
Pure Noise to the Rescue of Insufficient Data: Improving Imbalanced Classification by Training on Random Noise Images

Shiran Zada¹ Itay Benou¹ Michal Irani¹

Abstract

Despite remarkable progress on visual recognition tasks, deep neural-nets still struggle to generalize well when training data is scarce or highly imbalanced, rendering them extremely vulnerable to real-world examples. In this paper, we present a surprisingly simple yet highly effective method to mitigate this limitation: **using pure noise images as additional training data**. Unlike the common use of additive noise or adversarial noise for data augmentation, we propose an entirely different perspective by directly training on pure random noise images. We present a new Distribution-Aware Routing Batch Normalization layer (DAR-BN), which enables training on pure noise images in addition to natural images within the same network. This encourages generalization and suppresses overfitting. Our proposed method significantly improves imbalanced classification performance, obtaining state-of-the-art results on a large variety of long-tailed image classification datasets (CIFAR-10-LT, CIFAR-100-LT, ImageNet-LT, Places-LT, and CelebA-5). Furthermore, our method is extremely simple and easy to use as a general new augmentation tool (on top of existing augmentations), and can be incorporated in any training scheme. It does not require any specialized data generation or training procedures, thus keeping training fast and efficient.

1. Introduction

Large-scale annotated datasets play a vital role in the success of deep neural networks for visual recognition tasks. While popular benchmark datasets are usually well-balanced (e.g.,

¹Department of Computer Science and Applied Mathematics, Weizmann Institute of Science, Rehovot, Israel. Correspondence to: Shiran Zada <shiran.elyahuzada@weizmann.ac.il>, Itay Benou <itay.benou@weizmann.ac.il>, Michal Irani <michal.irani@weizmann.ac.il>.

CIFAR (Krizhevsky et al., 2009), Places (Zhou et al., 2017), ImageNet (Deng et al., 2009)), data in the real world often follows a *long-tail distribution*. Namely, most of the data belongs to several *majority classes*, while the rest is spread across a large number of *minority classes* (Buda et al., 2018; Reed, 2001; Liu et al., 2019). Training on such imbalanced datasets results in models that are biased towards majority classes, demonstrating poor generalization on minority classes. There are two common approaches to compensate for class imbalance during training: (i) *re-weighting* the loss term so that prediction errors on minority samples are given higher penalties (Huang et al., 2016; Cui et al., 2019; Hong et al., 2021; Ren et al., 2020), and (ii) *resampling* the dataset to re-balance the class distribution during training (Chawla et al., 2002; Kim et al., 2020; Mullick et al., 2019). This can be done by under-sampling majority classes (Drummond et al., 2003), or by over-sampling of minority classes (Shen et al., 2016; Haixiang et al., 2017; Kang et al., 2019).

However, *re-weighting* methods typically suffer from overfitting the minority classes (Kim et al., 2020). *Resampling* techniques also suffer from well-known limitations: under-sampling majority classes may impair classification accuracy due to loss of information, while over-sampling leads to overfitting on minority classes (Buda et al., 2018). Several methods have been proposed to alleviate these limitations, including augmentation-based methods (Chu et al., 2020; Mullick et al., 2019; Liu et al., 2020), learning from majority classes (Kim et al., 2020) and evolutionary under-sampling (Galar et al., 2013).

The data scarcity in minority classes thus poses a very challenging problem (Kim et al., 2020). This is especially true in highly imbalanced datasets, where minority classes contain very few samples (e.g., 5 images per class, vs. thousands of images in majority classes). In such cases, overfitting is almost inevitable, even with extensive data augmentation, since the ability to produce a significant variety of new observations from just a few samples is extremely limited.

In this work, we directly address this problem by taking a new perspective on data re-balancing for imbalanced classification. Unlike traditional resampling approaches, we do not restrict ourselves to training strictly on existing images and their augmentations, thus bypassing this limitation.




Figure 1. **Method overview.** (Left) OPeN re-balances an imbalanced dataset with pure-noise images, in addition to oversampled natural images. In OPeN, we replace the standard Batch Normalization layer with DAR-BN. (Right) “Distribution Aware Routing BN” (DAR-BN) handles the distribution gap between natural images and pure-noise images, by normalizing them separately. The affine parameters learned on the natural input only, are used to correctly scale and shift the noise input.




Figure 2. **Re-balanced dataset with OPeN.** To balance the dataset, the original images (blue) are oversampled (green) together with additional pure random noise images (orange). The amount of pure noise added to each class is inversely proportional to its size.

Specifically, we propose *generating pure random noise images and using them as additional training data* (especially for minority classes). We show that training on pure noise images can suppress overfitting and encourage generalization, leading to state-of-the-art results on commonly used imbalanced classification benchmarks (Section 4). We further provide an intuitive explanation as to why this counter-intuitive approach works in practice (Section 3.3). To facilitate learning on pure noise images, which are out-of-distribution of natural images, we present a new batch normalization layer called *Distribution-Aware Routing Batch Normalization* (DAR-BN). Unlike standard Batch-Normalization (BN) (Ioffe & Szegedy, 2015) that assumes that all inputs are drawn from the same distribution, DAR-BN is specifically designed to mitigate the distribution shift between two different input domains (namely, natural images and pure noise images).

We note that many previous works have used noise as a form of data augmentation to improve the accuracy and robustness of deep learning models (Koziarski & Cyganek, 2017; Lopes et al., 2019). These methods, however, show limited improvement when the training data is scarce (Koziarski & Cyganek, 2017). This is due to the fact that applying small doses of additive/multiplicative noise to existing images produces samples in close vicinity to the original ones, thus limiting the data variability. Adding large amounts of noise,

on the other hand, degrades models’ performance due to the large distribution shift from natural images (Appendix B). In contrast, in our method – OPeN (Oversampling with Pure Noise Images), the model is explicitly trained on pure noise images that are far-off the natural images manifold, while explicitly handling the distribution shift, thus promoting generalization.

Our contributions are therefore several fold:

- *State-of-the-art results* on multiple imbalanced classification benchmarks: CIFAR10-LT and CIFAR100-LT (Cao et al., 2019), ImageNet-LT (Liu et al., 2019), Places-LT (Liu et al., 2019), CelebA-5 (Kim et al., 2020).
- To our best knowledge, we are the first to successfully use pure-noise images for training deep image recognition models. We provide extensive empirical evidence to its improved generalization capabilities (and intuition why).
- We introduce a new *distribution-aware batch normalization layer* (DAR-BN), that can bridge the distribution-gap between different input domains of neural networks. While in this work we used DAR-BN to bridge the gap between real and pure-noise images, it can be applied as a general BN layer for handling any pair of different input domains.
- Our method is extremely simple to use as a *general new augmentation tool* (on top of existing augmentations), and can be incorporated in any training scheme. It does not require any specialized data generation or training procedures, thus training is fast and efficient. Our code is available at <https://github.com/shiranzada/pure-noise>.

2. Related Work

Imbalanced Classification:

Data resampling: Most data-based approaches for imbalanced classification aim to re-balance the dataset such that minority and majority classes are equally represented during training. This can be achieved by either over-sampling mi-

nity classes (Chawla et al., 2002; Wang et al., 2014; Shen et al., 2016) or under-sampling majority classes (Drummond et al., 2003; Liu et al., 2008; Galar et al., 2013). More recent works address class re-balancing using GANs (Mullick et al., 2019; Hao et al., 2020) and semi-supervised learning (Wei et al., 2021). An oversampling framework related to our work is *M2m* (Kim et al., 2020), in which majority samples are “transferred” to minor classes using adversarial noise perturbations. Our method also belongs to the data re-balancing category, however, it does so by adding pure random noise images as additional training samples rather than using additive noise augmentations. We also note that our method does not require any optimized data creation procedure or using an auxiliary classifier, which allows a simple and efficient training process ($\times 10$ faster than (Kim et al., 2020), with better results). Our OPeN framework belongs to the category of data resampling approaches, which are therefore most relevant to our work.

Loss re-weighting: aims to compensate for data imbalance by adjusting the loss function, e.g., by assigning minority samples with higher loss weights than majority samples (Li et al., 2021; Buda et al., 2018; Cui et al., 2019; Ren et al., 2018; Park et al., 2021). Recently, *BALMS* (Ren et al., 2020) and *LADe* (Hong et al., 2021) both suggested calibrating the predicted logits according to a prior distribution, by adjusting the softmax function or by adding a regularization term to the loss, respectively.

Margin loss: using a loss function that pushes the decision boundary further away from minority classes samples (Zhang et al., 2017; Dong et al., 2018). For example, (Cao et al., 2019) presented the Label Distribution Aware Margin (*LDAM*) loss, which is combined with deferred re-weighting (*DRW*) training schedule for improved results.

Decoupled training: A recent line of work showing that separating the feature representation learning from the final classification task can be beneficial for imbalanced classification (Kang et al., 2019; Zhou et al., 2020; Wang et al., 2021). E.g., the recently proposed *MiSLAS* method (Zhong et al., 2021) suggested using a shifted batch normalization layer between the two stages of the decoupling framework, in addition to calibrating the final model predictions using a label-aware smoothing scheme.

Noise-Based Augmentation:

Augmenting training data with additive or multiplicative noise has long been in use for training visual recognition models (Holmstrom & Koistinen, 1992; Bengio et al., 2011; Ding et al., 2016). The main motivation behind such augmentation techniques is improving the model robustness to noisy inputs and preventing its fixation on specific input features by randomly “occluding” parts of them (Lopes et al., 2019). While demonstrating some success in reducing overfit (Zur et al., 2009), these methods usually provide limited

improvement to deep models as they tend to overfit to the specific type of noise used during training (Yin et al., 2015).

Another group of methods that uses additive noise are adversarial training techniques, which aim to “fool” a deep model by perturbing images with small, optimized noise (Goodfellow et al., 2014; Kurakin et al., 2016). Particularly relevant is *M2m* (Kim et al., 2020), which suggests using adversarial noise to “transfer” images from major classes to minor classes in an imbalanced classification setting. Similarly, *AdvProp* (Xie et al., 2020) suggests utilizing adversarial examples for improving accuracy and robustness in a general (balanced) classification setting. They try to bridge the distribution gap between two types of inputs (real and adversarial images), for which they use an auxiliary batch normalization layer. In our work, however, the training data is enriched using pure noise images rather than adversarial examples. Additionally, *AdvProp* learns two completely separate sets of batch-norm parameters while in our proposed *DAR-BN* the affine parameters are learned only based on real images, and then applied to both data sources.

Normalization Layers:

Since the introduction of batch-normalization (Ioffe & Szegedy, 2015), various extensions have been proposed to further improve normalization within deep networks, including layer-norm (Ba et al., 2016), instance-norm (Ulyanov et al., 2016) and group-norm (Wu & He, 2018). In common to all these layers is that they normalize activation maps based on a single set of statistics (i.e., mean and variance) for the entire training set. While this may work well when all data samples are from the same underlying distribution, it is sub-optimal when the data is multi-modal or originates from several different domains (Xie et al., 2020). Several recent works which relate to ours have addressed this issue: adaptive instance-normalization (Huang & Belongie, 2017) was introduced for style transfer by adjusting the statistics of content and style inputs. Similarly, (Li et al., 2018; Xie et al., 2020) propose mitigating the domain shift by keeping separate sets of normalization terms for different domains. In our proposed *DAR-BN* layer we also normalize real and noise inputs using their internal statistics. However, unlike (Li et al., 2018; Xie et al., 2020), *DAR-BN* uses the affine parameters learned by the *natural activation maps* in order to scale and shift the *noise activation maps*. Finally, since at test time inputs are sampled only from the natural images domain, *DAR-BN* updates the batch statistics only using activation maps of natural images. We found those difference to be critical to the results.

3. Imbalanced Classification using OPeN

Figures 1 and 2 provide a schematic overview of our approach for imbalanced classification, which is detailed next.

$\mathcal{D} := \bigcup_{i \in [C]} \{(x_j, c_i)\}_{j=1}^{n_i}$ is a long-tailed imbalanced dataset

containing C classes $\{1, 2, \dots, C\}$, where each class c_i consist of n_i training samples. For simplicity we assume $n_1 \geq n_2 \geq \dots \geq n_C$ and $n_1 \gg n_C$. In some cases, the ratio between the largest and the smallest class is a factor of 1000. While the training-set \mathcal{D} is *class-imbalanced*, the test-set is *class-balanced*, and therefore classification of minority classes (with only few samples) is of equal importance to that of majority classes. To compensate for the lack of training data in minority classes, we adopt an oversampling approach that levels the number of samples in each class. However, in contrast to common oversampling techniques (Chawla et al., 2002; Kim et al., 2020), whose training images are solely based on original ones (i.e., their duplications and augmentations), we propose to use also *pure random Gaussian noise images as additional training samples*. As shown in Fig. 2, for each class c_i , we balance the data by adding $n_i^* = n_{max} - n_i$ new training images (where $n_{max} = n_1$ is the largest class), out of which $\delta \cdot n_i^*$ are pure noise images, and $(1 - \delta) \cdot n_i^*$ are real (oversampled) images. During training, we feed the network with mixed batches containing both natural images (with augmentations) and pure noise images. In Sec. 3.1 we further elaborate on this process. Since pure-noise images are out-of-distribution of natural images, we normalize them separately using a new distribution-aware normalization layer. This is explained in Section 3.2. Finally, Sec. 3.3 provides intuition why this improves accuracy of imbalanced classification, especially on minority classes.

3.1. Oversampling with Pure Noise Images (OPeN)

We define the *representation-ratio* of each class c_i in \mathcal{D} as $\rho_i := n_i / n_{max}$. By definition, minority classes have a smaller representation ratio than majority classes. Since standard oversampling results in overfitting of minority classes (Buda et al., 2018), we replace part of the oversampled images with pure random noise images, with the following probability:

$$\mathbb{P}(\text{replace } x \text{ with } x_{noise} | c_i) = (1 - \rho_i) \cdot \delta \quad (1)$$

where c_i is the associated class label of image x , ρ_i is the representation ratio of class c_i , and x_{noise} is randomly sampled from a normal distribution using the mean and variance of all images in the dataset. $\delta \in [0, 1]$ is a hyper-parameter defining the ratio between pure noise images and natural images. Each class in the dataset has a different number of samples, hence is prone to overfitting to a different extent. Equation (1) adjusts the number of noise images added per class accordingly. Lower ρ_i results in a higher probability to replace a sample from class c_i with a pure random noise image, and vice versa for larger ρ_i .

The pure random noise images are generated as follows. Let \mathcal{X} be the set of training images in the dataset \mathcal{D} . We first compute the mean and standard deviation for each color

channel $l \in \{1, 2, 3\}$:

$$\mu_{\mathcal{D},l} = \mathbb{E}(\mathcal{X}[l, :]) \quad , \quad \sigma_{\mathcal{D},l} = \sqrt{\text{Var}(\mathcal{X}[l, :])} \quad (2)$$

Noise images are then sampled from the following normal distribution and clipped to the feasible domain $[0, 1]$:

$$\hat{x}_{noise} \sim \mathcal{N}(\mu_{\mathcal{D}}, \sigma_{\mathcal{D}}) \quad (3)$$

$$x_{noise} = \min(\max(\hat{x}_{noise}, 0), 1) \quad (4)$$

At every epoch, we randomly sample new noise images, as this helps the network to avoid overfitting to specific noise images. A pseudo-code for OPeN is shown in Algorithm 1. Full PyTorch code is provided in Appendix E.

Algorithm 1 Oversampling with Pure Random Noise (OPeN)

Input: (i) Imbalanced dataset: $\mathcal{D} = \bigcup_{i \in [C]} \{(x_j, c_i)\}_{j=1}^{n_i}$
 (ii) noise ratio: $\delta \in [0, 1]$; (iii) dataset statistics: $\mu_{\mathcal{D}}, \sigma_{\mathcal{D}}$

Initialize $\rho = \{\rho_1, \dots, \rho_C\} := \{n_i / n_{max}\}_{i=1}^C$
 $L_D \leftarrow$ Balanced loader for \mathcal{D} using oversampling
 $B = \{(x_j, y_j)\}_{j=1}^{batch_size} \leftarrow$ Sample a batch from L_D
for all $(x_j, y_j) \in B$ **do**
 \triangleright Compute probability ξ of replacing x_j with noise
 $\xi \leftarrow (1 - \rho[y_j]) \cdot \delta$
 $r \sim \text{Bernoulli}(\xi)$
 if $r == 1$ **then**
 $x_{noise} \sim \mathcal{N}(\mu_{\mathcal{D}}, \sigma_{\mathcal{D}})$
 $x_j \leftarrow \min(\max(x_{noise}, 0), 1)$
 end if
end for
return B

3.2. Distribution-Aware Routing Batch Norm (DAR-BN)

Standard Batch Normalization (Ioffe & Szegedy, 2015) is designed to handle the change in the distribution of inputs to layers in deep neural networks, also known as internal covariate shift. However, it assumes that all input samples are taken from the same or similar distributions. Therefore, when inputs originate from several different distributions, BN fails to properly normalize them (He et al., 2019; Xie et al., 2020; Xie & Yuille, 2019). In our framework, OPeN uses pure random noise as additional training examples, which are clearly out of the distribution of natural images. As a result, the layer’s input consists of activation maps obtained both from natural images (where the similar distribution assumption holds) as well as pure noise images (whose distribution is very different from that of the natural images in the train or test datasets). We experimentally observe that using noise images with the standard BN layer, leads to a significant degradation in classification results (see Section 5), even below the baseline of not using noise at all. This may further suggest why this simple idea (of adding pure noise images as additional training examples) has not been previously proposed as a general tool to improve generalization

of deep neural networks. To handle the significant distribution gap between random noise images and natural images, we introduce a new normalization layer called ‘‘DAR-BN’’: *Distribution-Aware Routing Batch Normalization*.

We start by revisiting the standard Batch Normalization (BN) (Ioffe & Szegedy, 2015), and then explain how we extend it to our proposed DAR-BN. Let $X \in \mathbb{R}^{N \times d \times H \times W}$ denote an input to the normalization layer, where N is the batch size, $H \times W$ is the spatial dimension size, and d is the number of channels. The BN layer acts on each channel independently by first normalizing the input across the spatial and the batch dimensions, then applying an affine layer with trainable parameters. Formally, for each channel $j \in [d]$:

$$\hat{x}_j = \frac{x_j - \mathbb{E}(x_j)}{\sqrt{\text{Var}(x_j)}} \quad (5)$$

$$\text{out}_j = \gamma_j \cdot \hat{x}_j - \beta_j \quad (6)$$

where $x_j \in \{x_1, \dots, x_d\} \subseteq \mathbb{R}^{N \times H \times W}$ is an input channel (i.e., $x_j = X[:, j, :, :]$) and β_j, γ_j are trainable parameters per channel. At inference time, the input is normalized using the running mean and running variance that were computed during training using an exponential moving average (EMA) of the batch statistics.

$$\bar{\mu} = \eta \cdot \mathbb{E}(X) + (1 - \eta) \cdot \bar{\mu} \quad (7)$$

$$\bar{v} = \eta \cdot \text{Var}(X) + (1 - \eta) \cdot \bar{v} \quad (8)$$

where η is the momentum parameter. Then, at test time the data is normalized by the running mean and variance, i.e., we replace Equation (5) with:

$$\hat{x}_j = \frac{x_j - \bar{\mu}_j}{\sqrt{\bar{v}_j}} \quad (9)$$

To handle the significant distribution shift between random noise images and natural images, we propose DAR-BN, an extension to the standard BN layer. To this goal, DAR-BN normalizes the noise activation maps and the natural activation maps separately. Specifically, assume $X = X_{nat} \cup X_{noise}$ where X_{nat}, X_{noise} are activation maps of natural images and pure noise images in the batch, respectively. DAR-BN replaces Equation (5) with:

$$\hat{x}_{nat,j} = \frac{x_{nat,j} - \mathbb{E}(x_{nat,j})}{\sqrt{\text{Var}(x_{nat,j})}} \quad (10)$$

$$\hat{x}_{noise,j} = \frac{x_{noise,j} - \mathbb{E}(x_{noise,j})}{\sqrt{\text{Var}(x_{noise,j})}} \quad (11)$$

Then, motivated by AdaBN (Li et al., 2016) (which is designed to handle the covariate shift for domain adaptation/transfer learning), DAR-BN uses the affine parameters learned by the *natural activation maps* in order to scale and shift the *noise activation maps*. Specifically, DAR-BN replaces Equation (6) with:

$$\text{out}_{nat,j} = \gamma_{nat,j} \cdot \hat{x}_{nat,j} - \beta_{nat,j} \quad (12)$$

$$\text{out}_{noise,j} = \gamma_{nat,j} \cdot \hat{x}_{noise,j} - \beta_{nat,j} \quad (13)$$

Equation (13) is applied when the parameters $\beta_{nat}, \gamma_{nat}$ remain fixed, such that no update is applied to these parameters in the back-propagation step due to the operation in Equation (13). Finally, since at test time inputs are sampled only from the natural images domain, DAR-BN updates the batch statistics only using activation maps of natural images. Accordingly, equations Equations (7) and (8) are replaced with:

$$\bar{\mu} = \eta \cdot \mathbb{E}(X_{nat}) + (1 - \eta) \cdot \bar{\mu} \quad (14)$$

$$\bar{v} = \eta \cdot \text{Var}(X_{nat}) + (1 - \eta) \cdot \bar{v} \quad (15)$$

Pseudo-code of DAR-BN is found in Appendix A (Algorithm 2) and full PyTorch code in Appendix E. The distribution shift between pure-noise & real activation maps, and its implications on BN, is illustrated in Figure 4 (Appendix A).

3.3. Underlying Intuition

Training directly on random noise images may seem counter-intuitive. However, we claim it provides a unique regularization effect that can significantly improve generalization of minority classes.

Consider the average batch during training in the imbalanced classification setting described above. Each class is represented in the batch according to its relative size in the training set, i.e., its representation ratio, ρ_i . When backpropagating, the total gradient can be decomposed into the sum of C individual components, one per class. When applying conventional oversampling (using duplications and augmentations of original images), gradient components of minority classes will increase in magnitude, since they are now over-represented. However, their direction will remain relatively unchanged since oversampled images are usually similar to original ones, thus limiting generalization for these classes. This relates to another well-known drawback of such oversampling methods which tend to perform poorly when the number of samples in minority classes is very small, since the ability to synthesize new and varied samples for those classes is extremely limited (Kim et al., 2020).

In contrast, using the proposed OPeN resampling scheme alleviates both of these problems: (i) From a training point of view, oversampling with pure noise images also increases the magnitude of minority gradient components, but at the same time adds stochasticity to their direction. On CIFAR-100-LT, the mean gradient magnitude and the direction variance using OPeN are $\times 2$ and $\times 9$ larger (respectively) than using oversampling *without pure-noise images*. This stochasticity has a regularization effect on the training process, whose strength is *inversely proportional* to the class size. This way, overfitting of minority classes can be suppressed, and generalization is encouraged. (ii) By using random noise images, generation of new training samples

Methods	CIFAR-10-LT		CIFAR-100-LT		CelebA-5
	IR=100	IR=50	IR=100	IR=50	IR=10.7
ERM	79.6±0.2	84.9±0.4	47.0±0.5	52.4±0.4	78.6 ±0.1
Oversampling	75.1±0.4	82.2±0.4	42.5±0.3	48.0±0.2	76.4 ±0.2
LADM-DRW [§]	80.5±0.6	85.3±0.2	46.8±0.2	52.6±0.2	78.5±0.5
M2m [§]	81.3±0.4	85.5±0.3	46.5±0.5	52.9±0.2	76.9±0.4
MiSLAS [⊙]	82.1	85.7	47.0	52.3	-
OPeN	84.6±0.2	87.9±0.2	51.5±0.4	56.3±0.4	79.7±0.2
ERM + AA	81.4±0.3	86.4±0.2	49.9±0.4	55.7±0.4	79.3±0.5
BALMS + AA [⊙]	84.9	-	50.8	-	-
OPeN + AA	86.1±0.1	89.2±0.2	54.2±0.5	59.8±0.5	80.9±0.4

Table 1. **Comparisons on CIFAR-10-LT, CIFAR-100-LT, CelebA-5.** Rows with [⊙] denote results as reported in the original papers. Rows with [§] denote results reproduced with the same architecture as in our experiments, for fair comparison (new results are higher than reported in the original papers). AA stands for AutoAugment optimized on CIFAR-10. Missing results indicate datasets not evaluated in the cited papers. AA is not a legal augmentation for CIFAR-10-LT and ImageNet-LT (as it was optimized on their full balanced datasets). Top part of the table is thus without AA. However, since BALMS was trained with AA, we add such a comparison in the bottom part of the table.

is not limited by the variety of existing samples in the data. This way, we bypass the limitation posed by the small number of minority samples, and explicitly teach the network to handle inputs that are significantly out of its training-set distribution. In particular, the network learns to expect much higher variability and uncertainty in the test images of minority classes. Indeed, at test time, as our experiments suggest, this translates into increased generalization performance. We note that many previous works have used noise as a form of data augmentation. These methods, however, show slight improvement when training data is scarce (Kozlowski & Cyganek, 2017), mostly since applying small doses of noise produces new images that are in close vicinity to original ones, thus providing limited data variability.

One can also understand how our method mitigates data imbalance from another perspective. Since noise inputs are completely random and are class independent, they in fact carry no information except for the class labels we assign to them. Consequently, a key effect of using noise images is on the prior class probabilities learnt by the network. Since in the proposed re-sampling scheme more noise images are assigned to minority classes, we hypothesize that the network learns to implicitly encode these prior probabilities and correct its predictions accordingly.

4. Experiments: Imbalanced Classification

We evaluate our method on five benchmark datasets for imbalanced classification: CIFAR-10-LT, CIFAR-100-LT, ImageNet-LT, Places-LT, and CelebA-5. We follow the evaluation protocol used in (Liu et al., 2019; Kim et al., 2020) for imbalanced classification tasks: The model is trained on the *class-imbalanced* training set, but then evaluated on

Methods	ImageNet-LT
ERM	51.1
Oversampling	49.0
BALMS [◇]	52.1
LADE [◇]	53.0
MisLAS [⊙]	52.7
MisLAS [§]	53.7
OPeN	55.1

Table 2. **Comparison on ImageNet-LT using ResNeXt-50.** Rows with [◇] denote results reported in (Hong et al., 2021). The row with [⊙] denotes the result reported in the original paper (with ResNet-50). The row with [§] denotes the result reproduced by us with ResNeXt-50. OPeN outperforms all previous methods, obtaining state-of-the-art results.

Methods	Places-LT
ERM	29.9
Oversampling	38.1
BALMS [⊙]	38.7
LADE [⊙]	38.8
MiSLAS [⊙]	40.4
OPeN	40.5

Table 3. **Comparison on Places-LT using ImageNet pretrained ResNet-152.** Rows with [⊙] denote results reported in the original papers.

a *balanced class distribution* test set. Our results (summarized in Tables 1 to 3) exhibit state-of-the-art performance on all these datasets. Our implementation details can be found in Appendix D.

4.1. Experimental Setup

Imbalanced (Long-Tail) Datasets. The *Imbalance-Ratio* (IR) of a longtail dataset is defined as $IR = n_{max}/n_{min}$, where n_{max} and n_{min} are the number of training images in its largest and smallest class, respectively. The five long-tailed datasets we used are described below, and their information is summarized in Table 5 in Appendix D.

- **CIFAR-10-LT & CIFAR-100-LT** (Cao et al., 2019). The full balanced CIFAR-10/100 datasets (Krizhevsky et al., 2009) consist of 50,000 training images, and 10,000 test images (split uniformly into 10/100 classes, respectively). Their long-tailed versions, CIFAR-10/100-LT (Cao et al., 2019), were constructed by an exponential decay sampling of the number of training images per class, while their corresponding test sets remain unchanged (i.e., uniform class distribution). We evaluate our method on the challenging dataset settings (IR=50, IR=100).
- **ImageNet-LT & Places-LT** (Liu et al., 2019). Both datasets are long-tail subset of the original large-scale (balanced) ImageNet (Deng et al., 2009) and Places (Zhou et al., 2017). The imbalanced datasets were constructed by sampling the original dataset following the Pareto distribution (Reed, 2001) with $\alpha=6$ as its power value. The resulting long-tailed datasets have a smallest class of size 5, and a largest class of size 1280 for ImageNet-LT (IR=256), and 4980 for Places-LT (IR=996).
- **CelebA-5** (Kim et al., 2020). The standard CelebA

dataset (Liu et al., 2015) consists of face images with 40 binary attributes per image. CelebA-5 dataset was proposed in (Mullick et al., 2019) by selecting samples from non-overlapping attributes of hair color (*blonde, black, bald, brown, gray*). Naturally, the resulting dataset is imbalanced (with $IR=10.7$), as human hair colors are not uniformly distributed. The images were then resized to 64×64 pixels. Kim et al. (2020) constructed a smaller version of the imbalanced dataset by sampling each class with a ratio 1:20, preserving the IR .

Baseline methods. On top of comparing to recent leading methods (Cao et al., 2019; Hong et al., 2021; Ren et al., 2020; Zhong et al., 2021; Kim et al., 2020), we also compare ourselves under the same training parameters, augmentations and architectures to the following baselines: (i) Empirical Risk Minimization (ERM): training without any re-balancing scheme); (ii) Oversampling: re-balancing the dataset by oversampling minority classes with augmentations; (iii) ERM + AutoAugment (AA) (Cubuk et al., 2019).

4.2. Results

Tables 1 to 3 show the results for imbalanced image classification. **OPeN obtains SOTA results on all benchmark datasets.** For example, on CIFAR-10-LT and CIFAR-100-LT with an imbalance ratio of 100, OPeN outperforms the previous SOTA method, MiSLAS (Zhong et al., 2021) by **4.5%** on both datasets. On ImageNet-LT, OPeN is higher than previous SOTA method, LADE (Hong et al., 2021), by **2.1%**. On Places-LT, OPeN achieves comparable results (0.1% better) to previous SOTA, MiSLAS (Zhong et al., 2021). Since our method uses a stronger network (WideResNet-28-10), for fair comparison we reproduced the results of (Cao et al., 2019; Kim et al., 2020) with the same architecture. This improved their results compared to the original papers (Table 1), yet are still inferior to ours.

The above results were obtained *without* using AutoAugment (AA) (Cubuk et al., 2019). Using AutoAugment for training extremely small/longtailed subsets of CIFAR and ImageNet is unfair (Azuri & Weinshall, 2021), since AutoAugment was optimized using the *entire large and balanced* CIFAR-10 and ImageNet datasets. However, since BALMS (Ren et al., 2020) report results only with AutoAugment, we evaluated our method also with that setting. OPeN with AA outperforms BALMS by **1.1%** and **3.4%** on CIFAR-10-LT and CIFAR-100-LT, respectively. We further note that OPeN obtains state-of-the-art results even without using AutoAugment.

Generalization of minority classes. Besides improving the *mean* accuracy (reported in Tables 1 to 3), finer exploration reveals that most of this overall improvement stems from a *dramatic improvement in classification accuracy of minority classes*, while preserving the accuracy of majority classes.




Figure 3. **Ablation study: The added value of pure noise w.r.t. various augmentation methods.** Mean accuracy on CIFAR-10-LT with $IR=100$. We compare OPeN to ERM baseline and to deferred oversampling (Cao et al., 2019) (with same training parameters).

Specifically, OPeN improves the accuracy of the 20 smallest classes of CIFAR100-LT (with $IR=100$, where minority classes have 5-12 samples) by 13.9% above baseline ERM training, **from mean accuracy of 11.6% to 25.5%**. OPeN also outperforms the baseline deferred oversampling (Cao et al., 2019) (without noise images) by 4.3% on the same subset of minor classes. On CIFAR10-LT, OPeN improves generalization of the two smallest classes by 6.3% compared to deferred oversampling, and **by 15.6% above ERM training**. These findings provide empirical evidence to our hypothesis that adding pure noise to minority classes (as opposed to only augmenting the existing training images) significantly diminishes the overfitting problem and increases the generalization capabilities. Please see Appendix C for more detailed evaluations.

5. Ablation Studies & Observations

In this section we explore the added-value of training on pure-noise images under various different settings, and the importance of using DAR-BN for batch normalization.

Data augmentation. In this ablation study, we explore the added value of pure noise when using OPeN with different types of data augmentation methods. We evaluate several augmentation techniques with increasing power on CIFAR-10-LT: (i) random horizontal flip, followed by random crop with padding of four pixels; (ii) Cutout (DeVries & Taylor, 2017) (which zeros out a random fixed-size window in the image); (iii) SimCLR (Chen et al., 2020) (which includes in addition to the horizontal flip and crop, also color distortion and Gaussian blur) followed by Cutout; (iv) AutoAugment (Cubuk et al., 2019) (which is optimized on the entire balanced CIFAR-10 and ImageNet datasets, and considered to be a highly-powerful augmentation). Figure 3 shows that OPeN provides a significant improvement over all four augmentation types, even when the optimal augmentation for that dataset (AutoAugment (Cubuk et al., 2019))

Norm Layer	CIFAR-10-LT	CIFAR-100-LT
Standard BN (Ioffe & Szegedy, 2015)	81.45±0.70	49.18±0.54
Auxiliary BN (Xie et al., 2020)	83.38±0.16	50.13±0.06
DAR-BN	84.64±0.16	51.50±0.44

Table 4. **Ablation study: Comparing different Batch-Norm layers.** Mean accuracy on CIFAR-10/100-LT with $IR=100$. Each type of BN is plugged into OPeN (with same training parameters). DAR-BN outperforms the other normalization layers.

is used. This further supports our hypothesis that training on out-of-distribution pure noise images has a significant added value in suppressing overfitting, beyond augmentation of existing training images.

The impact of DAR-BN. Our distribution-aware normalization layer (DAR-BN) is an essential component for the success of our method, since it helps bridge the distribution gap between random pure noise images and natural images. (Xie & Yuille, 2019; Xie et al., 2020) already observed that natural images and adversarial images are drawn from two different domains. They addressed this using an “Auxiliary BN” layer, which separates adversarial examples and clean images into two separate standard BN layers, with two separate learnable sets of affine parameters. In contrast, in DAR-BN we use only one set of trainable parameters, which are learned by activation maps of natural images only, and use them to scale and shift both the natural activation maps and the activation maps of the pure noise. This normalization difference is important, since the test data in our case will contain only natural images and no pure noise images.

Table 4 compares the effect of plugging each of 3 different BN layers into OPeN: (i) Standard BN (Ioffe & Szegedy, 2015), (ii) Auxiliary BN (Xie et al., 2020), and (iii) DAR-BN (ours). Results show that DAR-BN outperforms other BN layers (surpassing standard BN by 3.2% and 2.3% on CIFAR-10/100-LT, respectively).

Pure Noise Images – a General Useful Augmentation. Our method and experiments are primarily focused on imbalanced classification. However, we observed that adding pure noise images is often effective as a *general data enrichment method, which complements existing augmentation methods*, even in standard balanced datasets. To use it as such, we simply add a fixed number of pure noise images to each class (e.g., some pre-defined percentage of the class size), and train the network using DAR-BN as described in Algorithm 2. We note that since our method does not modify existing training images, it can be easily applied in addition to any other augmentation technique.

While we did not perform extensive evaluations of this, we exemplify the potential power of training on pure-noise images (with DAR-BN) as an additional useful data augmentation method, on the two *full (balanced)* CIFAR datasets (Krizhevsky et al., 2009). To examine the power of this “complementary augmentation”, we measure its added

value *on top of* successful and commonly used augmentation techniques: (i) *Baseline augmentation*: using random horizontal flip and random cropping (with crop size of 32 and padding of 4); (ii) *AutoAugment* (Cubuk et al., 2019) using the corresponding dataset policy; (iii) *Our method*: adding pure noise images (normalized with DAR-BN) in addition to AutoAugment.

We perform our experiments on the *full (balanced)* CIFAR-10 and CIFAR-100, using a Wide-ResNet-28-10 architecture (Zagoruyko & Komodakis, 2016). All models were trained for 200 epochs using the Adam optimizer with $\beta_1 = 0.9$ and $\beta_2 = 0.999$, and with a standard cross-entropy loss. Noise images were sampled from a Gaussian distribution with mean and variance of the corresponding training set and with a noise-to-real ratio of 1:4 in each batch. Our proposed method (AutoAugment complemented with pure-noise images) achieves the best classification accuracy on both datasets. Specifically:

- **Improvement over the baseline augmentation:**
+2.38% on CIFAR-10, +6.34% on CIFAR-100.
- **Improvement over AutoAugment:**
+0.9% on CIFAR-10, +1.5% on CIFAR-100.

These results suggest that properly utilizing pure noise images (with our proposed DAR-BN), may serve as an additional useful augmentation method in general, without any elaborated data creation schemes. It has the potential to further improve classification accuracy, even when used on top of highly sophisticated augmentation methods such as AutoAugment (which was optimized for these specific datasets). Extensively verifying this observation on a large variety of datasets, architectures, and augmentation methods, is part of our future work.

6. Conclusion

We present a new framework (OPeN) for imbalanced image classification: re-balance the training set by using pure noise images as additional training samples, along with a special distribution-aware normalization layer (DAR-BN). Our method achieves *SOTA results on a large variety of imbalanced classification benchmarks*. In particular, it significantly improves generalization of tiny classes with very few training images. Our method is extremely simple to use, and can be incorporated in any training scheme. While we developed DAR-BN to bridge the distribution gap between real and pure-noise images, it may potentially serve as a new BN layer for bridging the gap between other pairs of different input domains in neural-nets. Our work may open up new research directions for harnessing noise, as well as other types of out-of-distribution data, both for imbalanced classification, and for data enrichment in general.

Acknowledgement: This project received funding from the D. Dan and Betty Kahn Foundation.

References

- Azuri, I. and Weinshall, D. Generative latent implicit conditional optimization when learning from small sample. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pp. 8584–8591. IEEE, 2021.
- Ba, J. L., Kiros, J. R., and Hinton, G. E. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- Bengio, Y., Bastien, F., Bergeron, A., Boulanger-Lewandowski, N., Breuel, T., Chherawala, Y., Cisse, M., Côté, M., Erhan, D., Eustache, J., et al. Deep learners benefit more from out-of-distribution examples. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pp. 164–172. JMLR Workshop and Conference Proceedings, 2011.
- Buda, M., Maki, A., and Mazurowski, M. A. A systematic study of the class imbalance problem in convolutional neural networks. *Neural Networks*, 106:249–259, 2018.
- Cao, K., Wei, C., Gaidon, A., Arechiga, N., and Ma, T. Learning imbalanced datasets with label-distribution-aware margin loss. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, pp. 1567–1578, 2019.
- Chawla, N. V., Bowyer, K. W., Hall, L. O., and Kegelmeyer, W. P. Smote: synthetic minority over-sampling technique. *Journal of artificial intelligence research*, 16:321–357, 2002.
- Chen, T., Kornblith, S., Norouzi, M., and Hinton, G. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pp. 1597–1607. PMLR, 2020.
- Chu, P., Bian, X., Liu, S., and Ling, H. Feature space augmentation for long-tailed data. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXIX 16*, pp. 694–710. Springer, 2020.
- Cubuk, E. D., Zoph, B., Mane, D., Vasudevan, V., and Le, Q. V. Autoaugment: Learning augmentation strategies from data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 113–123, 2019.
- Cui, Y., Jia, M., Lin, T.-Y., Song, Y., and Belongie, S. Class-balanced loss based on effective number of samples. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 9268–9277, 2019.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pp. 248–255. Ieee, 2009.
- DeVries, T. and Taylor, G. W. Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552*, 2017.
- Ding, J., Chen, B., Liu, H., and Huang, M. Convolutional neural network with data augmentation for sar target recognition. *IEEE Geoscience and remote sensing letters*, 13(3):364–368, 2016.
- Dong, Q., Gong, S., and Zhu, X. Imbalanced deep learning by minority class incremental rectification. *IEEE transactions on pattern analysis and machine intelligence*, 41(6):1367–1381, 2018.
- Drummond, C., Holte, R. C., et al. C4. 5, class imbalance, and cost sensitivity: why under-sampling beats over-sampling. In *Workshop on learning from imbalanced datasets II*, volume 11, pp. 1–8. Citeseer, 2003.
- Galar, M., Fernández, A., Barrenechea, E., and Herrera, F. Eusboost: Enhancing ensembles for highly imbalanced data-sets by evolutionary undersampling. *Pattern recognition*, 46(12):3460–3471, 2013.
- Gidaris, S. and Komodakis, N. Dynamic few-shot visual learning without forgetting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4367–4375, 2018.
- Goodfellow, I. J., Shlens, J., and Szegedy, C. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014.
- Haixiang, G., Yijing, L., Shang, J., Mingyun, G., Yuanyue, H., and Bing, G. Learning from class-imbalanced data: Review of methods and applications. *Expert Systems with Applications*, 73:220–239, 2017.
- Hao, J., Wang, C., Zhang, H., and Yang, G. Annealing genetic gan for minority oversampling. *arXiv preprint arXiv:2008.01967*, 2020.
- He, Z., Xie, L., Chen, X., Zhang, Y., Wang, Y., and Tian, Q. Data augmentation revisited: Rethinking the distribution gap between clean and augmented data. *arXiv preprint arXiv:1909.09148*, 2019.
- Holmstrom, L. and Koistinen, P. Using additive noise in back-propagation training. *IEEE transactions on neural networks*, 3(1):24–38, 1992.
- Hong, Y., Han, S., Choi, K., Seo, S., Kim, B., and Chang, B. Disentangling label distribution for long-tailed visual recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6626–6636, 2021.

- Huang, C., Li, Y., Loy, C. C., and Tang, X. Learning deep representation for imbalanced classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 5375–5384, 2016.
- Huang, X. and Belongie, S. Arbitrary style transfer in real-time with adaptive instance normalization. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 1501–1510, 2017.
- Ioffe, S. and Szegedy, C. Batch normalization: Accelerating deep network training by reducing internal covariate shift, 2015.
- Kang, B., Xie, S., Rohrbach, M., Yan, Z., Gordo, A., Feng, J., and Kalantidis, Y. Decoupling representation and classifier for long-tailed recognition. In *International Conference on Learning Representations*, 2019.
- Kim, J., Jeong, J., and Shin, J. M2m: Imbalanced classification via major-to-minor translation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 13896–13905, 2020.
- Koziarski, M. and Cyganek, B. Image recognition with deep neural networks in presence of noise—dealing with and taking advantage of distortions. *Integrated Computer-Aided Engineering*, 24(4):337–349, 2017.
- Krizhevsky, A., Hinton, G., et al. Learning multiple layers of features from tiny images. 2009.
- Kurakin, A., Goodfellow, I., and Bengio, S. Adversarial machine learning at scale. *arXiv preprint arXiv:1611.01236*, 2016.
- Li, M., Zhang, X., Thrampoulidis, C., Chen, J., and Oymak, S. Autobalance: Optimized loss functions for imbalanced data. *Advances in Neural Information Processing Systems*, 34, 2021.
- Li, Y., Wang, N., Shi, J., Liu, J., and Hou, X. Revisiting batch normalization for practical domain adaptation. *arXiv preprint arXiv:1603.04779*, 2016.
- Li, Y., Wang, N., Shi, J., Hou, X., and Liu, J. Adaptive batch normalization for practical domain adaptation. *Pattern Recognition*, 80:109–117, 2018.
- Liu, J., Sun, Y., Han, C., Dou, Z., and Li, W. Deep representation learning on long-tailed data: A learnable embedding augmentation perspective. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2970–2979, 2020.
- Liu, X.-Y., Wu, J., and Zhou, Z.-H. Exploratory undersampling for class-imbalance learning. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 39(2):539–550, 2008.
- Liu, Z., Luo, P., Wang, X., and Tang, X. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015.
- Liu, Z., Miao, Z., Zhan, X., Wang, J., Gong, B., and Yu, S. X. Large-scale long-tailed recognition in an open world. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2537–2546, 2019.
- Lopes, R. G., Yin, D., Poole, B., Gilmer, J., and Cubuk, E. D. Improving robustness without sacrificing accuracy with patch gaussian augmentation. *arXiv preprint arXiv:1906.02611*, 2019.
- Mullick, S. S., Datta, S., and Das, S. Generative adversarial minority oversampling. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 1695–1704, 2019.
- Park, S., Lim, J., Jeon, Y., and Choi, J. Y. Influence-balanced loss for imbalanced visual classification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 735–744, 2021.
- Reed, W. J. The pareto, zipf and other power laws. *Economics letters*, 74(1):15–19, 2001.
- Ren, J., Yu, C., Sheng, S., Ma, X., Zhao, H., Yi, S., and Li, H. Balanced meta-softmax for long-tailed visual recognition. *arXiv preprint arXiv:2007.10740*, 2020.
- Ren, M., Zeng, W., Yang, B., and Urtasun, R. Learning to reweight examples for robust deep learning. In *International Conference on Machine Learning*, pp. 4334–4343. PMLR, 2018.
- Shen, L., Lin, Z., and Huang, Q. Relay backpropagation for effective learning of deep convolutional neural networks. In *European conference on computer vision*, pp. 467–482. Springer, 2016.
- Ulyanov, D., Vedaldi, A., and Lempitsky, V. Instance normalization: The missing ingredient for fast stylization. *arXiv preprint arXiv:1607.08022*, 2016.
- Wang, K.-J., Makond, B., Chen, K.-H., and Wang, K.-M. A hybrid classifier combining smote with pso to estimate 5-year survivability of breast cancer patients. *Applied Soft Computing*, 20:15–24, 2014.
- Wang, P., Han, K., Wei, X.-S., Zhang, L., and Wang, L. Contrastive learning based hybrid networks for long-tailed image classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 943–952, 2021.
- Wei, C., Sohn, K., Mellina, C., Yuille, A., and Yang, F. Crest: A class-rebalancing self-training framework for

- imbalanced semi-supervised learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10857–10866, 2021.
- Wu, Y. and He, K. Group normalization. In *Proceedings of the European conference on computer vision (ECCV)*, pp. 3–19, 2018.
- Xie, C. and Yuille, A. Intriguing properties of adversarial training at scale. *arXiv preprint arXiv:1906.03787*, 2019.
- Xie, C., Tan, M., Gong, B., Wang, J., Yuille, A. L., and Le, Q. V. Adversarial examples improve image recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 819–828, 2020.
- Xie, S., Girshick, R., Dollár, P., Tu, Z., and He, K. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1492–1500, 2017.
- Yin, S., Liu, C., Zhang, Z., Lin, Y., Wang, D., Tejedor, J., Zheng, T. F., and Li, Y. Noisy training for deep neural networks in speech recognition. *EURASIP Journal on Audio, Speech, and Music Processing*, 2015(1):1–14, 2015.
- Zagoruyko, S. and Komodakis, N. Wide residual networks. In *British Machine Vision Conference 2016*. British Machine Vision Association, 2016.
- Zhang, X., Fang, Z., Wen, Y., Li, Z., and Qiao, Y. Range loss for deep face recognition with long-tailed training data. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 5409–5418, 2017.
- Zhong, Z., Cui, J., Liu, S., and Jia, J. Improving calibration for long-tailed recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 16489–16498, 2021.
- Zhou, B., Lapedriza, A., Khosla, A., Oliva, A., and Torralba, A. Places: A 10 million image database for scene recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017.
- Zhou, B., Cui, Q., Wei, X.-S., and Chen, Z.-M. Bbn: Bilateral-branch network with cumulative learning for long-tailed visual recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9719–9728, 2020.
- Zur, R. M., Jiang, Y., Pesce, L. L., and Drukker, K. Noise injection for training artificial neural networks: A comparison with weight decay and early stopping. *Medical physics*, 36(10):4810–4818, 2009.

A. DAR-BN & Distribution Shift

This section provides more details on Distribution-Aware Routing Batch Normalization (DAR-BN) layer discussed in Section 3.2. A detailed pseudo-code of DAR-BN layer is given in Algorithm 2. In addition, we provide an illustration of the difference between the vanilla Batch Normalization and DAR-BN when the input comes from two different domains (see Figure 4). While pure random noise is clearly out of the distribution of natural images, this distribution shift is kept also in deeper layers of the network (see Figure 5).

Algorithm 2 Distribution-Aware Routing BN (DAR-BN)

Input: (i) Batch of activation maps (per-channel) $\mathcal{X} = \{x_j\}_{j=1}^{batch_size}$ where x_j is the channel activation map of example j ; (ii) Function indicator $\mathbb{1}_{noise}$ satisfies: $\mathbb{1}_{noise}(x_j) = 1 \Leftrightarrow x_j$ is an activation map of a pure noise image.

Initialize $\gamma = 1, \beta = 0, \bar{\mu} = 0, \bar{v} = 0$

splits $\leftarrow \{ 'natural', 'noise' \}$

for all *split* **in** *splits* **do**

is_noise_split $\leftarrow \mathbb{1}_{\{split \neq 'noise'\}}$ (*split*)

 ▷ Split the batch

$\mathcal{X}_{split} \leftarrow \{x \in \mathcal{X} \mid \mathbb{1}_{noise}(x) == is_noise_split\}$

$\mu_{split} \leftarrow \mathbb{E}(\mathcal{X}_{split})$

$v_{split} \leftarrow Var(\mathcal{X}_{split})$

$\sigma_{split} \leftarrow \sqrt{v_{split}}$

$\mathcal{X}_{split} \leftarrow \frac{\mathcal{X}_{split} - \mu_{split}}{\sigma_{split}}$

if *is_noise_split* **then**

 ▷ Do not update β, γ as well as the batch statistics

with no gradient update:

$\mathcal{X}_{split} \leftarrow \gamma \cdot \mathcal{X}_{split} - \beta$

else

$\mathcal{X}_{split} \leftarrow \gamma \cdot \mathcal{X}_{split} - \beta$

 ▷ Update statistics according to the natural split

$\bar{\mu} \leftarrow \eta \cdot \mu_{split} + (1 - \eta) \cdot \bar{\mu}$

$\bar{v} \leftarrow \eta \cdot v_{split} + (1 - \eta) \cdot \bar{v}$

end if

end for




Figure 4. **Illustration of normalization with BN versus DAR-BN.** (Left) Input of mixed natural (blue) and pure noise (orange) activation maps to be normalized. The input is assumed to be distributed as a bimodal normal distribution. (Center) Output of vanilla Batch Normalization layer which estimates only one set of mean and std (see the green line in the left figure) for both natural and noise activation maps. (Right) Output of DAR-BN layer that separately normalizes the natural and the noise activation maps. As required, the output is distributed as one Gaussian with zero mean and unit variance.

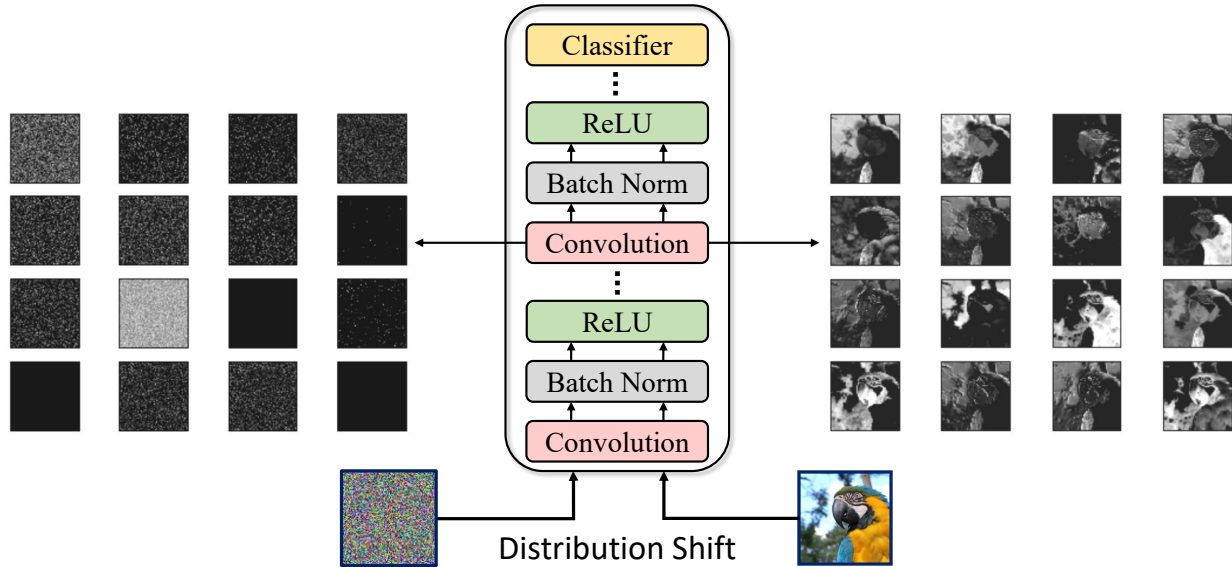


Figure 5. **Distribution shift between natural images and noise images.** Activation maps of a natural image and a pure random noise image. The activation maps are the output of the first convolution block of a pre-trained VGG16 model (we sample 16 out of 64 channels). We see that the features of the natural image and the pure-noise image were sampled from different distributions.

B. Using Additive Noise Instead of Pure-Noise Images

Using OPEN, we sample *pure-noise* image from a normal distribution with fixed mean and standard deviation, determined by the mean and the standard deviation of the natural-image training set (Section 3.1). We experimented with replacing the *pure-noise* images in OPeN with *Additive Gaussian noise* of varying strengths (std) added to real images from the training set. Results on CIFAR-100-LT are plotted in Figure 6, both with our DAR-BN (in blue) and vanilla BN (in orange). Each point on the graph relates to a model trained with a single noise strength (std) for all classes. Results show the increasing the strength of the added noise results in increased accuracy (in blue), until reaching a plateau at the point where the input image practically becomes pure noise ($\sigma \geq 75$), obtaining *the same accuracy as our pure-noise images*. The most significant improvement is achieved for minor classes (blue dashed line), which is also maximized when approaching pure noise levels. This experiment demonstrates that adding training samples that are *increasingly* ‘out-of-distribution’ (i.e., increasingly larger additive noise) allows the model to better handle unseen test images of minor classes. This further supports our use of pure noise images. Note that when using vanilla BN instead of our DAR-BN, accuracy decreases when using higher std since the distribution shift becomes larger.




Figure 6. **Additive Gaussian Noise Instead of Pure-Noise Images.** **Left:** The overall test accuracy (averaged over all classes - solid lines) and accuracy of the 20 smallest classes (dashed lines) of CIFAR100-LT (with IR=100, where minority classes have 5-12 samples). We use the same training scheme used by OPeN while using natural images with additive Gaussian noise instead of pure-noise images. Using DAR-BN results in significantly higher accuracy than vanilla BN. Best results are obtained where the additive noise practically becomes pure-noise (approximately where the standard deviation is higher than 75). **Right:** Example of a natural image and the same image with additive Gaussian noise with std=75.

C. Generalization as a Function of the Class Size

This section provides more details on the improvement provided by OPeN to the generalization of small (minority) classes (extending the evaluation in Section 4.2). To this goal, we perform a finer evaluation over the classes in CIFAR-10-LT and CIFAR-100-LT datasets. We divide the classes (according to their sample size) into five non-overlapping groups of equal size, i.e., each group consists of 20% of the classes. For example, for CIFAR-100-LT, *Group #1* consists of the twenty smallest classes in the training set, while *Group #5* consists of the twenty largest classes. Similarly, for CIFAR-10-LT, each group consists of two classes.

Figure 7 shows the classification results using OPeN compared to the ERM baseline and deferred oversampling (Cao et al., 2019), according to the classes division described above. OPeN provides a significant improvement over two methods on minority classes. We specifically note that **on CIFAR-10-LT, OPeN improves the accuracy over ERM for Group #1 (the two smallest classes) by 15.6%, for Group #2 by 8.6% and for Group #3 by 4.6% while degrades the accuracy for Group #4 & #5 in less the 2%**. This shows that using OPeN, besides improving the overall accuracy (discussed in Section 4.2), results in a more balanced classifier. These results support our claim that OPeN bypasses the limitation of using solely augmented images based on existing ones by employing out-of-distribution random images as additional training examples. Note that pure noise images add stochasticity to the training process. While this significantly improves generalization on severely and mildly overfitted minor classes (i.e., groups 1-3), noisy gradient may interfere with training on (non-overfitted) major classes (e.g., group 5).

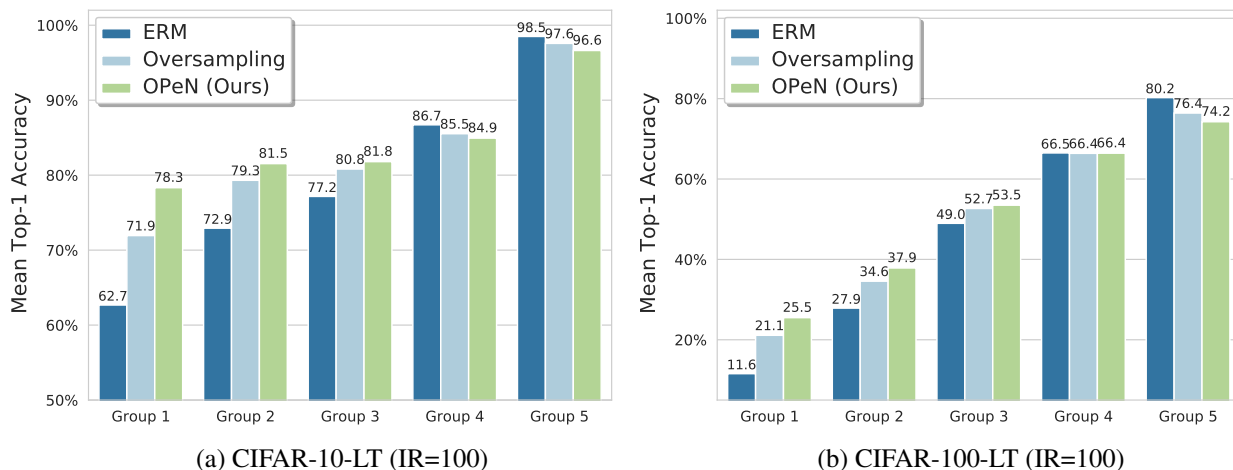


Figure 7. Mean accuracy for each group of classes. We divide the classes into five groups according to their size and report the mean accuracy for each group of classes. Group 1 consists of the 20% most minor classes, while Group 5 consists of the 20% most major classes. We compare OPeN to both the baseline ERM and deferred oversampling (Cao et al., 2019) (with the same training parameters). (a) Results on CIFAR-10-LT (IR=100), which has two classes in each group. (b) Results on CIFAR-100-LT (IR=100), which has twenty classes in each group. OPeN significantly improves the generalization of minority classes on both datasets.

D. Experimental Setup

D.1. Architectures & training

ImageNet-LT. Following (Hong et al., 2021) We use ResNeXt-50-32x4d (Xie et al., 2017) with a cosine classifier (Gidaris & Komodakis, 2018). We train it for 220 epochs using SGD optimizer with momentum 0.9 and weight decay $5e-4$. We use a step learning rate decay with an initial learning rate of $5e-2$, then decay by a factor of 0.1 at epochs 160 and 170. As in (Cao et al., 2019; Kim et al., 2020), we too defer our method (OPeN) to the last phase of the training. This allows the network to learn an initial representation of the data with natural images only. Only when the learning rate decays (which is when the model is exposed to overfitting), do we add the Oversampling + Pure-Noise images. For ImageNet OPeN is deferred to the last 40 epochs.

Places-LT. We follow the procedure of (Hong et al., 2021; Ren et al., 2020; Kang et al., 2019) which use ResNet152 pre-trained on ImageNet as a feature extractor. We use a randomly initialized cosine classifier on top of the backbone, and train the entire network end-to-end for additional 30 epochs using SGD optimizer with momentum 0.9 and weight decay $5e-4$. The initial learning-rate is set to $5e-2$ for the classifier and $1e-3$ for the backbone, then decay the learning rate by a factor of 0.1 at epochs 10 and 15 while OPeN is applied on epoch 15.

Pure Noise to the Rescue of Insufficient Data

Dataset	# of classes	Imbalance-ratio (IR)	Largest class size	Smallest class size	# of samples
CIFAR-10-LT (Cao et al., 2019)	10	{50, 100}	5,000	{100, 50}	{13,996, 12,406}
CIFAR-100-LT (Cao et al., 2019)	100	{50, 100}	500	{10, 5}	{12,608, 10,847}
ImageNet-LT (Liu et al., 2019)	1,000	256	1,200	5	115,846
Places-LT (Liu et al., 2019)	365	996	4,980	5	62,500
CelebA-5 (Kim et al., 2020)	5	10.7	2423	227	6651

Table 5. **Long-tailed datasets.** Summary of the long-tailed datasets we used for evaluation. (see Sec. 4 below for a detailed explanation).

CIFAR-10-LT, CIFAR-100-LT and CelebA-5. All experiment use WideResNet-28-10 (Zagoruyko & Komodakis, 2016). For CIFAR-10/100-LT We follow the setup in (Kim et al., 2020; Cao et al., 2019), that is, train for 200 epochs with Cross-Entropy loss using SGD optimizer with momentum 0.9 and weight decay of $2e-4$. We use a step learning rate decay with an initial learning rate of 0.1, then decay by a factor of 0.01 at epochs 160 and 180. OPeN is deferred to the last 40 epochs. For CelebA-5 (Kim et al., 2020), we follow (Kim et al., 2020) and train for 90 epochs, decay the learning rate by a factor of 0.1 at epochs 30 and 60.

The information of the five long-tailed datasets (ImageNet-LT, Places-LT, CIFAR-10-LT, CIFAR-100-LT, and CelebA-5) is summarized in Table 5.

Noise-ratio (δ): The noise-ratio δ was determined using the CIFAR validation set and set to $\frac{1}{3}$. We simply applied the same value to all other datasets.

Randomization: Since small datasets tend to present high variance, for CelebA-5, CIFAR-10-LT and CIFAR-100-LT, we repeat the experiments 4 times, reporting mean and standard error. For the rest of the datasets, for fair comparison, we use the same randomization seed across all experiments.

D.2. Image Augmentations

In each experiment, the following data augmentations were used:

- For the datasets with small images (i.e., CIFAR-10-LT (Cao et al., 2019), CIFAR-100-LT (Cao et al., 2019) and CelebA-5 (Kim et al., 2020)), we use random horizontal flip followed by random crop with padding of four pixels, then apply Cutout (DeVries & Taylor, 2017) (which zeros out a random 16×16 window in the image) and SimCLR (Chen et al., 2020) (which includes ColorJitter, random Grayscale and random GaussianBlur).
- For the datasets with a higher resolution images (i.e., ImageNet-LT (Liu et al., 2019) and Places-LT (Liu et al., 2019)), we apply random resize crop (with default parameters) to 224×224 pixels followed by SimCLR and random rotation.

When AutoAugment (Cubuk et al., 2019) is employed (see Section 4), it replaces all above augmentations. We use CIFAR-10 policy for CIFAR-10-LT, CIFAR-100-LT and CelebA-5 datasets, and ImageNet policy for ImageNet-LT and Places-LT datasets.

E. PyTorch Code

In this section, we provide code snippets of the core components of our method.

Full code is available at <https://github.com/shiranzada/pure-noise>.

- Listing 1 provides code implementation for the DAR-BN layer as described in the main paper. DAR-BN can be integrated into any neural network by replacing the standard BN layer. Specifically replace

```
x = self.batch_norm(x)
```

with

```
x = dar_bn(self.batch_norm, x, noise_mask)
```

where "noise_mask" is a Boolean array that indicates which activation map is obtained from noise.

- Listing 2 provides code implementation for OPeN together with pure noise images sampling function. The code is given in Python using the PyTorch library.

Pure Noise to the Rescue of Insufficient Data

```
def dar_bn(bn_layer, x, noise_mask):
    """Applies DAR-BN normalization to a 4D input (a mini-batch of 2D inputs with
        additional channel dimension)

    bn_layer : torch.nn.BatchNorm2d
        Batch norm layer operating on activation maps of natural images
    x : torch.FloatTensor of size: (N, C, H, W)
        2D activation maps obtained from both natural images and noise images
    noise_mask: torch.BoolTensor of size: (N)
        Boolean 1D tensor indicates which activation map is obtained from noise
    """
    # Batch norm for activation maps of natural images
    out_natural = bn_layer(x[torch.logical_not(noise_mask)])
    # Batch norm for activation maps of noise images
    # Do not compute gradients for this operation
    with torch.no_grad():
        adaptive_params = {"weight": bn_layer.weight, "bias": bn_layer.bias,
                           "eps": bn_layer.eps}
        out_noise = batch_norm_with_adaptive_parameters(x[noise_mask],
                                                       adaptive_params)

    # Concatenate activation maps in original order
    out = torch.empty_like(torch.cat([out_natural, out_noise], dim=0))
    out[torch.logical_not(noise_mask)] = out_natural
    out[noise_mask] = out_noise

    return out

def batch_norm_with_adaptive_parameters(x_noise, adaptive_parameters):
    """Applies batch normalization to x_noise according to adaptive_parameters

    x_noise : torch.FloatTensor of size: (N, C, H, W)
        2D activation maps obtained from noise images only
    adaptive_parameters:
        a dictionary containing:
        weight: scale parameter for the adaptive affine
        bias: bias parameter for the adaptive affine
        eps: a value added to the denominator for numerical stability.
    """
    # Calculate mean and variance for the noise activations batch per channel
    mean = x_noise.mean([0, 2, 3])
    var = x_noise.var([0, 2, 3], unbiased=False)
    # Normalize the noise activations batch per channel
    out = x_noise - mean[None, :, None, None]
    out = out / torch.sqrt(var[None, :, None, None] + adaptive_parameters["eps"])

    # Scale and shift using adaptive affine per channel
    out = out * adaptive_parameters["weight"][None, :, None, None]
        + adaptive_parameters["bias"][None, :, None, None]

    return out
```

Listing 1: PyTorch code for DAR-BN layer.

Pure Noise to the Rescue of Insufficient Data

```
def oversampling_with_pure_noise_train_epoch(model, balanced_loader, criterion,
                                             optimizer, delta, num_samples_per_class,
                                             dataset_mean, dataset_std, image_size):
    """Trains model for one epoch according to the OPeN scheme

    model : torch.nn.Module;
        Model to train
    balanced_loader: torch.utils.data.DataLoader
        A class balanced loader - samples each class with equal probability
    delta: float
        Hyper-parameter for OPeN (see description in paper)
    num_samples_per_class: torch.IntTensor
        Number of samples in each class in the original imbalanced dataset
    dataset_mean: torch.FloatTensor of size: (3)
        Dataset mean per color channel
    dataset_std: torch.FloatTensor of size: (3)
        Dataset standard deviation per color channel
    image_size: int
        Image size - Assumes squared images of size (image_size , image_size)
    """
    for images, targets in balanced_loader:
        # Compute representation ratio
        max_class_size = torch.max(num_samples_per_class)
        representation_ratio = num_samples_per_class[targets] / max_class_size
        # Compute probabilities to replace natural images with pure noise images
        noise_probs = (1 - representation_ratio) * delta
        # Sample indexes to replace with noise according to Bernoulli distribution
        noise_indices = torch.nonzero(torch.bernoulli(noise_probs)).view(-1)
        # Replace natural images with sampled pure noise images
        noise_images = sample_noise_images(image_size=image_size,
                                          mean=dataset_mean, std=dataset_std, count=len(noise_indices))
        images[noise_indices] = noise_images
        # Create mask for noise images - later used by DAR-BN
        noise_mask = torch.zeros(images.size(0), dtype=torch.bool)
        noise_mask[noise_indices] = True
        # Train model
        outputs = model(images, noise_mask)
        loss = criterion(outputs, targets)
        optimizer.zero_grad()
        loss.backward()
        optimizer.step()

def sample_noise_images(image_size, mean, std, count):
    """Samples pure noise images from the normal distribution N(mean,std)"""
    r = torch.normal(mean[0], std[0], size=(count, 1, image_size, image_size))
    g = torch.normal(mean[1], std[1], size=(count, 1, image_size, image_size))
    b = torch.normal(mean[2], std[2], size=(count, 1, image_size, image_size))
    pure_noise_images = torch.cat((r, g, b), 1)

    return pure_noise_images
```

Listing 2: PyTorch code for OPeN.