THE WEIGHTED MEAN TRICK – OPTIMIZATION STRATEGIES FOR ROBUSTNESS

Anonymous authors
Paper under double-blind review

ABSTRACT

We prove that minimizing a weighted mean results in optimizing the higher-order moments of the loss distribution such as the variance, skewness, and kurtosis. By optimizing the higher-order moments, one can tighten the upper bound of the loss mean deviating from the true expectation and improve the robustness against outliers. Such types of optimization problems often lead to non-convex objectives, therefore, we explore the extent to which the proposed weighted mean trick preserves convexity, albeit at times at a decrease in efficiency. Experimental results show that the weighted mean trick exhibits similar performance with other specialized robust loss functions when training on noisy datasets while providing a stronger theoretical background. The proposed weighted mean trick is a simple yet powerful optimization framework that is easy to integrate into existing works.

1 INTRODUCTION

The most common objective in machine learning is to minimize the average loss. However, by doing this it implies that all samples are equal. The goal of this work is to convince the reader that from the perspective of the model, not all samples are equal and optimizing a weighted mean is a more progressive approach. For instance, at the end of the training, few samples are left uncaptured by the model which yield large loss values. One might choose to put more weight on those samples to help the model learn (individual fairness), or consider them noise and assign them less weight (robustness). The decision will also be reflected by the loss variance; it will decrease when hard samples are weighted more and increase otherwise.

Before introducing more formally the weighted mean, we first justify theoretically the impact of variance penalization. By Bennett's inequality (Hoeffding, 1963; Maurer & Pontil, 2009) for i.i.d. loss values Z and bounded variance we have with probability $1-\delta$:

$$\mathbb{E}[Z] - \frac{1}{n}\sum_{i=1}^{n} Z_i \quad \leq \quad C_1\sqrt{\frac{2\mathbb{V}[Z]\ln 1/\delta}{n}} + \frac{C_2 \ln 1/\delta}{3n} \tag{1}$$

where $C_1,\ C_2$ are problem depended constants. This inequality reveals two things. First, we have with high probability that the empirical mean is close to the theoretical value. This bound along with similar PAC-Bayes bounds (Seldin et al., 2012; Tolstikhin & Seldin, 2013) justifies the practical success of the empirical risk minimization (ERM) which has the objective to minimize the mean loss value (Namkoong & Duchi, 2017). Secondly, the difference between the two is bounded by the variance. This second implication led to a large and growing body of studies that investigate variance penalization (Maurer & Pontil, 2009; Namkoong & Duchi, 2017; Duchi & Namkoong, 2019; Staib et al., 2019; Lam, 2019; Heinze-Deml & Meinshausen, 2021; Hu et al., 2018). Moreover, by penalizing the variance, the intrinsic bias-variance tradeoff of the ERM can be controlled. Two studies influential for this paper are that of Duchi & Namkoong (2019) and Li et al. (2021). Duchi & Namkoong (2019) proposed taking the expectation with respected to a different distribution which allowed penalizing variance while preserving the convexity of the objective. Similarly, Li et al. (2021) investigated optimizing a tilted empirical risk which is equivalent to penalizing all the higher-order moments simultaneously. In summary, previous methods either penalize only one moment or all the higher-order moments but are not flexible enough to penalize any desired combination of the higher-order moments.

Inspired by the above works, we propose to optimize a weighted mean and prove that for certain weights, it is equivalent to optimizing any higher-order moments of the loss distribution such as variance, skewness, and kurtosis. Our approach generalizes that of Duchi & Namkoong (2019); Li et al. (2021) while simplifying the optimization procedure and enabling separate penalization of the higher-order moments.

In particular, we will construct the weights such that optimizing the weighted mean is equivalent to applying a variance penalization:

$$\mathbb{E}[W\ell] = \mathbb{E}[\ell] + \lambda\mathbb{V}[\ell] \tag{2}$$

or penalization of other central moments (e.g., skewness, kurtosis). Construction of these weights are treated by Theorems 1 and 3 from section 2, which require minimal computational resources. Important note, penalizing directly the variance preserves convexity only for values of λ within a very narrow range (Maurer & Pontil, 2009). However, the weighted mean method yields a convex objective for any positive values of λ.

In detail, our contributions are as follows:

(**C1**) We build upon the work of Duchi & Namkoong (2019) and prove that optimizing a weighted mean is equivalent to reducing or amplifying both variance (Theorem 1) and higher-order moments (Theorem 3).

(**C2**) We derive the limits of the penalization interval which preserves convexity (Lemma 2) and also show how to penalize with values outside this interval while still maintaining convexity (Lemma 4).

(**C3**) We connect the variance and higher-order moments penalization using weighted mean to the robustification of loss functions (Lemma 5 and Lemma 6).

(**C4**) We develop a convex version of the variance penalized cross-entropy loss which provides a higher accuracy in high noise scenarios with class dependent noise.

(**C5**) We show experimentally that a negative variance penalization improves model accuracy when training with noisy labels.

The implications of the weighted mean are much broader than what we investigate in this work. We limit the scope of this paper to classification using deep neural networks trained with noisy labels. But the mathematical framework also covers the control of the bias-variance trade-off and can also be applicable to regression problems.

In the sequel, we introduce the moment penalization strategy and the weighted mean trick in section 2, which is a computational technique to make moment penalization practical. Next, in section 3, we illustrate the application of moment penalization using the weighted mean formulation to optimize for robustness when training with noisy labels.

2 Penalizing Moments

In ERM, the impact on the model parameters of a sample is determined by the balance between its loss value and the average of the training batch. Especially observed in the latter stages of the training when majority of samples have a small loss value except few which are left uncaptured by the model. Case in which we consider them as unlearned samples and amplify their impact or as outliers and suppress their impact on the model. The *weighted mean trick* consist in assigning weights to each sample based on the loss value to either amplify or suppress their impact on training. This is similar to dropout (Hinton et al., 2012) but for losses. However, *weighted mean trick* is a deterministic process and the weights are not restricted to 0 or 1 but can take any non-negative value.

In what follows, we apply the *weighted mean trick* to extend the ERM framework to multi-objective optimization of the mean and higher-order moments of the loss function. First, we show that for some distinct weights, optimizing the weighted mean is equivalent to a simultaneous optimization of the mean and variance.

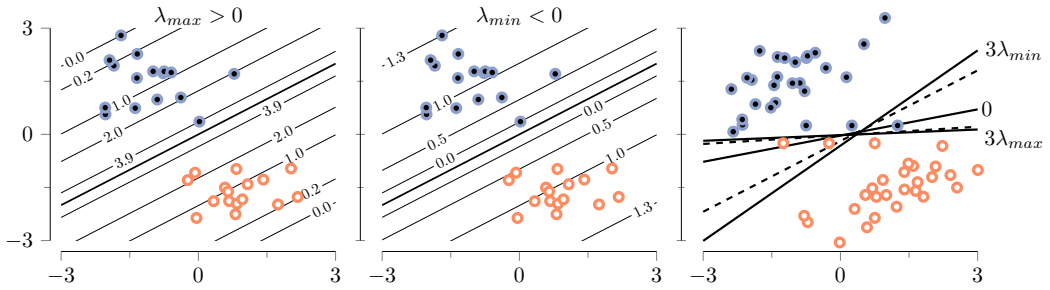

Figure 1: Variance penalization for classification problems. Contour lines of the weights distribution for positive λ are shown on the left and for negative on the center plot. The right plot shows the use of variance penalization for outlier suppression or amplification (robustification versus generalization).

Theorem 1 (Variance Expansion)**.** *Let ℓ be a random variable with finite first and second moments and $W = 1 + \lambda(\ell - \mathbb{E}[\ell])$, then taking the weighted mean of ℓ with weights W is equivalent to:*

$$\mathbb{E}[W\ell] = \mathbb{E}[\ell] + \lambda\mathbb{V}[\ell] \tag{3}$$

Proof. Replacing W with the definition and using the linearity property of the expectation along with Proposition 7 we get: $\mathbb{E}[W\ell] = \mathbb{E}[\ell] + \lambda\mathbb{E}\big[(\ell - \mathbb{E}[\ell])\ell\big] = \mathbb{E}[\ell] + \lambda\mathbb{E}\big[(\ell - \mathbb{E}[\ell])^2\big] = \mathbb{E}[\ell] + \lambda\mathbb{V}[\ell]$ □

Thus, switching from mean to weighted mean allows us to control the bias-variance tradeoff through λ to improve the distributional robustness of the model (Maurer & Pontil, 2009). However, the range of λ values that preserve the convexity depends on the statistical properties of ℓ, average and the minimum value, as shown by the next lemma.

Lemma 2. *The variance expansion using theorem 1 of a convex objective ℓ is also convex if $\lambda \in [0, \lambda_{max}]$, where $\lambda_{max} = 1/(\mathbb{E}[\ell] - \min(\ell))$.*

Proof is provided in Appendix A. This lemma precisely shows why directly penalizing the variance does not preserve convexity besides when λ takes values in a narrow interval. Moreover, the upper limit of this interval is not constant and changes with each iteration of the optimization algorithm. Note that the further the minimum value $\min(\ell)$ is from the average loss $\mathbb{E}[\ell]$, the narrower the λ interval is. However, when using the weighted mean trick to penalize the variance and clip the negative weights to 0 then the objective remains convex for any positive λ. Clipping negative weights to 0 also improves convergence when λ takes negative values and the resulting objective is non-convex, as we will see later.

Figure 1 shows the weights distribution for different values of λ. When λ is positive (left plot) samples closer to the decision boundary (which also means with larger loss values) receive more weight. On the other hand, when λ is negative (center plot) samples with larger loss values receive less weight. Of note, for $\lambda < 0$ the objective is not convex, however as we will show later will still converge to an optimal solution. The right plot shows a binary classification problem where each class consists of two clusters, one cluster is diagonally aligned containing the majority of points and another consisting of the three central points aligned across a horizontal line. The parameter λ controls the placement of the decision boundary with respect to these two clusters. Positive λ values place more weight on the three horizontal points as they are closer to the decision boundary and favors the model to learn how to perfectly separate the two classes. On the other hand, negative λ values place less weight on the horizontal points close to the boundary and more on the farther ones. Therefore, this aligns the decision boundary only with respect to the two large clusters.

Variance is not the only central moment that can be penalized using the weighted mean. In fact, any combination of the moments can be penalized. The next theorem generalizes theorem 1.

Theorem 3 (Moments Expansion)**.** *Let ℓ be a random variable with finite first n central moments and $W = \sum_{i=1}^{n} \lambda_i(\ell - \mathbb{E}[\ell])^{i-1}$, then taking the weighted mean of ℓ with weights W is equivalent*

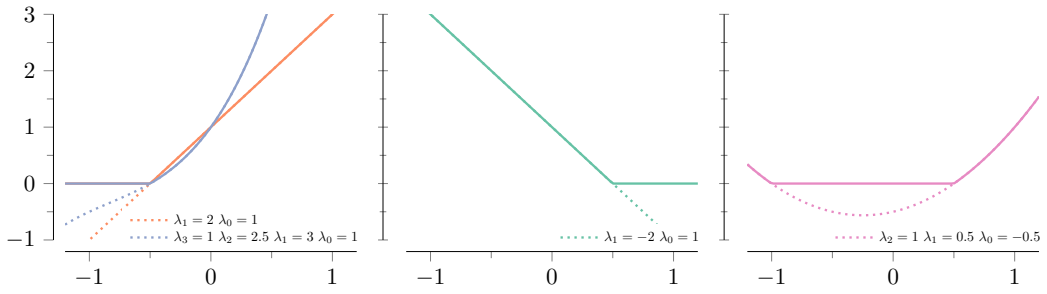

Figure 2: Polynomial functions for moments penalization. Dotted lines show the complete polynomial whereas solid lines the functions when the negative part is clipped to 0. Left plot shows two functions which are convex and non decreasing when their corresponding polynomial is clipped to 0. Center plot shows a convex polynomial that is non increasing and thus will result in a non convex objective. Right plot shows a polynomial which is convex and non decreasing only on the interval $[-1, +\infty]$, and thus will yield a convex objective when $\min(\ell) - \mathbb{E}[\ell] \geq -1$.

to:

$$\mathbb{E}[W\ell] = \lambda_1 \mathbb{E}[\ell] + \sum_{i=2}^{n} \tilde{\lambda}_i \mathbb{E}\left[(\ell - \mathbb{E}[\ell])^i\right] \tag{4}$$

where $\tilde{\lambda}_i = \lambda_i + \lambda_{i+1}\mathbb{E}[\ell]$ *for* $i < n$ *and* $\tilde{\lambda}_n = \lambda_n$

Proof is provided in Appendix A. We notice that penalizing moments higher than two incurs an additional penalization of the previous moment. For example, penalizing skewness by λ_3 incurs a variance penalization of $\lambda_3\mathbb{E}[\ell]$.

Theorem 3 also has an algebraic interpretation. Note that the formula for weights is nothing more than a polynomial in ℓ translated by $\mathbb{E}[\ell]$. When penalizing the variance, λ_2 controls the slope of the linear equation, and when penalizing the skewness, λ_3 controls the curvature of the quadratic equation. Moreover, the penalization factors λ_i also define the placement of the roots of the polynomial and the convexity of the weighted mean objective.

Lemma 4 (Convexity of Moments Expansion)**.** *Let* $\ell : \mathbb{R}^n \to \mathbb{R}_+$ *and* $p : \mathbb{R} \to \mathbb{R}_+$ *be two non negative convex functions and* $M \geq 0$*, then the* $W\ell(x)$ *with* $W = p(\ell(x) - M)$ *is convex if* p *is non decreasing.*

Proof is provided in Appendix A. The above theorem lists the two requirements on the polynomial function p such that the weighted mean objective remains convex. First, the function must be non-negative, thus, the negative weights must be clipped to 0. Secondly, when the function takes positive values, p must be non decreasing and convex. Figure 2 shows several examples of polynomials. Variance penalization implies p is an affine function, and thus, convex. However, for $\lambda_1 > 0$ the polynomial is non decreasing (left plot) and for $\lambda_1 < 0$ the polynomial is non increasing (center plot). Thus, only positive variance penalization will result in a convex objective. Of note, Lemma 2 upper bounds λ_1 if weights are not clipped to 0, however, if clipping is used λ_1 is not upper bounded. When penalizing skewness, p is a quadratic function (right plot). In this case, if $\lambda_2 > 0$ then p is convex and non decreasing only on a restricted interval instead of the entire real line. Thus, convexity can be preserved by appropriately adjusting the roots of the polynomial such that p is convex on $[\min(\ell) - \mathbb{E}[\ell], +\infty]$.

When clipping negative weights to 0 it prevents the optimization objective to switch from minimization to maximization, usually an undesirable behavior. Consequently, the samples with a corresponding zero weight reach their maximum contribution in the moments penalization. Further increasing the factors λ_i, will have no effect on those samples and only samples with non-zero weights will participate in the training. As a result, the efficiency of the moments penalization will slightly fall. Moreover, the samples with zero weight will be excluded from training, technique used in the past by multiple studies. Rockafellar & Uryasev (2000) used a similar clipping technique to optimize only for samples that are part of the tail of the distribution.

The Moments Expansion theorem extends the variance expansion proposed by Duchi & Namkoong (2019) to include higher-order moments. However, even when only the variance is penalized, the two methods are still slightly different. Moments expansion theorem computes the weights directly from the loss values, whereas the variance expansion of Duchi & Namkoong (2019) solves a secondary optimization problem to find the weights. The use of a secondary optimization problem has the advantage of penalizing the variance more consistently despite being more computationally expensive. On the contrary, when using the moments expansion, the penalized variance can be slightly lower depending on the number of weights that are 0. However, the advantage is that the direct computation of the weights makes it easier to include the method into existing analysis frameworks. A similar method that extends the optimization objective to include penalization factors for higher-order moments was proposed by Li et al. (2021). The proposed method replaces the ERM objective with a tilted version calculated as $\frac{1}{t}K(t) = \frac{1}{t}\log \mathbb{E}[e^{t\ell}]$ where $K(t)$ is the cumulant-generating function of the loss ℓ. The penalization of the higher order moments of the tilted ERM can be recovered from the power series expansion of $K(t)$. The distinction between the two is that the moments penalization introduced in this paper represents a generalization of the tilted ERM as it allows any combination of the higher order moments to be penalized whereas tilted ERM uses a single parameter that governs the penalization factors. In summary, the moments penalization implemented using the weighted mean trick is more flexible, however, it comes at a cost as there are more parameters to tune when penalizing multiple moments compared to tilted ERM of Li et al. (2021).

Weighted mean trick in practice. To apply the method in practice, the classical batch training algorithm must be extended to include an additional step, the weights calculation. Instead of directly calculating the average loss, the user will calculate the loss value for each sample in the batch and then use the expression from Theorem 3 to compute the weights and the weighted mean. The operations for computing weights are not part of the computation graph as weights must be treated as constants. This is achieved using the `detach()` function in Pytorch (Paszke et al., 2019) and `stop_gradient()` in Tensorflow (Abadi et al., 2015). The moments penalization factors λ_i are the hyper-parameters and are tuned in ascending order with respect to i. However, penalizing higher-order moments might affect the impact of the lower-order ones, and thus, it might require a few iterations to find the optimal combination. The implementation of this algorithm in Pytorch is available on GitHub.[1] Since the gradient of the weighted mean is the weighted gradient of the elements, this allows weights to control the impact of each sample on the model parameters. Of note, the theoretical results hold when switching from expectation to sample expectation, $\mathbb{E}_n[\ell] = \frac{1}{n}\sum_{i=1}^{n} \ell_i$.

Algorithm 1: Training with Moments Penalization

input :
- $f(x;\theta)$ – model to be trained
- $\{x_i, y_i\}_1^n$ – batch of training data
- $\{\lambda\}_1^k$ – penalization factors
- ℓ – loss function

while *stopping criteria not reached* **do**
  for $i \leftarrow 1$ **to** n **do**
    $\ell_i \leftarrow \mathcal{L}(f(x_i,\theta), y_i)$; `/* sample loss */`
  end
  $z \leftarrow \text{detach}(\ell)$; `/* detach from computational graph */`
  $w \leftarrow \left[\sum_{j=1}^{k} \lambda_j (z - \mathbb{E}_n[z])^{j-1}\right]_+$
  $\mathcal{L}_w \leftarrow \frac{1}{n}\sum_{i=1}^{n} w_i \ell_i$; `/* weighted mean loss */`
  $\theta \leftarrow \theta - \gamma \nabla_\theta \mathcal{L}_w$; `/* update model parameters */`
end

3 Robust Classification

In this section, we will explore the robustification of the cross-entropy function under label noise by bounding the loss values using negative variance penalization. In this case, the resulting objective will not be convex. However, for moderately small penalization factors the objective remains convex

[1] For this phase we submit the source code as part of supplementary materials to preserve anonymity, however, the final version will contain a link to our GitHub repository.

on almost the entirety of the domain and does not hinder the convergence. Moreover, we also develop a convex version of the variance penalized cross-entropy though for a minor price in performance.

The objective of robust classification is to learn an optimal classifier f^* that minimize the average loss for both clean and noisy data. Formally, $f^* = \operatorname{argmin} \mathbb{E}[\ell(f(x), \hat{y})]$ where $\hat{y}$ represents noisy labels. The noise considered in this paper corrupts the class labels with probability $\mathbb{P}(\hat{y}_i \neq y_i) = \eta$ or preserves it with probability $\mathbb{P}(\hat{y}_i = y_i) = 1 - \eta$, where $\eta \in [0, 1]$. We investigate two scenarios: when η does not depend on the class label (class independent or uniform noise), and when η depends on the class label (class dependent or asymmetric noise). In what follows, we show that using a negative variance penalization factor, $\lambda_2 < 0$, bounds the loss function. By bounding the loss function the impact of misclassification of noisy samples on the average loss will decrease, and thus reduce the noise impact on the model during training.

Ghosh et al. (2017) outlined the distribution independent sufficient conditions for a loss function to be robust under both, class dependent and class independent noise. Specifically, if the loss function satisfies the symmetry constraint:

$$\sum_{i=1}^{k} \ell(f(x), i) = C \tag{5}$$

where C is a constant and it holds $\forall f, \forall x \in \mathcal{X}$, then ℓ is noise tolerant for class independent noise when $\eta < \frac{k-1}{k}$ (see Theorem 1 of Ghosh et al. (2017)). Moreover, if the population risk of the optimum classifier $\mathbb{E}[\ell(f^*(x), \hat{y})] = 0$ then ℓ is also noise tolerant for class dependent noise. Among the commonly used losses, only the mean absolute error (MAE) satisfies the symmetry constraint. On the other end, the cross-entropy (CE) loss which is widely used for classification, not only does not satisfy the above constraint, but is also an unbounded loss making it extremely susceptible to noise. Specifically, as noise makes the predicted probability of the correct class approach zero, its cross-entropy loss will approach infinity.

Following the work of Ghosh et al. (2017), many studies (Zhang & Sabuncu, 2018; Feng et al., 2020; Wang et al., 2019b; Ma et al., 2020; Wang et al., 2019a) found that models trained with MAE struggle to converge and proposed novel losses that rely on boundedness to achieve robustness instead of the symmetry constraint to avoid convergence problems. Zhang & Sabuncu (2018) proposed a generalized cross-entropy loss parameterized by q that recovers CE loss for $q \to 0$ and MAE for $q = 1$. For intermediate values of q, the loss is bounded and trades robustness for convergence. Feng et al. (2020) found that the first two terms of the Taylor series expansion for the CE loss are the MAE and MSE, respectively. And suggested adjusting the number of terms in the Taylor expansion for the CE to balance noise robustness and convergence. Wang et al. (2019b) proposed using a two term loss, complimenting the CE loss with a secondary reverse cross-entropy term that satisfies the symmetry constraint. The research by Ma et al. (2020) extended the concept of two term loss combining a robust "active" loss and a robust "passive" loss.

In the following, we use the *weighted mean trick* to bound a loss function to improve its noise robustness. The next lemma specify the requirements for the weights:

Lemma 5. *Let ℓ represent unbounded loss values and W represent non-negative and bounded weights, then the product $W\ell$ is bounded if there exists a finite threshold ℓ_0 such that loss values greater than ℓ_0 have a corresponding weight of 0.*

Proof. For $\ell > \ell_0$ since the corresponding weights W are 0 the product $W\ell$ is also 0 and for $\ell \leq \ell_0$ since the weights are bounded the product $W\ell$ is also bounded and thus $0 \leq W\ell < B$ where B is a positive constant that depends on the weights. □

The simplest solution that satisfies the above requirement is the negative variance penalization, $\lambda_1 < 0$, as the next lemma shows:

Lemma 6. *For non-negative weights W computed using Theorem 3 (Moments Expansion) with $\lambda_0 > 0, \lambda_1 < 0$ and $\lambda_i = 0, \forall i > 1$, loss values ℓ greater or equal to $\ell_0 = \mathbb{E}[\ell] - \frac{\lambda_0}{\lambda_1}$ will have an associated weight of* 0.

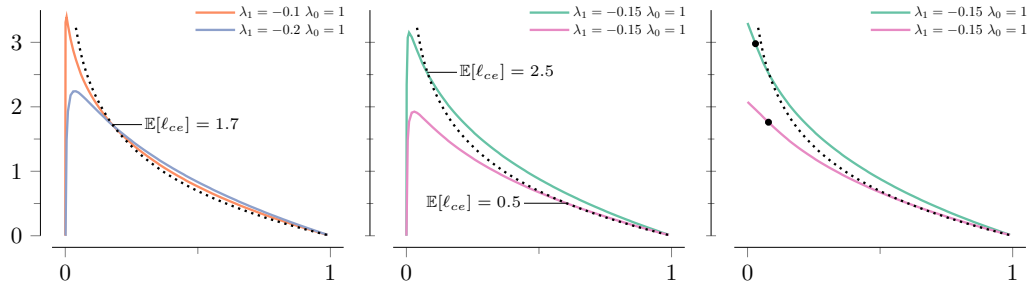

Figure 3: Bounding the cross-entropy loss function using negative variance penalization. The impact of the λ_1 parameter on the loss shape when the average loss is constant is shown on the left plot and the impact of the average loss on the shape in the center plot. Right plot shows the convex version of the variance penalized cross-entropy loss.

Proof. From Theorem 3 the weights are calculated as $W = \lambda_0 + \lambda_1(\ell - \mathbb{E}[\ell])$ from which we can determine the threshold value $\ell_0 = \mathbb{E}[\ell] - \frac{\lambda_0}{\lambda_1}$. Thus losses $\ell \geq \ell_0$ will have a 0 weight due to the non-negativity constraint. □

Of note, bounding an unbounded loss function is achieved through clipping, however, the resulting loss function is not convex as shown in Figure 3. With parameters λ_0 and λ_1 along with $\mathbb{E}[\ell]$ establishing which samples are considered noise and excluded from training. Since for $\lambda_1 < 0$ the threshold $\ell_0 = \mathbb{E}[\ell] - \frac{\lambda_0}{\lambda_1}$ is greater than $\mathbb{E}[\ell]$, thus all excluded samples have a loss value above average. Moreover, the higher the magnitude of λ_1, the closer the threshold ℓ_0 is to $\mathbb{E}[\ell]$. This can be seen in the left plot of Figure 3. Moreover, the magnitude of λ_1 dictates by how much the loss values below average are amplified and the ones above average suppressed. In practice, using larger values for λ_1 will decrease the impact of misclassified samples and those near the separating hyperplane on the placement of the decision boundary. This effect is shown in Figure 1 by the decision boundary $3\lambda_{min}$ which correspond to a negative variance penalization and which is aligned with the diagonal clusters far from the boundary. Lemma 6 can be extended to include other higher-order moments. In this case, the roots of the weights polynomial will determine which samples participate in the optimization problem and which are considered noise. For example, if the decision boundary should be decided by samples with average loss values then a possible solution is to penalize skewness. Since in this case the weights are defined by a quadratic equation, and using a negative value for λ_2 will make the parabola open downwards and thus assign non-zero weights only to samples around the mean.

Moreover, we also develop a convex version of the variance penalized cross-entropy by constraining the second derivative of the weighted loss function to be non-negative. We note that the second derivative of the weighted loss $w(x)\ell(x)$ reaches 0 for $x = x_c$ and $x_c = \exp\left(1 + \frac{\lambda_1}{2\lambda_2} - \frac{\mathbb{E}[\ell]}{2}\right)$ and is negative for $x < x_c$ where x is the predicted probability of the correct class. To constrain the second derivative be non-negative we linearly interpolate $w(x)\ell(x)$ for $x < x_c$ using the derivative at x_c. The weight function for $x < x_c$ is $w(x) = \left[w(x_c)\ell(x_c) + \nabla(w(u)\ell(u))|_{u=x_c}(x - x_c)\right]/\ell(x)$ and the resulting convex objective is shown on the right plot of Figure 3. The black dot marks the transition point from the weighted cross-entropy on the right to the linear interpolation on the left of the dot. Of note, the resulting objective is still bounded albeit with a higher upper bound.

In general, methods for training of neural networks in the presence of noise can be classified into two broad categories: noise model-free and noise model-based strategies. Noise model-free methods focus on reducing the impact of outliers and the two main subcategories are: robust losses (Bartlett et al., 2006; Ghosh et al., 2017; Wang et al., 2019a; Zhang & Sabuncu, 2018; Wang et al., 2019b; Natarajan et al., 2013; Mnih & Hinton, 2012; Xu et al., 2019; Patrini et al., 2016; Rooyen et al., 2015; Feng et al., 2020; Ma et al., 2020) and learning management such as meta-learning or regularization (Li et al., 2017; Szegedy et al., 2016; Hendrycks et al., 2019; Liu et al., 2020; Harutyunyan et al., 2020). On the other hand, noise model-based methods estimate the properties of the noise and use this information when train the model. Noise model-based methods can be further divided into: using

noise transition matrix (Patrini et al., 2017; Hendrycks et al., 2018; Chen & Gupta, 2015; Bekker & Goldberger, 2016; Goldberger & Ben-Reuven, 2017; Sukhbaatar et al., 2015; Xia et al., 2019; Yao et al., 2019; Wang et al., 2020; Xia et al., 2020; Yao et al., 2020; Lukasik et al., 2020), noise mitigation requiring a clean dataset (Jiang et al., 2018; Ren et al., 2018; Veit et al., 2017; Zhang et al., 2020; Yuan et al., 2018; Vahdat, 2017; Mirzasoleiman et al., 2020), or incrementally estimating the distribution of clean data (Liu et al., 2017; Zheng et al., 2020; Arazo et al., 2019; Zhang et al., 2017; Ghosh & Lan, 2021; Zhang et al., 2021; Wu et al., 2020).

Experimental results. The procedure adopted for all the experiments is similar to other studies investigating robust losses (Zhang & Sabuncu, 2018; Feng et al., 2020; Wang et al., 2019b; Ma et al., 2020). The source code to reproduce our results is available online.[1] After we corrupt the training datasets (CIFAR-10, CIFAR-100, Fashion-MNIST) with noise we retain 10% as a secondary validation dataset. This allows us to detect when the model is overfitting on noise by comparing the performance on clean versus noisy validation data. For the class independent noise, we flip the label to any other class with equal probability such that the ratio of noisy labels is η. For class dependent noise, for CIFAR-10 and Fashion-MNIST we only flip the label in the source classes $\{9, 2, 3, 5, 4\}$ to the corresponding target class $\{1, 0, 5, 3, 7\}$ given the noise ratio η. And for CIFAR-100 we flip between two randomly selected subclasses withing each superclass. We use the same modes as in Wang et al. (2019b) when training on CIFAR-10 which is an 8 layer network composed of 6 convolutional layers followed by 2 fully connected layers. For Fashion-MNIST the model contains 4 convolutional layers followed by 3 fully connected layers. For CIFAR-100 the model we use is ResNet-34 (He et al., 2016). We train using SGD with 0.9 momentum, 0.005 weight decay for the convolutional layers and 0.01 for the fully connected layers, and a starting learning rate of 0.01 which we divide by 10 every 20 epochs for a total of 60 training epochs for CIFAR datasets and every 5 epochs for a total of 15 training epochs for Fashion-MNIST. The training batch size is 128 samples.

We compare our solution for robustification of the loss function through moments penalization against five other state of the art methods: $i)$ Tilted empirical risk minimization (TERM) introduced by Li et al. (2021), $ii)$ Taylor expansion of cross entropy (Taylor) proposed by Feng et al. (2020), $iii)$ Normalized cross entropy coupled with reverse cross entropy (Normalized) investigated by Ma et al. (2020), $iv)$ Symmetric cross entropy (Symmetric) studied by Wang et al. (2019b), and $v)$ Generalized cross entropy (Generalized) explored by Zhang & Sabuncu (2018). For these methods, when possible we used the same parameters suggested by their authors, however, when the method underperfomed we used cross-validation to find better ones. When training on CIFAR-10 and Fashion-MNIST datasets we used: $\lambda_1 = 1$, $\lambda_2 = -0.5$ for Moments and its convex version, $t = -0.5$ for TERM, $t = 2$ for Taylor, $\alpha = 10$, $\beta = 1$ for Normalized, $\alpha = 0.1$, $\beta = 1$ for Symmetric, and $q = 0.7$ for Generalized. When training on CIFAR-100 the parameters where: $\lambda_1 = 1$, $\lambda_2 = -0.5$ for Moments and its convex version, $t = -0.5$ for TERM, $t = 6$ for Taylor, $\alpha = 10$, $\beta = 0.1$ for Normalized, $\alpha = 6.0$, $\beta = 0.1$ for Symmetric, and $q = 0.7$ for Generalized.

The model accuracy for clean data when training with both, class independent and class dependent noise, are summarized in Table 1. Of note, we report the accuracy results of the model at the end of the training to capture the overfitting on noise for some models. We observe that the training with moments penalization outperforms other methods in low to moderate noise level scenarios and falls behind in high noise scenarios but still anove the classical cross-entropy. The performances of the convex version and non-convex version are very similar, however, the convex version shows better results in high noise scenarios under class dependent noise.

In case of CIFAR datasets, for class independent noise and for $\eta = 0.2$ the proposed method improves the accuracy of the classical cross-entropy by 6%, the highest improvement among the investigated methods. However, this improvement decreases as the noise ratio increases and approaches the accuracy of the classical cross entropy loss for $\eta = 0.8$. To note, the convex version of the moments method for $\eta = 0.8$ underpeformed by 0.8% compared to the classical cross-entropy, however, in all other scenarios it had a higher accuracy. High noise scenarios turned out to be challenging for all losses with the Taylor and Generalized cross-entropy losses not converging on CIFAR-10 for $\eta = 0.8$ and the best accuracy was registered by Symmetric cross-entropy loss proposed by Wang et al. (2019b). For class dependent noise, the moments penalization registered the highest accuracy for CIFAR-10 when $\eta = 0.1$ and $\eta = 0.2$ but had the lowest accuracy for $\eta = 0.4$. To further

[1]For this phase we submit the source code as part of supplementary materials to preserve anonymity, however, the final version will contain a link to our GitHub repository.

Table 1: Mean accuracy (%) and standard deviation on clean data over 5 runs. The best result for each scenario is underlined.

	Loss	Noise rate η							
		Class independent				Class dependent			
		0.2	0.4	0.6	0.8	0.1	0.2	0.3	0.4
CIFAR-10	Classical	84.1 $_{0.2}$	77.1 $_{0.4}$	66.3 $_{0.4}$	36.0 $_{1.5}$	88.9 $_{0.2}$	87.2 $_{0.2}$	84.8 $_{0.4}$	<u>80.9 $_{0.7}$</u>
	Moments	<u>90.5 $_{0.1}$</u>	85.0 $_{0.6}$	69.1 $_{0.4}$	37.3 $_{2.3}$	<u>91.4 $_{0.2}$</u>	<u>89.9 $_{0.0}$</u>	85.0 $_{3.2}$	72.9 $_{3.3}$
	Moment-convex	86.8 $_{0.6}$	80.3 $_{0.3}$	67.6 $_{0.6}$	35.2 $_{1.5}$	89.1 $_{0.1}$	87.2 $_{0.2}$	83.5 $_{0.3}$	78.7 $_{0.4}$
	TERM	89.5 $_{0.1}$	83.2 $_{0.2}$	69.2 $_{0.5}$	37.2 $_{2.0}$	91.1 $_{0.2}$	88.1 $_{0.4}$	83.0 $_{0.7}$	77.2 $_{1.1}$
	Taylor	90.2 $_{0.2}$	86.2 $_{0.1}$	73.0 $_{0.5}$	10.0 $_{0.0}$	91.0 $_{0.2}$	88.4 $_{0.3}$	83.7 $_{0.4}$	78.0 $_{1.1}$
	Normalized	90.1 $_{0.1}$	<u>87.0 $_{0.2}$</u>	<u>80.8 $_{0.2}$</u>	32.6 $_{2.9}$	90.9 $_{0.1}$	89.7 $_{0.3}$	<u>86.8 $_{0.5}$</u>	79.8 $_{0.5}$
	Symmetric	90.0 $_{0.2}$	86.7 $_{0.2}$	80.2 $_{1.0}$	<u>44.8 $_{4.0}$</u>	90.8 $_{0.2}$	89.6 $_{0.1}$	86.5 $_{0.4}$	79.8 $_{0.9}$
	Generalized	90.4 $_{0.2}$	86.0 $_{0.2}$	67.3 $_{2.7}$	10.0 $_{0.0}$	91.3 $_{0.1}$	89.0 $_{0.1}$	83.7 $_{0.5}$	74.0 $_{3.2}$
CIFAR-100	Classical	40.0 $_{0.5}$	31.1 $_{1.1}$	20.7 $_{0.5}$	11.6 $_{0.4}$	46.8 $_{0.2}$	42.6 $_{0.4}$	37.2 $_{0.5}$	31.7 $_{0.3}$
	Moments	<u>46.1 $_{0.4}$</u>	39.8 $_{0.7}$	27.9 $_{0.8}$	13.8 $_{1.0}$	<u>47.8 $_{0.9}$</u>	43.7 $_{0.9}$	37.2 $_{0.4}$	30.6 $_{0.3}$
	Moment-convex	44.5 $_{0.8}$	36.7 $_{0.9}$	24.8 $_{0.8}$	13.4 $_{0.8}$	47.4 $_{0.5}$	42.9 $_{0.7}$	36.9 $_{0.5}$	31.1 $_{0.6}$
	TERM	45.5 $_{0.3}$	<u>41.5 $_{1.1}$</u>	<u>32.8 $_{1.4}$</u>	17.7 $_{0.4}$	47.1 $_{0.6}$	<u>44.6 $_{0.7}$</u>	37.9 $_{0.6}$	30.9 $_{0.6}$
	Taylor	37.3 $_{0.6}$	33.0 $_{1.3}$	27.2 $_{0.8}$	17.4 $_{0.2}$	38.7 $_{0.6}$	35.9 $_{0.7}$	31.0 $_{1.0}$	26.6 $_{0.5}$
	Normalized	32.9 $_{0.8}$	27.8 $_{0.7}$	22.7 $_{0.4}$	13.9 $_{0.2}$	35.0 $_{0.3}$	32.7 $_{0.3}$	30.2 $_{0.4}$	26.5 $_{0.6}$
	Symmetric	43.1 $_{0.4}$	37.9 $_{0.4}$	31.2 $_{0.8}$	<u>19.4 $_{0.5}$</u>	45.3 $_{0.5}$	43.1 $_{0.2}$	<u>41.0 $_{0.5}$</u>	<u>34.6 $_{0.3}$</u>
	Generalized	38.2 $_{0.4}$	33.7 $_{0.9}$	28.1 $_{0.6}$	18.5 $_{0.4}$	39.6 $_{0.9}$	37.9 $_{0.7}$	35.8 $_{0.6}$	29.9 $_{0.3}$
Fashion-MNIST	Classical	92.0 $_{0.1}$	90.9 $_{0.1}$	89.0 $_{0.2}$	78.4 $_{1.0}$	<u>92.9 $_{0.1}$</u>	<u>92.5 $_{0.3}$</u>	92.0 $_{0.7}$	90.6 $_{1.7}$
	Moments	92.1 $_{0.2}$	91.6 $_{0.2}$	89.8 $_{0.1}$	80.2 $_{1.1}$	87.0 $_{2.9}$	82.4 $_{6.7}$	79.3 $_{6.8}$	62.1 $_{3.0}$
	Moments-convex	<u>92.4 $_{0.1}$</u>	91.6 $_{0.1}$	89.6 $_{0.2}$	79.7 $_{1.0}$	92.6 $_{0.2}$	92.0 $_{0.4}$	<u>92.4 $_{0.2}$</u>	88.7 $_{1.5}$
	TERM	92.1 $_{0.1}$	91.6 $_{0.1}$	90.0 $_{0.1}$	80.3 $_{0.9}$	92.3 $_{0.1}$	92.0 $_{0.4}$	91.5 $_{0.9}$	91.3 $_{0.4}$
	Taylor	90.6 $_{0.0}$	89.4 $_{0.3}$	86.8 $_{0.2}$	75.7 $_{0.7}$	91.1 $_{0.1}$	90.7 $_{0.4}$	89.8 $_{1.4}$	85.1 $_{3.8}$
	Normalized	91.5 $_{0.1}$	90.8 $_{0.3}$	88.9 $_{0.3}$	79.8 $_{1.3}$	91.7 $_{0.1}$	91.7 $_{0.3}$	90.7 $_{0.7}$	91.1 $_{0.1}$
	Symmetric	92.3 $_{0.1}$	<u>91.8 $_{0.1}$</u>	<u>90.6 $_{0.3}$</u>	<u>82.8 $_{1.2}$</u>	92.5 $_{0.1}$	92.3 $_{0.3}$	91.9 $_{0.6}$	<u>91.5 $_{1.3}$</u>
	Generalized	91.2 $_{0.2}$	90.2 $_{0.2}$	88.1 $_{0.2}$	73.9 $_{1.1}$	91.6 $_{0.1}$	91.3 $_{0.3}$	90.7 $_{0.7}$	90.4 $_{0.3}$

investigate this behavior, we reran the experiments and monitored the accuracy for each individual class. The classical cross-entropy recorded similar accuracy as the robust methods on classes affected by noise and outperformed them on classes 6 and 8, both not targeted by class dependent noise. The poor performance of robust methods can be justified by using parameters tuned for low and moderate noise.

In case of Fashion-MNIST dataset, as similar behavior as with CIFAR datasets was observed. For low to moderate noise level scenarios the moments penalization methods had the highest or the second highest accuracy. Moreover, when the noise was class dependent the convex version outperformed the non-convex version by at least 6%.

4 Conclusion

The main goal of the current work was to investigate the optimization of a weighted mean and the flexibility it provides in enforcing properties such as robustness when training with noisy labels. In addition, we extended previous variance penalization methods to include higher-order moments while eliminating some of their limitations. One of the significant findings to emerge from this study was that we can control the final distribution of the loss values by penalizing higher-order moments. In particular, by enforcing the distribution of the cross-entropy to have a higher variance through negative variance penalization, we improved the models' accuracy when trained with noisy labels. Although this paper centers on classification problems, the framework can also control the bias-variance trade-off and can also be applicable to regression problems.

REFERENCES

Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. URL `https://www.tensorflow.org/`. Software available from tensorflow.org.

Eric Arazo, Diego Ortego, Paul Albert, Noel O'Connor, and Kevin McGuinness. Unsupervised label noise modeling and loss correction. In *ICML*, 2019.

Peter L Bartlett, Michael I Jordan, and Jon D McAuliffe. Convexity, classification, and risk bounds. *JASA*, 2006.

Alan Joseph Bekker and Jacob Goldberger. Training deep neural-networks based on unreliable labels. In *ICASSP*, 2016.

Xinlei Chen and Abhinav Gupta. Webly supervised learning of convolutional networks. In *ICCV*, 2015.

John Duchi and Hongseok Namkoong. Variance-based regularization with convex objectives. *The Journal of Machine Learning Research*, 20(1):2450–2504, 2019.

Lei Feng, Senlin Shu, Zhuoyi Lin, Fengmao Lv, Li Li, and Bo An. Can cross entropy loss be robust to label noise. In *IJCAI*, 2020.

Aritra Ghosh and Andrew Lan. Do we really need gold samples for sample weighting under label noise? In *WACV*, 2021.

Aritra Ghosh, Himanshu Kumar, and PS Sastry. Robust loss functions under label noise for deep neural networks. In *AAAI*, 2017.

J. Goldberger and E. Ben-Reuven. Training deep neural-networks using a noise adaptation layer. In *ICLR*, 2017.

Hrayr Harutyunyan, Kyle Reing, Greg Ver Steeg, and Aram Galstyan. Improving generalization by controlling label-noise information in neural network weights. In *ICML*, 2020.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.

Christina Heinze-Deml and Nicolai Meinshausen. Conditional variance penalties and domain shift robustness. *Machine Learning*, 110(2):303–348, 2021.

Dan Hendrycks, Mantas Mazeika, Duncan Wilson, and Kevin Gimpel. Using trusted data to train deep networks on labels corrupted by severe noise. In *NeurIPS*, 2018.

Dan Hendrycks, Kimin Lee, and Mantas Mazeika. Using pre-training can improve model robustness and uncertainty. In *ICML*, 2019.

Geoffrey E. Hinton, Nitish Srivastava, A. Krizhevsky, Ilya Sutskever, and R. Salakhutdinov. Improving neural networks by preventing co-adaptation of feature detectors. *NeurIPS*, 2012.

Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, pp. 13–30, 1963.

Weihua Hu, Gang Niu, Issei Sato, and Masashi Sugiyama. Does distributionally robust supervised learning give robust classifiers? In *International Conference on Machine Learning*, pp. 2029–2037. PMLR, 2018.

Lu Jiang, Zhengyuan Zhou, Thomas Leung, Li-Jia Li, and Li Fei-Fei. Mentornet: Learning data-driven curriculum for very deep neural networks on corrupted labels. In *ICML*, 2018.

Henry Lam. Recovering best statistical guarantees via the empirical divergence-based distributionally robust optimization. *Operations Research*, 67(4):1090–1105, 2019.

Tian Li, Ahmad Beirami, Maziar Sanjabi, and Virginia Smith. Tilted empirical risk minimization. *ICLR*, 2021.

Yuncheng Li, Jianchao Yang, Yale Song, Liangliang Cao, Jiebo Luo, and Li-Jia Li. Learning from noisy labels with distillation. In *ICCV*, 2017.

Sheng Liu, Jonathan Niles-Weed, Narges Razavian, and Carlos Fernandez-Granda. Early-learning regularization prevents memorization of noisy labels. In *NeurIPS*, 2020.

Xin Liu, Shaoxin Li, Meina Kan, Shiguang Shan, and Xilin Chen. Self-error-correcting convolutional neural network for learning with noisy labels. In *FG*, 2017.

Michal Lukasik, Srinadh Bhojanapalli, Aditya Menon, and Sanjiv Kumar. Does label smoothing mitigate label noise? In *ICML*, 2020.

Xingjun Ma, Hanxun Huang, Yisen Wang, Simone Romano, Sarah Erfani, and James Bailey. Normalized loss functions for deep learning with noisy labels. In *ICML*, 2020.

Andreas Maurer and Massimiliano Pontil. Empirical bernstein bounds and sample variance penalization. *Proc. Computational Learning Theory Conference (COLT)*, 2009.

Baharan Mirzasoleiman, Kaidi Cao, and Jure Leskovec. Coresets for robust training of deep neural networks against noisy labels. In *NeurIPS*, 2020.

Volodymyr Mnih and Geoffrey E Hinton. Learning to label aerial images from noisy data. In *ICML*, 2012.

Hongseok Namkoong and John C Duchi. Variance-based regularization with convex objectives. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.

Nagarajan Natarajan, Inderjit S Dhillon, Pradeep Ravikumar, and Ambuj Tewari. Learning with noisy labels. In *NIPS*, 2013.

Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, pp. 8024–8035. Curran Associates, Inc., 2019.

Giorgio Patrini, Frank Nielsen, Richard Nock, and Marcello Carioni. Loss factorization, weakly supervised learning and label noise robustness. In *ICML*, 2016.

Giorgio Patrini, Alessandro Rozza, Aditya Krishna Menon, Richard Nock, and Lizhen Qu. Making deep neural networks robust to label noise: A loss correction approach. In *CVPR*, 2017.

Mengye Ren, Wenyuan Zeng, Bin Yang, and Raquel Urtasun. Learning to reweight examples for robust deep learning. In *ICML*, 2018.

R Tyrrell Rockafellar and Stanislav Uryasev. Optimization of conditional value-at-risk. *Journal of Risk*, 2000.

Brendan van Rooyen, Aditya Krishna Menon, and Robert C Williamson. Learning with symmetric label noise: the importance of being unhinged. In *NeurIPS*, 2015.

Yevgeny Seldin, François Laviolette, Nicolo Cesa-Bianchi, John Shawe-Taylor, and Peter Auer. Pac-bayesian inequalities for martingales. *IEEE Transactions on Information Theory*, 58(12): 7086–7093, 2012.

Matthew Staib, Bryan Wilder, and Stefanie Jegelka. Distributionally robust submodular maximization. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 506–516. PMLR, 2019.

Sainbayar Sukhbaatar, Joan Bruna, Manohar Paluri, Lubomir Bourdev, and Rob Fergus. Training convolutional networks with noisy labels. In *ICLR*, 2015.

Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *CVPR*, 2016.

Ilya Tolstikhin and Yevgeny Seldin. Pac-bayes-empirical-bernstein inequality. *Advances in Neural Information Processing Systems 26 (NIPS 2013)*, pp. 1–9, 2013.

Arash Vahdat. Toward robustness against label noise in training deep discriminative neural networks. In *NeurIPS*, 2017.

Andreas Veit, Neil Alldrin, Gal Chechik, Ivan Krasin, Abhinav Gupta, and Serge Belongie. Learning from noisy large-scale datasets with minimal supervision. In *CVPR*, 2017.

X Wang, E Kodirov, Y Hua, and NM Robertson. Improved mean absolute error for learning meaningful patterns from abnormal training data. Technical report, Technical Report, 2019a.

Yisen Wang, Xingjun Ma, Zaiyi Chen, Yuan Luo, Jinfeng Yi, and James Bailey. Symmetric cross entropy for robust learning with noisy labels. In *ICCV*, 2019b.

Zhen Wang, Guosheng Hu, and Qinghua Hu. Training noise-robust deep neural networks via meta-learning. In *CVPR*, 2020.

Pengxiang Wu, Songzhu Zheng, Mayank Goswami, Dimitris Metaxas, and Chao Chen. A topological filter for learning with label noise. In *NeurIPS*, 2020.

Xiaobo Xia, Tongliang Liu, Nannan Wang, Bo Han, Chen Gong, Gang Niu, and Masashi Sugiyama. Are anchor points really indispensable in label-noise learning? In *NeurIPS*, 2019.

Xiaobo Xia, Tongliang Liu, Bo Han, Nannan Wang, Mingming Gong, Haifeng Liu, Gang Niu, Dacheng Tao, and Masashi Sugiyama. Part-dependent label noise: Towards instance-dependent label noise. In *NeurIPS*, 2020.

Yilun Xu, Peng Cao, Yuqing Kong, and Yizhou Wang. L_dmi: A novel information-theoretic loss function for training deep nets robust to label noise. In *NeurIPS*, 2019.

Jiangchao Yao, Hao Wu, Ya Zhang, Ivor W Tsang, and Jun Sun. Safeguarded dynamic label regression for noisy supervision. In *AAAI*, 2019.

Yu Yao, Tongliang Liu, Bo Han, Mingming Gong, Jiankang Deng, Gang Niu, and Masashi Sugiyama. Dual t: Reducing estimation error for transition matrix in label-noise learning. In *NeurIPS*, 2020.

Bodi Yuan, Jianyu Chen, Weidong Zhang, Hung-Shuo Tai, and Sara McMains. Iterative cross learning on noisy labels. In *WACV*, 2018.

Jing Zhang, Victor S Sheng, Tao Li, and Xindong Wu. Improving crowdsourced label quality using noise correction. *IEEE Trans. Neural Netw. Learn. Syst.*, 2017.

Yikai Zhang, Songzhu Zheng, Pengxiang Wu, Mayank Goswami, and Chao Chen. Learning with feature-dependent label noise: A progressive approach. In *ICLR*, 2021.

Zhilu Zhang and Mert R Sabuncu. Generalized cross entropy loss for training deep neural networks with noisy labels. In *NeurIPS*, 2018.

Zizhao Zhang, Han Zhang, Sercan O Arik, Honglak Lee, and Tomas Pfister. Distilling effective supervision from severe label noise. In *CVPR*, 2020.

Songzhu Zheng, Pengxiang Wu, Aman Goswami, Mayank Goswami, Dimitris Metaxas, and Chao Chen. Error-bounded correction of noisy labels. In *ICML*, 2020.

A PROOFS AND ADDITIONAL THEORETICAL RESULTS

First, we present a proposition introduced by Duchi & Namkoong (2019) and used in the proof of Theorem 1.

Proposition 7. *Let U and V be two random variables, and let $\overline{U}$ and $\overline{V}$ denote their expectation, then:*

$$\mathbb{E}\left[\left(U-\overline{U}\right)V\right] = \mathbb{E}\left[\left(U-\overline{U}\right)\left(V-\overline{V}\right)\right] \tag{6}$$

Proof. To prove we write V as $\left(V-\overline{V}\right)+\overline{V}$ and then use the linear property of the expectation operator:

$$\begin{aligned}\mathbb{E}\left[\left(U-\overline{U}\right)V\right] &= \mathbb{E}\left[\left(U-\overline{U}\right)\left(\left(V-\overline{V}\right)+\overline{V}\right)\right]\\ &= \mathbb{E}\left[\left(U-\overline{U}\right)\left(V-\overline{V}\right)\right]+\mathbb{E}\left[\left(U-\overline{U}\right)\overline{V}\right]\\ &= \mathbb{E}\left[\left(U-\overline{U}\right)\left(V-\overline{V}\right)\right]+\overline{V}\left(\mathbb{E}\left[U\right]-\overline{U}\right)\\ &= \mathbb{E}\left[\left(U-\overline{U}\right)\left(V-\overline{V}\right)\right]\end{aligned}$$

□

Next we present two propositions used in the proofs of the main theorems of this paper.

Proposition 8. *Let $\ell : \mathbb{R}^n \to \mathbb{R}_+$ and $p : \mathbb{R} \to \mathbb{R}$ two convex functions, then the composition $w(x) = p(\ell(x) - M)$ where $M \geq 0$ is convex if p is non decreasing.*

Proof. Using the convexity of ℓ for any $x, y \in \mathbb{R}^n$ and $\alpha \in [0, 1]$ we have:

$$\ell(\alpha x + (1-\alpha)y) \leq \alpha\ell(x) + (1-\alpha)\ell(y) \tag{7}$$

Using the fact that p is non decreasing along we the above inequality we obtain:

$$p(\ell(\alpha x + (1-\alpha)y)) \leq p(\alpha\ell(x) + (1-\alpha)\ell(y)) \tag{8}$$

and using convexity of p for the term on the right we get:

$$p(\alpha\ell(x) + (1-\alpha)\ell(y)) \leq \alpha p(\ell(x)) + (1-\alpha)p(\ell(y)) \tag{9}$$

Combining the two inequalities leads to the following result $(\ell(\alpha x + (1-\alpha)y)) \leq \alpha p(\ell(x)) + (1-\alpha)p(\ell(y))$ which proves the composition $p(\ell(x))$ is convex. Moreover, since p is convex and $\ell(x) - M$ is in the domain of p then $p(\ell(x) - M)$ is also convex as it is a composition with an affine mapping. This proves that $w(x)$ is convex.

□

Proposition 9. *Let $f, g : \mathbb{R}^n \to \mathbb{R}_+$ be two convex function taking non-negative values then their product $h(x) = f(x)g(x)$ is convex if $0 \leq [f(x) - f(y)][g(x) - g(y)], \forall x, y \in \mathbb{R}^n$.*

Proof. The function h is convex if it satisfies the following inequality where $\alpha \in [0, 1]$:

$$0 \leq \alpha h(x) + (1-\alpha)h(y) - h(\alpha x + (1-\alpha)y) \tag{10}$$

$$0 \leq \alpha(fg)(x) + (1-\alpha)(fg)(y) - (fg)(\alpha x + (1-\alpha)y) \tag{11}$$

Note that since the two functions are convex the following inequalities hold:

$$f(\alpha x + (1-\alpha)y) \leq \alpha f(x) + (1-\alpha)f(y) \tag{12}$$

$$g(\alpha x + (1-\alpha)y) \leq \alpha g(x) + (1-\alpha)g(y) \tag{13}$$

Given that the two functions take non-negative values we can multiply the above two inequalities.

$$f(\alpha x + (1-\alpha)y)g(\alpha x + (1-\alpha)y) \leq [\alpha f(x) + (1-\alpha)f(y)]\,[g(x) + (1-\alpha)g(y)] \tag{14}$$

$$(fg)(\alpha x + (1-\alpha)y) \leq [\alpha f(x) + (1-\alpha)f(y)]\,[g(x) + (1-\alpha)g(y)] \tag{15}$$

Substituting $(fg)(\alpha x + (1-\alpha)y)$ from 15 in inequality 10 we obtain:

$$0 \leq \alpha(fg)(x) + (1-\alpha)(fg)(y) - [\alpha f(x) + (1-\alpha)f(y)]\,[g(x) + (1-\alpha)g(y)] \tag{16}$$

Multiplying the two square brackets and then grouping the terms we get:

$$0 \leq \alpha(1-\alpha)\,[f(x) - f(y)]\,[g(x) - g(y)] \tag{17}$$

We complete the proof by noting that $\alpha \geq 0$ and thus the product of two convex functions is convex if $0 \leq [f(x) - f(y)]\,[g(x) - g(y)]$. □

Lemma 2. *The variance expansion using theorem 1 of a convex objective ℓ is also convex if $\lambda \in [0, \lambda_{max}]$, where $\lambda_{max} = 1/(\mathbb{E}[\ell] - \min(\ell))$.*

Proof. We note that $W = p(\ell(x) - \mathbb{E}[\ell])$ where p is a polynomial of degree 1 with $p(x) = \lambda x + 1$. First we use the result of Proposition 8 and get that p is convex if λ is non negative thus the lower bound of the interval. Next, we use Proposition 8 which states that $W(x)\ell(x) = p(\ell(x) - \mathbb{E}[\ell])\ell(x)$ is convex if $[w(x) - w(y)]\,[\ell(x) - \ell(y)] \geq 0$ and p is non negative. Since p is convex we replace p in the previous inequality with its the first order condition $p(a) - p(b) \geq \lambda(a - b)$ where $a = \ell(x) - M$ and $b = \ell(y) - M$ to obtain $\lambda[\ell(x) - \ell(y)]^2 \geq 0$. This is always true for $\lambda > 0$, however, $p(\ell(x) - \mathbb{E}[\ell])$ is non negative for $\lambda \leq \frac{1}{\min(\ell) - \mathbb{E}[\ell]}$ which proves the upper bound of the interval. □

Theorem 3 (Moments Expansion)**.** *Let ℓ be a random variable with finite first n central moments and $W = \sum_{i=1}^{n} \lambda_i(\ell - \mathbb{E}[\ell])^{i-1}$, then taking the weighted mean of ℓ with weights W is equivalent to:*

$$\mathbb{E}[W\ell] = \lambda_1\mathbb{E}[\ell] + \sum_{i=2}^{n} \tilde{\lambda}_i \mathbb{E}\left[(\ell - \mathbb{E}[\ell])^i\right] \tag{4}$$

where $\tilde{\lambda}_i = \lambda_i + \lambda_{i+1}\mathbb{E}[\ell]$ for $i < n$ and $\tilde{\lambda}_n = \lambda_n$

Proof. In this case, for $i \geq 2$ we cannot apply Lemma 7 since the expression $(\ell - \mathbb{E}[\ell])^i$ is not guaranteed to have zero mean and as a result it incurs an additional penalization of the previous moment. The proof follows the same steps as Theorem 1:

$$\begin{aligned}
\mathbb{E}[W\ell] &= \mathbb{E}\left[\sum_{i=1}^{n} \lambda_i(\ell - \mathbb{E}[\ell])^{i-1}\ell\right] \\
&= \sum_{i=1}^{n} \lambda_i \mathbb{E}\left[(\ell - \mathbb{E}[\ell])^{i-1}\ell\right] \\
&= \lambda_1\mathbb{E}[\ell] + \sum_{i=2}^{n} \lambda_i \mathbb{E}\left[(\ell - \mathbb{E}[\ell])^{i-1}\ell\right] \\
&= \lambda_1\mathbb{E}[\ell] + \sum_{i=2}^{n} \lambda_i \mathbb{E}\left[(\ell - \mathbb{E}[\ell])^{i-1}(\ell - \mathbb{E}[\ell] + \mathbb{E}[\ell])\right] \\
&= \lambda_1\mathbb{E}[\ell] + \sum_{i=2}^{n} \lambda_i \mathbb{E}\left[(\ell - \mathbb{E}[\ell])^{i}\right] + \sum_{i=2}^{n} \lambda_i \mathbb{E}[\ell]\mathbb{E}\left[(\ell - \mathbb{E}[\ell])^{i-1}\right] \\
&= \lambda_1\mathbb{E}[\ell] + \sum_{i=1}^{n-1} \tilde{\lambda}_i \mathbb{E}\left[(\ell - \mathbb{E}[\ell])^{i}\right] + \lambda_n \mathbb{E}\left[(\ell - \mathbb{E}[\ell])^{n}\right]
\end{aligned}$$

For the last step, we combine the two sums by matching the $\mathbb{E}\left[(\ell - \mathbb{E}[\ell])^i\right]$ terms and consolidate the penalization factors as $\tilde{\lambda}_i = \lambda_i + \lambda_{i+1}\mathbb{E}[\ell]$. □

Lemma 4 (Convexity of Moments Expansion)**.** *Let $\ell : \mathbb{R}^n \to \mathbb{R}_+$ and $p : \mathbb{R} \to \mathbb{R}_+$ be two non negative convex functions and $M \geq 0$, then the $W\ell(x)$ with $W = p(\ell(x) - M)$ is convex if p is non decreasing.*

Proof. We set $w(x) = p(\ell(x) - M)$ and using the result of Proposition 8 we obtain that $w(x)$ is convex. Next, from Proposition 9 we see that the product $w(x)\ell(x)$ is convex if $[w(x) - w(y)]\,[\ell(x) - \ell(y)] \geq 0$. Since p is convex we use the first order condition $p(a) - p(b) \geq p'(b)(a - b)$ and set $a = \ell(x) - M$ and $b = \ell(y) - M$ to obtain:

$$p(\ell(x) - M) - p(\ell(y) - M) \geq p'(\ell(y) - M)[\ell(x) - M - \ell(y) + M] \tag{18}$$

$$w(x) - w(y) \geq p'(\ell(y) - M)[\ell(x) - \ell(y)] \tag{19}$$

Substituting this result in the requirement from Proposition 9 we obtain:

$$p'(\ell(y) - M)[\ell(x) - \ell(y)]^2 \geq 0 \tag{20}$$

Given that p is non decreasing and $p'(\ell(y) - M) \geq 0$ proves that the inequality always holds. □