

CLIP-QDA: An Explainable Concept Bottleneck Model

Anonymous authors

Paper under double-blind review

Abstract

In this paper, we introduce an explainable algorithm designed from a multi-modal foundation model, that performs fast and explainable image classification. Drawing inspiration from CLIP-based Concept Bottleneck Models (CBMs), our method creates a latent space where each neuron is linked to a specific word. Observing that this latent space can be modeled with simple distributions, we use a Mixture of Gaussians (MoG) formalism to enhance the interpretability of this latent space. Then, we introduce CLIP-QDA, a classifier that only uses statistical values to infer labels from the concepts. In addition, this formalism allows for both **sample-wise** and **dataset-wise** explanations. These explanations come from the inner design of our architecture, our work is part of a new family of greybox models, combining performances of opaque foundation models and the interpretability of transparent models. Our empirical findings show that in instances where the MoG assumption holds, CLIP-QDA achieves similar accuracy with state-of-the-art methods CBMs. Our explanations compete with existing XAI methods while being faster to compute.

1 Introduction

The field of artificial intelligence is advancing rapidly, driven by sophisticated models like Deep Neural Networks (LeCun et al., 2015) (DNNs). These models find extensive applications in various real-world scenarios, including conversational chatbots (Ouyang et al., 2022), neural machine translation (Liu et al., 2020), and image generation (Rombach et al., 2021). Although these systems demonstrate remarkable accuracy, the process behind their decision-making often remains obscure. Consequently, deep learning encounters certain limitations and drawbacks. The most notable one is the lack of transparency regarding their behavior, which leaves users with limited insight into how specific decisions are reached. This lack of transparency becomes particularly problematic in high-stakes situations, such as medical diagnoses or autonomous vehicles.

The imperative to scrutinize the behavior of DNNs has become increasingly compelling as the field gravitates towards methods of larger scale in terms of both data utilization and number of parameters involved, culminating in what is commonly referred to as “foundation models” (Brown et al., 2020; Radford et al., 2021; Ramesh et al., 2021). These models have demonstrated remarkable performance, particularly in the domain of generalization, while concurrently growing more intricate and opaque. Additionally, there is a burgeoning trend in the adoption of multimodality (Reed et al., 2022), wherein various modalities such as sound, image, and text are employed to depict a single concept. This strategic use of diverse data representations empowers neural networks to transcend their reliance on a solitary data format. Nonetheless, the underlying phenomena that govern the amalgamation of these disparate inputs into coherent representations remain shrouded in ambiguity and require further investigation.

The exploration of latent representations is crucial for understanding the internal dynamics of a DNN. DNNs possess the capability to transform input data into a space, called latent space, where inputs representing the same semantic concept are nearby. For example, in the latent space of a DNN trained to classify images, two different images of cats would be mapped to points that are close to each other (Johnson et al., 2016). This capability is further reinforced through the utilization of multimodality (Akkus et al., 2023), granting access to neurons that represent abstract concepts inherent to multiple types of data signals.

A class of networks that effectively exploits this notion is Concept Bottleneck Models (CBMs) (Koh et al., 2020). CBMs are characterized by their deliberate construction of representations for high-level human-understandable concepts, frequently denoted as words. Remarkably, there is a growing trend in employing CLIP (Radford et al., 2021), a foundation model that establishes a shared embedding space for both text and images, to generate concept bottleneck models in an unsupervised manner.

Unfortunately, while CLIP embeddings represent tangible concepts, the derived values, often termed “CLIP scores” pose challenges in terms of interpretation. Furthermore, to the best of our knowledge, there is a notable absence of studies that seek to formally characterize CLIP’s latent space. The underlying objective here is to gain insights into how the pre-trained CLIP model organizes a given input distribution. Consequently, there is an opportunity to develop mathematically rigorous methodologies for elucidating the behavior of CLIP.

Then, our contributions are summarized as follows:

- We propose to represent the distribution of CLIP scores by a mixture of Gaussians. This representation enables a mathematically interpretable classification of images using human-understandable concepts.
- Utilizing the modeling approach presented in this study, we use Quadratic Discriminant Analysis (QDA) to classify the labels from the concepts, we name this method CLIP-QDA. CLIP-QDA demonstrates competitive performance when compared to existing CBMs based on CLIP. Notably, CLIP-QDA achieves this level of performance while utilizing a reduced set of parameters, limited solely to statistical values, including means, covariance matrices, and label probabilities.
- We propose two efficient and mathematically grounded XAI methods for model explanation, named **CLIP-QDA^{local}** and **CLIP-QDA^{global}**. These methods encompass both global and local assessments of why the model behaves. The global metric directly emanates from our Gaussian modeling approach, providing a comprehensive evaluation of CLIP-QDA’s performance. Additionally, our local metric draws inspiration from counterfactual analysis, furnishing insights into individual data points.
- We extend two established post-hoc XAI methods, LIME and SHAP, to formulate a novel XAI approach specifically tailored for CBMs. Departing from the conventional application of these methods, which typically produce explanations on the image level, CLIP-LIME and CLIP-SHAP generate explanations on the concept level allowed by CBMs.
- We propose a new evaluation protocol, specifically designed for the unique characteristics of CBMs to assess the effectiveness of explanations. This protocol includes a deletion metric, which examines faithfulness to the model, and a detection metric, which evaluates faithfulness to the data.

2 Background and Related Work

2.1 Contrastive Image Language Pre-training (CLIP)

CLIP (Radford et al., 2021) is a state-of-the-art model that can jointly process image and text inputs. The model was pre-trained on a large dataset of images and texts to learn to associate visual and textual concepts. Then, the capacity of CLIP to create a semantically rich encoding induced the creation of many emergent models in detection, few-shot learning, or image captioning.

The widespread adoption of CLIP stems from the remarkable robustness exhibited by its pre-trained model. Through training on an extensive multimodal dataset, such as Schuhmann et al. (2022), the model achieves impressive performance. Thus, on few- and zero-shot learning, for which it was designed, it obtains impressive results across a wide range of datasets. Notably, CLIP provides a straightforward and efficient means of obtaining semantically rich representations of images in low-dimensional spaces. This capability enables researchers and practitioners to divert the original use of CLIP to various other applications (Luo et al., 2022; Menon & Vondrick, 2022; Gabeff et al., 2023).




Figure 1: **Overview of our modeling method.** By considering the whole dataset CLIP scores \mathbf{z} as class conditioned distributions $\mathbf{Z} = [\mathbf{Z}^1 \dots \mathbf{Z}^N]$, we model the CLIP latent space as a mixture of Gaussians, allowing for mathematically grounded explanations.

2.2 CLIP-based Concept Bottleneck Models (CLIP-based CBMs)

The term Concept Bottleneck Model (CBM), as outlined in Koh et al. (2020), pertains to the implementation of a bottleneck reliant on human-specified concepts to execute a task, predominantly image classification. Consequently, the resultant algorithm inherently facilitates enhanced interpretability. While the term itself is relatively recent, it characterizes a lineage of methods that were employed in earlier research (Kumar et al., 2009; Lampert et al., 2009; Koh et al., 2020; Losch et al., 2019). However, despite the advantages offered by CBMs in terms of better understanding, early implementations encountered challenges stemming from the requirement for dedicated datasets. These datasets needed to encompass not only inputs and labels but also incorporate human-specified concepts for each sample.

In this context, the emergence of multimodal foundational models has opened up novel opportunities. Recent research (Yang et al., 2023; Oikarinen et al., 2023) has leveraged large language models to directly construct concepts from CLIP text embeddings, opening the door to a family of CLIP-based CBMs. Additionally, efforts have been made to create sparse CLIP-based CBMs (Panousis et al., 2023; Feng et al., 2023). Yan et al. (2023a) explore methods to achieve superior representations with minimal labels. Yuksekgonul et al. (2022) capitalize on the CLIP embedding spaces, considering concepts as activation vectors. Finally, Kim et al. (2023) build upon the idea of activation vectors to discover counterfactuals.

2.3 Explainable AI

According to Arrieta et al. (2020), we can define an explainable model as a computational model, that is designed to provide specific details or reasons to ensure clarity and ease of understanding regarding its functioning. In broader terms, an explanation denotes the information or output that an explainable model delivers to elucidate its operational processes.

The literature shows a clear distinction between non-transparent (or blackbox) and transparent (or whitebox) models. Transparent models are characterized by their inherent explainability. These models can be readily explained due to their simplicity and easily interpretable features. Examples of such models include linear regression (Galton, 1886), logistic regression (McCullagh, 2019), and decision trees (Quinlan, 1986). In contrast, non-transparent models are inherently non-explainable. This category encompasses models that could have been explainable if they possessed simpler and more interpretable features (Galton, 1886; Quinlan,

1986), as well as models that inherently lack explainability, including deep neural networks. The distinction between these two types of models highlights the trade-off between model complexity and interpretability (Arrieta et al., 2020), with transparent models offering inherent explainability while non-transparent models allow for better performance, but require the use of additional techniques for explanations, named post-hoc methods. Commonly used post-hoc methods include visualization techniques, such as saliency maps, which highlight the influential features in an image that contribute to the model’s decision-making. Within this category of methods, notable examples include approaches such as Grad-CAM (Selvaraju et al., 2017), which generates activation maps by computing the gradients of the output labels. Sensitivity analysis (Cortez & Embrechts, 2011) represents another avenue, involving the analysis of model predictions by varying input data. Sample-wise explanation techniques are also used to explain the model from a local simplification of the model around a point of interest (Ribeiro et al., 2016; Plumb et al., 2018). Finally, feature relevance techniques aim at estimating the impact of each feature on the decision (Lundberg & Lee, 2017).

In an endeavor to integrate the strengths of both black and whitebox models, the concept of greybox XAI has been introduced by Bennetot et al. (2022). These models divide the overall architecture into two distinct components. Initially, a blackbox model is employed to process high-entropy input signals, such as images, and transform them into a lower-entropy latent space that is semantically meaningful and understandable by humans. By leveraging the blackbox model’s ability to simplify complex problems, a whitebox model is then used to deduce the final result based on the output of the blackbox model. This approach yields a partially explainable model that outperforms traditional whitebox models while retaining partial transparency, in a unified framework.

Feature Attribution Methods, a category of techniques employed to address the complexity of DNNs’ output for explainability, strive to identify crucial features in the input. These methods leverage a mapping function to reduce input complexity. Notable examples in this family include DeepSHAP and KernelSHAP, as proposed by Lundberg & Lee (2017). This approach resonates with the concept of greybox models, wherein the input is initially simplified to be explainable by a transparent classifier. However, greyboxes differ from Feature Attribution Methods in that the mapping function is independent from the input under consideration.

3 A Greybox Concept Bottleneck Model: CLIP-QDA

3.1 General Framework

For our experimental investigations, we consider a general framework based on prior work on CLIP-based CBMs (Yang et al., 2023; Oikarinen et al., 2023). This framework consists of two core components. The first component centers on the extraction of multi-modal features, enabling the creation of connections between images and text. The second component encompasses a classifier head. A visual depiction of this process is presented in Figure 2.

We build upon CLIP DNN (Radford et al., 2021), which enables the creation of a multi-modal latent space through the fusion of image and text information. Rather than relying on a single text or prompt, we employ a set of diverse prompts, each representing distinct concepts. These concepts remain consistent across the dataset and are not subject to alterations. The purpose of CLIP’s representation is to gauge the similarities between each concept and an image, thereby giving rise to a latent space. To prevent ambiguity, we denote the resulting space of CLIP scores as the “CLIP latent space”, while the spaces generated by the text and image encoders are respectively referred to as the “CLIP text embedding space” and the “CLIP image embedding space”. Here, “CLIP score” denotes the value derived from a cosine distance computation between the image and text encodings.

The selection of concepts is guided by expert input and acts as a hyperparameter within our framework. For comprehensive examples of concept sets, please refer to Section A.2. Notably, there is no requirement for individual image annotation with these concepts. This is due to CLIP’s inherent design, which allows it to score concepts in a zero-shot manner.

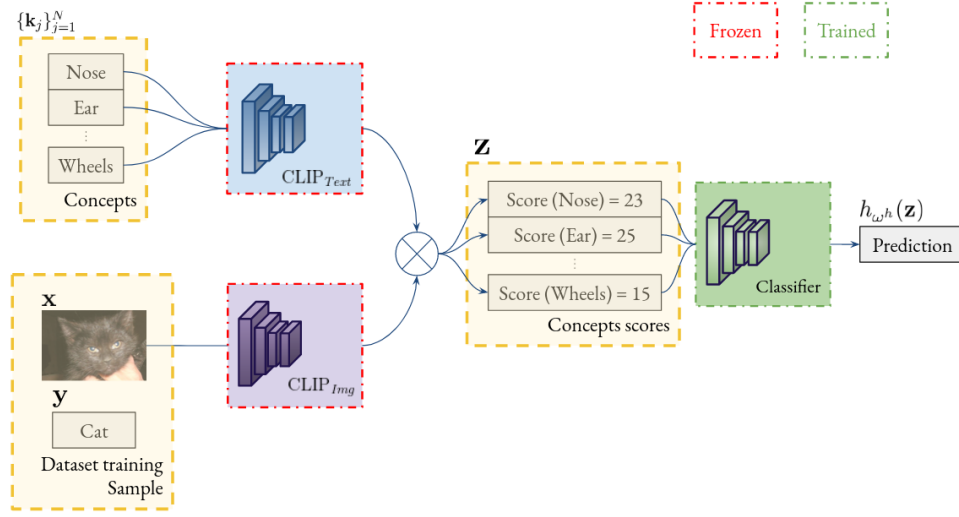


Figure 2: **Training procedure of the general framework.** First, CLIP scores \mathbf{z} are computed for each of the concepts $\{\mathbf{k}^j\}_{j=1}^N$, then a classifier $h_{\omega^h}(\cdot)$, with parameters ω^h is trained to classify the label from the concatenation of the CLIP scores.

Following the acquisition of the CLIP latent space, it is given as an input to a classifier head, which is responsible for learning to predict the class. Thanks to the low dimension of the latent space and the clear semantics of each component (concepts), it is possible to design simple and explainable classifiers.

3.2 CLIP Latent Space Analysis

3.2.1 Notations and formalism

Let us introduce the following notations used in the rest of the paper. X and Y represent two random variables (RVs) with joint distribution $(X, Y) \sim \mathcal{P}_{X,Y}$. A realization of this distribution is a pair (\mathbf{x}, y) that concretely represents one image and its label. In particular, y takes values in $\llbracket 1, C \rrbracket$, with $C \in \mathbb{N}$ the number of classes. From this distribution, we can deduce the marginal distributions $X \sim \mathcal{P}_X$ and $Y \sim \mathcal{P}_Y$. We can also describe for each class c , an RV $X_{Y=c} \sim \mathcal{P}_{X_{Y=c}}$ that represents the conditional distribution of images that have the class c .

Let $\{\mathbf{k}^j\}_{j=1}^N$ denote a set of $N \in \mathbb{N}$ concepts, where each \mathbf{k}^j is a character string representing the concept in natural language. We consider “CLIP’s DNN” to refer to the vector of its pre-trained weights, denoted as ω^g , and a function g that represents the architecture of the deep neural network (DNN). Given an image \mathbf{x} and a concept \mathbf{k}^j , the output of CLIP’s DNN is represented as $z^j = g_{\omega^g}(\mathbf{x}, \mathbf{k}^j)$. The projection in the multi-modal latent space of an image \mathbf{x} is the vector $\mathbf{z} = [g_{\omega^g}(\mathbf{x}, \mathbf{k}^1) \dots g_{\omega^g}(\mathbf{x}, \mathbf{k}^N)]$. We define Z^j as the random variable associated with the observation z^j . It should be noted that $\mathbf{Z} = [Z^1 \dots Z^N]$ is the random variable representing the concatenation of the CLIP scores associated with the N concepts. Furthermore, we denote the conditional distributions of \mathbf{z} having class c as $\mathbf{Z}_{Y=c} = [Z_{Y=c}^1 \dots Z_{Y=c}^N]$.

Finally, we define the classifier as a function $h_{\omega^h}(\mathbf{z})$ with parameters ω^h that, given a vector \mathbf{z} , concepts the predicted class.

3.2.2 Gaussian modeling of CLIP’s latent space

To analyze the behavior of the CLIP latent space, we conduct a thorough examination of the distribution of CLIP scores. To elucidate our modeling approach, we suggest to visualize a large set of samples from Z^j by observing the CLIP scores of an entire set of images taken from a toy example, which consists of images

representing only cats and cars (see the Cats/Cars dataset in Section 4.1). In this instance, the concept denoted by “ j ” corresponds to “Pointy-eared”.

In Figure 3, which illustrates the histogram of CLIP scores, we observe that the distribution exhibits characteristics that can be effectively modeled as a mixture of two Gaussians. The underlying intuition here is that the distribution Z^j represents two types of images: those without pointy ears, resulting in the left mode (low scores) of CLIP scores, and those with pointy ears, resulting in the right mode (high scores). Since this concept uniquely characterizes the classes – cats have ears but not cars – we can assign each mode to a specific class. This intuition is corroborated by the visualization of the distribution $Z_{Y=1}^j$ (Car) in violet and the distribution $Z_{Y=2}^j$ (Cat) in red. Since the extracted distributions exhibit similarities to normal distributions, we are motivated to describe \mathbf{Z} as a mixture of Gaussians. Yet, we also discuss the validity and limitations of this modeling approach in Section 4.3.




Figure 3: **Normalized histogram of scores Z^j specifically for the concept “Pointy-eared”.** On the left, we observe that the different classes can be modeled as weighted Gaussians. On the right, we show the resulting Gaussian mixture modeling.

Mathematically, the Gaussian prior assumption is equivalent to:

$$p(\mathbf{Z} = \mathbf{z} \mid Y = c) = \mathcal{N}(\mathbf{z} \mid \boldsymbol{\mu}_c, \boldsymbol{\Sigma}_c), \quad (1)$$

where $\boldsymbol{\Sigma}_c$ and $\boldsymbol{\mu}_c$ are the mean vectors and the covariance matrices, different for each class. Moreover, given the multinomial distribution of Y , with the notation $p_c = P(Y = c)$, we can model \mathbf{Z} as a mixture of Gaussians:

$$p(\mathbf{Z} = \mathbf{z}) = \sum_{c_i=1}^C p_{c_i} \mathcal{N}(\mathbf{z} \mid \boldsymbol{\mu}_{c_i}, \boldsymbol{\Sigma}_{c_i}). \quad (2)$$

3.3 CLIP Quadratic Discriminant Analysis (CLIP-QDA)

Based on the Gaussian distribution assumption described in Section 3.2.2, a natural choice for h_{ω^h} (the classifier in Figure 2) is the Quadratic Discriminant Analysis (QDA) as defined in Hastie et al. (2009). To compute it, we need to estimate the parameters $(\boldsymbol{\Sigma}_c, \boldsymbol{\mu}_c, p_c)$ of the probability distributions $\mathbf{Z}_{Y=c}$ and Y , which is done by computing the maximum likelihood estimators on the training data.

Subsequently, with the knowledge of the functions $p(\mathbf{Z} = \mathbf{z} \mid Y = c)$ and $p(Y = c)$, we can apply Bayes theorem to make an inference on $p(Y = c \mid \mathbf{Z} = \mathbf{z})$:

$$p(Y = c \mid \mathbf{Z} = \mathbf{z}) = \frac{p_c \mathcal{N}(\mathbf{z} \mid \boldsymbol{\mu}_c, \boldsymbol{\Sigma}_c)}{\sum_{c_i=1}^C p_{c_i} \mathcal{N}(\mathbf{z} \mid \boldsymbol{\mu}_{c_i}, \boldsymbol{\Sigma}_{c_i})}. \quad (3)$$

Then, the output of the QDA classifier can be described as:

$$h_{\omega^h}(\mathbf{z}) = \arg \max_c \frac{p_c}{(2\pi)^{N/2} |\Sigma_c|^{1/2}} e^{-\frac{1}{2}(\mathbf{z}-\mu_c)^T \Sigma_c^{-1} (\mathbf{z}-\mu_c)}. \quad (4)$$

In practice, we leverage the training data to estimate $\omega^h = (\Sigma_c, \mu_c, p_c)$, which enables us to bypass the standard stochastic gradient descent process, resulting in an immediate “training time”. Furthermore, this classifier offers the advantage of transparency, akin to the approach outlined by Arrieta et al. (2020), with its parameters comprising identifiable statistical values and its output values representing probabilities.

3.4 Explainable AI for Concept Bottleneck Models

Our CLIP-QDA model is founded upon a transparent probabilistic framework, hence, we have at our disposal a variety of statistical tools to explain the functioning of our classifier as illustrated in Figure 4. In this section, we present two distinct types of explanations: CLIP-QDA^{global}, offering a *global* perspective that sheds light on the classifier’s behavior across the entire dataset (refer to Section 3.4.1), and CLIP-QDA^{local}, providing a *local* explanation tailored to elucidate the model’s actions on individual samples (refer to Section 3.4.2).

Furthermore, we propose the adaptation of two well-established XAI post-hoc methods, LIME (Ribeiro et al., 2016) and SHAP (Lundberg & Lee, 2017), to the unique characteristics of CBMs. These methods are denoted as CLIP-SHAP and CLIP-LIME, respectively (see Section 3.4.3). An overview of the methods incorporated in our investigation is depicted in Figure 4.

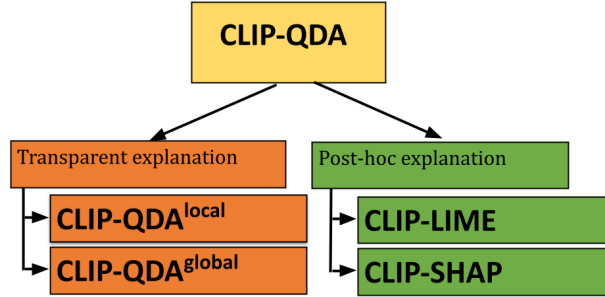


Figure 4: **The different source of explanation of CLIP-QDA.** Global (CLIP-QDA^{global}) vs. Local (CLIP-QDA^{local}) explanations, offering insights into classifier behavior across the dataset and on individual samples, respectively. Post-hoc explanations with CLIP-SHAP and CLIP-LIME offering using traditional techniques.

3.4.1 Dataset-wise explanation with CLIP-QDA^{global}

As we have access to priors that describe the distribution of each class, a valuable insight to gain an understanding of which concept our classifier aligns with is the measurement of distances between these distributions. Specifically, we focus on the conditional distributions of two classes of interest c_1 and c_2 , that we denote by $Z_{Y=c_1}^j$ and $Z_{Y=c_2}^j$. The underlying intuition behind measuring the distance between these distributions is that the larger the distance, the more the attribute j can differentiate between the classes c_1 and c_2 .

To accomplish this, we propose to use the Wasserstein-2 distance (Ramdas et al., 2017) as a metric for quantifying the separation between the two conditional distributions. It is worth noting that calculating the Wasserstein-2 distance can be a complex task in general. However, for Gaussian distributions, there exists a closed-form solution for computing the Wasserstein-2 distance. In addition, we sign the distance to keep the information of the position of c_1 relative to c_2 :

$$\tilde{W}_2(Z_{Y=c_1}^j, Z_{Y=c_2}^j) = \text{sign}([\mu_{c_1}]_{(j)} - [\mu_{c_2}]_{(j)}) ([\mu_{c_1}]_{(j)} - [\mu_{c_2}]_{(j)})^2 + \Lambda_{c_1, c_2}^j),$$

where $\Lambda_{c_1, c_2}^j = [\Sigma_{c_1}]_{(j,j)} + [\Sigma_{c_2}]_{(j,j)} - 2\sqrt{[\Sigma_{c_1}]_{(j,j)}[\Sigma_{c_2}]_{(j,j)}}$.

Note that the resulting value is no longer a distance since we lost the commutativity property. Examples of explanations based on this metric are given in Sections 4.5 and A.6. Also, an alternative way to produce dataset-wise explanations, oriented on example based explanations, are also available in Section A.3.

3.4.2 Sample-wise explanation with CLIP-QDA^{local}

One would like to identify the key concepts associated with a particular image that plays a pivotal role in achieving the task’s objective. To delve deeper into the importance and relevance of concepts in the decision-making process of the classifier, a widely accepted approach is to generate counterfactuals (Plumb et al., 2022; Luo et al., 2023; Kim et al., 2023). If a small perturbation of a concept score changes the class, the concept is considered important. We now formalize this mathematically.

Consider a pre-trained classifier denoted as $h_{\omega^h}(\cdot)$. In this context, ω^h represents the set of weights associated with the CLIP-QDA, specifically $\omega^h = (\Sigma_c, \mu_c, p_c)$. Given a score vector \mathbf{z} , we define counterfactuals as hypothetical values $\mathbf{z} + \epsilon_s^j$, ϵ_s^j being called the perturbation. This perturbation aims to be of minimal magnitude and is obtained by solving the following optimization problem:

$$\min \|\epsilon_s^j\|^2 \quad \text{s.t.} \quad h_{\omega^h}(\mathbf{z} + \epsilon_s^j) \neq h_{\omega^h}(\mathbf{z}). \quad (5)$$

The idea behind this equation is to find the minimal perturbation ϵ_s^j of the input \mathbf{z} that makes the classifier produce a different label than $h_{\omega^h}(\mathbf{z})$. However, in our case, two important restrictions are applied to ϵ_s^j :

1. **Sparsity:** for interpretability, we only change one attribute at a time, indicated by the index j . Then $\epsilon_s^j = [0, \dots, 0, \epsilon_s^j, 0, \dots, 0]$.
2. **Sign:** we take into account the sign $s \in \{-, +\}$ of the perturbation. Then, we separate the positive counterfactuals $\epsilon_+^j = [0, \dots, 0, \epsilon_+^j, 0, \dots, 0]$, $\epsilon_+^j \in \mathbb{R}^+$ and the negative counterfactuals. $\epsilon_-^j = [0, \dots, 0, \epsilon_-^j, 0, \dots, 0]$, $\epsilon_-^j \in \mathbb{R}^-$.

These two constraints are imposed to generate concise and, consequently, more informative counterfactuals. In this context, if a solution to equation 5, denoted as $\epsilon_{s,*}^j$, exists, it represents the minimal modification (addition or subtraction) to the coordinate j of the original vector \mathbf{z} that results in a change from $h_{\omega^h}(\mathbf{z})$ to $h_{\omega^h}(\mathbf{z} + \epsilon_{s,*}^j)$. Note that this approach allows for an explicit evaluation of the effect of an *intervention*, denoted as $do(\mathbf{Z} = \mathbf{z} + \epsilon_s^j)$ using a common notation in causal inference (Peters et al., 2017). Concretely, this emulates answers to questions of the form: “Would the label of my cat’s image change if I removed a certain amount of its pointy ears?”

Another important point to notice is that to obtain all possible counterfactuals, this equation must be solved for all concepts j and both signs s . A practical way to compute counterfactuals is given below.

Proposition 1. *Let us consider a pre-trained QDA classifier $h_{\omega^h}(\cdot)$ with parameters ω^h . Assume that the input data is drawn from the corresponding Gaussian Mixture model, as defined in 4, and that ϵ_s^j a perturbation with the above sparsity and sign restrictions. Then, there is a closed-form solution to problem 5, which is a function of the parameters $(\Sigma_c, \mu_c, p_c)_{c=1}^C$.*

The proof and expression are given in Appendix A.4. We illustrate the behavior of our classifier and our local metric with a toy example which consists of two Gaussians ($C = 2$) among two concepts Z_1 and Z_2 ($N = 2$). We find the counterfactuals for both signs following the first concept ($\epsilon_{-,*}^1$ and $\epsilon_{+,*}^1$). Results are presented in Figure 5.

It is worth noting that these counterfactual values are initially expressed in CLIP score units, which may not inherently provide meaningful interpretability. To mitigate this limitation, we introduce scaled counterfactuals, denoted as $\epsilon_{s,*}^{j,scaled}$, obtained by dividing each counterfactual by the standard deviation associated with its respective distribution:

$$\epsilon_{s,*}^{j,scaled} = \frac{\epsilon_{s,*}^j}{\sqrt{[\Sigma_c]_{(j,j)}}}. \quad (6)$$

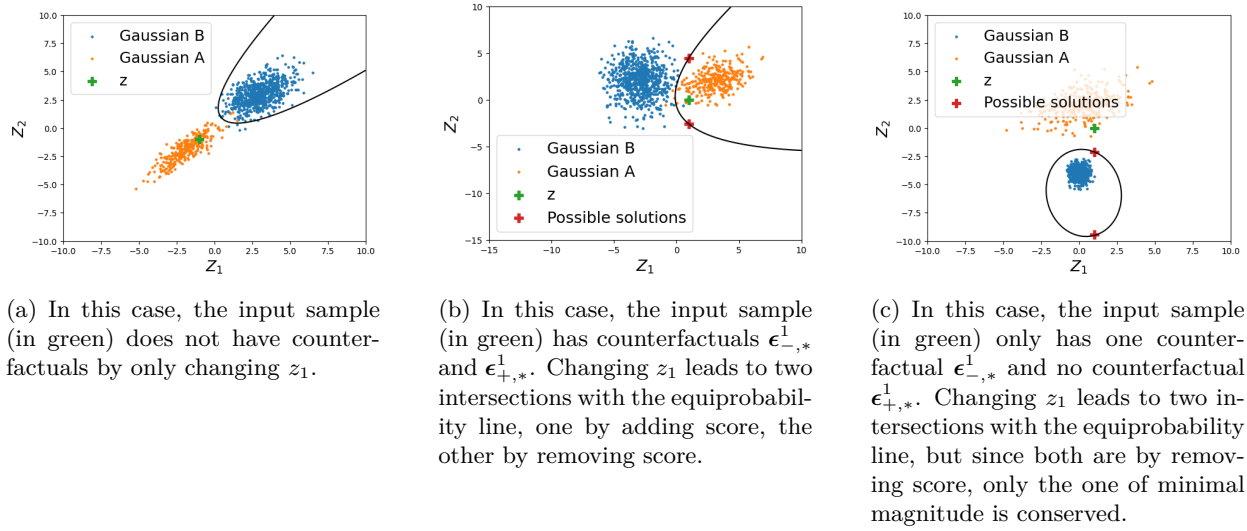


Figure 5: **Visualization of the counterfactuals in the two Gaussians toy example.** Samples of the two distributions are plotted in blue and orange. The equiprobability line is plotted in black.

Then, the value of each counterfactual can be expressed as “the addition (or subtraction) of standard deviations in accordance to $\mathbf{Z}_{Y=h_{\omega^h}(\mathbf{z})}$ that changes the label”. Examples of such explanations are given in Sections 4.5 and A.6.

3.4.3 CLIP-LIME and CLIP-SHAP

CLIP-LIME. To adapt LIME to the operation of CBMs, we begin with the image input x . From this input, we calculate the projection in the latent space $\mathbf{z} = [g_{\omega^g}(\mathbf{x}, \mathbf{k}^1) \dots g_{\omega^g}(\mathbf{x}, \mathbf{k}^N)]$. Subsequently, following the LIME method, we train a surrogate model to approximate h_{ω^h} in the vicinity of \mathbf{z} by training it on a dataset comprised of perturbed inputs around \mathbf{z} . Finally, the explanation is derived from the importance weights of the resulting surrogate model.

CLIP-SHAP. To adapt SHAP to CBMs, we also consider the projection $\mathbf{z} = [g_{\omega^g}(\mathbf{x}, \mathbf{k}^1) \dots g_{\omega^g}(\mathbf{x}, \mathbf{k}^N)]$. To compute statistical values relevant to CLIP-SHAP, we incorporate the projections of all the images in the training set. Given that the number of concepts can be high, we employ the Kernel version of SHAP (Kernel SHAP) to reduce the computational cost. The resulting explanation comprises the Shapley values associated with each of the concepts.

4 Experiments

4.1 Experimental Setup

We evaluate our methods on ImageNet (Deng et al., 2009), PASCAL-Part (Donadello & Serafini, 2016), MIT Indoor Scenes dataset (Quattoni & Torralba, 2009), **MonuMAI** (Lamas et al., 2021) and Flowers102 (Nilsback & Zisserman, 2008). In addition to these well-established datasets, we introduce a custom dataset named Cats/Dogs/Cars dataset. To construct this dataset, we concatenated three widely recognized datasets, namely, the Kaggle Cats and Dogs Dataset (Cukierski, 2013) and the Stanford Cars Dataset (Krause et al., 2013). Subsequently, we filtered the dataset to contain images of white and black animals and cars exclusively. This curation resulted in six distinct subsets: “Black Cars”, “Black Dogs”, “Black Cats”, “White Cars”, “White Dogs”, “White Cats”. The primary objective of this dataset is to facilitate experiments under conditions of substantial data bias, such as classifying white cats when the training data has only encountered white dogs and black cats. Specifically, we refer to two distinct scenarios (Table 1): one containing cats and cars of both colors (referred to as the unbiased setup) and the other one with only black cats and white cars

Table 1: **Cats/Dogs/Cars compositions used in our study**

	Complete dataset	Biased setup	Unbiased setup
Composition	Black Cats, White Cats, Black Cars, White Cats, Black Dogs, White Dogs	Black Cats, White Cars	Black Cats, White Cats, Black Cars, White Cars
Num samples	6436	2536	4031

(referred to as the biased setup). When none of these scenarios are referred (for example in the *Del* values of Table 4 and 5), the complete dataset is used. In its final form, the dataset comprises 6,436 images. The dataset will be shared after the reviewing process.

In addition to CLIP-QDA, we also test our method with linear regression as a classifier from CLIP scores, frequently named linear probe classifier. The objective here is to provide a comparison with existing CLIP-based CBMs (Yang et al., 2023; Yan et al., 2023a;b).

4.2 CBM Quantitative Evaluation Process

Assessing the quality of explanations has long been a challenge, given its subjective nature. While quantitative evaluation processes exist for evaluating performance in a general setup (Hedström et al., 2023), they may not be well-suited to the unique characteristics of CBMs. Consequently, we present a novel method designed to evaluate XAI solutions tailored to CBMs. This method comprises two metrics: the Deletion metric, gauging faithfulness to the model, and the Detection metric, quantifying faithfulness to the data.

Deletion metric. The deletion metric is adapted from the methodology introduced by Petsiuk et al. (2018). The procedure involves taking each sample from the test set and nullifying (i.e., setting to the average value of the score across classes) a certain number, N_{null} , of concepts. We nullify the concepts based on their importance order as determined by each explanation method. If nullifying the concepts leads to a significant decrease in performance, we consider it a successful selection of concepts that influenced the classifier’s decision. The intuition behind this idea is that if the concept is important, its absence will result in a loss of performance. Given the different values of accuracy obtained by nullifying N_{null} concepts, we note $Acc(N_{null})$ we deduce the deletion score *Del* by computing the area under the curve of $Acc(N_{null})$. The interest is to probe the ability to correctly order the important concept in the explanation:

$$Del = \frac{1}{Acc(0)} \sum_{i=1}^{N_{max}} \frac{Acc(i-1) + Acc(i)}{2}. \quad (7)$$

Here, N_{max} represents a hyperparameter that defines the maximum number of deletions considered. The selection of this hyperparameter is critical: it balances between capturing the metric’s capacity to fit the model and ensuring that the computed inputs remain plausible. A default value of 9 is assigned to maintain this equilibrium. Note that we normalize the result by the maximum accuracy $Acc(0)$ to allow better comparisons among different classifiers.

The evaluation framework is structured into two setups: the first one uses the concept set outlined in Table 6 (referred to as Set 1), while the second one uses an equivalent number of concepts randomly chosen from a dictionary of words (referred to as Set 2).


Detection metric. The detection metric evaluates the model’s ability to identify relevant concepts. For each sample s , a set of ground truth concepts \mathcal{S}_s to detect is defined by an oracle. . Then, we construct a set of concepts \mathcal{T}_s consisting of the top $|\mathcal{S}_s|$ concepts in the explanation to be tested. The resulting detection metric is the average ratio of agreements among all samples in the test dataset:

$$Det = \frac{1}{S} \sum_{s=1}^S \frac{|\mathcal{S}_s \cap \mathcal{T}_s|}{|\mathcal{S}_s|}, \quad (8)$$


where S is the total number of samples in the dataset. The selection of ground truth concepts depends on the available options within each dataset, as elaborated in Section 4.6.

4.3 Gaussian Prior Hypothesis

In this section, we investigate to which extent the Gaussian prior hypothesis (Equation 1) holds. To assess this, we use Chi-Square Q-Q plots (Chambers, 2018; Mahalanobis, 2018), a normality assessment method adapted to multidimensional data. We display Chi-Square Q-Q plots on the conditional distribution ($\mathbf{Z} \mid Y = c$), with data sourced from PASCAL-Part and class $c = \text{“aeroplane”}$. First, we compute a set of concepts adapted to the output class, following the procedure described in Appendix A.2. Subsequently, we perform the same experiment with a set of words specifically dedicated to the class of interest (Figure 6a). In addition, we show a visualization with randomly chosen concepts from the PASCAL-Part set of concepts, as depicted in Figure 6b.



(a) CBM with a set of concepts related to the label “Aeroplane”: [Winged, Jet engines, Tail fin, Fuselage, Landing gear]



(b) CBM with a set of concepts unrelated to the label “Aeroplane”: [Furry, Equine, Container or pot, Saddle or seat, Multi-doored]

Figure 6: Multivariate Q-Q plot for a Gaussian fit for the PASCAL-Part images that have the label “Aeroplane”.

Notably employing a less precise set of concepts can introduce disturbances, as evidenced in the observations. As indicated in Section 3.2.2, the ambiguity associated with certain concepts, such as “Multi-doored”, can lead to bimodal distributions (an airplane having one, multiple, or no doors). In Figure 7, we show the histogram of the clip scores z of the images that have the class “aeroplane”. Compared to the histogram of less ambiguous cases (like in 3), we observe that the histogram presents anomalies, especially around the mean.

Additionally, we performed a similar experiment with larger sets of concepts. We selected random subsets containing 10, 15, and 20 concepts from the PASCAL-Part set listed in Table 6. The outcomes are displayed in Figure 8. In this scenario, it becomes obvious that the Gaussian assumption is increasingly violated as the number of concepts grows. Indeed, as the number of concepts increases, the likelihood of encountering ambiguous concepts in each sample significantly rises, which undermines the feasibility of modeling the data as an unimodal Gaussian distribution.

4.4 Assessing the Accuracy of XAI Methods

Comparison with other classifiers. In this section, we undertake a comparative evaluation of the performance of our CLIP-QDA classifier in contrast to other classifiers commonly employed in contemporary




Figure 7: **Histogram and Gaussian fitting of CLIP scores z for images that have the class “aeroplane” of the attribute “Multi-doored”.**




Figure 8: **Multivariate Q-Q plot for a Gaussian fit for the PASCAL-Part images that have the label “Aeroplane”.**

research. Precisely, we compare our method to the use of CLIP as a zero-shot classifier, by making predictions based on the class that produces the maximum CLIP score with the image embedding. We also use the Yan et al. (2023a) method, which consists of using a linear layer as the classifier of the CBM (see Figure 2). Additionally, we use Yang et al. (2023)’s method, involving the use of a class-concept matrix as the classifier of the CBM. Our assessment is based on the computation of accuracy metrics across three diverse datasets: PASCAL-Part, ImageNet, MIT scenes, and MonuMAI. A comprehensive description of the procedure for concept set acquisition is provided in Appendix A.2. All the CLIP-based CBMs are trained on the same sets of concepts denoted in Table 6. Furthermore, for contextualization, we include in the table the test accuracy of ResNet (He et al., 2016) and ViT (Dosovitskiy et al., 2021) trained in a supervised manner by using images as inputs. We also include Greybox XAI (Bennetot et al., 2022), X-NeSyL (Díaz-Rodríguez et al., 2022) and Morales Rodríguez et al. into our study, which offer explainability but require training data with annotated concepts. It is important to note that these methods, which require training from images, have significantly longer training times.

Results of these experiments are available in Table 2. The image encoder used in all our experiments based on CLIP-based CBMs is *ViT - L/14@336px* provided from OpenAI’s public repository.

Our findings reveal that fine-tuning using CBM, either as a linear or a QDA probe, significantly improves performance, as evidenced by the increase in accuracy compared to using CLIP as a zero-shot classifier.

Table 2: **Test set accuracy.** On the top, are methods that require full training on images, and on the bottom, CBMs. Greybox XAI, X-NeSyL and Morales Rodríguez et al. requiring concept annotations, results on MIT scenes and ImageNet are not available.

	Method	PASCAL-Part \uparrow	MIT scenes \uparrow	MonuMAI \uparrow	ImageNet \uparrow
Non XAI	Resnet 50 (He et al., 2016)	0.84	0.86	0.95	0.80
	ViT-L 336px (Dosovitskiy et al., 2021)	0.95	0.94	0.94	0.85
CBMs	<i>Greybox XAI</i> (Bennetot et al., 2022)	0.88	-	0.94	-
	<i>X-NeSyL</i> (Díaz-Rodríguez et al., 2022)	0.82	-	0.90	-
	<i>Morales Rodríguez et al.</i>	0.86	-	0.98	-
	<i>CLIP (zero-shot)</i> (Radford et al., 2021)	0.81	0.63	0.52	0.76
	<i>LaBo</i> (Yang et al., 2023)	0.83	0.75	0.74	0.69
	<i>Yan et al. (2023b),</i> <i>Yan et al. (2023a)</i>	0.91	0.77	0.77	0.81
	<i>CLIP-QDA (ours)</i>	0.90	0.81	0.89	0.60

This improvement is particularly pronounced on datasets dedicated to specialized tasks, such as **MonuMAI**. Additionally, CLIP CBMs tend to achieve performances comparable to networks trained from raw images, making these models appealing for image classification due to their reduced training cost in both time and resources, as well as their interpretability. Notably, CLIP-QDA demonstrates competitive performance, when compared to linear probe techniques. However, a slight decrease in performance is observed, especially in the PASCAL-Part dataset. This decline could be attributed to the use of a larger set of concepts specifically for this dataset, posing challenges to the Gaussian assumption and potentially affecting the applicability of our classifier.

Influence of the number of concepts. To assess the influence of the number of concepts C on the accuracy, we conduct experiments on the PASCAL-Part dataset. These experiments involve testing accuracy for both QDA and linear probe with concept sets of different lengths, all generated following the methodology described in Section A.2. As seen in Figure 9, CLIP-QDA performs better than linear probe when the number of concepts is relatively low. In contrast, the linear probe outperforms CLIP-QDA as the number of concepts increases. This observation aligns with the insights gained from the discussion on Gaussian modeling in Section 4.3, where a higher number of concepts challenges the grounding assumptions of CLIP-QDA.

4.5 Assessing the Dataset-wise Interpretability of XAI Methods

4.5.1 Evaluation of CLIP-QDA^{local} and CLIP-QDA^{global} on the Cats/Dogs/Cars dataset

To evaluate the relevance of our explanations, we present a toy example from our Cats/Dogs/Cars dataset by constructing a CBM consisting of the concepts Table 6, plus the concepts “Black” and “White”. Next, we display the 10 most influential concepts according to our global metric (the top 10 concepts that have the highest Wasserstein distance) in both biased and unbiased setups. Results are presented in Figure 10.




Figure 9: **Accuracy of the classifiers on PASCAL-Part for different concept sizes.** QDA refers to the use of CLIP-QDA as a classifier. Linear probe refers to the use of a linear layer as a classifier.

We can observe that the concepts “Black” and “White” hold significantly higher importance in the biased setup, indicating that the classifier is likely to be biased about these concepts. This shows that our global explanation method has the potential for detecting biases in datasets (Tommasi et al., 2017). Additional explanation samples for various use cases are available in Appendix A.6.




Figure 10: **Global explanation on subsets of Cats/Cars.** Here, c_1 =“Cat” and c_2 =“Car”. Positive values indicate concepts that are more prevalent in cat images than car images, while negative values indicate concepts that are more common in car images compared to cat images. We display here only the top 10 concepts that have the highest Wasserstein distance (the concept “Black” is positioned 15th in the unbiased setup).

We show an application of our local metric within the framework of our biased setup, as previously described. Subsequently, we feed an image of a white cat into our classifier (Figure 11). It is noteworthy that the image is misclassified as a car. Our local metric demonstrates sensitivity to the dataset’s color bias, corroborating the warning issued by the global explanation.

4.6 Assessing the Sample-wise Interpretability of XAI Methods

The objective of this subsection is to present our method within the context of current XAI methods applied to CBMs. We aim to compare CLIP-QDA^{global}, CLIP-QDA^{local}, CLIP-LIME and CLIP-SHAP with existing approaches such as SHAP (Lundberg & Lee, 2017), LIME (Ribeiro et al., 2016), LaBo (Yang et al., 2023),

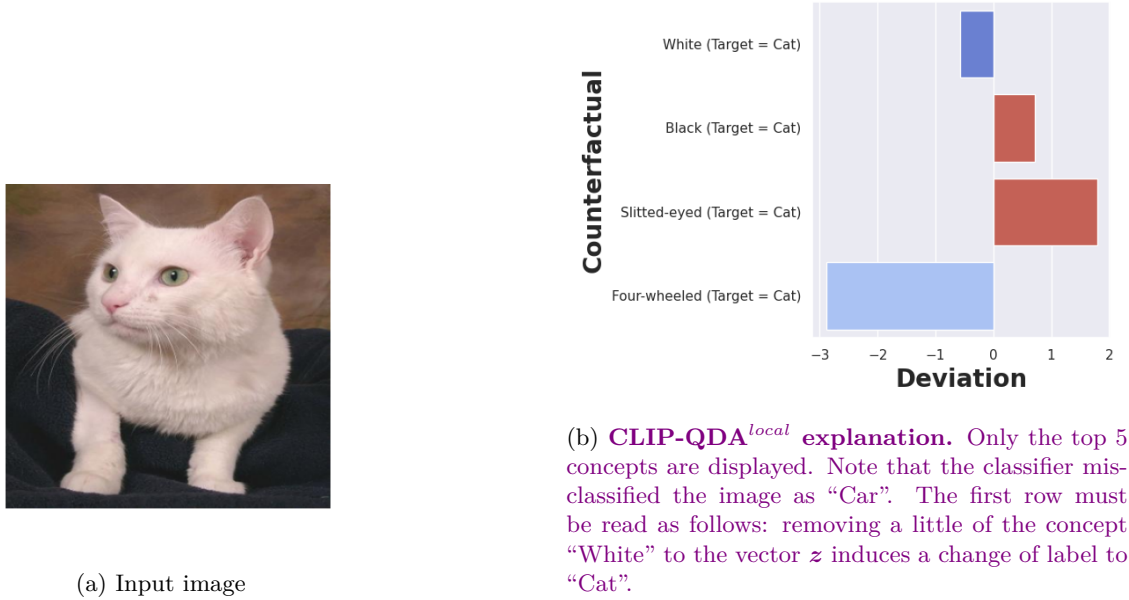


Figure 11: **Local explanation on subsets of Cats/Cars.** On the right figure, the blue/red scale represents the scaled counterfactuals as in equation 6, the text in each box corresponds to the label predicted after the perturbation, followed to the concept changed to obtain its result (in parentheses).

Table 3: **Comparison of the different features of the XAI methods.** Top: existing methods. Bottom: ours. *Dataset-wise* refers to methods that provide dataset-wise explanations. *CLIP-QDA^{local}* refers to methods that provide sample-wise explanations. *Closed-form solution* refers to methods that do not require an optimization process to produce explanations. *Computable from weights* refers to methods that produce explanations from the model parameters and input values. *Image level* refers to methods that produce explanations from the image input. *Concept level* refers to methods that produce explanations from the concept input. *Compatible with CLIP-QDA* refers to methods that can produce explanations with CLIP-QDA as a classifier. *Compatible with CLIP linear probe* refers to methods that can produce explanations with CLIP linear probe as a classifier.

Metric	Dataset wise	Sample wise	Closed-form solution	Computable from weights	Image level	Concept level	Compatible with CLIP-QDA	Compatible with CLIP linear probe
GradCAM		✓			✓		✓	✓
LIME		✓			✓		✓	✓
SHAP		✓			✓		✓	✓
LaBo	✓		✓	✓		✓		
Yan et al	✓	✓	✓	✓		✓		✓
Greybox XAI		✓			✓	✓		
X-NeSyL		✓			✓	✓		
Morales et al		✓				✓		
CLIP-LIME		✓				✓	✓	✓
CLIP-SHAP		✓				✓	✓	✓
QDA-CBM	✓	✓	✓	✓		✓	✓	

and Yan et al. (Yan et al., 2023b). A comparative overview of the distinctive features of these methods is presented in Table 3. Since Greybox XAI, X-NeSyL and Morales Rodríguez et al. require additional annotations, we chose to omit these methods from our study.

Qualitative analysis. To facilitate the comparison of various explanations generated for the same sample in the dataset, we computed the results for one image from PASCAL-Part (Donadello & Serafini, 2016). The results are categorized into three sections: sample-wise explanations at the image level, sample-wise explanations at the concept level, and dataset-wise explanations at the concept level. The subsequent section focuses on explaining the model’s prediction process given the image 12a as the input, which is labeled as “person”. The top two predicted labels by the classifier for this image are “person” and “potted plant”. Note that additional samples are available in Section A.6.

Image-level explanations are presented in Figure 12. For each of these samples, a weight is assigned to each pixel. All the results are computed by applying these post-hoc explanation methods on CLIP-QDA. Additional implementation and computation details can be found in A.1. While the results are easy to comprehend, criticism may arise due to potential misunderstandings regarding the specific focus of the method. While these methods successfully highlight the object of interest, it is challenging to discern the exact pattern the method emphasizes. Another observation is that, compared to classical classifiers, the results on CLIP-based CBMs are more imprecise, highlighting large areas of interest.



Figure 12: **Sample-wise explanations (image level).** Note that the classifier classified the image correctly.

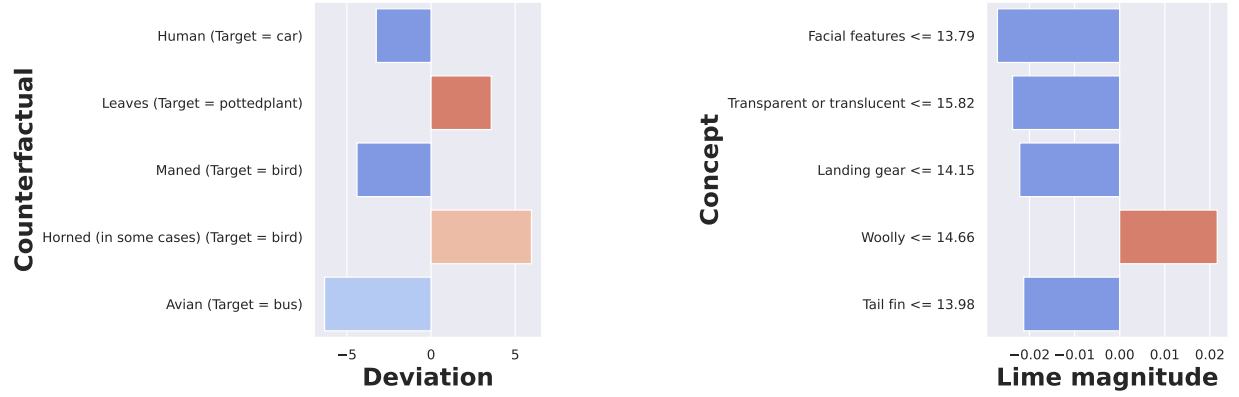
Concept-level sample explanations are calculated in Figure 13. On a sample-wise basis, CLIP-LIME and CLIP-SHAP are computed using CLIP-QDA. Given that the method is model-specific, the Yan et al. method is computed on a linear classifier.

First, we notice that there is a high diversity of concepts displayed among the methods. Overall, there are concepts associated with the prediction label (“Posture or stance”, “Human”) and concepts that arise from the specificities of the image (“Leaves”). Some concepts seem irrelevant (“Landing Gear”, “Windshield”, “Wooly”), which may be explained by the fact that their scores are useful to the classification process despite their absence in the image. Plus, it is worth noting that the context provided by the counterfactual analysis helps to unveil insights into the reason behind the explanation. For example, when taking into account the concept “Leaves” seems inappropriate, CLIP-QDA^{local} explanations suggest that adding that concept influences the prediction of the class “potted plant,” indicating that this concept is important because it reveals the decision process behind predicting a person rather than a plant.

In the realm of dataset-wise explanations (Figure 14), we present the explanations for the CLIP-QDA^{global} method, computed on CLIP-QDA, the LaBo explanation, computed on the LaBo classifier, and the Yan et al. explanation, computed on a linear classifier.


We observe that the concepts obtained from all the dataset-wise methods are generally coherent, including concepts commonly associated with persons, such as “Posture or stance” or “Facial features”. However, in contrast to the two other methods that impose positive magnitudes, the CLIP-QDA^{global} explanation allows negative values, providing insights into whether the impact is negative or positive. Additionally, our CLIP-QDA^{global} method emphasizes the comparison between two classes of interest, enabling a detailed exploration of specific disparities, like here between “potted plant” and “person”.

Quantitative analysis. To conduct a more comprehensive analysis beyond qualitative observations, we subject the various methods to quantitative evaluations. We apply the procedure outlined in Section 4.2




(a) **CLIP-QDA^{local} explanation.** The first row can be read as follows: removing a little of the concept “Human” to the vector \mathbf{z} would change the label to “Car”.

(b) **CLIP-LIME explanation.** The first row can be read as follows: the fact that the concept score of “Facial features” is below 13.79 has a negative impact on the predicted label.




(c) **Yan et al. explanation (sample).** The first row can be read as follows: the concept “Posture or stance” has a positive impact on the predicted label.



(d) **CLIP-SHAP explanation.** The highest value can be read as follows: the concept “Fuselage” has a positive impact on the predicted label.


Figure 13: **Sample-wise explanations (concept level) of the image 12a.** Except for CLIP-SHAP, only the top 5 concepts are displayed. Note that the classifier correctly labeled the image.

to the PASCAL-Part and Cats/Dogs/Cars datasets. In addition to CLIP-LIME, CLIP-SHAP, and QDA-CBM explanations, we also tested the explanations provided by LaBo and Yan et al. For the selection of S_s , in the case of the PASCAL-Part dataset, we opted to designate the concepts to detect as the concepts that are associated with the two highest probabilities of the model’s inference (refer to the Classes/Concepts



(a) **CLIP-QDA^{global}** explanation associated with the top two predictions (“person” and “potted plant”). The first row can be read as follows: the distribution of the concept “Container or pot” for the class “person” is mostly to the left (smaller values) of the one for the class “potted plant”.

(b) **LaBo** explanation associated with the top prediction (“person”). The first row can be read as follows: the weight associated with the concept “Displaying an image” has the highest value among the weights related to the class “person”.



(c) **Yan et al.** explanation (dataset) associated with the top prediction (“person”). The first row can be read as follows: the weight associated with the concept “Posture or stance” has the highest value among the weights related to the class “person”.

Figure 14: **Dataset-wise explanations of the image 12a.** Only the top 5 concepts are displayed. Note that the classifier correctly labeled the image.

association in Table 6). In the Cats/Dogs/Cars dataset, we chose to identify the concepts "Black" and "White" in the biased setup to mitigate potential biases. The resulting *Det* score for this dataset is calculated as the average across all possible biased binary classification tasks. Additionally, we incorporated inference time as a parameter in our experiments referring to the time taken to produce explanations inferences on the entire validation set. Results on PASCAL-Part are displayed in Table 5 and results on Cats/Dogs/Car are displayed in Table 4.

Regarding the Cats/Dogs/Cars dataset, our methods exhibit superior performance compared to state of the art, emphasizing the reliability of CLIP-QDA on datasets with a low number of concepts. Particularly,

Table 4: **Quantitative results for different XAI methods (Cats/Dogs/Cars dataset).** Top: existing methods. Bottom: ours. Del refers to the deletion score (7) on a random set of concepts (Set 1) and the set of concepts defined in 6 (Set 2). Det refers to the detection score (8). The explanations of the three upper methods being direct, they present no computation time.

Method	Del (Set 1) ↓	Del (Set 2) ↓	Det ↑	Inference time (im/s) ↓
<i>Yan et al.</i>	0.6253	0.7522	0.3183	/
<i>LaBo</i>	0.5971	0.6224	0.2140	/
<i>Random</i>	0.8278	0.7474	0.1171	/
CLIP-QDA ^{local} (ours)	0.7609	0.5820	0.2724	3.01
CLIP-LIME (ours)	0.5397	0.5646	0.4042	76.69
CLIP-SHAP (ours)	0.4821	0.3831	0.3696	256.76

Table 5: **Quantitative results for different XAI methods (PASCAL-Part dataset).** Top: existing methods. Bottom: ours. Del refers to the deletion score (7) on a random set of concepts (Set 1) and the usual set of concepts defined in 6 (Set 2). Det refers to the detection score (8). The explanations of the three upper methods being direct, they present no computation time.

Method	Del (Set 1) ↓	Del (Set 2) ↓	Det ↑	Inference time (im/s) ↓
<i>Yan et al.</i>	0.6968	0.8824	0.4157	/
<i>LaBo</i>	0.4446	0.8822	0.4213	/
<i>Random</i>	0.8452	0.8999	0.1258	/
CLIP-QDA ^{local} (ours)	0.7983	0.8189	0.1564	2341.12
CLIP-LIME (ours)	0.7313	0.8511	0.0873	2857.28
CLIP-SHAP (ours)	0.5698	0.5510	0.2679	7207.58

QDA-CBM demonstrates significantly faster inference times than CLIP-LIME and CLIP-SHAP, albeit at the expense of slightly lower deletion and detection scores. On datasets with a higher number of concepts and classes, such as PASCAL-Part, existing methods maintain higher scores.

5 Conclusion

In this paper, we introduce a modeling approach for the embedding space of CLIP, a foundation model trained on a vast dataset. Our observations reveal that CLIP can organize information in a distribution that exhibits similarities with a mixture of Gaussians. Building upon this insight, we develop an adapted concept bottleneck model that demonstrates competitive performance along with transparency. While the model that we have presented offers the advantage of simplicity and a limited number of parameters, it does encounter challenges when dealing with a broader and more ambiguous set of concepts. As a suggestion for future research, we propose to extend this modeling approach to incorporate other priors, such as Laplacian distributions, and to explore more complex models, including those with multiple components, *i.e.*, using more than one Gaussian to describe a class for example. **Another avenue to potentially enhance the performance of our model is to explore guiding the latent space towards a Gaussian distribution.** Additionally, our research is centered around a specific embedding space (CLIP scores). Exploring similar work on other latent spaces, particularly those associated with multimodal foundation models, could be valuable to determine if similar patterns exist in those spaces.

References

- Cem Akkus, Luyang Chu, Vladana Djakovic, Steffen Jauch-Walser, Philipp Koch, Giacomo Loss, Christopher Marquardt, Marco Moldovan, Nadja Sauter, Maximilian Schneider, et al. Multimodal deep learning. *arXiv preprint arXiv:2301.04856*, 2023.
- Alejandro Barredo Arrieta, Natalia Díaz-Rodríguez, Javier Del Ser, Adrien Bennetot, Siham Tabik, Alberto Barbado, Salvador García, Sergio Gil-López, Daniel Molina, Richard Benjamins, et al. Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion*, 58:82–115, 2020.
- Adrien Bennetot, Gianni Franchi, Javier Del Ser, Raja Chatila, and Natalia Diaz-Rodriguez. Greybox XAI: A neural-symbolic learning framework to produce interpretable predictions for image classification. *Knowledge-Based Systems*, 258:109947, 2022.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D. Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33:1877–1901, 2020.
- John M. Chambers. *Graphical Methods for Data Analysis*. CRC Press, 2018.
- Paulo Cortez and Mark J. Embrechts. Opening black box data mining models using sensitivity analysis. In *2011 IEEE Symposium on Computational Intelligence and Data Mining (CIDM)*, pp. 341–348. IEEE, 2011.
- Will Cukierski. Dogs vs. cats, 2013. URL <https://kaggle.com/competitions/dogs-vs-cats>.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 248–255, 2009.
- Natalia Díaz-Rodríguez, Alberto Lamas, Jules Sanchez, Gianni Franchi, Ivan Donadello, Siham Tabik, David Filliat, Policarpo Cruz, Rosana Montes, and Francisco Herrera. Explainable neural-symbolic learning (x-nesyl) methodology to fuse deep learning representations with expert knowledge graphs: The monumai cultural heritage use case. *Information Fusion*, 79:58–83, 2022.
- Ivan Donadello and Luciano Serafini. Integration of numeric and symbolic information for semantic image interpretation. *Intelligenza Artificiale*, 10(1):33–47, 2016.
- 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*, 2021.
- Zhili Feng, Anna Bair, and J. Zico Kolter. Leveraging multiple descriptive features for robust few-shot image learning. *arXiv preprint arXiv:2307.04317*, 2023.
- Valentin Gabeff, Marc Russwurm, Devis Tuia, and Alexander Mathis. Wildclip: Scene and animal attribute retrieval from camera trap data with domain-adapted vision-language models. *bioRxiv*, pp. 2023–12, 2023.
- Francis Galton. Regression towards mediocrity in hereditary stature. *The Journal of the Anthropological Institute of Great Britain and Ireland*, 15:246–263, 1886.
- Trevor Hastie, Robert Tibshirani, and Jerome H. Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, volume 2. Springer, 2009.
- 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*, pp. 770–778, 2016.

- Anna Hedström, Leander Weber, Daniel Krakowczyk, Dilyara Bareeva, Franz Motzkus, Wojciech Samek, Sebastian Lapuschkin, and Marina 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. URL <http://jmlr.org/papers/v24/22-0142.html>.
- Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part II 14*, pp. 694–711. Springer, 2016.
- Siwon Kim, Jinoh Oh, Sungjin Lee, Seunghak Yu, Jaeyoung Do, and Tara Taghavi. Grounding counterfactual explanation of image classifiers to textual concept space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10942–10950, 2023.
- Pang Wei Koh, Thao Nguyen, Yew Siang Tang, Stephen Mussmann, Emma Pierson, Been Kim, and Percy Liang. Concept bottleneck models. In *International Conference on Machine Learning*, pp. 5338–5348, 2020.
- Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In *4th International IEEE Workshop on 3D Representation and Recognition (3dRR-13)*, Sydney, Australia, 2013.
- Neeraj Kumar, Alexander C. Berg, Peter N. Belhumeur, and Shree K. Nayar. Attribute and simile classifiers for face verification. In *2009 IEEE 12th International Conference on Computer Vision*, pp. 365–372. IEEE, 2009.
- Alberto Lamas, Siham Tabik, Policarpo Cruz, Rosana Montes, Álvaro Martinez-Sevilla, Teresa Cruz, and Francisco Herrera. Monumai: Dataset, deep learning pipeline and citizen science based app for monumental heritage taxonomy and classification. *Neurocomputing*, 420:266–280, 2021.
- Christoph H. Lampert, Hannes Nickisch, and Stefan Harmeling. Learning to detect unseen object classes by between-class attribute transfer. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 951–958. IEEE, 2009.
- Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015.
- Xiaodong Liu, Kevin Duh, Liyuan Liu, and Jianfeng Gao. Very deep transformers for neural machine translation. *arXiv preprint arXiv:2008.07772*, 2020.
- Max Losch, Mario Fritz, and Bernt Schiele. Interpretability beyond classification output: Semantic bottleneck networks. *arXiv preprint arXiv:1907.10882*, 2019.
- Scott M. Lundberg and Su-In Lee. A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems*, 30, 2017.
- Huaishao Luo, Lei Ji, Ming Zhong, Yang Chen, Wen Lei, Nan Duan, and Tianrui Li. Clip4clip: An empirical study of clip for end to end video clip retrieval and captioning. *Neurocomputing*, 508:293–304, 2022.
- Jinqi Luo, Zhaoning Wang, Chen Henry Wu, Dong Huang, and Fernando De la Torre. Zero-shot model diagnosis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11631–11640, 2023.
- Prasanta Chandra Mahalanobis. On the generalized distance in statistics. *Sankhyā: The Indian Journal of Statistics, Series A (2008-)*, 80:S1–S7, 2018.
- Peter McCullagh. *Generalized Linear Models*. Routledge, 2019.
- Sachit Menon and Carl Vondrick. Visual classification via description from large language models. In *The Eleventh International Conference on Learning Representations*, 2022.

- David Morales Rodríguez, Manuel Pegalajar Cuellar, and Diego P Morales. On the fusion of soft-decision-trees and concept-based models. *Available at SSRN 4402768*.
- Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In *Indian Conference on Computer Vision, Graphics and Image Processing*, 12 2008.
- Tuomas Oikarinen, Subhro Das, Lam Nguyen, and Lily Weng. Label-free concept bottleneck models. In *International Conference on Learning Representations*, 2023.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744, 2022.
- Konstantinos P. Panousis, Dino Ienco, and Diego Marcos. Sparse linear concept discovery models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 2767–2771, 2023.
- Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. *Elements of Causal Inference: Foundations and Learning Algorithms*. The MIT Press, 2017.
- Vitali Petsiuk, Abir Das, and Kate Saenko. RISE: Randomized input sampling for explanation of black-box models. *arXiv preprint arXiv:1806.07421*, 2018.
- Gregory Plumb, Denali Molitor, and Ameet S Talwalkar. Model agnostic supervised local explanations. *Advances in neural information processing systems*, 31, 2018.
- Gregory Plumb, Marco Tulio Ribeiro, and Ameet Talwalkar. Finding and fixing spurious patterns with explanations. *Transactions on Machine Learning Research*, 2022. ISSN 2835-8856. URL <https://openreview.net/forum?id=whJPugmP5I>. Expert Certification.
- Ariadna Quattoni and Antonio Torralba. Recognizing indoor scenes. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 413–420. IEEE, 2009.
- J. Ross Quinlan. Induction of decision trees. *Machine Learning*, 1:81–106, 1986.
- 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*, pp. 8748–8763, 2021.
- Aaditya Ramdas, Nicolás García Trillos, and Marco Cuturi. On wasserstein two-sample testing and related families of nonparametric tests. *Entropy*, 19(2):47, 2017.
- Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *International Conference on Machine Learning*, pp. 8821–8831, 2021.
- Scott Reed, Konrad Zolna, Emilio Parisotto, Sergio Gómez Colmenarejo, Alexander Novikov, Gabriel Barthmaron, Mai Giménez, Yury Sulsky, Jackie Kay, Jost Tobias Springenberg, et al. A generalist agent. *Transactions on Machine Learning Research*, 2022.
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. “Why should I trust you?”: Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, August 13-17, 2016*, pp. 1135–1144, 2016.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 10674–10685, 2021.

- Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. Laion-5b: An open large-scale dataset for training next generation image-text models. *Advances in Neural Information Processing Systems*, 35:25278–25294, 2022.
- 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*, pp. 618–626, 2017.
- Tatiana Tommasi, Novi Patricia, Barbara Caputo, and Tinne Tuytelaars. A deeper look at dataset bias. *Domain Adaptation in Computer Vision Applications*, pp. 37–55, 2017.
- An Yan, Yu Wang, Yiwu Zhong, Chengyu Dong, Zexue He, Yujie Lu, William Yang Wang, Jingbo Shang, and Julian McAuley. Learning concise and descriptive attributes for visual recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 3090–3100, 2023a.
- An Yan, Yu Wang, Yiwu Zhong, Zexue He, Petros Karypis, Zihan Wang, Chengyu Dong, Amilcare Gentili, Chun-Nan Hsu, Jingbo Shang, et al. Robust and interpretable medical image classifiers via concept bottleneck models. *arXiv preprint arXiv:2310.03182*, 2023b.
- Yue Yang, Artemis Panagopoulou, Shenghao Zhou, Daniel Jin, Chris Callison-Burch, and Mark Yatskar. Language in a bottle: Language model guided concept bottlenecks for interpretable image classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 19187–19197, 2023.
- Mert Yuksekgonul, Maggie Wang, and James Zou. Post-hoc concept bottleneck models. In *The Eleventh International Conference on Learning Representations*, 2022.

A Appendix

A.1 Implementation details

GradCAM. To compute GradCAM explanations, we applied the method to the 21th block of the image encoder in CLIP-QDA.

LIME. The LIME image-level explanations are generated using the superpixel approach defined in (Ribeiro et al., 2016). This method involves treating the input as a set of superpixels obtained through a segmentation process. LIME is then applied, considering these superpixels as binary values corresponding to whether they contribute to displaying the related zone. Following the original implementation, we use quickshift clustering to segment the image. We display the top 10 superpixels that exert the most influence on the prediction, highlighting superpixels in blue that have a negative impact on the label and superpixels in red that have a positive impact on the label.

SHAP. Following the advised method from the SHAP library, the image-level explanation from SHAP is computed using the partition explainer, which creates from the input a subset of images by applying a set of perturbations. In all our experiments, the perturbations are a blur of kernel (128,128).

CLIP-LIME. For the CLIP-LIME, we consider inputs as a set of tabular values. Then, each CLIP-LIME explanation is obtained by running the lime tabular explainer provided by the author’s repo with *num features* = 100000, *num samples* = 1000, and *batchsize* = 10.

CLIP-SHAP. For the CLIP-SHAP, we consider inputs as a set of tabular values. Then, each CLIP-SHAP explanation is obtained by running the shap kernel explainer provided by the author’s repo with a background of 5 samples using k-means.

LaBo. For LaBo, we train a LaBo classifier by using Adam optimizer with a learning rate of $5e-1$ a weight decay of 0, and a batch size of 8192. The resulting explanations follow the resulting weight matrix.

Yan et al. For the Yan et al. method, we train a linear classifier by using Adam optimizer with a learning rate of $5e-3$ a weight decay of $1e-4$, and a batch size of 512. The sample-wise (concept level) explanation results from the product between the concept score and its associated weight. The dataset-wise explanation results from the weights alone.

A.2 Set of concepts

Inspired by Yang et al. (2023), we use large language models to provide concept sets. Concretely, we use GPT-3 (Brown et al., 2020) with the following preprompt: “In this task, you have to give visual descriptions that describe an image. Respond as a list. Each item being a word.” Then, we generate the sets of words by the following prompt: “What are [N] useful visual descriptors to distinguish a [class] in a photo?” By doing so, we generated 5 concepts per class, presented in Table 6. An additional ordering of concepts by subcategories is available in Table 7.

Table 6: List of concepts used, ordered by classes of interest.

Dataset	Set of concepts
<i>PASCAL-Part</i>	'aeroplane':[Winged, Jet engines, Tail fin, Fuselage, Landing gear], 'bicycle':[Two-wheeled, Pedals, Handlebars, Frame, Chain-driven], 'bird':[Feathery, Beaked, Wingspread, Perched, Avian], 'bottle':[Glass or plastic, Cylindrical, Necked, Cap or cork, Transparent or translucent], 'bus':[Large, Rectangular, Windows, Wheels, Multi-doored], 'cat':[Furry, Whiskered, Pointy-eared, Slitted-eyed, Four-legged], 'car':[Metallic, Four-wheeled, Headlights, Windshield, License plate], 'dog':[Snout, Wagging-tailed, Snout-nosed, Floppy-eared, Tail-wagging], 'cow':[Bovine, Hooved, Horned (in some cases), Spotted or solid-colored, Grazing (if in a field)], 'horse':[Equine, Hooved, Maned, Tailed, Galloping (if in motion)], 'motorbike':[Two-wheeled, Engine, Handlebars, Exhaust, Saddle or seat], 'person':[Human, Facial features, Limbs (arms and legs), Clothing, Posture or stance], 'potted plant':[Potted, Green, Leaves, Soil or potting mix, Container or pot], 'sheep':[Woolly, Hooved, Grazing, Herded (if in a group), White or colored fleece], 'train':[Locomotive, Railcars, Tracks, Wheels, Carriages], 'tvmonitor':[Screen, Rectangular, Frame or bezel, Stand or wall-mounted, Displaying an image]
<i>MonuMAI</i>	'Baroque':[Ornate, Elaborate sculptures, Intricate details, Curved or asymmetrical design, Historical or aged appearance], 'Gothic':[Pointed arches, Ribbed vaults, Flying buttresses, Stained glass windows, Tall spires or towers], 'Hispanic muslim':[Mudejar style, Intricate geometric patterns, Horseshoe arches, Decorative tilework (azulejos), Islamic-inspired motifs], 'Renaissance':[Classical proportions, Symmetrical design, Columns and pilasters, Human statues and sculptures, Dome or dome-like structures]
<i>MIT scenes</i>	'Store':[Building or structure, Signage or banners, Glass windows or doors, Displayed products or merchandise, People entering or exiting (if applicable)], 'Home':[Residential, Roofed, Windows, Landscaping or yard, Front entrance or door], 'Public space':[Open area, Crowds (if people are present), Benches or seating, Pathways or walkways, Architectural features (e.g., buildings, statues)], 'Leisure':[Recreational, Play equipment (if applicable), Greenery or landscaping, Picnic tables or seating, Relaxing atmosphere], 'Working space':[Office equipment (e.g., desks, computers), Task-oriented, Office chairs, Organized or structured, People working (if applicable)]
<i>Cats/Dogs/Cars</i>	'Cat':[Furry, Whiskered, Pointy-eared, Slitted-eyed, Four-legged], 'Car':[Metallic, Four-wheeled, Headlights, Windshield, License plate], 'Dog':[Snout, Wagging-tailed, Snout-nosed, Floppy-eared, Tail-wagging]

Table 7: List of concepts used, ordered by subcategories.

Dataset	Set of concepts
<i>PASCAL-Part</i>	'Objects':[Screen, Rectangular, Frame or bezel, Stand or wall-mounted, Displaying an image, Glass or plastic, Cylindrical, Necked, Cap or cork, Transparent or translucent, Potted, Green, Leaves, Soil or potting mix, Container or pot], 'Transportation-related':[Two-wheeled, Engine, Handlebars, Exhaust, Saddle or seat, Four-wheeled, Headlights, Windshield, License plate, Wheels, Multi-doored, Locomotive, Railcars, Tracks, Carriages, Chain-driven], 'Aircraft-related':[Winged, Jet engines, Tail fin, Fuselage, Landing gear], 'Building/Structure-related':[Large, Rectangular, Windows] 'Human Characteristics':[Human, Facial features, Limbs (arms and legs), Clothing, Posture or stance], 'Avian Characteristics':[Feathery, Beaked, Wingspread, Perched, Avian], 'Animal Characteristics':[Furry, Whiskered, Pointy-eared, Slitted-eyed, Four-legged, Equine, Hooved, Maned, Tailed, Galloping (if in motion), Snout, Wagging-tailed, Snout-nosed, Floppy-eared, Tail-wagging, Woolly, Grazing, Herded (if in a group), White or colored fleece, Bovine, Horned (in some cases), Spotted or solid-colored]
<i>MonuMAI</i>	'Architectural Styles and Elements':[Mudejar style, Intricate geometric patterns, Horse-shoe arches, Decorative tilework (azulejos), Islamic-inspired motifs, Stained glass windows, Pointed arches, Ribbed vaults], 'Artistic Details and Features':[Ornate, Elaborate sculptures, Intricate details, Curved or asymmetrical design, Human statues and sculptures, Historical or aged appearance], 'Architectural Components':[Flying buttresses, Classical proportions, Symmetrical design, Columns and pilasters, Dome or dome-like structures, Tall spires or towers]
<i>MIT scenes</i>	'Work Environment':[Office equipment (e.g., desks, computers), Task-oriented, Office chairs, Organized or structured, People working (if applicable)], 'Leisure Environment':[Play equipment (if applicable), Greenery or landscaping, Picnic tables or seating, Relaxing atmosphere, Recreational], 'Community Environment':[Roofed, Windows, Landscaping or yard, Front entrance or door, Building or structure, Signage or banners, Glass windows or doors, Displayed products or merchandise, People entering or exiting (if applicable), Open area, Crowds (if people are present), Benches or seating, Pathways or walkways, Architectural features (e.g., buildings, statues), Residential]
<i>Cats/Dogs/Cars</i>	'Vehicle Features':[License plate, Headlights, Windshield, Four-wheeled], 'Animal Features':[Wagging-tailed, Tail-wagging, Snout, Whiskered, Pointy-eared, Slitted-eyed, Floppy-eared, Furry, Snout-nosed, Four-legged], 'Material/Texture':[Metallic]

A.3 Study of images according to their position in the latent space

In addition to the CLIP-QDA^{global} method, we introduce another approach leveraging multivariate Gaussian modeling to generate dataset-wise visual explanations for our model. Specifically, for a given class c , we compute the Mahalanobis distance with respect to $\mathcal{N}(\mathbf{z} \mid \boldsymbol{\mu}_c, \boldsymbol{\Sigma}_c)$ for each sample's latent space projection \mathbf{z} labeled as c within a dataset of interest:

$$d_M(\mathbf{z}, \boldsymbol{\mu}_c, \boldsymbol{\Sigma}_c) = \sqrt{(\mathbf{z} - \boldsymbol{\mu}_c)^\top \boldsymbol{\Sigma}_c^{-1} (\mathbf{z} - \boldsymbol{\mu}_c)}. \quad (9)$$

Then, we plot in Figures 15, 16, 17 and 18, the closest and farthest images with the dataset, according to this distance for PascalPART and MonuMAI.

For PascalPART, the closest images tend to be predominantly close-up shots of the object in question. However, this occasionally presents challenges due to the inherent pixelation of the images, leading to somewhat counterintuitive results. Conversely, the furthest images are typically those where the object is distant. Another scenario of failure arises when there are two classes present within the image, which complicates

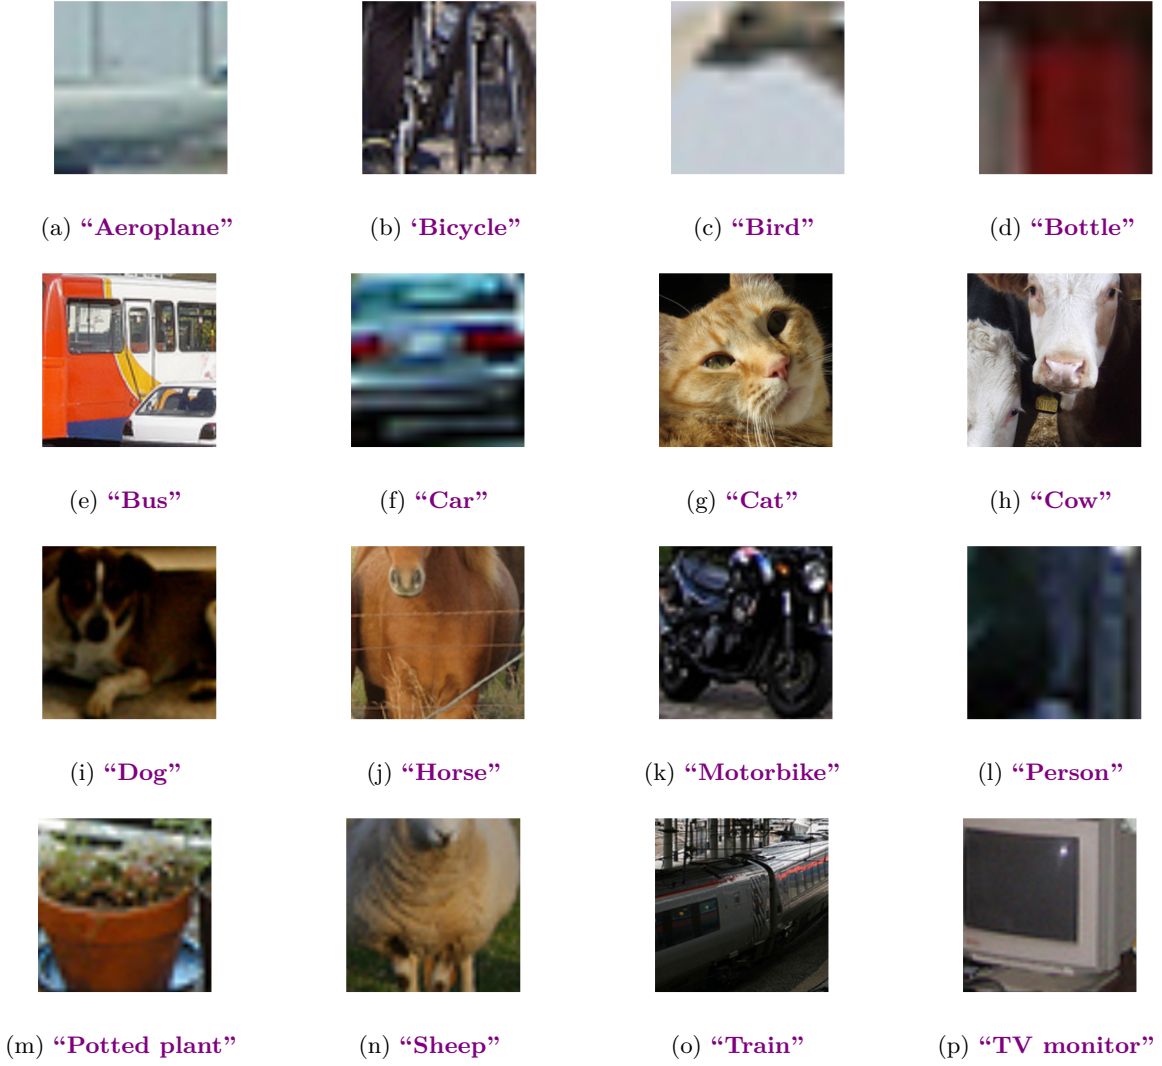


Figure 15: Images of minimal distance for each classes of PascalPART

the retrieval process. On the other hand, for MonumAI, the closest images appear to be the simplest ones, without sign of disturbance. In contrast, the farthest images may include extraneous objects, such as signs.

A.4 Closed-form solution of counterfactuals for CLIP QDA

First, we will derive the resolution of equation 5 for the binary case before extending it to the multiclass case. Let us begin with a binary classifier, where $Y \in \{h_{\omega^h}(z), \bar{h}_{\omega^h}(z)\}$ (for convenience, we note $h_{\omega^h}(z) = c_h$ and $\bar{h}_{\omega^h}(z) = c_h^-$).

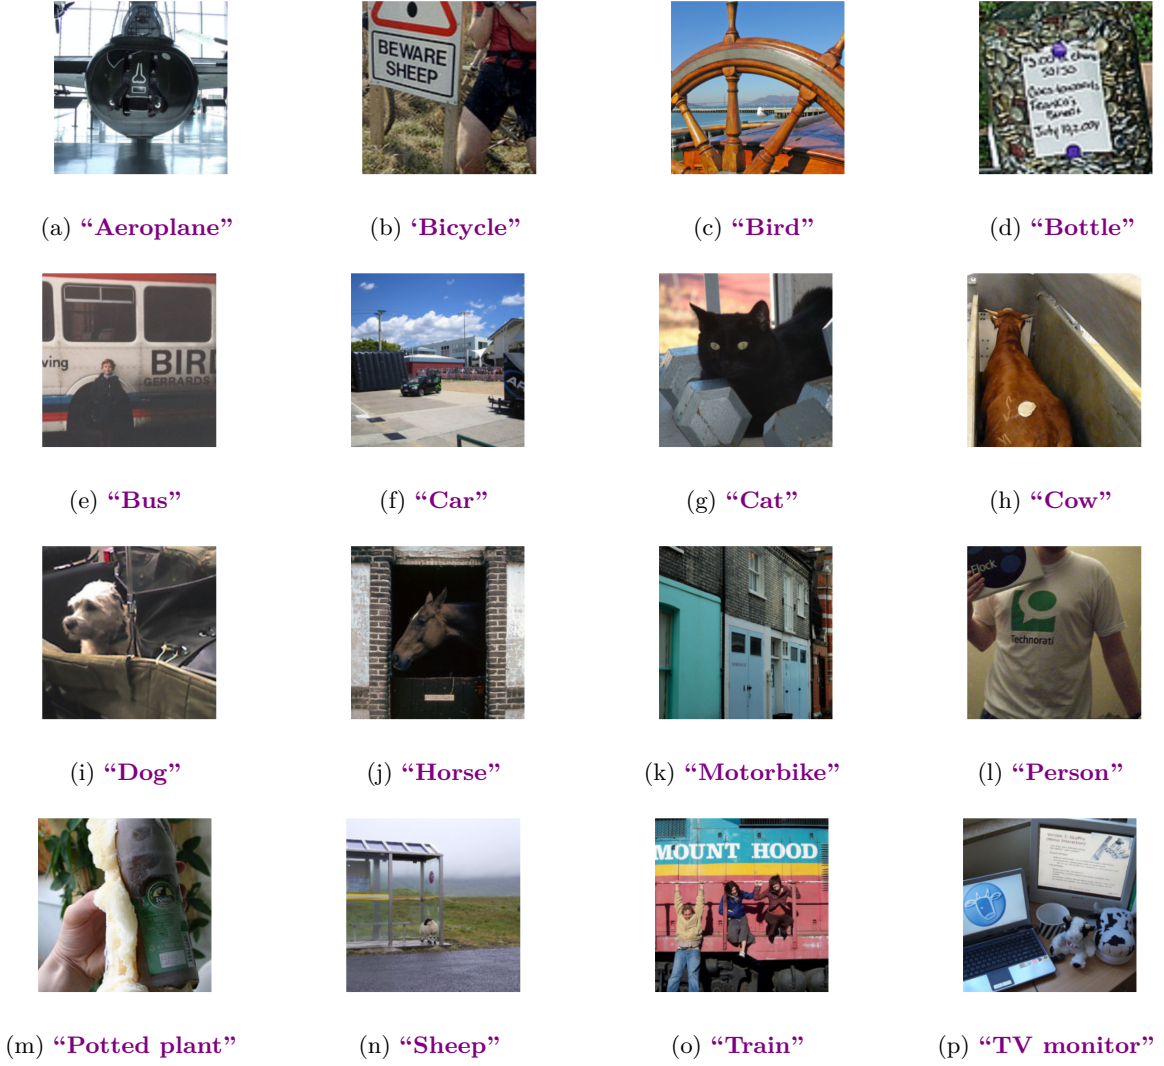



Figure 16: Images of maximal distance for each classes of PascalPART



Figure 17: Images of minimal distance for each classes of MonuMAI

Proposition 2. Let h_{ω^h} a binary classifier and \mathbf{Z} following a mixture of Gaussians such as $\mathbf{Z}_{Y=c_h} \sim \mathcal{N}(\boldsymbol{\mu}_{c_h}, \boldsymbol{\Sigma}_{c_h})$ and $\mathbf{Z}_{Y=c_h^-} \sim \mathcal{N}(\boldsymbol{\mu}_{c_h^-}, \boldsymbol{\Sigma}_{c_h^-}^{-1})$ and $\boldsymbol{\epsilon}_s^j$ a perturbation with the above sparsity and sign restrictions.

Figure 18: **Images of maximal distance for each classes of MonuMAI**

Then, there is a closed form solution to problem 5, given by:

$$\epsilon_s^j = \begin{cases} \emptyset & \text{if } b^2 - 4Pc < 0 \\ & \text{or } (s \neq \text{sign}(b_1) \text{ and } s \neq \text{sign}(b_2)) \\ b_1 & \text{if } b^2 - 4Pc > 0 \text{ and } \text{sign}(b_1) = s \\ & \text{and } (\text{sign}(b_2) \neq s \text{ or } |b_2| \geq |b_1|) \\ b_2 & \text{if } b^2 - 4Pc > 0 \text{ and } \text{sign}(b_2) = s \\ & \text{and } (\text{sign}(b_1) \neq s \text{ or } |b_1| > |b_2|), \end{cases} \quad (10)$$

with:

$$\begin{aligned} P &= \frac{1}{2} \left([\Sigma_{c_h^-}^{-1}]_{(j,j)} - [\Sigma_{c_h}^{-1}]_{(j,j)} \right) \\ b &= \sum_{k=1}^N \left(([z]_{(k)} - [\mu_{c_h^-}]_{(k)}) [\Sigma_{c_h^-}^{-1}]_{(j,k)} - ([z]_{(k)} - [\mu_{c_h}]_{(k)}) [\Sigma_{c_h}^{-1}]_{(j,k)} \right) \\ c &= \frac{1}{2} \log \left[\frac{|\Sigma_{c_h^-}|}{|\Sigma_{c_h}|} \right] + \log \left[\frac{p_{c_h}}{p_{c_h^-}} \right] + \frac{1}{2} (\mathbf{z} - \mu_{c_h^-})^\top \Sigma_{c_h^-}^{-1} (\mathbf{z} - \mu_{c_h^-}) - \frac{1}{2} (\mathbf{z} - \mu_{c_h})^\top \Sigma_{c_h}^{-1} (\mathbf{z} - \mu_{c_h}) \\ b_1 &= \frac{-b - \sqrt{b^2 - 4Pc}}{2P} \\ b_2 &= \frac{-b + \sqrt{b^2 - 4Pc}}{2P}. \end{aligned} \quad (11)$$

Proof. For the binary case, equation 5 can be written as :

$$\min \|\epsilon_s^j\|^2 \quad \text{s.t.} \quad \frac{p_{c_h^-}}{(2\pi)^{N/2} |\Sigma_{c_h^-}|^{\frac{1}{2}}} e^{-\frac{1}{2} (\mathbf{z} + \epsilon_s^j - \mu_{c_h^-})^\top \Sigma_{c_h^-}^{-1} (\mathbf{z} + \epsilon_s^j - \mu_{c_h^-})} \leq \frac{p_{c_h}}{(2\pi)^{N/2} |\Sigma_{c_h}|^{\frac{1}{2}}} e^{-\frac{1}{2} (\mathbf{z} + \epsilon_s^j - \mu_{c_h})^\top \Sigma_{c_h}^{-1} (\mathbf{z} + \epsilon_s^j - \mu_{c_h})}. \quad (12)$$

Let us focus on the inequality constraint:

$$\begin{aligned}
& \frac{p_{c_h^-}}{|\Sigma_{c_h^-}|^{\frac{1}{2}}} e^{-\frac{1}{2}(\mathbf{z} + \boldsymbol{\epsilon}_s^j - \boldsymbol{\mu}_{c_h^-})^T \Sigma_{c_h^-}^{-1} (\mathbf{z} + \boldsymbol{\epsilon}_s^j - \boldsymbol{\mu}_{c_h^-})} \leq \frac{p_{c_h}}{|\Sigma_{c_h}|^{\frac{1}{2}}} e^{-\frac{1}{2}(\mathbf{z} + \boldsymbol{\epsilon}_s^j - \boldsymbol{\mu}_{c_h})^T \Sigma_{c_h}^{-1} (\mathbf{z} + \boldsymbol{\epsilon}_s^j - \boldsymbol{\mu}_{c_h})} \\
& \iff \log(p_{c_h^-}) - \frac{1}{2} \log |\Sigma_{c_h^-}| - \frac{1}{2} (\mathbf{z} + \boldsymbol{\epsilon}_s^j - \boldsymbol{\mu}_{c_h^-})^T \Sigma_{c_h^-}^{-1} (\mathbf{z} + \boldsymbol{\epsilon}_s^j - \boldsymbol{\mu}_{c_h^-}) \\
& \quad \leq \log(p_{c_h}) - \frac{1}{2} \log |\Sigma_{c_h}| - \frac{1}{2} (\mathbf{z} + \boldsymbol{\epsilon}_s^j - \boldsymbol{\mu}_{c_h})^T \Sigma_{c_h}^{-1} (\mathbf{z} + \boldsymbol{\epsilon}_s^j - \boldsymbol{\mu}_{c_h}) \\
& \iff \log(p_{c_h^-}) - \frac{1}{2} \log |\Sigma_{c_h^-}| - \frac{1}{2} (\mathbf{z} - \boldsymbol{\mu}_{c_h^-})^T \Sigma_{c_h^-}^{-1} (\mathbf{z} - \boldsymbol{\mu}_{c_h^-}) - \frac{1}{2} \boldsymbol{\epsilon}_s^{jT} \Sigma_{c_h^-}^{-1} \boldsymbol{\epsilon}_s^j - \boldsymbol{\epsilon}_s^{jT} \Sigma_{c_h^-}^{-1} (\mathbf{z} - \boldsymbol{\mu}_{c_h^-}) \\
& \quad \leq \log(p_{c_h}) - \frac{1}{2} \log |\Sigma_{c_h}| - \frac{1}{2} (\mathbf{z} - \boldsymbol{\mu}_{c_h})^T \Sigma_{c_h}^{-1} (\mathbf{z} - \boldsymbol{\mu}_{c_h}) - \frac{1}{2} \boldsymbol{\epsilon}_s^{jT} \Sigma_{c_h}^{-1} \boldsymbol{\epsilon}_s^j - \boldsymbol{\epsilon}_s^{jT} \Sigma_{c_h}^{-1} (\mathbf{z} - \boldsymbol{\mu}_{c_h}) \\
& \iff \log(p_{c_h^-}) - \frac{1}{2} \log |\Sigma_{c_h^-}| - \frac{1}{2} (\mathbf{z} - \boldsymbol{\mu}_{c_h^-})^T \Sigma_{c_h^-}^{-1} (\mathbf{z} - \boldsymbol{\mu}_{c_h^-}) - \frac{1}{2} [\Sigma_{c_h^-}^{-1}]_{(j,j)} (\epsilon_s^j)^2 - \epsilon_s^j \sum_{k=1}^N ([\mathbf{z}]_{(k)} - [\boldsymbol{\mu}_{c_h^-}]_{(k)}) [\Sigma_{c_h^-}^{-1}]_{(j,k)} \\
& \leq \log(p_{c_h}) - \frac{1}{2} \log |\Sigma_{c_h}| - \frac{1}{2} (\mathbf{z} - \boldsymbol{\mu}_{c_h})^T \Sigma_{c_h}^{-1} (\mathbf{z} - \boldsymbol{\mu}_{c_h}) - \frac{1}{2} [\Sigma_{c_h}^{-1}]_{(j,j)} (\epsilon_s^j)^2 - \epsilon_s^j \sum_{k=1}^N ([\mathbf{z}]_{(k)} - [\boldsymbol{\mu}_{c_h}]_{(k)}) [\Sigma_{c_h}^{-1}]_{(j,k)} \\
& \iff P(\epsilon_s^j)^2 + b\epsilon_s^j + c \geq 0.
\end{aligned}$$

Then we can rewrite the problem as :

$$\min (\epsilon_s^j)^2 \text{ s.t. } P(\epsilon_s^j)^2 + b\epsilon_s^j + c \leq 0. \quad (13)$$

To solve this problem, we introduce slack variables λ and I , and define a Lagrangian as:

$$L(\epsilon_s^j, \lambda, I) = (\epsilon_s^j)^2 + \lambda(P(\epsilon_s^j)^2 + b\epsilon_s^j + c + I^2). \quad (14)$$

Then, we can find the possible solutions by solving the system :

$$\begin{cases} \frac{\partial L}{\partial \epsilon_s^j} = 0 \\ \frac{\partial L}{\partial \lambda} = 0 \\ \frac{\partial L}{\partial I} = 0, \end{cases}$$

which corresponds to:

$$\begin{cases} 2(\lambda P + 1)\epsilon_s^j + \lambda b = 0 \\ P(\epsilon_s^j)^2 + b\epsilon_s^j + c + I^2 = 0 \\ 2\lambda I = 0. \end{cases} \quad (15)$$

The third equation of 15 indicates whether the inequality condition is saturated or not. If it is not saturated ($\lambda = 0$), then the label c_h^- is already achieved for \mathbf{z} , resulting in a solution of $\epsilon_{s,*}^j = 0$. This being impossible by construction, we only focus on the case where $\lambda \neq 0$.

If $\lambda \neq 0$, the condition is saturated, the second equation leads to:

$$P(\epsilon_s^j)^2 + b\epsilon_s^j + c = 0,$$

whose solutions are:

$$\epsilon_s^j \in \left\{ \frac{-b - \sqrt{b^2 - 4Pc}}{2P}, \frac{-b + \sqrt{b^2 - 4Pc}}{2P} \right\} \text{ if } b^2 - 4Pc > 0,$$

which we rewrite as:

$$\epsilon_s^j \in \{b_1, b_2\} \text{ if } b^2 - 4Pc > 0. \quad (16)$$

Considering 16, the validity of the results is conditioned by the sign s and the condition of the magnitude of $\epsilon_{s,*}^j$. Then, the final result of the problem 5 is either \emptyset , b_1 or b_2 depending on the conditions:

$$\left\{ \begin{array}{l} \emptyset \text{ if } b^2 - 4Pc < 0 \\ \text{or } (s \neq \text{sign}(b_1) \text{ and } s \neq \text{sign}(b_2)) \\ b_1 \text{ if } b^2 - 4Pc > 0 \text{ and } \text{sign}(b_1) = s \\ \text{and } (\text{sign}(b_2) \neq s \text{ or } |b_2| \geq |b_1|) \\ b_2 \text{ if } b^2 - 4Pc > 0 \text{ and } \text{sign}(b_2) = s \\ \text{and } (\text{sign}(b_1) \neq s \text{ or } |b_1| > |b_2|). \end{array} \right. \quad (17)$$

□

To expand problem 5 to multiclass classification $C > 2$, we consider it as a succession of $C - 1$ binary classifications between each $i' \neq c_h$ and c_h . Given a concept index j and a sign s , if we denote the solutions of these problems as the set $\{\epsilon_{+,*,1}^j, \dots, \epsilon_{s,*,c_h-1}^j, \epsilon_{s,*,c_h+1}^j, \dots, \epsilon_{s,*,C}^j\}$, the final solution $\epsilon_{s,*}^j$ is if it exists, the value of minimal magnitude among this set.

Examples of explanations based on this metric are given in Sections 4.5 and A.6.

A.5 Sample-wise explanation time

In Table 8, we display the amount of time taken to produce a sample-wise explanation for each image of the test set of PASCAL-Part and Cats/Dogs/Cars.

Table 8: Time (in seconds) to produce explanations on the entire test set.

Method	CLIP-QDA ^{local}	CLIP-LIME	CLIP-SHAP
<i>Cats/Dogs/Cars</i>	3.01	76.69	256.76
<i>PASCAL-Part</i>	2341.12	2857.28	7207.58
<i>MonuMAI</i>	2.19	19.94	56.25
<i>MITscenes</i>	33.41	509.33	1398.64

In this table, we can observe that using the local explanation is the fastest, especially for the cases where the number of concepts and classes are low. This is justified by the fact that our method consists of using Proposition 2 for each concept, sign, and class other than the inference one. Hence the complexity of this computation for each sample is $O(2(C - 1)N)$, where LIME’s complexity does not depend on C .

A.6 Additional samples

We display here additional samples of both **sample-wise** and **dataset-wise** explanations from the PASCAL-Part and Cats/Dogs/Cars dataset.

(a) Input image (b) GradCAM Explanation (c) LIME Explanation (d) SHAP explanation




Figure 19: **Sample-wise explanations (image level)**. The first two examples come from the PascalPART dataset, and the last two samples come from the Cats/Dogs/Cars dataset in the biased setup. Note that the classifier mislabeled the 2nd example as “car” and the 4th example as “car”.

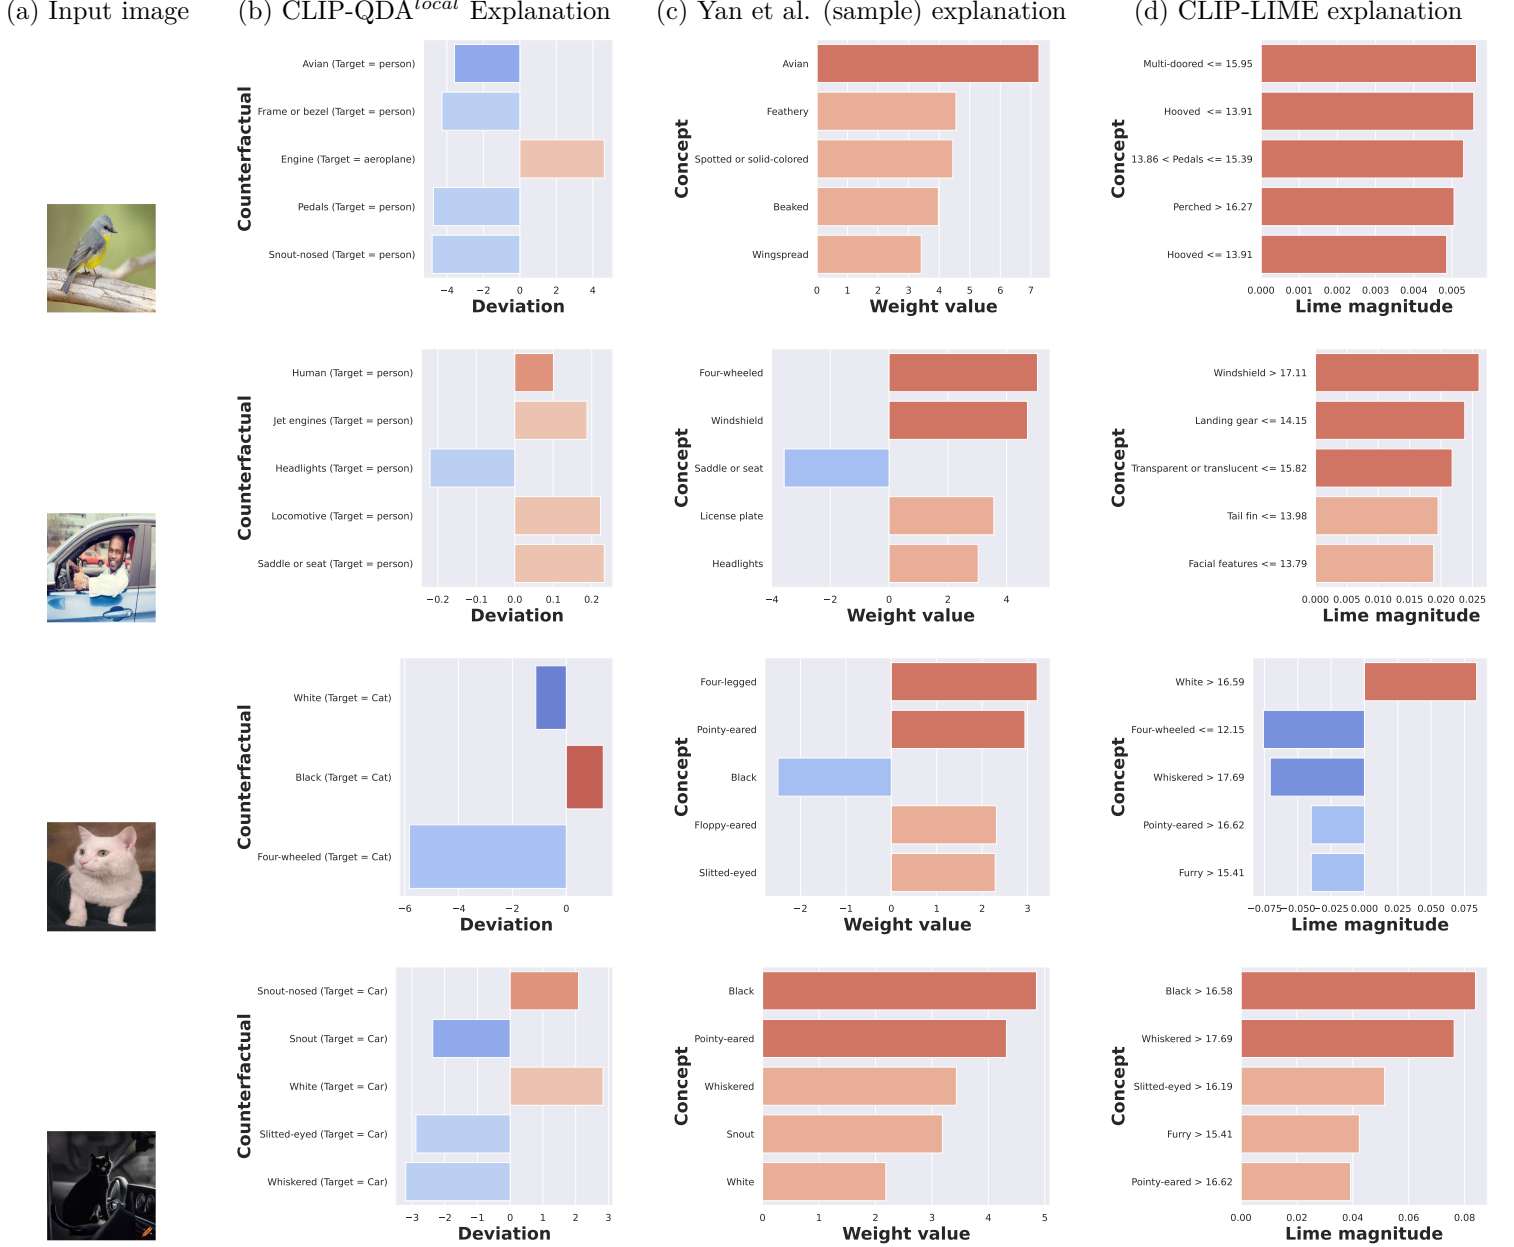


Figure 20: **Sample-wise explanations (concept level)**. The first two examples come from the Pascal-PART dataset, and the last two samples come from the Cats/Dogs/Cars dataset in the biased setup. Note that the classifier mislabeled the 2nd example as “car” and the 4th example as “car”.

(a) Input image

(b) CLIP-SHAP Explanation




Figure 21: **Sample-wise explanations (concept level).** The first two examples come from the Pascal-PART dataset, and the last two samples come from the Cats/Dogs/Cars dataset in the biased setup. Note that the classifier mislabeled the 2nd example as “car” and the 4th example as “car”.

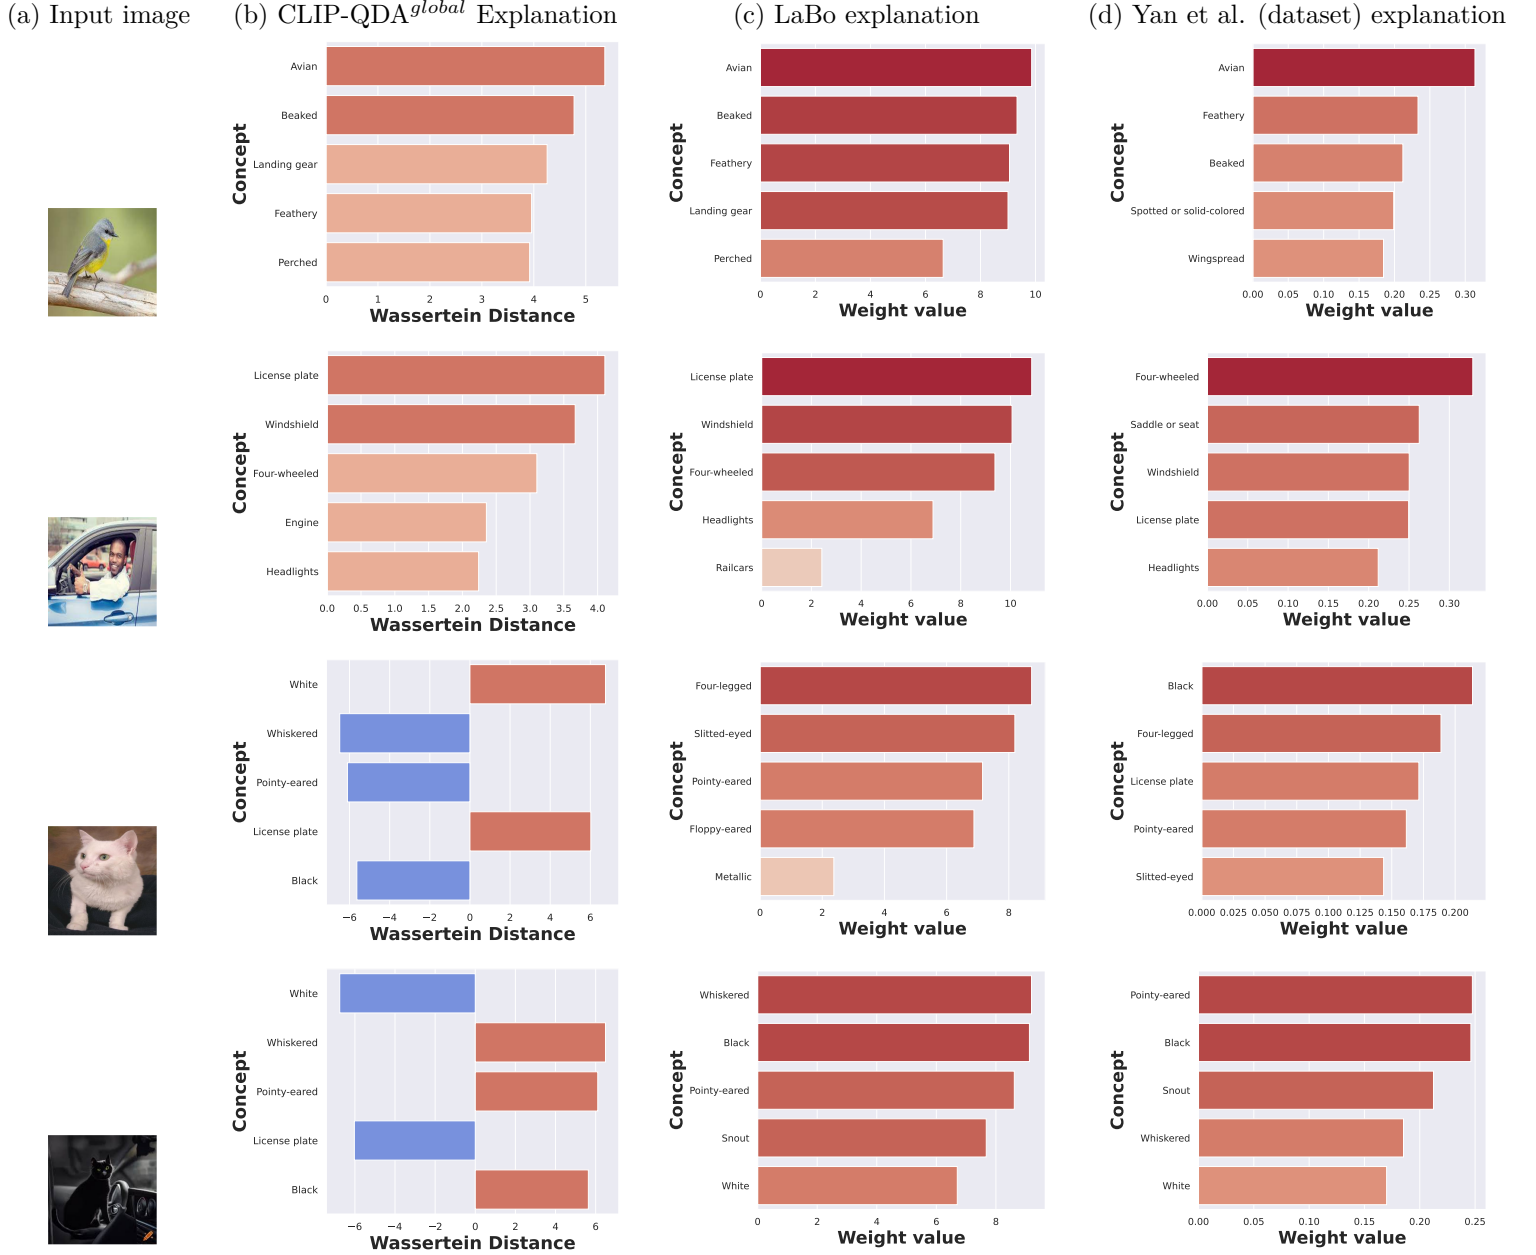


Figure 22: **Dataset-wise explanations.** The first two examples come from the PascalPART dataset, and the last two samples come from the Cats/Dogs/Cars dataset in the biased setup. Note that the classifier mislabeled the 2nd example as “car” and the 4th example as “car”.