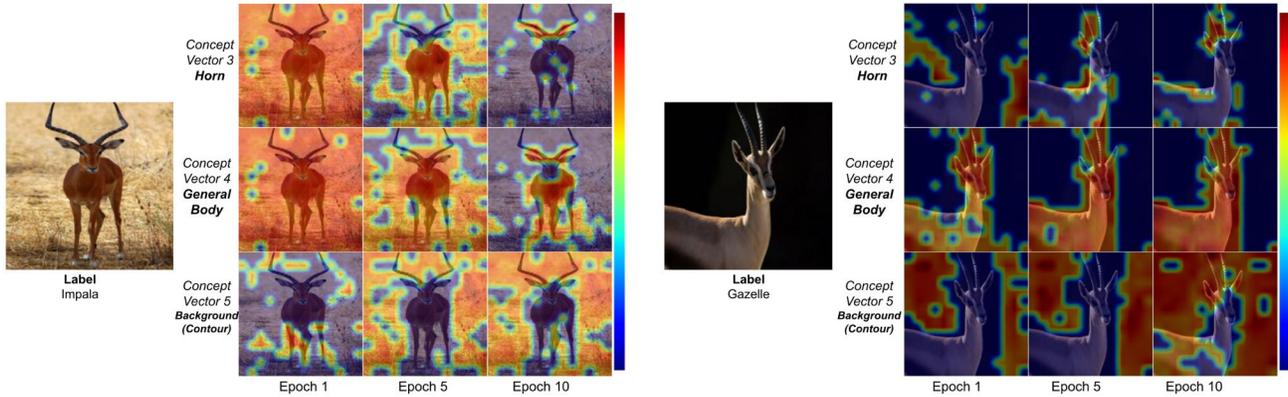


(a) Instance: Amphibians (Salamander and Triturus vulgaris) – Concept Vector Convergence and Refinement. This panel shows the evolution of concept vectors from epoch 1 to 10 for an image labeled as “Salamander” and “Triturus vulgaris.” Each concept vector identifies distinct parts of the image: both head, its general body, and the background (contours).

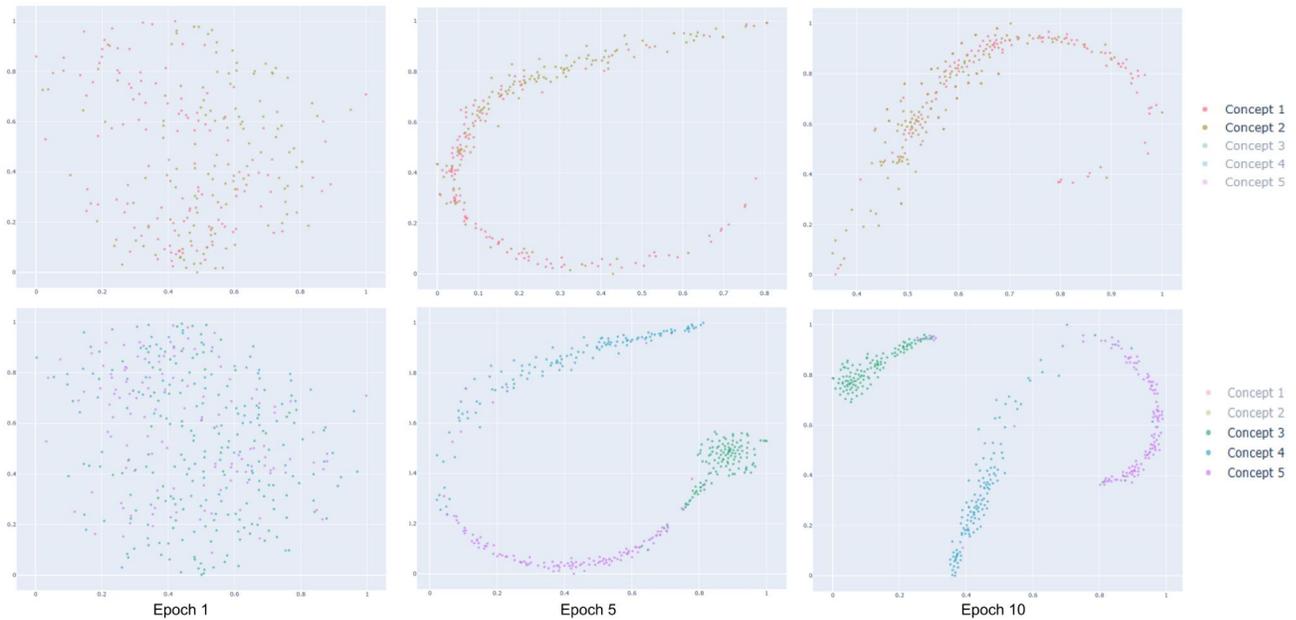


(b) t-SNE Visualization: Concept Vector Separation Across Epochs. This panel displays the t-SNE visualization of concept vectors for the same instance above. As epochs progress, the separation of concept vectors improves, with minimal overlap between clusters. This demonstrates the model’s ability to effectively disentangle and organize features without any ground truth guidance, although some overlap between concept vectors remains.

Figure A-10. Convergence and Refinement of Concept Vectors Across Epochs on ImageNet instances. This figure explores how concept vectors evolve and refine over epochs for two instances, alongside t-SNE projections highlighting separable and overlapping internal representations. The submodel presented here was trained on amphibian classes (e.g., Salamander and Triturus vulgaris), isolating features like the head, body, and background contours.

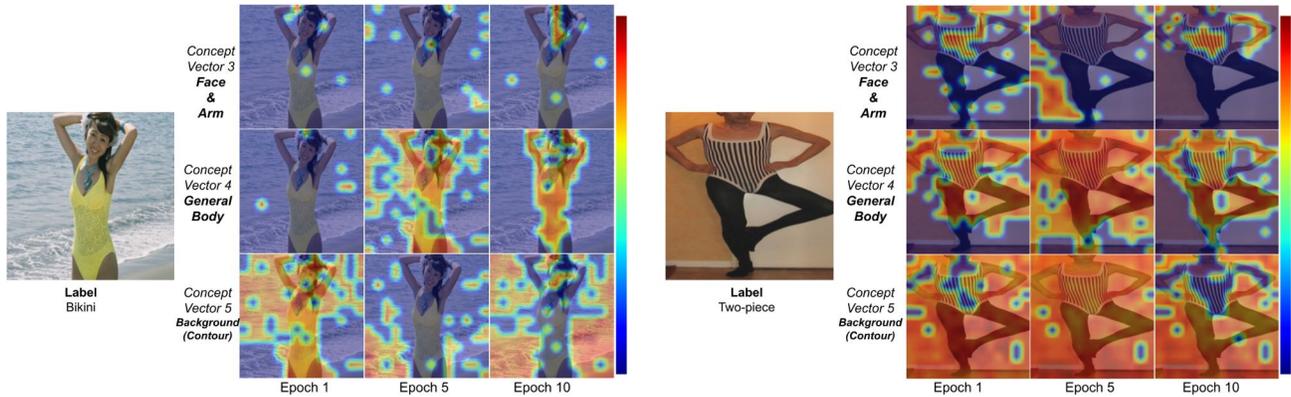


(a) Instance : Impala and Gazelle – Concept Vector Convergence and Refinement. This panel shows the evolution of concept vectors from epoch 1 to 10 for an image labeled as "Impala" and "Gazelle". Each concept vector isolates distinct aspects of the image: the horns, the general body, and the background (contours).

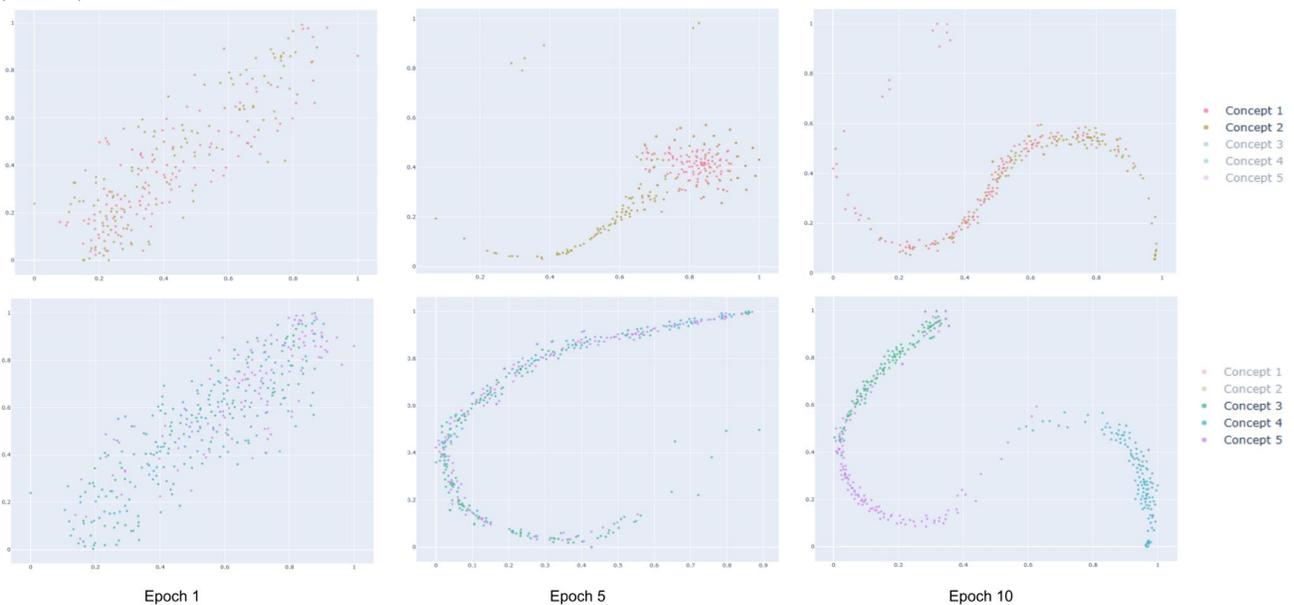


(b) t-SNE Visualization of Concept Vector Clustering. This panel highlights the t-SNE projections of concept vectors across epochs. Concept vectors 1 and 2 overlap significantly, indicating redundancy and focus on less interpretable or irrelevant image areas. In contrast, concept vectors 3, 4, and 5 cluster separately, aligning well with meaningful features and demonstrating effective disentanglement over training epochs.

Figure A-11. Convergence and Refinement of Concept Vectors Across Epochs on ImageNet instances. This figure explores how concept vectors converge and refine over epochs for two instances, alongside t-SNE projections highlighting separable and overlapping internal representations. The submodel presented here was trained on antelope-like classes (e.g., Impala and Gazelle), focusing on features such as horns, the general body, and the background.

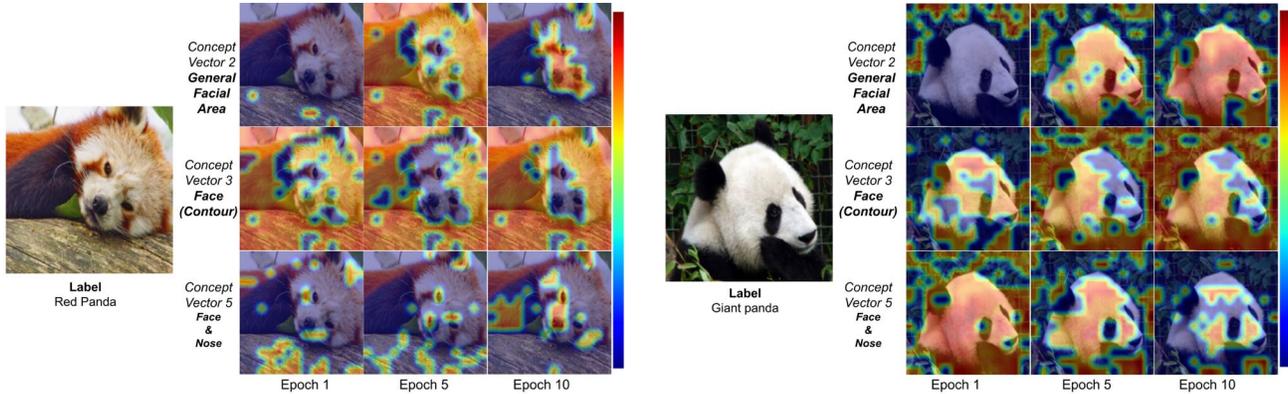


(a) Instance : Bikini and Two-Piece – Concept Vector Convergence and Refinement. This panel shows the evolution of concept vectors from epoch 1 to 10 for an image labeled as "Bikini" and "Two-Piece." The vectors isolate distinct components: the face and arms, the general body, and the background (contours).

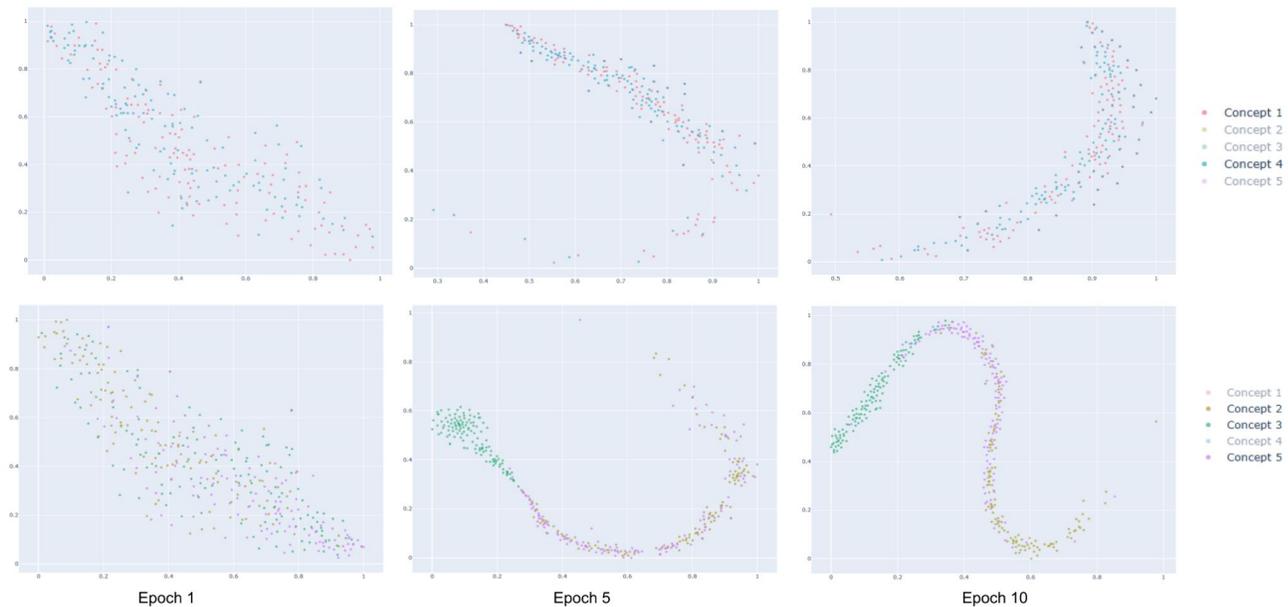


(b) t-SNE Visualization of Concept Vector Clustering. This panel highlights the t-SNE projections of concept vectors across epochs. Concept vectors 1 and 2 overlap significantly, indicating redundancy and focus on less interpretable or irrelevant image areas. In contrast, concept vectors 3, 4, and 5 cluster separately, aligning well with meaningful features and demonstrating effective disentanglement over training epochs.

Figure A-12. Convergence and Refinement of Concept Vectors Across Epochs on ImageNet instances. This figure showcases the refinement of concept vectors over epochs for two instances involving human-centered images, alongside t-SNE visualizations illustrating the clustering behavior of these internal representations. The submodel presented here was trained on human-centered image classes (e.g., Bikini and Two-Piece) to isolate specific body parts and their surrounding context.



(a) Instance : Red Panda and Giant Panda – Concept Vector Convergence and Refinement. This panel shows the evolution of concept vectors from epoch 1 to 10 for an image labeled as "Red Panda" and "Giant Panda." The model captures the general facial area, facial contour, and specific features such as the face and nose.



(b) t-SNE Visualization of Concept Vector Clustering. This panel highlights the t-SNE projections of concept vectors across epochs. Concept vectors 1 and 4 overlap significantly, indicating redundancy and focus on less interpretable or irrelevant image areas. In contrast, concept vectors 2, 3, and 5 cluster separately, aligning well with meaningful features and demonstrating effective disentanglement over training epochs.

Figure A-13. Convergence and Refinement of Concept Vectors Across Epochs on ImageNet instances. This figure showcases the refinement of concept vectors over epochs for two instances, alongside t-SNE visualizations illustrating the clustering behavior of these internal representations. The submodel presented here was trained on panda-like classes (e.g., Red Panda and Giant Panda) to capture distinctive features such as facial contours and specific markings.

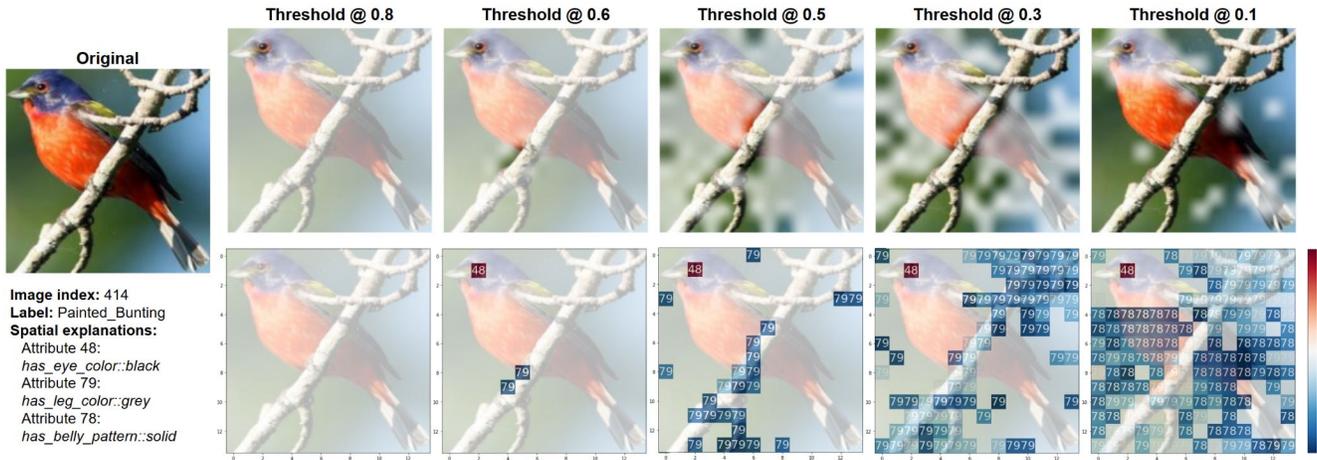


Figure A-14. Visualization of thresholding variation for concept activations. The top row shows whether concepts are activated at different thresholds, while the bottom row illustrates the specific attributes and their activation strengths (red = stronger, blue/transparent = weaker). At 0.8, relevant attributes (e.g., eye color) are barely visible due to very low intensity, while at lower thresholds (e.g., 0.3 or 0.1), activations spread excessively and become less interpretable. Empirically, we found that 0.6 provides the most human-interpretable balance, which motivated our choice of this threshold in the main experiments.

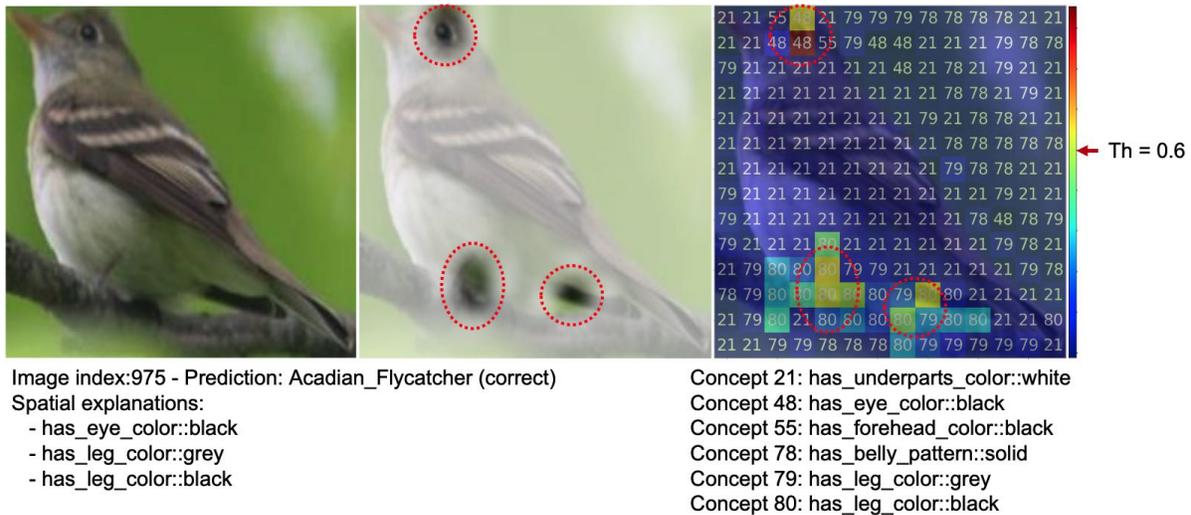


Figure A-15. The detail of the example in Fig. 7 in the main text. **Left:** The original image. **Middle:** The resultant image highlighting the activated patches. The active patches are displayed as their original image, while the rest are blurred. **Right:** The image shows the resultant concept localization with a 14 × 14 grid overlay. Each cell is color-coded with a colorbar, where red indicates a high concept contribution score, and blue indicates its low value. Each number in the cell indicates the corresponding concept of the highest concept importance value. The threshold determines whether each patch is active, contributing to the image in the middle.



Image index:5615, 5617, and 5624
 Prediction: Carolina_Wren (correct)
 Spatial explanations:
 - has_eye_color::black
 - has_leg_color::buff

Concept 21: has_underparts_color::white
 Concept 48: has_eye_color::black
 Concept 52: has_forehead_color::brown
 Concept 78: has_belly_pattern::solid
 Concept 79: has_leg_color::grey
 Concept 81: has_leg_color::buff

(a) The set of examples of class 'Carolina Wren'

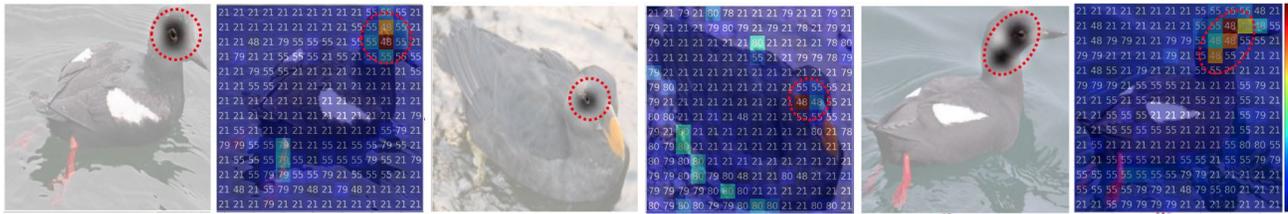


Image index:1599, 1601, and 1613
 Prediction: Pigeon_Guillemot (correct)
 Spatial explanations:
 - has_eye_color::black
 - has_forehead_color::black

Concept 21: has_underparts_color::white
 Concept 48: has_eye_color::black
 Concept 55: has_forehead_color::black
 Concept 78: has_belly_pattern::solid
 Concept 79: has_leg_color::grey
 Concept 80: has_leg_color::black

(b) The set of examples of class 'Pigeon Guillemot'

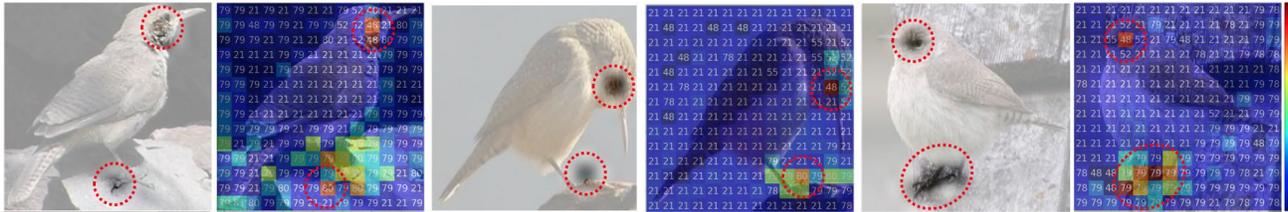


Image index:5706, 5707, and 5715
 Prediction: Rock_Wren (correct)
 Spatial explanations:
 - has_eye_color::black
 - has_forehead_color::brown
 - has_leg_color::grey
 - has_leg_color::black
 - has_leg_color::buff

Concept 21: has_underparts_color::white
 Concept 48: has_eye_color::black
 Concept 52: has_forehead_color::brown
 Concept 55: has_forehead_color::black
 Concept 78: has_belly_pattern::solid
 Concept 79: has_leg_color::grey
 Concept 80: has_leg_color::black

(c) The set of examples of class 'Rock Wren'

Figure A-16. The additional concept localization results for three bird classes (Carolina Wren, Pigeon Guillemot, and Rock Wren) on CUB, highlighting specific regions associated with various concepts. Overall, the samples in the same class included consistent spatial explanations corresponding to semantically meaningful regions, highlighted by the dashed red circles.

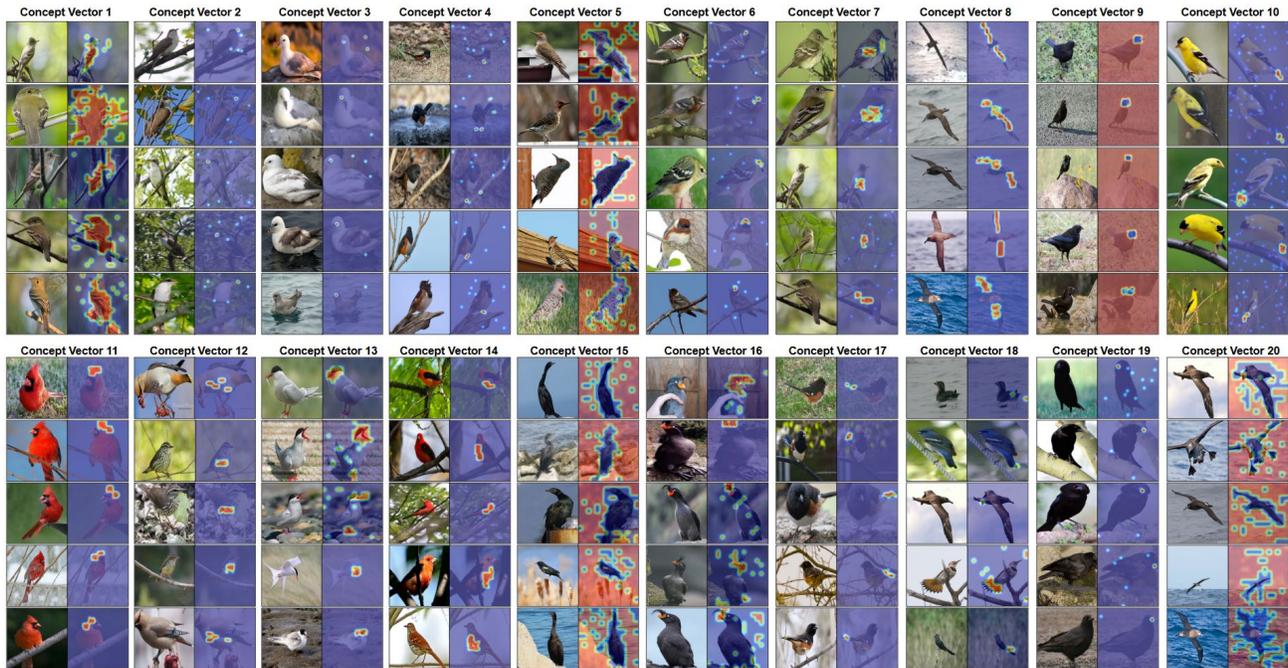


Figure A-17. Top-5 activated samples for twenty discovered concept vectors in the unsupervised setting on CUB200 dataset. Each column corresponds to one concept vector, and each row shows images that most strongly activate the corresponding vector. Highlighted regions (blue = low focus, red = stronger focus) indicate the spatial contributions to the activation. These examples illustrate the raw outputs of unsupervised concept discovery, which motivate the need for a user study: while the visual patterns appear meaningful, the crucial question is whether humans reach consistent and shared interpretations of these concepts.

Section 2 of 22

Identify from Highlighted Bird Images #1

You will see bird images with highlighted areas (blue = no focus, red = stronger focus).
 For each question, select the option that best matches the highlighted part in the following categories:

- Body Parts
- Color
- Action

Choose only **one option per category**. If it's not clear, select "None of Them / No Clear Option."



Sample 1.



Sample 2.



Sample 5.



Sample 3.



Sample 4.

Body Parts *

- Head
- Beak (Mouth)
- Wing
- Leg (Foot)
- Breast (Chest)
- Tail
- Back
- Bottom body part (Belly)
- Eye
- Whole body shape
- Background
- None of Them / No Clear Body Parts

Color

- Red
- Grey
- Beige
- Black
- Yellow
- Brown
- White
- Blue
- Green
- Colorful (Many colors)
- None of Them / No Clear Color

Action

- Flying
- Swimming
- Climbing
- Sitting on a branch / or somewhere
- None of Them / No Clear Action

Figure A-18. Example survey interface shown to participants. Each question presents bird images with highlighted regions (blue = low focus, red = stronger focus). Participants are asked to select the most appropriate label in three categories: Body Parts, Color, and Action, or choose "None of Them / No Clear Option" when no clear match is present.

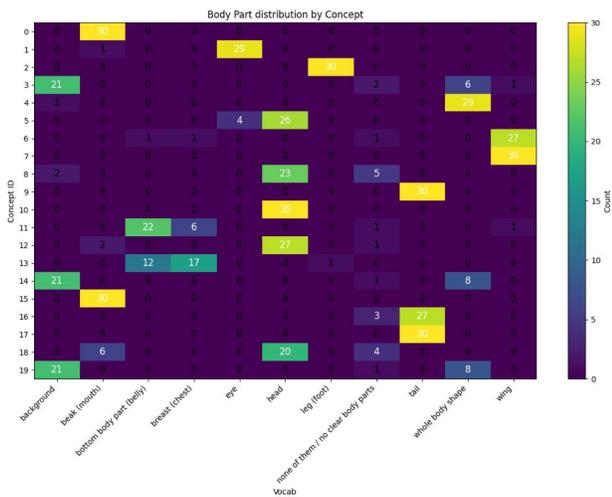


Figure A-19. Contingency heatmap of body part responses by concept example. Most concepts are dominated by a single vocabulary term (e.g., *beak*, *wing*, *tail*), producing highly concentrated response patterns. This strong alignment accounts for the extremely large χ^2 statistic ($\chi^2(190) = 4171.64$) and effect size (Cramer's $V = 0.82$).

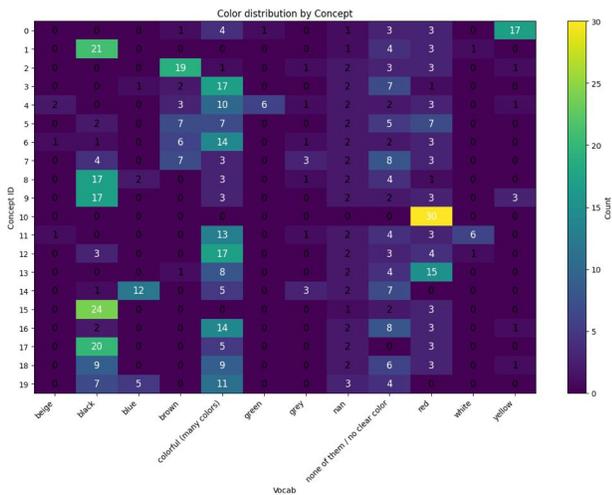


Figure A-20. Contingency heatmap of color responses by concept example. While some concepts are clearly associated with single terms (e.g., *black*, *blue*), others show more distributed choices across multiple terms (e.g., *grey*, *brown*). This greater variability explains the moderate effect size ($\chi^2(209) = 1308.05$, Cramer's $V = 0.41$).

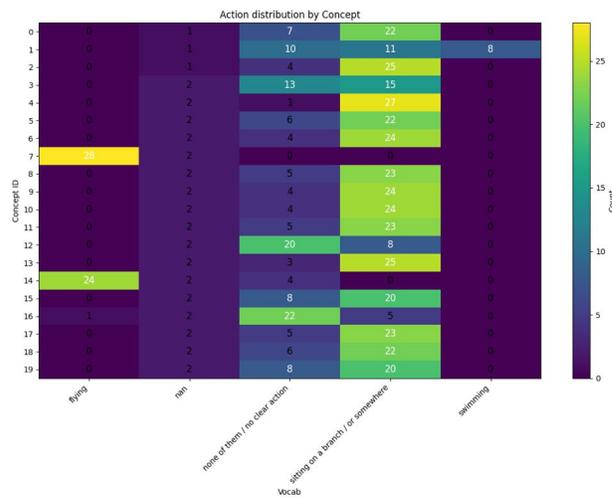


Figure A-21. Contingency heatmap of action responses by concept example. Many concepts are dominated by *sitting on a branch*, while others show clear associations with *flying* or *swimming*. Some ambiguity is present (e.g., overlap with *climbing*), producing a strong but not maximal association ($\chi^2(76) = 777.16$, Cramer's $V = 0.54$).

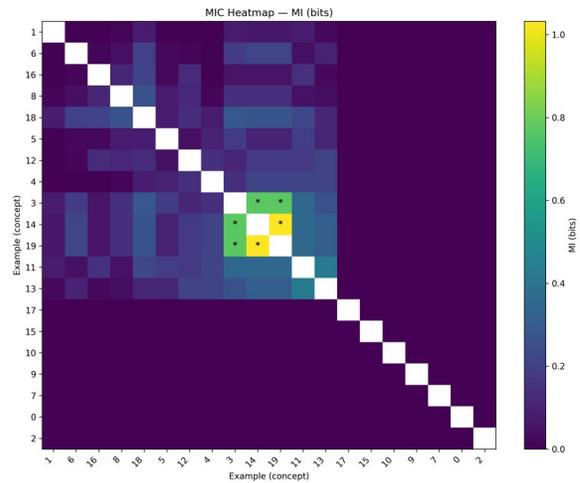


Figure A-22. MIC heatmap (Mutual Information in bits) between concept examples using **Body part** annotations. Most example pairs show low MI values (< 0.3 bits, dark cells), indicating that body part labels (e.g., *beak*, *wing*, *tail*) provide highly distinct information across examples. A localized high-MI cluster (yellow and green cells with significance markers) emerges, corresponding to examples that participants consistently labeled as *background* or *whole body parts*. This pattern highlights the dual role of body part vocabulary: it achieves strong global discriminability while also revealing interpretable redundancy in fallback categories.