

CLIPScope: Enhancing Zero-Shot OOD Detection with Bayesian Scoring

Hao Fu, Naman Patel, Prashanth Krishnamurthy, Farshad Khorrani
 {hf881, nkp269, prashanth.krishnamurthy, khorrani}@nyu.edu
 Department of Electrical and Computer Engineering*
 New York University, Brooklyn, NY, 11201

Abstract

Detection of out-of-distribution (OOD) samples is crucial for safe real-world deployment of machine learning models. Recent advances in vision language foundation models have made them capable of detecting OOD samples without requiring in-distribution (ID) images. However, these zero-shot methods often underperform as they do not adequately consider ID class likelihoods in their detection confidence scoring. Hence, we introduce CLIPScope, a zero-shot OOD detection approach that normalizes the confidence score of a sample by class likelihoods, akin to a Bayesian posterior update. Furthermore, CLIPScope incorporates a novel strategy to mine OOD classes from a large lexical database. It selects class labels that are farthest and nearest to ID classes in terms of CLIP embedding distance to maximize coverage of OOD samples. We conduct extensive ablation studies and empirical evaluations, demonstrating state of the art performance of CLIPScope across various OOD detection benchmarks. Code is available at <https://github.com/fu1001hao/CLIPScope>.

1. Introduction

Machine learning systems often encounter challenges when dealing with out-of-distribution (OOD) samples, which are not represented in the training distribution. These systems are built on the assumption that the data encountered during testing will mirror the training distribution. OOD detection [37] is vital for reliable deployment of machine learning systems. This detection process employs a confidence scoring mechanism, in alignment with prior work, that assesses whether a data point is part of the known training distribution or an OOD instance. It designates higher values to in-distribution (ID) samples and lower values to OOD samples, thus enabling their detection.

*This paper is supported in part by the Army Research Office under grant number W911NF-21-1-0155 and by the New York University Abu Dhabi (NYUAD) Center for Artificial Intelligence and Robotics, funded by Tamkeen under the NYUAD Research Institute Award CG010.

Traditional OOD detection methods predominantly focus on image data, overlooking the potential benefits of incorporating textual information. The concept of zero-shot OOD detection, which utilizes both textual and image information, was introduced by ZOC [4]. This approach involves using image captioning to generate potential OOD labels for input instances, followed by a zero-shot classification using CLIP [19] with both ID and generated OOD labels. A noted limitation of ZOC, as highlighted by [34] and [10], is its reduced effectiveness in generating relevant OOD labels for large ID class datasets, such as ImageNet-1K, resulting in decreased OOD detection accuracy. MCM [16] addresses OOD detection by employing the maximum logit of scaled softmax as a confidence score. However, MCM's reliance solely on ID class labels, without fully leveraging open-world textual information, limits its performance, particularly with challenging OOD samples, as reported by [34]. CLIPN [34] introduces an advancement by incorporating a CLIP-based negative-text encoder trained with additional datasets, including 3 million image-text pairs. This encoder is designed to understand negative prompts. For an input categorized into a specific ID class, CLIPN calculates the confidence scores using both the negative-text encoder and the original CLIP text encoder, with the final OOD confidence score being the product of these two. NegLabel [10] proposes a mining algorithm for extracting candidate OOD labels and a grouping strategy to compute the OOD score. However, the performance of NegLabel is sensitive to the size of negative label space and the percentile distance used in the mining procedure. This work proposes a method that is robust against these parameters.

In this paper, we introduce CLIPScope, a novel zero-shot method that leverages Bayesian inference to enhance confidence scoring for OOD detection. The model evidence or marginal likelihood for our Bayesian scoring is calculated based on the class membership likelihood, determined by the proportion of number of instances in specific ID classes to prior instances. CLIPScope operates on the principle that without prior knowledge of OOD sample distribution, the chance of an instance being OOD should be consistent

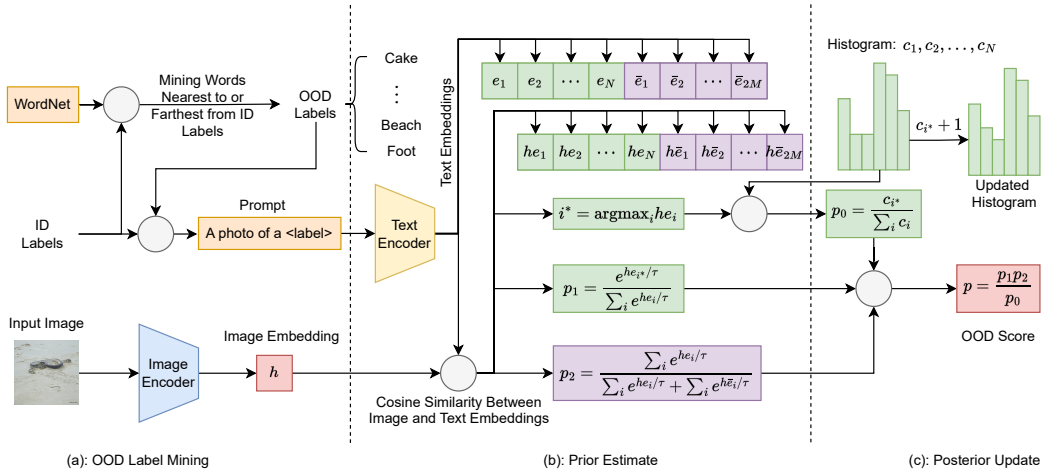



Figure 1. The process of CLIPScope involves three key stages: OOD label mining, prior estimate, and posterior update.

across ID classes. Therefore, classes with a higher frequency of instances are deemed more likely to include OOD samples. CLIPScope adjusts confidence scores accordingly, reducing them for OOD instances in these “high-frequency” classes while keeping scores for ID instances high. Furthermore, CLIPScope introduces a novel approach for mining potential OOD labels from WordNet lexical database [5] by considering both the closest and farthest words to ID labels. This strategy aims to maximize the coverage of the OOD sample space, thereby enhancing the robustness and effectiveness of zero-shot OOD detection, in contrast to methods like NegMining [10] that only consider the most distant words from ID labels. After all, OOD samples close to ID space are more challenging to detect than OOD samples far from ID space. To address the issue that certain nearest OOD labels may fall within ID label space by being synonyms of ID labels, CLIPScope uses another two components that do not require any mined OOD labels.

The CLIPScope framework for OOD detection, illustrated in Fig. 1 begins by mining potential OOD labels from a vocabulary database like WordNet [5], based on their lexical proximity to ID labels. This proximity is determined by the text embeddings of the words, with CLIPScope calculating the negative cosine similarities between each word’s embedding and the entire ID label space. Next, CLIPScope identifies the 5th percentile of these similar words for each ID label and selects the top M farthest words and top M nearest words as potential OOD labels, as illustrated in Fig. 2. This mining procedure conducted prior to the inference phase, establishes the OOD labels, which are thereafter fixed and do not require updates in subsequent operations. Next, using identified ID and OOD label text embeddings, and extracted image embedding for a given input, CLIPScope estimates the prior likelihood for an input being ID. It first identifies the ID class, i^* , closest to the image embedding, and then calculates p_1 , the confidence that the



input belongs to class i^* , and p_2 , the confidence that a given input belongs to any ID class, as shown in Fig. 1(b). The final step given in Fig. 1(c) involves updating the posterior with Bayesian inference. CLIPScope uses a histogram of prior instance occurrences to calculate the class likelihood p_0 of instances being classified into class i^* . The histogram is updated after every inference step of the input by incrementing occurrences of i^* . The final OOD confidence score $p = \frac{p_1 p_2}{p_0}$, thus is determined by the posterior update rule.

A key innovation of CLIPScope is it leverages the posterior information, particularly through a histogram of prior instance occurrences. This strategy significantly diverges from previous methods, which underutilized posterior information for improving OOD detection. Typically, for certain OOD samples, despite their anomalous nature, might be assigned high confidence scores by existing methods, making them difficult to detect. By incorporating posterior information, our approach greatly enhances the detection of these hard-to-distinguish OOD samples in subsequent analyses. Particularly, if hard-to-distinguish OOD samples initially bypass detection, they will contribute to an increase in the corresponding histogram bar, denoted as c_{i^*} . Since our overall confidence score is formulated as $p = \frac{p_1 p_2}{p_0}$ with $p_0 = \frac{c_{i^*}}{\sum_i c_i}$, an increment in c_{i^*} results in a proportional decrement in the overall confidence score p . This dynamic ensures that the confidence scores assigned to these chal-

lenging OOD samples gradually decrease over time, leading to their eventual detection. $p = \frac{p_1 p_2}{p_0}$ can be interpreted as a posterior update using Bayesian inference, where $p_1 p_2$ represents the joint confidence that an input instance is an ID sample and belongs to class i^* , and p_0 is the marginal likelihood, enhancing the effectiveness of OOD detection in complex and evolving data environments. It is important to mention that instead of depending on the full test samples, the calculation of p_0 for a particular sample relies on only the samples prior to the moment. Such a setting ensures CLIPScope to be realistic in real-world applications.

Our contributions are manifold and can be summarized as: 1) we introduce CLIPScope, a novel zero-shot OOD detection method that leverages the prior instances to perform posterior updates, thereby enhancing the efficacy of OOD sample detection; 2) we interpret CLIPScope as a Bayesian inference problem, providing a theoretical foundation that aligns with probabilistic reasoning; 3) we propose a novel criterion for the selection of OOD labels, which contributes to the improved accuracy of our method by maximizing coverage of OOD samples; and 4) we conduct comprehensive comparisons with prior work and ablation studies.

The remainder of this paper is organized as follows: Sec. 2 delves into the background of OOD detection and CLIP. In Sec. 3, we introduce CLIPScope and provide the motivation. Sec. 4 presents a series of comparisons and ablation studies. Sec. 5 provides a further discussion of our approach. Finally, Sec. 6 offers a conclusion to the paper.

2. Background

OOD Detection: Classic OOD detectors have included methods such as one-class SVM [25], decision-tree algorithms [2], and one-class nearest neighbor approaches [31]. In more recent developments, some studies have focused on training neural networks to serve as OOD score functions. Notable examples include Deep SVDD [22], OCGAN [17], GradCon [11], and Deep SAD [23]. Additionally, other research efforts have utilized neural networks for extracting feature embeddings, subsequently applying these embeddings to enhance OOD score functions. These post-hoc methods [18, 21, 24] do not change model parameters and use confidence scoring based on predicted softmax probabilities [7], energy functions [12, 15, 27, 28], or feature space distance [13]. With the rising popularity of multimodal vision-language models, they [4, 10, 16, 26, 34] have also been leveraged for OOD detection. This recent trend signifies an evolution in the field, integrating vision-language processing capabilities to improve OOD detection.

CLIP and Zero-Shot Classification: CLIP [19] is a vision-language model, trained on a vast dataset comprising approximately 400 million text-image pairs. The model is structured with two core components: a text encoder and an image encoder. For zero-shot classification tasks, CLIP

employs its image encoder to extract embeddings from a given image. Concurrently, it utilizes the text encoder to generate embeddings for candidate class labels. The model then calculates the cosine similarities between the image embedding and each of the text embeddings. CLIP selects the ID class with the highest cosine similarity as the most likely classification for the input image. This zero-shot classification mechanism also extends to other CLIP-based tasks. These include zero-shot object detection and zero-shot OOD detection, demonstrating CLIP’s versatility and broad applicability across various domains.

3. CLIPScope

Problem Formulation: Given a threshold γ , our goal is to design a score function g to determine if an input x is ID or OOD. Specifically, x is considered as ID if and only if $g(x) \geq \gamma$, i.e., $x \in \text{ID}$ if $g(x) \geq \gamma$; otherwise $x \in \text{OOD}$.

Ability: We assume the availability of the CLIP f , the ID labels $\mathcal{Y} = \{y_1, y_2, \dots, y_N\}$ with N representing the number of ID labels, and text from a large lexical database, such as WordNet. However, we do not have access to training ID images, ground-truth OOD labels, or prior knowledge about the distributions of ID or OOD samples in our approach.

Notation: Let $f(x, \mathcal{Y}) = y_i$ represent the process where the CLIP model f receives an image x and a set of ID labels $\mathcal{Y} = \{y_1, y_2, \dots, y_N\}$ as inputs for the zero-shot classification. Among all the ID labels, CLIP identifies y_i as the label that most closely matches the ground-truth class of x .

3.1. Bayes’ Theorem and Intuition of CLIPScope

Lemma 3.1. *With Bayes rule, the class likelihood of OOD samples is proportional to the global class likelihood, i.e., $\mathbb{P}(f(x, \mathcal{Y}) = y_i \mid x \in \text{OOD}) \propto \mathbb{P}(f(x, \mathcal{Y}) = y_i)$.*

Lemma 3.1 suggests that if certain ID classes are more frequently chosen for classification during inference than others, then OOD samples are also more likely to be classified into these same ID classes. Without additional information, it’s reasonable to expect these classes will continue to attract more OOD samples incorrectly as ID samples in the future. To improve future OOD sample rejection, we propose adapting the current OOD score function g to make it harder for samples classified into these higher-risk classes to achieve high confidence scores. We address this by introducing a Bayesian inference-aided confidence scorer that adjusts the likelihood of a sample being OOD based on its classification into specific ID classes, offering a more nuanced and effective approach for OOD detection.

Scoring with Bayesian Inference: Consider a random variable x that represents an instance, which can be either ID or OOD. At a specific moment t , x becomes a concrete

input x_t , for which we do a posterior update as follows:

$$\begin{aligned} \mathbb{P}(x_t \in \text{ID} \mid f(x_t, \mathcal{Y}) = y_i) &= \\ \frac{\mathbb{P}(f(x_t, \mathcal{Y}) = y_i \mid x_t \in \text{ID}) \mathbb{P}(x_t \in \text{ID})}{\mathbb{P}(f(x_t, \mathcal{Y}) = y_i)} &= \\ \frac{\mathbb{P}(f(x_t, \mathcal{Y}) = y_i \mid x_t \in \text{ID}) \mathbb{P}(x_t \in \text{ID})}{\mathbb{P}(f(x, \mathcal{Y}) = y_i)} \times C(x_t). \end{aligned} \quad (1)$$

Here, $C(x_t) = \frac{\mathbb{P}(f(x, \mathcal{Y}) = y_i)}{\mathbb{P}(f(x_t, \mathcal{Y}) = y_i)}$ which accounts for the inherent unpredictability of the sequence of incoming samples. This adjustment ensures a consistent and fair evaluation framework across different instances, crucial for developing generalized and reliable OOD detection.

Prior Estimation: The prior, $\mathbb{P}(x_t \in \text{ID})$, is computed based on similarities between image embedding and text embedding of both ID and OOD labels [4, 10].

Likelihood Estimation: We estimate the likelihood, $\mathbb{P}(f(x_t, \mathcal{Y}) = y_i \mid x_t \in \text{ID})$ based on similarities between image embedding and text embedding of only ID labels.

Marginal Likelihood Estimation: The marginal likelihood, $\mathbb{P}(f(x, \mathcal{Y}) = y_i)$ is updated based on prior sample predictions by maintaining a histogram. We update the frequency of the predicted ID labels. Consequently, $\mathbb{P}(f(x, \mathcal{Y}) = y_i)$ is calculated by dividing the number of samples for label y_i by the total samples. Thus, our confidence score function for x_t is formulated as

$$g(x_t) = \frac{\hat{\mathbb{P}}(f(x_t, \mathcal{Y}) = y_i \mid x_t \in \text{ID}) \hat{\mathbb{P}}(x_t \in \text{ID})}{\hat{\mathbb{P}}(f(x, \mathcal{Y}) = y_i)}. \quad (2)$$

This formulation is distinct from previous zero-shot methods, as it incorporates the class likelihood, $\hat{\mathbb{P}}(f(x, \mathcal{Y}) = y_i)$, into the confidence score. This inclusion of class likelihood adjusts the confidence score in a manner that makes it challenging for samples classified into classes with high likelihood to achieve high confidence scores.

Consider an OOD instance x_t with $g(x_t) > \gamma$ that is initially misclassified as an ID sample. As x_t reappears, the marginal likelihood $\hat{\mathbb{P}}(f(x, \mathcal{Y}) = y_i)$ associated with it incrementally increases. Since the prior and likelihood components of $g(x_t)$ remain constant for the same x_t , the increase in the marginal likelihood results in a gradual decrease of the confidence score $g(x_t)$, eventually causing it to fall below the threshold γ . While it is possible for x_t to maintain a high confidence score through infrequent appearances, this sporadic occurrence inherently limits the negative impact of such hard-to-distinguish OOD samples. In essence, the confidence score function $g(x_t)$ in (2) effectively mitigates the adverse effects of hard-to-distinguish OOD samples by either making them easier to identify over time or reducing their frequency of occurrence.

3.2. CLIPScope

The algorithm of CLIPScope, as detailed in Alg. 1, is structured into three main steps: OOD label mining (covered in lines 7-15), prior estimation (lines 18-21), and posterior updating (lines 22-26). Implementing CLIPScope requires only the ID labels \mathcal{Y} with $N = |\mathcal{Y}|$, CLIP model (the text encoder f^t and image encoder f^i), a user-defined threshold γ , and a histogram $\{c_i\}_{i=1}^N$ initialized such that each bin starts with a count of one, expressed as $c_i = 1 \forall i$.

Algorithm 1 CLIPScope

```

1: Input: Text encoder  $f^t$  and image encoder  $f^i$  (CLIP),
   ID labels  $\mathcal{Y}$ , out/in-distribution threshold  $\gamma$ , & class
   likelihood histogram  $\{c_i\}_{i=1}^N$  initialized with  $c_i = 1$ 
   for  $i = 1, 2, \dots, N$ , dataset  $\mathcal{X}$  where input image  $x \in \mathcal{X}$ 
2: /* Extract text embeddings of ID labels */
3: for  $y_i \in \mathcal{Y}$  do
4:   prompt = a photo of a <  $y_i$  >
5:    $e_i = f^t(\text{prompt}) / \|f^t(\text{prompt})\|_{L_2}$ 
6: end for
7: /* OOD label mining */
8:  $\mathcal{Y}^c = \{\text{nouns} \in \text{WordNet}\}$  // candidate OOD labels
9: for  $y_i^c \in \mathcal{Y}^c$  do
10:  prompt = a photo of a <  $y_i^c$  >
11:   $\bar{e}_i = f^t(\text{prompt}) / \|f^t(\text{prompt})\|_{L_2}$ 
12:   $d_i = \text{percentile}_\eta(\{-\text{sim}(\bar{e}_i, e_j)\}_{j=1}^N})$  // distance to  $\mathcal{Y}$ 
13: end for
14:  $\mathcal{Y}^- = \text{TopK}(\{d_i\}_{i=1}^{|\mathcal{Y}^c|}, \mathcal{Y}^c, M)$ 
15:  $\mathcal{Y}^- = \mathcal{Y}^- \cup \text{TopK}(\{-d_i\}_{i=1}^{|\mathcal{Y}^c|}, \mathcal{Y}^c, M)$ 
16: /* Bayesian OOD inference */
17: for all  $x \in \mathcal{X}$  do
18:   $h = f^i(x) / \|f^i(x)\|_{L_2}$ 
19:   $p_1 = \max_i \frac{e^{\text{sim}(h, e_i)/\tau}}{\sum_{j=1}^N e^{\text{sim}(h, e_j)/\tau}}$ 
20:   $p_2 = \frac{\sum_{j=1}^N e^{\text{sim}(h, e_j)/\tau}}{e^{\text{sim}(h, e_j)/\tau} + \sum_{j=1}^{2M} e^{\text{sim}(h, \bar{e}_j)/\tau}}$ 
21:   $p = p_1 p_2$ 
22: /* Posterior update with class likelihoods */
23:   $i^* = \arg \max_i \text{sim}(h, e_i)$  // nearest ID label  $y_{i^*}$ 
24:   $p_0 = \frac{c_{i^*}}{\sum_{i=1}^N c_i}$  // class likelihood of  $y_{i^*}$ 
25:   $p = p/p_0$  // adapting the confidence score
26:   $c_{i^*} = c_{i^*} + 1$  // update the histogram bar
27:  if  $p \geq \gamma$  then
28:     $x$  is ID
29:  else
30:     $x$  is OOD
31:  end if
32: end for

```

OOD Label Mining: This stage identifies OOD labels that are most distant from or closest to the ID labels in the embedding space to ensure maximal coverage of potential OOD samples. CLIPScope starts by gathering a set of

nouns from open-world sources (e.g., WordNet¹) as potential OOD labels, denoted as \mathcal{Y}^c with a cardinality of $|\mathcal{Y}^c|$.

The subsequent step involves extracting text embeddings for the OOD labels. A commonly used prompt format for this task is “a photo of a $\langle \rangle$ ”, with the blank filled by specific labels. For instance, if the label is {cat}, the generated prompts would be “a photo of a cat”². CLIPScope inputs each of these prompts into the text encoder f^t of the CLIP model to obtain the corresponding text embeddings. These embeddings are then normalized using the L_2 norm to focus solely on their directional attributes.

To determine the distance between OOD and ID labels, our work employs a percentile distance approach, similar to the one used in NegMining [10]. Specifically, for each candidate OOD label, we calculate the negative cosine similarities (denoted as sim) between its text embedding and the embeddings of all ID labels. This results in a total of N negative cosine similarities for each OOD label. The distance for each OOD label to the ID labels is then defined as the 100η -th percentile of these negative cosine similarities. In our approach, we set η to 0.05, aligning with the practice in NegMining. This choice of $\eta = 0.05$, as opposed to $\eta = 0$ which represents the minimum distance, enhances robustness against outlier ID labels. After computing these distances for all OOD labels, we select those whose distances rank in the top M largest or top M smallest, resulting in a total of $2M$ OOD labels, defined as:

$$\mathcal{Y}^- = \text{TopK}(\{d_i\}_{i=1}^{|\mathcal{Y}^c|}, \mathcal{Y}^c, M) \cup \text{TopK}(\{-d_i\}_{i=1}^{|\mathcal{Y}^c|}, \mathcal{Y}^c, M).$$

The goal of deriving these OOD labels, \mathcal{Y}^- , is to facilitate the approximation of $\hat{\mathbb{P}}(x_t \in \text{ID})$ as specified in (2).

Prior Estimate: This stage computes the confidence score of the input image with only prior knowledge, which includes the ID labels \mathcal{Y} and the OOD labels \mathcal{Y}^- . Denote $p_1 = \max_i \hat{\mathbb{P}}(f(x_t, \mathcal{Y}) = y_i \mid x_t \in \text{ID})$ as the confidence level of the zero-shot classification with only the ID labels. We begin by scaling the cosine similarities between the image embedding h of the input instance and the text embeddings e_j of each ID label y_j by a factor τ . Following this, we apply the softmax function to these scaled similarities to compute p_1 . The equation for this computation is:

$$p_1 = \max_i \frac{e^{\text{sim}(h, e_i)/\tau}}{\sum_{j=1}^N e^{\text{sim}(h, e_j)/\tau}}. \quad (3)$$

Denote $p_2 = \hat{\mathbb{P}}(x_t \in \text{ID})$ as the confidence level that an input image belongs to ID classes, distinct from the OOD classes, as defined in (2). Unlike p_1 , which is concerned with specific ID classes, the calculation of p_2 is broader, encompassing the likelihood of the image belonging to any

ID class rather than specific OOD classes. To calculate p_2 , we scale the cosine similarities between the image embedding h and the text embeddings of both ID and OOD labels by τ . Subsequently, these scaled similarities are processed through the softmax function. p_2 is obtained by summing the softmax values corresponding to all ID labels. The formula for this calculation is as follows:

$$p_2 = \frac{\sum_{j=1}^N e^{\text{sim}(h, e_j)/\tau}}{\sum_{j=1}^N e^{\text{sim}(h, e_j)/\tau} + \sum_{j=1}^{2M} e^{\text{sim}(h, \bar{e}_j)/\tau}} \quad (4)$$

where \bar{e}_j is the embedding of the OOD label with index j . This approach differs from the one used in NegLabel [10] in that it does not rely on any grouping strategy. This simplifies the preprocessing steps, making our method more straightforward and efficient.

The overall OOD confidence score p is the product of p_1 and p_2 . This method of combining two components is also seen in the CLIPN approach [34]. In CLIPN, the first component is identical to our p_1 . However, the second component in CLIPN, $1 - p^{no}$, is derived from the agreement level of their negative encoder with the CLIP model. A notable distinction between CLIPScope and CLIPN lies in the resources required. CLIPN necessitates an additional dataset containing image-text pairs for its operation. In contrast, CLIPScope relies primarily on WordNet, which provides only textual information. This difference makes our CLIPScope approach less demanding in terms of computational resources and memory usage.

Posterior Update: This step refines the confidence score by incorporating posterior knowledge during the inference stage. Previous methods have not fully considered the behavior of the CLIP model across all historical instances, a factor we identify as critical for enhancing detection accuracy. CLIPScope capitalizes on this insight by analyzing the observed behavior of CLIP on historical instances, using this data to fine-tune the confidence score. This process allows for more accurate OOD detection, as it adapts to the model’s performance over time and incorporates this evolving understanding into the assessment of new instances.

Initially, CLIPScope identifies the specific ID class i^* that is most closely aligned with the input image in terms of the cosine similarity between their respective embeddings,

$$i^* = \arg \max_i \text{sim}(h, e_i). \quad (5)$$

CLIPScope then proceeds to compute the likelihood p_0 . This likelihood represents the probability of the ID class i^* being the nearest to the input images based on historical instances. To achieve this, CLIPScope utilizes a histogram that tallies how frequently each ID class has historically been identified as the nearest to input images. Initially, this histogram is set with a count of one in all bins,

¹Other similar databases can also be employed.

²The prompt style can be adapted as needed.

reflecting the lack of prior posterior knowledge and implying an equal likelihood for all ID classes. As the system processes more input images, it accumulates posterior knowledge about the behavior of the CLIP model across these historical instances. The likelihood p_0 for the ID class i^* being the nearest to input images is approximated by the ratio of the count c_{i^*} (the number of times class i^* has been deemed the nearest) to the total count $\sum_{i=1}^N c_i$ of all historical instances, i.e., $p_0 = \frac{c_{i^*}}{\sum_{i=1}^N c_i}$. The confidence score p is divided by p_0 , i.e.,

$$p = \frac{p_1 p_2}{\frac{c_{i^*}}{\sum_{i=1}^N c_i}} = \frac{p_1 p_2}{p_0}. \quad (6)$$

(6) incorporates the evolving understanding of the CLIP model’s behavior, making the CLIPScope more adaptive and accurate in its assessments over time. The histogram is updated to reflect the latest input instance’s classification. This is achieved by incrementing the count for the identified nearest ID class i^* by one, i.e., $c_{i^*} = c_{i^*} + 1$. CLIPScope with (6) makes it challenging for instances (particularly hard-to-distinguish OOD instances) classified into classes with high likelihood to receive high OOD scores.

4. Experimental Results

Datasets and Metrics: We employ ImageNet-1K [9] as the primary ID dataset. For OOD datasets, we utilize iNaturalist [32], SUN [36], Places [38], and Textures [1]. These datasets are widely used in zero-shot OOD detection research for method evaluation. They encompass large-scale visual images across diverse scenes and semantics, offering a comprehensive testing ground. The OOD datasets are carefully curated to ensure no class overlap with the ID dataset, thus maintaining distinct separations between ID and OOD categories. Our method’s efficacy is evaluated using standard metrics prevalent in the field: the area under the receiver operating characteristic curve (AUROC) and the false positive rate at a 95% true positive rate (FPR95). A higher AUROC indicates greater detection accuracy, while a lower FPR95 denotes superior performance.

Implementation Details: For our zero-shot OOD detection, we utilize the CLIP model [19], specifically the CLIP-B/16 variant, which is a common choice in related studies. During the OOD label mining phase, WordNet serves as our chosen corpus, with $M = 5000$ set as the parameter, leading to a selection of $2M = 10000$ words as OOD labels. Consistent with the methodology in NegLabel, we set the percentile distance to $\eta = 0.05$ and the scaled temperature to $\tau = 0.01$. Our experiments are conducted using an NVIDIA Quadro RTX 5000 GPU.

Baseline Methods: The compared zero-shot OOD detection methods include Mahalanobis distance [13], energy score [15], ZOC [4], MCM [16], CLIPN [34], and

NegLabel [10]. We also compare our approach with methods that train or fine-tune OOD detection models with ID data, including MSP [7], ODIN [14], GradNorm [9], ViM [35], KNN [29], VOS [3], and NPOS [30]. MSP uses the highest softmax output as a confidence measure. The KNN approach adopts a contrastive learning framework tailored for OOD detection. ODIN enhances detection by merging temperature scaling with a gradient-based input preprocessing. GradNorm differentiates ID from OOD data by comparing the magnitudes of gradient norms, positing that ID data typically results in larger gradients. ViM introduces an additional logit in the output layer, which signifies the confidence level of an input being OOD after applying softmax. VOS generates virtual OOD samples for training effective detectors. NPOS fine-tunes CLIP’s image encoder with synthetic OOD data to enhance detection capabilities.

Sequence Order: The performance of CLIPScope is influenced by the sequence order in which ID and OOD samples are processed, as p_0 is calculated based on historical instances. We consider three different scenarios: forwarding, reversing, and random orders. In forwarding order, ID instances are processed before all OOD instances. In reversing order, ID instances follow after all OOD instances. In random order, ID and OOD instances are intermixed and appear in a random sequence. Note that in practical applications, the random order is more likely to happen without knowing the positions of OOD inputs. Therefore, we assessed CLIPScope across various random sequences by conducting five independent trials and calculating the average. However, we also provide the results on forwarding and reversing orders for certain cases to show the potential range of our approach’s performance.

Large-Scale Datasets: The performance of CLIPScope in the random order scenario is presented in Table 1. CLIPScope shows the highest AUROC and lowest FPR95 for all cases. Table 1 also shows the performance of CLIPScope in the forwarding and reversing order scenarios. CLIPScope demonstrates its highest performance in the reversing order, while the lowest performance is observed in the forwarding order. The reason behind this trend is that when calculating confidence scores for OOD samples in the forwarding order, the histogram c_i has been initialized by the distribution of ID samples. Therefore, the deviation of c_i from the actual OOD distribution is the largest among the three cases, leading to the worst performance. In contrast, the reversing order has the smallest distributional deviation shift, leading to the best performance. The distributional deviation shift of the random order case is in the middle and so is the performance. Notably, even under the forwarding order scenario, CLIPScope surpasses existing approaches across most cases, demonstrating its efficacy.

Domain-Shifted ID Datasets: CLIPScope demonstrates notable robustness when applied to domain-shifted

Table 1. Performance (%) of CLIPScope (ours) and other methods. The utilized CLIP model is ViT-B/16. The ID dataset is ImageNet-1k. \uparrow denotes that a higher value implies a better performance, whereas \downarrow signifies that a lower value represents a better performance.

OOD Dataset	iNaturalist		SUN		Places		Textures		Average	
Metric	AUROC \uparrow	FPR95 \downarrow	AUROC \uparrow	FPR95 \downarrow	AUROC \uparrow	FPR95 \downarrow	AUROC \uparrow	FPR95 \downarrow	AUROC \uparrow	FPR95 \downarrow
Methods That Require Training										
MSP	87.44	58.36	79.73	73.72	79.67	74.41	79.69	71.93	81.633	69.605
ODIN	94.65	30.22	87.17	54.04	85.54	55.06	87.85	51.67	88.803	47.748
GradNorm	72.56	81.50	72.86	82.00	73.70	80.41	70.26	79.36	72.345	80.818
ViM	93.16	32.19	87.19	54.01	83.75	60.67	87.18	53.94	87.820	50.203
KNN	94.52	29.17	92.67	35.62	91.02	39.61	85.67	64.35	90.970	42.188
VOS	94.62	28.99	92.57	36.88	91.23	38.39	86.33	61.02	91.188	41.320
NPOS	96.19	16.58	90.44	43.77	89.44	45.27	88.80	46.12	91.218	37.935
Zero-Shot (No Training Required)										
Mahalanobis	55.89	99.33	59.94	99.41	65.96	98.54	64.23	98.46	61.505	98.935
Energy	85.09	81.08	84.24	79.02	83.38	75.08	65.56	93.65	79.568	82.208
ZOC	86.09	87.30	81.20	81.51	83.39	73.06	76.46	98.90	81.785	85.193
MCM	94.59	32.20	92.25	38.80	90.31	46.20	86.12	58.50	90.818	43.925
CLIPN	95.27	23.94	93.93	26.17	92.28	33.45	90.93	40.83	93.103	31.098
NegLabel	99.49	1.91	95.49	20.53	91.64	35.59	90.22	43.56	94.210	25.398
Ours (Random)	99.61	1.29	96.77	15.56	93.54	28.45	91.41	38.37	95.301	20.883
Ours (Forward)	99.64	1.42	96.1	17.6	92.73	30.16	89.76	42.35	94.558	22.883
Ours (Reverse)	99.60	1.28	97.34	13.52	94.20	26.32	93.04	34.41	96.045	18.883

ImageNet datasets. In our evaluations, we utilized ImageNet-A [8], ImageNet-R [6], ImageNet-Sketch [33], and imageNet-V2 [20] as the ID datasets. The results of these tests are visually represented in Fig.3. In these domain-shifted scenarios, CLIPScope shows comparable performance to NegLabel. This highlights the method’s capability to adapt to various data distributions effectively. In contrast, MCM exhibits a noticeable decrease in performance on these domain-shifted ID datasets.

Contribution of p_0 , p_1 , and p_2 : We conducted an ablation analysis to assess the individual and combined effectiveness of the components prior (p_0), likelihood (p_1), and marginal likelihood (p_2). The results of this analysis are detailed in Table 2. The findings reveal that the incorporation of p_0 significantly enhances OOD detection accuracy, even when it is the sole component used in conjunction with either p_1 or p_2 as the numerator in the confidence score calculation. This underscores the vital role of p_0 in the algorithm. Furthermore, the results indicate that employing all three components in the confidence score function leads to best performance for most cases, both in terms of AUROC and FPR95. The data in the row $g = p_1/p_0$ illustrate that CLIPScope distinguishes itself from NegLabel by remaining effective without using any negative OOD labels.

Contribution of Nearest and Farthest OOD Labels: We explored variations of the label selection process. These variations included selecting only the nearest OOD labels and only the farthest OOD labels. The results are presented in Table 2. The results indicate that both farthest and nearest OOD labels help contribute to detecting OOD samples.

Therefore, combining them achieves the best or second-best performance across all tested scenarios. As stated previously, we include labels close to the ID labels within the OOD labels to maximize coverage of the OOD label space. Table 2 demonstrates that using both nearest and farthest OOD labels outperforms using only the farthest OOD labels in three out of four OOD datasets (iNaturalist, SUN, and Places). This is likely because the model’s ability to distinguish between ID and OOD samples is enhanced by capturing a more diverse range of OOD examples.

Noisy OOD Labels: Our confidence scorer contains three components: p_0 , p_1 , and p_2 . Neither p_0 nor p_1 is influenced by OOD labels. As illustrated in Table 2, our approach is effective with the confidence score $g = p_1/p_0$, being inherently robust against the presence of noisy OOD labels that might fall within ID space. Moreover, we conducted targeted experiments by constructing OOD labels from a mix of 5000 nearest OOD, 5000 farthest OOD, and 1000 ID labels, intentionally creating conditions where OOD labels intersect with ID labels. The result is shown in the bottom of Table 2. Performance degradation is not apparent because (1) the number of OOD labels is reasonably higher than the number of ID labels. And (2) only p_2 is affected when ID labels are added to the OOD labels. However, the potential inaccuracy of p_2 is mitigated by the other two components, p_0 and p_1 , since they are not affected and the performance of the scorer p_1/p_2 is reasonably accurate. A similar outcome was observed with an alternative label set comprising 10000 nearest OOD and 1000 ID labels, further reinforcing our method’s robustness. **Additional Eval-**




Figure 3. AUROC (%) on domain-shifted ID datasets. A higher AUROC implies a better performance.

Table 2. Performance (%) of CLIPScope with various confidence scores g and OOD label space \mathcal{Y}^- . The ID dataset is ImageNet-1k.

OOD Dataset	iNaturalist		SUN		Places		Textures	
Metric	AUROC \uparrow	FPR95 \downarrow	AUROC \uparrow	FPR95 \downarrow	AUROC \uparrow	FPR95 \downarrow	AUROC \uparrow	FPR95 \downarrow
Confidence Score g								
$g = p_1$	80.19	70.43	85.01	55.52	83.41	57.03	79.75	63.24
$g = p_1/p_0$	97.98	8.23	95.79	18.38	91.81	30.68	92.36	31.52
$g = p_2$	99.31	2.18	91.13	43.51	88.07	53.81	78.32	73.01
$g = p_2/p_0$	99.54	1.58	96.33	17.31	92.54	31.67	90.41	41.58
$g = p_1p_2$	99.19	3.15	91.49	34.65	88.89	43.09	81.03	60.9
$g = p_1p_2/p_0$	99.61	1.29	96.77	15.56	93.54	28.45	91.41	38.37
Different OOD Label Space \mathcal{Y}^-								
$M = 5000$, Nearest & Farthest	99.61	1.29	96.77	15.56	93.54	28.45	91.41	38.37
$M = 5000$, Nearest Only	99.69	1.01	96.56	15.98	93.06	30.02	91.18	39.27
$M = 5000$, Farthest Only	97.39	12.95	96.41	17.57	92.91	29.35	92.08	36.97
OOD Label Space \mathcal{Y}^- Overlap With ID Label Space \mathcal{Y}								
$M = 5000$, Nearest & Farthest $\cup \mathcal{Y}$	99.61	1.36	96.79	15.27	93.48	27.92	91.84	36.93
$M = 10000$, Nearest Only $\cup \mathcal{Y}$	99.66	1.08	96.87	14.56	93.31	28.91	91.35	38.61

uations are provided in our supplementary material.

5. Discussion

Comparisons with NegLabel: The core innovation of our approach lies in the Bayesian OOD scorer. This scorer leverages the mined coverage-maximizing OOD labels in conjunction with historical data through the marginal likelihood term p_0 . By dynamically adapting confidence scores in real-time, this unique feature enables our method to significantly improve OOD detection performance. As demonstrated in Table 2, the inclusion of p_0 alone contributes to a substantial increase in AUROC, highlighting the effectiveness and novelty of our approach. While p_2 bears similarity to the NegLabel score, it is not identical in that it does not rely on any grouping strategy, therefore simplifying the preprocessing steps and removing need for heuristic computations. Moreover, our method remains robust even without the p_2 term, as evidenced by the results in the $g = p_1/p_0$ row of Table 2. The strategic adaptation enabled by p_0 represents a significant advancement in the field of zero-shot OOD detection by introducing real-time adaptability to previously static detection frameworks.

Benefits over Fine-Tuning Methods: While fine-tuning approaches for OOD detection can leverage historical test samples to improve performance, our method offers signif-

icant advantages in both efficiency and practicality. Unlike training-based methods, our approach eliminates the need for ground-truth labels from historical test data and operates with minimal memory requirements by utilizing only histogram information from CLIP’s empirical outputs. Furthermore, by avoiding the computational overhead of fine-tuning numerous model parameters, our method achieves faster processing times, making it particularly suitable for applications that demand frequent updates. **Additional discussions** are given in our supplementary material.

6. Conclusion

We have presented CLIPScope to detect OOD samples in a zero-shot manner, leveraging Bayesian inference to enhance confidence scoring. It uses information from prior instances as evidence for posterior updates to effectively identify OOD samples. Furthermore, we have developed a unique OOD label mining strategy for CLIPScope scoring, enhancing OOD sample coverage by utilizing the closest and farthest WordNet labels based on their CLIP embedding proximity to ID classes. We perform extensive comparisons with prior approaches, demonstrating excellent performance, and conduct ablation studies to validate the effectiveness of CLIPScope.

References

- [1] Mircea Cimpoi, Subhansu Maji, Iasonas Kokkinos, Sammy Mohamed, and Andrea Vedaldi. Describing textures in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3606–3613, 2014. 6
- [2] Francesco De Comit e, Franois Denis, R emi Gilleron, and Fabien Letouzey. Positive and unlabeled examples help learning. In *Proceedings of International Conference on Algorithmic Learning Theory*, pages 219–230, 1999. 3
- [3] Xuefeng Du, Zhaoning Wang, Mu Cai, and Sharon Li. Towards unknown-aware learning with virtual outlier synthesis. In *Proceedings of International Conference on Learning Representations*, 2022. 6
- [4] Sepideh Esmailpour, Bing Liu, Eric Robertson, and Lei Shu. Zero-shot out-of-distribution detection based on the pre-trained model clip. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 6568–6576, 2022. 1, 3, 4, 6
- [5] Christiane Fellbaum. *WordNet: An electronic lexical database*. MIT press, 1998. 2
- [6] Dan Hendrycks, Steven Basart, Norman Mu, Saurav Kadavath, Frank Wang, Evan Dorundo, Rahul Desai, Tyler Zhu, Samyak Parajuli, Mike Guo, et al. The many faces of robustness: A critical analysis of out-of-distribution generalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8340–8349, 2021. 7
- [7] Dan Hendrycks and Kevin Gimpel. A baseline for detecting misclassified and out-of-distribution examples in neural networks. In *Proceedings of International Conference on Learning Representations*, 2017. 3, 6
- [8] Dan Hendrycks, Kevin Zhao, Steven Basart, Jacob Steinhardt, and Dawn Song. Natural adversarial examples. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15262–15271, 2021. 7
- [9] Rui Huang, Andrew Geng, and Yixuan Li. On the importance of gradients for detecting distributional shifts in the wild. In *Proceedings of Advances in Neural Information Processing Systems*, pages 677–689, 2021. 6
- [10] Xue Jiang, Feng Liu, Zhen Fang, Hong Chen, Tongliang Liu, Feng Zheng, and Bo Han. Negative label guided OOD detection with pretrained vision-language models. In *Proceedings of International Conference on Learning Representations*, 2024. 1, 2, 3, 4, 5, 6
- [11] Gukyeong Kwon, Mohit Prabhushankar, Dogancan Temel, and Ghassan AlRegib. Backpropagated gradient representations for anomaly detection. In *Proceedings of European Conference on Computer Vision*, pages 206–226, 2020. 3
- [12] Yann LeCun, L eon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. 3
- [13] Kimin Lee, Kibok Lee, Honglak Lee, and Jinwoo Shin. A simple unified framework for detecting out-of-distribution samples and adversarial attacks. In *Proceedings of Advances in Neural Information Processing Systems*, volume 31, pages 1–9, 2018. 3, 6
- [14] Shiyu Liang, Yixuan Li, and R Srikant. Enhancing the reliability of out-of-distribution image detection in neural networks. In *Proceedings of International Conference on Learning Representations*, 2018. 6
- [15] Weitang Liu, Xiaoyun Wang, John Owens, and Yixuan Li. Energy-based out-of-distribution detection. In *Proceedings of Advances in Neural Information Processing Systems*, volume 33, pages 21464–21475, 2020. 3, 6
- [16] Yifei Ming, Ziyang Cai, Jiuxiang Gu, Yiyu Sun, Wei Li, and Yixuan Li. Delving into out-of-distribution detection with vision-language representations. In *Proceedings of Advances in Neural Information Processing Systems*, volume 35, pages 35087–35102, 2022. 1, 3, 6
- [17] Pramuditha Perera, Ramesh Nallapati, and Bing Xiang. Ogan: One-class novelty detection using gans with constrained latent representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2898–2906, 2019. 3
- [18] Chen Qiu, Aodong Li, Marius Kloft, Maja Rudolph, and Stephan Mandt. Latent outlier exposure for anomaly detection with contaminated data. In *Proceedings of International Conference on Machine Learning*, pages 18153–18167, 2022. 3
- [19] 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 *Proceedings of International Conference on Machine Learning*, pages 8748–8763, 2021. 1, 3, 6
- [20] Benjamin Recht, Rebecca Roelofs, Ludwig Schmidt, and Vaishal Shankar. Do imagenet classifiers generalize to imagenet? In *Proceedings of International Conference on Machine Learning*, pages 5389–5400, 2019. 7
- [21] Tal Reiss, Niv Cohen, Liron Bergman, and Yedid Hoshen. Panda: Adapting pretrained features for anomaly detection and segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2806–2814, 2021. 3
- [22] Lukas Ruff, Robert Vandermeulen, Nico Goernitz, Lucas Deecke, Shoaib Ahmed Siddiqui, Alexander Binder, Emmanuel M uller, and Marius Kloft. Deep one-class classification. In *Proceedings of International Conference on Machine Learning*, pages 4393–4402, 2018. 3
- [23] Lukas Ruff, Robert A. Vandermeulen, Nico G ornitz, Alexander Binder, Emmanuel M uller, Klaus-Robert M uller, and Marius Kloft. Deep semi-supervised anomaly detection. In *Proceedings of International Conference on Learning Representations*, 2020. 3
- [24] Mohammadreza Salehi, Niousha Sadjadi, Soroosh Baselizadeh, Mohammad H Rohban, and Hamid R Rabiee. Multiresolution knowledge distillation for anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14902–14912, 2021. 3
- [25] Bernhard Sch olkopf, John C Platt, John Shawe-Taylor, Alex J Smola, and Robert C Williamson. Estimating the support of a high-dimensional distribution. *Neural Computation*, 13(7):1443–1471, 2001. 3

- [26] Yang Shu, Xingzhuo Guo, Jialong Wu, Ximei Wang, Jianmin Wang, and Mingsheng Long. Clipood: Generalizing clip to out-of-distributions. In *Proceedings of International Conference on Machine Learning*, 2023. 3
- [27] Yiyu Sun, Chuan Guo, and Yixuan Li. React: Out-of-distribution detection with rectified activations. In *Proceedings of Advances in Neural Information Processing Systems*, pages 144–157, 2021. 3
- [28] Yiyu Sun and Yixuan Li. Dice: Leveraging sparsification for out-of-distribution detection. In *Proceedings of the European Conference on Computer Vision*, pages 691–708, 2022. 3
- [29] Yiyu Sun, Yifei Ming, Xiaojin Zhu, and Yixuan Li. Out-of-distribution detection with deep nearest neighbors. In *Proceedings of International Conference on Machine Learning*, pages 20827–20840, 2022. 6
- [30] Leitian Tao, Xuefeng Du, Jerry Zhu, and Yixuan Li. Non-parametric outlier synthesis. In *Proceedings of The International Conference on Learning Representations*, 2022. 6
- [31] David Martinus Johannes Tax. *One-class classification: Concept learning in the absence of counter-examples*. PhD thesis, Technische Universiteit Delft, 2002. 3
- [32] Grant Van Horn, Oisin Mac Aodha, Yang Song, Yin Cui, Chen Sun, Alex Shepard, Hartwig Adam, Pietro Perona, and Serge Belongie. The inaturalist species classification and detection dataset. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8769–8778, 2018. 6
- [33] Haohan Wang, Songwei Ge, Zachary Lipton, and Eric P Xing. Learning robust global representations by penalizing local predictive power. In *Proceedings of Advances in Neural Information Processing Systems*, 2019. 7
- [34] Hualiang Wang, Yi Li, Huifeng Yao, and Xiaomeng Li. Clipn for zero-shot ood detection: Teaching clip to say no. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1802–1812, 2023. 1, 3, 5, 6
- [35] Haoqi Wang, Zhizhong Li, Litong Feng, and Wayne Zhang. Vim: Out-of-distribution with virtual-logit matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4921–4930, 2022. 6
- [36] Jianxiong Xiao, James Hays, Krista A Ehinger, Aude Oliva, and Antonio Torralba. Sun database: Large-scale scene recognition from abbey to zoo. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3485–3492, 2010. 6
- [37] Jingkan Yang, Kaiyang Zhou, Yixuan Li, and Ziwei Liu. Generalized out-of-distribution detection: A survey. *arXiv preprint arXiv:2110.11334*, 2021. 1
- [38] 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*, 40(6):1452–1464, 2017. 6