

GHOST: Gaussian Hypothesis Open-Set Technique

Ryan Rabinowitz¹, Steve Cruz², Manuel Günther³, and Terrance E. Boult¹

¹Vision and Security Technology Lab, University of Colorado Colorado Springs,

²Computer Vision Research Lab, University of Notre Dame,

³Department of Informatics, University of Zurich

{rrabinow, tboult}@uccs.edu, stevecruz@nd.edu, guenther@ifi.uzh.ch

Abstract

Evaluations of large-scale recognition methods typically focus on overall performance. While this approach is common, it often fails to provide insights into performance across individual classes, which can lead to fairness issues and misrepresentation. Addressing these gaps is crucial for accurately assessing how well methods handle novel or unseen classes and ensuring a fair evaluation. To address fairness in Open-Set Recognition (OSR), we demonstrate that per-class performance can vary dramatically. We introduce Gaussian Hypothesis Open Set Technique (GHOST), a novel hyperparameter-free algorithm that models deep features using class-wise multivariate Gaussian distributions with diagonal covariance matrices. We apply Z-score normalization to logits to mitigate the impact of feature magnitudes that deviate from the model’s expectations, thereby reducing the likelihood of the network assigning a high score to an unknown sample. We evaluate GHOST across multiple ImageNet-1K pre-trained deep networks and test it with four different unknown datasets. Using standard metrics such as AUOSCR, AUROC and FPR95, we achieve statistically significant improvements, advancing the state-of-the-art in large-scale OSR. Source code is provided online.

Code Repository — <https://github.com/Vastlab/GHOST>

Introduction

When deploying deep neural networks (DNNs) in real-world environments, they must handle a wide range of inputs. The “closed-set assumption,” prevalent in most evaluations, represents a significant limitation of traditional recognition-oriented machine learning algorithms (Scheirer et al. 2012). This assumption presumes that the set of possible classes an algorithm will encounter is known a priori, meaning that these algorithms are not evaluated for robustness against samples from previously unseen classes. Open-Set Recognition (OSR) challenges this assumption by requiring systems to manage errors from unknown classes during testing.

Often, OSR is performed by thresholding on confidence (Hendrycks and Gimpel 2017; Vaze et al. 2022) or having an explicit “other” class (Ge, Demyanov, and Garnavi 2017) and computing overall performance, ignoring the effects of per-class performance differentials (Li, Wu, and Su 2023).

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

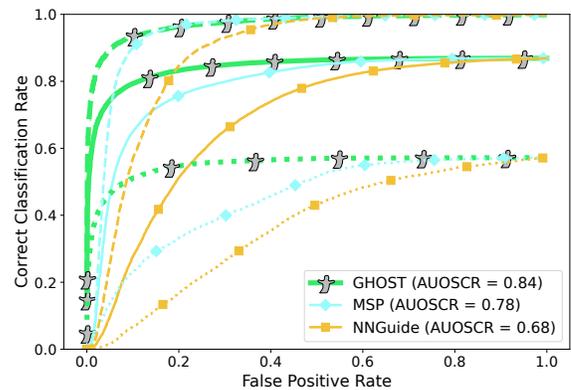


Figure 1: CLASS-WISE OPEN-SET RECOGNITION. OSR comparison using the MAE-H architecture with OpenImage-O as unknowns. Overall performance is the solid line; Average performance on easy (top 10%) and hard (bottom 10%) classes shown as dashed/dotted lines, respectively. We compare GHOST with Maximum Softmax Probability (MSP) and NNGuide. Also, we show the area under the curve (AUC) of each method’s overall OSR. In each setting, GHOST maintains correct classification as FPR rate decreases while others fall dramatically; hence, GHOST maintains fairness in difficult cases while improving OSR.

However, evaluating recognition systems under OSR conditions is crucial for understanding their behavior in real-world scenarios. This paper shows that as more unknowns are rejected, there is great variation in per-class accuracy, which could lead to unfair treatment of underperforming classes, see Fig. 1.

Recently, research has followed two primary methodologies for adapting DNNs to OSR problems: (1) training processes that enhance feature spaces and (2) post-processing techniques applied to pre-trained DNNs to adjust their outputs for identifying known and unknown samples (Roady et al. 2020). Although OSR training methods have occasionally proven effective (Zhou, Ye, and Zhan 2021; Miller et al. 2021; Dhamija, Günther, and Boult 2018), their application is complex due to the evolving nature of DNNs and the specific, often costly training requirements for each. If different DNNs are trained in various ways, why should a single OSR training

technique be universally applicable? Furthermore, if an OSR technique is specific to a particular DNN, its value diminishes as state-of-the-art DNNs evolve. In contrast, post-processing methods, such as leveraging network embeddings (Bendale and Boulton 2016), are more straightforward to implement and can be applied to almost any DNN. These methods avoid the complexities associated with training techniques and focus instead on evaluating performance. Thus, the challenge becomes: *how can various DNNs designed with a closed-set assumption be adapted for OSR problems?*

Initial post-processing OSR algorithms (Bendale and Boulton 2016; Rudd et al. 2017) used distance metrics in high-dimensional feature spaces to relate inference samples to known class data from training. However, choosing appropriate hyperparameters, such as a distance metric, is not straightforward, particularly for networks trained without distance metric learning, leading to an expensive parameter search. Further, large-scale datasets like ImageNet (Deng et al. 2009) and small-scale splits (Neal et al. 2018; Perera et al. 2020; Yang et al. 2020; Geng, Huang, and Chen 2020; Zhou, Ye, and Zhan 2021) often lack suitable train-validation-test splits for a fair parameter search.

Additionally, a major limitation with prior evaluations is their emphasis on overall performance, rather than ensuring robust performance for each individual class. This focus can obscure significant disparities between classes, leading to an incomplete understanding of the algorithm’s effectiveness and potentially resulting in unfair treatment of some classes. For example, an algorithm might achieve high overall accuracy but fail to recognize rare or challenging classes, which is critical for applications requiring high precision across all classes. Such evaluations can misrepresent the algorithm’s performance for underrepresented or underperforming classes, which may be overlooked when only aggregate metrics are considered. This lack of detailed analysis can lead to skewed evaluations, where the model’s weaknesses in specific areas are not addressed, ultimately affecting its real-world applicability and fairness. While fairness is not a major concern for ImageNet-1K, the dataset used herein, we consider it a reasonable proxy for operational open-set problems due to its size and widespread use as a feature extractor or for fine-tuning domain-specific models.

We propose a novel post-processing OSR algorithm, the Gaussian Hypothesis Open-Set Technique (GHOST), which uses per-class multivariate Gaussian models with diagonal covariance of DNN embeddings to reduce network overconfidence for unknown samples. The use of per-class modeling is crucial for ensuring fairness across all classes. By modeling each feature dimension separately for each class, GHOST evaluates each class on its own merits, rather than grouping them together. This technique helps address the challenge of handling the worst-performing classes fairly and reduces the risk of the model being overly confident about samples from these difficult classes. Importantly, GHOST eliminates the need for hyperparameters, simplifying the application of OSR techniques for end-users. Our novel GHOST algorithm improves traditional OSR measures and fairness, achieving a win-win outcome in line with recent fairness goals presented by Islam, Pan, and Foulds (2021); Li, Wu, and Su (2023).

In summary, our main contributions are:

- We introduce GHOST, a novel, state-of-the-art, hyperparameter-free post-processing algorithm that models per-feature, per-class distributions to improve per-class OSR.
- We present an extensive experimental analysis that adapts both the previous and recent state-of-the-art methods while evaluating multiple state-of-the-art DNNs, with results showing that GHOST is statistically significantly better on both global OSR and OOD metrics.
- We provide the first fairness analysis in OSR, identify significant per-class differences in large-scale OSR evaluations, and demonstrate that GHOST improves fairness.

Related Work

Some methods have been proposed to improve the training of DNNs for OSR (Zhang et al. 2022; Xu, Shen, and Zhao 2023; Wan et al. 2024; Wang et al. 2024; Li et al. 2024a,b; Sensoy, Kaplan, and Kandemir 2018). We do not consider these as direct competitors, as they go beyond statistical inference and train reconstruction models and use generative techniques or other additional training processes. Post-processing methods, including GHOST, can all use better features, but as Vaze et al. (2022) pointed out, better closed-set classifiers improve performance more and are continuing to evolve rapidly, so our focus is on post-processing algorithms.

Post-hoc approaches are well-explored in out-of-distribution (OOD) detection. Moreover, they are used in various practical settings requiring large pre-trained networks. The first attempt to adapt pre-trained DNNs for OSR using statistical inference on representations extracted from a pre-trained backbone was made by Bendale and Boulton (2016). They sought to replace the SoftMax layer, which is problematic for OSR, with OpenMax. OpenMax computes the centroid for each known class from training data and uses Extreme Value Theory to fit Weibull distributions over the distance from the centroid to the training samples. During inference, the probabilities that a sample belongs to a known class are converted to probabilities of unknown, which are summed and effectively form an additional category representing the probability of unknown. The Extreme Value Machine (EVM) proposed by Rudd et al. (2017) is another OSR system based on statistical inference using distance between samples. It finds a set of extreme vectors in each training-set class and fits a Weibull distribution on distances between them and closest samples of other “negative” classes in high-dimensional feature space. Both systems compute distances in high-dimensional space, so a practitioner must select a distance metric that applies to their DNN backbone. This process often requires a search over possible metrics and other algorithm-related hyperparameters. We might consider these methods to be direct competitors as they employ straightforward statistical measures to recognize known samples, but large scale evaluation shows they are not as effective as some simple baselines (Bisgin et al. 2024).

Using network outputs to reject unknowns is widely used, and Hendrycks and Gimpel (2017); Hendrycks et al. (2022) showed that thresholding on Maximum Softmax Probability (MSP) or Maximum Logits (MaxLogit) from a closed-set

DNN provides good baselines for OSR. In addition, Vaze et al. (2022) went so far as to argue that good closed-set classifiers with logit-based thresholding are sufficient for OSR. We also consider the popular energy-based OOD detection technique (Liu et al. 2020), which computes energy based on the logit vector of a DNN (this method’s performance is subpar, and so it is relegated to the supplemental). A recent collection of OOD methods, OpenOOD (Yang et al. 2022; Zhang et al. 2023), has compared many of these post-hoc methods using recent, large-scale datasets. Herein, we consider only the best performing: Nearest Neighbor Guidance (NNGuide) (Park, Jung, and Teoh 2023) for OOD detection (others in the supplemental). This method scales the confidence output from a DNN’s classification layer by the sample’s cosine similarity to a subset of training data, and is currently leading in the OpenOOD ImageNet-1K leaderboard¹ and so we use it as a primary comparison. We show that GHOST normalization, which does not need a reference set, improves performance overall, setting a new standard for large-scale OSR and OOD.

Approach

A Gaussian Hypothesis

The first works on open-set recognition and open-set deep networks (Scheirer et al. 2012; Scheirer, Jain, and Boult 2014; Bendale and Boult 2016; Rudd et al. 2017) all focused on the most distant points within a class or the values at the class boundaries. Hence, it is natural that they employed extreme-value theory as their underlying model. Having evaluated many of these EVT-based approaches in practical settings, we found a few significant difficulties: These methods are sensitive to outliers/mislabeled data (due to their reliance on a small percentage of extreme data) and have a high cost and sensitivity of tuning their hyperparameters. A final difficulty with this approach is reducing the high-dimensional features into a 1-dimensional distance, typically Euclidean or Cosine.

Features within a DNN are learned using large amounts of data. Various papers have shown that, with some mild assumptions, convergence in a two-layer network follows a central-limit theory (Sirignano and Spiliopoulos 2020), and using a mean-field analysis that was extended to some older deep network architectures (Lu et al. 2020) – so there are inherently some reasons to hypothesize Gaussian models.

We start by summarizing the main simple NN central-limit theory of (Sirignano and Spiliopoulos 2020), which indicates that for a large number M of neurons, the empirical distribution of the neural network’s parameters behaves as Gaussian distribution. Their theorems show that, given their assumptions, the empirical distribution of the parameters behaves as a Gaussian distribution with a specific variance-covariance structure. Central to the proofs of these theories is mean-field theory, and the convergence of the parameters to the mean follows from the central limit theorem. These mean-field distributional convergence results were then extended to some older deep networks (Lu et al. 2020), but extending to new networks is complex. We believe that empirical testing, as we do in our experiments, is a sufficient and much easier way to evaluate the Gaussian hypothesis for any new network.

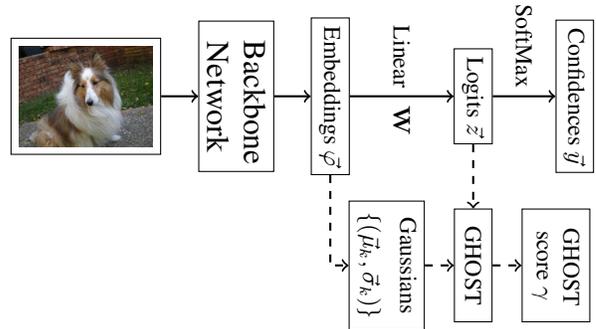


Figure 2: GHOST SCORES. In a pre-trained network indicated with solid arrows, an image is presented to the backbone network, which extracts deep feature embeddings $\vec{\varphi}$ that are then processed with a Linear layer to logits \vec{z} , and further with SoftMax to probabilities \vec{y} . For training GHOST, we extract embeddings from training data, from which we model class-wise multivariate Gaussian distributions. During evaluation, the Gaussian of the predicted class are used to turn the embeddings $\vec{\varphi}$ into z-score, which is used together with the maximum logit z_k to compute the GHOST score γ .

Inspired by those theories, we hypothesize that similarly, when input is from a class seen in training, each value in the network can be reasonably approximated by a multivariate Gaussian and that, importantly, out-of-distribution samples would be more likely to be inconsistent with the resulting Gaussian model. While the theories of Sirignano and Spiliopoulos (2020); Lu et al. (2020) are about the learnable network parameters, we hypothesize that with a Gaussian distribution per parameter, after many layers of computation, for a set of inputs from a given class, the distribution of each embedding value may also be well modeled with a Gaussian. Critical in this hypothesis is that for at least the embedding $\vec{\varphi}$ as shown in Fig. 2, which is used to compute the per-class logits, Gaussian models are class-specific; Fig. 3 shows an example model with sample values for a known and outlier.

Due to the complexities and variations of modern DNN architectures, formally proving that this hypothesis is valid for every DNN is impractical and unlikely. Instead, we derive a technique from this hypothesis and apply it to the most well-performing, publicly available architectures to show its utility.

GHOST Training. Consider the network processing shown in Fig. 2. Given a training dataset $\mathcal{X} = \{(x_n, t_n) \mid 1 \leq n \leq N\}$ with N samples x_n and their class labels $1 \leq t_n \leq K$ representing K known classes. Here, we apply post-processing, so we assume the backbone network to be trained on the same K classes contained in \mathcal{X} . For each correctly-classified training sample, we use the backbone to extract its D -dimensional embedding $\vec{\varphi} \in \mathbb{R}^D$. For each class k , we model a multivariate Gaussian distribution with mean $\vec{\mu}_k$ and diagonal covariance $\vec{\sigma}_k$ from the samples of that class and collect these Gaussians for all classes as

¹<http://zjysteven.github.io/OpenOOD>

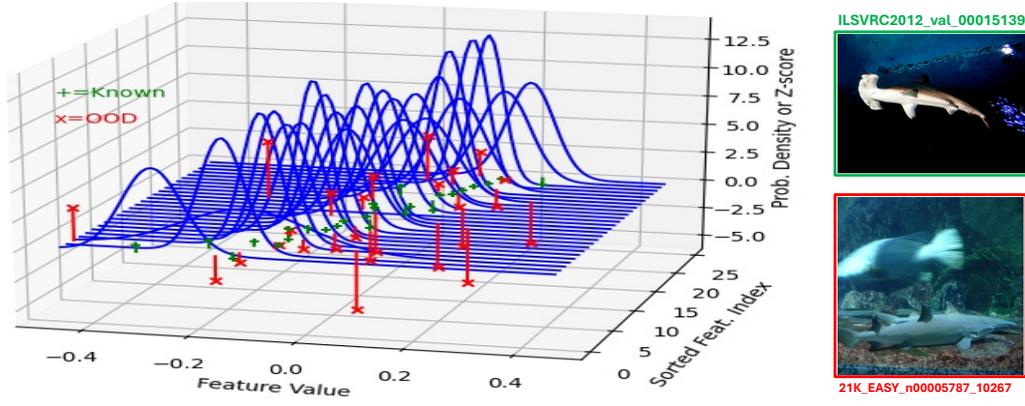


Figure 3: GHOST MODELING OF A MULTIVARIATE GAUSSIAN PER CLASS. Samples of Gaussians from the MAE-H network are shown on the left, sampled once every 30 dimensions. Dimensions were sorted on mean value to improve visibility, and the spread shows how some dimensions have greater variance than others. The plot also shows the value of per-dimension z-scores associated with a correctly classified hammerhead image (known in green) and an OOD example with a shark (red) misclassified as a hammerhead. The z-scores of the OOD example are much larger than those of the known.

$\mathcal{G} = \{(\vec{\mu}_k, \vec{\sigma}_k) \mid 1 \leq k \leq K\}$ via:

$$\begin{aligned} \vec{\mu}_k &= \frac{1}{N_k} \sum_{(x_n, t_n) \in \mathcal{X}} \mathbb{I}(k, t_n) \cdot \vec{\varphi}_n \\ \vec{\sigma}_k^2 &= \frac{1}{N_k - 1} \sum_{(x_n, t_n) \in \mathcal{X}} \mathbb{I}(k, t_n) \cdot (\vec{\mu}_k - \vec{\varphi}_n)^2 \end{aligned} \quad (1)$$

where the indicator function $\mathbb{I}(k, t_n)$ is 1 if the sample x_n belongs to class k and is correctly classified,² and N_k is the number of correctly-classified samples for class k . Hence, each feature dimension of each known class is modeled from its training data. Together, these Gaussian models are useful for differentiating unknown samples.

GHOST Inference. Building from the open-set theory by Bendale and Boulton (2016), we know a provably open-set algorithm is produced if the confidence scores are monotonically decreasing with the distance of a feature from a mean. To achieve this with our Gaussian Hypothesis, we combine our model with the intuition that there are significant deviations of DNN’s embedding magnitude when an unknown sample is encountered (Dhamija, Günther, and Boulton 2018; Cruz et al. 2024) and, thus, the embedding $\vec{\varphi}$ deviates from all class means, even though the angular direction might overlap with a certain class’ mean $\vec{\mu}_k$. For a given test sample, we first compute embedding $\vec{\varphi}$, logits \vec{z} and the predicted class \hat{k} as:

$$\hat{k} = \arg \max_k z_k. \quad (2)$$

We select the associated Gaussian $(\vec{\mu}_{\hat{k}}, \vec{\sigma}_{\hat{k}})$ to compute our z-score:

$$s = \sum_{d=1}^D \frac{|\varphi_d - \mu_{\hat{k},d}|}{\sigma_{\hat{k},d}} \quad (3)$$

which is small if the embedding is close to – and larger the more it deviates from – the mean. Unlike Euclidean or

²This restriction reduces the influence of mislabeled samples.

cosine distance to reduce dimensionality in a geometrically fixed way, this z-score-based deviation measure adapts to the inherent “shape” of the embeddings differently for each class. Some classes may have large variations in some dimension φ_d , whereas others have minor variations.

The most obvious way to use the z-score s to ensure monotonically decreasing score and generate an open-set algorithm is by dividing the predicted class’ logit $z_{\hat{k}}$:

$$\gamma = \frac{z_{\hat{k}}}{s}. \quad (4)$$

If the sample is close to the class mean of the predicted class \hat{k} , γ increases in scale, whereas a large z-score s lead to a reduction of γ . Thus, thresholding on γ to reject items as unknown or out-of-distribution is consistent with formal open-set theory.

Note that we are not normalizing the predicted γ score, but basically threshold this score directly, comparable to MaxLogit (Hendrycks et al. 2022). As compared to OpenMax, an advantage of GHOST is that the Gaussian models \mathcal{G} are less sensitive to outliers or mislabeled data in the training set, as the contribution of any input is only to the mean and standard deviation, which reduces noise. In contrast, even a single outlier can dominate the computation of Weibulls (Scheirer et al. 2012; Scheirer, Jain, and Boulton 2014; Bendale and Boulton 2016; Rudd et al. 2017). Furthermore, GHOST removes the necessity of selecting a tail size for Weibull fitting, and there is no need to choose a distance metric.

Class-Based Evaluation

When evaluating the performance of Open-Set Recognition, we make use of the Open-set Classification Rate (OSCR) (Dhamija, Günther, and Boulton 2018) as our primary metric since it was specifically designed to evaluate OSR performance at a given operational threshold. While OSCR was designed for Open-Set Recognition and the effect of unknown samples, it is related to Accuracy-Rejection Curves which

examine performance of systems with respect to uncertainty of new samples from known classes (Nadeem, Zucker, and Hanczar 2009). We split our test dataset in known samples $\mathcal{K} = \{(x_n, t_n) \mid 1 \leq n \leq N_K\}$ and unknown samples $\mathcal{U} = \{x_n \mid 1 \leq n \leq N_U\}$ that do not have class labels. An OSCR plot shows the Correct Classification Rate (CCR) versus the False Positive Rate (FPR):

$$\begin{aligned} \text{CCR}(\theta) &= \frac{|\{(x_n, t_n) \in \mathcal{K} \wedge \hat{k} = t_n \wedge \gamma \geq \theta\}|}{|\mathcal{K}|}, \\ \text{FPR}(\theta) &= \frac{|\{x_n \in \mathcal{U} \wedge \gamma \geq \theta\}|}{|\mathcal{U}|}. \end{aligned} \quad (5)$$

For other algorithms, we replace γ by $z_{\hat{k}}$ (MaxLogit), $y_{\hat{k}}$ (MSP), or other prediction scores for maximum class \hat{k} . To plot the CCR at a specific $\text{FPR}=\tau$, one can invert the FPR to compute $\theta_\tau = \text{FPR}^{-1}(\tau)$, which allows us to line up different algorithms/classes. We also utilize the area under the OSCR curve (AUOSCR) to compare algorithms across all thresholds, but we wish to emphasize that this suffers from many of the same problems as AUROC because it combines all possible thresholds, which is not how systems operate.

A fact that is overlooked by all OSR evaluations is the difference in performance for different classes. Only few researchers evaluated the variance of closed-set accuracy across classes. Since this is related to algorithmic fairness, we go a step further and compute the variances and coefficients of variation of CCR values across classes. First, we split our test dataset into samples from certain classes $\mathcal{K}_k = \{x_n \mid (x_n, k) \in \mathcal{K}\}$ and compute per-class CCR at $\text{FPR}=\tau$:

$$\text{CCR}_k(\theta_\tau) = \frac{|\{x_n \in \mathcal{K}_k \wedge \hat{k} = k \wedge \gamma \geq \theta_\tau\}|}{|\mathcal{K}_k|} \quad (6)$$

Note that we do not compute per-class thresholds/FPR here; the same set of thresholds θ_τ is used in all classes, but is different for each algorithm. We follow the idea of Atkinson et al. (1970); Formby, Smith, and Zheng (1999); Xinying Chen and Hooker (2023) and compute the mean, variance and coefficient of variation of per-class CCR at $\text{FPR}=\tau$:

$$\begin{aligned} \mu_{\text{CCR}}(\theta_\tau) &= \frac{1}{K} \sum_k \text{CCR}_k(\theta_\tau) \\ \sigma_{\text{CCR}}^2(\theta_\tau) &= \frac{1}{K-1} \sum_k (\text{CCR}_k(\theta_\tau) - \mu_{\text{CCR}}(\theta_\tau))^2 \\ \mathcal{V}_{\text{CCR}}(\theta_\tau) &= \frac{\sigma_{\text{CCR}}(\theta_\tau)}{\mu_{\text{CCR}}(\theta_\tau)} \end{aligned} \quad (7)$$

where \mathcal{V}_{CCR} provides a commonly used measure of inequality (unfairness) that facilitates comparisons by normalizing for changes in mean values. We evaluate μ_{CCR} , σ_{CCR} and \mathcal{V}_{CCR} at various FPR values. Since the CCR at $\text{FPR}=1$ represents closed-set accuracy, $\mathcal{V}_{\text{CCR}}(\theta_1)$ corresponds to the unfairness in closed-set accuracy, while the associated $\mu_{\text{CCR}}(\theta_1)$ represents closed-set accuracy (because the same number of samples exist per known class). To highlight some of the performance differentials, we also sort the classes by their closed-set accuracy and show the average CCR over the top-10 and bottom-10 percent of classes.

In prior OSR works, the Area Under the Receiver Operating Characteristic (AUROC) has been identified as an important metric for OSR evaluations. Binary unknown rejection is rather conceptually related to OOD, but we include AUROC as a secondary metric and additionally present ROC curves in the supplemental material. For both metrics, the curves themselves must accompany reported Area Under-statistics because area alone cannot distinguish if competing curves cross and characterize performance at specific ranges of sensitivities. We present more OSCR curves in the supplement. Additionally, in a problem where 90% of data (or risk) comes from potentially unknown inputs, having very low FPR is important and may not be easily discernable in linear plots. For highlighting differences for high-security applications, we follow Dhamija, Günther, and Boulton (2018) and plot OSRC and ROC curves with logarithmic x-axes (as in Fig. 5). For additional quantitative insight into high-security performance, we report FPR95 in our overall results (Tab. 1).

We also introduce a new OSR measure that avoids integrating over thresholds. Our goal is to determine the lowest FPR that maintains a set classification accuracy level. We call this F@C95, where we compute the FPR at the point where CCR is 95% of the closed-set accuracy. This measure, analogous to FPR95 in binary out-of-distribution detection, uses CCR and can be applied overall or per class.

Experiments

Our main evaluation relies on large-scale datasets that cover both known and unknown samples. Specifically, we draw from recent large-scale settings (Vaze et al. 2022; Hendrycks and Gimpel 2017; Bitterwolf, Mueller, and Hein 2023) that differ from other evaluations which use only small-scale data with few classes and low-resolution images. We include additional results in the supplemental material.

In particular, we use ImageNet-1K (Russakovsky et al. 2015) pre-trained networks and the validation set as our test set for knowns. For unknowns, we consider multiple datasets from the literature. We utilize a recent purpose-built OOD dataset called No ImageNet Class Objects (NINCO) (Bitterwolf, Mueller, and Hein 2023), which consists of images specifically excluding any semantically overlapping or background ImageNet objects. Additionally, we use the ImageNet-21K-P Open-Set splits (*Easy & Hard*) proposed by Vaze et al. (2022) in their semantic shift benchmark. We also include OpenImage-O (Wang et al. 2022), a dataset constructed from a public image database. Further details on each dataset and comparisons on additional datasets such as Places (Zhou et al. 2017), SUN (Xiao et al. 2010), and Textures (Cimpoi et al. 2014) are provided in the supplemental.

Experimental Setup

Methods. We compare GHOST with Maximum Logit (MaxLogit) (Hendrycks et al. 2022; Vaze et al. 2022), which is currently the state-of-the-art in large-scale OSR according to Vaze et al. (2022), and Maximum Softmax Probability (MSP) (Hendrycks and Gimpel 2017; Vaze et al. 2022). For completeness, we also compare with the current state-of-the-art in large-scale OOD according to the OpenOOD (Yang

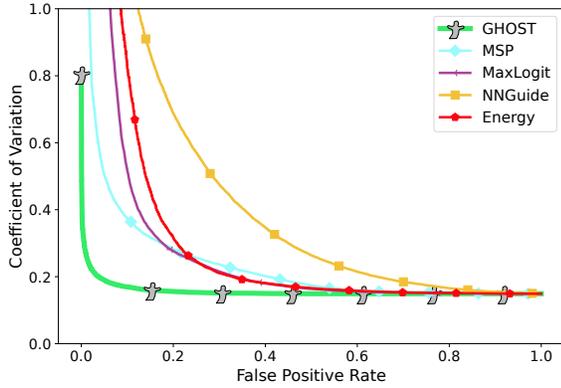


Figure 4: UNFAIRNESS (COEFFICIENT OF VARIATION). This figure shows the unfairness of OSR algorithms across False Positive Rates using MAE-H network with OpenImages as unknowns. All algorithms include the inherent unfairness from the base classifier on the right, but GHOST maintains its level much better as FPR rates are decreased to the left.

et al. 2022; Zhang et al. 2023) leaderboard,³ NNGuide (Park, Jung, and Teoh 2023). Note that NNGuide has been adapted to more recent architectures than those used in their paper, as we have observed that this adaptation significantly impacts performance. Also, GHOST results on OpenOOD’s ImageNet-1K benchmark are found in the supplemental, as well as a comparison with SCALE (Xu et al. 2024), REACT (Sun, Guo, and Li 2021) and KNN (Sun et al. 2022).

Architectures. We utilize two architectures: Masked Auto Encoder-trained Vision Transformer MAE-H (He et al. 2022) and ConvNeXtV2-H (Woo et al. 2023). MAE-H is a ViT-H network trained with a masked autoencoder; it is competitive with the state-of-the-art, PeCo (Dong et al. 2023), which does not have publicly available code or checkpoints. ConvNeXtV2-H is a recent, high-performing convolutional neural network (CNN). It is included to show that GHOST performance gains are not limited to transformer-based networks. Both networks were trained exclusively with ImageNet-1K by their respective authors. We report results on additional networks in the supplemental, offering evidence for generalizability to other architectures.

Results and Discussion

Global Performance. We present some of our quantitative results in Tab. 1, while more datasets are found in the supplemental material. On open-set AUOSCR, Tab. 1 shows that GHOST outperforms other methods on all datasets with an absolute gain of at least 2%. While OOD is not the primary focus, GHOST also outperforms other methods on the AUROC measure (Tab. 1) with a lead of 4%. It is important to note that on the NINCO (Bitterwolf, Mueller, and Hein 2023) dataset, which was specifically designed to avoid overlap with ImageNet-1K, GHOST shows clear and convincing performance gains in terms of AUROC and AUOSCR, and

³<https://zjstevens.github.io/OpenOOD>

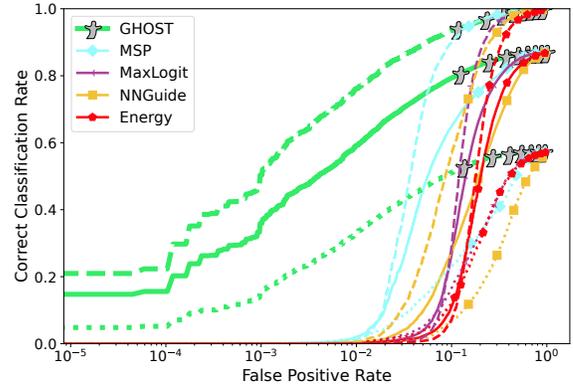


Figure 5: OSKR IN LOGSCALE. In applications with the high cost of false-positives or those with many potential unknowns, it is more important to focus on low FPR performance, in which case log FPR as shown here are more useful. The global performance is presented as a solid line, while top-10 % is dashed, and bottom-10 % is dotted. In all cases, GHOST is significantly better at low FPR levels, and below FPR of 0.1 GHOST’s bottom-10 % performance is better than most algorithms’ top-10 %.

some of the reduced performance for others may be a sign of overlap. We present results on our proposed F@C95 in Tab. 2, where each method reports the effective FPR it achieves while maintaining 95% of the closed-set accuracy. On each dataset, GHOST achieves far lower F@C95 rates than other methods.

Naturally, statistical testing should be used to validate the hypothesis of superior performance. We present statistical analysis in the supplemental material and summarize it here. For AUROC, GHOST is very significantly outperforms all methods on all datasets with $p < 10^{-6}$. For AUOSCR, GHOST is significantly better overall and on most datasets with $p < 10^{-3}$.

Fairness and Class-Based Evaluation. Fairness in OSR has two components – the inherent differential ability of the base network and the ability of the OSR to maintain the accuracy of classes in a fair/balanced manner. Previous work has ignored how individual classes are impacted by OSR thresholding. In Fig. 4, we use the \mathcal{V}_{CCR} coefficient of (7). As this is a common measure of unfairness, lower values are more equitable. At the right-hand side for FPR=1, we see the network’s baseline unfairness, and while GHOST maintains that for the majority of lower FPR values, other algorithms quickly degrade. Notably, the traditional “recognition” algorithms MSP/MaxLogit maintain their fairness better than the OOD-oriented algorithms, though none are even close to GHOST at lower FPRs. Given that we cannot compute the area under this curve as many of the curves are not bounded, we show quantitative values at 10% FPR in Tab. 3.

To provide a more detailed analysis on the class differentials, we plot the OSKR for different classes. For each algorithm, we select the top-10 % of best-performing and bottom-10 % of worst-performing classes based on the closed-set class-based accuracy, which is identical to $CCR_k(\theta_1)$ as de-

Unknowns	↑ AUOSCR ↑ AUROC ↓ FPR95								
	MAE-H				ConvNeXtV2-H				
	GHOST (ours)	MSP	MaxLogit	NNGuide	GHOST (ours)	MSP	MaxLogit	NNGuide	
21K-P <i>Easy</i>	↑.75 ↑.84 ↓.58	↑.72 ↑.80 ↓.65	↑.67 ↑.75 ↓.63	↑.62 ↑.69 ↓.80	↑.74 ↑.83 ↓.60	↑.72 ↑.79 ↓.65	↑.68 ↑.75 ↓.64	↑.70 ↑.79 ↓.70	
21K-P <i>Hard</i>	↑.73 ↑.81 ↓.62	↑.69 ↑.75 ↓.75	↑.65 ↑.71 ↓.74	↑.47 ↑.52 ↓.89	↑.72 ↑.80 ↓.65	↑.68 ↑.74 ↓.76	↑.65 ↑.72 ↓.74	↑.60 ↑.67 ↓.83	
NINCO	↑.81 ↑.91 ↓.47	↑.78 ↑.83 ↓.65	↑.73 ↑.79 ↓.62	↑.49 ↑.55 ↓.88	↑.79 ↑.89 ↓.50	↑.75 ↑.83 ↓.64	↑.73 ↑.82 ↓.60	↑.74 ↑.74 ↓.78	
OpenImage-O	↑.84 ↑.95 ↓.26	↑.76 ↑.87 ↓.52	↑.71 ↑.82 ↓.49	↑.68 ↑.77 ↓.64	↑.83 ↑.94 ↓.32	↑.79 ↑.88 ↓.49	↑.77 ↑.87 ↓.44	↑.66 ↑.83 ↓.64	

Table 1: OVERALL QUANTITATIVE RESULTS. On two state-of-the-art pre-trained architectures, GHOST achieves new state-of-the-art performance across all three metrics. In the supplemental, we demonstrate that the improvements provided by GHOST are statistically significant and consistent across additional unknowns and architectures. Methods such as Energy, SCALE, and others that are less effective are found in the supplemental.

Unknowns	GHOST	MSP	MaxLogit	NNGuide	Energy
21K-P <i>Easy</i>	0.48	0.53	0.54	0.76	0.65
21K-P <i>Hard</i>	0.53	0.64	0.64	0.86	0.74
NINCO	0.35	0.51	0.51	0.86	0.62
OpenImage-O	0.17	0.39	0.39	0.60	0.50

Table 2: F@C95. The corresponding minimum FPR at 95% of closed set accuracy (\downarrow), computed on a pre-trained MAE-H. Energy is shown here as space permitted, with additional comparisons in the supplemental

Unknowns	GHOST	MSP	MaxLogit	NNGuide	Energy
21K-P <i>Easy</i>	0.32	0.55	0.68	1.35	0.83
21K-P <i>Hard</i>	0.36	0.60	0.61	2.28	0.69
NINCO	0.21	0.45	0.50	2.18	0.68
OpenImage-O	0.17	0.38	0.52	1.16	0.82

Table 3: COEFFICIENT OF VARIANCE. The unfairness measure \mathcal{V}_{CCR} coefficients (\downarrow) of all methods, computed on a pre-trained MAE-H at 10% FPR and evaluated on various unknown datasets. Energy is shown here as space permitted, with additional comparisons in the supplemental.

fined in (6). We combine these classes and plot OSCR curves. In Fig. 1 and Fig. 5, these best- and worst-performing classes are shown together with the global OSCR curve that includes all classes. While Fig. 1 presents a linear FPR-axis, Fig. 5 shows a logarithmic FPR axis that allows investigation of very low FPR ranges. Especially in Fig. 1 it is obvious that with decreasing FPR, GHOST provides the same drop of CCR for all three lines, whereas other algorithms have different behavior. Especially MSP has superior CCR for well-classified classes (top-10%), while dropping much quicker for difficult bottom-10% classes. Furthermore, Fig. 5 shows that this behavior of GHOST extends to very low FPRs, levels that are not even reached by other methods.

From our results in Tab. 1 and Tab. 3, GHOST dominates performance across the board, in both well accepted OSR and OOD metrics (AUOSCR, FPR95, AUROC) as well as in class-wise fairness. We present additional OSCR and ROC curves in the supplemental for interested readers to verify the consistency of our results on area-based metrics.

Testing the Gaussian Hypothesis

To empirically test the Gaussian hypothesis, we use the Shapiro-Wilk test for normality with Holm’s step-down procedure for family-wise error control (Trawiński et al. 2012). For MAE-H and ConvNeXtV2-H, pretrained networks, only

2.79% and 3.13% of per-class distributions rejected the null hypothesis (normality), consistent with expected 95% confidence. Tests on Swin-T and DenseNet-121 rejected normality for 0.56% and 8.27% of the distributions, indicating the Gaussian assumption does not hold for every class in every network, but it still generally holds. We include the results for these additional networks in the supplemental material.

Additionally, GHOST could be adapted to use full covariance matrices and mahalanobis distance, but we leave this adaptation for future work.

Conclusion

In this paper, we propose GHOST, our Gaussian Hypothesis Open-Set Technique, which follows the formal definition of provable open-set theory for deep networks, based on the theorems by Bendale and Boulton (2016). We hypothesize that using per-class, per-dimension Gaussian models of feature vectors to normalize raw network logits can effectively differentiate unknown samples and improve OSR performance. Although this remains a hypothesis, it may be valuable for future work to explore mean-field theory as a means to formally prove it. By utilizing Gaussian models, we move away from traditional assumptions that rely on distance metrics in high-dimensional spaces. Instead, we normalize logits through a sum of z-scores. These Gaussian models are more robust to outliers, which can significantly affect extreme value-based statistics. We demonstrate this on two distinct architectures, providing strong support for our assumption.

Our experiments provide compelling evidence, setting a new state-of-the-art performance. Using both networks, we achieve superior results in AUOSCR and AUROC with ImageNet-1K as knowns and four datasets (and more in the supplemental) as unknowns. In nearly all cases, GHOST outperforms all methods, with statistically very significant performance gains (shown in the supplemental). Furthermore, GHOST is computationally efficient and easy to store, requiring only the mean and standard deviation (i.e., two floats per feature per class). A pre-trained network requires just one pass over the validation or training data to compute the GHOST model, and its test-time complexity is $\mathcal{O}(1)$.

We are the first to investigate fairness in OSR by examining class-wise performance differences, and we hope to encourage research that incorporates more fairness-related metrics for OSR. We have shown that GHOST maintains the closed-set unfairness of the original classifier across most FPRs, whereas other algorithms struggle significantly, increasing unfairness even at moderate FPRs.

References

- Atkinson, A. B.; et al. 1970. On the measurement of inequality. *Journal of economic theory*, 2(3): 244–263.
- Bendale, A.; and Boulton, T. E. 2016. Towards open set deep networks. In *Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE.
- Bisgin, H.; Palechor, A.; Suter, M.; and Günther, M. 2024. Large-Scale Evaluation of Open-Set Image Classification Techniques. *arxiv*.
- Bitterwolf, J.; Mueller, M.; and Hein, M. 2023. In or Out? Fixing ImageNet Out-of-Distribution Detection Evaluation. In *ICLR Workshop on Trustworthy and Reliable Large-Scale Machine Learning Models*.
- Cimpoi, M.; Maji, S.; Kokkinos, I.; Mohamed, S.; and Vedaldi, A. 2014. Describing Textures in the Wild. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Cruz, S.; Rabinowitz, R.; Günther, M.; and Boulton, T. E. 2024. Operational Open-Set Recognition and PostMax Refinement. In *European Conference on Computer Vision (ECCV)*.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. ImageNet: A large-scale hierarchical image database. In *Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE.
- Dhamija, A. R.; Günther, M.; and Boulton, T. 2018. Reducing network agnostophobia. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Dong, X.; Bao, J.; Zhang, T.; Chen, D.; Zhang, W.; Yuan, L.; Chen, D.; Wen, F.; Yu, N.; and Guo, B. 2023. Peco: Perceptual codebook for bert pre-training of vision transformers. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 552–560.
- Formby, J. P.; Smith, W. J.; and Zheng, B. 1999. The coefficient of variation, stochastic dominance and inequality: a new interpretation. *Economics Letters*, 62(3): 319–323.
- Ge, Z.; Demyanov, S.; and Garnavi, R. 2017. Generative OpenMax for Multi-Class Open Set Classification. In *British Machine Vision Conference (BMVC)*.
- Geng, C.; Huang, S.-j.; and Chen, S. 2020. Recent advances in open set recognition: A survey. *Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 43(10).
- He, K.; Chen, X.; Xie, S.; Li, Y.; Dollár, P.; and Girshick, R. 2022. Masked autoencoders are scalable vision learners. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Hendrycks, D.; Basart, S.; Mazeika, M.; Zou, A.; Kwon, J.; Mostajabi, M.; Steinhardt, J.; and Song, D. 2022. Scaling Out-of-Distribution Detection for Real-World Settings. In *International Conference on Machine Learning (ICML)*.
- Hendrycks, D.; and Gimpel, K. 2017. A Baseline for Detecting Misclassified and Out-of-Distribution Examples in Neural Networks. *International Conference on Learning Representations (ICLR)*.
- Islam, R.; Pan, S.; and Foulds, J. R. 2021. Can we obtain fairness for free? In *AAAI/ACM Conference on AI, Ethics, and Society*.
- Li, C.; Zhang, E.; Geng, C.; and Chen, S. 2024a. All Beings Are Equal in Open Set Recognition. In *AAAI Conference on Artificial Intelligence*, volume 38.
- Li, H.; Song, J.; Gao, L.; Zhu, X.; and Shen, H. 2024b. Prototype-based aleatoric uncertainty quantification for cross-modal retrieval. *Advances in Neural Information Processing Systems*, 36.
- Li, X.; Wu, P.; and Su, J. 2023. Accurate fairness: Improving individual fairness without trading accuracy. In *AAAI Conference on Artificial Intelligence*, volume 37.
- Liu, W.; Wang, X.; Owens, J.; and Li, Y. 2020. Energy-based out-of-distribution detection. *Advances in Neural Information Processing Systems (NeurIPS)*.
- Lu, Y.; Ma, C.; Lu, Y.; Lu, J.; and Ying, L. 2020. A mean field analysis of deep resnet and beyond: Towards provably optimization via overparameterization from depth. In *International Conference on Machine Learning (ICML)*.
- Miller, D.; Sunderhauf, N.; Milford, M.; and Dayoub, F. 2021. Class anchor clustering: A loss for distance-based open set recognition. In *Winter Conference on Applications of Computer Vision (WACV)*.
- Nadeem, M. S. A.; Zucker, J.-D.; and Hanczar, B. 2009. Accuracy-rejection curves (ARCs) for comparing classification methods with a reject option. In *Machine Learning in Systems Biology*, 65–81. PMLR.
- Neal, L.; Olson, M.; Fern, X.; Wong, W.-K.; and Li, F. 2018. Open set learning with counterfactual images. In *European Conference on Computer Vision (ECCV)*.
- Park, J.; Jung, Y. G.; and Teoh, A. B. J. 2023. Nearest Neighbor Guidance for Out-of-Distribution Detection. In *International Conference on Computer Vision (ICCV)*.
- Perera, P.; Morariu, V. I.; Jain, R.; Manjunatha, V.; Wigginton, C.; Ordonez, V.; and Patel, V. M. 2020. Generative-Discriminative Feature Representations for Open-Set Recognition. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Roady, R.; Hayes, T. L.; Kemker, R.; Gonzales, A.; and Kanan, C. 2020. Are open set classification methods effective on large-scale datasets? *Plos one*, 15(9): e0238302.
- Rudd, E. M.; Jain, L. P.; Scheirer, W. J.; and Boulton, T. E. 2017. The extreme value machine. *Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*.
- Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.; Berg, A. C.; and Fei-Fei, L. 2015. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3): 211–252.
- Scheirer, W. J.; de Rezende Rocha, A.; Sapkota, A.; and Boulton, T. E. 2012. Toward open set recognition. *Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 35(7).
- Scheirer, W. J.; Jain, L. P.; and Boulton, T. E. 2014. Probability models for open set recognition. *Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 36(11).

- Sensoy, M.; Kaplan, L.; and Kandemir, M. 2018. Evidential deep learning to quantify classification uncertainty. *Advances in neural information processing systems*, 31.
- Sirignano, J.; and Spiliopoulos, K. 2020. Mean field analysis of neural networks: A central limit theorem. *Stochastic Processes and their Applications*, 130(3): 1820–1852.
- Sun, Y.; Guo, C.; and Li, Y. 2021. React: Out-of-distribution detection with rectified activations. *Advances in Neural Information Processing Systems*, 34: 144–157.
- Sun, Y.; Ming, Y.; Zhu, X.; and Li, Y. 2022. Out-of-distribution detection with deep nearest neighbors. In *International Conference on Machine Learning*, 20827–20840. PMLR.
- Trawiński, B.; Smetek, M.; Telec, Z.; and Lasota, T. 2012. Nonparametric statistical analysis for multiple comparison of machine learning regression algorithms. *International Journal of Applied Mathematics and Computer Science*, 22(4): 867–881.
- Vaze, S.; Han, K.; Vedaldi, A.; and Zissermann, A. 2022. Open-Set Recognition: A Good Closed-Set Classifier is All You Need? In *International Conference on Learning Representations (ICLR)*.
- Wan, W.; Wang, X.; Xie, M.-K.; Li, S.-Y.; Huang, S.-J.; and Chen, S. 2024. Unlocking the power of open set: A new perspective for open-set noisy label learning. In *AAAI Conference on Artificial Intelligence*, volume 38.
- Wang, H.; Li, Z.; Feng, L.; and Zhang, W. 2022. Vim: Out-of-distribution with virtual-logit matching. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Wang, Y.; Mu, J.; Zhu, P.; and Hu, Q. 2024. Exploring diverse representations for open set recognition. In *AAAI Conference on Artificial Intelligence*, volume 38.
- Woo, S.; Debnath, S.; Hu, R.; Chen, X.; Liu, Z.; Kweon, I. S.; and Xie, S. 2023. ConvNeXt V2: Co-designing and Scaling ConvNets with Masked Autoencoders. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Xiao, J.; Hays, J.; Ehinger, K. A.; Oliva, A.; and Torralba, A. 2010. Sun database: Large-scale scene recognition from abbey to zoo. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Xinying Chen, V.; and Hooker, J. N. 2023. A guide to formulating fairness in an optimization model. *Annals of Operations Research*, 326(1): 581–619.
- Xu, B.; Shen, F.; and Zhao, J. 2023. Contrastive open set recognition. In *AAAI Conference on Artificial Intelligence*, volume 37.
- Xu, K.; Chen, R.; Franchi, G.; and Yao, A. 2024. Scaling for Training Time and Post-hoc Out-of-distribution Detection Enhancement. In *International Conference on Learning Representations (ICLR)*.
- Yang, H.-M.; Zhang, X.-Y.; Yin, F.; Yang, Q.; and Liu, C.-L. 2020. Convolutional prototype network for open set recognition. *Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 44(5).
- Yang, J.; Wang, P.; Zou, D.; Zhou, Z.; Ding, K.; Peng, W.; Wang, H.; Chen, G.; Li, B.; Sun, Y.; et al. 2022. OpenOOD: Benchmarking Generalized Out-of-Distribution Detection. *Advances in Neural Information Processing Systems (NeurIPS)*.
- Zhang, J.; Yang, J.; Wang, P.; Wang, H.; Lin, Y.; Zhang, H.; Sun, Y.; Du, X.; Zhou, K.; Zhang, W.; Li, Y.; Liu, Z.; Chen, Y.; and Li, H. 2023. OpenOOD v1.5: Enhanced Benchmark for Out-of-Distribution Detection. In *NeurIPS Workshop on Distribution Shifts: New Frontiers with Foundation Models*.
- Zhang, X.; Cheng, X.; Zhang, D.; Bonnington, P.; and Ge, Z. 2022. Learning Network Architecture for Open-Set Recognition. In *AAAI Conference on Artificial Intelligence*, volume 36.
- Zhou, B.; Lapedriza, A.; Khosla, A.; Oliva, A.; and Torralba, A. 2017. Places: A 10 million image database for scene recognition. *Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 40(6).
- Zhou, D.-W.; Ye, H.-J.; and Zhan, D.-C. 2021. Learning Placeholders for Open-Set Recognition. In *CVPR 2021*, 4401–4410.