

Enhancing Noise-Robust Losses for Large-Scale Noisy Data Learning

Max Staats^{1,2*}, Matthias Thamm², Bernd Rosenow²

¹ScaDS.AI Dresden/Leipzig, Germany

²Institut für Theoretische Physik, Universität Leipzig, Brüderstrasse 16, 04103 Leipzig, Germany
{staats,thamm}@itp.uni-leipzig.de, rosenow@physik.uni-leipzig.de

Abstract

Large annotated datasets inevitably contain noisy labels, which poses a major challenge for training deep neural networks as they easily memorize the labels. Noise-robust loss functions have emerged as a notable strategy to counteract this issue, but it remains challenging to create a robust loss function which is not susceptible to underfitting. Through a quantitative approach, this paper explores the limited overlap between the network output at initialization and regions of non-vanishing gradients of bounded loss functions in the initial learning phase. Using these insights, we address underfitting of several noise robust losses with a novel method denoted as *logit bias*, which adds a real number ϵ to the logit at the position of the correct class. The *logit bias* enables these losses to achieve state-of-the-art results, even on datasets like WebVision, consisting of over a million images from 1000 classes. In addition, we demonstrate that our method can be used to determine optimal parameters for several loss functions – without having to train networks. Remarkably, our method determines the hyperparameters based on the number of classes, resulting in loss functions which require zero dataset or noise-dependent parameters.

1 Introduction

Supervised deep learning depends on high-quality labeled data for effective pattern recognition and accurate predictions (Goodfellow, Bengio, and Courville 2016). In real-world datasets, however, there is often label noise - erroneous or unclear labels due to human error or incomplete annotations (Liang, Liu, and Yao 2022). Such noise can drastically impair the effectiveness of deep learning models, which often operate under the assumption of pristine labels (Song et al. 2022). Therefore, it is important to develop robust deep-learning algorithms that can efficiently learn from noisy datasets.

One effective approach to navigate label noise lies in employing noise-robust loss functions. These loss functions, notable for their model-agnostic nature, seamlessly integrate with any deep learning paradigm. The existing literature highlights their ability to improve the robustness and generalization ability of deep learning models under noisy conditions (Ghosh, Kumar, and Sastry 2017; Zhang and Sabuncu

2018; Wang et al. 2019; Amid et al. 2019; Ma et al. 2020; Zhou et al. 2021; Englesson and Azizpour 2021).

A majority of these loss functions are bounded to prevent the learning of mislabeled examples. From a theoretical point of view, bounded losses have a higher robustness to noise if they belong to the class of symmetric losses (Ghosh, Kumar, and Sastry 2017). Nonetheless, it has been suggested that such symmetry could be overly constraining (Zhou et al. 2021), with functions like the Mean Absolute Error (MAE) leaning towards underfitting. Reflecting this, many contemporary loss functions do not satisfy this symmetry condition (Zhou et al. 2021; Englesson and Azizpour 2021).

In this paper, we quantitatively explore how the vanishing derivatives of bounded loss functions impact their learning behavior. According to our findings, the cause of underfitting is the limited overlap between the output values of an initialized network and the region where the derivative of a particular bounded loss function is nonzero. To counteract this, we add a real number, ϵ , to the logit corresponding to the correct class label. This subtle adjustment restores the overlap between network outputs and the region of sufficiently large derivatives of the loss, enabling e.g. the MAE loss to surpass the Cross Entropy loss on datasets like Cifar-100, even in the absence of label noise. Impressively, this approach requires only a single constant ϵ , which is determined by the number of classes, providing an effectively parameter-free method. Other loss functions like the generalized cross entropy (genCE) (Zhang and Sabuncu 2018) and NF-MAE (a combination of normalized focal loss (NF) with MAE) (Ma et al. 2020) are also able to learn the WebVision dataset with the help of the logit bias.

Furthermore, our description of the early learning phase enables us to calculate suitable hyperparameters for other loss functions like genCE and NCE-AGCE (a combination of Normalized Cross Entropy and Asymmetric Generalized Cross Entropy) (Zhou et al. 2021), allowing them to show their potential for an arbitrary number of classes without the need of tuning their parameters first, e.g., by an expensive grid search. The method we propose is intended as a first step towards a universal framework that is capable of noise robust learning across varied numbers of classes without requiring hyperparameter fine-tuning. The need for such a method is underscored by our observation that none of the proposed loss functions that are noise resistant on the Cifar-

*Staats is the corresponding author.

10 dataset are capable of learning the WebVision dataset.

In summary, our paper (i) quantitatively describes how the initial learning phase of a newly initialized network is contingent upon the dataset’s class count; (ii) explores the limitations of bounded losses in multi-class datasets and introduces the *logit bias* technique, enabling MAE to consistently deliver competitive or even superior results across benchmarks like Fashion-MNIST, Cifar-10, Cifar-100, and Web-Vision – without hyperparameters; (iii) enables other noise-robust loss functions to learn the WebVision dataset using either the logit bias or by determining their hyperparameters without training a single network.

All code for reproducing the data and creating the figures in this paper is open source and available under (Author s).

2 Related Work

Label noise in training data is a pervasive challenge that has attracted much attention in recent years (Liang, Liu, and Yao 2022; Song et al. 2022). One strategy for addressing it is data cleaning, aiming to filter out mislabeled samples from the training dataset. To identify noisy instances, (Xiao et al. 2015) employs a probabilistic model to capture the relationship between images, labels, and noise. Other approaches utilize an auxiliary neural network, trained on curated data, to clean the main dataset (Veit et al. 2017; Lee et al. 2018). Yet, an overzealous curation can sometimes be counterproductive, as eliminating too many samples might degrade model performance (Khetan, Lipton, and Anandkumar 2017), compared to retaining some corrupted instances.

Another approach is estimating the noise transition matrix, which depicts the likelihood of mislabeling across classes. This matrix can be incorporated directly into the loss function (Han et al. 2018) or inferred throughout training (Goldberger and Ben-Reuven 2017; Sukhbaatar and Fergus 2014), mitigating the consequences of mislabeling. A variation of this strategy involves dynamically adjusting the weights of samples during training. For example, (Reed et al. 2014) adjust labels based on the network’s current predictions, and (Ren et al. 2018) evaluate label trustworthiness based on the gradient induced by a given example. Furthermore, (Thulasidasan et al. 2019) prompt the network to predict the likelihood of the example being correct, enabling fine-tuned loss adjustments.

Another avenue entails constraining the minimum training loss, emphasizing that an optimal scenario that avoids learning from incorrect labels will incur a finite loss (Toner and Storkey 2023). Self-supervised methods that iteratively adjust labels by extracting information directly from data structures — via dynamic label learning (Chen et al. 2020) and contrastive learning (Hendrycks et al. 2019; Zheltonozhskii et al. 2022; Ghosh and Lan 2021; Xue, Whitecross, and Mirzasoleiman 2022; Yi et al. 2022) — have also been shown to improve model generalization in noisy conditions.

Opting for a noise-robust loss function can complement and enhance many of the strategies described above. In (Ghosh, Kumar, and Sastry 2017) it has been shown that bounded and symmetric losses are noise-tolerant. The belief that symmetric losses are prone to underfitting (Zhang and Sabuncu 2018; Zhou et al. 2021) led to the development

of alternative loss functions. (Feng et al. 2021) augment the Cross Entropy loss to render it bounded, and (Lyu and Tsang 2019) introduce the *curriculum loss*, acting as a tight upper bound to the 0-1 loss. Meanwhile, (Xu et al. 2019) advocate for a non-bounded loss based on information-theoretic arguments.

Various combinations have emerged to take advantage of the strengths of multiple losses. The generalized Cross Entropy (genCE) (Zhang and Sabuncu 2018) synergizes MAE with the Cross Entropy loss (CE) using a Box-Cox transformation, in an attempt to combine the fast initial learning performance of CE with the noise robustness of MAE. On the other hand, symmetric Cross Entropy (symCE) uses a linear combination of these two losses (Wang et al. 2019). Active-Passive Losses employ linear combinations of MAE with a normalized version of either CE or the Focal loss (Ma et al. 2020). Lastly, the Bi-Tempered Logistic loss (biTemp) *tempers* the exponential and logarithmic functions to create robust loss function similar to the cross entropy (Amid et al. 2019).

3 Theoretical Considerations

Consider a classification task with K classes, an example space defined as X , and associate labels $\{1, 2, \dots, K\}$. We define a loss function, \mathcal{L} , as symmetric if

$$\sum_{y=1}^K \mathcal{L}(f(x), y) = C \quad \forall x \in X \quad \forall f: \mathbb{R}^n \mapsto \mathbb{R}^K, \quad (1)$$

where $C \in \mathbb{R}$ is a fixed constant. Symmetric loss functions have been demonstrated to exhibit robustness against uniform and class-conditional noise (Ghosh, Kumar, and Sastry 2017). The MAE, for instance, serves as a notable example of a symmetric loss. Such losses are often bounded, meaning that constant $C \in \mathbb{R}$ exists such that $\mathcal{L}(\mathbf{a}, \mathbf{y}) < C$, regardless of label \mathbf{y} and network output \mathbf{a} .

A heuristic understanding of this theory can be gained by recognizing that wrongly labeled examples tend to deviate significantly from their class in feature space. In a scenario where these outliers can accumulate unbounded loss, the gradient adjusts the decision boundaries to accommodate even the most extreme outliers. In contrast, a bounded loss ensures that the gradient decreases at greater distances from the decision boundary, preventing overfitting these outliers.

The downside of this approach is a potential slowdown in learning speed, as legitimate examples could be misinterpreted as outliers, with no gradient available to adjust the network weights for accurate classification. The symmetric MAE loss embodies this scenario — it is robust to noise (Ghosh, Kumar, and Sastry 2017), yet underperforms on datasets like Cifar-100 (Zhang and Sabuncu 2018).

This section focuses on the early learning stage of a generic neural network and attempts to determine why certain noise-robust loss functions underfit on multiclass datasets (Zhang and Sabuncu 2018). In addition, we want to identify strategies to address this problem. To explore the underlying causes, we imagine a freshly initialized network with random weights \mathbf{W} drawn from a zero-mean distribu-




Figure 1: Analysis of the average error for the neuron corresponding to the correct label, determined for the MAE loss. Increasing the number of classes from 10 to 100 leads to a notable shift in the range where the average error $\langle \partial_{z_k} \mathcal{L} \rangle$ is non-vanishing. However, the logit distribution of a newly initialized network (shown in the histogram at the bottom) remains class-count invariant. This mismatch leads to diminished gradients, stalling the learning due to tiny errors.

tion. We focus on the logit distribution for each neuron z_i in the final layer, denoted as $p_z(z_i)$.

Owing to inherent symmetry, we obtain $p_z(z_i) = p_z(z_j) \equiv p_z(z)$. The central limit theorem suggests that when the width N of the last hidden layer is infinite, the logit distribution p_z corresponds to a normal distribution with zero mean (Neal 1994; Lee et al. 2017). As N is typically quite large, $p_z(z)$ can indeed be well approximated by a zero-mean normal distribution.

Under conventional training regimes, where the weights are initialized with variances given by the inverse of the layer output size, initial neuron pre-activations have approximately a unit variance¹. Therefore, we adopt a $\mathcal{N}(0, 1)$ distribution for the logits z_i . Though our results primarily consider this distribution, they can be easily generalized for other variances. For an empirical demonstration of the distribution under standard initializations, we refer to the appendix.

When training neural networks using backpropagation, the error δ associated with neuron n can be expressed as

$$\delta_n = \frac{\partial \mathcal{L}}{\partial z_n} = \sum_{j=1}^K \frac{\partial \mathcal{L}}{\partial a_j} a_j (\delta[j-n] - a_n), \quad (2)$$

where K denotes the number of classes and $\delta[j-n]$ the Kronecker delta, $\delta[0] = 1$ and $\delta[m] = 0$ for $m \neq 0$. Moreover, we assume that the network output a_i at position i is defined as $a_i = \exp(z_i) / \sum_j \exp(z_j)$ where z_i is the logit at output position i meaning that we assume the softmax activation for the final layer.

This error plays a central role in backpropagation. Specifically, it determines the magnitude of the gradient descent updates, expressed as $\nabla_{W_i} \mathcal{L} = \delta \cdot \mathbf{a}_{i-1}^T$. An error of zero results in a zero gradient, causing learning to stall.

From our previous discussions, we established that the distribution of z_i corresponds to a normal distribution,

¹It is worth noting that there have been instances where smaller initializations have been recommended (Yang et al. 2021).

$\mathcal{N}(0, 1)$. Consequently, in the initial learning phases the average error of neuron n as a function of the logit z_n of that neuron is given by:

$$\langle \delta_n \rangle(z_n) = \left\langle \sum_{j=1}^K \frac{\partial \mathcal{L}}{\partial a_j} \frac{\exp(z_j)}{\sum_i \exp(z_i)} \times \left(\delta[j-n] - \frac{\exp(z_n)}{\sum_i \exp(z_i)} \right) \right\rangle_{z_i, i \neq n}. \quad (3)$$

To gain insights, we visualize the average error, $\langle \delta_k \rangle$, in the neuron k corresponding to the correct label as a function of its pre-activation. This involves computing the average with respect to $z_{i \neq k}$ and assigning a static value to z_k . This perspective offers a snapshot of typical error values in a newly initialized network. Using the MAE as an example, Fig. 1 shows that with 10 classes, the range over which the average error $\langle \delta_k \rangle$ notably deviates from zero overlaps significantly with the logit distribution of a network at initialization. However, this overlap diminishes with 100 classes as the error depends on the number of classes.

This illustration underscores the challenge: as the number of classes increases, gradients vanish under the MAE loss. This phenomenon has implications for all bounded loss functions, which assign small gradients for outliers. Importantly, what is deemed an “outlier” varies with the class count. We will later illustrate how this makes learning on the WebVision dataset infeasible for certain noise-robust loss functions optimized for up to 100 classes.

Relevant to this discussion, (Ma et al. 2020) categorize the active component of a loss function as the contribution that optimizes the output in the correct neuron. This component has been identified as essential to successful fitting. The optimization within the correct neuron turns out to be the linchpin of the fitting process, especially since the errors in the incorrect neurons decrease with increasing number of classes: already for 100 classes, the initial activation of an “incorrect” neuron in the final layer is very close to the optimal value of zero, with an average value of 0.01, leading to rather small errors in these neurons. This accentuates the need to focus primarily on the error in the “correct” neuron, δ_k .

3.1 Comparison of various loss functions

To address the underfitting described above for datasets with a moderate number of classes, prior work has proposed two main strategies: (i) Implementing an unbounded loss of the form $\mathcal{L} = \alpha \mathcal{L}_{\text{bound}} + \beta \mathcal{L}_{\text{unbound}}$ (Wang et al. 2019), and (ii) designing bounded losses that decay slowly towards outliers, thus assigning them non-zero gradients (Ma et al. 2020; Zhou et al. 2021; Zhang and Sabuncu 2018; Amid et al. 2019). In Tab. 1, we show a comprehensive list of noise-robust loss functions that we will use to identify potential problems that arise when learning with these losses. The parameter settings for each of these loss functions, as per their respective publications, are documented in the appendix.

In Fig. 2 we show the average error $\langle \delta_k \rangle$ in the “correct” neuron as a function of its pre-activation across a set

loss	definition	parameters	output error δ_n	Ref.
CE	$-\log(a_k)$	0	$a_n - \delta[n - k]$	-
MAE	$2(1 - a_k)$	0	$2a_k(a_n - \delta[n - k])$	-
genCE	$q^{-1}(1 - (a_k)^q)$	1	$(a_k)^q(a_n - \delta[n - k])$	(Zhang and Sabuncu 2018)
biTemp	no analytic form	0	no analytic form	(Amid et al. 2019)
NCE	$\frac{\log(a_k)}{\sum_i \log(a_i)}$	0	see appendix	(Ma et al. 2020)
NF	$\frac{\log((1-a_k)^{0.5} a_k)}{\sum_i \log((1-a_i)^{0.5} a_i)}$	0	see appendix	(Ma et al. 2020)
AGCE	$[(b+1)^q - (b+a_k)^q]/q$	2	$a_k(a_n - \delta[n - k])(b + a_n)^{q-1}$	(Zhou et al. 2021)
NF-MAE	α NF + β MAE	1	see appendix	(Ma et al. 2020)
NCE-MAE	α NCE + β MAE	1	see appendix	(Ma et al. 2020)
NCE-AGCE	α NCE + β AGCE	4	see appendix	(Zhou et al. 2021)

Table 1: Overview of the examined loss functions, with the neuron associated with the correct label indexed by k . The reported number of parameters excludes values that we keep constant for all datasets. Instead, it focuses on values actively adjusted based on the dataset under consideration. Parameter values utilized in our simulations, sourced from respective publications, are detailed in the appendix.

of noise-robust loss functions. Panel (a) shows a scenario with 10 classes for a neural network at initialization. The lower section of this panel represents the anticipated logit distribution $\mathcal{N}(0, 1)$ for the neuron k . It is obvious that these noise-robust loss functions deviate from the traditional Cross Entropy loss. Specifically, these functions assign gradients of reduced magnitude to examples with higher uncertainty (represented by small z_k) than the Cross Entropy loss. Despite this difference, the errors significantly deviate from zero at points of expected initial learning, as the histogram shows. However, in a 1000-class scenario, the overlap between the initial distribution and the regions where noise-robust losses record substantial errors is almost non-existent. Subsequent sections will show how this characteristic leads to negligible learning in the context of the 1000-class WebVision dataset. Already in the case of 100 classes (not shown), the overlap is significantly reduced, slowing down learning on the Cifar-100 dataset.

Instead of using the unbounded Cross Entropy loss, which indiscriminately learns from outliers (indicated by very small z_k values), we suggest an alternative: adjusting the logit distribution to $p_z(z_k) = \mathcal{N}(\epsilon, 1)$, to reestablish the crucial overlap with the error, as visually represented by the green histogram.

3.2 Boosting learning with example-dependent logit bias

In Fig. 1, we have illustrated the challenge of using the MAE loss in scenarios with a large number of classes. The root of the challenge is the reduced overlap between the initial logit distribution and the region where the derivative of the loss is substantial. To tackle this, we introduce an example-dependent bias, denoted as ϵ , to the logit, adjusting it from z_k to $z_k + \epsilon$. This adjustment ensures that the achieved learning speed is preserved irrespective of the class count. As before, k refers to the class specified by the label.

The consequence of this adjustment for a newly initial-

ized network is demonstrated in Fig. 2(b). Here, the logit distribution, $p_z(z_k)$, shifts from $\mathcal{N}(0, 1)$ to $\mathcal{N}(\epsilon, 1)$, reinstating the overlap with the region where significant gradients can emerge. A noteworthy aspect of the MAE error is its implicit dependency on class count through the relation $\delta_n = 2(1 - a_k)$. Since the a_k distribution is influenced by the class number, maintaining a consistent a_k distribution would ensure a steady error. To realize this, we aim to solve the following implicit equation for $\epsilon(K)$:

$$\left\langle \frac{\exp(z_k + \epsilon(K))}{\sum_{i=1}^K \exp(z_i + \delta[i - k]\epsilon(K))} \right\rangle = C, \quad (4)$$

where $C \in [0, 1]$ represents a freely chosen constant which specifies the average activation of a_k after the logit bias is applied. For the MAE loss our choice is $C = 0.15$. This particular value ensures a substantial overlap between the network’s initial logit distribution and the region where substantial gradients can occur, whilst excluding outliers with respect to the initial distribution.

We can evaluate Eq. (4) numerically for varying class counts. This is achieved by sampling Gaussian numbers, z_i , for a given ϵ , and subsequently refining our estimate using techniques such as binary search. Results for the dependence of ϵ on the class number are presented in the appendix. To clarify, when a logit bias is incorporated into a loss function, we append a “*” to the name of the loss.

3.3 Calculation of hyperparameters

An alternative approach is the adjustment of hyperparameters for different learning scenarios and datasets. However, choosing the right parameters can be highly non trivial and resource intensive. To reduce the number of hyperparameters, we propose to instead use the overlap shown in Fig. 2 to compute parameters without training. Using hyperparameters empirically chosen for a specific number of classes (e.g. for Cifar-10), one computes the value of the backpropagation error $\langle \delta_k \rangle(z_k)$ via Eq. (3) at the position $z_k = 0$

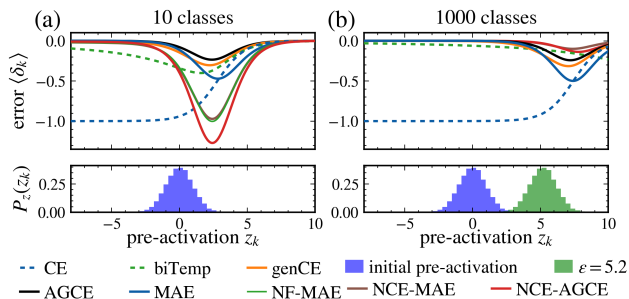


Figure 2: The average error $\langle \delta_k \rangle$ of the final layer’s correct neuron k , a determinant factor for the magnitude of a gradient descent update, plotted against the pre-activation z_k , for a network at initialization. Panel (a) depicts a ten-class learning scenario, showcasing a pronounced overlap between the region where learning is possible (characterized by large negative $\langle \delta_k \rangle$) and the logit distribution from a randomly initialized network (blue histogram). In contrast, in a scenario with 1000 classes shown in panel (b), this overlap is not present for the bounded loss functions. Rather than employing a tailed loss — exemplified by biTemp (delineated by the dashed green line) — our approach involves shifting the z_k distribution into the learning-possible range by adding a bias to the correct neuron’s pre-activation (green histogram). We note that the average error of NCE-AGCE is rescaled by a factor of 0.5 in panel a) to enhance visual clarity.

for this number of classes. When changing the number of classes, the hyperparameters are adjusted such that $\langle \delta_k \rangle(0)$ coincides with the value of the baseline for which it is known that learning works well. In the case of a single unknown this method yields the parameter as a function of the number of classes. In Fig. 3 we demonstrate how this optimization changes the learning curve of genCE when going from 100 to 1000 classes. We find that the parameter $q = 0.48$ provides the same $\langle \delta_k \rangle(0)$ for $K = 1000$ classes as in the case of $K = 100$ with $q = 0.7$. This is empirically verified in Sec. 4 while further details on the calculation are shown in the appendix.

4 Empirical Results

In the following, we present empirical evidence highlighting the efficiency of the logit bias in enhancing the learning capability of MAE on all datasets, demonstrating competitive or superior performance in the presence of label noise. In addition, we enable the loss functions genCE, NCE-AGCE and NF-MAE, to learn on the more complicated WebVision dataset using either new optimized hyperparameters or the logit bias, showcasing that the overlap between the initial output distribution and $\langle \delta_k \rangle$ is indeed the crucial requirement for learning. Datasets employed in this study include the publicly available datasets Cifar-10, Cifar-100 (Krizhevsky, Hinton et al. 2009), Fashion-MNIST (Xiao, Rasul, and Vollgraf 2017), and the web-scraped WebVision dataset (Li et al. 2017). The noise profile in Cifar datasets is inherently min-




Figure 3: Computation of the hyperparameter q for genCE when changing from 100 to 1000 classes. Instead of performing an expensive hyperparameter search, we suggest to adjust the hyperparameter such that $\langle \delta_k \rangle$ at position $z_k = 0$ is unchanged for the larger class size.

imal, enabling us to precisely adjust the label noise level. In contrast, WebVision naturally offers a more complex, potentially asymmetric noise structure. ResNet architectures (He et al. 2016), tailored in size to the specific dataset, are used for most of our experiments. We create symmetric label noise by assigning random labels to each example with probability p . In the case of asymmetric label noise we flip the labels of specific classes to different classes with probability p . The classes for the asymmetric noise are chosen identical to (Zhang and Sabuncu 2018). Details on the training, together with further results from alternative network architectures, are available in the appendix.

Table 2 summarizes the results on the datasets Cifar-10, Cifar-100 and FashionMNIST providing mean accuracies and associated errors (across five seeds) under various levels of label noise. The best performing loss functions are emphasized in bold for every combination of network architecture and noise level. For Cifar-10, we find that in the absence of label noise, MAE is unable to sufficiently learn the dataset, performing worse than all other loss functions. Yet, the introduction of a small logit bias (specifically, $\epsilon = 0.5$) in MAE* noticeably enhances the learning, bringing it on par with Cross Entropy loss. In scenarios with symmetric label noise, MAE* maintains its noise resilience often performing best out of all loss functions. Except for genCE and MAE*, most proposed loss functions either attain diminished accuracies in noise-free scenarios relative to the Cross Entropy or fit the noise to a high degree. Overfitting is especially present for the Cross Entropy and biTemp loss, a trait anticipated from our theoretical discussions — especially noting biTemp’s pronounced outlier tail in Fig. 2. Similar results are found for the FashionMNIST dataset.

For the Cifar-100 dataset MAE* emerges as superior in the noiseless case, suggesting that even the minimal noise inherent in the Cifar-100 dataset (Northcutt, Athalye, and Mueller 2021) puts bounded loss functions in a favored position relative to CE. It is worth noting that ϵ is not optimized, but derived from Eq. (4), based on the results of the ten-class tasks. For the case of Cifar-100 with added noise, a different picture emerges: The loss functions CE, biTemp, and MAE*

Dataset, Network	Loss	Noise: 0 %	10 %	20 %	40 %	60 %
Sym. Noise	CE	91.30 ± 0.33	83.80 ± 0.36	78.21 ± 0.42	65.65 ± 0.23	50.12 ± 0.66
	MAE	84.90 ± 1.86	82.59 ± 1.22	80.45 ± 0.24	74.08 ± 4.17	59.90 ± 6.19
Cifar-10	MAE* $\epsilon = 0.5$	91.27 ± 0.13	90.67 ± 0.09	89.95 ± 0.06	87.33 ± 0.07	81.43 ± 0.19
	AGCE	88.32 ± 1.86	87.83 ± 1.87	87.06 ± 1.86	84.77 ± 1.76	82.24 ± 1.37
	genCE	91.70 ± 0.10	91.12 ± 0.04	89.92 ± 0.11	86.81 ± 0.12	78.88 ± 0.30
ResNet-32	NCE-AGCE	64.70 ± 1.51	59.21 ± 1.13	55.77 ± 2.44	46.19 ± 3.56	35.76 ± 1.85
	NF-MAE	88.47 ± 0.26	87.64 ± 0.18	86.91 ± 0.07	84.05 ± 0.37	75.98 ± 0.49
	NCE-MAE	88.57 ± 0.09	87.73 ± 0.14	86.81 ± 0.08	84.00 ± 0.19	75.67 ± 0.57
	biTemp	92.07 ± 0.05	86.84 ± 0.17	81.33 ± 0.16	68.06 ± 0.63	51.58 ± 0.40
	CE	94.90 ± 0.07	92.41 ± 0.10	90.72 ± 0.07	85.60 ± 0.50	78.15 ± 0.71
	MAE	90.72 ± 3.37	93.74 ± 0.06	91.71 ± 1.88	88.24 ± 2.11	86.50 ± 2.16
Fashion-MNIST	MAE* $\epsilon = 0.5$	94.77 ± 0.03	94.42 ± 0.02	94.23 ± 0.05	93.50 ± 0.11	92.04 ± 0.12
	AGCE	94.39 ± 0.06	94.24 ± 0.07	94.01 ± 0.06	93.35 ± 0.06	92.28 ± 0.09
	genCE	94.79 ± 0.06	94.40 ± 0.06	94.28 ± 0.02	93.46 ± 0.07	91.11 ± 0.15
ResNet-32	NCE-AGCE	92.73 ± 0.10	92.57 ± 0.10	92.12 ± 0.05	91.58 ± 0.09	90.09 ± 0.13
	NF-MAE	93.92 ± 0.07	93.74 ± 0.06	93.71 ± 0.06	92.84 ± 0.07	91.75 ± 0.10
	NCE-MAE	93.98 ± 0.04	93.73 ± 0.07	93.51 ± 0.06	93.00 ± 0.07	91.73 ± 0.12
	biTemp	94.78 ± 0.03	93.38 ± 0.03	91.49 ± 0.09	86.65 ± 0.50	78.05 ± 0.47
	CE	75.88 ± 0.15	69.21 ± 0.12	63.07 ± 0.15	46.46 ± 0.24	24.69 ± 0.68
	MAE	6.84 ± 0.66	5.30 ± 0.36	4.93 ± 0.43	2.54 ± 0.29	1.67 ± 0.17
Cifar-100	MAE* $\epsilon = 1.5$	74.81 ± 0.13	72.33 ± 0.13	68.30 ± 0.24	55.71 ± 0.45	35.51 ± 0.19
	MAE* $\epsilon = 3.0$	76.49 ± 0.09	69.11 ± 0.14	61.13 ± 0.16	44.20 ± 0.41	26.45 ± 0.16
	AGCE	76.32 ± 0.12	73.23 ± 0.11	67.91 ± 0.19	52.64 ± 0.15	32.48 ± 0.12
ResNet-34	genCE	73.90 ± 0.17	72.15 ± 0.09	69.37 ± 0.20	61.15 ± 0.10	41.52 ± 0.64
	NCE-AGCE	72.46 ± 0.09	70.00 ± 0.30	66.20 ± 0.35	46.91 ± 0.69	10.86 ± 0.42
	NF-MAE	74.66 ± 0.17	72.91 ± 0.14	70.25 ± 0.06	60.97 ± 0.14	25.23 ± 0.68
	NCE-MAE	74.85 ± 0.10	72.71 ± 0.07	70.03 ± 0.10	60.51 ± 0.10	26.77 ± 0.81
	biTemp	76.19 ± 0.14	68.40 ± 0.12	61.36 ± 0.19	46.15 ± 0.23	26.67 ± 0.16
Asym. Noise	CE	91.98 ± 0.11	91.62 ± 0.20	91.63 ± 0.27	85.28 ± 6.82	64.86 ± 6.71
	MAE	83.41 ± 1.49	83.62 ± 1.48	80.02 ± 1.79	77.25 ± 6.32	65.35 ± 6.40
Cifar-10	MAE* $\epsilon = 0.5$	91.14 ± 0.08	91.15 ± 0.09	91.33 ± 0.15	85.82 ± 5.68	65.51 ± 6.48
	AGCE	89.45 ± 1.57	87.80 ± 1.87	89.43 ± 1.54	83.78 ± 5.51	68.40 ± 5.63
	genCE	91.84 ± 0.06	91.65 ± 0.06	91.69 ± 0.10	84.86 ± 6.74	64.51 ± 6.72
ResNet-32	NCE-AGCE	61.21 ± 1.26	64.00 ± 1.97	62.74 ± 1.65	55.29 ± 2.94	49.10 ± 2.96
	NF-MAE	88.50 ± 0.05	88.49 ± 0.17	88.71 ± 0.07	82.03 ± 6.29	63.12 ± 6.31
	NCE-MAE	88.60 ± 0.11	88.41 ± 0.08	88.40 ± 0.15	82.22 ± 6.40	63.03 ± 6.27
	biTemp	92.24 ± 0.04	92.19 ± 0.08	92.27 ± 0.06	85.52 ± 6.86	65.07 ± 6.78
	CE	94.91 ± 0.07	92.38 ± 0.13	90.38 ± 0.16	85.61 ± 0.20	78.07 ± 0.72
	MAE	91.94 ± 1.84	93.81 ± 0.10	93.38 ± 0.16	87.91 ± 3.26	88.83 ± 1.85
FashionMNIST	MAE* $\epsilon = 0.5$	94.67 ± 0.08	94.48 ± 0.06	94.19 ± 0.06	93.33 ± 0.04	91.86 ± 0.04
	AGCE	94.44 ± 0.12	94.31 ± 0.11	94.01 ± 0.07	93.31 ± 0.07	92.25 ± 0.10
	genCE	94.81 ± 0.06	94.47 ± 0.03	94.23 ± 0.09	93.27 ± 0.03	90.75 ± 0.15
ResNet-32	NCE-AGCE	92.69 ± 0.04	92.63 ± 0.08	92.28 ± 0.17	91.71 ± 0.09	90.42 ± 0.17
	NF-MAE	93.92 ± 0.06	93.81 ± 0.06	93.55 ± 0.04	92.82 ± 0.10	91.75 ± 0.17
	NCE-MAE	94.07 ± 0.05	93.81 ± 0.06	93.57 ± 0.05	92.93 ± 0.05	91.74 ± 0.06
	biTemp	94.90 ± 0.08	93.37 ± 0.10	91.47 ± 0.11	86.64 ± 0.34	77.56 ± 0.62

Table 2: Final test accuracies when training with different loss functions on the Cifar-10, Fashion-MNIST and Cifar-100 dataset using various amounts of label noise. In the case of symmetric label noise with ten-classes (first two panels) we find that MAE* improves significantly over MAE, performing as well or superior to the CE in the noiseless case while often achieving the top result in the case of label noise. For the Cifar-100 dataset, the logit bias allows MAE* to learn exceptionally well compared to MAE, even achieving the best test-accuracy under all loss functions on the pristine dataset. For larger label noises MAE* with the calculated logit bias $\epsilon = 3.0$ overfits, showcasing that a smaller logit bias (e.g. $\epsilon = 1.5$) might be preferable. In the case of asymmetric noise (bottom panels) MAE* again improves significantly over MAE on all noise level. When compared to the other loss functions we find that counted over all datasets and noise levels MAE* achieves the most top accuracies.

	NF-MAE	NCE-AGCE	genCE	CE	biTemp	MAE*
Top-1	1.04	7.48	13.28	62.55 ± 0.16	63.60 ± 0.14	58.72 ± 0.11
Top-5	1.17	10.08	15.39	79.63 ± 0.10	80.66 ± 0.13	79.68 ± 0.04
	NF-MAE**	NCE-AGCE+	genCE**	genCE+	MAE**	
Top-1	62.30 ± 0.01	63.90 ± 0.09	62.03 ± 0.07	64.85 ± 0.04	62.70 ± 0.07	
Top-5	80.95 ± 0.07	80.22 ± 0.09	80.70 ± 0.03	80.93 ± 0.02	81.11 ± 0.05	

Table 3: Final accuracies on the validation images of the 1000-class WebVision dataset after training Wide-ResNet-101 with different loss functions. If an error is provided, accuracies for five different seeds are averaged. For the losses that do not learn, we report the accuracies of a single training run. The upper row demonstrates that all loss function which where noise robust on Cifar-10 and Cifar-100 are unable to learn on WebVision, with the exception of MAE*. The second row demonstrates that adding either a scheduled logit bias (indicated with **) or calculating new parameters for the loss as described by our theory (denoted by +), all loss functions are capable of learning the WebVision dataset. The exact parameters in the second row are provided in the appendix.

tend to fit the noise, while the loss functions which underfitted the pristine dataset now display superior generalisation. This is especially pronounced for genCE, which lacks 2.5% behind in the noiseless case but performs best in the case of larger noise levels. For demonstration purpose we also add $\epsilon = 1.5$ to the table, showcasing that the logit bias ϵ can be used to tune a balance between learning speed and noise-robustness. The case of asymmetric noise is summarized in the bottom two panels of Tab. 2. In the case of MAE* we find very similar results to the case of symmetric label noise in terms of improvements over MAE, as well as in cases where it performs best. Considering all datasets and noise variants, MAE* with logit bias is the loss function with the most top accuracies.

In the 100 class scenario discussed above the majority of the studied loss functions still demonstrate good performance, albeit at a marginally reduced learning speed. Their design incorporates “tails” that encompass the logit distribution at initialization. These tails ensure non-zero gradients, enabling the learning of the dataset. However, when scaling to 1000 classes, most tails appear to fall flat. This is demonstrated in the first row of Tab. 3, where we find that all loss function which where noise robust on Cifar-10 and Cifar-100 are unable to learn on the WebVision dataset, with the exception of MAE*. The second row demonstrates that if we restore the overlap with the initial distribution using either the logit bias or by calculating new parameters for these losses as described in Sec. 3.3 (denoted by +), these losses are now capable of learning on the WebVision dataset.

Though learning works well with the calculated value of the logit bias $\epsilon = 5.2$, removal of the logit bias during testing reveals reduced confidence in the correct label, induced by the elevated ϵ value during the training phase. This can be mitigated by weaning the network from its high confidence prediction by adopting smaller values of ϵ throughout training using an ϵ schedule (denoted by **), as elaborated in the appendix. Comparing the losses against each other, we find that the ϵ enhanced losses NF-MAE**, genCE**, and MAE** perform best in the top-5 accuracy, while genCE+ and NCE-AGCE+ with calculated new hyperparameters perform best in the top-1 accuracy, showcasing that our method

provides the missing ingredient to these losses.

Limitations— Our theoretical insights indicate that our primary findings should exhibit general applicability across numerous learning tasks, however, the current scope is confined to image classification. Furthermore, it seems that our present computation of ϵ based on the number of classes tends to overestimate the required ϵ , which can compromise noise robustness. To address this issue, future theoretical studies could explore a more sophisticated understanding of this phenomenon, while empirical analyses could configure ϵ as a tunable hyperparameter to circumvent this limitation.

5 Conclusion

In this study, we analyzed the dynamics of early-stage learning by computing the average backpropagation error, offering quantitative insights into how increasing the class count influences initial learning, especially in the context of bounded loss functions. Capitalizing on these insights, we showed how learning with bounded losses can be enhanced by calculating how specific hyperparameters depend on the number of classes, effectively eradicating the need of extensive hyperparameter searching. For the hyperparameter-free case of the MAE-loss, we enabled learning using the *logit bias*. Both techniques use the method of realigning the distribution of a newly initialized network with the region where the loss function has a nonzero derivative, enabling effective learning even in scenarios with a multitude of classes. Empirical evidence underscores the effectiveness of this method. The loss functions with calculated hyperparameters emerged as the best performing loss functions on WebVision regarding the top-1 accuracy, while the logit bias enhanced MAE loss demonstrates comparable, if not superior, performance across datasets spanning ten to a thousand classes. Before our approach, such outcomes were largely exclusive to the Cross Entropy or biTemp loss. However, both these methods tend to overfit. In summary, we argue that our method is a first step towards a comprehensive framework – one that allows for noise robust learning, regardless of the number of classes, and without an over-reliance on fine-tuned hyperparameters.

Acknowledgments

Computations for this work were done (in part) using resources of the Leipzig University Computing Center.

References

- Amid, E.; Warmuth, M. K.; Anil, R.; and Koren, T. 2019. Robust bi-tempered logistic loss based on bregman divergences. *Advances in Neural Information Processing Systems*, 32.
- Author(s), A. 2023. All code, scripts, and data used in this work are included in a Zenodo archive: <https://zenodo.org/records/13150928>. *Zenodo*.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, 1597–1607. PMLR.
- Engleson, E.; and Azizpour, H. 2021. Generalized jensen-shannon divergence loss for learning with noisy labels. *Advances in Neural Information Processing Systems*, 34: 30284–30297.
- Feng, L.; Shu, S.; Lin, Z.; Lv, F.; Li, L.; and An, B. 2021. Can cross entropy loss be robust to label noise? In *Proceedings of the Twenty-Ninth International Conference on International Joint Conferences on Artificial Intelligence*, 2206–2212.
- Ghosh, A.; Kumar, H.; and Sastry, P. S. 2017. Robust loss functions under label noise for deep neural networks. In *Proceedings of the AAAI conference on artificial intelligence*, volume 31.
- Ghosh, A.; and Lan, A. 2021. Contrastive learning improves model robustness under label noise. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2703–2708.
- Goldberger, J.; and Ben-Reuven, E. 2017. Training deep neural-networks using a noise adaptation layer. In *International conference on learning representations*.
- Goodfellow, I.; Bengio, Y.; and Courville, A. 2016. *Deep learning*. MIT press.
- Han, B.; Yao, J.; Niu, G.; Zhou, M.; Tsang, I.; Zhang, Y.; and Sugiyama, M. 2018. Masking: A new perspective of noisy supervision. *Advances in neural information processing systems*, 31.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Hendrycks, D.; Mazeika, M.; Kadavath, S.; and Song, D. 2019. Using self-supervised learning can improve model robustness and uncertainty. *Advances in neural information processing systems*, 32.
- Khetan, A.; Lipton, Z. C.; and Anandkumar, A. 2017. Learning from noisy singly-labeled data. *arXiv preprint arXiv:1712.04577*.
- Krizhevsky, A.; Hinton, G.; et al. 2009. Learning multiple layers of features from tiny images. *Tech Report*.
- Lee, J.; Bahri, Y.; Novak, R.; Schoenholz, S. S.; Pennington, J.; and Sohl-Dickstein, J. 2017. Deep neural networks as gaussian processes. *arXiv preprint arXiv:1711.00165*.
- Lee, K.-H.; He, X.; Zhang, L.; and Yang, L. 2018. Cleanet: Transfer learning for scalable image classifier training with label noise. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5447–5456.
- Li, W.; Wang, L.; Li, W.; Agustsson, E.; and Van Gool, L. 2017. Webvision database: Visual learning and understanding from web data. *arXiv preprint arXiv:1708.02862*.
- Liang, X.; Liu, X.; and Yao, L. 2022. Review—a survey of learning from noisy labels. *ECS Sensors Plus*, 1(2): 021401.
- Lyu, Y.; and Tsang, I. W. 2019. Curriculum loss: Robust learning and generalization against label corruption. *arXiv preprint arXiv:1905.10045*.
- Ma, X.; Huang, H.; Wang, Y.; Romano, S.; Erfani, S.; and Bailey, J. 2020. Normalized loss functions for deep learning with noisy labels. In *International conference on machine learning*, 6543–6553. PMLR.
- Neal, R. M. 1994. Priors for infinite networks (tech. rep. no. crg-tr-94-1). *University of Toronto*, 415.
- Northcutt, C. G.; Athalye, A.; and Mueller, J. 2021. Pervasive label errors in test sets destabilize machine learning benchmarks. *arXiv preprint arXiv:2103.14749*.
- Reed, S.; Lee, H.; Anguelov, D.; Szegedy, C.; Erhan, D.; and Rabinovich, A. 2014. Training deep neural networks on noisy labels with bootstrapping. *arXiv preprint arXiv:1412.6596*.
- Ren, M.; Zeng, W.; Yang, B.; and Urtasun, R. 2018. Learning to reweight examples for robust deep learning. In *International conference on machine learning*, 4334–4343. PMLR.
- Song, H.; Kim, M.; Park, D.; Shin, Y.; and Lee, J.-G. 2022. Learning from noisy labels with deep neural networks: A survey. *IEEE Transactions on Neural Networks and Learning Systems*.
- Sukhbaatar, S.; and Fergus, R. 2014. Learning from noisy labels with deep neural networks. *arXiv preprint arXiv:1406.2080*, 2(3): 4.
- Thulasidasan, S.; Bhattacharya, T.; Bilmes, J.; Chennupati, G.; and Mohd-Yusof, J. 2019. Combating label noise in deep learning using abstention. *arXiv preprint arXiv:1905.10964*.
- Toner, W.; and Storkey, A. 2023. Label Noise: Correcting a Correction. *arXiv preprint arXiv:2307.13100*.
- Veit, A.; Alldrin, N.; Chechik, G.; Krasin, I.; Gupta, A.; and Belongie, S. 2017. Learning from noisy large-scale datasets with minimal supervision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 839–847.
- Wang, Y.; Ma, X.; Chen, Z.; Luo, Y.; Yi, J.; and Bailey, J. 2019. Symmetric cross entropy for robust learning with noisy labels. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 322–330.
- Xiao, H.; Rasul, K.; and Vollgraf, R. 2017. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*.

- Xiao, T.; Xia, T.; Yang, Y.; Huang, C.; and Wang, X. 2015. Learning from massive noisy labeled data for image classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2691–2699.
- Xu, Y.; Cao, P.; Kong, Y.; and Wang, Y. 2019. L_{dmi}: A novel information-theoretic loss function for training deep nets robust to label noise. *Advances in neural information processing systems*, 32.
- Xue, Y.; Whitecross, K.; and Mirzasoleiman, B. 2022. Investigating why contrastive learning benefits robustness against label noise. In *International Conference on Machine Learning*, 24851–24871. PMLR.
- Yang, G.; Hu, E.; Babuschkin, I.; Sidor, S.; Liu, X.; Farhi, D.; Ryder, N.; Pachocki, J.; Chen, W.; and Gao, J. 2021. Tuning large neural networks via zero-shot hyperparameter transfer. *Advances in Neural Information Processing Systems*, 34: 17084–17097.
- Yi, L.; Liu, S.; She, Q.; McLeod, A. I.; and Wang, B. 2022. On learning contrastive representations for learning with noisy labels. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 16682–16691.
- Zhang, Z.; and Sabuncu, M. 2018. Generalized cross entropy loss for training deep neural networks with noisy labels. *Advances in neural information processing systems*, 31.
- Zheltonozhskii, E.; Baskin, C.; Mendelson, A.; Bronstein, A. M.; and Litany, O. 2022. Contrast to divide: Self-supervised pre-training for learning with noisy labels. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 1657–1667.
- Zhou, X.; Liu, X.; Jiang, J.; Gao, X.; and Ji, X. 2021. Asymmetric loss functions for learning with noisy labels. In *International conference on machine learning*, 12846–12856. PMLR.