

# MOS: Towards Scaling Out-of-distribution Detection for Large Semantic Space

Rui Huang

Department of Computer Sciences  
University of Wisconsin-Madison

huangrui@cs.wisc.edu

Yixuan Li

Department of Computer Sciences  
University of Wisconsin-Madison

sharonli@cs.wisc.edu

## Abstract

*Detecting out-of-distribution (OOD) inputs is a central challenge for safely deploying machine learning models in the real world. Existing solutions are mainly driven by small datasets, with low resolution and very few class labels (e.g., CIFAR). As a result, OOD detection for large-scale image classification tasks remains largely unexplored. In this paper, we bridge this critical gap by proposing a group-based OOD detection framework, along with a novel OOD scoring function termed MOS. Our key idea is to decompose the large semantic space into smaller groups with similar concepts, which allows simplifying the decision boundaries between in- vs. out-of-distribution data for effective OOD detection. Our method scales substantially better for high-dimensional class space than previous approaches. We evaluate models trained on ImageNet against four carefully curated OOD datasets, spanning diverse semantics. MOS establishes state-of-the-art performance, reducing the average FPR95 by 14.33% while achieving 6x speedup in inference compared to the previous best method.*

## 1. Introduction

Out-of-distribution (OOD) detection has become a central challenge in safely deploying machine learning models in the open world, where the test data may be distributionally different from the training data. A plethora of literature has emerged in addressing the problem of OOD detection [3, 14, 16, 24, 27, 29, 35, 4, 20, 31]. However, existing solutions are mainly driven by small, low-resolution datasets such as CIFAR [23] and MNIST [25]. Deployed systems like autonomous vehicles often operate on images that have far greater resolution and perceive environments with far more categories. As a result, a critical research gap exists in developing and evaluating OOD detection algorithms for large-scale image classification tasks.

While one may be eager to conclude that solutions for small datasets should transfer to a large-scale setting, we argue that this is far from the truth. The main challenges posed in OOD detection stem from the fact that it is impos-

sible to comprehensively define and anticipate anomalous data in advance, resulting in a large space of uncertainty. As the number of semantic classes increases, the plethora of ways that OOD data may occur increases exponentially. For example, our analysis reveals that the average false positive rate (at 95% true positive rate) of a common baseline [16] would rise from 17.34% to 76.94% as the number of classes increases from 50 to 1,000 on ImageNet-1k [8]. Very few works have studied OOD detection in the large-scale setting, with limited evaluations and effectiveness [41, 15]. This begs the following question: *how can we design an OOD detection algorithm that scales effectively for classification with large semantic space?*

Motivated by this, we take an important step to bridge this gap and propose a group-based OOD detection framework that is effective for large-scale image classification. Our key idea is to decompose the large semantic space into smaller groups with similar concepts, which allows simplifying the decision boundary and reducing the uncertainty space between in- vs. out-of-distribution data. Intuitively, for OOD detection, it is simpler to estimate whether an image belongs to one of the coarser-level semantic groups than to estimate whether an image belongs to one of the finer-grained classes. For example, consider a model tasked with classifying 200 categories of plantations and another 200 categories of marine animals. A `truck` image can be easily classified as OOD data since it does not resemble either the plantation group or the marine animal group.

Formally, our proposed method leverages group softmax and derives a novel OOD scoring function. Specifically, the group softmax computes probability distributions within each semantic group. A key component is to utilize a category `others` in each group, which measures the probabilistic score for an image to be OOD with respect to the group. Our proposed OOD scoring function, *Minimum Others Score (MOS)*, exploits the information carried by the `others` category. As illustrated in Figure 1, MOS is higher for OOD inputs as they will be mapped to `others` with high confidence in all groups, and is lower for in-distribution inputs.

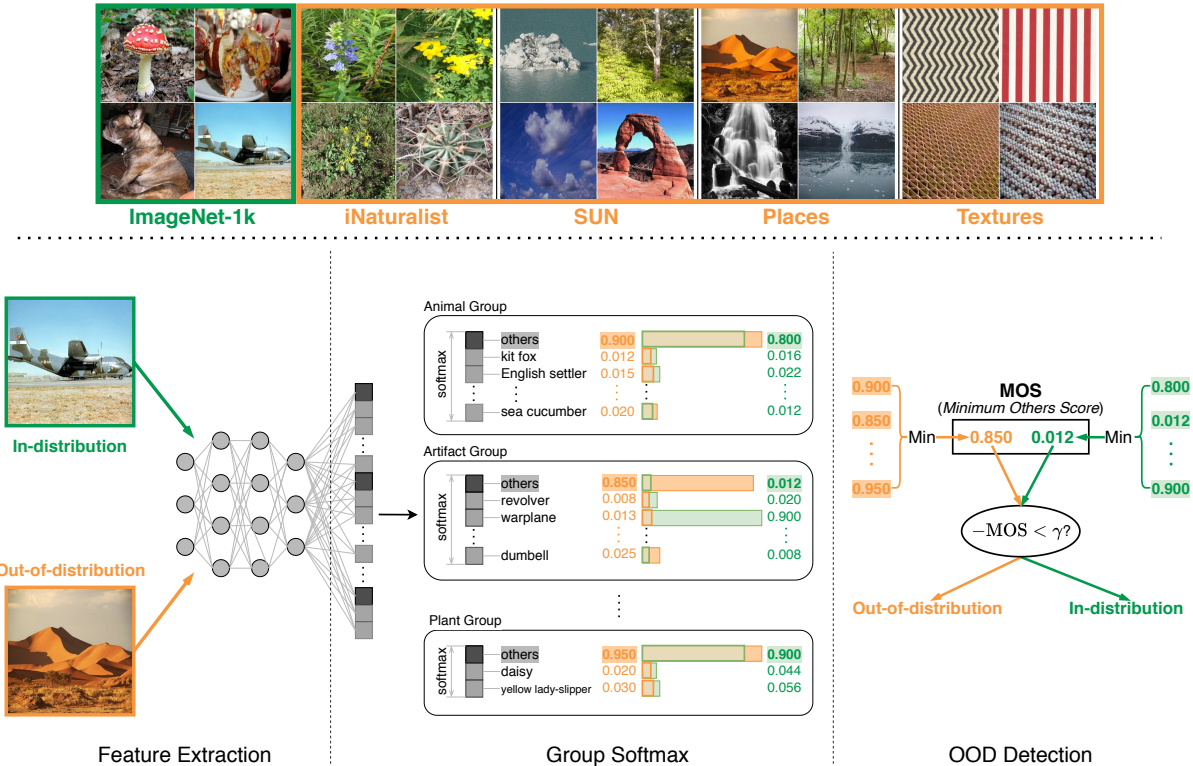


Figure 1: *Top*: Examples of in-distribution images sampled from ImageNet-1k (in green) and OOD images sampled from 4 datasets described in Section 4.1 (in orange). *Bottom*: Overview of the proposed group-based OOD detection framework. The key idea is to decompose the large semantic space into smaller groups, which allows simplifying the decision boundary between in- and out-of-distribution data. A category `others` is added to each group. An OOD image is mapped to `others` with high confidence for all groups, whereas an in-distribution image will have a lower score for `others` in the group it belongs to (e.g., artifact group). The minimum score on category `others` among all groups, MOS, allows effective differentiation of OOD data.

We extensively evaluate our approach on models trained with the ImageNet-1k dataset, leveraging the state-of-the-art pre-trained BiT-S models [22] as backbones. We explore label space of size 10-100 times larger than that of previous works [16, 29, 27, 31, 4, 17]. Compared to the best baseline [15], our method improves the average performance of OOD detection by 14.33% (FPR95) over four diverse OOD test datasets. More importantly, our method achieves improved OOD detection performance while preserving the classification accuracy on in-distribution datasets. We note that while group-based learning has been used for improving tasks such as long-tail object detection [28], our objective and motivation are very different—we are interested in reducing the uncertainty between in- and out-of-distribution data, rather than reducing the confusion among in-distribution data themselves. Below we summarize our **key results and contributions**:

- We propose a group-based OOD detection framework, along with a novel OOD scoring function MOS, that scales substantially better for large label space. Our method establishes the new state-of-the-art performance, reducing the average FPR95 by **14.33%** while achieving **6x** speedup in inference time compared to

the best baseline.

- We conduct extensive ablations which improve understandings of our method for large-scale OOD detection under (1) different grouping strategies, (2) different sizes of semantic class space, (3) different backbone architectures, and (4) varying fine-tuning capacities.
- We curate diverse OOD evaluation datasets from four real-world high-resolution image databases, which enables future research to evaluate OOD detection methods in a large-scale setting<sup>1</sup>.

## 2. Preliminary and Analysis

**Preliminaries** We consider a training dataset drawn i.i.d. from the in-distribution  $P_X$ , with label space  $Y = \{1, 2, \dots, C\}$ . For OOD detection problem, it is common to train a classifier  $f(x)$  on the in-distribution  $P_X$ , and evaluate on samples that are drawn from a different distribution  $Q_X$ . An OOD detector  $G(x)$  is a binary classifier:

$$G(x) = \begin{cases} \text{in}, & \text{if } S(x) \geq \gamma \\ \text{out}, & \text{if } S(x) < \gamma, \end{cases}$$

<sup>1</sup>Code and data for reproducing our results are available at: [https://github.com/deeplearning-wisc/large\\_scale\\_ood](https://github.com/deeplearning-wisc/large_scale_ood)

where  $S(\mathbf{x})$  is the scoring function, and  $\gamma$  is the threshold chosen so that a high fraction (e.g., 95%) of in-distribution data is correctly classified.

**Effect of Number of Classes on OOD Detection** We first revisit the common baseline approach [16], which uses the maximum softmax probability (MSP),  $S(\mathbf{x}) = \max_i \frac{e^{f_i(\mathbf{x})}}{\sum_{j=1}^C e^{f_j(\mathbf{x})}}$ , for OOD detection. We investigate the effect of label space size on the OOD detection performance. In particular, we use a ResNetv2-101 architecture [13] trained on different subsets<sup>2</sup> of ImageNet with varying numbers of classes  $C$ . As shown in Figure 2, the performance (FPR95) degrades rapidly from 17.34% to 76.94% as the number of in-distribution classes increases from 50 to 1,000. This trend signifies that current OOD detection methods are indeed challenged by the increasingly large label space, which motivates our work.




Figure 2: OOD detection performance of a common baseline MSP [16] decreases rapidly as the number of ImageNet-1k classes increases (left: AUROC; right: FPR95).

### 3. Method

Our novel group-based OOD detection framework is illustrated in Figure 1. In what follows, we first provide an overview and then describe the group softmax training technique in Section 3.1. We introduce our proposed OOD detection algorithm MOS in Section 3.2, followed with grouping strategies in Section 3.3.

**Method Overview: A Conceptual Example** As aforementioned, OOD detection performance can suffer notably from the increasing number of in-distribution classes. To mitigate this issue, our key idea is to decompose the large semantic space into smaller groups with similar concepts, which allows simplifying the decision boundary and reducing the uncertainty space between in- vs. out-of-distribution data. We illustrate our idea with a toy example in Figure 3, where the in-distribution data consists of class-conditional Gaussians. Without grouping (left), the decision boundary between in- vs. OOD data is determined by *all* classes and becomes increasingly complex as the number of classes

<sup>2</sup>To create the training subset, we first randomly select  $C$  ( $C \in \{50, 200, 300, 400, 500, 600, 700, 800, 900, 1000\}$ ) labels from the 1,000 ImageNet classes. For each of the chosen label, we then sample 700 images for training.




Figure 3: A toy example in 2D space of group-based OOD detection framework. *Left*: without grouping, the decision boundary between in- vs. out-of-distribution data becomes increasingly complex with more classes. *Right*: our group-based method simplifies the decision boundary and reduces the uncertainty space for OOD data.

grows. In contrast, with grouping (right), the decision boundary for OOD detection can be significantly simplified, as shown by the dotted curves.

In other words, by way of grouping, the OOD detector only needs to make a small number of relatively simple estimations about *whether an image belongs to this group*, as opposed to making a large number of hard decisions about *whether an image belongs to this class*. An image will be classified as OOD if it belongs to none of the groups. We proceed with describing the training mechanism that achieves our novel conceptual idea.

#### 3.1. Group-based Learning

We divide the total number of  $C$  categories into  $K$  groups,  $\mathcal{G}_1, \mathcal{G}_2, \dots, \mathcal{G}_K$ . We calculate the standard group-wise softmax for each group  $\mathcal{G}_k$ :

$$p_c^k(\mathbf{x}) = \frac{e^{f_c^k(\mathbf{x})}}{\sum_{c' \in \mathcal{G}_k} e^{f_{c'}^k(\mathbf{x})}}, \quad c \in \mathcal{G}_k, \quad (1)$$

where  $f_c^k(\mathbf{x})$  and  $p_c^k(\mathbf{x})$  denote the output logit and the softmax probability for class  $c$  in group  $\mathcal{G}_k$ , respectively.

**Category “Others”** Standard group softmax is insufficient as it can only discriminate classes within the group, but cannot estimate the OOD uncertainty between inside vs. outside the group. To this end, a new category `others` is introduced to every group, as shown in Figure 1. The model can predict `others` if the input  $\mathbf{x}$  does not belong to this group. In other words, the `others` category allows explicitly learning the decision boundary between inside vs. outside the group, as illustrated by the dashed curves surrounding classes C1/C2/C3 in Figure 3. This is desirable for OOD detection, as an OOD input can be mapped to `others` for all groups, whereas an in-distribution input will be mapped to one of the semantic categories in some group with high confidence.

Importantly, our use of the category `others` creates “virtual” group-level outlier data without relying on any external data. Each training example  $\mathbf{x}$  not only helps estimate the decision boundary for the classification problem,

but also effectively improves the OOD uncertainty estimation for groups to which it does not belong. We show the formulation can in fact achieve the dual objective of in-distribution classification, as well as OOD detection.

**Training and Inference** During training, the ground-truth labels are re-mapped in each group. In groups where  $c$  is not included, class `others` will be defined as the ground-truth class. The training objective is a sum of cross-entropy losses in each group:

$$\mathcal{L}_{GS} = -\frac{1}{N} \sum_{n=1}^N \sum_{k=1}^K \sum_{c \in \mathcal{G}_k} y_c^k \log(p_c^k(\mathbf{x})), \quad (2)$$

where  $y_c^k$  and  $p_c^k$  represent the label and the softmax probability of category  $c$  in  $\mathcal{G}_k$ , and  $N$  is the total number of training samples.

We denote the set of all valid (non-`others`) classes in each group as  $\mathcal{G}'_k = \mathcal{G}_k \setminus \{\text{others}\}$ . During inference time, we derive the group-wise class prediction in the valid set for each group:

$$\hat{p}^k = \max_{c \in \mathcal{G}'_k} p_c^k(\mathbf{x}), \quad \hat{c}^k = \arg \max_{c \in \mathcal{G}'_k} p_c^k(\mathbf{x}).$$

Then we use the maximum group-wise softmax score and the corresponding class for final prediction:

$$k_* = \arg \max_{1 \leq k \leq K} \hat{p}^k.$$

The final prediction is category  $\hat{c}^{k_*}$  from group  $\mathcal{G}_{k_*}$ .

### 3.2. OOD Detection with MOS

For a classification model trained with the group softmax loss, we propose a novel OOD scoring function, **Minimum Others Score (MOS)**, that allows effective differentiation between in- vs. out-of-distribution data. Our key observation is that category `others` carries useful information for how likely an image is OOD with respect to each group.

As discussed in Section 3.1, an OOD input will be mapped to `others` with high confidence in all groups, whereas an in-distribution input will have a low score on category `others` in the group it belongs to. Therefore, the lowest `others` score among all groups is crucial for distinguishing between in- vs. out-of-distribution data. This leads to the following OOD scoring function, termed as *Minimum Others Score*:

$$S_{\text{MOS}}(\mathbf{x}) = -\min_{1 \leq k \leq K} p_{\text{others}}^k(\mathbf{x}). \quad (3)$$

Note that we negate the sign to align with the conventional notion that  $S_{\text{MOS}}(\mathbf{x})$  is higher for in-distribution data and lower for out-of-distribution.

To provide an interpretation and intuition behind MOS, we show in Figure 4 the average scores for the category `others` in each group for both in-distribution and OOD images. For in-distribution, we select all validation images

from the `animal` group in the ImageNet-1k dataset. The minimum `others` score among all groups is significantly lower for in-distribution data than that for OOD data, allowing for effective differentiation between them.

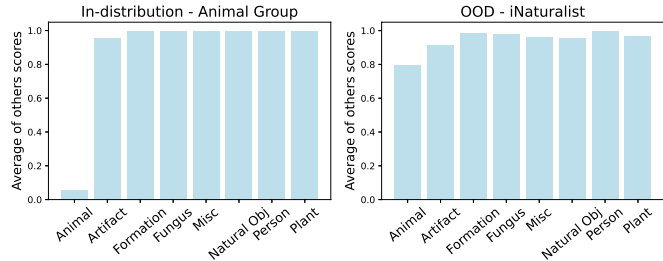


Figure 4: Average of `others` scores in each group for both in-distribution data (left) and OOD data (right).

### 3.3. Grouping Strategies

Given the dependency on the group structure, a natural question arises: *how do different grouping strategies affect the performance of OOD detection?* To answer this, we systematically consider three grouping strategies: (1) taxonomy, (2) feature clustering, and (3) random grouping.

**Taxonomy** The first grouping strategy is applicable when the taxonomy of the label space is known. For example, in the case of ImageNet, each class is associated with a synset in WordNet [34], from which we can build the taxonomy as a hierarchical tree. In particular, we adopt the 8 super-classes defined by ImageNet<sup>3</sup> as our groups and map each category into one of the 8 groups: `animal`, `artifact`, `geological formation`, `fungus`, `misc`, `natural object`, `person`, and `plant`.

**Feature Clustering** When taxonomy is not available, we can approximately estimate the structure of semantic classes through feature clustering. Specifically, we extract feature representations for each training image from a pre-trained feature extractor. Then, the feature representation of each class is the average of feature embeddings in that class. Finally, we perform a K-Means clustering [32] on categorical feature representations, one for each class.

**Random Grouping** Lastly, we contrast the taxonomy and the feature clustering strategies with random grouping, where each class is randomly assigned to a group. This allows us to estimate the lower bound of OOD detection performance with MOS.

By default, we use taxonomy as the grouping strategy if not specified otherwise. In Section 4.3.3, we experimentally compare the OOD detection performance using all three grouping strategies.

<sup>3</sup><http://image-net.org/explore>

Method	Test Time (min)	iNaturalist		SUN		Places		Textures		Average	
		AUROC ↑	FPR95 ↓	AUROC ↑	FPR95 ↓	AUROC ↑	FPR95 ↓	AUROC ↑	FPR95 ↓	AUROC ↑	FPR95 ↓
MSP [16]	3.1	87.59	63.69	78.34	79.98	76.76	81.44	74.45	82.73	79.29	76.96
ODIN [29]	23.6	89.36	62.69	83.92	71.67	80.67	76.27	76.30	81.31	82.56	72.99
Mahalanobis [27]	145.4	46.33	96.34	65.20	88.43	64.46	89.75	72.10	52.23	62.02	81.69
Energy [31]	3.1	88.48	64.91	85.32	65.33	81.37	73.02	75.79	80.87	82.74	71.03
KL Matching [15]	20.6	93.00	27.36	78.72	67.52	76.49	72.61	<b>87.07</b>	<b>49.70</b>	83.82	54.30
<b>MOS (ours)</b>	3.2	<b>98.15</b>	<b>9.28</b>	<b>92.01</b>	<b>40.63</b>	<b>89.06</b>	<b>49.54</b>	81.23	60.43	<b>90.11</b>	<b>39.97</b>

Table 1: OOD detection performance comparison between MOS and baselines. All methods are fine-tuned from the same pre-trained BiT-S-R101x1 backbone with ImageNet-1k as in-distribution dataset. The description of 4 OOD test datasets is provided in Section 4.1. ↑ indicates larger values are better, while ↓ indicates smaller values are better. All values are percentages. **Bold** numbers are superior results. Test time for all methods are evaluated with the same in- and out-of-distribution datasets (60k images in total).

## 4. Experiments

We first describe the evaluation datasets (Section 4.1) and experimental setups (Section 4.2). In Section 4.3, we show that MOS achieves state-of-the-art OOD detection performance, followed by extensive ablations that improve the understandings of MOS for large-scale OOD detection.

### 4.1. Datasets

#### 4.1.1 In-distribution Dataset

We use ImageNet-1k [8] as the in-distribution dataset, which covers a wide range of real-world objects. ImageNet-1k has at least 10 times more labels compared to CIFAR datasets used in prior literature. In addition, the image resolution is also significantly higher than CIFAR ( $32 \times 32$ ) and MNIST ( $28 \times 28$ ).

#### 4.1.2 Out-of-distribution Datasets

To evaluate our approach, we consider a diverse collection of OOD test datasets, spanning various domains including fine-grained images, scene images, and textural images. We carefully curate the OOD evaluation benchmarks to make sure concepts in these datasets do not overlap with ImageNet-1k. Below we describe the construction of each evaluation dataset in detail. Samples of each OOD dataset are provided in Figure 1. We provide the list of concepts chosen for each OOD dataset in Appendix A.

**iNaturalist** iNaturalist [46] is a fine-grained dataset containing 859,000 images across more than 5,000 species of plants and animals. All images are resized to have a max dimension of 800 pixels. We manually select 110 plant classes not present in ImageNet-1k, and randomly sample 10,000 images for these 110 classes.

**SUN** SUN [48] is a scene database of 397 categories and 130,519 images with sizes larger than  $200 \times 200$ . SUN and ImageNet-1k have overlapping categories. Therefore, we carefully select 50 nature-related concepts that are unique in SUN, such as *forest* and *iceberg*. We randomly sample 10,000 images for these 50 classes.

**Places** Places365 [50] is another scene dataset with similar concept coverage as SUN. All images in this dataset have been resized to have a minimum dimension of 512. We manually select 50 categories from this dataset that are not present in ImageNet-1k and then randomly sample 10,000 images for these 50 categories.

**Textures** Textures [6] consists of 5,640 images of textural patterns, with sizes ranging between  $300 \times 300$  and  $640 \times 640$ . We use the entire dataset for evaluation.

### 4.2. Experiment Setup

**Pre-trained Backbone** We use Google BiT-S models [22] as our feature extractor in all experiments. The models are trained on ImageNet-1k, with ResNetv2 architectures [13] at varying capacities. Pre-trained models allow extracting high-quality features with minimal time and energy consumption. In practice, one can always choose to train from scratch.

For the main results, we use the BiT-S-R101x1 model with depth 101 and width factor 1, unless specified otherwise. We provide a comparison of using feature extractors of varying model sizes in Section 4.3.4. For efficiency, we fix the backbone and only fine-tune the last fully-connected (FC) layer in the main experiments. We additionally explore the effect of fine-tuning more layers beyond the last FC layer in Section 4.3.5.

**Training Details** We follow the procedure in BiT-HyperRule [22] and fine-tune the pre-trained BiT-S model for 20k steps with a batch size of 512. We use SGD with an initial learning rate of 0.003 and a momentum of 0.9. The learning rate is decayed by a factor of 10 at 30%, 60%, and 90% of the training steps. During training, all images are resized to  $512 \times 512$  and randomly cropped to  $480 \times 480$ . At test time, all images are resized to  $480 \times 480$ . A learning rate warm-up is used for the first 500 steps. We perform all experiments on NVIDIA GeForce RTX 2080Ti GPUs.

**Evaluation Metrics** We measure the following metrics that are commonly used for OOD detection: (1) the false positive rate of OOD examples when the true positive rate

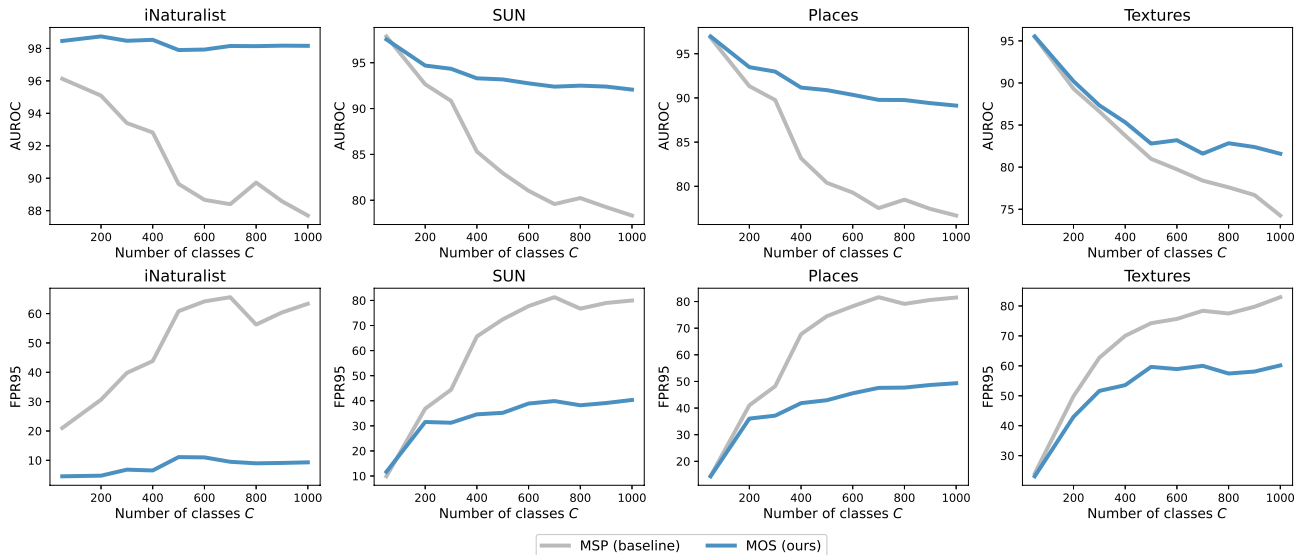


Figure 5: OOD detection performance of MOS (blue) and the MSP baseline (gray). MOS exhibits more stabilized performance as the number of in-distribution classes increases. For each OOD dataset, we show AUROC (*top*) and FPR95 (*bottom*).

of in-distribution examples is at 95% (FPR95); (2) the area under the receiver operating characteristic curve (AUROC). We additionally report the area under the precision-recall curve (AUPR) in Appendix D.

### 4.3. Results

#### 4.3.1 MOS vs. Existing Methods

The main results are shown in Table 1. We report performance for each dataset described in Section 4.1, as well as the average performance. For fair evaluation, we compare with competitive methods in the literature that derive OOD scoring functions from a model trained on in-distribution data and do not rely on auxiliary outlier data. We first compare with approaches driven by small datasets, including MSP [16], ODIN [29], Mahalanobis [27], as well as Energy [31]. All these methods rely on networks trained with flat softmax. Under the same network backbone (BiT-S-R101x1), MOS outperforms the best baseline Energy [31] by **31.06%** in FPR95. It is also worth noting that fine-tuning with group softmax maintains competitive classification accuracy (75.16%) on in-distribution data compared with its flat softmax counterpart (75.20%).

We also compare our method with KL matching [15], a competitive baseline evaluated on large-scale image classification. MOS reduces FPR95 by **14.33%** compared to KL matching. Note that for each input, KL matching needs to calculate its KL divergence to all class centers. Therefore, the running time of KL matching increases linearly with the number of in-distribution categories, which can be computationally expensive for a very large label space. As shown in Table 1, our method achieves a **6x** speedup compared to KL matching.

#### 4.3.2 MOS with Increasing Numbers of Classes

In Figure 5, we show the OOD detection performance as we increase the number of in-distribution classes  $C \in \{50, 200, 300, 400, 500, 600, 700, 800, 900, 1000\}$  on ImageNet-1k. For each  $C$ , we create training data by first randomly sampling  $C$  labels from the entire 1k classes, and then sampling 700 images for each chosen label. Importantly, we observe that MOS (in blue) exhibits more stabilized performance as  $C$  increases, compared to MSP [16] (in gray). For example, on the iNaturalist OOD dataset, FPR95 rises from 21.02% to 63.36% using MSP, whilst MOS degrades by only 4.76%. This trend signifies that MOS is an effective approach for scaling OOD detection towards a large semantic space.

We also explore an alternative setting where we fix the total number of training images, as we vary the number of classes  $C$ . In this setting, the model is trained on fewer images per class as the number of classes increases, making the problem even more challenging. We report those results in Appendix B.1. Overall, MOS remains less sensitive to the number of classes compared to the MSP baseline.

#### 4.3.3 MOS with Different Grouping Strategies

In this ablation, we contrast the performance of three different grouping strategies described in Section 3.3. For a fair comparison, we use the number of groups  $K = 8$  for all methods, since ImageNet taxonomy has 8 super-classes. For feature clustering, we first extract the feature vector from the penultimate layer of the pre-trained BiT-S model for each training image. The feature representation for each category is the average feature vector among all images in that category. We then perform a K-Means clustering on the 1,000 categorical feature vectors (one for each class) with




Figure 6: OOD detection performance comparison between MOS with different grouping strategies and the MSP baseline on 4 OOD datasets. (top: AUROC; bottom: FPR95).

$K = 8$ . For random grouping, we randomly split 1,000 classes into 8 groups with equal sizes (125 classes each).

We compare the performance of MOS under different grouping strategies in Figure 6. We observe that feature clustering works substantially better than the MSP baseline [16] while maintaining similar in-distribution classification accuracy (-0.16%) to the taxonomy-based grouping. Interestingly, random grouping achieves better performance than the MSP baseline [16] on 3 out of 4 OOD datasets. However, we do observe a drop of in-distribution classification accuracy (-0.98%) using random grouping, compared to the taxonomy-based grouping. We argue that feature clustering is a viable strategy when taxonomy is unavailable, as it outperforms MSP by 18.2% (FPR95) on average. We additionally report how different numbers of groups  $K$  affect the OOD detection performance for all three grouping strategies in Appendix B.2.

#### 4.3.4 MOS with Different Feature Extractors

We investigate how the performance of OOD detection changes as we employ different pre-trained feature extractors. In Figure 7, we compare the performance of using a family of 5 feature extractors (in increasing size): BiT-S-R50x1, BiT-S-R101x1, BiT-S-R50x3, BiT-S-R152x2, BiT-S-R101x3<sup>4</sup>. All models are ResNetv2 architectures with varying depths and width factors. It is important to note that since we fix the entire backbone and only fine-tune the last FC layer, this ablation is about the effect of the quality of feature extractors rather than model capacities.

As we use feature extractors trained on larger capacities, the classification accuracy increases, with comparable per-

<sup>4</sup>[https://github.com/google-research/big\\_transfer](https://github.com/google-research/big_transfer)

formance between using the flat vs. group softmax. Overall the OOD detection performance improves as the capacity of feature extractors increases. More importantly, MOS consistently outperforms MSP [16] in all cases. These results suggest that using pre-trained models with better feature representations will not only improve classification accuracy but also benefit OOD detection performance.

#### 4.3.5 MOS with Varying Fine-tuning Capacities

In this ablation, we explore the efficacy of fine-tuning more layers. Concretely, we go beyond the FC layer and fine-tune different numbers of residual blocks in BiT-S-R101x1. Figure 8 shows the classification accuracy and OOD detection performance under different fine-tuning capacities. Noticeably, MOS consistently outperforms MSP [16] in OOD detection under all fine-tuning capacities. As expected, we observe that fine-tuning more layers leads to better classification accuracy. However, increasing the number of fine-tuning layers would adversely affect OOD detection in some cases. We hypothesize that fine-tuning with more layers will result in more label-overfitted predictions, and undesirably produce higher confidence scores for OOD data. This suggests that only fine-tuning the top FC layers is not only computationally efficient but also in fact desirable for OOD detection performance.

## 5. Related Work

**OOD Detection with Pre-trained Models** Hendrycks and Gimpel [16] establish a common baseline for OOD detection by using maximum softmax probability (MSP). Several works attempt to improve the OOD uncertainty estimation by using ODIN score [29], deep ensembles [24], Mahalanobis distance-based confidence score [27], generalized ODIN score [20], and energy score [31]. However, previous methods driven by small datasets are either sub-optimal or inefficient in a large-scale setting. In contrast, MOS scales significantly better with large label space, both in algorithmic accuracy and computational efficiency.

**OOD Detection with Model Fine-tuning** An orthogonal line of work explores training with auxiliary outlier data for model regularization [3, 11, 33, 35, 43, 31]. Auxiliary outlier data can either be realistic images [17, 35, 37, 31, 4] or synthetic images generated by GANs [26]. Several loss functions have been designed to regularize model predictions of the auxiliary outlier data towards uniform distributions [26], a background class for OOD data [4, 35], or higher energies [31]. In this work, our model is fine-tuned only on in-distribution data, as we do not assume the availability of auxiliary outlier data. Different from previous settings, it can be prohibitive to construct an auxiliary outlier dataset in large-scale image classification, since the in-distribution data has a much wider coverage of concepts.

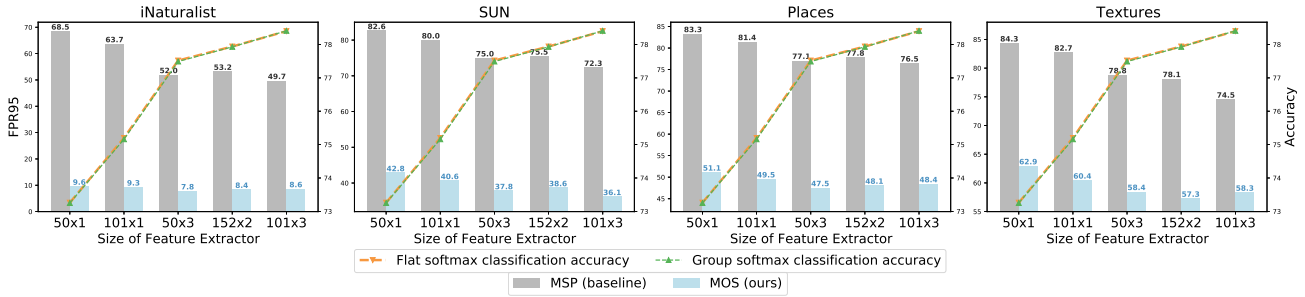


Figure 7: Effect of using different pre-trained feature extractors. The x-axis indicates feature extractors with larger capacities from left to right. Only the top FC layer is fine-tuned in all experiments. Both OOD detection (*bars*) and image classification (*dashed lines*) benefit from improved feature extractors.

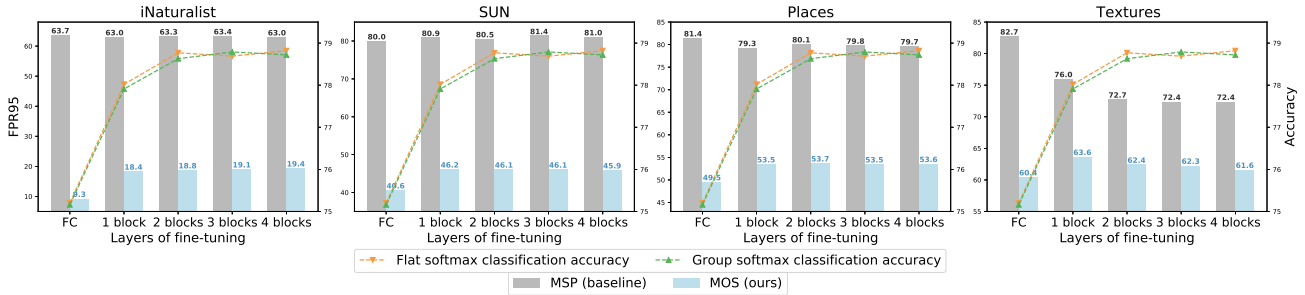


Figure 8: Effect of fine-tuning different numbers of residual blocks in BiT-S-R101x1. We show both OOD detection (*bars*) and image classification (*dashed lines*) performance.

**Generative Modeling Based OOD Detection** Generative models [21, 44, 40, 10, 45] estimate the probability density of input data and can thus be directly utilized as OOD detectors with high density indicating in-distribution and low density indicating out-of-distribution. However, as shown in [36], deep generative models can undesirably assign a high likelihood to OOD data. Several strategies have been proposed to mitigate this issue, such as improved metrics [5], likelihood ratio [39, 42], and modified training techniques [17]. In this work, we mainly focus on the discriminative-based approaches. It is important to note that generative models [19] can be prohibitively challenging to train and optimize on large-scale real-world datasets.

**OOD Detection for Large-scale Classification** Several works make pioneering efforts in large-scale OOD detection. Roody *et al.* [41] sample half of the classes from ImageNet-1k as in-distribution data, and evaluate the other half as OOD test data. They use a one-vs-rest training strategy and background class regularization, which requires access to an auxiliary dataset. KL matching was employed as the OOD scoring function in [15]. In this work, we propose a novel group-based solution that scales more effectively and efficiently for large-scale OOD detection. We also perform evaluations on more diverse real-world OOD datasets and conduct thorough ablations that improve understanding of the problem and solutions in many aspects.

**Learning with Hierarchical Labels** The hierarchical structure of the class categories has been utilized for efficient inference [9, 30], improved classification accuracy [7], and stronger object detection performance [38]. Some works aim to learn a label tree structure when taxonomy is unavailable [2, 9, 30]. As a typical hierarchy, group-based learning has been widely adopted in image classification tasks [18, 1, 49, 47, 12]. Recently, a group softmax classifier is proposed to tackle the problem of long-tail object detection, where categories are grouped according to the number of training instances [28]. We contribute to this field by showing the promise of using a group label structure for effective OOD detection.

## 6. Conclusion

In this paper, we propose a group-based OOD detection framework, along with a novel OOD scoring function, MOS, that effectively scales the OOD detection to a real-world setting with a large label space. We curate four diverse OOD evaluation datasets that allow future research to evaluate OOD detection methods in a large-scale setting. Extensive experiments show our group-based framework can significantly improve the performance of OOD detection in this large-scale setting compared to existing approaches. We hope our research can raise more attention to expand the view of OOD detection from small benchmarks to a large-scale real-world setting.



## References

- [1] Karim Ahmed, Mohammad Haris Baig, and Lorenzo Torresani. Network of experts for large-scale image categorization. In *European Conference on Computer Vision*, pages 516–532. Springer, 2016. 8
- [2] Samy Bengio, Jason Weston, and David Grangier. Label embedding trees for large multi-class tasks. In *Advances in Neural Information Processing Systems*, pages 163–171, 2010. 8
- [3] Petra Bevanđić, Ivan Krešo, Marin Oršić, and Siniša Šegvić. Discriminative out-of-distribution detection for semantic segmentation. *arXiv preprint arXiv:1808.07703*, 2018. 1, 7
- [4] Jiefeng Chen, Yixuan Li, Xi Wu, Yingyu Liang, and Somesh Jha. Robust out-of-distribution detection via informative outlier mining. *arXiv preprint arXiv:2006.15207*, 2020. 1, 2, 7
- [5] Hyunsun Choi and Eric Jang. Generative ensembles for robust anomaly detection. 2018. 8
- [6] Mircea Cimpoi, Subhransu 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. 5, 11
- [7] Jia Deng, Nan Ding, Yangqing Jia, Andrea Frome, Kevin Murphy, Samy Bengio, Yuan Li, Hartmut Neven, and Hartwig Adam. Large-scale object classification using label relation graphs. In *European conference on computer vision*, pages 48–64. Springer, 2014. 8
- [8] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 1, 5, 11
- [9] Jia Deng, Sanjeev Satheesh, Alexander C Berg, and Fei Li. Fast and balanced: Efficient label tree learning for large scale object recognition. In *Advances in Neural Information Processing Systems*, pages 567–575, 2011. 8
- [10] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real NVP. In *5th International Conference on Learning Representations, ICLR 2017*, 2017. 8
- [11] Yonatan Geifman and Ran El-Yaniv. Selectivenet: A deep neural network with an integrated reject option. In *International Conference on Machine Learning*, pages 2151–2159. PMLR, 2019. 7
- [12] Sam Gross, Marc’Aurelio Ranzato, and Arthur Szlam. Hard mixtures of experts for large scale weakly supervised vision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6865–6873, 2017. 8
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *European conference on computer vision*, pages 630–645. Springer, 2016. 3, 5
- [14] Matthias Hein, Maksym Andriushchenko, and Julian Bitterwolf. Why relu networks yield high-confidence predictions far away from the training data and how to mitigate the problem. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 41–50, 2019. 1
- [15] Dan Hendrycks, Steven Basart, Mantas Mazeika, Mohamadreza Mostajabi, Jacob Steinhardt, and Dawn Song. A benchmark for anomaly segmentation. *arXiv preprint arXiv:1911.11132*, 2019. 1, 2, 5, 6, 8
- [16] Dan Hendrycks and Kevin Gimpel. A baseline for detecting misclassified and out-of-distribution examples in neural networks. In *5th International Conference on Learning Representations, ICLR 2017*, 2017. 1, 2, 3, 5, 6, 7, 12, 14
- [17] Dan Hendrycks, Mantas Mazeika, and Thomas Dietterich. Deep anomaly detection with outlier exposure. In *International Conference on Learning Representations*, 2018. 2, 7, 8
- [18] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. 8
- [19] Tobias Hinz, Stefan Heinrich, and Stefan Wermter. Generating multiple objects at spatially distinct locations. In *International Conference on Learning Representations*, 2018. 8
- [20] Yen-Chang Hsu, Yilin Shen, Hongxia Jin, and Zsolt Kira. Generalized odin: Detecting out-of-distribution image without learning from out-of-distribution data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10951–10960, 2020. 1, 7
- [21] Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In Yoshua Bengio and Yann LeCun, editors, *2nd International Conference on Learning Representations, ICLR 2014*. 8
- [22] Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Joan Puigcerver, Jessica Yung, Sylvain Gelly, and Neil Houlsby. Big transfer (bit): General visual representation learning. In *ECCV 2020*, 2020. 2, 5
- [23] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 1
- [24] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. In *Advances in neural information processing systems*, pages 6402–6413, 2017. 1, 7
- [25] Yann LeCun, Corinna Cortes, and CJ Burges. Mnist handwritten digit database. *ATT Labs [Online]*. Available: <http://yann.lecun.com/exdb/mnist>, 2, 2010. 1
- [26] Kimin Lee, Honglak Lee, Kibok Lee, and Jinwoo Shin. Training confidence-calibrated classifiers for detecting out-of-distribution samples. In *International Conference on Learning Representations*, 2018. 7
- [27] Kimin Lee, Kibok Lee, Honglak Lee, and Jinwoo Shin. A simple unified framework for detecting out-of-distribution samples and adversarial attacks. In *Advances in Neural Information Processing Systems*, pages 7167–7177, 2018. 1, 2, 5, 6, 7
- [28] Yu Li, Tao Wang, Bingyi Kang, Sheng Tang, Chunfeng Wang, Jintao Li, and Jiashi Feng. Overcoming classifier imbalance for long-tail object detection with balanced group softmax. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10991–11000, 2020. 2, 8
- [29] Shiyu Liang, Yixuan Li, and Rayadurgam Srikant. Enhancing the reliability of out-of-distribution image detection in

- neural networks. In *6th International Conference on Learning Representations, ICLR 2018*, 2018. 1, 2, 5, 6, 7
- [30] Baoyuan Liu, Fereshteh Sadeghi, Marshall Tappen, Ohad Shamir, and Ce Liu. Probabilistic label trees for efficient large scale image classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 843–850, 2013. 8
- [31] Weitang Liu, Xiaoyun Wang, John Owens, and Yixuan Li. Energy-based out-of-distribution detection. *Advances in Neural Information Processing Systems*, 2020. 1, 2, 5, 6, 7
- [32] James MacQueen et al. Some methods for classification and analysis of multivariate observations. In *Proceedings of the fifth Berkeley symposium on mathematical statistics and probability*, volume 1, pages 281–297. Oakland, CA, USA, 1967. 4
- [33] Andrey Malinin and Mark Gales. Predictive uncertainty estimation via prior networks. In *Advances in Neural Information Processing Systems*, pages 7047–7058, 2018. 7
- [34] George A Miller. Wordnet: a lexical database for english. *Communications of the ACM*, 38(11):39–41, 1995. 4
- [35] Sina Mohseni, Mandar Pitale, JBS Yadawa, and Zhangyang Wang. Self-supervised learning for generalizable out-of-distribution detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 5216–5223, 2020. 1, 7
- [36] Eric Nalisnick, Akihiro Matsukawa, Yee Whye Teh, Dilan Gorur, and Balaji Lakshminarayanan. Do deep generative models know what they don’t know? In *International Conference on Learning Representations*, 2018. 8
- [37] Aristotelis-Angelos Papadopoulos, Mohammad Reza Rajati, Nazim Shaikh, and Jiamian Wang. Outlier exposure with confidence control for out-of-distribution detection. *arXiv preprint arXiv:1906.03509*, 2019. 7
- [38] Joseph Redmon and Ali Farhadi. Yolo9000: better, faster, stronger. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7263–7271, 2017. 8
- [39] Jie Ren, Peter J Liu, Emily Fertig, Jasper Snoek, Ryan Poplin, Mark Depristo, Joshua Dillon, and Balaji Lakshminarayanan. Likelihood ratios for out-of-distribution detection. In *Advances in Neural Information Processing Systems*, pages 14680–14691, 2019. 8
- [40] Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. In *International conference on machine learning*, pages 1278–1286. PMLR, 2014. 8
- [41] Ryne Roady, Tyler L. Hayes, Ronald Kemker, Ayesha Gonzales, and Christopher Kanan. Are out-of-distribution detection methods effective on large-scale datasets?, 2019. 1, 8
- [42] Joan Serrà, David Álvarez, Vicenç Gómez, Olga Slizovskaia, José F. Núñez, and Jordi Luque. Input complexity and out-of-distribution detection with likelihood-based generative models. In *International Conference on Learning Representations*, 2020. 8
- [43] Akshayvarun Subramanya, Suraj Srinivas, and R Venkatesh Babu. Confidence estimation in deep neural networks via density modelling. *arXiv preprint arXiv:1707.07013*, 2017. 7
- [44] Esteban G Tabak and Cristina V Turner. A family of non-parametric density estimation algorithms. *Communications on Pure and Applied Mathematics*, 66(2):145–164, 2013. 8
- [45] Aaron Van den Oord, Nal Kalchbrenner, Lasse Espeholt, Oriol Vinyals, Alex Graves, et al. Conditional image generation with pixelcnn decoders. In *Advances in neural information processing systems*, pages 4790–4798, 2016. 8
- [46] Grant Van Horn, Oisín 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. 5, 11
- [47] David Warde-Farley, Andrew Rabinovich, and Dragomir Anguelov. Self-informed neural network structure learning. *arXiv preprint arXiv:1412.6563*, 2014. 8
- [48] Jianxiong Xiao, James Hays, Krista A Ehinger, Aude Oliva, and Antonio Torralba. Sun database: Large-scale scene recognition from abbey to zoo. In *2010 IEEE computer society conference on computer vision and pattern recognition*, pages 3485–3492. IEEE, 2010. 5, 11
- [49] Zhicheng Yan, Hao Zhang, Robinson Piramuthu, Vignesh Jagadeesh, Dennis DeCoste, Wei Di, and Yizhou Yu. Hd-cnn: hierarchical deep convolutional neural networks for large scale visual recognition. In *Proceedings of the IEEE international conference on computer vision*, pages 2740–2748, 2015. 8
- [50] 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. 5, 11