

# MOOD: Multi-level Out-of-distribution Detection

Ziqian Lin\*, Sreya Dutta Roy\*, Yixuan Li  
 Department of Computer Sciences  
 University of Wisconsin-Madison  
 zlin284, duttaroy, sharonli@wisc.edu

## Abstract

*Out-of-distribution (OOD) detection is essential to prevent anomalous inputs from causing a model to fail during deployment. While improved OOD detection methods have emerged, they often rely on the final layer outputs and require a full feedforward pass for any given input. In this paper, we propose a novel framework, multi-level out-of-distribution detection (MOOD), which exploits intermediate classifier outputs for dynamic and efficient OOD inference. We explore and establish a direct relationship between the OOD data complexity and optimal exit level, and show that easy OOD examples can be effectively detected early without propagating to deeper layers. At each exit, the OOD examples can be distinguished through our proposed adjusted energy score, which is both empirically and theoretically suitable for networks with multiple classifiers. We extensively evaluate MOOD across 10 OOD datasets spanning a wide range of complexities. Experiments demonstrate that MOOD achieves up to 71.05% computational reduction in inference, while maintaining competitive OOD detection performance.*

## 1. Introduction

Out-of-distribution (OOD) detection has become a central building block for safely deploying machine learning models in the real world, where the testing data may be distributionally different from the training data. Existing OOD detection methods commonly rely on a scoring function that derives statistics from the penultimate layer or output layer of the neural network [13, 31, 34, 32, 41, 15]. As a result, existing solutions require a full feedforward pass for any given test-time input and utilize a fixed amount of computation. This can be undesirable for safety-critical applications such as self-driving cars, where higher computational cost directly translates into higher latency for the model to take prevention in the presence of OOD driving scenes. Fur-

ther, the computational cost of OOD detection can be exacerbated by the over-parameterization of neural networks, which nowadays have reached unprecedented depth and capacity. For example, recent computer vision models [18] can have over 900 million parameters, which unavoidably incurs high computational demand during inference time. This motivates the following unexplored question: *how can we enable out-of-distribution detection that can adjust and save computations adaptively on-the-fly?*

In this paper, we take the first step to explore the feasibility and efficacy of an adaptive OOD detection framework based on intermediate classifier outputs. Adaptive OOD detection offers several compelling yet untapped advantages. It allows exploiting the intrinsic complexity of OOD examples which may vary at a wide spectrum (see Figure 1). Ideally, easy samples could be detected at the early layer, while more complex ones could still be propagated to the deeper layers for more confident decisions. For example, it might be sufficient to utilize coarse-level features such as color to detect an OOD image of grayscale MNIST digit [25], for a model tasked to classify color images of animals. Secondly, adaptive OOD detection allows flexibility in controlling the computational cost. This is a valuable property in many scenarios, where the computational budget may change over time or vary across different devices.

Formally, we propose a novel framework, Multi-level Out-of-distribution Detection (**MOOD**), which exploits the aforementioned benefits and allows resource-efficient OOD detection. In designing MOOD, we identify three key technical challenges: (1) how to set up an adaptive network with intermediate exits for both OOD detection and classification; (2) how to dynamically choose the optimal exit conditioned on test-time inputs; and (3) at a given exit, how to derive effective OOD scoring function to differentiate between in- vs. out-of-distribution data?

This paper contributes the following technical components by carefully addressing the three challenges above.

- First, we propose *intermediate OOD detectors* operating at varying depths of the network and enabling dynamic OOD inference. Each intermediate detector is

\* Authors contributed equally.

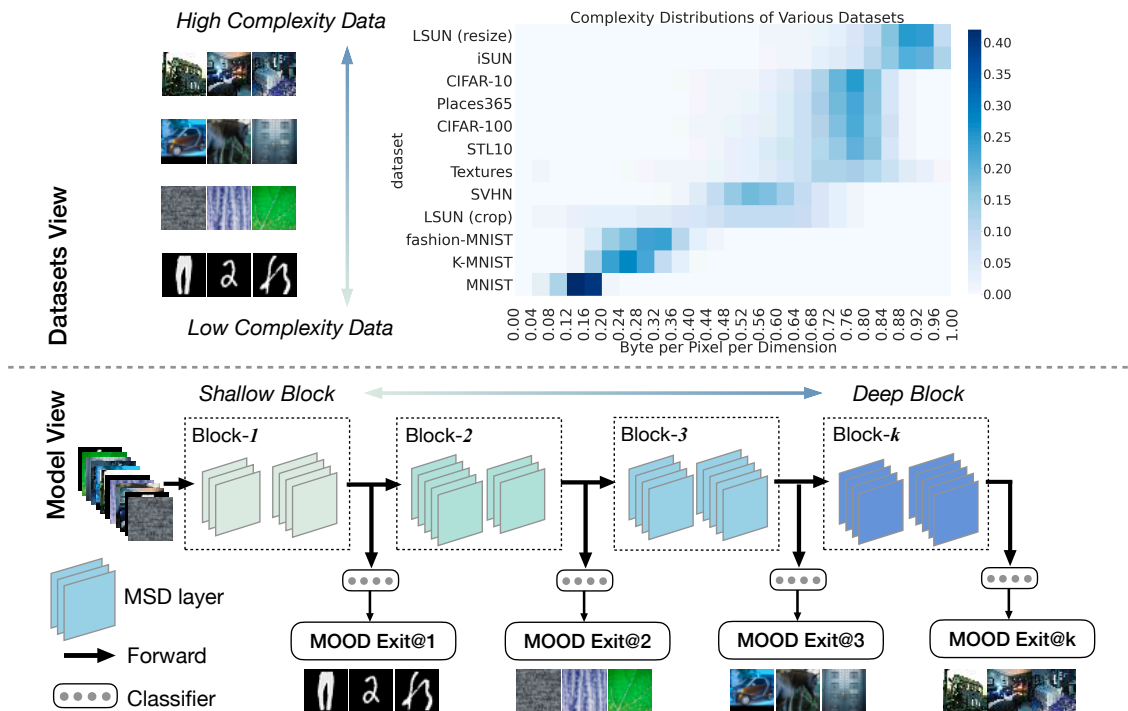


Figure 1. Overview of proposed *Multi-level Out-of-distribution Detection* (MOOD) framework. MOOD exploits the intrinsic complexity of OOD examples which vary at a wide spectrum (top). The adaptive inference network is composed of  $k$  OOD detectors, operating at different depths of the network (bottom). For a given input, a complexity score is used to dynamically determine the exit during inference time. An OOD detector is attached at each exit for differentiating between in- vs. out-of-distribution data.

also referred to as an exit. Whilst our study is certainly inspired by prior works on resource-efficient learning [16, 29], the problem we explore differs substantially: most prior solutions are designed for optimizing the classification accuracy for the in-distribution (ID) task, rather than OOD detection.

- Second, we exploit a novel *complexity-based exit strategy*, which uses model-agnostic complexity scores for determining the intrinsic easiness/hardness of the input data. We show a direct relationship between the OOD data complexity and the optimal exit and show that MOOD can detect easy examples early without propagating to deeper layers. Our setting is more challenging than choosing an optimal exit for the image classification task since OOD data are unexposed to the model during training.
- Thirdly, we introduce an *adjusted energy score* as the OOD scoring function for each intermediate classifier. Our method effectively mitigates the issue of [32], where the non-probabilistic energy score can fluctuate and are not comparable across different exits. We show both empirically and mathematically that adjusted energy scores are suitable for MOOD with multiple classifiers operating at different depths of the network.

We extensively evaluate MOOD on a collection of 10 OOD test datasets, spanning a wide range of complexities. We show a relationship between OOD data complexity and the optimal exit level, where the early exit is more favorably chosen for less complex datasets such as MNIST while deeper exits are preferred for complex datasets such as iSUN. Under the same network architecture and capacity, MOOD with dynamic exit reduces the computational cost by up to **71.05%**, with an average computation deduction by **22.02%** across all 10 datasets. Moreover, MOOD maintains competitive OOD detection performance. Our code and models are available at <https://github.com/deeplearning-wisc/MOOD>.

## 2. Method

Our proposed adaptive OOD detection framework, *Multi-level Out-of-distribution Detection* (MOOD), addresses three key questions and challenges in designing an adaptive OOD detection framework. Firstly, in Section 2.1, we describe the adaptive inference model, which allows early exits for OOD detection at varying depths of the network. In Section 2.2, we then introduce a novel complexity-based exit strategy, which allows dynamically determining the optimal exit level for each input during inference time. Lastly, in Section 2.3, we introduce the OOD inference method operating at different levels of classifiers.

## 2.1. OOD Detector at Early Exits

We begin by addressing the first challenge of *how do we set up an adaptive inference model with early exit for OOD detection?* Inspired by recent works on adaptive neural networks [2, 16, 10, 30, 47, 20, 33, 48, 29], we consider an adaptive inference model composed of  $k$  classifiers. As illustrated in Figure 1, the model can be viewed as a conventional CNN with  $k - 1$  intermediate classifiers attached at varying blocks of the network, where each block contains multiple layers. The model can generate a set consisting of  $k$  predictions, one from each of the exits:

$$f(\mathbf{x}; \theta) = [f_1(\mathbf{x}; \theta_1), \dots, f_k(\mathbf{x}; \theta_k)], \quad (1)$$

where  $\theta$  is the parameterizations of the neural network, and  $f_i(\mathbf{x}; \theta_i)$  represents the output from the classifier at exit  $i \in \{1, 2, \dots, k\}$ . However, prior adaptive networks are designed for optimizing the in-distribution task such as classification or segmentation, and therefore miss the critical component of OOD detection.

In our framework, we introduce intermediate OOD detectors that operate at each level of classifier and enable dynamic OOD inference. Our key idea is that “easy” OOD examples can be captured by early layers, whereas more complex ones should be propagated to the deeper layers for more confident decisions. Specifically, each detector  $G_i(\mathbf{x})$  can be viewed as a binary classifier:

$$G_i(\mathbf{x}; \theta_i) = \begin{cases} \text{in}, & \text{if } S_i(\mathbf{x}; \theta_i) \geq \gamma_i \\ \text{out}, & \text{if } S_i(\mathbf{x}; \theta_i) < \gamma_i, \end{cases}$$

where  $S_i(\mathbf{x}; \theta_i)$  is the scoring function defined for classifier at exit  $i$ .  $\gamma_i$  is the threshold at exit  $i$ , which is chosen so that a high fraction (e.g., 95%) of in-distribution data is correctly classified by the detector. Our multi-level OOD framework allows the full flexibility of early exits at any level, and more importantly, does not incur additional parameterizations on top of the classification networks.

## 2.2. Complexity-based Exit for OOD Detection

In this subsection, we address the second challenge of *how to dynamically choose the optimal exit, conditioned on the test-time inputs?* At test time, we define an exit function  $I(\mathbf{x}) \in \{1, \dots, k\}$ , which determines the index of the exit classifier for a given input  $\mathbf{x}$ . While previous work [29, 16] has relied on prediction confidence to determine the exit for classifying in-distribution data, we argue that such a mechanism is not suitable for OOD detection since the model is not exposed to any OOD data during training. As we show in Section 3, deeper and more confident layers do not necessarily perform better OOD detection. This suggests the need for an exit strategy that does not depend on the model parameterizations.

To this end, we propose a complexity-based exit strategy and derive an exit function  $I(\mathbf{x})$  by exploiting the intrinsic complexity of images that naturally varies at a wide spectrum, from easy to hard. In an attempt to quantify this notion of “easiness”, we consider a *complexity score* as a proxy measurement [42]. Specifically, the complexity of a given image  $\mathbf{x}$  can be upper bounded by a lossless compression algorithm. An input of high complexity will require more bits while a less complex one will be compressed with fewer bits. The complexity for  $\mathbf{x}$  can be defined as the number of bits used to encode the compressed image:

$$L(\mathbf{x}) = \text{Bit length}(c(\mathbf{x})). \quad (2)$$

For the compression function  $c$ , one can use common methods such as PNG [39] and JPEG2000 [43]. We further normalize the score by the maximum complexity score of ID dataset, where  $L_{\text{normalized}} = L(\mathbf{x})/L_{\text{max}}$ . In Figure 1, we show the distribution of complexity scores for different datasets, ranging from the lowest (e.g., MNIST [24]) to the highest (e.g., LSUN [52]).

To make use of the complexity score for inference, we divide the spectrum of complexity scores into  $k$  sub-ranges. For any given example  $\mathbf{x}$  (ID and OOD), we assign exit  $I(\mathbf{x})$  based on the complexity range a sample belongs to:

$$I(\mathbf{x}) = \min(\lceil L_{\text{normalized}}(\mathbf{x}) * k \rceil, k).$$

## 2.3. OOD Scoring Function

OOD detection is a binary classification problem, which commonly relies on a scoring function to distinguish between in- vs. out-of-distribution data. In this subsection, we address the last key challenge in designing MOOD: *what is the suitable scoring function that allows deriving statistics at intermediate classifiers for OOD detection?*

We begin by exploring the energy-based method for OOD detection, inspired by recent work [32]. Specifically, we derive the **exit-wise energy score** based on classifier

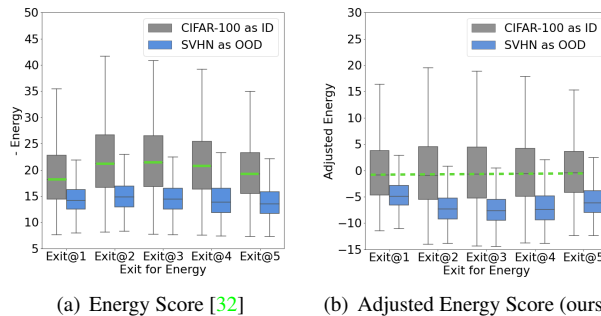


Figure 2. *Left*: Energy scores are not comparable across exits. *Right*: Adjusted energy scores are more comparable across exits (shown in dashed green line).

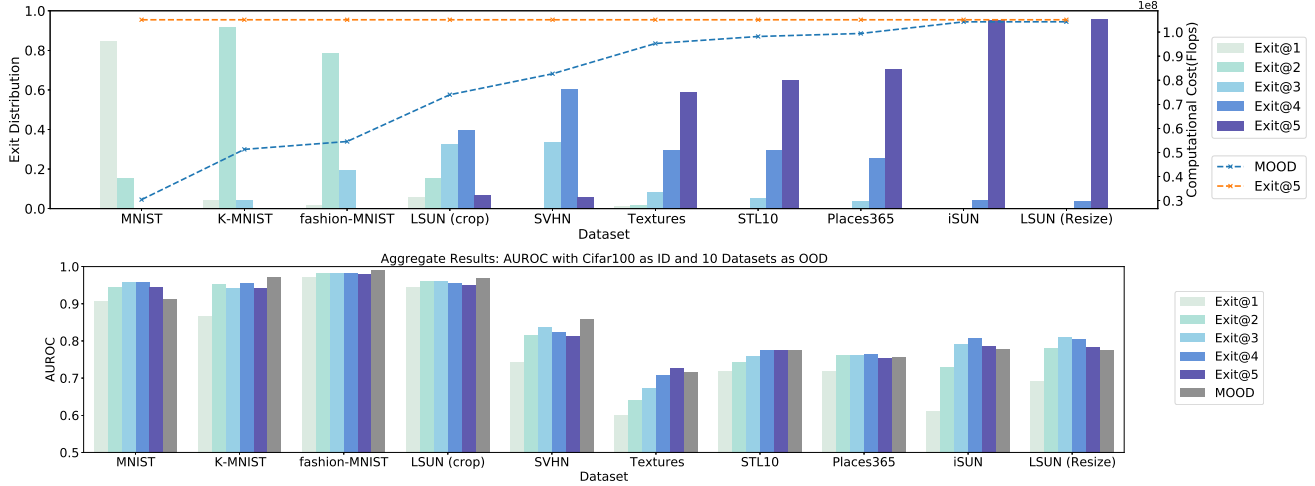


Figure 3. *Top*: Average computational cost (FLOPs) with MOOD, and the normalized frequency distribution of exits chosen by MOOD. Gap in between the orange and blue lines indicates the computational savings. *Bottom*: Average AUROC by taking constant exits at different levels. Model is trained on CIFAR-100 as in-distribution, and evaluated on 10 OOD test datasets described in Section 3.1.

outputs at different exits:

$$E(\mathbf{x}; \theta_i) = -\log \sum_{j=1}^C e^{f_i^{(j)}(\mathbf{x}; \theta_i)}, \quad (3)$$

where  $C$  is the number of classes for in-distribution data, and  $f_i^{(j)}(\mathbf{x}; \theta_i)$  is the logit output corresponding to class  $j \in \{1, 2, \dots, C\}$  for intermediate classifier at exit  $i$ . The likelihood function can be expressed in terms of energy function [25]:

$$p(\mathbf{x}|\theta_i) = \frac{e^{-E(\mathbf{x}; \theta_i)}}{\int_{\mathbf{x}} e^{-E(\mathbf{x}; \theta_i)}}. \quad (4)$$

Taking the logarithm on both sides, we have

$$\log p(\mathbf{x}|\theta_i) = -E(\mathbf{x}; \theta_i) - \log \int_{\mathbf{x}} e^{-E(\mathbf{x}; \theta_i)} \quad (5)$$

$$= -E(\mathbf{x}; \theta_i) - \log Z_i. \quad (6)$$

For a single classifier at exit  $i$ , the energy score  $-E(\mathbf{x}; \theta_i)$  is indicative of the log likelihood  $\log p(\mathbf{x}|\theta_i)$ , since the second term  $\log Z_i$  is a constant for all  $\mathbf{x}$  [32].

However, in our case with multiple classifiers, the second term  $\log Z_i$  cannot be ignored as it depends on exit  $i$ , and can cause variation across classifiers owing to different input features at each exit. To see this, we depict in Figure 2(a) the energy score distributions at each exit. In particular, the energy scores for CIFAR-100 (ID) shift significantly among exits. Ideally, the scores should be comparable regardless of which exit they come from.

**Adjusted Energy Score** To mitigate this issue, we introduce a new scoring function *Adjusted Energy Score*:

$$E_{\text{adjusted}}(\mathbf{x}; \theta_i) = -E(\mathbf{x}; \theta_i) - \mathbb{E}_{\mathbf{x} \in D_{\text{in}}}[-E(\mathbf{x}; \theta_i)], \quad (7)$$

---

**Algorithm 1: MOOD: MULTI-LEVEL OUT-OF-DISTRIBUTION DETECTION**

---

**Input:**  $\mathbf{x}$ , neural network  $f(\mathbf{x}; \theta)$ , number of exits  $k$ , threshold  $\gamma$  chosen on in-distribution data;

**Compute normalized complexity:**

$$L_{\text{normalized}} = L(\mathbf{x})/L_{\text{max}};$$

**Choose exit classifier:**

$$I(\mathbf{x}) = \min(\lceil L_{\text{normalized}}(\mathbf{x}) * k \rceil, k);$$

**if**  $E_{\text{adjusted}}(\mathbf{x}; \theta_{I(\mathbf{x})}) \geq \gamma$  **then**

$$\begin{aligned} & G(\mathbf{x}) = in; \\ & f(\mathbf{x}) = f_{I(\mathbf{x})}(\mathbf{x}; \theta_{I(\mathbf{x})}); \end{aligned}$$

**else**

$$G(\mathbf{x}) = out$$

**end**

---

where  $\mathbb{E}_{\mathbf{x} \in D_{\text{in}}}[-E(\mathbf{x}; \theta_i)]$  is the mean of energy scores derived from exit  $i$ , and in practice can be estimated empirically on the test set of in-distribution data. We show that *Adjusted Energy Score* produces comparable values across exits. To see this, we can rewrite adjusted energy as:

$$\begin{aligned} E_{\text{adjusted}}(\mathbf{x}; \theta_i) &= (\log e^{-E(\mathbf{x}; \theta_i)} - \log Z_i) \\ &\quad - (\mathbb{E}_{\mathbf{x}}[\log e^{-E(\mathbf{x}; \theta_i)} - \log Z_i]) \end{aligned} \quad (8)$$

$$\begin{aligned} &= \log \frac{e^{-E(\mathbf{x}; \theta_i)}}{\int_{\mathbf{x}} e^{-E(\mathbf{x}; \theta_i)}} - \mathbb{E}_{\mathbf{x}} \left[ \log \frac{e^{-E(\mathbf{x}; \theta_i)}}{\int_{\mathbf{x}} e^{-E(\mathbf{x}; \theta_i)}} \right] \end{aligned} \quad (9)$$

$$= \log p(\mathbf{x}|\theta_i) - \mathbb{E}_{\mathbf{x}}[\log p(\mathbf{x}|\theta_i)]. \quad (10)$$

*Adjusted Energy Score* is more comparable across exits for all samples  $\mathbf{x}$ , since the second term is the estimation of average log-likelihood. Depending on how well each

In-distribution (ID)	Architecture	Method	FLOPs	AUROC	FPR95	ID Acc	
			$\downarrow (\times 10^8)$	$\uparrow$	$\downarrow$	$\uparrow (\%)$	
CIFAR-10	WideResNet-40-4	MSP [13]	13.00	0.8898	0.5681	94.93	
		ODIN [31]	13.00	0.9011	0.3531	94.93	
		Mahalanobis [28]	13.00	0.8933	0.3548	94.93	
		Energy [32]	13.00	0.9004	0.3526	94.93	
	MSDNet Exit@last	MSP [13]	1.05	0.8972	0.4987	94.09	
		ODIN [31]	1.05	0.9033	0.2930	94.09	
		Mahalanobis [28]	1.05	0.8284	0.7519	94.09	
		Energy [32]	1.05	0.9048	0.3362	94.09	
	MSDNet (dynamic exit)	MOOD ( <i>ours</i> )	<b>0.79</b>	<b>0.9126</b> ( $\pm 0.0016$ )	<b>0.2805</b> ( $\pm 0.0051$ )	94.13	
	CIFAR-100	WideResNet-40-4	MSP [13]	13.00	0.7710	0.7751	76.90
			ODIN [31]	13.00	0.8466	0.5722	76.90
			Mahalanobis [28]	13.00	0.8319	0.5352	76.90
Energy [32]			13.00	0.8369	0.6271	76.90	
MSDNet Exit@last		MSP [13]	1.05	0.7833	0.7671	75.43	
		ODIN [31]	1.05	0.8489	0.5745	75.43	
		Mahalanobis [27]	1.05	0.7380	0.7806	75.43	
		Energy [32]	1.05	0.8451	0.5915	75.43	
MSDNet (dynamic exit)		MOOD ( <i>ours</i> )	<b>0.79</b>	<b>0.8497</b> ( $\pm 0.0026$ )	<b>0.5722</b> ( $\pm 0.0068$ )	75.26	

Table 1. OOD detection performance comparison between MOOD and baselines. All results here are averaged across 10 datasets. **Bold** numbers are superior results. For MOOD, the complexity calculation takes a negligible amount of time. Therefore, the computations are dominated by the neural network inference cost, measured by FLOPs. The mean and variance of AUROC / FPR95 is reported based on 5 runs. See supplementary A for detailed results for each OOD test dataset.

classifier approximates the log-likelihood, the second term  $\mathbb{E}_{\mathbf{x} \sim D_{in}}[\log p(\mathbf{x}|\theta_i)]$  can be arbitrarily similar across exits as evidenced in Figure 2(b). The score comparability of in-distribution data allows us to use a *single threshold*  $\gamma$  across all exits in inference, which is chosen so that 95% in-distribution data is above the threshold.

We summarize the complete inference-time algorithm for MOOD in Algorithm 1. For input detected as in-distribution, MOOD also performs classification using the classifier of the current exit. We note that previous methods may require tuning hyper-parameters [28], or use auxiliary outlier training data [14]. In contrast, MOOD inference is parameter-free and is easy to use and deploy.

### 3. Experiments

We discuss our experimental setup in Section 3.1, and show that MOOD achieves improved OOD detection performance while reducing computational cost by a large margin in Section 3.2. We also conduct extensive ablation analysis to explore different aspects of our algorithm.

#### 3.1. Setup

**In-distribution Datasets** We use CIFAR-10 and CIFAR-100 [21] as in-distribution datasets, which are common benchmarks for OOD detection. We use the standard split, with 50,000 training images and 10,000 test images. All the images are of size  $32 \times 32$ .

**Out-of-distribution Datasets** For the OOD detection evaluation, we consider a total of 10 datasets with a diverse spectrum of image complexity. In order of increasing complexity, we use MNIST [24], K-MNIST [7], fashion-MNIST [49], LSUN (crop) [52], SVHN [37], Textures [6], STL10 [8], Places365 [54], iSUN [50] and LSUN (resize) [52]. The complexity distribution of each dataset is shown in Figure 1 (top). All images are resized to  $32 \times 32$ . For each OOD dataset, we evaluate on the entire test split. See supplementary A for details.

**Evaluation Metrics** We evaluate MOOD and baseline methods using the following metrics: (1) Number of computational FLOPs during inference time; (2) False Positive Rate (FPR95) on OOD data when the true positive rate for in-distribution data is 95%; and (3) The area under the receiver operating characteristic curve (AUROC).

**Training Details** For training the classification model on in-distribution data, we follow the settings in [16, 29], and use default MSDNet with  $k = 5$  blocks with 4 layers each. We use the default growth rate of 6, with scale factors [1, 2, 4]. We train the model for 300 epochs, using Gradient Equilibrium (GE) proposed in [29]. The initial learning rate is 0.1, which is decayed by a factor of 10 after 150 and 225 epochs. We use SGD (stochastic gradient descent) [3, 45] with a mini-batch size of 64. We use Nesterov momentum [36] with weight 0.9, and a weight decay of

$10^{-4}$ . In Section 3.2.4, we further explore the performance of MOOD with different model capacities by varying the number of layers within each block to be  $\{2, 3, 4, 6, 8, 12\}$ .

## 3.2. Results

### 3.2.1 How does MOOD compare with existing OOD detection methods?

The evaluation results on CIFAR-10 and CIFAR-100 are summarized in Table 1. All numbers are averaged across 10 OOD datasets described in Section 3.1. 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 does not require any auxiliary outlier data. In particular, we compare with MSP [13], ODIN [31], Mahalanobis [28] as well as Energy [32]. All methods above rely on outputs or features extracted from the final exit and consume the same amount of computations measured by FLOPs. Under the same network (MSDNet), MOOD with dynamic exit reduces the computational cost (FLOPs) by up to **71.05%** on the low-complexity dataset, and reduces the average FLOPs by **22.02%** across all 10 datasets. In addition to the computational saving, MOOD achieves better OOD detection performance than the baseline methods. This demonstrates that taking early exit is not only computationally beneficial but also algorithmically desirable for OOD inference.

We note that WideResNet [53] used in prior works is much deeper and computationally more expensive (with 8.9 million parameters), whereas MSDNet with 2.8 million parameters achieves similar OOD detection performance, as shown in Table 1.

Method	SVHN	CIFAR-100	CelebA
Glow [17]	0.64	0.65	0.54
JEM [11]*	0.89	<b>0.87</b>	0.79
Serrà et al [42]*	0.95	0.74	0.86
MOOD	<b>0.96</b>	0.84	<b>0.88</b>

Table 2. Comparison with generative-based models for OOD detection. Values are AUROC. In-distribution dataset is CIFAR-10, which is common setting used in all baselines. \* indicates the variant with best results for the methods.

We also compare with generative-based OOD detection approaches, namely Glow [17], JEM [11], and likelihood ratio-based method [42]. As shown in Table 2, MOOD achieves competitive OOD detection performance with only  $0.79 \times 10^8$  FLOPs, as opposed to Glow-based models [17] ( $4.09 \times 10^9$  FLOPs) or PixelCNN-based model in [42] ( $2.78 \times 10^{10}$  FLOPs).

### 3.2.2 What is the effect of complexity-based dynamic exiting strategy?

To better understand MOOD’s exiting behavior, we show in Figure 3 (top) the normalized frequency distribution of

	FLOPs ↓ ( $\times 10^8$ )	AUROC ↑	FPR95 ↓	ID Acc ↑ (%)
Exit@1	0.267	0.7769	0.7083	65.37
Exit@2	0.516	0.8313	0.5719	71.24
Exit@3	0.689	0.8471	0.5776	74.26
Exit@4	0.884	0.8531	0.5712	75.10
Exit@5	1.051	0.8451	0.5915	75.43
MOOD (ours)	0.794	0.8507	0.5709	75.26

Table 3. Performance comparison between MOOD (dynamic exit) vs. taking constant exits at different levels (non-dynamic exit). Numbers are averaged over 10 OOD datasets with CIFAR-100 as ID dataset.

exits chosen by MOOD, for each OOD dataset. We show that early exits are more frequently chosen for less complex datasets such as MNIST, while deeper exits are chosen more often for complex datasets such as iSUN. The blue curve shows the average FLOPs consumed by MOOD for each dataset. This trend signifies a direct relationship between computational cost (FLOPs) and OOD data complexity, which is effectively exploited by MOOD.

Figure 3 (bottom) depicts the average AUROC by taking constant exits at  $\{1, 2, \dots, 5\}$  for each OOD dataset. The best AUROC chosen among 5 exits can be viewed as the performance upper bound of MOOD, which chooses exit dynamically based on input complexity. We show that early exit is optimal for less complex datasets such as fashion-MNIST while deeper exits are preferred by more complex datasets such as iSUN. Overall, MOOD with dynamic exiting achieves comparable performance as the best possible exit, on all 10 datasets. This trend signifies a direct relationship between the complexity score and an optimal exit. We also show the average results on 10 OOD datasets in Table 3. MOOD overall achieves similar OOD detection performance compared to the best constant exit (upper bound performance), while reducing the computational cost in terms of FLOPs and maintaining a high ID classification accuracy.

We additionally experiment with an alternative lossless compressor JPEG2000 [43] for estimating complexity, which results in similar OOD detection performance (see supplementary B for details).

### 3.2.3 Is MOOD compatible with other scoring functions?

We show that MOOD is a flexible framework that is compatible with alternative scoring functions  $S(\mathbf{x}; \theta)$ . To see this, we replace the adjusted energy score with MSP [13] and ODIN score [31] derived from outputs of intermediate classifiers, and report performance in Table 4. The dynamic exit strategy remains the same, where we utilize the complexity score for determining the exit  $I(\mathbf{x})$ . On

In-distribution	MOOD w/ score	AUROC ( $\uparrow$ )	FPR95 ( $\downarrow$ )
CIFAR-10	MSP	0.8984	0.5001
	ODIN	0.9065	0.2906
	Energy	0.9070	0.3009
	Adjusted Energy	0.9129	0.2834
CIFAR-100	MSP	0.7974	0.7336
	ODIN	0.8548	0.5688
	Energy	0.8433	0.6041
	Adjusted Energy	0.8507	0.5709

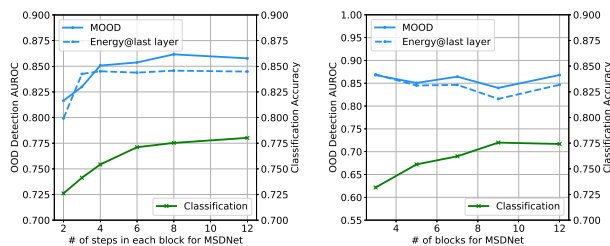
Table 4. Effect of scoring functions in MOOD. All numbers aggregated over 10 datasets.

CIFAR-10, using MOOD with adjusted energy score results in the optimal performance with average AUROC 91.29%. While the performance using the ODIN score on CIFAR-100 is slightly better, we note that it requires tuning hyper-parameters. In contrast, MOOD with adjusted energy scores is completely parameter-free and is easy to use and deploy.

Further, by contrasting the performance with Table 1, we observe that MOOD consistently improves the performance compared to using final outputs (*i.e.*, MSDNet Exit@last), regardless of the scoring function.

### 3.2.4 How does model capacity affect MOOD performance?

We explore the effect of model capacity on MOOD performance by increasing the depth of MSDNet. First, we fix the number of blocks  $k = 5$  and vary the number of steps for each block to be  $\{2, 3, 4, 6, 8, 12\}$ , resulting in network configurations with varying total depths. As shown in Figure 4(a), we observe that MOOD with dynamic exit works better than the best baseline [32] using the last layer output. Additionally, we explore the effect of varying number of blocks  $k \in \{3, 5, 7, 9, 12\}$ , as shown in Figure 4(b). In both cases, the improvement is more pronounced as the network gets deeper. This suggests the applicability of MOOD at different model capacities.



(a) varying layers per block ( $k = 5$ ) (b) varying number of blocks  $k$   
 Figure 4. Effect of model capacity on MOOD performance (CIFAR-100 as ID). *Left*: we fix the number of blocks  $k = 5$  and vary the number of steps in each block. *Right*: we vary the number of blocks. Blue curves indicate OOD detection performance (AUROC). Green curves indicate classification accuracy on in-distribution data.

### 3.2.5 How does MOOD affect in-distribution classification accuracy?

We show in Table 1 that MOOD enables efficient OOD detection while maintaining the classification accuracy on in-distribution data. During inference time, we apply complexity-based strategy to determine the exit  $I(\mathbf{x})$  for any given in-distribution test input  $\mathbf{x}$ , and use the classification predictions made by the corresponding classifier  $f(\mathbf{x}; \theta_{I(\mathbf{x})})$ . This confirms that our MOOD framework provides several desiderata *altogether*, including (1) safety guarantee against OOD examples, (2) accurate and comparable predictions on in-distribution examples, and (3) computational efficiency during inference.

### 3.2.6 Comparison with Greedy-based and Random Exit Strategy

As an ablation on the effect of alternative exit strategies, we attempt a greedy-based method that takes the exit  $i$  whenever the adjusted energy score is below the threshold  $\gamma_i$ . The algorithm is summarized in Algorithm 2. In contrast to our complexity-based method, the greedy-based strategy depends on the model parameterizations and thresholds for each intermediate OOD exit.  $\gamma_i$  is chosen for each exit  $i$  so that 95% of the in-distribution data is correctly classified.

In Table 5, we show results comparison between complexity- vs. greedy-based exiting strategies. The greedy-based method is suboptimal for the task of OOD detection. For our experiment on MSDNet with 5 blocks, the greedy method leads to a 9.26% higher FPR95 compared to our complexity-based method. Since the model is not trained with OOD data, the thresholds are less indicative of the optimal OOD exit.

Lastly, we also consider a randomized exit strategy, where we randomly select an exit  $\{1, 2, \dots, k\}$  for a given sample during test time. This can be viewed as lower bounding the performance. As expected, the results in Table 5 show a substantial performance drop in both OOD detection as well as in-distribution classification accuracy.

Exit strategy	FLOPs $\downarrow (\times 10^8)$	AUROC $\uparrow$	FPR95 $\downarrow$	ID Acc
Complexity	0.7943	<b>0.8507</b>	<b>0.5709</b>	<b>0.7526</b>
Greedy	0.6890	0.8400	0.6635	0.7384
Randomized	<b>0.6880</b>	0.8298	0.6078	0.7299

Table 5. Comparison between complexity-based, greedy-based and randomized exiting strategies. Numbers are averaged over 10 OOD datasets with CIFAR-100 as ID dataset.

## 4. Related Work

**Adaptive Inference Networks** The concept of anytime exit has been previously explored by FractalNets [23],

---

**Algorithm 2:** MOOD WITH GREEDY-BASED EXIT

---

**Input:**  $\mathbf{x}$ , neural network  $f(\mathbf{x}; \theta)$ ;  
**while**  $i \leq k$  **do**  
  Calculate adjusted energy score at exit  $i$ :  
     $E_{\text{adjusted}}(\mathbf{x}; \theta_i)$ ;  
    **if**  $E_{\text{adjusted}}(\mathbf{x}; \theta_i) \leq \gamma_i$  **then**  
       $G(\mathbf{x}) = \text{out}$ ;  
    **else**  
       $i++$ ;  
    **end**  
**end**  
 $G(\mathbf{x}) = \text{in}$

---

deeply supervised networks [26] and MSDNet [16]. FractalNets [23] introduced multiple paths that were evaluated during inference in order of increasing computational requirement. Big-Little Net [4] used multiple branches with different computational complexities and merged them to use features of different scales. Dynamic inference methods [2, 10, 12, 51, 10, 48, 46, 47] adapted the inference to each test instance through gating functions or learning policies, which reduced computation by skipping units or even entire layers. MSDNet [16, 29] also allowed classification for early exiting. Kong and Fowlkes [20] constructed Pixel-wise Attentional Gating units (PAG) for adaptive inference by learning a dynamic computation path for each pixel. McIntosh *et al.* [33] considered the property of RNNs, varied the number of iterations of the RNN to achieve a flexible range of computational budgets during Segmentation. However, prior solutions are designed for optimizing the in-distribution task such as classification. Our work instead focuses on adaptive inference for OOD detection and introduces intermediate detectors operating at different depths of the network.

**OOD Detection Using Final Outputs** A baseline method for OOD detection was introduced in [13], which uses maximum softmax probability (MSP) from a pre-trained network. Several works attempt to improve the OOD uncertainty estimation using the ODIN score [31], deep ensembles [22], generalized ODIN score [15], the energy score [32], and confidence score estimated from a special branch in the model [9]. Several loss functions have been proposed to regularize model predictions of the auxiliary outlier data towards uniform distributions [27], a background class for OOD data [5, 34], or higher energies [32]. However, previous methods have mostly relied on the last layer or the penultimate layer of the neural network. Our paper explores OOD detection by dynamically leveraging intermediate exits, which offers compelling advantages both in computation and detection accuracy.

**OOD Detection with Intermediate Information** Several works explored OOD detection using feature representations [28] or outputs [1] at different intermediate layers of the networks, which signify the usefulness of early layers. Particularly, Lee *et al.* [28] proposed a scoring function that combined the Mahalanobis distances derived from intermediate features of the neural networks. Abdelzad *et al.* [1] explored exiting at an early layer for OOD detection and required auxiliary dataset to train intermediate OOD detectors. However, in all three methods, there is no adaptive decision made while exiting, and hence computational bottleneck remains unsolved. In contrast, our work takes the first step to explore the feasibility and efficacy of an adaptive OOD detection framework based on intermediate classifier outputs. It is also worth noting that many existing approaches have a number of hyper-parameters that need to be tuned [28], sometimes with the help of outlier or additional data [1, 28]. MOOD with the complexity and adjusted energy scores is parameter-free and does not require any training with auxiliary data.

**Complexity and OOD Detection** Sabeti *et al.* [40] combined concepts of typicality and minimum description length to perform novelty detection. They considered atypical sequences that can be described (coded) with fewer bits. Serrà *et al.* [42] used input complexity to derive OOD score for generative models such as Glow [17], which improved likelihood-based methods [35, 38]. Since it is intractable to compute Kolmogorov complexity [19], one can use a lossless compression algorithms such as PNG [39], JPEG2000 [43] and FLIF [44] as a proxy. Different from [42], we use complexity in a novel context with multi-level exits and establish the relationship between the complexity of OOD data and suitable exit for dynamic inference.

## 5. Conclusion

In this work, we propose a novel framework, *Multi-level Out-of-distribution Detection* (MOOD), which dynamically exploits early classifier outputs for OOD inference. The key idea is to detect easy OOD examples earlier in the network while propagating hard OOD examples to deeper layers. We establish a relationship between OOD data complexity and optimal exit level, which is effectively exploited by MOOD. We extensively evaluate MOOD across 10 OOD datasets spanning a wide range of complexities. We show that MOOD achieves better OOD detection performance than previous approaches relying on final outputs, and at the same time reduces the computations by up to 71.05%. We hope that future research will increase the attention towards a broader view of the computational efficiency aspect of OOD detection.



## References

- [1] Vahdat Abdelzad, Krzysztof Czarnecki, Rick Salay, Taylor Denouden, Sachin Vernekar, and Buu Phan. Detecting out-of-distribution inputs in deep neural networks using an early-layer output. *arXiv preprint arXiv:1910.10307*, 2019. 8
- [2] Tolga Bolukbasi, Joseph Wang, Ofer Dekel, and Venkatesh Saligrama. Adaptive neural networks for efficient inference. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*, pages 527–536, 2017. 3, 8
- [3] Léon Bottou. Online learning and stochastic approximations. *On-line learning in neural networks*, 17(9):142, 1998. 5
- [4] Chun-Fu (Richard) Chen, Quanfu Fan, Neil Mallinar, Tom Sercu, and Rogério Schmidt Feris. Big-little net: An efficient multi-scale feature representation for visual and speech recognition. In *Proceedings of the 7th International Conference on Learning Representations (ICLR)*, 2019. 8
- [5] Jiefeng Chen, Yixuan Li, Xi Wu, Yingyu Liang, and Somesh Jha. Informative outlier matters: Robustifying out-of-distribution detection using outlier mining, 2020. 8
- [6] Mircea Cimpoi, Subhansu Maji, Iasonas Kokkinos, Sammy Mohamed, and Andrea Vedaldi. Describing textures in the wild. In *Proceedings of 2014 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3606–3613, 2014. 5
- [7] Tarin Clanuwat, Mikel Bober-Irizar, Asanobu Kitamoto, Alex Lamb, Kazuaki Yamamoto, and David Ha. Deep learning for classical japanese literature. In *Proceedings of Neural Information Processing Systems 2018 Workshop on Machine Learning for Creativity and Design (NeurIPS)*, 2018. 5
- [8] Adam Coates, Andrew Ng, and Honglak Lee. An analysis of single-layer networks in unsupervised feature learning. In *Proceedings of the 14th International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 215–223, 2011. 5
- [9] Terrance DeVries and Graham W Taylor. Learning confidence for out-of-distribution detection in neural networks. *arXiv preprint arXiv:1802.04865*, 2018. 8
- [10] Michael Figurnov, Maxwell D Collins, Yukun Zhu, Li Zhang, Jonathan Huang, Dmitry Vetrov, and Ruslan Salakhutdinov. Spatially adaptive computation time for residual networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1039–1048, 2017. 3, 8
- [11] Will Grathwohl, Kuan-Chieh Wang, Joern-Henrik Jacobsen, David Duvenaud, Mohammad Norouzi, and Kevin Swersky. Your classifier is secretly an energy based model and you should treat it like one. In *Proceedings of the 8th International Conference on Learning Representations (ICLR)*, 2020. 6
- [12] Alex Graves. Adaptive computation time for recurrent neural networks. *arXiv preprint arXiv:1603.08983*, 2016. 8
- [13] Dan Hendrycks and Kevin Gimpel. A baseline for detecting misclassified and out-of-distribution examples in neural networks. In *Proceedings of the 5th International Conference on Learning Representations (ICLR)*, 2017. 1, 5, 6, 8
- [14] Dan Hendrycks, Mantas Mazeika, and Thomas G. Dietterich. Deep anomaly detection with outlier exposure. In *Proceedings of the 7th International Conference on Learning Representations (ICLR)*, 2019. 5
- [15] 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 (CVPR)*, pages 10951–10960, 2020. 1, 8
- [16] Gao Huang, Danlu Chen, Tianhong Li, Felix Wu, Laurens van der Maaten, and Kilian Q. Weinberger. Multi-scale dense networks for resource efficient image classification. In *Proceedings of the 6th International Conference on Learning Representations (ICLR)*, 2018. 2, 3, 5, 8
- [17] Diederik P. Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. In *Proceedings of Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems (NeurIPS)*, pages 10236–10245, 2018. 6, 8
- [18] Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Joan Puigcerver, Jessica Yung, Sylvain Gelly, and Neil Houlsby. Big transfer (bit): General visual representation learning. In *Proceedings of the 16th European Conference on Computer Vision (ECCV)*, pages 491–507, 2020. 1
- [19] Andrei N. Kolmogorov. On tables of random numbers (reprinted from "sankhya: The indian journal of statistics", series a, vol. 25 part 4, 1963). *Theor. Comput. Sci.*, 207(2):387–395, 1998. 8
- [20] Shu Kong and Charless Fowlkes. Pixel-wise attentional gating for parsimonious pixel labeling. *arXiv preprint arXiv:1805.01556*, 2018. 3, 8
- [21] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 5
- [22] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. In *Proceedings of Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems (NeurIPS)*, pages 6402–6413, 2017. 8
- [23] Gustav Larsson, Michael Maire, and Gregory Shakhnarovich. Fractalnet: Ultra-deep neural networks without residuals. In *Proceedings of the 5th International Conference on Learning Representations (ICLR)*, 2017. 7, 8
- [24] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. 3, 5
- [25] Yann LeCun, Sumit Chopra, Raia Hadsell, M Ranzato, and F Huang. A tutorial on energy-based learning. *Predicting structured data*, 1(0), 2006. 1, 4
- [26] Chen-Yu Lee, Saining Xie, Patrick W. Gallagher, Zhengyou Zhang, and Zhuowen Tu. Deeply-supervised nets. In *Proceedings of the 18th International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2015. 8
- [27] Kimin Lee, Honglak Lee, Kibok Lee, and Jinwoo Shin. Training confidence-calibrated classifiers for detecting out-of-distribution samples. In *Proceedings of the 6th International Conference on Learning Representations (ICLR)*, 2018. 5, 8
- [28] Kimin Lee, Kibok Lee, Honglak Lee, and Jinwoo Shin. A simple unified framework for detecting out-of-distribution

- samples and adversarial attacks. In *Proceedings of Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems (NeurIPS)*, pages 7167–7177, 2018. 5, 6, 8
- [29] Hao Li, Hong Zhang, Xiaojuan Qi, Ruigang Yang, and Gao Huang. Improved techniques for training adaptive deep networks. In *Proceedings of IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 1891–1900, 2019. 2, 3, 5, 8
- [30] Zhichao Li, Yi Yang, Xiao Liu, Feng Zhou, Shilei Wen, and Wei Xu. Dynamic computational time for visual attention. In *Proceedings of IEEE International Conference on Computer Vision Workshops (ICCV)*, pages 1199–1209, 2017. 3
- [31] Shiyu Liang, Yixuan Li, and Rayadurgam Srikant. Enhancing the reliability of out-of-distribution image detection in neural networks. In *Proceedings of the 6th International Conference on Learning Representations (ICLR)*, 2018. 1, 5, 6, 8
- [32] Weitang Liu, Xiaoyun Wang, John D. Owens, and Yixuan Li. Energy-based out-of-distribution detection. In *Proceedings of Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2020. 1, 2, 3, 4, 5, 6, 7, 8
- [33] Lane McIntosh, Niru Maheswaranathan, David Sussillo, and Jonathon Shlens. Recurrent segmentation for variable computational budgets. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPR)*, pages 1648–1657, 2018. 3, 8
- [34] Sina Mohseni, Mandar Pitale, JBS Yadawa, and Zhangyang Wang. Self-supervised learning for generalizable out-of-distribution detection. In *Proceedings of the 34 AAAI Conference on Artificial Intelligence (AAAI)*, pages 5216–5223, 2020. 1, 8
- [35] Eric T. Nalisnick, Akihiro Matsukawa, Yee Whye Teh, Dilan Görür, and Balaji Lakshminarayanan. Do deep generative models know what they don’t know? In *Proceedings of the 7th International Conference on Learning Representations (ICLR)*, 2019. 8
- [36] Yurii Nesterov. A method for unconstrained convex minimization problem with the rate of convergence  $O(1/k^2)$ . In *Doklady an ussr*, volume 269, pages 543–547, 1983. 5
- [37] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bisacco, Bo Wu, and Andrew Y Ng. Reading digits in natural images with unsupervised feature learning. In *Proceedings of NIPS Workshop on Deep Learning and Unsupervised Feature Learning (NeurIPS)*, 2011. 5
- [38] 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 *Proceedings of Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems (NeurIPS)*, pages 14680–14691, 2019. 8
- [39] Greg Roelofs and Richard Koman. *PNG: the definitive guide*. O’Reilly & Associates, Inc., 1999. 3, 8
- [40] Elyas Sabeti and Anders Høst-Madsen. Data discovery and anomaly detection using atypicality for real-valued data. *Entropy*, 21(3):219, 2019. 8
- [41] Vikash Sehwal, Arjun Nitin Bhagoji, Liwei Song, Chawin Sitawarin, Daniel Cullina, Mung Chiang, and Prateek Mittal. Analyzing the robustness of open-world machine learning. In *Proceedings of the 12th ACM Workshop on Artificial Intelligence and Security (ACM)*, pages 105–116, 2019. 1
- [42] Joan Serra, 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 *Proceedings of the 8th International Conference on Learning Representations (ICLR)*, 2020. 3, 6, 8
- [43] Athanassios Skodras, Charilaos Christopoulos, and Touradj Ebrahimi. The jpeg 2000 still image compression standard. *IEEE Signal processing magazine*, 18(5):36–58, 2001. 3, 6, 8
- [44] Jon Sneyers and Pieter Wuille. Flif: Free lossless image format based on maniac compression. In *Proceedings of the IEEE International Conference on Image Processing (ICIP)*, pages 66–70, 2016. 8
- [45] Ilya Sutskever, James Martens, George Dahl, and Geoffrey Hinton. On the importance of initialization and momentum in deep learning. In *Proceedings of the 30th International Conference on Machine Learning (ICML)*, pages 1139–1147, 2013. 5
- [46] Andreas Veit and Serge Belongie. Convolutional networks with adaptive inference graphs. In *Proceedings of the 15th European Conference on Computer Vision (ECCV)*, pages 3–18, 2018. 8
- [47] Xin Wang, Fisher Yu, Zi-Yi Dou, Trevor Darrell, and Joseph E Gonzalez. Skipnet: Learning dynamic routing in convolutional networks. In *Proceedings of the 15th European Conference on Computer Vision (ECCV)*, pages 409–424, 2018. 3, 8
- [48] Zuxuan Wu, Tushar Nagarajan, Abhishek Kumar, Steven Rennie, Larry S Davis, Kristen Grauman, and Rogerio Feris. Blockdrop: Dynamic inference paths in residual networks. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8817–8826, 2018. 3, 8
- [49] Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017. 5
- [50] Pingmei Xu, Krista A Ehinger, Yinda Zhang, Adam Finkelstein, Sanjeev R Kulkarni, and Jianxiong Xiao. Turkergaze: Crowdsourcing saliency with webcam based eye tracking. *arXiv preprint arXiv:1504.06755*, 2015. 5
- [51] Chris Ying and Katerina Fragkiadaki. Depth-adaptive computational policies for efficient visual tracking. In *Proceedings of Energy Minimization Methods in Computer Vision and Pattern Recognition - 11th International Conference (EMMCVPR)*, pages 109–122, 2017. 8
- [52] Fisher Yu, Ari Seff, Yinda Zhang, Shuran Song, Thomas Funkhouser, and Jianxiong Xiao. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365*, 2015. 3, 5
- [53] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. In *Proceedings of the British Machine Vision Conference (BMVC)*, 2016. 6

- [54] 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](#)