A Novel Data Augmentation Technique for Out-of-Distribution Sample Detection using Compounded Corruptions

Ramya Hebbalaguppe Soumya Ghosal Jatin Prakash Harshad Khadilkar
Chetan Arora

Indian Institute of Technology Delhi and TCS Research

Abstract. Modern deep neural network models are known to erroneously classify out-of-distribution (OOD) test data into one of the in-distribution (ID) training classes with high confidence. This can have disastrous consequences for safety-critical applications. A popular mitigation strategy is to train a separate classifier that can detect such OOD samples at the test time. In most practical settings OOD examples are not known at the train time, and hence a key question is how to augment the ID data with synthetic OOD samples for training such an OOD detector. In this paper, we propose a novel ***C**ompou**n**ded **C**orruption* technique for the OOD data augmentation. One of the major advantages of our technique, named CnC, is that it does not require any hold-out data apart from the training set. Further, unlike current state-of-the-art (SOTA) techniques, CnC does not require backpropagation or ensembling at the test time, making our method much faster at inference. Our extensive comparison with 20 methods from the major conferences in last 4 years show that a model trained using CnC based data augmentation, significantly outperforms SOTA, both in terms of OOD detection accuracy as well as inference time. We include a detailed post-hoc analysis to investigate the reasons for the success of our method and identify higher relative entropy and diversity of CnC samples as probable causes. We also provide useful insights via a piece-wise decomposition analysis on a two-dimensional dataset to reveal (visually and quantitatively) that our approach leads to a tighter boundary around ID classes, leading to better detection of OOD samples. The source-code of our work can be found at: https://github.com/cnc-ood.

1 Introduction

Deep neural network (DNN) models generalize well when the test data is independent and identically distributed (IID) with respect to training data [45]. However, the condition is difficult to enforce in the real world due to distributional drifts, covariate shift, and/or adversarial perturbations. Thus, a *reliable* system based on a DNN model must be able to detect an OOD sample, and either abstain from making any decision on such samples, or flag them for human intervention. We assume that the in-distribbution (ID) samples belong to one of the K known classes, and club all OOD samples into a new class called a *reject*/OOD class. We do not attempt to identify which specific class (unseen label) the unknown sample belongs to. Our goal is to build a classifier that can accurately detect OOD samples as the $(K+1)^{\text{th}}$ OOD class, with an objective of rejecting samples belonging to any novel class.

Most techniques for OOD detection assume the availability of validation samples from the OOD set for tuning model hyper-parameters [34,32,2,19]. Based on the samples, the techniques either update the model weights so as to predict lower scores for the OOD samples, or try to learn correlation between activations and the output score vector [32]. Such approaches have limited utility as in most practical scenarios, either the OOD samples are not available, or cover a tiny fraction of OOD sample space. Yet, other class of techniques learn the threshold on the uncertainty of the output score using deep ensembling [29] or MC dropout [10]. Understandably, OOD detection capability of these techniques suffer when the samples from a different OOD domain are presented.

The other popular class of OOD detectors do not use representative samples from OOD domain, but generate them synthetically [8,17,37,40]. The synthetic samples can be used to train any of the earlier mentioned SOTA models in lieu of the real OOD samples. This obviates the need for any domain specific OOD validation set. Such methods typically use natural corruptions (e.g. blur, noise, and geometric transformations etc.) or adversarial perturbations to generate samples near decision boundary of a classifier. This class also have limited accuracy on real OOD datasets, as the synthetic images generated in such a way are visually much similar to the ID samples, and the behavior of a DNN when shown natural OOD images much farther (in terms of ℓ_2 distance in RGB space) from the ID samples still remains unknown.

Recent theoretical works towards estimating or minimizing open set loss recommend training with OOD samples covering as much of the probable input space as possible. For example, [25] show that a piece-wise DNN model shatters the input space into a polyhedral complex, and prove that empirical risk of a DNN model in a region of input space scales inversely with the density of training samples lying inside the polytope corresponding to the region. Similarly, [9] show that under an unknown OOD distribution, the best way to minimize the open set loss is by choosing OOD samples uniformly from the support set in the input space. Encouraged by such theoretical results, we propose a data augmentation technique which does not focus on generating samples visually similar to the ID samples but synthesizing OOD samples in two key regions of the input space: (i) finely distributed at the boundary of ID classes, and (ii) coarsely distributed in the inter-ID sample space (See Sec. 3.3 for details).

Contributions: We make the following key contributions:

1. We propose a novel data augmentation strategy, CompouNded Corruptions (CnC) for OOD detection. Unlike contemporary techniques [13,19,32,34] the proposed approach does not need a separate OOD train or validation dataset.
2. Unlike SOTA techniques which detect OOD samples by lowering the confidence of ID classes [1,18,32,36], we classify OOD samples into a separate reject class. We show empirically that our approach leads to clearer separation between ID and OOD samples in the embedding space (c.f. Fig. 5).
3. Our method does not require any input pre-processing at the test time, or a second forward pass with perturbation/noise. This makes it significantly faster in inference as compared to the other SOTA methods [22,34].
4. Visualization and analysis of our results indicate that finer granularity of the polyhedral complex around the ID regions learnt by a model is a good indicator of performance of a OOD data augmentation technique. Based on our analysis, we also

recommend higher entropy and diversity of generated OOD samples as good predictors for OOD detection performance.
5. Our approach yields SOTA results on multiple benchmark datasets. Using DenseNet [23] feature extractor, CIFAR-100 as ID dataset, and SVHN [41] as OOD, we achieve TNR@TPR95 of 98.7%, AUROC of 99.7% and Detection Error of 2%. This outperforms Generalized-ODIN [22] (CVPR'20) by a large margin of 18.4% in TNR@TPR95. Our method also outperforms Gram Matrices (ICML'20) (c.f. Tab. 1).

2 Related Work

Our approach is a hyper-parameter-free OOD detection technique, which does not need access to a validation OOD dataset. Below, we give a detailed review of the contemporary works in terms of the above requirements.

Hyper-parameter tuning using OOD data: This class comprises of OOD detection methods that fine-tune hyper-parameters on a validation set. ODIN [34] utilizes temperature scaling with input perturbations using the OOD validation dataset to tune hyper-parameters for calibrating the neural networks. However, hyper-parameters tuned with one OOD dataset may not generalize to other datasets. Lee et al.[32] propose training a logistic regression detector on the Mahalanobis distance vectors calculated between test images' feature representations and class conditional Gaussian distribution at each layer.

Retraining a model using OOD data: G-ODIN [22] decompose confidence score along with modified input pre-processing for detecting OOD, whereas ATOM [2] essentially makes a model robust to the small perturbations, and hard negative mining for OOD samples. MOOD [35] introduce multi-level OOD detection based on the complexity of input data, and exploit simpler classifier for faster OOD inference. Most recently, MOS [24] suggests decomposing the large semantic space into similar concepts to simplify the decision boundaries between ID-OOD data and demonstrate on large scale datasets.

Using a pre-trained model's score for OOD detection: Hendrycks and Gimpel [18] use maximum confidence scores from a softmax output to detect OOD. Liu et al.[36] use energy as a scoring function for OOD detection without tuning hyper-parameters. Shastry and Oore [44] leverage p^{th}-order Gram matrices to identify anomalies between activity patterns and the predicted class. Blundell et al.[1] focus on a closed world assumption which forces a DNN to choose from one of the ID classes, even for the OOD data. They introduce a new model layer, *OpenMax*, that estimates the probability of an input being from an unknown class using a Weibull distribution. G-OpenMax[11] explicitly model OOD samples and its decision score and report findings on small datasets like MNIST.

OOD detection using uncertainty estimation: OOD samples can be rejected by thresholding on the uncertainty measure. Graves et al.[12] and Wen et al.[50] propose anomaly detection based on stochastic Bayesian inference. Gal et al.[10] propose MC-dropout to measure uncertainty of a model using multiple inferences. Deep Ensembles [29] use multiple networks trained independently to improve uncertainty estimation.

Data augmentation for OOD detection: This line of research augments the training set to improve OOD detection. Data augmentations like flipping and cropping generate samples that can be easily classified by a pre-trained classifier. Generative techniques based on VAEs, and GANs try to synthesize data samples near the decision boundary

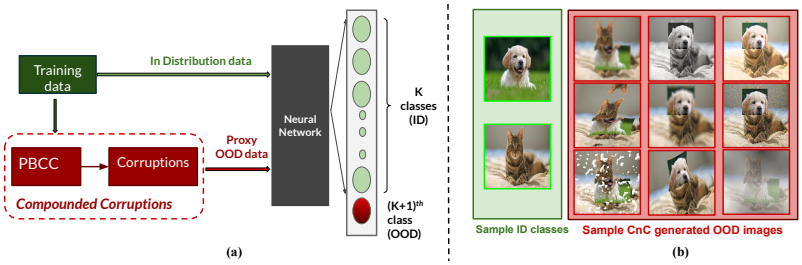

Fig. 1: Creating augmented data samples using Compounded Corruptions (CnC). Pane (a) shows block diagram of the training procedure: first we take a patch based convex combination (PBCC) of patches chosen from image pair belonging to $\binom{K}{2}$ labels; second, we apply corruptions on the data points obtained using PBCC. This proxy OOD data is then used to train a $(K+1)$ way classifier, where, first K classes correspond to the ID classes and $(K+1)^{\text{th}}$ class contains synthesized OOD samples corresponding to reject/OOD class. Pane (b) shows CnC synthesized sample images from cat and dog classes. Intuitively, CnC gives two knobs for generating OOD samples: a coarse exploration ability through linear combination of two ID classes achieved through PBCC operation, and a finer warping capability through corruption of these images. The order of the two operations (PBCC before corruption) is important, as we show later.

[8,31,33,42,51,49]. Other data augmentation strategies do not directly target OOD detection, but domain generalization: SaliencyMix [47], CutOut[7], GridMask[3], AugMix [20], RandomErase [56], PuzzleMix [27], RandAugment [5], SuperMix [6]. Mixup [55] generates new data through convex combination of training samples and labels to improve DNN generalization. CutMix [52] which generates samples by replacing an image region with a patch from another training image. The approach is not directly suitable for OOD detection, as the generated samples lie on the line joining the training samples, and may not cover the large input space[25,9].

3 Proposed Approach

3.1 Problem Formulation

We consider a training set, $\mathcal{D}_{\text{in}}^{\text{train}}$, consisting of N training samples: $(x_n, y_n)_{n=1}^{N}$, where samples are drawn independently from a probability distribution: $\mathcal{P}_{X,Y}$. Here, $X \in \mathcal{X}$ is a random variable defined in the image space, and $Y \in \mathcal{Y} = \{1, \ldots, K\}$ represents its label. Traditionally, a classifier $f_\theta : \mathcal{X} \to \mathcal{Y}$ is trained on in-distribution samples drawn from a marginal distribution $\mathcal{P}_X$ of X derived from the joint distribution $\mathcal{P}_{X,Y}$. Here θ refers to model parameters. Let $\mathcal{Q}_X$ be another distinct data distribution defined on the image space $\mathcal{X}$. During testing phase, input images are drawn from a conditional mixture distribution $\mathcal{M}_{X|Z}$ where $Z \in \{0, 1\}$, such that $\mathcal{M}_{X|Z=0} = \mathcal{P}_X$, and $\mathcal{M}_{X|Z=1} = \mathcal{Q}_X$. We define all $\mathcal{Q}_X \nsim \mathcal{P}_X$ as OOD distributions, and Z is a latent (binary) variable to denote ID if $Z = 0$ and OOD if $Z = 1$.

One possible approach to detecting an OOD sample is if confidence of f_θ for a given input is low for all elements of $\mathcal{Y}$. However, we use an alternative approach where we learn to map OOD samples generated using our technique to an additional label $(K+1)$. Given any two ID samples $x_1, x_2 \sim \mathcal{P}_X$, we generate the synthetic data using the CnC operation $C(x_1, x_2) : \mathcal{X} \times \mathcal{X} \to \mathcal{X}$. We then define an extended label set

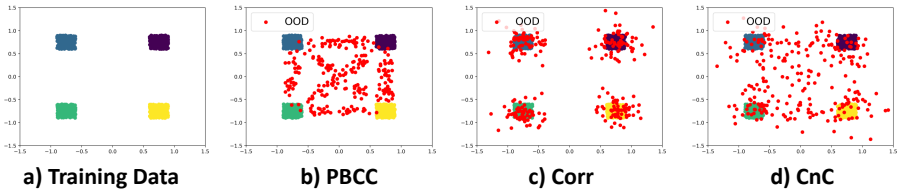

Fig. 2: Intuition and illustrative plot of OOD synthesis on a toy dataset with four ID classes. Each sample is in $\mathbb{R}^2$. Consider $\mathbf{p}_1 = (x_1, y_1)$, and $\mathbf{p}_2 = (x_2, y_2)$ to be the two input samples belonging to distinct classes 1 and 2, then $\mathbf{p}_3 = (x_3, y_3)$ is the geometric convex combination of $\mathbf{p}_1$ and $\mathbf{p}_2$ such that: $\mathbf{p}_3 = \lambda \mathbf{p}_1 + (1-\lambda)\mathbf{p}_2$, $0 \leq \lambda \leq 1$. (a) training data corresponding to 4 distinct classes in shades of teal, green, purple, and yellow; Synthesised OOD points are in red; (b) PBCC generates OOD points through a convex combination of ID points from different classes in $\binom{4}{2}$ ways, whereas corruptions depicted in (c) can generate OOD points around each cluster. Observe that points generated by CnC spans wider OOD space including inter-ID-cluster area as well as outside the convex hull of ID points.

$\mathcal{Y}^+ = \{1, \ldots, K+1\}$, and train a classifier f_θ^+ over $\mathcal{Y}^+$. The goal is to train f_θ^+ to implicitly build an estimate $\hat{Z}$ of Z, such that the output of f_θ^+ is $(K+1)$ if $\hat{Z} = 1$, and one of the elements of $\mathcal{Y}$ if $\hat{Z} = 0$.

3.2 Synthetic OOD Data Generation

Our synthetic sample generation strategy consists of following two steps.

Patch Based Convex Combination (PBCC): The first step of our method generates a synthetic sample by convex combination of two input images. Let $x \in \mathbb{R}^{W \times H \times C}$, and y denote a training image and its label respectively. Here, W, H, C denote width, height, channels of the image respectively. A new sample, $\tilde{x}$, is generated by a convex combination of two training samples (x_A, y_A), and (x_B, y_B):

$$\tilde{x} = \mathbf{M} \odot x_A + (\mathbf{1} - \mathbf{M}) \odot x_B. \tag{1}$$

Here, x_A and x_B do not belong to a same class ($y_A \neq y_B$), and $\mathbf{M} \in \{0,1\}^{W \times H}$ denotes a rectangular binary mask that indicates which region to drop, or use from the two images. $\mathbf{1}$ is a binary mask filled with ones, and $\odot$ is element-wise multiplication. To sample $\mathbf{M}$, we first sample the bounding box coordinates $\mathbf{B} = (r_x, r_y, r_w, r_h)$, indicating the top-left coordinates, and width, and height of the box. The region $\mathbf{B}$ in x_A is cut-out and filled in with the patch cropped from $\mathbf{B}$ of x_B. The coordinates of $\mathbf{B}$ is uniformly sampled according to: $r_x \sim \mathrm{U}(0, W), r_w = W\sqrt{1-\lambda}$ and similarly, $r_y \sim \mathrm{U}(0, H), r_h = H\sqrt{1-\lambda}$. Here, $\lambda \in [0, 1]$ denotes the crop area ratio, and is fixed at different values for generating random samples. The cropping mask $\mathbf{M}$ is generated by filling zeros within the bounding box $\mathbf{B}$ and ones outside. We generate the samples by choosing each pair of labels in $\binom{K}{2}$ ways, and then randomly selecting input images corresponding to the chosen labels. This generates OOD samples spread across various inter-class regions in the embedding space. For ablation on range of λ to ensure that a large number of OOD samples are generated outside the ID clusters see supplementary. We label all generated samples as that of the $(K+1)^{\text{th}}$ reject class.

PBCC and CutMix [52]: We note that PBCC, and CutMix[52] both rely on the same basic operation **convex combination of images**, but for two very different objectives. Whereas, CutMix uses the combination step to guide a model to attend on less discriminative parts of objects (e.g. leg as opposed to head of a person), thereby letting the network generalize better on object detection. On the other hand, we use PBCC as a first step for OOD data generation, where the operation generates samples in a large OOD space between a pair of classes in $\binom{K}{2}$ ways.

PBCC Shortcomings: Note that PBCC performs a convex combination of the two ID images belonging to two distinct classes. Hence, unlike adversarial perturbations, it is able to generate sample points far from the ID points in the RGB space. However, still it can generate samples from only within the convex hull of the ID points corresponding to all classes.Thus, as we show in our ablation studies, sample generated using this step alone are insufficient to train a good OOD detector. Below we show how to improve upon the shortcoming of PBCC.

Compounded Corruptions: We aim to address the above shortcomings by using corruptions on top of PBCC generated samples, thus increasing the sample density in inter-class regions as well as generating samples outside the convex hull. We reason that such compounded corruptions increase the spread of the augmented data to a much wider region. Thus, a reasoning based on "per sample" generalisation error bound from [25]:[Fig. 1, Equation 11] could be utilized for our problem. [25] constructs an input-dependent generalization error bound by analysing the subfunction membership of each input, and show that generalisation error bound improves with smoother training sample density (as defined by number of samples in each region). Intuitively, corruptions over PBCC produces a smoother approximation of ID classes with a finer fit at the ID class boundary. A detailed analysis is given in Sec. 3.3. To give an intuitive understanding, Fig. 2 shows visualizations of the generated OOD samples in red using a 4 class toy dataset in two dimensions.

Hendrycks et al. [17] benchmark robustness of a DNN using 15 algorithmically generated image corruptions that mimic natural corruptions. Each corruption severity ranges from 1 to 5 based on the intensity of corruption, where 5 is most severe. The corruptions can be seen as perturbing a sample point in its local neighborhood, while remaining in the support space of the probability distribution of valid images. We apply these corruptions on the samples generated using PBCC step described earlier. Together, PBCC, and corruptions, allow us to generate a synthetic sample far from, and outside the convex hull of ID samples. At the same time, unlike pure random noise images, the process maintains plausibility of the generated samples. Specifically we apply following corruptions: Gaussian noise, Snow, Fog, Contrast, Shot noise/Poisson noise, Elastic transform, JPEG compression, and blur such as Defocus, Motion etc.

Fig. 1 gives a pictorial overview of the overall proposed scheme with a few OOD image samples generated by our approach. CnC formulates the problem as $(K+1)$ class classification which improves the model representation of underlying distribution, and at the same time improves DNN calibration as seen in Sec. 5.2. Please see Suppl. for the precise steps of our algorithm outlined in a pseudo-code format.

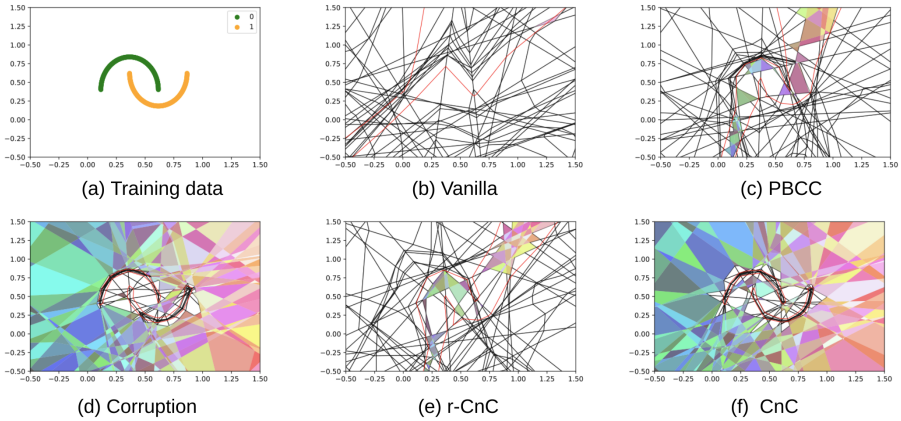

(a) Training data (b) Vanilla (c) PBCC

(d) Corruption (e) r-CnC (f) CnC

Fig. 3: Visualization of trained classifiers as a result of OOD augmentation. A ReLU type DNN is trained on the two-dimensional half-moon data set shown in (a). The shattered neural networks [16] show that CnC has the tightest fit around the ID regions, as measured by the area of the (white colored) polytopes in which no training ID point is observed but a network predicts a point in that region as ID. The measured areas for such polytopes are (b)**Vanilla training without data augmentation**:5.65, (c)**PBCC**: 8.20, (d) **Corruption**:0.40, (e)**r-CnC**: 5.66, (f)**CnC**: 0.37. Note: [25] state that the more densely supported a polytope is by the training set, the more reliable the network is in that region. Hence, the samples declared ID in the regions where no ID sample is observed may actually be OOD with high probability. We observe that PBCC/r-CnC/Vanilla, all predict ID in many such polytopes. Note: r-CnC we reverse the order of PBCC and corruptions, i.e., we first apply corruptions over the input images followed by PBCC over the corrupted images.

3.3 CnC Analysis via Polyhedral Decomposition of Input Space

While we validate the improved performance of CnC in Sec. 5, in this section we seek to provide a plausible explanation for the CnC's performance. We draw inspiration from theoretical support provided in recent work by [25] who formally derive and empirically test prediction unreliability for ReLU based neural networks.

Consider a ReLU network with n inputs and m neurons in total. [25] show that parameters of a trained model partition the input space into a polyhedral complex (PC) consisting of individual *convex* polytopes (also called *activation regions* in [25]). See Fig. 3 for an example with a 2D input space. Each possible input corresponds to a unique state (active or inactive) of each of the m ReLU neurons, and the interior of each polytope corresponds to a unique combination of states of all m neurons. Thus a trained network behaves linearly in the interior of corresponding polytopes. Each edge in the PC corresponds to the state flip of a single neuron (active to inactive, or vice versa).

For the purpose of classification based on the final layer activation, a key corollary from [25] is that *the decision boundary between two classes must be a straight line within a polytope, and can only turn at the vertices*. This is an immediate consequence of the observation that the decision boundary is the locus along which the two highest activations (most probable labels) in the output layer remain equal to each other. This implies that smaller polytopes near the decision boundary are needed for finer control over the boundary between training samples from different classes. Note also that

the authors in [25]:[equation (11)] infer that (paraphrased) "the more a subfunction (polytope) is surrounded by samples with accurate predictions, the lower its empirical error and bound on generalization gap, and thus the lower its expected error bound".

The key question from OOD detection perspective is, how do we force a network to create tighter polytopes at the ID class decision boundaries? We believe the answer is to distribute a large number of the augmented samples (over which we have control) with contrasting OOD and ID labels all around each ID region, forcing the decision boundary to form a tight bounding surface. At the same time, we must also retain a good fraction of the augmented samples in the open space between ID classes, which can be covered by relatively large polytopes (recall that the maximum number of polytopes is bounded by the number of neurons, and thus small polytopes in one region may need to be traded off by larger polytopes in another region). Neglecting the inter-ID space entirely would run the risk of creating very large polytopes in this region, which increases the empirical error bound ([25]:[equation (5) and (11), large subfunctions have low probability mass and hence higher error bound]. CnC lets us achieve this dual objective by using compounding to sample the space between ID classes, and corruption to pepper the immediate neighborhoods around ID classes (especially for λ values near 0 and 1).

In Fig. 3, we show polyhedral complex corresponding to the DNN models trained on two-dimensional half-moon dataset [16,26], and OOD samples generated using various techniques. The first plot shows the input space with training samples from two ID classes (green and yellow semicircles). The learnt polytope structure for vanilla uses a neural network of size $[2, 32, 32, 2]$, while the remaining three plots use $[2, 32, 32, 3]$ (with an additional *reject*/OOD class).

Recall from Fig. 2 that PBCC produces samples sparsely between the ID classes, but not around the ID class boundaries. Pure corruptions produce samples only near and on ID classes, but not in the inter-ID space. On the other hand, CnC produces samples both near the ID boundaries as well as in the inter-ID space.

In Fig. 3, we define any polytope that is fully or partially (decision boundary crosses through it) classified as ID, as an "ID classified polytope" and mark it in white color. **Visually, we can see that the white polytopes occupy a smaller total area when we compare Vanilla to CnC, with the actual values noted in the caption. This indicates that the CnC produces the tightest approximation of ID classes in our example, which in turn leads to better OOD detection.** Though we show for two-dimensional data, we posit that the same generalizes to higher dimensional input data as well, and is the reason for success of CnC based OOD detection.

CnC and Robustness to Adversarial Attacks: Note that, small polytopes in the input space partitioned by a DNN may also provide better safety against black box adversarial attacks as suggested by [16,26]. This is because the black box adversarial attacks extrapolate the gradients based upon a particular test sample. Since the linearity of the output, and thus the gradients is only valid inside a polytope, smaller polytopes near the ID or in the OOD region makes it difficult for an adversary to extrapolate an output to a large region. However, since adversarial robustness is not the focus of this paper, we do not further explore this direction.

3.4 Training Procedure

We train a $(K+1)$ class classifier network f_θ^+, where first K classes correspond to the multi-classification ID classes, and the $(K+1)^{th}$ class label indicates the OOD class. Our training objective takes the form:

$$\begin{aligned}\mathcal{L} = \underset{\theta}{\text{minimize}} \quad & \mathbb{E}_{(x,y)\sim D_{\text{in}}^{\text{train}}}[\mathcal{L}_{\text{CE}}(x,y;f_\theta^+(x))] \\ & +\alpha\cdot\mathbb{E}_{(x,y)\sim D_{\text{pbcc}}^{\text{corr}}}[\mathcal{L}_{\text{CE}}(x,K+1;f_\theta^+(x))],\end{aligned} \tag{2}$$

where $\mathcal{L}_{\text{CE}}$ is the cross entropy loss, $f_\theta^+(x)$ denotes the softmax output of neural network for an input sample x. We use $\alpha = 1$ in our experiments based on the ablation study reported in the supplementary material.

3.5 Inference

After training, we obtain a trained model F^+. We use $F^+(x)[K+1]$ as the OOD score of x during testing, and define an OOD detector $D(x)$ as:

$$D(x) = \begin{cases} 0, & \text{if } F^+(x)[K+1] > \delta \\ 1, & \text{if } F^+(x)[K+1] \leq \delta \end{cases} \tag{3}$$

where, $D(x) = 0$ indicates an OOD prediction, and $D(x) = 1$ implies an ID sample prediction. δ is a threshold learnt from the ID train set such that TPR, i.e., fraction of ID images correctly classified as ID is 95%. For images which are characterized as ID by $D(x)$, the labels are given as:.

$$\hat{y} = \underset{i\in 1,\ldots,K}{\arg\max}\, F^+(x)_i \tag{4}$$

4 Dataset and Evaluation Methodology

In-Distribution Datasets: For ID samples, we use SVHN (10 classes) [41], CIFAR-10 (10 classes), CIFAR-100 (100 classes)[28] containing images of size 32×32. We also use TinyImageNet (200 classes) [30] containing images of resolution 64×64 images.

Out-of-Distribution Datasets: For comparison, we use the following OOD datasets: TinyImageNet-crop (TINc), TinyImageNet-resize (TINr), LSUN-crop (LSUNc), LSUN-resize (LSUNr), iSUN, SVHN. Various works [43,39,4] have shown that the generative models trained using CIFAR-10 /100 have a higher tendency to confuse SVHN dataset as in-distribution, indicating SVHN as a hard case for OOD detection. Hence, for a model trained on SVHN (ID), we use CIFAR-10 as the OOD dataset.

Evaluation Metrics: We compare the performance of various approaches using TNR@TPR95, AUROC and Detection Error. See Supplementary for description on evaluation metrics.

5 Experiments and Results

To show that our data augmentation is effective across different feature extractors, we train using both DenseNet-BC [23] and ResNet-34 [15]. DenseNet has 100 layers with a growth rate of 12. WideResNet [53] models have the same training configuration as in [36]. We provide detailed training procedure in Supplementary.

$\mathcal{D}_{\text{in}}^{\text{train}}$	Method	TNR@TPR95 ↑	AUROC ↑	DetErr ↓	ID Acc. ↑
CIFAR-10 DenseNet-BC	MSP (ICLR'17) [18]	56.1	93.5	12.3	95.3
	ODIN (ICLR'18)[34]	92.4	98.4	5.8	95.3
	Maha(NeurIPS'18)[32]	83.9	93.5	10.2	95.3
	Gen-ODIN (CVPR'20)[22]	94.0	98.8	5.4	94.1
	Gram Matrices(ICML'20)[44]	96.4	99.3	3.6	95.3
	ATOM(ECML'21) [2]	98.3	99.2	1.2	94.5
	CnC(Proposed)	**98.4** $\pm$ **0.8**	**99.5** $\pm$ **1.2**	**2.7** $\pm$ **0.2**	94.7
CIFAR-100 DenseNet-BC	MSP (ICLR'17) [18]	21.7	75.2	31.4	77.8
	ODIN (ICLR'18)[34]	61.7	90.6	16.7	77.8
	Gen-ODIN (CVPR'20)[22]	86.5	97.4	8.0	74.6
	Maha (NeurIPS'18)[32]	68.3	92.8	13.4	77.8
	Gram Matrices(ICML'20)[44]	88.8	97.3	7.3	77.8
	ATOM(ECML'21)[2]	67.7	93	5.6	75.9
	CnC(Proposed)	**97.1** $\pm$ **1.4**	**98.5** $\pm$ **0.4**	**4.6** $\pm$ **0.6**	76.8
TIN RN50	MSP (ICLR'17) [18]	53.15	85.3	22.1	57.0
	ODIN (ICLR'18)[34]	68.5	93.7	12.3	57.0
	CnC(Proposed)	**97.8** $\pm$ **0.8**	**99.6** $\pm$ **0.2**	**2.1** $\pm$ **0.2**	60.5
C-10 WRN	OE (ICLR'19) [19]	93.23	98.64	5.32	94.8
	EBO (NeurIPS'20)[36]	**96.7**	99.0	**3.83**	95.2
	CnC(Proposed)	96.2 $\pm$ 1.5	**99.02** $\pm$ **0.1**	4.5 $\pm$ 0.8	94.3
C-100 WRN	OE (ICLR'19) [19]	47.35	86.02	21.24	75.6
	EBO (NeurIPS'20)[36]	54.0	86.65	19.7	75.7
	CnC(Proposed)	**97.6** $\pm$ **0.9**	**99.5** $\pm$ **0.1**	**2.2** $\pm$ **0.3**	75.1

Table 1: Comparison of competing OOD detectors. TIN stands for TinyImageNet, and RN50 for ResNet50, WRN stands for WideResNet-40-2 The values are averaged over all OOD benchmark datasets. We give individual dataset-wise results in the supplementary. We would like to emphasize that ATOM[2], and OE [19] require large image datasets like 80-Million Tiny Images [46] as representative of OOD samples. However, our method (CnC) generates its own OOD dataset using the ID training data. **Note**: OE[19] and EBO[36] numbers are as reported by the respective papers using their WideResNet trained models. CnC performance is averaged on 3 evaluation runs. **See suppl. for additional results.**

5.1 Comparison with State-of-the-art

OOD Detection Performance: Tab. 1 shows comparison of CnC with recent state-of-the-art. The numbers indicate averaged OOD detection performance on 6 datasets (TinyImagenet, TinyImageNet-crop (TINc), TinyImageNet-resize (TINr), LSUN-crop (LSUNc), LSUN-resize (LSUNr), iSUN, SVHN). As noted in Sec. 4, for SVHN as ID we use CIFAR-10 as the OOD dataset, along with 5 other datasets listed above (more details included in the supplementary). We would like to emphasize that CnC does not need any validation OOD data for fine-tuning. But ODIN [34] and Mahalanobis [32] require OOD data for fine-tuning the hyper-parameters; the hyper-parameters for ODIN and Mahalanobis methods [34,32] are set by validating on 1K images randomly sampled from the test set $\mathcal{D}_{\text{in}}^{\text{test}}$. Tab. 1 clearly shows that CnC outperforms the existing methods. In supplementary material, (a) we show the summary performance of various approaches along with their key conceptual differences; (b) ROC plot of CnC vs. other OOD detectors in addition to reporting AUROC in Tab. 1.

Comparison with Other Data Generation Methods: Tab. 2 shows how CnC fairs against recent OOD data generation methods. In each case we train a $(K+1)$ way classier where first K classes correspond to ID and $(K+1)^{\text{th}}$ class comprised of OOD data generated by corresponding method. As seen from the table, CnC outperforms the recent data augmentation schemes. First, we consider *SaliencyMix* [47] which chooses a representative image patch with the help of a saliency map. Other methods include

Table 2: Comparison with other synthetic data generation methods. We consider CIFAR10 as ID. The values are averaged over all OOD benchmarks. We have used DenseNet[23] as the architecture for all methods trained for $(K+1)$ class classification. Samples obtained through the listed data augmentation schemes were assumed to be of $(K+1)^{\text{th}}$ class. Observe that CnC has superior OOD detection performance. We report average and standard deviation of CnC trained models computed over 3 runs.

Data Augmentation Methods	TNR (95% TPR) ↑	AUROC ↑	Detection Err ↓
Mixup (ICLR'18) [55]	60.6	90.9	15.5
CutOut (arXiV'17) [7]	80.8	94.8	10
CutMix (ICCV'19) [52]	83.2	92.7	8.6
GridMask (arXiV'20) [3]	50.3	79.1	23.6
SaliencyMix (ICLR'21) [47]	85.3	95.7	8.0
AugMix (ICLR'20) [20]	81.3	94.6	11.2
RandomErase (AAAI'20) [56]	41.9	68.1	24.2
Corruptions (ICLR'19) [17]	98.0	99.4	2.8
PuzzleMix (ICML'20) [27]	66.8	84.1	15.2
RandAugment (NeurIPS'20) [5]	89.5	97.9	4.7
Fmix (ICLR'21) [14]	73	90.3	12.6
Standard Gaussian Noise	71.5	93.2	11.7
CnC(Proposed)	**98.4** $\pm$ **0.8**	**99.5** $\pm$ **1.2**	**2.7** $\pm$ **0.2**

Mixup[55], that generates new data through convex combination of training samples and labels. *CutOut* and *GridMask* work on the principle of random and structured information dropping [3]. *RandomErase* [56] selects a rectangular region in an image and replaces its pixels with random values. *AugMix* [20] generates OOD samples through introduction of data drift preserving local image statistics and successive compositions of corruptions. *RandAugment* overcomes the problem of large-scale adoption of these methods but requires a separate and expensive search phase. Other methods who have not open-sourced their implementation like [48] could not be compared. We also could not use the numbers reported in their paper as they report OOD performance using GANs on MNIST/Fashion MNIST datasets and give no results on natural OOD benchmarks: CIFAR-100/Imagenet-R/O/C/Tinyimagenet etc.

5.2 Other Benefits of CnC

Detecting Domain Shift as OOD: We analyze if a model trained with CnC augmented data can detect non-semantic domain shift, i.e. images with the same label but different distribution. For the experiments we use a model trained using CIFAR-100 as ID, and ImageNet-O/ImageNet-R/Corrupted-ImageNet [21] as the OOD. While testing, we downsample the images from ImageNet-O, ImageNet-R and TinyImageNet-C to a size of 32×32. Tab. 3 shows results on ImageNet-R OOD dataset. We outperform the next best technique by 14% on TNR@0.95TPR, 2.9% in AUROC, 3.1% in detection error. See supplementary for results on ImageNet-O and Corrupted ImageNet.

Method	TNR@0.95TPR	AUROC	DetErr
MSP (ICLR'17) [18]	24.4	80.1	26.5
ODIN (ICLR"18) [34]	46.0	88.6	18.9
Gen-ODIN (CVPR'20) [22]	45.0	88.7	18.8
Mahalanobis (NeurIPS'18) [32]	14.0	56.2	41.6
Gram Matrices (ICML'20) [44]	35.0	81.5	25.8
CnC (Proposed)	**60.0**	**91.6**	**15.7**

Table 3: Detecting domain shift using our technique. A model trained with CnC data on CIFAR-100 as the ID using DenseNet-BC [23] feature extractor can successfully detect the domain shift when observing ImageNet-R at the test time. **See Suppl. for additional datasets.**

Model Calibration: Another benefit of training with CnC is model calibration on ID data as well. A classifier is said to be calibrated if the confidence probabilities matches the empirical frequency of correctness [13], hence a crucial to measure of trust in classification models. To measure the calibration error, we compute RMS Calibration error (RMS CE) and Expected Calibration Error (ECE) as suggested in [38].

$\mathcal{D}_{\text{in}}^{\text{train}}$	RMS CE ↓		ECE ↓	
	TS	CnC	TS	CnC
CIFAR-10	11.1	**6.2**	5.1	**3.2**
CIFAR-100	26.1	**19.3**	19.6	**15.4**

Table on the right shows the calibration error for a model trained on CIFAR-10, and CIFAR-100 as the ID data, with CnC samples as the $(K+1)^{\text{th}}$ class. Note that the calibration error is measured only for the ID test samples. We compare the error for a similar model, trained using only ID train data, and calibrated using temperature scaling (TS) [13].

Time Efficiency: For applications demanding real-time performance, it is crucial to have low latency in systems using DNN for inference. To analyze the comparative efficiency, we run the experiments on a desktop system with Intel Core i7 (9th Gen) processor, 16GB RAM, and NVidia RTX 2060 GPU . We report inference time in *ms* averaged over $10K$ images at test time. CnC took the least time when compared to other SOTA techniques. We report that both MSP [18] and CnC model took 0.7ms whereas ODIN [34] and Generalised-ODIN [22] took 3.1ms and 3.3ms respectively. Inference time of Mahalanobis [32] was 8.8ms while GRAM matrices[44] took 10ms. We use CIFAR10 trained models with [23] as the feature extractor for this experiment.

5.3 Ablation Studies

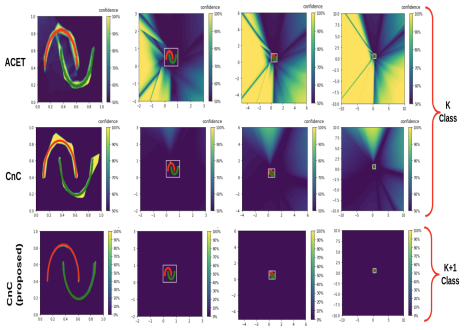

Fig. 4: Confidence plot of a neural network trained using CnC data on data points far away from ID data. **Yellow** color shows high confidence, and **Purple** shows low confidence. **Row 1:** is output of ACET (K class classifier), **Row 2:** is for the same K-class classifier, but using CnC data during training, whereas **Row 3:** is for CnC trained $(K+1)$ class classifier. We observe that the confidence plots for CnC are in general better, and those for $(K+1)$ class classifier are better than K-class one.

Rationale for Design choice of K vs. (K+1) Classifier: We empirically verify having a separate class helps in better optimization/learning during training a model using CnC augmentation. Fig. 5 shows the advantages of using a $(K+1)$ way classifier as compared to standard K class training with better ID-OOD separation. ACET [16] report that despite adversarial training confidence of a ReLU based network far away from the ID points remains high. For a given input $\mathbf{x}$, we obtain $\mathbf{y}$, where $\mathbf{y}$ is the Softmax output from a trained NN model $f_\theta(\mathbf{x})$. We define confidence measure as $\max(\mathbf{y}[:2])$ (max of ID softmax values). First row of the Fig. 4 shows the confidence plot for the half moon dataset, with the training data for positive(Red) and negative classes (Green). Yellow color shows high confidence, and purple show the low confidence. Note that the scale of the inputs changes from left to right. As one can see for ACET the classifier indeed outputs low confidence near the ID region but as the scale increases (far from the ID points), we see more and more high confidence regions, showing the fragility of the OOD detection based on the confidence scores. In the second row we show the K class classifier trained with CnC generated data. As we can that the confidence is much lower even in the far off points. We see that the classifier confidence further improves (low in far away regions) with $(K+1)$ class classifier and CnC based training. Note that for $(K+1)$ class classification, though the classification is K-way, the confidence depicted

is still $\max(\mathbf{y}[:K])$. The Fig. 4 serves to illustrate (1) the benefit of CnC for K class classification, and (2) merit of using $(K+1)$ over a K class classifier.

Comparison with other data synthesis configurations: The table below compares performance of models trained using CnC data against the ones trained using only PBCC, only Corruptions, and r-CnC generated data. Note that in r-CnC we reverse the order of PBCC and corruptions operations. In supplementary, we also compare the UMAP visualization of OOD samples generated using CnC and r-CnC. Intuitively, the embedding vectors after applying only corruptions remain close to the clusters of ID samples. Hence applying PBCC on top of the corrupted point simply generates OOD samples in the line connecting the two points that does not span a wider area outside ID clusters, and are not much useful for training an OOD detector. In contrast, CnC synthesized OOD data exhibits higher diversity in the embedding space making CnC desirable (See Tab. 4).

$\mathcal{D}_{\text{out}}^{\text{test}}$	Method	TNR@95TPR ↑	AUROC ↑	DetErr ↓
TINc	PBCC	93.7	98.7	5.4
	Corruption	96.7	99.2	4.0
	r-CnC	94.8	98.7	4.5
	CnC	**99.7**	**99.9**	**1.3**
LSUNc	PBCC	67.2	90.3	18.6
	Corruption	96.7	99.2	4.1
	r-CnC	95.4	98.7	5.9
	CnC	**98.3**	**99.6**	**3.1**

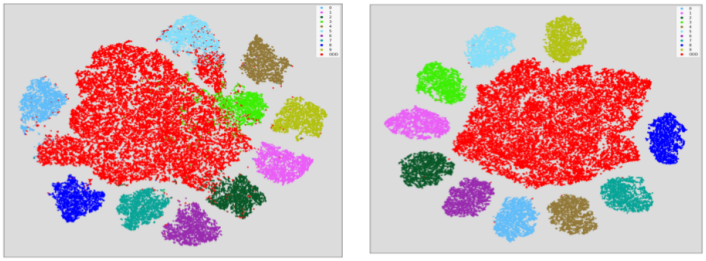

Fig. 5: We show sample t-SNE plots for K Vs. $(K+1)$ classifiers, where CIFAR-10 is used as ID and SVHN is used as OOD (marked in red). The K-class classifier uses temperature scaling [13], where the temperature parameter is tuned on SVHN test set. On the other hand, the $(K+1)$ class classifier uses SVHN data for $(K+1)^{\text{th}}$ class during training. The visualization shows that the OOD data (marked in red) is better separated in a $(K+1)$-class classifier as compared to a K-class classifier

Recommendation for a Good OOD detector: We performed detailed comparison of various configurations of our technique to understand the quantitative scores which can predict the quality of an OOD detector. For the experiment we keep the input images used same across configs, PBCC and corruptions applied are also fixed to remove any kind of randomness.

We use ResNet34 as feature extractor for all methods. CIFAR-10 is used as ID dataset and TinyImageNet-crop as OOD dataset. We observe that the quality of OOD detection improves as the diversity, and entropy of the synthesized data increases (c.f. Tab. 4). Here, entropy is com-

Method	TNR@ 0.95TPR ↑	AUROC ↑	DetErr ↓	Mean Diversity ↑	Mean Entropy ↑
PBCC	93.7	98.6	6.2	2.30	0.33
Corruptions	95.5	97.4	3.5	2.68	0.38
CnC	**98.3**	**99.6**	**2.6**	**3.40**	**0.80**

Table 4: Using entropy/diversity of synthesized data to predict quality of OOD detection. Please refer to text for more details.

puted as the average entropy of the predicted probability vectors by the K class model for the synthesized data. We adapt data diversity from Zhang et al.[54] to measure diversity of OOD data. Refer supplementary for Algorithm for diversity computation.

Limitations of CnC data augmentation: Introduction of additional synthetic data indeed increases training time. For eg., training a model with CnC data on TinyImageNet dataset takes 10 mins. 23 secs./epoch, whereas without CnC data it takes 5 mins 30 secs./epoch on the same Nvidia V100 GPU. Given the performance gain the overhead of training time can be discounted as there is no increase in inference time. We assume the absence of adversarial intentions in this approach, Our method fails when tested against L_∞-norm bounded perturbed image. In future we intend to look at OOD detection using CnC variants for non-visual domains.

6 Conclusions

We have introduced ***C**ompou**n**ded **C**orruptions* (CnC), a novel data augmentation technique for OOD detection in image classifiers. CnC outperforms all the SOTA OOD detectors on standard benchmark datasets tested upon. The major benefit of CnC over SOTA is absence of OOD exposure requirement for training or validation. We also show additional results for robustness to distributional drift, and calibration for CnC trained models. CnC requires just one inference pass at the test time, and thus has much faster inference time compared to SOTA. Finally, we also recommend high diversity and entropy of the synthesized data as good measures to predict quality of OOD detection using it.

References

1. Bendale, A., Boult, T.E.: Towards open set deep networks. In: Proceedings of the IEEE CVPR. pp. 1563–1572 (2016) ii, iii
2. Chen, J., Li, Y., Wu, X., Liang, Y., Jha, S.: Atom: Robustifying out-of-distribution detection using outlier mining. In Proceedings of European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML PKDD) (2021) ii, iii, x
3. Chen, P., Liu, S., Zhao, H., Jia, J.: Gridmask data augmentation (2020) iv, xi
4. Choi, H., Jang, E., Alemi, A.A.: Waic, but why? generative ensembles for robust anomaly detection (2018) ix
5. Cubuk, E.D., Zoph, B., Shlens, J., Le, Q.: Randaugment: Practical automated data augmentation with a reduced search space. In: Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M.F., Lin, H. (eds.) Advances in NeurIPS (2020) iv, xi
6. Dabouei, A., Soleymani, S., Taherkhani, F., Nasrabadi, N.M.: Supermix: Supervising the mixing data augmentation. In: Proceedings of the IEEE/CVF CVPR. pp. 13794–13803 (2021) iv
7. DeVries, T., Taylor, G.W.: Improved regularization of convolutional neural networks with cutout (2017) iv, xi
8. Du, X., Wang, Z., Cai, M., Li, Y.: VOS: learning what you don't know by virtual outlier synthesis. ICLR (2022) ii, iv
9. Fang, Z., Lu, J., Liu, A., Liu, F., Zhang, G.: Learning bounds for open-set learning. In: International Conference on Machine Learning. pp. 3122–3132. PMLR (2021) ii, iv
10. Gal, Y., Ghahramani, Z.: Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In: ICML (2016) ii, iii
11. Ge, Z., Demyanov, S., Chen, Z., Garnavi, R.: Generative openmax for multi-class open set classification. In: BMVC (2017) iii

12. Graves, A.: Practical variational inference for neural networks. In: Shawe-Taylor, J., Zemel, R.S., Bartlett, P.L., Pereira, F., Weinberger, K.Q. (eds.) Advances in Neural Information Processing Systems 24, pp. 2348–2356 (2011) iii
13. Guo, C., Pleiss, G., Sun, Y., Weinberger, K.Q.: On calibration of modern neural networks. In: International Conference on Machine Learning. pp. 1321–1330. PMLR (2017) ii, xi, xii, xiii
14. Harris, E., Marcu, A., Painter, M., Niranjan, M., Prügel-Bennett, A., Hare, J.: Fmix: Enhancing mixed sample data augmentation (2021) xi
15. He, K., Zhang, X., Ren, S., Sun, J.: Identity mappings in deep residual networks (2016) ix
16. Hein, M., Andriushchenko, M., Bitterwolf, J.: Why relu networks yield high-confidence predictions far away from the training data and how to mitigate the problem. In: Proceedings of the IEEE/CVF CVPR. pp. 41–50 (2019) vii, viii, xii
17. Hendrycks, D., Dietterich, T.: Benchmarking neural network robustness to common corruptions and perturbations. In: ICLR (2018) ii, vi, xi
18. Hendrycks, D., Gimpel, K.: A baseline for detecting misclassified and out-of-distribution examples in neural networks. ICLR (2017) ii, iii, x, xi, xii
19. Hendrycks, D., Mazeika, M., Dietterich, T.: Deep anomaly detection with outlier exposure. International Conference on Learning Representations (ICLR) (2019) ii, x
20. Hendrycks*, D., Mu*, N., Cubuk, E.D., Zoph, B., Gilmer, J., Lakshminarayanan, B.: Augmix: A simple method to improve robustness and uncertainty under data shift. In: International Conference on Learning Representations (2020) iv, xi
21. Hendrycks, D., Zhao, K., Basart, S., Steinhardt, J., Song, D.: Natural adversarial examples. In: Proceedings of the IEEE/CVF CVPR. pp. 15262–15271 (2021) xi
22. Hsu, Y.C., Shen, Y., Jin, H., Kira, Z.: Generalized odin: Detecting out-of-distribution image without learning from out-of-distribution data. In: Proceedings of the IEEE/CVF CVPR. pp. 10951–10960 (2020) ii, iii, x, xi, xii
23. Huang, G., Liu, Z., Van Der Maaten, L., Weinberger, K.Q.: Densely connected convolutional networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4700–4708 (2017) iii, ix, xi, xii
24. Huang, R., Li, Y.: Mos: Towards scaling out-of-distribution detection for large semantic space. In: Proceedings of the IEEE/CVF CVPR. pp. 8710–8719 (2021) iii
25. Ji, X., Pascanu, R., Hjelm, D., Lakshminarayanan, B., Vedaldi, A.: Test sample accuracy scales with training sample density in neural networks (2021) ii, iv, vi, vii, viii
26. Jordan, M., Lewis, J., Dimakis, A.G.: Provable certificates for adversarial examples: Fitting a ball in the union of polytopes. 33rd Conference on Neural Information Processing Systems (NeurIPS) (2019) viii
27. Kim, J.H., Choo, W., Song, H.O.: Puzzle mix: Exploiting saliency and local statistics for optimal mixup. In: ICML (2020) iv, xi
28. Krizhevsky, A.: Learning multiple layers of features from tiny images (2009) ix
29. Lakshminarayanan, B., Pritzel, A., Blundell, C.: Simple and scalable predictive uncertainty estimation using deep ensembles. In: Advances in NeurIPS. pp. 6402–6413 (2017) ii, iii
30. Le, Y., Yang, X.: Tiny imagenet visual recognition challenge. CS 231N **7**(7), 3 (2015) ix
31. Lee, K., Lee, H., Lee, K., Shin, J.: Training confidence-calibrated classifiers for detecting out-of-distribution samples. arXiv preprint arXiv:1711.09325 (2017) iv
32. Lee, K., Lee, K., Lee, H., Shin, J.: A simple unified framework for detecting out-of-distribution samples and adversarial attacks. Advances in NeurIPS **31** (2018) ii, iii, x, xi, xii
33. Li, D., Chen, D., Goh, J., Ng, S.k.: Anomaly detection with generative adversarial networks for multivariate time series. ACM KDD iv
34. Liang, S., Li, Y., Srikant, R.: Enhancing the reliability of out-of-distribution image detection in neural networks. In: ICLR (2018) ii, iii, x, xi, xii
35. Lin, Z., Roy, S.D., Li, Y.: Mood: Multi-level out-of-distribution detection. In: Proceedings of the IEEE/CVF CVPR. pp. 15313–15323 (2021) iii

36. Liu, W., Wang, X., Owens, J., Li, Y.: Energy-based out-of-distribution detection. Advances in Neural Information Processing Systems (NeurIPS) (2020) ii, iii, ix, x
37. Mohseni, S., Pitale, M., Yadawa, J., Wang, Z.: Self-supervised learning for generalizable out-of-distribution detection. AAAI pp. 5216–5223 (Apr 2020) ii
38. Naeini, M.P., Cooper, G.F., Hauskrecht, M.: Obtaining well calibrated probabilities using bayesian binning. In: AAAI. p. 2901–2907. AAAI'15 (2015) xi
39. Nalisnick, E., Matsukawa, A., Teh, Y.W., Gorur, D., Lakshminarayanan, B.: Do deep generative models know what they don't know? In: ICLR (2018) ix
40. Neal, L., Olson, M., Fern, X., Wong, W.K., Li, F.: Open set learning with counterfactual images. In: ECCV. pp. 613–628 (2018) ii
41. Netzer, Y., Wang, T., Coates, A., Bissacco, A., Wu, B., Ng, A.Y.: Reading digits in natural images with unsupervised feature learning. In: NeurIPS'21 (2011) iii, ix
42. Perera, P., Nallapati, R., Xiang, B.: Ocgan: One-class novelty detection using gans with constrained latent representations. In: Proceedings of the IEEE/CVF CVPR. pp. 2898–2906 (2019) iv
43. Ren, J., Liu, P.J., Fertig, E., Snoek, J., Poplin, R., Depristo, M., Dillon, J., Lakshminarayanan, B.: Likelihood ratios for out-of-distribution detection. Advances in NeurIPS **32** (2019) ix
44. Sastry, C.S., Oore, S.: Detecting out-of-distribution examples with gram matrices. In: International Conference on Machine Learning. pp. 8491–8501. PMLR (2020) iii, x, xi, xii
45. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 (2014) i
46. Torralba, A., Fergus, R., Freeman, W.T.: 80 million tiny images: A large data set for nonparametric object and scene recognition. IEEE Transactions on PAMI **30**(11), 1958–1970 (2008) x
47. Uddin, A.F.M.S., Monira, M.S., Shin, W., Chung, T., Bae, S.H.: Saliencymix: A saliency guided data augmentation strategy for better regularization. In: ICLR (2021) iv, x, xi
48. Vernekar, S., Gaurav, A., Abdelzad, V., Denouden, T., Salay, R., Czarnecki, K.: Out-of-distribution detection in classifiers via generation. NeurIPS 2019, Safety and Robustness in Decision Making Workshop **abs/1910.04241** (2019) xi
49. Wang, W., Wang, A., Tamar, A., Chen, X., Abbeel, P.: Safer classification by synthesis. arXiv preprint arXiv:1711.08534 (2017) iv
50. Wen, Y., Vicol, P., Ba, J., Tran, D., Grosse, R.: Flipout: Efficient pseudo-independent weight perturbations on mini-batches. In: ICLR (2018) iii
51. Xiao, Z., Yan, Q., Amit, Y.: Likelihood regret: An out-of-distribution detection score for variational auto-encoder. Advances in Neural Information Processing Systems (2020) iv
52. Yun, S., Han, D., Oh, S.J., Chun, S., Choe, J., Yoo, Y.: Cutmix: Regularization strategy to train strong classifiers with localizable features. In: Proceedings of the IEEE/CVF ICCV. pp. 6023–6032 (2019) iv, vi, xi
53. Zagoruyko, S., Komodakis, N.: Wide residual networks. BMVC (2016) ix
54. Zhang, C., Öztireli, C., Mandt, S., Salvi, G.: Active mini-batch sampling using repulsive point processes. In: AAAI. vol. 33, pp. 5741–5748 (2019) xiv
55. Zhang, H., Cisse, M., Dauphin, Y.N., Lopez-Paz, D.: mixup: Beyond empirical risk minimization. In: ICLR (2018) iv, xi
56. Zhong, Z., Zheng, L., Kang, G., Li, S., Yang, Y.: Random erasing data augmentation. In: AAAI. vol. 34, pp. 13001–13008 (2020) iv, xi