
CoSy: Evaluating Textual Explanations of Neurons

Laura Kopf^{1,2} Philine Lou Bommer^{1,3} Anna Hedström^{1,3,4} Sebastian Lapuschkin⁴
Marina M.-C. Höhne^{3,5} Kirill Bykov^{1,2,3}
¹TU Berlin, Germany ²BIFOLD, Germany ³UMI Lab, ATB Potsdam, Germany
⁴Fraunhofer Heinrich-Hertz-Institute, Germany ⁵University of Potsdam, Germany
kopf@tu-berlin.de
{pbommer, ahedstroem, mhoehe, kbykov}@atb-potsdam.de
sebastian.lapuschkin@hhi.fraunhofer.de

Abstract

A crucial aspect of understanding the complex nature of Deep Neural Networks (DNNs) is the ability to explain learned concepts within their latent representations. While methods exist to connect neurons to human-understandable textual descriptions, evaluating the quality of these explanations is challenging due to the lack of a unified quantitative approach. We introduce COSY (Concept Synthesis), a novel, architecture-agnostic framework for evaluating textual explanations of latent neurons. Given textual explanations, our proposed framework uses a generative model conditioned on textual input to create data points representing the explanations. By comparing the neuron’s response to these generated data points and control data points, we can estimate the quality of the explanation. We validate our framework through sanity checks and benchmark various neuron description methods for Computer Vision tasks, revealing significant differences in quality. We provide an open-source implementation on GitHub¹.

1 Introduction

One of the key obstacles to the wider adoption of Machine Learning methods across various fields is the inherent opacity of modern Deep Neural Networks (DNNs)—in essence, we often lack an understanding of why these models make certain predictions. To address this problem, various explainability methods [1, 2] have been developed to make the decision-making processes of DNNs more understandable to humans. Explainability methods have broadened their focus from interpreting the decision-making of DNNs *locally*—for instance, interpreting specific inputs using saliency maps [3, 4, 5, 6]—to understanding the *global* behavior of models by analyzing individual model components and their functional purpose [7]. Following the latter global explainability approach, often referred to as *mechanistic interpretability* [8, 9, 10], some methods aim to describe the specific concepts neurons have learned to detect [11, 12, 13, 14, 15, 16], enabling analysis of how these high-level concepts influence network predictions.

A popular approach for explaining the functionality of latent representations of a network is to describe neurons using human-understandable textual concepts. A textual description is assigned to a neuron based on the concepts that the neuron has learned to detect or is significantly activated by. Over time, these methods have evolved from providing label-specific descriptions [11] to more complex compositional [12, 16] and open-vocabulary explanations [13, 15]. However, a significant challenge remains: the lack of a universally accepted quantitative evaluation measure for open-vocabulary neuron descriptions. As a consequence, different methods devised their own evaluation criteria, making it difficult to perform general-purpose, comprehensive cross-comparisons.

¹<https://github.com/lkopf/cosy>

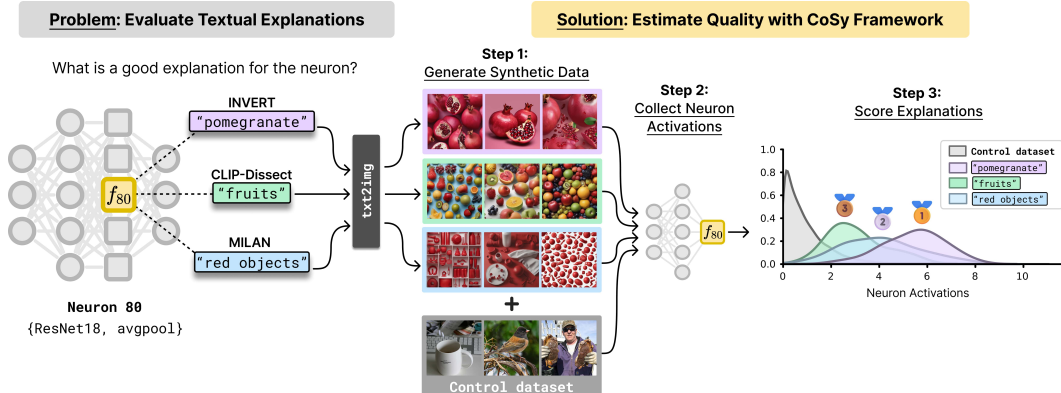


Figure 1: A schematic illustration of the CoSY evaluation framework for Neuron 80 in ResNet18’s avgpool layer. The current challenge lies in the absence of general-purpose, quantitative evaluation measures to benchmark textual explanations of neurons. To address this, we propose CoSY, a framework consisting of three steps: first, a generative model translates textual concepts into the visual domain, creating synthetic images for each explanation using a text-to-image model. Then, inference is performed on these synthetic images alongside a control image dataset to collect neuron activations. Finally, by comparing activations from the synthetic images with those from the control dataset, we quantitatively assess the quality of the textual explanation and compare results across different explanation methods. The implementation details of this example can be found in Appendix A.2.

In this work, we aim to bridge this gap by introducing Concept Synthesis (CoSY), the first automatic evaluation framework for textual explanations of neurons in Computer Vision (CV) models (illustrated in Figure 1). Our approach builds on recent advancements in Generative AI, which enable the generation of synthetic images that align with provided neuron explanations. We use a set of available text-to-image models to synthesize data points that are prototypical for specific target explanations. These data points allow us to evaluate how neurons differentiate between concept-related images and non-concept-related images combined in a control dataset. We summarize our contributions as below:

- (C1) We provide the first general-purpose, quantitative evaluation framework CoSY (Section 3) that enables the evaluation of individual or a set of textual explanation methods for CV models.
- (C2) In a series of sanity checks (Section 4), we analyze the choice of generative models and prompts for synthetic image generation, demonstrating framework reliability.
- (C3) We benchmark existing explanation methods (Section 5) and extract novel insights, revealing substantial variability in the quality of explanations. Generally, textual explanations for lower layers are less accurate compared to those for higher layers.

2 Related Works

Activation Maximization Activation Maximization is a commonly used methodology to understand what a neuron has learned to detect [17]. Such methods work by identifying input signals that trigger the highest activation in a neuron. This can be achieved synthetically, where an optimization process is employed to create the optimal input that maximizes the neuron’s activation [18, 19, 20], or naturally, by finding such inputs within a data corpus [21]. Activation Maximization has been employed for explaining latent representations of models [22, 23], including probabilistic models [24], detection of backdoor attacks [25] and spurious correlations [26]. However, one of the key limitations of this methodology lies in its inability to scale, as it relies on users to manually audit maximization signals.

Automatic Neuron Interpretation A more scalable alternative approach involves linking neurons with human-understandable concepts through textual descriptions. Network Dissection [11] (NetDissect) is a pioneering method in this field, associating convolutional neurons with a concept based on the Intersection over Union (IoU) of neuron activation maps and ground truth segmentation masks.

Table 1: Comparison of characteristics of neuron description methods. The columns (from left to right) represent the explanation method used, its textual output type (fixed-label, compositional, or open-vocabulary), the type of neuron targeted for analysis (convolutional, scalar, or predetermined), the target metric the method optimizes (IoU, WPMI, AUC, etc.), whether the method relies on auxiliary black-box models for finding or generating explanations (img2txt model, CLIP), and whether the explanation method is architecture-agnostic, meaning it can be applied to any CV model. For a more detailed description of each method, refer to Appendix A.1.

Method	Explanation	Neuron Type	Target	Black-Box Dependency	Architecture-Agnostic
NetDissect [11]	fixed-label	conv.	IoU	—	✓
CompExp [12]	compositional	conv.	IoU	—	✓
MILAN [13]	open-vocabulary	conv.	WPMI	img2txt model	✓
FALCON [14]	open-vocabulary	predetermined	avg. CLIP score	CLIP	—
CLIP-Dissect [15]	open-vocabulary	scalar	SoftWPMI	CLIP	✓
INVERT [16]	compositional	scalar	AUC	—	✓

Building on this, Compositional Explanations of Neurons (CompExp) [12] enhanced explanation detail by enabling the use of compositional concepts—i.e., concepts constructed with logical operators. MILAN [13] further expanded this by allowing for open-vocabulary explanations, permitting the generation of descriptions beyond predefined labels. INVERT [16] adopted a compositional concept approach, enabling explanations for general neuron types without the need for segmentation masks. It assigns compositional labels based on a neuron’s ability to distinguish concepts, using the Area Under the Receiver Operating Characteristic Curve (AUC). FALCON [14] and CLIP-Dissect [15] compute image-text similarity with a CLIP model [27] for the most activating images and their corresponding captions or concept sets. Each method defines its optimization criteria, lacking a unified consensus on what constitutes a good explanation. For detailed descriptions of the methods and their optimization objectives, please refer to Appendix A.1. An overview of the different techniques is illustrated in Table 1.

Prior Methods for Evaluation While significant effort has been made towards developing approaches and tools for evaluating *local* explanations [28, 29, 30], there has been relatively limited focus on evaluating *global* methods, in particular neuron description methods. Currently, to the best of our knowledge, there is no unified approach that allows for benchmarking across models and explanation methods. In their respective papers, the INVERT and CLIP-Dissect explanation methods evaluated the accuracy of their explanations by comparing the generated neuron labels with ground truth descriptions provided for neurons in the output layer of a network. However, this evaluation is limited to output neurons and fixed labels only. CLIP-Dissect additionally evaluates the quality of explanations by computing the Cosine Similarity in a sentence embedding space between the ground truth class name for each neuron and the explanation generated by the method. FALCON employs a human study conducted on Amazon Mechanical Turk to evaluate the concepts generated by the method. Participants are tasked with selecting the best explanation for each target feature from a selection of explanation methods, considering a given set of highly and lowly activating images. MILAN evaluates the performance of neuron labeling methods relative to human annotations using BERTScores [31]. While human studies are generally beneficial, the conventional setup can be misleading and may fail to fully capture the intended evaluation criteria, introducing potential biases. Typically, annotators describe the images that most strongly activate a neuron, and these descriptions are then compared to an automatic explanation. However, this approach primarily evaluates the alignment with the most activating images rather than the accuracy of the explanation in describing the neuron’s function. Moreover, these highly activating images may not accurately represent the neuron’s overall behavior, as they only reflect the maximum tail of the distribution.

3 Method

In the following section, we introduce COSY—a first automatic evaluation procedure for open-vocabulary textual explanations for neurons. We first define preliminary notations in Section 3.1, then describe COSY formally in Section 3.2.

3.1 Preliminaries

Consider a Deep Neural Network (DNN) represented by the function $g : \mathcal{X} \rightarrow \mathcal{Z}$, where $\mathcal{X} \subset \mathbb{R}^{h \times w \times c}$ denotes the input image domain and $\mathcal{Z} \subset \mathbb{R}^l$ represents the model’s output domain. We can view the model as a composition of two functions, $F : \mathcal{X} \rightarrow \mathcal{Y}$, and $L : \mathcal{Y} \rightarrow \mathcal{Z}$, such that $g = L \circ F$. Here $\mathcal{Y} \subset \mathbb{R}^{d \times w^* \times h^*}$, where $d \in \mathbb{N}$ is the number of neurons in the layer, and $w^*, h^* \in \mathbb{N}$ represent the width and height of the feature map, respectively. The function F , which we refer to as the *feature extractor*, can be chosen based on the layer of the model we aim to inspect. This could be an existing layer within the model or a concept bottleneck layer [32]. We refer to the i -th neuron within the layer as $f_i(\mathbf{x}) = F_i(\mathbf{x}) : \mathcal{X} \rightarrow \mathbb{R}^{w^* \times h^*}$. Within the scope of this paper, we refer to *explanation method* as an operator \mathcal{E} that maps a neuron to the textual description $s = \mathcal{E}(f_i) \in \mathcal{S}$, where \mathcal{S} is a set of potential textual explanations. The specific set of explanations depends on the implementation of the particular method (see Appendix A.1).

3.2 CoSy: Evaluating Open-Vocabulary Explanations

We assume that a good textual explanation for a neuron should provide a human-understandable description of an input that strongly activates the neuron. However, modern methods for explaining the functional purpose of neurons often provide open-vocabulary textual explanations, complicating the quantitative collection of natural data that represents the explanation. To address this issue, COSY utilizes recent advancements in generative models to synthesize data points that correspond to the textual explanation. The response of a neuron to a set of synthetic images is measured and compared to the neuron’s activation on a set of control natural images representing random concepts. This comparison allows for a quantitative evaluation of the alignment between the explanation and the target neuron.

The parameters of the proposed method include a control dataset $\mathbb{X}_0 = \{\mathbf{x}_1^0, \dots, \mathbf{x}_n^0\} \subset \mathcal{X}$, $n \in \mathbb{N}$, which consists of natural images representing the concepts on which the model was originally trained. Additionally, it incorporates a generative model p_M used for synthesizing images, along with a specified number of generated images $m \in \mathbb{N}$. The control dataset typically includes a balanced selection of validation class images. Given a neuron f_i and explanation $s \in \mathcal{S}$, COSY evaluates the alignment between the explanation and a neuron in three consecutive steps, which are illustrated in Figure 1.

1. **Generate Synthetic Data.** The first step involves generating synthetic images for a given explanation $s \in \mathcal{S}$, which we use as a prompt to a generative model p_M to create a collection of synthetic images, denoted as $\mathbb{X}_1 = \{\mathbf{x}_1^1, \dots, \mathbf{x}_m^1\} \sim p_M(\mathbf{x} | s)$. This collection consists of $m \in \mathbb{N}$ images, where m is adjustable as a parameter of the evaluation procedure.
2. **Collect Neuron Activations.** Given the control dataset \mathbb{X}_0 and the set of generated synthetic images \mathbb{X}_1 , we collect activations as follows:

$$\begin{aligned} \mathbb{A}_0 &= \{\sigma(f_i(\mathbf{x}_1^0)), \dots, \sigma(f_i(\mathbf{x}_n^0))\} \in \mathbb{R}^n, \\ \mathbb{A}_1 &= \{\sigma(f_i(\mathbf{x}_1^1)), \dots, \sigma(f_i(\mathbf{x}_m^1))\} \in \mathbb{R}^m, \end{aligned} \quad (1)$$

where $\sigma : \mathbb{R}^{w^* \times h^*} \rightarrow \mathbb{R}$ is an aggregation function for multi-dimensional neurons. Within the scope of our paper, we use Average Pooling as aggregation function

$$\sigma(\mathbf{y}) = \frac{1}{w^* h^*} \sum_{k \in [1, w^*], l \in [1, h^*]} \mathbf{y}_{k,l}, \quad \mathbf{y} \in \mathcal{Y} \subset \mathbb{R}^{w^* \times h^*}. \quad (2)$$

3. **Score Explanations.** The final step of the proposed method relies on the evaluation of the difference between neuron activations on the control dataset \mathbb{A}_0 and neuron activations given the synthetic dataset \mathbb{A}_1 . To quantify this difference, we utilize a *scoring function* $\Psi : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}$ to measure the difference between the distributions of activations.

In the context of our paper, we employ the following scoring functions:

- **Area Under the Receiver Operating Characteristic (AUC)**

AUC is a widely used non-parametric evaluation measure for assessing the performance of binary classification. In our method, AUC measures the neuron’s ability to distinguish

between synthetic and control data points

$$\Psi_{\text{AUC}}(\mathbb{A}_0, \mathbb{A}_1) = \frac{\sum_{a \in \mathbb{A}_0} \sum_{b \in \mathbb{A}_1} \mathbf{1}[a < b]}{|\mathbb{A}_0| \cdot |\mathbb{A}_1|}. \quad (3)$$

- **Mean Activation Difference (MAD)**

MAD is a parametric measure that quantifies the difference between the mean activation of the neuron on synthetic images and the mean activation on control data points

$$\Psi_{\text{MAD}}(\mathbb{A}_0, \mathbb{A}_1) = \frac{\frac{1}{m} \sum_{b \in \mathbb{A}_1} b - \frac{1}{n} \sum_{a \in \mathbb{A}_0} a}{\sqrt{\frac{1}{n-1} \sum_{a \in \mathbb{A}_0} (a - \bar{a})^2}}, \quad (4)$$

with mean control activation $\bar{a} = \frac{1}{n} \sum_{a \in \mathbb{A}_0} a$.

These two chosen metrics complement each other. AUC, being non-parametric and stable to outliers, evaluates the classifier’s ability to rank synthetic images higher than control images (with scores ranging from 0 to 1, where 1 represents a perfect classifier and 0.5 is random). On the other hand, MAD allows us to parametrically measure the extent to which images corresponding to explanations maximize neuron activation.

4 Sanity Checks

To ensure the reliability of our proposed evaluation measure, all steps within our framework need to be subject to sanity checks [33]. In this section, we analyze the following: (1) which generative models and prompts provide the best similarity to natural images, (2) whether the model’s behavior on synthetic and natural images differs for the same class, and (3) validating that COSY provides appropriate evaluation scores for true and random explanations, given a known ground truth class for the neuron.

4.1 Synthetic Image Reliability

One of the key features of COSY is its reliance on generative models to translate textual explanations of neurons into the visual domain. Thus, it is essential that the generated images reliably resemble the textual concepts. In the following section, we present an experiment where we varied several parameters of the generation procedure and evaluated the visual similarity between generated images and natural ones, focusing on concepts for which we have a collection of natural images.

For our analysis, we used only open-source and freely available text-to-image models, namely Stable Diffusion XL 1.0-base (SDXL) [34] and Stable Cascade (SC) [35]. We also varied the prompts for image generation. To measure the similarity between synthetic images and natural images corresponding to the same concept, we used Cosine Similarity (CS) in the CLIP embedding space with the CLIP-ViT-B/32 model [27]. We select a set of 50 random concepts from the 1,000 classes in the ImageNet validation dataset [36]. For each [concept] we use five different prompts and employ them with SDXL and SC models, generating 50 images per class. We then measure the CS between image pairs of the same class.

Figure 2 illustrates the comparison across all generative models and prompts in terms of CS of generated images to natural images of the same class. The results indicate that when using Prompt 5 as input to SDXL, the synthetic images show the highest similarity to natural images. The performance is generally best with the most detailed prompt (5) and closely aligns with prompts 1, 3, and 4. Moreover, SDXL appears to be slightly more effectively realizing detailed prompts than SC. As anticipated, the poorly constructed prompt (2) results in the lowest similarity to natural images for both models. To address prompt bias and dataset dependency, we compare the object-focused ImageNet with the scene-focused Places365 (see Appendix A.3). We find that close-up prompts work well for object-centric datasets, while general prompts like “photo of” are better for scene-based datasets. If not stated otherwise, for all following experiments, Prompt 5 together with SDXL model was employed for image generation.

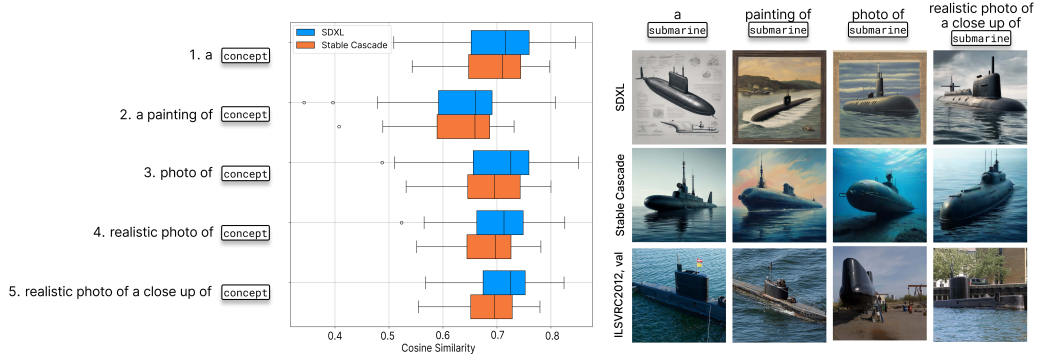


Figure 2: An overview of the impact of varying the prompt on the similarity between natural and synthetic images, using two text-to-image models. Left: average Cosine Similarity (CS) across all natural and synthetic images over all classes are reported. Higher CS values are better, indicating greater similarity between the images. Right: an illustration of the visual differences produced by the SDXL and SC models in response to diverse prompts for the concept “submarine”, and natural images from the ImageNet validation dataset [36]. Our results show that both SDXL and SC generate similar images, with SDXL generally being more closely aligned with natural images than SC.

4.2 Do Models Respond Differently to Synthetic and Natural Images?

Given the visual similarity between natural and synthetic images of the same class, we investigate whether CV models respond differently to these groups and if the activation differences indicate adversarial behavior. To this end, we employed four different models pre-trained on ImageNet: ResNet18 [37], DenseNet161 [38], GoogLeNet [39], and ViT-B/16 [40]. For each model, we randomly selected 50 output classes and generated 50 images per class using the class descriptions. We pass both synthetic and natural images through the models, collecting the activations of the output neuron corresponding to each class.

Figure 3 (a) illustrates the distributions of the MAD between synthetic and natural images for the same class across the 50 classes. Across all models, we observe that the median activation of synthetic images is slightly higher than that of natural images of the same class. However, this difference is small, given the 0 value lies within 1 standard deviation. We also illustrate the activations of neuron 504 in the ResNet18 output layer for the “coffee mug” class in Figure 3 (b). The results indicate a strong overlap in the neural response to both synthetic and natural images. While synthetic images activate the neuron slightly more, this does not constitute an artifactual behavior or affect our framework, which we demonstrate in the following experiment.




Figure 3: An overview of analyses performed to study the similarity between natural and synthetic images. From left to right: (a) an overview of MAD scores between synthetic and natural image activations of the output neuron’s ground truth classes for each model studied in this work, (b) activations collected for neuron 504 in ResNet18 for the class “coffee mug”, showcasing the difference between the natural and synthetic distributions and (c) examples of natural versus synthetic images. In both analyses, we observe a substantial overlap in the activations of synthetic and natural images, suggesting that the models respond similarly to both types of images.

Table 2: Comparison of true and random explanations on output neurons with known ground truth labels. This table presents the average quality scores (with standard deviations) for true explanations, derived from target class labels, and random explanations, derived from randomly selected synthetic image classes (including the target class), across four models pre-trained on ImageNet. Higher values are better. Our results consistently show high scores for true explanations and low scores for random ones.

Model	AUC (\uparrow)		MAD (\uparrow)	
	True	Random	True	Random
ResNet18	0.98 \pm 0.09	0.47 \pm 0.21	6.46 \pm 2.07	-0.11 \pm 0.79
DenseNet161	0.99 \pm 0.08	0.44 \pm 0.22	7.11 \pm 1.82	-0.19 \pm 0.73
GoogLeNet	0.99 \pm 0.07	0.48 \pm 0.23	7.74 \pm 2.14	-0.06 \pm 0.78
ViT-B/16	0.99 \pm 0.05	0.49 \pm 0.22	13.12 \pm 3.19	0.09 \pm 1.10

4.3 Random Baseline

A robust evaluation metric should reliably discern between random explanations resulting in low scores and non-random explanations resulting in high scores. To assess our evaluation framework regarding this requirement, we evaluated the results of the COSY evaluation by comparing the scores of ground truth explanations with those of randomly selected explanations.

Following the experimental setup in Section 4.2, we selected a set of 50 output neurons and compared the COSY scores of the ground truth explanations, given by the neuron label, with those of randomly selected explanations. The results, presented in Table 2, consistently demonstrate high scores for true explanations and low scores for random explanations. This experiment provides further evidence supporting the correctness of the proposed evaluation procedure. An additional experiment that excludes the target class from the control dataset is presented in Appendix A.4, along with an analysis of the robustness of the evaluation measure detailed in Appendix A.6.

5 Evaluating Explanation Methods

Within the scope of this section, we produce a comprehensive cross-comparison of various methods for the textual explanations of neurons. For this comparison, we employed models trained on different datasets, and we conducted our analysis on the latent layers of the models, where no ground truth is known.

5.1 Benchmarking Explanation Methods

In this section, we evaluated three recent textual explanation methods, namely MILAN, INVERT, and CLIP-Dissect. Our analysis involves six distinct models: four pre-trained on the ImageNet dataset [36] (ResNet18 [37], ResNet50 [37], ViT-B/16 [40], DINO ViT-S/8 [41]) and two pre-trained on the Places365 dataset [42] (DenseNet161 [38], ResNet50 [37]). The ImageNet dataset focuses on objects, whereas the Places365 dataset is designed for scene recognition. Consequently, we customized our prompts accordingly: Prompt 5 performs best for object recognition, while for scene recognition, we found that Prompt 3 is more effective. Therefore, Prompt 3 was utilized in the Places365 experiment.

For generating explanations with the explanation methods, we use a subset of 50,000 images from the training dataset on which the models were trained. For evaluation with COSY, we use the corresponding validation datasets the models were pre-trained on as the control dataset. Additionally, for CLIP-Dissect, we define concept labels by combining the 20,000 most common English words with the corresponding dataset labels. For INVERT we set the compositional length of the explanation as $L = 1$, where $L \in \mathbb{N}$. For more details on compute resources, refer to Appendix A.7.

Results of the evaluation can be found in Table 3. Overall, INVERT achieves the highest AUC scores across all models and datasets, except for DINO ViT-S/8 and ResNet18 applied to ImageNet, where CLIP-Dissect achieves a higher or similar score. Also across other models and datasets, CLIP-Dissect demonstrates consistently good results. Since INVERT optimizes AUC in explanation generation, it

Table 3: Benchmarking of neuron description methods, for neurons in the second to last layer across different models. Explanations are generated for a randomly selected set of 50 neurons, with average scores for both AUC and MAD reported alongside standard deviations. Higher values indicate better performance; **bold** numbers represent the highest scores.

Dataset	Model	Layer	Method	AUC (\uparrow)	MAD (\uparrow)
ImageNet	ResNet18	Avgpool	MILAN	0.61 \pm 0.23	0.69 \pm 1.35
			CLIP-Dissect	0.93\pm0.11	3.85\pm1.88
			INVERT	0.93\pm0.11	3.23 \pm 1.72
	ResNet50	Avgpool	MILAN	0.44 \pm 0.23	-0.08 \pm 0.72
			CLIP-Dissect	0.95 \pm 0.08	4.98\pm2.57
			INVERT	0.96\pm0.06	4.62 \pm 2.26
	ViT-B/16	Features	MILAN	0.53 \pm 0.19	0.12 \pm 0.76
			CLIP-Dissect	0.78 \pm 0.19	1.29 \pm 1.01
			INVERT	0.89\pm0.17	1.67\pm0.82
	DINO ViT-S/8	Layer 11	MILAN	0.59 \pm 0.21	0.37 \pm 0.91
			CLIP-Dissect	0.95\pm0.08	4.59\pm2.62
			INVERT	0.73 \pm 0.27	2.70 \pm 3.48
Places365	DenseNet161	Features	MILAN	0.56 \pm 0.28	0.44 \pm 1.30
			CLIP-Dissect	0.82 \pm 0.21	2.52\pm2.33
			INVERT	0.85\pm0.16	2.21 \pm 1.95
	ResNet50	Avgpool	MILAN	0.65 \pm 0.28	1.11 \pm 1.67
			CLIP-Dissect	0.92 \pm 0.11	3.73\pm2.39
			INVERT	0.94\pm0.08	3.54 \pm 1.99

may be biased towards AUC in our evaluation, leading to higher scores. MILAN generally performs poorly, with an average AUC below 0.65 across all tasks, indicating performance close to random guessing. MILAN tends to generate highly abstract explanations, such as “white areas”, “nothing” or “similar patterns”. These abstract concepts are particularly challenging for a text-to-image model to generate accurately, likely contributing significantly to the low scores of MILAN. Contrary to the AUC scores, the MAD scores suggest that CLIP-Dissect outperforms INVERT for convolutional neural networks applied to both datasets. Nonetheless, in these cases, INVERT concepts also achieve consistently high scores. Otherwise, we find similar outcomes for both metrics Ψ , with MILAN achieving poor scores in all experimental settings.

5.2 Explanation Methods Struggle to Explain Lower Layer Neurons

In addition to the general benchmarking, we aimed to study the quality of explanations for neurons in different layers of a model. Since it is well known that lower-layer neurons usually encode lower-level concepts [43], it is interesting to see whether explanation methods can capture the concepts these neurons detect. To investigate this, we examined the quality of explanations across layers 1 to 4 and the output layer of an ImageNet pre-trained ResNet18. In addition to three prior explanation methods, we included the FALCON method in our analysis. For more details on the implementation of FALCON see Appendix A.1.4, for additional results of the original FALCON implementation see Appendix A.8, and for qualitative examples and a discussion of lower-level concepts see Appendix A.9. For each layer, we randomly selected 50 neurons for analysis.

In Figure 4 we present the AUC and MAD results for all explanation methods across layers 1 to 4 and the output layer of ResNet18. While less pronounced for the AUC metric, in general, we find increasing scores for later layers across all methods and both metrics Ψ , which suggest higher concept quality in later layers. Furthermore, we find that similar to the benchmarking experiments, MILAN achieves lower scores across metrics. Both MILAN and FALCON consistently show lower performance, with AUC scores of 0.5 indicating random guessing. Nonetheless, we point out that these methods typically output semantically high-level concepts. Potentially, this is related to the inherent difficulty in describing low-level abstractions in natural language given their complexity (see Appendix A.10).

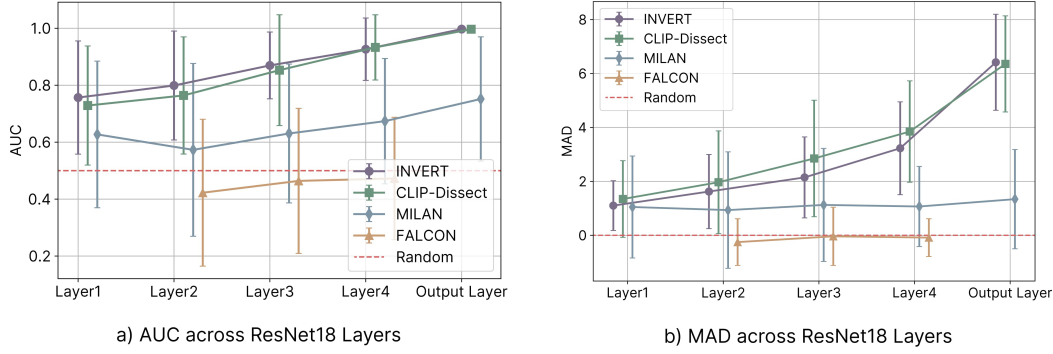


Figure 4: A comparison of how different explanation methods vary in their quality, as measured by (a) AUC and (b) MAD, across different layers in ResNet18. INVERT and CLIP-Dissect maintain high AUC and MAD scores across all layers, while MILAN and FALCON have lower scores. Overall, performance declines in the lower layers for all methods.

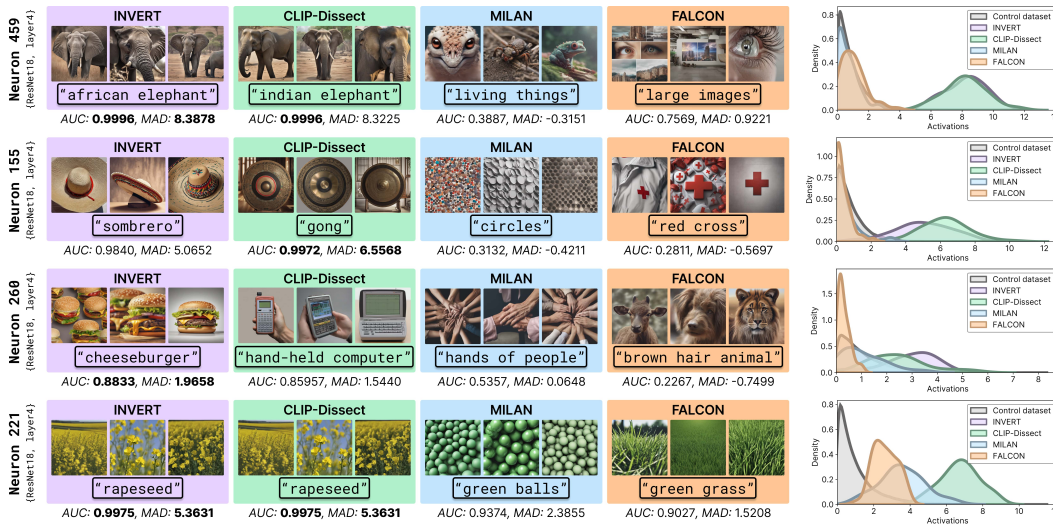


Figure 5: A qualitative example, of neuron explanations across four neurons. The first four panels include the textual explanation across INVERT, FALCON, CLIP-Dissect, and MILAN alongside three corresponding generated images. The respective AUC and MAD scores are reported below each panel. The last panel shows the activation distributions across 50 generated images for each method and the distribution of the control data.

5.3 What are Good Explanations?

In our approach, we propose that testing visual representations of textual explanations on neurons can provide insights into what constitutes good explanations. Building on this premise, we observe consistently high results from CLIP-Dissect and INVERT. The qualitative examples in Figure 5 demonstrate that their explanations share visually similar concepts (neurons 155 and 459) or even identical concepts (neuron 221) while both achieving high AUC and MAD scores. It is important to note that although INVERT performs slightly better in several tasks, the explanations are constrained to the input data labels. In contrast, CLIP-Dissect can generate descriptions from a broader selection of concepts, though its reliance on a black-box model reduces interpretability compared to INVERT.

There are instances, such as neuron 260 in Figure 5, where all explanations vary significantly. In these cases, we find that the explanation activation distributions of FALCON and MILAN often overlap with or even match the control dataset, providing the user with nearly random explanations. This observation aligns with our overall findings: both the AUC and MAD scores consistently reflect the

low performance of FALCON and MILAN explanations in the CoSY evaluation. Also, neurons 459 and 155 demonstrate the gap between consistently higher and lower-performing explanation methods.

6 Conclusion

In this work, we propose the first automatic evaluation framework for textual explanations of neurons. Unlike existing ad-hoc evaluation methods, we can now quantitatively compare different neuron description methods against each other and test, whether the given explanation describes the neuron accurately, based on its activations. We can evaluate the quality of individual neuron explanations by examining how accurately they align with the generated concept data points, without requiring human involvement.

Our comprehensive sanity checks demonstrate that CoSY guarantees a reliable explanation evaluation. In several experiments, we show that neuron description methods are most applicable for the last layers, where high-level concepts are learned. In these layers, INVERT and CLIP-Dissect provide high-quality neuron concepts, whereas MILAN and FALCON explanations have lower quality and can present close to random concepts, which might lead to wrong conclusions about the network. Thus, the results highlight the importance of evaluation when using neuron description methods.

Limitations The use of generative models involves a distinct set of limitations. For instance, text-to-image models may not include certain concepts in their training data, which can reduce their generative performance. This limitation, however, can often be addressed by analyzing the pre-training datasets and assessing model performance. Moreover, the model’s capabilities of generating highly abstract concepts like “white objects” can be limited. However, the challenges with abstract concepts also reflect the descriptive quality of the provided explanations—explanations should be inherently understandable to humans. In both cases, exploring more sophisticated, specialized, or constrained models may offer improvement.

Future Work Evaluation of non-local explanation methods is still a largely neglected research area, where CoSY plays an important yet preliminary part. In the future, we need additional, complementary definitions of explanation quality that extend our precise definition of AUC and MAD, e.g., that involve humans to assess plausibility [44] or evaluate explanation quality via the success of a downstream task [45]. Furthermore, we plan to extend the application of our evaluation framework to additional domains including NLP and healthcare. In particular, it would be interesting to analyze the quality of more recent autointerpretable explanation methods given by highly opaque, large language models (LLMs) [46, 9]. We believe that applying CoSY to healthcare datasets, where high-quality explanations are crucial, represents an impactful next step.

Acknowledgements

This work was partly funded by the German Ministry for Education and Research (BMBF) through the project Explaining 4.0 (ref. 01IS200551). Additionally, this work was supported by the European Union’s Horizon Europe research and innovation programme (EU Horizon Europe) as grant TEMA (101093003); the European Union’s Horizon 2020 research and innovation programme (EU Horizon 2020) as grant iToBoS (965221); and the state of Berlin within the innovation support programme ProFIT (IBB) as grant BerDiBa (10174498).

References

- [1] Wojciech Samek and Klaus-Robert Müller. Towards explainable artificial intelligence. *Explainable AI: interpreting, explaining and visualizing deep learning*, pages 5–22, 2019.
- [2] Feiyu Xu, Hans Uszkoreit, Yangzhou Du, Wei Fan, Dongyan Zhao, and Jun Zhu. Explainable AI: A brief survey on history, research areas, approaches and challenges. In *Natural Language Processing and Chinese Computing: 8th CCF International Conference, NLPCC 2019, Dunhuang, China, October 9–14, 2019, Proceedings, Part II 8*, pages 563–574. Springer, 2019.

- [3] Sebastian Bach, Alexander Binder, Grégoire Montavon, Frederick Klauschen, Klaus-Robert Müller, and Wojciech Samek. On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. *PLoS one*, 10(7):e0130140, 2015.
- [4] K Simonyan, A Vedaldi, and A Zisserman. Deep inside convolutional networks: visualising image classification models and saliency maps. In *Proceedings of the International Conference on Learning Representations (ICLR)*. ICLR, 2014.
- [5] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pages 618–626, 2017.
- [6] Daniel Smilkov, Nikhil Thorat, Been Kim, Fernanda Viégas, and Martin Wattenberg. Smoothgrad: removing noise by adding noise. *arXiv preprint arXiv:1706.03825*, 2017.
- [7] Chris Olah, Nick Cammarata, Ludwig Schubert, Gabriel Goh, Michael Petrov, and Shan Carter. Zoom in: An introduction to circuits. *Distill*, 5(3):e00024–001, 2020.
- [8] Kevin Ro Wang, Alexandre Variengien, Arthur Conmy, Buck Shlegeris, and Jacob Steinhardt. Interpretability in the Wild: a Circuit for Indirect Object Identification in GPT-2 Small. In *The Eleventh International Conference on Learning Representations*, 2022.
- [9] Steven Bills, Nick Cammarata, Dan Mossing, Henk Tillman, Leo Gao, Gabriel Goh, Ilya Sutskever, Jan Leike, Jeff Wu, and William Saunders. Language models can explain neurons in language models. <https://openaipublic.blob.core.windows.net/neuron-explainer/paper/index.html>, 2023.
- [10] Neel Nanda, Lawrence Chan, Tom Lieberum, Jess Smith, and Jacob Steinhardt. Progress measures for grokking via mechanistic interpretability. In *The Eleventh International Conference on Learning Representations*, 2022.
- [11] David Bau, Bolei Zhou, Aditya Khosla, Aude Oliva, and Antonio Torralba. Network Dissection: Quantifying Interpretability of Deep Visual Representations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [12] Jesse Mu and Jacob Andreas. Compositional Explanations of Neurons. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 17153–17163. Curran Associates, Inc., 2020.
- [13] Evan Hernandez, Sarah Schwettmann, David Bau, Teona Bagashvili, Antonio Torralba, and Jacob Andreas. Natural Language Descriptions of Deep Features. In *International Conference on Learning Representations*, 2022.
- [14] Neha Kalibhat, Shweta Bhardwaj, C. Bayan Bruss, Hamed Firooz, Maziar Sanjabi, and Soheil Feizi. Identifying Interpretable Subspaces in Image Representations. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett, editors, *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 15623–15638. PMLR, 23–29 Jul 2023.
- [15] Tuomas Oikarinen and Tsui-Wei Weng. CLIP-Dissect: Automatic Description of Neuron Representations in Deep Vision Networks. *International Conference on Learning Representations*, 2023.
- [16] Kirill Bykov, Laura Kopf, Shinichi Nakajima, Marius Kloft, and Marina MC Höhne. Labeling Neural Representations with Inverse Recognition. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- [17] Dumitru Erhan, Yoshua Bengio, Aaron Courville, and Pascal Vincent. Visualizing higher-layer features of a deep network. *University of Montreal*, 1341(3):1, 2009.
- [18] Chris Olah, Alexander Mordvintsev, and Ludwig Schubert. Feature visualization. *Distill*, 2(11):e7, 2017.

- [19] Anh Nguyen, Alexey Dosovitskiy, Jason Yosinski, Thomas Brox, and Jeff Clune. Synthesizing the preferred inputs for neurons in neural networks via deep generator networks. *Advances in neural information processing systems*, 29, 2016.
- [20] Thomas Fel, Thibaut Boissin, Victor Boutin, Agustin Martin Picard, Paul Novello, Julien Colin, Drew Linsley, Tom ROUSSEAU, Remi Cadene, Lore Goetschalckx, Laurent Gardes, and Thomas Serre. Unlocking Feature Visualization for Deep Network with MAgnitude Constrained Optimization. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- [21] Judy Borowski, Roland Simon Zimmermann, Judith Schepers, Robert Geirhos, Thomas SA Wallis, Matthias Bethge, and Wieland Brendel. Natural images are more informative for interpreting cnn activations than state-of-the-art synthetic feature visualizations. In *NeurIPS 2020 Workshop SVRHM*, 2020.
- [22] Gabriel Goh, Nick Cammarata, Chelsea Voss, Shan Carter, Michael Petrov, Ludwig Schubert, Alec Radford, and Chris Olah. Multimodal neurons in artificial neural networks. *Distill*, 6(3):e30, 2021.
- [23] Naoya Yoshimura, Takuya Maekawa, and Takahiro Hara. Toward understanding acceleration-based activity recognition neural networks with activation maximization. In *2021 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE, 2021.
- [24] Dennis Grinwald, Kirill Bykov, Shinichi Nakajima, and Marina MC Höhne. Visualizing the Diversity of Representations Learned by Bayesian Neural Networks. *Transactions on Machine Learning Research*, 2023.
- [25] Stephen Casper, Tong Bu, Yuxiao Li, Jiawei Li, Kevin Zhang, Kaivalya Hariharan, and Dylan Hadfield-Menell. Red teaming deep neural networks with feature synthesis tools. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- [26] Kirill Bykov, Mayukh Deb, Dennis Grinwald, Klaus Robert Muller, and Marina MC Höhne. DORA: Exploring Outlier Representations in Deep Neural Networks. *Transactions on Machine Learning Research*, 2023.
- [27] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021.
- [28] Chirag Agarwal, Satyapriya Krishna, Eshika Saxena, Martin Pawelczyk, Nari Johnson, Isha Puri, Marinka Zitnik, and Himabindu Lakkaraju. OpenXAI: Towards a Transparent Evaluation of Model Explanations. In *Thirty-sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2022.
- [29] Anna Hedström, Leander Weber, Daniel Krakowczyk, Dilyara Bareeva, Franz Motzkus, Wojciech Samek, Sebastian Lapuschkin, and Marina M.-C. Höhne. Quantus: An Explainable AI Toolkit for Responsible Evaluation of Neural Network Explanations and Beyond. *Journal of Machine Learning Research*, 24(34):1–11, 2023.
- [30] Anna Hedström, Leander Weber, Sebastian Lapuschkin, and Marina Höhne. Sanity Checks Revisited: An Exploration to Repair the Model Parameter Randomisation Test. In *XAI in Action: Past, Present, and Future Applications*, 2023.
- [31] Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. BERTScore: Evaluating Text Generation with BERT. In *International Conference on Learning Representations*, 2019.
- [32] Mert Yuksekgonul, Maggie Wang, and James Zou. Post-hoc Concept Bottleneck Models. In *The Eleventh International Conference on Learning Representations*, 2022.
- [33] Anna Hedström, Philine Bommer, Kristoffer K. Wickstrøm, Wojciech Samek, Sebastian Lapuschkin, and Marina M. C. Höhne. The Meta-Evaluation Problem in Explainable AI: Identifying Reliable Estimators with MetaQuantus. *Transactions on Machine Learning Research*, 2023.

- [34] Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. SDXL: Improving Latent Diffusion Models for High-Resolution Image Synthesis. In *The Twelfth International Conference on Learning Representations*, 2023.
- [35] Pablo Pernias, Dominic Rampas, Mats Leon Richter, Christopher Pal, and Marc Aubreville. Würstchen: An Efficient Architecture for Large-Scale Text-to-Image Diffusion Models. In *The Twelfth International Conference on Learning Representations*, 2023.
- [36] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015.
- [37] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [38] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017.
- [39] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1–9, 2015.
- [40] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *International Conference on Learning Representations*, 2020.
- [41] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2021.
- [42] Bolei Zhou, Agata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba. Places: A 10 million Image Database for Scene Recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017.
- [43] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *nature*, 521(7553):436–444, 2015.
- [44] David Cheng-Han Chiang and Hung-yi Lee. A Closer Look into Using Large Language Models for Automatic Evaluation. In Houda Bouamor, Juan Pino, and Kalika Bali, editors, *Findings of the Association for Computational Linguistics: EMNLP 2023, Singapore, December 6-10, 2023*, pages 8928–8942. Association for Computational Linguistics, 2023.
- [45] Satyapriya Krishna, Jiaqi Ma, Dylan Z Slack, Asma Ghandeharioun, Sameer Singh, and Himabindu Lakkaraju. Post Hoc Explanations of Language Models Can Improve Language Models. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- [46] Nicholas Kroeger, Dan Ley, Satyapriya Krishna, Chirag Agarwal, and Himabindu Lakkaraju. Are Large Language Models Post Hoc Explainers? *CoRR*, abs/2310.05797, 2023.
- [47] Christoph Molnar. *Interpretable Machine Learning*. 2 edition, 2022.
- [48] Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *International conference on machine learning*, pages 2048–2057. PMLR, 2015.
- [49] Thomas H Cormen, Charles E Leiserson, Ronald L Rivest, and Clifford Stein. *Introduction to algorithms*. MIT press, 2022.

- [50] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [51] Christoph Schuhmann, Richard Vencu, Romain Beaumont, Robert Kaczmarczyk, Clayton Mullis, Aarush Katta, Theo Coombes, Jenia Jitsev, and Aran Komatsuzaki. LAION-400M: Open Dataset of CLIP-Filtered 400 Million Image-Text Pairs. In *NeurIPS Workshop Datacentric AI, online (online), 14 Dec 2021 - 14 Dec 2021*, page 5 p., Dec 2021.
- [52] Tuomas Oikarinen and Tsui-Wei Weng. CLIP-Dissect: Automatic Description of Neuron Representations in Deep Vision Networks. In *The Eleventh International Conference on Learning Representations*, 2022.
- [53] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018.
- [54] George A Miller. WordNet: a lexical database for English. *Communications of the ACM*, 38(11):39–41, 1995.

A Appendix

A.1 Neuron Description Methods

Neuron description methods aim to provide insights into human-understandable concepts learned by DNNs, enabling a deeper understanding of their decision-making mechanisms. These methods provide textual descriptions for neurons in CV models. This creates a connection between the abstract representation of a concept by the neural network and a human interpretation. In general, a concept can be any abstraction, such as a color, an object, or even an idea [47]. Textual explanations of a neuron f_i can originate from various spaces depending on their generation process.

As defined in Section 3.1, we refer to *explanation method* as an operator \mathcal{E} that maps a neuron to the textual description $s = \mathcal{E}(f_i) \in \mathcal{S}$, where \mathcal{S} is a set of potential textual explanations. The specific set of explanations depends on the implementation of the particular method. We define the following subsets of textual descriptions $s \in \mathcal{S}$:

- \mathcal{C} represents the space of individual concepts,
- \mathcal{L} represents the space of logical combinations of concepts,
- \mathcal{N} represents the space of open-ended natural language concept descriptions.

These textual descriptions serve as explanations for f_i generated by explanation methods.

Examples for such explanation methods are MILAN [13], FALCON [14], CLIP-Dissect [15], and INVERT [16]. Figure 6 shows the general principle of how \mathcal{E} works. In Table 4 we outline the origin of textual descriptions and their corresponding set memberships for each \mathcal{E} .




Figure 6: Neuron Description Methods. A neuron f_i is selected, and a neuron description method \mathcal{E} is applied to generate a textual description s explaining f_i .

Table 4: Set Membership and Origin of Descriptions. Generated textual descriptions s have varying set membership and origin across all \mathcal{E} . These descriptions can originate from distinct spaces: individual concepts \mathcal{C} , logical combinations of concepts \mathcal{L} , and open-ended natural language concept descriptions \mathcal{N} . A labeled dataset refers to a collection of images paired with individual concept labels. Generated captions are produced by image-to-text models, such as Show-Attend-Tell [48]. An image caption dataset consists of image-caption pairs. A concept set consists of textual concept labels.

Method	Set	Origin
NetDissect	\mathcal{C}	labeled dataset
CompExp	\mathcal{L}	labeled dataset
MILAN	\mathcal{N}	generated caption
FALCON	\mathcal{N}	image caption dataset
CLIP-Dissect	\mathcal{C}	concept set
INVERT	\mathcal{L}	labeled dataset

A.1.1 NetDissect

Network Dissection (NetDissect) [11] is a method designed to explain individual neurons of DNNs, particularly convolutional neural networks (CNNs) within the domain of CV. This approach systematically analyzes the network’s learned concepts by aligning individual neurons with given semantic

concepts. To perform this analysis, annotated datasets with segmentation masks are required, where these masks label each pixel in an image with its corresponding object or attribute identity. The Broadly and Densely Labeled Dataset (Broden) [11] combines a set of densely labeled image datasets that represent both low-level concepts, such as colors, and higher-level concepts, such as objects. It provides a comprehensive set of ground truth examples for a broad range of visual concepts such as objects, scenes, object parts, textures, and materials in a variety of contexts.

A concept $s \in \mathcal{C} \subset \mathcal{S}$ is defined as a visual concept in NetDissect and is provided by the pixel-level annotated Broden dataset. Given a CNN and the Broden dataset as input, NetDissect explains a neuron f_i by searching for the highest similarity between concept image segmentation masks and neuron activation masks. Concept image segmentation masks are provided by the Broden dataset $B_s(\mathbf{x}) \in \{0, 1\}^{H \times W}$, where a value of 1 signifies the pixel-level presence of s , and 0 denotes its absence. Neuron activation masks are obtained by thresholding the continuous neuron activations of f_i into binary masks $A(\mathbf{x}) \in \{0, 1\}^{H \times W}$. Then the similarity δ_{IoU} between image segmentation masks and binary neuron masks can be evaluated using the Intersection over Union score (IoU) for an individual neuron within a layer:

$$\delta_{\text{IoU}}(f_i, s) = \frac{\sum_{\mathbf{x} \in \mathcal{X}} \mathbf{1}(B_s(\mathbf{x}) \cap A(\mathbf{x}))}{\sum_{\mathbf{x} \in \mathcal{X}} \mathbf{1}(B_s(\mathbf{x}) \cup A(\mathbf{x}))}. \quad (5)$$

The NetDissect method is optimized to identify the concept that yields the highest IoU score between binary masks and image segmentation masks. This can be formalized as:

$$\mathcal{E}_{\text{NetDissect}}(f_i) = \arg \max_{s \in \mathcal{C} \subset \mathcal{S}} \delta_{\text{IoU}}(f_i, s). \quad (6)$$

NetDissect is constrained to segmentation datasets, relying on pixel-level annotated images with segmentation masks. Moreover, its labeling capabilities are confined to concepts provided within a labeled dataset. Furthermore, only individual concepts can be associated with each neuron.

A.1.2 CompExp

To overcome the limitation of explaining neurons with only a single concept, the Compositional Explanations of Neurons (CompExp) method was later introduced [12], enabling the labeling of neurons with compositional concepts. The method obtains its explanations by merging individual concepts into logical formulas using composition operators AND, OR, and NOT. The formula length $L \in \mathbb{N}$ is defined beforehand. The initial stage of explanation generation is similar to NetDissect, a set of images is taken as input, and convolutional neuron activations are converted into binary masks. The explanations are constructed through a beam search algorithm [49], beginning with individual concepts and gradually building them into more complex logical formulas. Throughout the beam search stages, the existing formulas in the beam are combined with new concepts. These new formulas are measured by the IoU. The maximization of the IoU score is desired to get a high explanation quality.

The approach for obtaining δ_{IoU} is the same as in Equation 5. In contrast to NetDissect, the explanations can be a combination of concepts, where $s \in \mathcal{L} \subset \mathcal{S}$. The procedure of finding the best neuron description can be formalized as:

$$\mathcal{E}_{\text{CompExp}}(f_i) = \arg \max_{s \in \mathcal{L} \subset \mathcal{S}} \delta_{\text{IoU}}(f_i, s). \quad (7)$$

Similar to NetDissect, CompExp requires datasets containing segmentation masks and is primarily applicable to convolutional neurons.

A.1.3 MILAN

MILAN [13] is a method that aims to describe neurons within a DNN through open-ended natural language descriptions. First, a dataset of fine-grained human descriptions of image regions (Milanotations) is collected. These descriptions can be defined as concepts that are open-ended natural language descriptions, where $s \in \mathcal{N} \subset \mathcal{S}$. Given a DNN and input images $\mathbf{x} \in \mathcal{X}$, neuron masks $M(\mathbf{x}) \in \mathbb{R}^{H \times W \times C}$ are collected of highly activated image regions for f_i .

Two distributions are then derived: the probability $p(s|M(\mathbf{x}))$ that a human would describe an image region with s , and the probability $p(s)$ that a human would use the description s for any neuron. The probability $p(s|M(\mathbf{x}))$ is approximated with the Show-Attend-Tell [48] image-to-text model trained on the Milannotations dataset. Additionally, $p(s)$ is approximated with a two-layer LSTM language model [50] trained on the Milannotations dataset.

These distributions are then utilized to find a description that has high pointwise mutual information with $M(\mathbf{x})$. A hyperparameter $\lambda \in \mathbb{R}$ adjusts the significance of $p(s)$ during the computation of pointwise mutual information (PMI) between descriptions s and $M(\mathbf{x})$ sets, where the similarity δ_{WPMI} is weighted PMI (WPMI). The objective for WPMI is given by:

$$\delta_{\text{WPMI}}(s) = \log p(s|M(\mathbf{x})) - \lambda \log p(s). \quad (8)$$

MILAN aims to optimize high pointwise mutual information between s and $M(\mathbf{x})$ to find the best description for f_i :

$$\mathcal{E}_{\text{MILAN}}(f_i) = \arg \max_{s \in \mathcal{N} \subset \mathcal{S}} \delta_{\text{WPMI}}(f_i, s). \quad (9)$$

The requirement of collecting the curated labeled dataset, Milannotations, limits MILAN’s capabilities when applied to tasks beyond this specific dataset. Additionally, another drawback is the requirement for model training.

A.1.4 FALCON

The FALCON [14] explainability method has a similar approach to MILAN. Initially, it gathers the most highly activating images corresponding to a neuron. GradCam [5] is subsequently applied to identify highlighted features in these images, which are then cropped to focus on these regions. These cropped images, along with large captioning dataset LAION-400m [51] with concepts $s \in \mathcal{N} \subset \mathcal{S}$, are input to CLIP (Contrastive Language-Image Pre-training) [27], which computes the image-text similarity between the text embeddings of captions and the input cropped images. The top 5 captions are then extracted. Conversely, the least activating images are collected, and concepts are extracted and removed from the top-scoring concepts, ultimately yielding the explanation of the neuron.

The similarity $\delta_{\text{CLIPScore}}$ is obtained by calculating the CLIP confidence matrix, which is essentially a Cosine Similarity matrix. The aim is to find the maximum image-text similarity score between image embeddings and their closest text embeddings from a large captioning dataset:

$$\mathcal{E}_{\text{FALCON}}(f_i) = \arg \max_{s \in \mathcal{N} \subset \mathcal{S}} \delta_{\text{CLIPScore}}(f_i, s). \quad (10)$$

This restriction significantly narrows down the range of models suitable for analysis, setting it apart considerably from other explanation methods.

FALCON Implementation In its original implementation, FALCON restricts the set of “explainable neurons” based on specific parameters. These include the parameter $\alpha \in \mathbb{N}$, which determines the set of highly activating images for a given feature by requiring $\alpha > 10$. Additionally, it employs a threshold $\gamma \in \mathbb{R}$ for CLIP Cosine Similarity, with a set value of $\gamma > 0.8$.

These parameter settings significantly restrict the number of explainable neurons, resulting to fewer than 50 explainable neurons. This constraint prevents the necessary randomization for comparison with other methods. To address this, we set $\alpha = 0$ and $\gamma = 0$. However, for the original FALCON implementation, we retain the original settings of α and γ and calculate Ψ across all “explainable neurons”. In our experiments on ResNet18, FALCON can only be applied to layers 2 to 4.

A.1.5 CLIP-Dissect

CLIP-Dissect [52] is an explanation method that describes neurons in vision DNNs with open-ended concepts, eliminating the need for labeled data or human examples. This method integrates CLIP [27], which efficiently learns deep visual representations from natural language supervision. It utilizes both the image encoder and text encoder components of a CLIP model to compute the text embedding for

each concept $s \in \mathcal{C} \subset \mathcal{S}$ from a concept dataset and the image embeddings for the probing images in the dataset, subsequently calculating a concept-activation matrix.

The activations of a target neuron f_i are then computed across all images in the probing dataset \mathbf{X} . However, as this process is designed for scalar neurons, these activations are summarized by a function that calculates the mean of the activation map over spatial dimensions. The concept corresponding to the target neuron is determined by identifying the most similar concept s based on its activation vector. The most highly activated images are denoted as $\mathbf{X}_s \subset \mathbf{X}$.

SoftWPMI is a generalization of WPMI where the probability $p(\mathbf{x} \in \mathbf{X}_s)$ denotes the chance an image \mathbf{x} belongs to the example set \mathbf{X}_s . Standard WPMI corresponds to cases where $p(\mathbf{x} \in \mathbf{X}_s)$ is either 0 or 1 for all $\mathbf{x} \in \mathbf{X}$, while SoftWPMI relaxes this binary setting to real values between 0 and 1. The function can be formalized as:

$$\delta_{\text{SoftWPMI}}(s) = \log \mathbb{E}[p(s|\mathbf{X}_s)] - \lambda \log p(s). \quad (11)$$

The similarity function δ_{SoftWPMI} aims to identify the highest pointwise mutual information between the most highly activated images \mathbf{X}_s and a concept s . This optimization search is expressed as:

$$\mathcal{E}_{\text{CLIP-Dissect}}(f_i) = \arg \max_{s \in \mathcal{C} \subset \mathcal{S}} \delta_{\text{SoftWPMI}}(f_i, s). \quad (12)$$

A drawback of CLIP-Dissect lies in its interpretability; descriptions are generated by the CLIP model, which itself is challenging to interpret.

A.1.6 INVERT

Labeling Neural Representations with Inverse Recognition (INVERT) [16] shares the capability of constructing complex explanations like CompExp [12] but with the added advantage of not relying on segmentation masks and only needing labeled data. The method obtains its explanations by merging individual concepts into logical formulas using composition operators AND, OR, and NOT. It also exhibits greater versatility in handling various neuron types and is computationally less demanding compared to previous methods such as NetDissect [11] and CompExp [12]. Additionally, INVERT introduces a transparent metric for assessing the alignment between representations and their associated explanations. The non-parametric Area Under the Receiver Operating Characteristic (AUC) measure evaluates the relationship between representations and concepts based on the representation’s ability to distinguish the presence from the absence of a concept, with statistical significance. The probing dataset with the concept present is labeled as \mathbf{X}_1 , while the dataset without the concept is labeled as \mathbf{X}_0 .

The goal of INVERT is to identify the concept $s \in \mathcal{L} \subset \mathcal{S}$ that maximizes δ_{AUC} with the neuron f_i . Here, s can be a combination of concepts. The optimization process resembles that of CompExp, employing beam search [49] to find the optimal compositional concept. The top-performing concepts are iteratively selected until the predefined compositional length $L \in \mathbb{N}$ is reached.

The similarity measure δ_{AUC} is defined as:

$$\delta_{\text{AUC}}(f_i) = \frac{\sum_{\mathbf{x}_0 \in \mathbf{X}_0} \sum_{\mathbf{x}_1 \in \mathbf{X}_1} \mathbf{1}[f_i(\mathbf{x}_0) < f_i(\mathbf{x}_1)]}{|\mathbf{X}_0| \cdot |\mathbf{X}_1|}. \quad (13)$$

The objective of INVERT is to maximize the similarity δ_{AUC} between a concept s and the neuron f_i , which can be described as:

$$\mathcal{E}_{\text{INVERT}}(f_i) = \arg \max_{s \in \mathcal{L} \subset \mathcal{S}} \delta_{\text{AUC}}(f_i, s). \quad (14)$$

INVERT is constrained by the requirement of a labeled dataset and is computationally more expensive compared to CLIP-Dissect.

A.2 Schematic Illustration of CoSy Implementation Details

In the example shown in Figure 1, we used the default settings of the explanation methods to generate explanations for neuron 80 in the avgpool layer of ResNet18. For CLIP-Dissect, we used the 20,000

most common English words as the concept dataset and the ImageNet validation dataset [36] as the probing dataset. We employed Stable Diffusion XL 1.0-base (SDXL) [34] as the text-to-image model, using the prompt “realistic photo of a close up of [concept]” to generate concept images, with [concept] being replaced by the textual explanation from the methods. We generated 50 images per concept for 50 randomly chosen neurons from the avgpool layer of ResNet18. For evaluation, we also used the ImageNet validation dataset as the control dataset.

A.3 Prompt Bias Analysis

To address prompt bias and dataset dependency, we extended our analysis by comparing results on the object-focused ImageNet with the scene-focused Places365 dataset. Our results show a significant difference based on prompt selection and dataset: close-ups work well for object-centric datasets like ImageNet, while more general prompts, such as “photo of [concept]”, are more suitable for scene-based datasets like Places365 (see Figure 7).

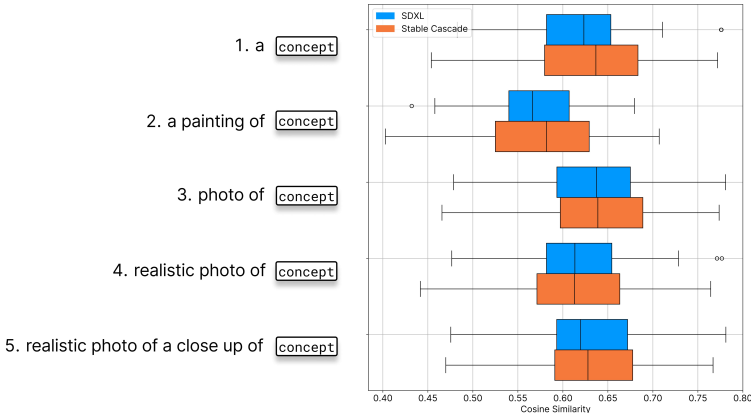


Figure 7: Cosine Similarity between synthetic and natural concept images for 50 concepts in the scene-based dataset Places365 [42], using different input prompts for the text-to-image models Stable Diffusion XL 1.0-base (SDXL) and Stable Cascade. The third prompt, “photo of [concept]”, performs best for generating scene images.

Additionally, we evaluated model performance using different similarity metrics across both datasets. Alongside Cosine Similarity (CS), we introduced two more distance measures: Learned Perceptual Image Patch Similarity (LPIPS) [53], which calculates perceptual similarity between two images and aligns well with human perception, using deep embeddings from a VGG model. Additionally, we included “Euclidean Distance” (ED) to capture the absolute differences in pixel values. The results of these evaluations are presented in Table 5.

A.4 Sanity Check Class Exclusion

In addition to the results presented in Table 2, we conducted the same experiment with the ground truth images excluded from the control dataset. The corresponding results, shown in Table 6, closely align with those obtained when the ground truth class was included. The findings are largely consistent with those obtained when the ground truth class was included.

A.5 Intraclass Image Similarity

In addition to comparing natural and synthetic images as in Section 4.1, we also analyze the intraclass distance to compare the similarity among synthetic images. Intraclass distance refers to the degree of diversity or dissimilarity observed within a set of images of the same class. It quantifies how much the individual images deviate from the average or central tendency of the image set. In this context, intraclass distance is desirable, reflecting how visual concepts can appear in natural images. Higher similarity scores indicate greater divergence from natural occurrences of concepts.

Cosine Similarity (CS), “Euclidean distance” (ED), and Learned Perceptual Image Patch Similarity (LPIPS) are commonly used metrics for measuring image similarity because they capture different

Table 5: Synthetic-to-Natural Image Similarity. This table illustrates the impact of varying parameters within our method pipeline. We evaluate five different prompts as input to two different text-to-image models, Stable Diffusion XL 1.0-base (SDXL) and Stable Cascade (SC). We randomly selected 50 from both ImageNet and Places365, using each class name as input to the prompt, denoted as [concept], generating 50 images per prompt for each model. We compute the average similarity metrics across all classes and report the standard deviation. Higher values of *CS* (Cosine Similarity) and lower values of *ED* (“Euclidean Distance”) and LPIPS (Learned Perceptual Image Patch Similarity) indicate greater similarity between synthetic and natural images. High similarity is desirable; **bold** values represent the highest scores for each prompt.

Dataset	Prompt	Text-to-image	Similarity Measure		
			<i>CS</i> (\uparrow)	<i>ED</i> (\downarrow)	LPIPS (\downarrow)
ImageNet	1. “a [concept]”	SDXL	0.70 \pm 0.08	8.0 \pm 1.09	0.72\pm0.03
		SC	0.69 \pm 0.08	8.19 \pm 1.02	0.74 \pm 0.04
	2. “a painting of [concept]”	SDXL	0.64 \pm 0.07	8.74 \pm 0.81	0.73 \pm 0.03
		SC	0.64 \pm 0.07	8.81 \pm 0.76	0.76 \pm 0.04
	3. “photo of [concept]”	SDXL	0.71 \pm 0.08	7.94\pm1.11	0.72\pm0.03
		SC	0.69 \pm 0.08	8.25 \pm 1.01	0.75 \pm 0.04
	4. “realistic photo of [concept]”	SDXL	0.70 \pm 0.08	8.07 \pm 1.09	0.72\pm0.03
		SC	0.69 \pm 0.08	8.34 \pm 0.99	0.73 \pm 0.03
	5. “realistic photo of a close up of [concept]”	SDXL	0.72\pm0.08	7.97 \pm 1.11	0.72\pm0.03
		SC	0.69 \pm 0.08	8.31 \pm 1.01	0.73 \pm 0.04
Places365	1. “a [concept]”	SDXL	0.62 \pm 0.09	9.02 \pm 1.04	0.71\pm0.02
		SC	0.63 \pm 0.09	8.96 \pm 1.04	0.72 \pm 0.03
	2. “a painting of [concept]”	SDXL	0.57 \pm 0.08	9.51 \pm 0.83	0.72 \pm 0.02
		SC	0.58 \pm 0.08	9.43 \pm 0.83	0.72 \pm 0.03
	3. “photo of [concept]”	SDXL	0.64\pm0.09	8.85\pm1.09	0.71\pm0.02
		SC	0.64\pm0.09	8.89 \pm 1.08	0.72 \pm 0.03
	4. “realistic photo of [concept]”	SDXL	0.62 \pm 0.09	9.08 \pm 1.02	0.71\pm0.02
		SC	0.62 \pm 0.08	9.21 \pm 0.98	0.72 \pm 0.03
	5. “realistic photo of a close up of [concept]”	SDXL	0.63 \pm 0.09	9.04 \pm 1.05	0.71 \pm 0.03
		SC	0.63 \pm 0.09	9.14 \pm 1.01	0.72 \pm 0.03

aspects of similarity and complement each other. We compute the average *CS*, *ED*, and LPIPS for each class and determine the overall class average. Table 7 provides a detailed overview of the results quantifying the similarity within synthetic images using *CS* and *ED*. When evaluating these results, it is important to note that high scores do not necessarily indicate optimal outcomes, as they suggest nearly identical images, which may lack intraclass distance. Conversely, very low scores imply significant differences among images, which might not capture the essence of the concept adequately. Ideally, we aim for somewhat similar yet slightly varied images representing the same class. The results show that the Stable Cascade (SC) model consistently achieves higher scores across all prompts compared to the Stable Diffusion XL 1.0-base (SDXL) model. Notably, it obtains the highest score for the two most elaborate prompts (4, 5). This indicates that the SC model tends to offer less intraclass distance in visually representing concepts.

Table 6: Comparison of true and random explanations on output neurons with known ground truth labels. This table presents the average quality scores (with standard deviation) for true explanations, derived from target class labels, and random explanations, derived from randomly selected synthetic image classes (excluding the target class), across four models pre-trained on ImageNet. Higher values are better. Our results consistently show high scores for true explanations and low scores for random ones.

Model	AUC (\uparrow)		MAD (\uparrow)	
	True	Random	True	Random
ResNet18	0.98 \pm 0.09	0.52 \pm 0.24	6.59 \pm 2.14	0.08 \pm 0.91
DenseNet161	0.99 \pm 0.08	0.52 \pm 0.24	7.33 \pm 1.91	0.04 \pm 0.84
GoogLeNet	0.99 \pm 0.07	0.49 \pm 0.24	8.01 \pm 2.28	-0.01 \pm 0.81
ViT-B/16	0.99 \pm 0.05	0.52 \pm 0.22	14.7 \pm 3.88	0.13 \pm 1.11

Table 7: Intraclass Image Similarity. This table illustrates the impact of varying parameters within our CoSY framework. We evaluate five different prompts as input to two different text-to-image models. A random selection of 50 ImageNet classes is made, with each class name used as input to the prompt, denoted as [concept], resulting in 50 images generated per prompt using a text-to-image model. We compute the average intraclass similarity across all classes and report the standard deviation. Higher *CS* and lower *ED* values indicate greater similarity between the images. In intraclass image similarity, neither excessively high nor excessively low scores are desirable.

Dataset	Prompt	Text-to-image	Similarity Measure		
			<i>CS</i> (\uparrow)	<i>ED</i> (\downarrow)	LPIPS (\downarrow)
ImageNet	1. "a [concept]"	SDXL	0.85±0.07	5.61±1.40	0.61±0.11
		SC	0.93±0.03	3.76±0.92	0.49±0.10
	2. "a painting of [concept]"	SDXL	0.87±0.05	4.94±1.13	0.62±0.10
		SC	0.93±0.03	3.67±0.87	0.54±0.10
	3. "photo of [concept]"	SDXL	0.84±0.06	5.65±1.34	0.62±0.10
		SC	0.92±0.04	4.03±0.99	0.52±0.10
	4. "realistic photo of [concept]"	SDXL	0.87±0.06	5.15±1.31	0.60±0.10
		SC	0.94±0.03	3.53±0.86	0.45±0.09
	5. "realistic photo of a close up of [concept]"	SDXL	0.89±0.04	4.82±1.16	0.60±0.10
		SC	0.94±0.03	3.58±0.86	0.50±0.09
Places365	1. "a [concept]"	SDXL	0.84±0.06	5.75±1.35	0.61±0.09
		SC	0.92±0.03	4.02±0.95	0.54±0.09
	2. "a painting of [concept]"	SDXL	0.88±0.04	4.89±1.05	0.61±0.09
		SC	0.93±0.03	3.79±0.84	0.56±0.09
	3. "photo of [concept]"	SDXL	0.85±0.06	5.61±1.27	0.61±0.10
		SC	0.93±0.03	3.94±0.92	0.54±0.09
	4. "realistic photo of [concept]"	SDXL	0.88±0.05	5.05±1.15	0.59±0.09
		SC	0.94±0.03	3.70±0.88	0.52±0.08
	5. "realistic photo of a close up of [concept]"	SDXL	0.87±0.05	5.25±1.24	0.60±0.09
		SC	0.94±0.03	3.68±0.91	0.54±0.09

A.6 Model Stability

In this experiment, our goal is to evaluate the stability of the image generation method employed, aiming to ensure consistent results within our CoSY framework. We achieve this by varying the seed of the image generator and observing the impact on image generation. We anticipate consistent image representations across different model initializations, thus ensuring the stability of our framework.

For our analysis, we utilize ResNet18 and focus on its output neurons, as the ground-truth labels associated with these neurons are known. We randomly select six classes s from the ImageNet validation dataset [36] and examine the corresponding f_i class output neurons using CoSY. Here, we exclude the s class from \mathbb{A}_0 and let \mathbb{A}_1 represent the s class. To ensure robustness, we initialize the text-to-image model across a random set of 10 seeds. Our analysis involves calculating the first (mean) and second moment (STD) using Ψ_{AUC} , as well as evaluating the intraclass image similarity (refer to Section A.5) within each synthetic ground truth class.

The results for our experiment, as shown in Table 8, demonstrate remarkably high AUC scores, indicating near-perfect detection of synthetic ground truth classes across all image model initializations. Furthermore, the standard deviation is exceptionally low, suggesting consistent image generation regardless of the chosen seed. The intraclass similarity values indicate a certain degree of distance in the generated images, indicating high similarity yet distinctiveness. This intraclass distance is desirable, ensuring that the images are not identical but share common characteristics.

These findings underscore the reliability and consistency of our image generation pipeline within our CoSY framework. The high stability of text-to-image generation across different seeds and the diversity of image similarity contribute to the robustness of our approach.

A.7 Compute Resources

For running the task of image generation for CoSY we use distributed inference across multiple GPUs with PyTorch Distributed, enabling image generation with multiple prompts in parallel. We run our script on three Tesla V100S-PCIE-32GB GPUs in an internal cluster. Generating 50 images for 3 prompts in parallel takes approximately 12 minutes.

Table 8: Model Stability. A comparison of various model initializations across 10 random seeds using SDXL. The results represent the average scores with standard deviations for each class, calculated across 10 seeds.

Concept	AUC (\uparrow)	Similarity Measure		
		CS (\uparrow)	ED (\downarrow)	LPIPS (\downarrow)
bulbul	0.9996 \pm 0.0002	0.91 \pm 0.03	3.99 \pm 0.66	0.65 \pm 0.06
china cabinet	0.9999 \pm 0.0001	0.89 \pm 0.04	5.00 \pm 0.90	0.58 \pm 0.04
leatherback turtle	0.9994 \pm 0.0001	0.91 \pm 0.04	4.65 \pm 0.87	0.63 \pm 0.04
beer bottle	0.9919 \pm 0.0038	0.80 \pm 0.08	6.79 \pm 1.41	0.69 \pm 0.08
half track	0.9998 \pm 0.0000	0.88 \pm 0.04	5.12 \pm 0.91	0.61 \pm 0.04
hard disc	1.0000 \pm 0.0001	0.90 \pm 0.05	4.64 \pm 1.17	0.60 \pm 0.05
Overall Mean	0.9984\pm0.0007	0.88\pm0.02	5.03\pm0.26	0.63\pm0.05

A.8 Additional Results for Method Comparison across ResNet18 Layers

Given that the original implementation of FALCON only provides results for their defined “explainable neurons” (see Appendix A.1.4), we included additional results comparing all methods based on this subset of neurons. Specifically, there are 7 explainable neurons in layer 2, 5 in layer 3, and 15 in layer 4. Figure 8 presents these results.

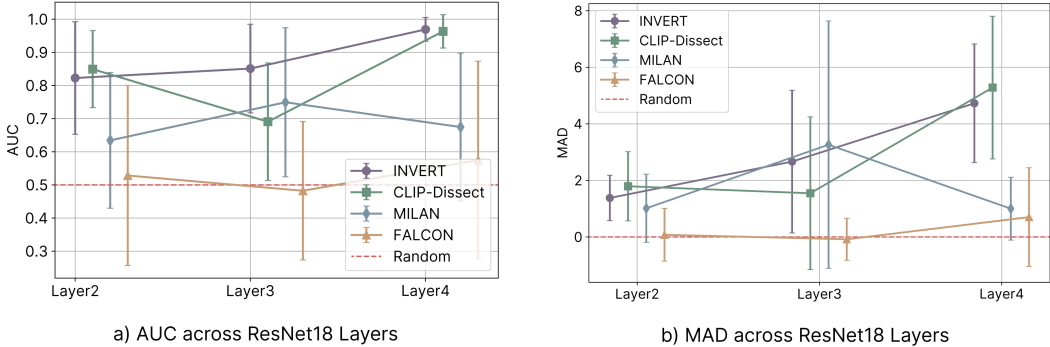


Figure 8: A comparison of explanation methods in ResNet18 shows that INVERT and CLIP-Dissect maintain high MAD scores across all layers, while MILAN and FALCON have lower scores. Overall, performance declines in the lower layers for all methods.

A.9 Qualitative Examples of Lower-Level Concepts

The performance of generative models may present challenges when dealing with abstract concepts. To assess this concern and provide a basis for further discussion, we have included an additional example of lower-layer concepts in Figure 9.

A.10 Concept Broadness

While COSY focuses on measuring the explanation quality, another open question is how broad or abstract are the concepts provided as textual explanations. This question of how specific or general an individual neuron is described by the explanation, might be relevant to different interpretability applications. For example, research fields where the user aims to deploy the same network for multiple tasks with varying image domains. In this case, describing a neuron’s more general concept such as “a round object” might be more informative than a more (domain-)specific concept such as “a tennis ball” for the network assessment. To provide insight into the broadness of concepts, we assessed whether the similarity between images generated based on the same concept changes from more general to more specific concepts.



Figure 9: Examples of various explanations from four different methods applied to the low-layer neurons in layer 2 of the ResNet18 model are provided. From the illustrated Feature Visualizations, we can observe that these neurons detect low-level abstractions. However, the methods studied generally fail to provide low-level explanations, instead attributing more complex explanations.

In our experiment, we define the broadness of a concept based on the number of hypernyms in the WordNet hierarchy [54]. The more specific a concept the larger the number of hypernyms. We choose two ImageNet classes (“ladybug”, “pug”) and generate 50 images for each concept as well as each hypernym of both concepts (with the most general concept being “entity”). Then, we measure the Cosine Similarity (CS) of all images generated based on the same concept. The box plot of the CS across both concepts and all hypernyms, in Figure 10 indicates that we do not find a correlation. Thus, we hypothesize that the chosen temperature of the diffusion model has a stronger effect on image similarity than the broadness of the prompt used for image generation.




Figure 10: The figure demonstrates the independence of the concept broadness measured by the number of hypernyms as defined in WordNet [54] to the inter-image similarity of corresponding generated images.

A.11 Prompt and Text-to-Image Model Comparison

Figure 11 showcases additional examples of synthetically generated images using both SDXL and SC across various prompts, highlighting the diversity and accuracy of concept representation.



Figure 11: Example images for “coffee mug” generated by the text-to-image models SDXL and SC across various prompts. (1) and (3) present examples of synthetic images with relatively low intraclass similarity and relatively high natural-to-synthetic similarity scores. (2) shows examples of synthetic images with the lowest similarity to natural images. (4) illustrates examples of synthetic images with the highest similarity to other synthetic images within the same class. (5) showcases examples of synthetic images with the highest similarity to natural images. (6) displays examples of natural images from the ImageNet validation dataset [36] belonging to the class “coffee mug”.

NeurIPS Paper Checklist

1. Claims

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

Answer: [Yes]

Justification: We state our contributions at the end of in Section 1 and refer to the sections where each contribution is demonstrated.

Guidelines:

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

2. Limitations

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

Answer: [Yes]

Justification: We mention limitations throughout the text and discuss them further in Paragraph 6

Guidelines:

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

3. Theory Assumptions and Proofs

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

Answer: [NA]

Justification: We do not include theoretical results.

Guidelines:

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

4. Experimental Result Reproducibility

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

Answer: [Yes]

Justification: We describe our algorithm, the datasets, models, methods, metrics, and parameters we use.

Guidelines:

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

5. Open access to data and code

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

Answer: [Yes]

Justification: We provided our code and use publicly available models, datasets and methods.

Guidelines:

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

6. Experimental Setting/Details

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

Answer: [Yes]

Justification: We specify and justify the hyperparameters, models, method restrictions, type and subsets of datasets.

Guidelines:

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

7. Experiment Statistical Significance

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

Answer: [Yes]

Justification: Our box plots and line plots include error bars.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).

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

8. Experiments Compute Resources

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

Answer: [Yes]

Justification: In Appendix A.7 we give detailed information on the type of computer resources, inference and running time.

Guidelines:

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

9. Code Of Ethics

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

Answer: [Yes]

Justification: We mention biases in models and methods we use.

Guidelines:

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

10. Broader Impacts

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

Answer: [Yes]

Justification: Our method provides more transparency of XAI methods and guides safer use.

Guidelines:

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

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

11. Safeguards

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

Answer: [Yes]

Justification: We mention the bias in and limitations of generative models.

Guidelines: We do not provide new data or a new pre-trained model.

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

12. Licenses for existing assets

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

Answer: [Yes]

Justification: All models, datasets, and methods used are cited.

Guidelines:

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

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

13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets? Answer: [\[Yes\]](#)

Justification: We provided our code and use publicly available models, datasets and methods and reference them.

Guidelines:

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

14. **Crowdsourcing and Research with Human Subjects**

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

Answer: [\[NA\]](#)

Justification: No human studies were conducted.

Guidelines:

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

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

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

Answer: [\[NA\]](#)

Justification: No crowdsourcing or research with human subjects was conducted.

Guidelines:

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