

Well-classified Examples are Underestimated in Classification with Deep Neural Networks

Guangxiang Zhao¹, Wenkai Yang², Xuancheng Ren¹, Lei Li¹, Xu Sun^{1,2}

¹ MOE Key Laboratory of Computational Linguistics, School of EECS, Peking University

² Center for Data Science, Peking University

{zhaoguangxiang, renxc, xusun}@pku.edu.cn, {wkyang, lilei}@stu.pku.edu.cn

Abstract

The conventional wisdom behind learning deep classification models is to focus on bad-classified examples and ignore well-classified examples that are far from the decision boundary. For instance, when training with cross-entropy loss, examples with higher likelihoods (i.e., well-classified examples) contribute smaller gradients in back-propagation. However, we theoretically show that this common practice hinders representation learning, energy optimization, and the growth of margin. To counteract this deficiency, we propose to reward well-classified examples with additive bonuses to revive their contribution to learning. This counterexample theoretically addresses these three issues. We empirically support this claim by directly verify the theoretical results or through the significant performance improvement with our counterexample on diverse tasks, including image classification, graph classification, and machine translation. Furthermore, this paper shows that because our idea can solve these three issues, we can deal with complex scenarios, such as imbalanced classification, OOD detection, and applications under adversarial attacks. Code is available at <https://github.com/lancopku/well-classified-examples-are-underestimated>.

1 Introduction

In common practice, classification with deep neural networks (DNNs) down-weights the contribution from well-classified examples. DNNs have achieved leading performance in mainstream classification tasks (He et al. 2016; Kipf and Welling 2016; Vaswani et al. 2017; Devlin et al. 2019). Usually, the training of DNNs relies on optimizing the designed metrics between the target and the prediction through back-propagation (Rumelhart, Hinton, and Williams 1986). Mean-Square Error (MSE) calculates a quadratic distance between the target and the probabilistic prediction of each example (Rumelhart, Hinton, and Williams 1986). Cross-Entropy (CE) loss measures the distance between the target distribution and the probability distribution (Baum and Wilczek 1988). CE loss is preferred as compared to MSE since CE loss encourages accurate predictions by bringing steep gradients to well-classified examples (Baum and Wilczek 1988; Goodfellow, Bengio, and Courville 2016). Therefore, CE loss shows better generalization ability due to the enlarged steepness (Solla, Levin, and Fleisher 1988). During training with CE loss, well-classified examples contribute

Preprint.

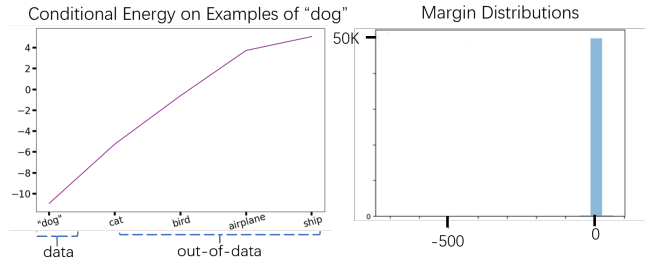


Figure 1: Illustration of energy and margins of CE loss on CIFAR-10. Left: averaged conditional energy $E(y|x)$ for examples of class "dog". Right: margin distributions of the trained classification models, there are some examples that have negative margins.

less to the gradient as compared to bad-classified ones. The wisdom behinds the operation of overlooking well-classified examples is well-classified examples have relatively less information in the learning. The improved variants of CE loss still comply with such wisdom (Liu et al. 2016; Szegedy et al. 2016; Lin et al. 2017).

We doubt the above practice with the following three facts: (1) Recent studies in imbalanced learning indicate that down-weighting the learning of relatively well-classified data-rich classes severely impairs the representation learning (Kang et al. 2020; Zhou et al. 2020). These studies inspire us to reflect on whether this is also the case at the sample level, and we validate that down-weighting the learning of well-classified samples also decreases the performance (Table 3). (2) As to Energy-Based Models (EBM) (LeCun et al. 2006), a sharper energy surface is desired.¹ However, we find that energy surface trained with CE loss is not sharp, as plotted in Figure 1. The possible reason is that CE loss has insufficient power to push down the energy of the positive examples as long as it is lower than the energy of negative examples. Our validation in Figure 5 shows that up-weighting the well-classified examples returns us a sharper surface. (3) As to classification, it is acknowledged that building classification mod-

¹Refer to the talk "The Future is Self-Supervised Learning" by Yann LeCun at ICLR 2020

els with large margins leads to good generalization (Bartlett 1997; Jiang et al. 2019) and good robustness (Elsayed et al. 2018; Matyasko and Chau 2017; Wu and Yu 2019), but we find that the learning with CE loss leads to smaller margins (as plotted Figure 1). The reason may be the limited incentive to further enlarge margins since well-classified examples are less optimized. Our results in Figure 6 and 7 show up-weighting classified samples enlarges the margin and helps to improve the adversarial robustness.

Contributions: We systematically study the role of well-classified examples in classification learning with DNNs, and the results challenge the common practice. First, we theoretically identify issues of CE loss with back-propagation in terms of the learning of representations, the learning of energy functions, and the growth speed of margins. (Refer to § 2). Second, we propose Encouraging Loss (EL), which can be viewed as a counterexample to the common practice since it up-weights the learning of well-classified examples compared to CE loss. Besides, we theoretically demonstrate that by paying attention to well-classified examples, learning with EL revives the representation learning from the part of well-classified examples, reduces the energy on the data manifold, and enlarges the classification margin. (Refer to § 3). Third, we conduct extensive experiments to empirically verify the effectiveness of paying more attention to the well-classified instances for training canonical and state-of-the-art models. Results on image recognition and machine translation based on heterogeneous data types show that the counterexample can consistently improve learning performance. We also empirically verify that well-classified examples regarding representation learning, energy, and margins. Further analysis shows that enhancing the learning of well-classified examples improves models’ capacity under complex scenarios, such as imbalanced classification, OOD detection, and application under adversarial attacks. (Refer to § 4).

2 Exploring Theoretical Issues of CE Loss

2.1 Setup and Notations

Classification In classification tasks, we have input data $\mathbf{x} \in \mathbb{R}^D$ and a label $y \in \mathbb{Y}$ which belongs to one of all K classes, and use $\mathbb{Y} = \{1, 2, \dots, K\}$ to denote the set of all class indices. We aim to learn a parametric function $f_\theta(\cdot) : \mathbb{Y}$ that predicts K -dim logits (before normalization) for the data \mathbf{x} , i.e., $f_\theta : \mathbb{R}^D \rightarrow \mathbb{R}^K$. Let $f_\theta(\mathbf{x})[y]$ denotes the y -th value of the predicted K -dim logits $f_\theta(\mathbf{x})[\mathbb{Y}]$, and $p_\theta(\mathbf{x})[y]$ is the normalized probability for the class y . We adopt the general softmax normalization which transforms the logit $f_\theta(\mathbf{x})[y]$ to the probability $p_\theta(\mathbf{x})[y]$ since it can generalize to both two-class and multi-class classification: $p_\theta(\mathbf{x})[y] = \text{softmax}(f_\theta(\mathbf{x})[y]) = \frac{\exp(f_\theta(\mathbf{x})[y])}{\sum_{y' \in \mathbb{Y}} \exp(f_\theta(\mathbf{x})[y'])}$.

CE loss Typically, the parametric classifier $f_\theta(\mathbf{x})[\mathbb{Y}]$ is estimated using Maximum Likelihood Estimation (MLE), which is equivalent to minimizing the Cross-Entropy (CE) between the predicted probability $p_\theta(\mathbf{x})[y]$ and the signal that whether the class y is the true label. When we get probabilities through the softmax normalization among all classes, we can simplify CE loss as minimizing negative log-likelihood (NLL) of $-\log p_\theta(\mathbf{x})[y]$, in which y is the true label class

for the input \mathbf{x} . The NLL loss is:

$$\mathcal{L}_{NLL} = -\log p_\theta(y | \mathbf{x}) = -\log p_\theta(\mathbf{x})[y]. \quad (1)$$

In Eq. (1), we get the predicted probability from the model θ and want to maximize the log probability of the target class y . We use the term **steepness of loss** to denote $\partial \mathcal{L} / \partial p$. For the NLL loss, the steepness of loss is $-\frac{1}{p}$, which means incorrect predicted examples with small p are learned with priority, i.e., embodying a sharper loss and a larger steepness of loss. In this section, we refer CE loss to NLL loss and discuss gradients regarding NLL loss and softmax normalization. However, our results can also easily generalize to BCE loss with sigmoid since the opposite class of BCE loss in case of the second class in NLL loss, and the derivative of the sigmoid is the same as softmax. In the NLL loss, we can only consider gradients from the index of the label class y . For simplicity, we use p to denote the correctness of predictions.

Back-propagation

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial \theta} &= \sum_{y' \in \mathbb{Y}} \frac{\partial \mathcal{L}}{\partial \sigma_\theta(\mathbf{x})[y']} \frac{\partial \sigma(f_\theta(\mathbf{x})[y'])}{\partial \theta} \\ &= \sum_{y' \in \mathbb{Y}} \frac{\partial \mathcal{L}}{\partial \sigma_\theta(\mathbf{x})[y']} \frac{\partial \sigma(f_\theta(\mathbf{x})[y'])}{\partial f_\theta(\mathbf{x})[y']} \frac{\partial f_\theta(\mathbf{x})[y']}{\partial \theta}. \end{aligned} \quad (2)$$

In Eq. (2), gradients depend on the loss function \mathcal{L} , logits normalization function σ and the current model f_θ .

2.2 Limitation of CE loss in Three Aspects

Normalization function brings a gradient vanishing problem to CE loss and hinders the representation learning At the beginning of the introduction of back-propagation to train deep neural networks, the loss function for back-propagation is MSE with measures the L_2 distance between probabilities and labels (Rumelhart, Hinton, and Williams 1986). However, the steepness of MSE gets to zero when the prediction gets close to the target. Baum and Wilczek (1988) introduce CE loss to back-propagation and point out it addresses the above issue in MSE and makes predictions more accurate. Combining the derivative of NLL loss and the derivative of the normalization, gradients for the model parameters with the CE loss are:

$$\frac{\partial \mathcal{L}_{NLL}}{\partial \theta} = (p - 1) \frac{\partial f_\theta(\mathbf{x})[y]}{\partial \theta}. \quad (3)$$

The normalization function brings the gradient saturation back as the prediction becomes correct. Since DNNs are regarded as a pipeline for extracting features (Zeiler and Fergus 2014; Tenney, Das, and Pavlick 2019), and it is in line with our perception that well-classified examples share pipeline with other examples, gradient vanishing hinders the part of representation learning from well-classified examples.

CE loss has insufficient power in reducing the energy on the data manifold Energy-based models (EBM) (LeCun et al. 2006) calculate the probability density of $p(x)$ for $x \in \mathbb{R}^D$ by: $p(x) = \frac{\exp(-E_\theta(x))}{\int_x \exp(-E_\theta(x))}$. Here, we re-interpret the classifier as a conditional EBM $p_\theta(y | \mathbf{x}) = \frac{\exp f_\theta(\mathbf{x})[y]}{\sum_{y' \in \mathbb{Y}} \exp(f_\theta(\mathbf{x})[y'])}$, in which the conditional energy is

$E_{\theta}(y | \mathbf{x}) = -f_{\theta}(\mathbf{x})[y]$. Thus, the CE loss in Eq. (1) can be written as:

$$\mathcal{L}_{NLL} = E_{\theta}(y | \mathbf{x}) + \log[\exp(-E_{\theta}(y | \mathbf{x})) + \sum_{y' \neq y} \exp(-E_{\theta}(y' | \mathbf{x}))]. \quad (4)$$

Minimizing the CE loss can push up the energy out of the data $E_{\theta}(y' | \mathbf{x}), y' \neq y$, but for the energy on the data manifold $E_{\theta}(y | \mathbf{x})$, although the nominator of the softmax function pulls that energy down, numerator which contains that term pushes the energy on the data up. Therefore, the learning with CE loss gets into a dilemma and has insufficient power to reduce the data's energy.

CE loss is not effective in enlarging margins Previous studies prove that large minimum margin (Bartlett 1997; Bartlett, Foster, and Telgarsky 2017; Neyshabur, Bhojanapalli, and Srebro 2018) or large overall margins (Zhang and Zhou 2017; Jiang et al. 2019) on the training set indicates good generalization ability. Though the margin $\gamma(\mathbf{x}, y) = f_{\theta}(\mathbf{x})[y] - \max_{y' \neq y} f_{\theta}(\mathbf{x})[y']$ is defined on the logits, since the softmax function generates probabilities against each other and has the exponential term that mimics the max operation, a larger likelihood is likely to lead to a larger margin. However, as we have deduced in the Eq. (3), CE loss is not good at increasing the likelihood when it becomes larger. Hence it is likely to be limited in increasing the margin. Following previous work (Wang et al. 2018; Cao et al. 2019; Menon et al. 2021), we view the gap $f_{\theta}(\mathbf{x})[y] - \max_{y' \neq y} f_{\theta}(\mathbf{x})[y']$ between the logit at the label position and the logit at any other position $y' \neq y$ as the approximated margin. Note that the NLL loss can then be written as $\mathcal{L}_{NLL} = \log[1 + \sum_{y' \neq y} \exp(f_{\theta}(\mathbf{x})[y'] - f_{\theta}(\mathbf{x})[y])]$. We use A to denote $\exp(f_{\theta}(\mathbf{x})[y'] - f_{\theta}(\mathbf{x})[y])$, the gradients of the NLL loss w.r.t. the parameter θ is:

$$\frac{\partial \mathcal{L}_{NLL}}{\partial \theta} = \frac{\sum_{y' \neq y} A \left(\frac{\partial (f_{\theta}(\mathbf{x})[y'] - f_{\theta}(\mathbf{x})[y])}{\partial \theta} \right)}{1 + \sum_{y' \neq y} A}. \quad (5)$$

The above formula interprets the training procedure of CE loss as increasing the logit for the label $f_{\theta}(\mathbf{x})[y]$, but decreasing the logits for other classes $f_{\theta}(\mathbf{x})[y']$, thus it enlarges the gap, and the standard margin $f_{\theta}(\mathbf{x})[y] - \max_{y' \neq y} f_{\theta}(\mathbf{x})[y']$ is likely to be larger. However, when the prediction gets close to the target during training, A gets close to 0, but the numerator has a constant 1, so the incentive for further enlarging the margin gets close to 0. Therefore, CE loss is not effective in enlarging margins to a certain extent.

3 Gaining from Reviving the Learning of Well-classified Examples

In this section, we propose a counterexample **Encouraging Loss** (EL) which increases the relative importance of well-classified examples in optimization compared to that in CE loss. We first define EL and then demonstrate that it can mitigate issues of CE loss mentioned above.

3.1 The Counterexample: Encouraging Loss

As plotted in the Figure 2, encouraging loss is the addition of CE loss and an additional loss (we term it as bonus) that



Figure 2: Illustration the counterexample the encouraging loss. CE loss: $-\log p$; Bonus: $\log(1 - p)$; Encouraging loss: CE loss + bonus. Bonus strengthens the learning of well-classified examples.

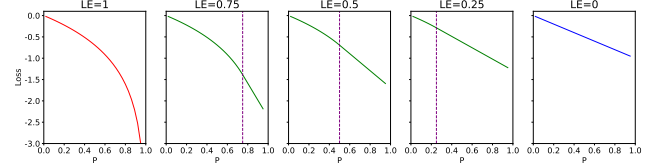


Figure 3: These are variations of bonuses. Left: normal bonus $\log(1 - p)$; Others: conservative bonus. LE is the position where Log curve Ends in conservative log bonus, loss curve continues after LE with the tangent at $p = LE$ of the log curve. These variations gradually increase steepness for well-classified examples from right to left. The original CE loss is the case $bonus = constant$.

makes the loss steeper again when p goes higher. The **normal bonus** is the mirror flip of the CE loss: $bonus = \log(1 - p)$, and the encouraging loss with normal bonus is:

$$\mathcal{L}_{EL} = -\log p_{\theta}(\mathbf{x})[y] + \log(1 - p_{\theta}(\mathbf{x})[y]) \quad (6)$$

We name it **encouraging loss** since it encourages the model to give more accurate predictions by rewarding these near-correct predictions with a bonus.

As long as the additional bonus is concave, its steepness for larger p is larger, indicating that the EL with that bonus pays more attention to the well-classified examples than the CE loss. To study the relative importance of the learning for well-classified examples compared to other examples, we can adjust the relative steepness of the additional bonus. We design many variants that approximate the normal bonus but are more conservative and show them in Figure 3. The relative importance of optimization for well-classified examples in these losses is larger than CE and gradually increases from the right to the left. The bonus can also be designed to be more aggressive than the normal bonus.

3.2 The counterexample can solve the identified issues in CE loss

Taking the normal bonus as a typical example, we analyze the benefits of encouraging loss.

Encouraging loss enhances the representation learning from the part of well-classified examples The steepness of the encouraging loss is $-\frac{1}{p} - \frac{1}{1-p}$, so the gradients are:

$$\frac{\partial \mathcal{L}_{EL}}{\partial \theta} = -1 \cdot \frac{\partial f_{\theta}(\mathbf{x})[y]}{\partial \theta}. \quad (7)$$

Contrast to gradients of NLL/CE loss in Eq. (3), here gradients are independent of the likelihood p (this is a bit like RELU in DNNs (Glorot, Bordes, and Bengio 2011)). As to the EL with conservative bonus, since the bonus is concave, gradients are also larger for well-classified examples.

Encouraging loss makes the model have smaller obstacle for reducing the energy on the data The conditional energy form of the EL is:

$$\mathcal{L}_{EL} = E_{\theta}(y | \mathbf{x}) - \log\left[\sum_{y' \neq y} \exp(-E_{\theta}(y' | \mathbf{x}))\right]. \quad (8)$$

The difference between it and the conditional energy form of the CE loss in Eq. (4) lies in the second term. Notice that training with EL does not need to push up the energy on the data to minimize the second term, so there is more incentive to lower the energy on the data. Although conservative bonus $< \log 1$ and $> \log(1 - p)$ does not remove the obstacle in the second term, the obstacle is smaller than CE loss.

Encouraging loss makes margins grow faster The margin perspective of gradients w.r.t EL is:

$$\frac{\partial \mathcal{L}_{EL}}{\partial \theta} = \frac{\sum_{y' \neq y} A \left(\frac{\partial (f_{\theta}(\mathbf{x})[y'] - f_{\theta}(\mathbf{x})[y])}{\partial \theta} \right)}{\sum_{y' \neq y} A}. \quad (9)$$

Now the growth speed for the margin is $\frac{1 + \sum_{y' \neq y} A}{\sum_{y' \neq y} A}$ times faster than that during the training with the CE loss. When the model is getting better, the exponent of negative gap A gets close to 0, the ratio becomes large and helps further increase the margin. The EL with a conservative bonus has a smaller ratio, but the ratio is still large than 1.

4 Practical Effect of Encouraging the Learning of Well-classified Examples

This section analyzes the practical effect of encouraging learning well-classified examples by applying the counterexample to various classification tasks and settings.

4.1 Experiment Setup

In here we briefly clarify the experiment setup, please refer to the Appendix B and the code for more details. For reliability, each result is the mean result of 5 different runs with error bars. Especially, on graph datasets, each run contains 50 different train, valid, test splits of the data (proportion is 0.8, 0.1, 0.1 respectively) since a recent study indicates that different dataset splits largely affect the test performance (Shchur et al. 2019). For other tasks, we use their official data splits.

Image Recognition It is a typical application of multi-class classification. In these tasks we need to predict the category an image belongs to. We adopt four tasks MNIST (LeCun and Cortes 2010), CIFAR-10, CIFAR-100 (Krizhevsky, Hinton et al. 2009), and ImageNet (Russakovsky et al. 2015), the descriptions of dataset are in Appendix B. For training, we borrow code from repositories with good reproduced accuracy and keep all their default settings unchanged. Specifically, we train the CNN model from Liu et al. (2016) on MNIST, train ResNet-50 (He et al. 2016) and EfficientNet-B0 (Tan and Le 2019) on CIFAR-10 and CIFAR-100 using

the code by Narumiruna, train the ResNet-50 on ImageNet with the example code from pytorch, train the EfficientNet-B0 on ImageNet using the code from timm (Wightman 2019). We choose ResNet-50 and EfficientNet-B0 because they are canonical and SoTA parameter-efficient models, respectively. For evaluation, we report the best top-1 accuracy on the test set following common practice.

Graph classification Typical applications of graph classification are to binary classify the functionality of the graph structured biological data, we do the experiments on PROTEINS (1113 graphs of protein structures) (Dobson and Doig 2003; Borgwardt et al. 2005) and NCI1 (4110 graphs of chemical compounds) (Wale, Watson, and Karypis 2008). The model is node feature learning model GCN (Kipf and Welling 2016) with the pooling method SAGPooling (Lee, Lee, and Kang 2019). We report test accuracy on the early stopping model with the best valid accuracy.

Machine translation In this task, we need to sequentially select one word class from the vocabulary, which contains tens of thousands of word classes. We perform experiments on IWSLT De-En (160K training sentence pairs) and IWSLT Fr-En (233K training sentence pairs). The base model is Transformer (Vaswani et al. 2017), and the evaluation metric is BLEU which calculates how many N -grams (N -gram is a contiguous sequence of N classification predictions) both exist in the predicted sequence and the generated sequence (Papineni et al. 2002). We adopt most settings from fairseq, including training and evaluation. The only modification is that we tune the best hyper-parameters for the default loss (CE loss with label smoothing) and then use them for training the model with encouraging loss for fair comparison.

Imbalanced classification We perform experiments on a large-scale natural imbalanced classification dataset—iNaturalist 2018, which has 8,142 classes, 438K training samples and 24K valid samples. The setting for training and evaluation is the same as Kang et al. (2020), including separately evaluate results on subsets of ‘many-shot’, ‘medium-shot’ and ‘few-shot’.

4.2 Improving Classification Performance

Image recognition of multi-class classification on the pixel data We plot the results in Figure 4 and summarize the improvements in Table 1. For EfficientNet-B0 on Imagenet, we directly run the baseline and encouraging loss with a conservative bonus (LE=0.75) to save energy since EfficientNet is not time-efficient. From all results, we point out that strengthening the learning of well-classified examples with an additional bonus can indeed help improve the accuracy. Take ImageNet as an example, by strengthening the learning of well-classified examples, the accuracy of the canonical model ResNet-50 is improved by 0.94 points, and the accuracy of the SOTA parameter efficient model EfficientNet-B0 can also be improved from 77.8 to 78.28. We want to point out that the improvement by EL is actually remarkable compared with other methods, considering that they either need 250x training data (Xie et al. 2020) or several times of epochs (Zhang et al. 2018) to achieve accuracy improvements similar to ours.

Binary graph classification on graph-structured data

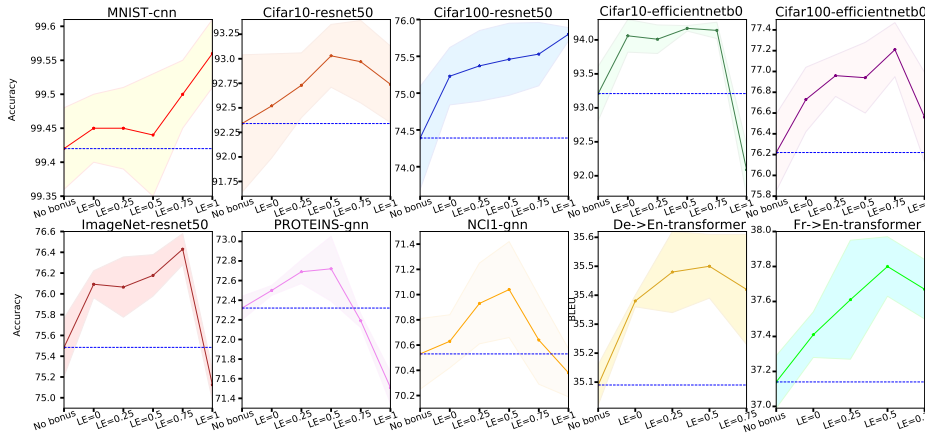


Figure 4: Enhancing the learning of well-classified examples can improve performance.

Setting	Cifar10-resnet50	Cifar100-resnet50
CE	92.34±0.70	74.39±0.70
CE (2xLR)	91.53±0.19	73.57±0.66
EL	92.74±0.38	75.80±0.09
EL (0.5xLR)	93.69±0.25	76.37±0.29

Table 2: Training with EL (LE=1) can benefit from reducing the amount of gradients, since the overall gradients of EL are larger than CE in theory.

We can see from the Figure 4 that strengthening the learning of well-classified examples with a conservative bonus can bring an increase in the accuracy of 0.44 on Proteins, and accuracy of 0.51 on NCI1.

Machine translation of sequential multi-class classification on the text data As we can see from the last two sub-figures of Figure 4, on machine translation, rewarding correct predictions can bring BLEU score improvement of 0.41 on De-En translation and 0.63 on Fr-En translation. We show in Appendix C.1 and Table 8 that our improvements are additive label smoothing (Szegedy et al. 2016) and on par with it.

Trade-off Results from Figure 4 indicate that in many conservative settings ($LE \leq 0.5$) where encouraging loss already pay more attention to well-classified examples than CE loss, the performance can be consistently improved. However, encouraging loss with a steeper bonus can not further improve the accuracy of deep classification models in some settings. The reason may be that CE loss actually implicitly decreases the learning rate along the training as the gradient norm of CE loss decays during training, and existing methods for optimization which adapt to CE loss should be modified. In our experiments, we choose to adopt the best setting in baselines for EL, which may not be the most suitable for encouraging loss. To verify this, we first show in Table 2 that when we enlarge the learning rate for CE loss, accuracy also drops. Then, we find re-normalizing the gradients of

Setting	MNIST	C10-r50	C10-eb0
CE	99.42±0.06	92.34±0.70	93.21±0.40
EL	99.56±0.05	92.97±0.42	94.24±0.17
Setting	C100-r50	C100-eb0	Img-r50
CE	74.39±0.70	76.22±0.37	75.49±0.28
EL	75.80±0.09	77.21±0.26	76.43±0.15
Setting	Img-eb0	Proteins	NCI1
CE	77.80±0.15	72.32±0.12	70.53±0.28
EL	78.28±0.13	72.76±0.18	71.04±0.38
Setting	De-En	Fr-En	
CE	35.09±0.07	37.10±0.06	
EL	35.50±0.11	37.73±0.17	

Table 1: We summarize the results of CE and EL on various tasks. Encouraging the Learning of well-classified examples brings consistent improvements. We abbreviated the names, e.g., Img refers to ImageNet.

encouraging loss by decreasing the global learning rate can help learn well-classified examples better. For example, we can continue improving the accuracy gap between CE loss and encouraging loss to 1.35 on Cifar10-resnet50 and 1.98 on Cifar100-resnet50 respectively.

Easy to apply Besides all promising results shown above, our method has another advantage that it can be widely applicable as a direct plug-in module without changing all the original hyper-parameters.

4.3 Addressing the Issues of CE Loss in Practice

In this subsection, we demonstrate that strengthening the learning of well-classified examples addresses the issues that CE loss faces as we discussed in §2.2. Due to the space limit, we compare results between CE loss and EL with a normal bonus for training ResNet-50 on CIFAR-10. Please refer to Appendix for more results in other settings.

Representation Learning Besides the overall accuracy improvement illustrated in §4.2, we also perform experiments to evaluate the learned representations by two losses. Specifically, we first use CE loss and EL to train the whole model, and then only train the re-initialized output layer with CE loss but fix the representations. The procedure is similar to Kang et al. (2020) that decouples the training into the representation learning phase and classification phase. We observe that representation learning with EL achieves an accuracy of 92.98 ± 0.01 , but representation learning with CE loss only gets an accuracy of 91.69 ± 0.04 .

Energy optimization We plot the conditional energy $E(y | x)$ averaged over the examples of class "cat" in Figure 5. We can see that the energy around the data becomes sharper with the help of encouraging loss since it pushes down the energy on the data more than CE loss.

Growth of margins We can see from Figure 6 that margins of encouraging loss are several times larger than margins of CE loss. These results demonstrate that learning well-classified examples with the additional bonus greatly improve the classification models by enlarging margins, which further makes the model have great generalization and robustness as

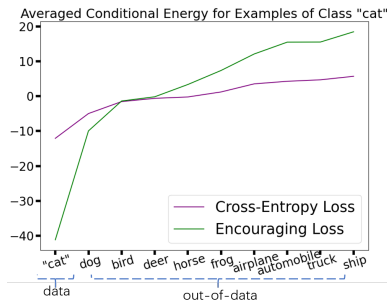


Figure 5: Conditional energy $E(y | x)$ averaged over the examples (X_{y^*}) of target class (here it is 'cat'), energy around the data ($y = y^*$) become sharper with the help of EL.

we will discuss in the following.

4.4 Coping with Complex Application Scenarios

This section shows that enhancing the learning of well-classified examples by EL can cope well with three complex application scenarios since this idea mitigates the three issues.

Imbalanced classification In imbalanced classification, we improve the classification performance of rare classes with the help of representation learning from samples of data-rich classes (Kang et al. 2020). Because enhancing the learning of easy samples also enhances their representation learning, we believe this property benefits imbalanced learning. We conduct validation experiments on the iNaturalist 2018 dataset, and the results are in Table 3. We find that encouraging the learning of well-classified samples makes models' performance outperform that trained with CE loss, both for the case with conservative bonus or the case with aggressive bonus (only to reward the highly well-classified samples of $p > 0.5$ with the normal bonus). We can also combine our idea with other advanced methods (Cao et al. 2019; Kang et al. 2020). For the additive experiment with "Decoupling classifier and representation learning" (Kang et al. 2020), in the representation learning phase, we not only remain the learning of data-rich classes as they do but also revive the learning of well-classified examples. In the classifier learning phase, we keep all the settings unchanged. Our method improves them by 1.9 points, **which empirically verify that the traditional re-weighting at the sample level (CE loss down-weights the importance of well-classified samples) is also harmful to the representation learning.**

OOD detection We can improve the performance of OOD detection since EL sharpens the energy surface around the data. For OOD detection, we use an indicator to detect whether samples are in the distribution (In-D) of the training set or out of the distribution (OOD). We use the minimum conditional energy $\min_{y' \in \mathcal{Y}} E_{\theta}(y' | x)$ as the indicator. We show in the Appendix that our indicator is better than the maximum probability and free energy (Liu et al. 2020). We conduct experiment on a small-scale OOD detection task (In-D: MNIST, OOD: Fashion-MNIST) and a large-scale task (In-

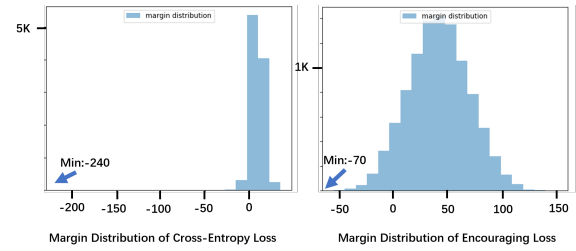


Figure 6: Left: Margins of CE loss distributed among -240 to 50, centered by 15. Right: Margins of EL distributed among -70 to 170, centered by 50. The learning of well-classified examples help double the margins in classification.

Method	Overall	iNaturalist2018		
		Many	Medium	Few
CE loss	64.3	74.1	65.9	59.8
+ Conservative Bonus (LE=0.5)	65.3	74.3	66.6	61.2
+ Normal Bonus	65.8	74.4	66.6	62.4
+ Aggressive Bonus	66.3 (+2.0)	75.1 (+1.0)	67.4(+1.6)	62.6(+2.8)
Decoupling Reps&Cls (CRT)	64.9	71.4	65.9	61.9
+ Normal Bonus	66.8(+1.9)	71.7(+0.3)	67.6(+1.7)	64.6(+2.7)
Deferred Re-weighting	68.1	71.0	68.3	67.1
+ Normal Bonus	70.3(+2.2)	69.0(-2.0)	70.1(+1.8)	70.9(+3.8)

Table 3: Comparison between CE loss and encouraging loss (add a bonus to CE loss) on imbalanced classification dataset iNaturalist 2018, the average standard deviation for results on iNaturalist 2018 is 0.4. Additional bonuses that revive the learning of well-classified examples bring improvements both on CE loss and its advanced methods.

D: ImageNet, OOD: iNaturalist2018). Samples in Fashion-MNIST are clothes, while samples in MNIST are digital numbers; samples in ImageNet are instances of objects, while samples in iNaturalist2018 are instances of species (refer to Appendix C.2 for detailed settings). We can see from Table 4 that EL leads to better performance than CE loss. For example, we reduce the metric FPR95 from 8.10% to 0.95% on MNIST vs. Fashion MNIST. These results confirm that strengthening the learning of well-classified examples leads to discriminative energy distribution.

Robustness to adversarial attacks We have shown in theory and practice that paying more attention to the classified samples can significantly increase margins. This property is likely to increase the robustness on adversarial exam-

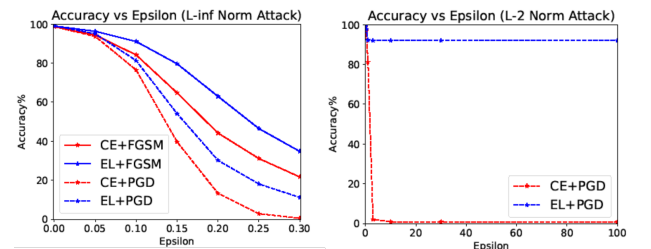


Figure 7: Accuracy of CE loss and encouraging loss under adversarial attacks of FGSM and PGD.

Setting Metric	MNIST vs. F. MNIST		Img vs. iNaturalist2018	
	AUROC↑	FPR95↓	AUROC↑	FPR95↓
CE	95.68	20.92	76.54	75.40
EL	98.04	8.69	78.41	71.86

Table 4: EL significantly improve OOD detection performance. We use F. MNIST to denote Fashion MNIST. Higher AUROC and lower FPR95 are better.

Name	Origin Loss	+Mirror Bonus
CE loss + softmax	99.42 ± 0.06	99.56 ± 0.05
MSE + softmax	99.56 ± 0.05	99.66 ± 0.04
MSE + sigmoid	99.63 ± 0.04	99.69 ± 0.02

Table 5: Well-classified examples also help improve MSE, which is better than CE loss on MNIST.

ples because the predictions of disturbed samples are likely to remain far away from the decision boundary. To verify this assumption, on MNIST, we use the FGSM (Goodfellow, Shlens, and Szegedy 2015), PGD (Madry et al. 2018) with L_∞ bound, and PGD with L_2 bound to construct adversarial samples to attack the models trained with CE loss and with EL, respectively. The results in Figure 7 validate our assumption that the model trained with EL is more robust on predictions when receiving a specific range of adversarial disturbances. We show in Appendix C.2 that our idea can also improve robustness under other powerful attack methods and is additive with large margin variants of CE loss to further enhance robustness and margins.

4.5 Well-classified Examples Are also Underestimated in MSE

It is well-known to the community that CE loss outperforms MSE for training deep neural networks, and the superiority has also been proven to come from the steeper gradient of CE (Solla, Levin, and Fleisher 1988). However, Table 5 shows that MSE beat CE loss on MNIST. The possible reason is that CE loss suffers from a much higher sensitivity to bad-classified examples. Nevertheless, we find that by adding a mirror bonus $-(y * p)^2 - ((1 - y) * (1 - p))^2$ to MSE $(y - p)^2$ to encourage the learning of well-classified examples, the performance improves. This indicates that, MSE also has the problem of underestimating well-classified examples.

4.6 The Idea is Additive with Improvements on CE Loss

We have show in the paper that our idea improves MSE, CE loss, and variations of CE loss in imbalanced classification. We also discuss previous improvements in CE loss in Appendix C.1, and demonstrate our idea is additive with including focal loss (Lin et al. 2017), label smoothing (Szegedy et al. 2016), adversarial training (Goodfellow, Shlens, and Szegedy 2015), and large margin softmax (Liu et al. 2016).

5 Limitations

To facility future research, we analyze the difficulties and the possible solutions in this new area. First, EL with the normal bonus can not improve the accuracy of deep classification models in some settings. One possible reason is that CE loss reduces the overall gradient norm, and existing optimization which adapt to CE loss should be modified. We show in Table 2 that reducing the learning rate when employing encouraging loss is a possible solution. Second, models trained with EL tend to give more confident predictions, and the Expected Calibration Error (ECE) (Guo et al. 2017) is slightly increased. Nevertheless, we can mitigate the issue by combining EL with label smoothing (Szegedy et al. 2016) to ECE. We show in Appendix C.2 that the combination of EL and label smoothing gets lower ECE than CE or CE with label smoothing.

6 Related work

Several studies are relevant to ours in different aspects.

Representation learning There are several studies that direly or indirectly mitigate the shortage of representation learning from well-classified examples. Relying on various augmentation techniques, contrastive learning is one of the methods which directly mitigates the issue and has made outstanding progress in learning representations (He et al. 2020; Chen et al. 2020; Grill et al. 2020). Kang et al. (2020) mitigate the representation learning issue caused by class-level re-weighting, which indirectly down-weight well-classified examples, by turning it off in representation learning. However, they do not improve CE loss in representation learning.

Using energy-based models to interpret the classifier Recent studies re-interpret the classifier as EBM (Grathwohl et al. 2020) or conditional EBM (Xie et al. 2016; Du and Mordatch 2019) for generative models. However, our focus is to separately investigate the energy on the data and the data optimized by CE loss in classification models.

Enlarging the margin A typical example of utilizing the idea of enlarging the minimum margin is the hinge loss (Suykens and Vandewalle 1999) that is proposed to help learn a considerate margin in SVMs (Cortes and Vapnik 1995) by focusing only on the data points close to the decision boundary. Recently, the idea of enlarging margins has been introduced to the case of CE loss (Liu et al. 2016; Wang et al. 2018; Li et al. 2018; Menon et al. 2021; Elsayed et al. 2018). Since these methods are based on logits adjustment before softmax, they can be combined with our method by substituting the original CE loss with EL loss after softmax. We show in Appendix C.1 that our idea can be combined with their large margin idea.

7 Conclusion

In this paper, we theoretically and empirically show that well-classified examples are very helpful for the further improvement of deep classification models. To illustrate this finding, we first directly analyze the failure of common practice, which weaken the learning of these examples, and then verify the positive effect of our proposed counterexamples, which value the learning of well-classified examples.

References

- Bartlett, P. 1997. For Valid Generalization the Size of the Weights is More Important than the Size of the Network. In Mozer, M. C.; Jordan, M.; and Petsche, T., eds., *Advances in Neural Information Processing Systems*, volume 9. MIT Press.
- Bartlett, P. L.; Foster, D. J.; and Telgarsky, M. J. 2017. Spectrally-normalized margin bounds for neural networks. In Guyon, I.; Luxburg, U. V.; Bengio, S.; Wallach, H.; Fergus, R.; Vishwanathan, S.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 30, 6240–6249. Curran Associates, Inc.
- Baum, E.; and Wilczek, F. 1988. Supervised Learning of Probability Distributions by Neural Networks. In Anderson, D., ed., *Neural Information Processing Systems*. American Institute of Physics.
- Borgwardt, K. M.; Ong, C. S.; Schönauer, S.; Vishwanathan, S.; Smola, A. J.; and Kriegel, H.-P. 2005. Protein function prediction via graph kernels. *Bioinformatics*, 21(suppl_1): i47–i56.
- Cao, K.; Wei, C.; Gaidon, A.; Arechiga, N.; and Ma, T. 2019. Learning Imbalanced Datasets with Label-Distribution-Aware Margin Loss. In Wallach, H.; Larochelle, H.; Beygelzimer, A.; d'Alché-Buc, F.; Fox, E.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020. A Simple Framework for Contrastive Learning of Visual Representations. In III, H. D.; and Singh, A., eds., *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, 1597–1607. PMLR.
- Cortes, C.; and Vapnik, V. 1995. Support-vector networks. *Machine learning*, 20(3): 273–297.
- Croce, F.; and Hein, M. 2020. Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks. In III, H. D.; and Singh, A., eds., *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, 2206–2216. PMLR.
- Cubuk, E. D.; Zoph, B.; Shlens, J.; and Le, Q. 2020. RandAugment: Practical Automated Data Augmentation with a Reduced Search Space. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M. F.; and Lin, H., eds., *Advances in Neural Information Processing Systems*, volume 33, 18613–18624. Curran Associates, Inc.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 4171–4186. Minneapolis, Minnesota: Association for Computational Linguistics.
- Dobson, P. D.; and Doig, A. J. 2003. Distinguishing enzyme structures from non-enzymes without alignments. *Journal of molecular biology*, 330(4): 771–783.
- Du, Y.; and Mordatch, I. 2019. Implicit Generation and Modeling with Energy Based Models. In Wallach, H.; Larochelle, H.; Beygelzimer, A.; d'Alché-Buc, F.; Fox, E.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Elsayed, G.; Krishnan, D.; Mobahi, H.; Regan, K.; and Bengio, S. 2018. Large Margin Deep Networks for Classification. In Bengio, S.; Wallach, H.; Larochelle, H.; Grauman, K.; Cesa-Bianchi, N.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc.
- Fey, M.; and Lenssen, J. E. 2019. Fast Graph Representation Learning with PyTorch Geometric. In *ICLR Workshop on Representation Learning on Graphs and Manifolds*.
- Glorot, X.; Bordes, A.; and Bengio, Y. 2011. Deep Sparse Rectifier Neural Networks. In Gordon, G.; Dunson, D.; and Dudík, M., eds., *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, volume 15 of *Proceedings of Machine Learning Research*, 315–323. Fort Lauderdale, FL, USA: PMLR.
- Goodfellow, I.; Bengio, Y.; and Courville, A. 2016. *Deep Learning*. MIT Press. <http://www.deeplearningbook.org>.
- Goodfellow, I. J.; Shlens, J.; and Szegedy, C. 2015. Explaining and Harnessing Adversarial Examples. arXiv:1412.6572.
- Grathwohl, W.; Wang, K.-C.; Jacobsen, J.-H.; Duvenaud, D.; Norouzi, M.; and Swersky, K. 2020. Your classifier is secretly an energy based model and you should treat it like one. In *International Conference on Learning Representations*.
- Grill, J.-B.; Strub, F.; Althé, F.; Tallec, C.; Richemond, P.; Buchatskaya, E.; Doersch, C.; Avila Pires, B.; Guo, Z.; Gheshlaghi Azar, M.; Piot, B.; kavukcuoglu, k.; Munos, R.; and Valko, M. 2020. Bootstrap Your Own Latent - A New Approach to Self-Supervised Learning. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M. F.; and Lin, H., eds., *Advances in Neural Information Processing Systems*, volume 33, 21271–21284. Curran Associates, Inc.
- Guo, C.; Pleiss, G.; Sun, Y.; and Weinberger, K. Q. 2017. On Calibration of Modern Neural Networks. In Precup, D.; and Teh, Y. W., eds., *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, 1321–1330. PMLR.
- He, K.; Fan, H.; Wu, Y.; Xie, S.; and Girshick, R. 2020. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9729–9738.
- 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.; Mu, N.; Cubuk, E. D.; Zoph, B.; Gilmer, J.; and Lakshminarayanan, B. 2019. AugMix: A Simple Data Processing Method to Improve Robustness and Uncertainty. In *International Conference on Learning Representations*.
- Jiang, Y.; Krishnan, D.; Mobahi, H.; and Bengio, S. 2019. Predicting the Generalization Gap in Deep Networks with Margin Distributions. In *International Conference on Learning Representations*.

- Jiang, Z.; Hou, Q.; Yuan, L.; Zhou, D.; Jin, X.; Wang, A.; and Feng, J. 2021. Token Labeling: Training a 85.4 with 56M Parameters on ImageNet. *arXiv:2104.10858*.
- Kang, B.; Xie, S.; Rohrbach, M.; Yan, Z.; Gordo, A.; Feng, J.; and Kalantidis, Y. 2020. Decoupling Representation and Classifier for Long-Tailed Recognition. In *International Conference on Learning Representations*.
- Kiefer, J.; Wolfowitz, J.; et al. 1952. Stochastic estimation of the maximum of a regression function. *The Annals of Mathematical Statistics*, 23(3): 462–466.
- Kim, H. 2020. Torchattacks: A pytorch repository for adversarial attacks. *arXiv preprint arXiv:2010.01950*.
- Kingma, D. P.; and Ba, J. 2015. Adam: A Method for Stochastic Optimization. In *ICLR (Poster)*.
- Kipf, T. N.; and Welling, M. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*.
- Kipf, T. N.; and Welling, M. 2017. Semi-Supervised Classification with Graph Convolutional Networks. *arXiv:1609.02907*.
- Krizhevsky, A.; Hinton, G.; et al. 2009. Learning multiple layers of features from tiny images.
- LeCun, Y.; Chopra, S.; Hadsell, R.; Ranzato, M.; and Huang, F. 2006. A tutorial on energy-based learning. *Predicting structured data*, 1(0).
- LeCun, Y.; and Cortes, C. 2010. MNIST handwritten digit database.
- Lee, J.; Lee, I.; and Kang, J. 2019. Self-attention graph pooling. In *International Conference on Machine Learning*, 3734–3743. PMLR.
- Li, Y.; Gao, F.; Ou, Z.; and Sun, J. 2018. Angular softmax loss for end-to-end speaker verification. In *2018 11th International Symposium on Chinese Spoken Language Processing (ISCSLP)*, 190–194. IEEE.
- Lin, T.; Goyal, P.; Girshick, R. B.; He, K.; and Dollár, P. 2017. Focal Loss for Dense Object Detection. In *IEEE International Conference on Computer Vision*, 2999–3007.
- Liu, W.; Wang, X.; Owens, J.; and Li, Y. 2020. Energy-based Out-of-distribution Detection. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M. F.; and Lin, H., eds., *Advances in Neural Information Processing Systems*, volume 33, 21464–21475. Curran Associates, Inc.
- Liu, W.; Wen, Y.; Yu, Z.; and Yang, M. 2016. Large-margin softmax loss for convolutional neural networks. In *ICML*, volume 2, 7.
- Madry, A.; Makelov, A.; Schmidt, L.; Tsipras, D.; and Vladu, A. 2018. Towards Deep Learning Models Resistant to Adversarial Attacks. In *International Conference on Learning Representations*.
- Matyasko, A.; and Chau, L.-P. 2017. Margin maximization for robust classification using deep learning. In *2017 International Joint Conference on Neural Networks (IJCNN)*, 300–307.
- Menon, A. K.; Jayasumana, S.; Rawat, A. S.; Jain, H.; Veit, A.; and Kumar, S. 2021. Long-tail learning via logit adjustments. In *International Conference on Learning Representations*.
- Neyshabur, B.; Bhojanapalli, S.; and Srebro, N. 2018. A PAC-Bayesian Approach to Spectrally-Normalized Margin Bounds for Neural Networks. In *International Conference on Learning Representations*.
- Ott, M.; Edunov, S.; Baevski, A.; Fan, A.; Gross, S.; Ng, N.; Grangier, D.; and Auli, M. 2019. fairseq: A Fast, Extensible Toolkit for Sequence Modeling. In *Proceedings of NAACL-HLT 2019: Demonstrations*.
- Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W.-J. 2002. Bleu: a Method for Automatic Evaluation of Machine Translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, 311–318. Philadelphia, Pennsylvania, USA: Association for Computational Linguistics.
- Rumelhart, D. E.; Hinton, G. E.; and Williams, R. J. 1986. Learning representations by back-propagating errors. *nature*, 323(6088): 533–536.
- Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M. S.; Berg, A. C.; and Li, F. 2015. ImageNet Large Scale Visual Recognition Challenge. *Int. J. Comput. Vis.*, 115(3): 211–252.
- Shchur, O.; Mumme, M.; Bojchevski, A.; and Günnemann, S. 2019. Pitfalls of Graph Neural Network Evaluation. *arXiv:1811.05868*.
- Solla, S. A.; Levin, E.; and Fleisher, M. 1988. Accelerated Learning in Layered Neural Networks. *Complex Syst.*, 2(6): 625–639.
- Sutton, R. S.; Barto, A. G.; et al. 1998. *Introduction to reinforcement learning*, volume 135. MIT press Cambridge.
- Suykens, J. A.; and Vandewalle, J. 1999. Least squares support vector machine classifiers. *Neural processing letters*, 9(3): 293–300.
- Szegedy, C.; Vanhoucke, V.; Ioffe, S.; Shlens, J.; and Wojna, Z. 2016. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2818–2826.
- Tan, M.; and Le, Q. 2019. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International Conference on Machine Learning*, 6105–6114. PMLR.
- Tenney, I.; Das, D.; and Pavlick, E. 2019. BERT Rediscovered the Classical NLP Pipeline. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 4593–4601. Florence, Italy: Association for Computational Linguistics.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L. u.; and Polosukhin, I. 2017. Attention is All you Need. In Guyon, I.; Luxburg, U. V.; Bengio, S.; Wallach, H.; Fergus, R.; Vishwanathan, S.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 30, 5998–6008. Curran Associates, Inc.

Wale, N.; Watson, I. A.; and Karypis, G. 2008. Comparison of descriptor spaces for chemical compound retrieval and classification. *Knowledge and Information Systems*, 14(3): 347–375.

Wang, F.; Cheng, J.; Liu, W.; and Liu, H. 2018. Additive Margin Softmax for Face Verification. *IEEE Signal Processing Letters*, 25(7): 926–930.

Wightman, R. 2019. PyTorch Image Models. <https://github.com/rwightman/pytorch-image-models>.

Williams, R. J. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4): 229–256.

Wu, K.; and Yu, Y. 2019. Understanding Adversarial Robustness: The Trade-off between Minimum and Average Margin. arXiv:1907.11780.

Xie, J.; Lu, Y.; Zhu, S.-C.; and Wu, Y. 2016. A Theory of Generative ConvNet. In Balcan, M. F.; and Weinberger, K. Q., eds., *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, 2635–2644. New York, New York, USA: PMLR.

Xie, Q.; Luong, M.-T.; Hovy, E.; and Le, Q. V. 2020. Self-Training With Noisy Student Improves ImageNet Classification. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 10684–10695.

Zeiler, M. D.; and Fergus, R. 2014. Visualizing and understanding convolutional networks. In *European conference on computer vision*, 818–833. Springer.

Zhang, H.; Cisse, M.; Dauphin, Y. N.; and Lopez-Paz, D. 2018. mixup: Beyond Empirical Risk Minimization. In *International Conference on Learning Representations*.

Zhang, T.; and Zhou, Z.-H. 2017. Multi-Class Optimal Margin Distribution Machine. In Precup, D.; and Teh, Y. W., eds., *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, 4063–4071. PMLR.

Zhou, B.; Cui, Q.; Wei, X.-S.; and Chen, Z.-M. 2020. Bbn: Bilateral-branch network with cumulative learning for long-tailed visual recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9719–9728.

A Complete Calculation of Equations

Recall that the NLL loss is:

$$\mathcal{L}_{NLL} = -\log p_{\theta}(y | \mathbf{x}) = -\log p_{\theta}(\mathbf{x})[y]$$

In the following paragraphs, we simply use p to denote the likelihood of the target $p_{\theta}(\mathbf{x})[y]$.

In Eq. (3), we calculate gradients for the model parameters with the CE loss, its complete calculation is:

$$\begin{aligned} \frac{\partial \mathcal{L}_{NLL}}{\partial \theta} &= \frac{\partial \mathcal{L}_{NLL}}{\partial p} \frac{\partial \sigma(f_{\theta}(\mathbf{x})[y])}{\partial f_{\theta}(\mathbf{x})[y]} \frac{\partial f_{\theta}(\mathbf{x})[y]}{\partial \theta} \\ &= -\frac{1}{p} [p(1-p)] \frac{\partial f_{\theta}(\mathbf{x})[y]}{\partial \theta} = (p-1) \frac{\partial f_{\theta}(\mathbf{x})[y]}{\partial \theta}. \end{aligned}$$

In Eq. (4), we re-write the CE loss, its complete calculation is:

$$\begin{aligned} \mathcal{L}_{NLL} &= -f_{\theta}(\mathbf{x})[y] \\ &+ \log[\exp(f_{\theta}(\mathbf{x})[y]) + \sum_{y' \neq y} \exp(f_{\theta}(\mathbf{x})[y'])] \\ &= E_{\theta}(y | \mathbf{x}) + \log[\exp(-E_{\theta}(y | \mathbf{x})) \\ &+ \sum_{y' \neq y} \exp(-E_{\theta}(y' | \mathbf{x}))], \end{aligned}$$

in which the we regard the negative of the logit as the conditional energy: $E_{\theta}(y | \mathbf{x}) = -f_{\theta}(\mathbf{x})[y]$.

Note that from the view of margin, the NLL loss can also be written as $\mathcal{L}_{NLL} = \log[\frac{\exp(f_{\theta}(\mathbf{x})[y])}{\exp(f_{\theta}(\mathbf{x})[y])} + \frac{\sum_{y' \neq y} \exp(f_{\theta}(\mathbf{x})[y'])}{\exp(f_{\theta}(\mathbf{x})[y])}] = \log[1 + \sum_{y' \neq y} \exp(f_{\theta}(\mathbf{x})[y'] - f_{\theta}(\mathbf{x})[y])]$.

In Eq. (5), we use A to denote $\exp(f_{\theta}(\mathbf{x})[y'] - f_{\theta}(\mathbf{x})[y])$, and calculate the gradients of the NLL loss w.r.t. the parameter θ :

$$\frac{\partial \mathcal{L}_{NLL}}{\partial \theta} = \frac{\sum_{y' \neq y} A (\frac{\partial (f_{\theta}(\mathbf{x})[y'] - f_{\theta}(\mathbf{x})[y])}{\partial \theta})}{1 + \sum_{y' \neq y} A}$$

The normal bonus is the mirror flip of the CE loss, it is $bonus = \log(1 - p)$, and the Encouraging Loss (EL) with normal bonus is:

$$\mathcal{L}_{EL} = -\log p_{\theta}(\mathbf{x})[y] + \log(1 - p_{\theta}(\mathbf{x})[y])$$

In Eq. (7), we calculate the gradients of the encouraging loss with normal bonus, since the steepness $\frac{\partial \mathcal{L}}{\partial p}$ of this loss is $-\frac{1}{p} - \frac{1}{1-p}$, its complete calculation is:

$$\begin{aligned} \frac{\partial \mathcal{L}_{EL}}{\partial \theta} &= \frac{\partial \mathcal{L}_{EL}}{\partial p} \frac{\partial \sigma(f_{\theta}(\mathbf{x})[y])}{\partial f_{\theta}(\mathbf{x})[y]} \frac{\partial f_{\theta}(\mathbf{x})[y]}{\partial \theta} \\ &= (-\frac{1}{p} - \frac{1}{1-p}) [p(1-p)] \frac{\partial f_{\theta}(\mathbf{x})[y]}{\partial \theta} \\ &= -(1-p+p) \cdot \frac{\partial f_{\theta}(\mathbf{x})[y]}{\partial \theta} = -1 \cdot \frac{\partial f_{\theta}(\mathbf{x})[y]}{\partial \theta} \end{aligned}$$

As to the EL with conservative bonus, since the bonus is concave, gradients are also larger for well-classified examples.

In Eq. (8), we write the conditional energy form of the EL, its complete calculation is:

$$\begin{aligned} \mathcal{L}_{EL} &= -\log p_{\theta}(\mathbf{x})[y] + \log(1 - p_{\theta}(\mathbf{x})[y]) \\ &= -f_{\theta}(\mathbf{x})[y] + \log[\sum_{y' \in \mathcal{Y}} \exp(f_{\theta}(\mathbf{x})[y'])] \\ &+ \log[\sum_{y' \neq y} \exp(f_{\theta}(\mathbf{x})[y'])] - \log[\sum_{y' \in \mathcal{Y}} \exp(f_{\theta}(\mathbf{x})[y'])] \\ &= -f_{\theta}(\mathbf{x})[y] + \log[\sum_{y' \neq y} \exp(f_{\theta}(\mathbf{x})[y'])] \\ &= E_{\theta}(y | \mathbf{x}) - \log[\sum_{y' \neq y} \exp(-E_{\theta}(y' | \mathbf{x}))] \end{aligned}$$

The difference between Eq. (8) and the conditional energy form of the CE loss in Eq. (4) lies in the second term, notice

now we do not need to push up energy on the data to minimize the second term, so there is more incentive to lower the energy on the data. Since the conservative bonus is $< \log 1 = \log[\sum_{y' \in \mathbb{Y}} \exp(f_{\theta}(\mathbf{x})[y'])]$ and $> \log(1 - p) = \log[\sum_{y' \neq y} \exp(f_{\theta}(\mathbf{x})[y'])]$ and $\log[\sum_{y' \in \mathbb{Y}} \exp(f_{\theta}(\mathbf{x})[y'])]$, although it does not remove the obstacle in the second term, the obstacle is smaller than CE loss.

Note that from the view of margin, the EL loss with normal bonus can also be written as $\mathcal{L}_{EL} = \frac{\sum_{y' \neq y} \exp(f_{\theta}(\mathbf{x})[y'])}{\exp(f_{\theta}(\mathbf{x})[y])}$ $\log[\sum_{y' \neq y} \exp(f_{\theta}(\mathbf{x})[y']) - f_{\theta}(\mathbf{x})[y]]$.

In Eq. (9), we calculate the gradients w.r.t EL from the view of margin:

$$\frac{\partial \mathcal{L}_{EL}}{\partial \theta} = \frac{\sum_{y' \neq y} A \left(\frac{\partial (f_{\theta}(\mathbf{x})[y'] - f_{\theta}(\mathbf{x})[y])}{\partial \theta} \right)}{\sum_{y' \neq y} A}$$

Growth speed for the margin is ratio $\frac{1 + \sum_{y' \neq y} A}{\sum_{y' \neq y} A}$ times faster than that during the training with the CE loss. When the model is getting better, the exponent of negative gap $A = \exp(f_{\theta}(\mathbf{x})[y'] - f_{\theta}(\mathbf{x})[y])$ gets close to 0, the ratio becomes large and helps further increase the margin. The EL with a conservative bonus has a smaller ratio, but the ratio is larger than 1 since the denominator of the ratio is larger.

B Details of the Experiment Setups

In this section, we describe the details of the experiments. Please refer to the code for reproduction. We provide the code in the supplementary material.

B.1 Experiment Setups for Main Results

In this subsection, we give details of experiment setups for main results and other experiments with the same settings.

For the results in image recognition, machine translation, imbalanced classification, each result is the mean result of 5 different runs with error bars to ensure reliable results. For graph classification, each run contains 50 different train, valid, test splits of the data (proportion is 0.8, 0.1, 0.1 respectively) since a recent study indicates that different dataset splits largely affect the test performance (Shchur et al. 2019). For other tasks, we use their official data splits.

We perform all experiments on the commonly used datasets and get data from their public release. Except for the graph data set, which ended training through the early exit, our method’s number of training rounds is the same as the baseline for other data sets.

Image recognition It is a typical application of multi-class classification. In these tasks, we need to predict the category an image belongs to. In this subsection, we mainly discuss tasks of balanced image recognition in which every class has almost the same number of samples.

For datasets, we adopt four tasks MNIST (LeCun and Cortes 2010)², CIFAR-10 (Krizhevsky, Hinton et al.

²MNIST data: <http://yann.lecun.com/exdb/mnist/>

2009)³(It’s license is unknown in paperwithcode), CIFAR-100, and ImageNet (Russakovsky et al. 2015)⁴, they have 10, 10, 100 and 1000 classes respectively and 60K, 50K, 50K, 1.2M training samples respectively.

For models, we borrow the code from several repositories with good reproduced accuracy. Specifically, we adopt the CNN model from Liu et al. (2016) on MNIST, ResNet-50 (He et al. 2016) and EfficientNet-B0 (Tan and Le 2019) on CIFAR-10 and CIFAR-100 using the code by Narumiruna⁵, we also ResNet-50 on ImageNet with the code from the official PyTorch examples for imagenet⁶, train the EfficientNet-B0 on ImageNet using the code from timm (Wightman 2019) since this code-base publicly releases scripts that can replicate a good performance for Efficientnet-B0. We choose ResNet-50 and EfficientNet-B0 because they are the common-used model and the SoTA parameter-efficient model, respectively.

We keep all the default hyper-parameters unchanged and consistent with that in the source repositories for preprocessing, training, and evaluation. Some of these hyper-parameters are: For training, we use RMSProp⁷ to minimize the CE loss or the encouraging loss for EfficientNet-B0 on ImageNet since RMSProp is the default optimizer for training EfficientNet on ImageNet, and use SGD (Kiefer, Wolfowitz et al. 1952) for other settings. The learning rate is 0.1 with the batch size of 256 where we follow the common practice, and readers can find the learning rate schedule in the code.

Graph Classification Typical applications of graph classification are to binary classify the functionality of the graph structured biological data, we do the experiments on PROTEINS (1113 graphs of protein structures) (Dobson and Doig 2003; Borgwardt et al. 2005) and NCI1 (4110 graphs of chemical compounds) (Wale, Watson, and Karypis 2008). The model is the combination of the Graph Convolution (GCN) (Kipf and Welling 2017) model for node feature learning and the pooling method SAGPooling (Lee, Lee, and Kang 2019) for selecting representative nodes in each graph. The details about the architecture can be found in the paper of SAGPooling (Lee, Lee, and Kang 2019) and our code. Our code and hyper-parameter settings are borrowed from pytorch-geometric (Fey and Lenssen 2019), and we do not change the hyper-parameters there. Specifically, for training, we adopt the default optimizer Adam (Kingma and Ba 2015) with a batch size of 128 and a learning rate of 5×10^{-4} . For evaluation, we report test accuracy on the early stopping model with the best valid accuracy. The patience is set to 50 following SAGPooling (Lee, Lee, and Kang 2019). The best epochs on the valid set are 35 and 140 respectively, so 85 and 190 epochs are executed.

³CIFAR-10 / 100 data: <https://www.cs.toronto.edu/~kriz/cifar.html>

⁴ImageNet data: <https://image-net.org/challenges/LSVRC/2012/>

⁵Environments of experiments on CIFAR 10/100: <https://github.com/narumiruna/efficientnet-pytorch>

⁶Environments of experiments of ResNet-50 on ImageNet: <https://github.com/pytorch/examples/tree/master/imagenet>

⁷RMSprop is an unpublished, adaptive learning rate method proposed by Geoffrey Hinton in Lecture 6e of his Coursera Class

Machine Translation In this task, we need to sequentially select one word class from the vocabulary which contains tens of thousands of word classes at each position. We perform experiments on IWSLT De-En⁸ (160K training sentence pairs, the target language has a vocabulary of 6632 words) and IWSLT Fr-En⁹ (233K training sentence pairs, the target language has a vocabulary of 16240 words). The model is Transformer (Vaswani et al. 2017). We adopt most settings from fairseq (Ott et al. 2019), including preprocessing, training, and evaluation. The only difference is that we tune the best hyper-parameters for the default loss (It is the CE loss with label smoothing.) and then use them to train with encouraging loss. Specifically, we use Adam optimizer, the highest learning rate is $1e - 3$, on De-En translation, we train the model for 20K updates with each batch has 16K tokens, on Fr-En translation, we train the model for 100K updates with each batch has 4K tokens. The number of updates is determined on the valid set, and the criteria is that the last 10 epochs have similar valid performance. For evaluation, the metric is BLEU which calculates how many N -grams (N -gram is a contiguous sequence of N classification predictions) both exist in the predicted sequence and the generated sequence (Papineni et al. 2002), we use tokenized BLEU with the default tokenize method. Following common practice (Vaswani et al. 2017), we report the BLEU of the averaged checkpoint over checkpoints saved in the last 10 epochs.

B.2 Experiment Setups for Verifying the Properties

In Sections 4.3, we empirically verify the theoretical results that well-classified examples help representation learning, energy optimization, and the growth of margins. Since the last two are straightforward, we describe the details for the first verification.

Representation Learning in Balanced Classification In Section 4.3 of the main paper, we report the representation learning performance on balanced classification. Specifically, we first use CE loss and EL to train the whole model and then only train the re-initialized output layer with CE loss but fix the representations. The procedure is similar to Kang et al. (2020) that decouples the training into the representation learning phase and classification phase. The additional stage of retraining the classifier took 20 epochs, starting from 0.1 and decreasing by 0.1 times every 5 epochs. We selected these hyper-parameters because they are the most straightforward hyper-parameter setting from our perspectives. In our initial experiment, there were similar results in other settings.

B.3 Experiment Setups for Complex Application Scenarios

Imbalanced Classification We also show that well-classified examples help in representation learning on imbalanced classification. We perform experiments on a large-scale

⁸De-En translation data: We get the data from <https://workshop2014.iwslt.org/>

⁹Fr-En translation data: We get the data from <https://workshop2017.iwslt.org/>

natural imbalanced classification dataset—iNaturalist 2018¹⁰, which has 8,142 classes, 438K training samples and 24K valid samples. The setting for training and evaluation is the same as Kang et al. (2020), including using the SGD optimizer with batch size 512 and learning rate 0.2. We reported the separately evaluated results on subsets of 'many-shot' (each class has more than 100 samples), 'medium-shot' (each class has more than 20 but less than 100 samples), and 'few-shot' (each class has less than 20 samples). The number of overall training epochs is 200. As to the method "Deferred Re-weighting" (Cao et al. 2019), this baseline method up-weights the learning of the rare classes and down-weights the learning of the common classes to create a relatively balanced training distribution in the latter stages of training. We defer the training with this class-balanced re-weighting method after 180 epochs of regular training. As to the method, "Decoupling Reps&Cls (CRT)" (Kang et al. 2020), this baseline method first trains the model with the regular training, and then fix representations but only re-initializes and re-train the classifier with a re-sampling strategy that up-weights the learning of rare classes and down-weights the learning of common classes for additional 30 epochs.

OOD detection We can improve the performance of OOD detection since EL sharpens the energy surface around the data. For OOD detection, we use an indicator to detect whether samples are in the distribution (In-D) of the training set or out of the distribution (OOD). We use the minimum conditional energy $\min_{y' \in \mathbb{Y}} E_{\theta}(y' | \mathbf{x})$ as the indicator. We conduct experiment on a small-scale OOD detection task (In-D: MNIST, OOD: Fashion-MNIST) and a large-scale task (In-D: ImageNet, OOD: iNaturalist2018). Samples in Fashion-MNIST are clothes (10 classes), while samples in MNIST are digital numbers (from 0 to 9); samples in ImageNet are instances of 1,000 objects, while samples in iNaturalist2018 are instances of 8,142 species. Samples in Fashion-MNIST and MNIST are Gray-scale pictures, but samples in ImageNet and iNaturalist2018 are RGB pictures. MNIST and Fashion-MNIST both have 60,000 training examples. ImageNet and iNaturalist2018 have 1.2M and 437,513 training examples respectively.

Adversarial Robustness In Section 4.4 of the main paper, to verify the claim that well-classified examples help improve robustness, we use FGSM (Goodfellow, Shlens, and Szegedy 2015) and PGD (Madry et al. 2018) to construct adversarial samples to attack the models trained with CE loss and the models trained with encouraging loss, respectively. In the Appendix, we also test robustness under Auto Attack (Croce and Hein 2020). We perform experiments on MNIST, the implementations of these three methods and the base model are from TorchAttacks (Kim 2020). We adopt their settings for training and evaluation on MNIST since they get 99% accuracy on MNIST with significantly less training time than (Liu et al. 2016), but (Liu et al. 2016) get higher accuracy with 60x epochs. To report improvements on strong baselines, we take the setting of (Liu et al. 2016) for experiments in

¹⁰iNaturalist 2018 data: https://github.com/visipedia/inat_comp/tree/master/2018

Table 1 and Table 5 in the main paper. The short training epochs of TorchAttacks can save much time in adversarial training too since we generate inputs in every step. Thus, to be consistent with experiments in adversarial training, we adopt the setting from TorchAttacks.

B.4 Experiment Setups for New Experiments in the Appendix

Object Detection For experiments on object detection, we perform experiments on the common-used COCO dataset¹¹. we adopt the configuration of “RetinaNet-R-50-FPN-1x” from the github repository Detectron2¹² as our default setting. In this setting, the one-stage detector of RetinaNet (Lin et al. 2017) with the backbone ResNet50 is trained for 90k updates, and each batch has 8 images.

Adversarial Training In the Appendix, we demonstrate our idea is additive to Adversarial Training. We first train the model with the input images generated by adversarial attack methods iteratively and then test the model with corresponding attack methods. The implementation for adversarial training is from TorchAttack (Kim 2020).

B.5 Training Costs

The training costs are summarized in Table 6, and we estimate that the total cost for this study is running a GPU of RTX 2080TI for 300 days.

C Supplementary Results

C.1 Well-classified examples are also underestimated in the improved variants of CE loss

Adversarial Training We also perform experiments on adversarial training with FGSM (Linf), PGD(Linf). We do not do that for AutoAttack since it is too slow. Following the script in “implementation from ”[https://github.com/Harry24k/adversarial-attacks-pytorch/demos/Adversarial Training \(MNIST\).ipynb](https://github.com/Harry24k/adversarial-attacks-pytorch/demos/Adversarial Training (MNIST).ipynb)”, we generate adversarial training examples at every step and use them for training the model. The results in Table 7 show that EL is additive with adversarial training. We note that EL does not further benefit from adversarial training under the attack of L2-norm PGD. The reason may be that models trained with EL and clean data are already robust to L2-norm PGD.

We plot margin distributions on MNIST in the above three adversarial scenarios and demonstrate EL can also enlarge margin in these scenarios. We can see from Figure 8 that EL increases the margin for most samples compared to CE, even with adversarial training.

Large margin softmax On MNIST, we add encouraging loss with large margin softmax (Liu et al. 2016), the accuracy is improved from 99.62 ± 0.04 to 99.64 ± 0.02 . Since they do not provide code for other tasks and do not work for other

¹¹COCO detection data: <https://cocodataset.org>

¹²Environment of detection experiments:<https://github.com/facebookresearch/detectron2>

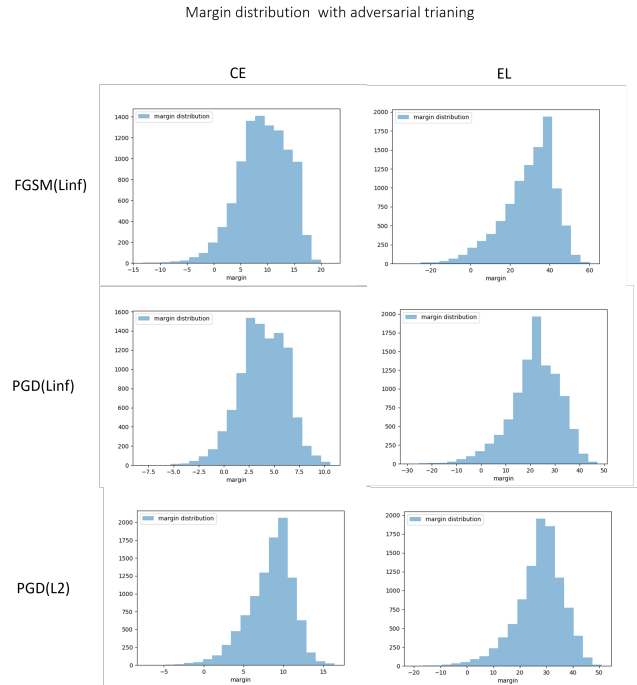


Figure 8: Margin distribution under adversarial training.

tasks with the provided code. The accuracy under attack is improved from 92.89 to 93.81 with epsilon(perturbation amount) 0.2, and 82.68 to 86.02 with epsilon 0.3. In addition, methods of increasing margins do not show the improvement on ImageNet, but we make it, as shown in the main paper.

Label Smoothing The translation experiments are with label smoothing. We show that our method is also additive to pure CE loss, and the improvements are on par with label smoothing. We show these results in Table 8.

Focal loss Focal Loss (FL) (Lin et al. 2017) propose that well-classified examples should be more overlooked than CE loss in imbalanced classification, its form is:

$$\mathcal{L} = -(1 - p_y)^{\gamma_f} \cdot (\mathbf{y} \cdot \log \mathbf{p}), \quad (10)$$

We propose a variant of the focal loss, the Halted Focal Loss (HFL), such that the high-likelihood area is not prioritized. The HFL reverts the focal loss to the cross-entropy Loss when the likelihood is high enough:

$$\mathcal{L} = \begin{cases} -(1 - p_y)^{\gamma_f} \cdot (\mathbf{y} \cdot \log \mathbf{p}), & \text{if } p_y \leq \varphi \\ -\mathbf{y} \cdot [\log \mathbf{p} + b], & \text{otherwise} \end{cases} \quad (11)$$

Where p_y is the prediction probability of the correct label, $b = \alpha(1 - (1 - \varphi)^{\gamma_f}) \log \varphi$ to ensure monotonicity and continuity, and φ is the boundary between the low-likelihood and the high-likelihood area, which we set as 0.5, i.e., a likelihood higher than which renders a definitely well-classified example. This halted focal loss is plotted in Figure 9, which has the same gradients as the cross-entropy in the high-likelihood area and remains the same as the Focal Loss in the low-likelihood area.

Data	Infrastructure	Mem/GPU	Time	Epochs	Samples	Model
MNIST	RTX 2080Ti * 1	3G	0.3h	100	60K	CNN-Arch (Liu et al. 2016)
CIFAR-10	RTX 2080Ti * 1	7G	5h	90	50K	ResNet50
CIFAR-10	RTX 2080Ti * 1	5G	5h	90	50K	EfficientNet-B0
CIFAR-100	RTX 2080Ti * 1	7G	5h	90	50K	ResNet50
CIFAR-100	RTX 2080Ti * 1	5G	5h	90	50K	EfficientNet-B0
ImageNet	TITAN RTX * 2	13G	48h	90	1.2M	ResNet50
ImageNet	TITAN RTX * 2	11G	158h	450	1.2M	EfficientNet-B0
PROTEINS	RTX 2080Ti * 1	0.7G	27s	85	1113	GCN+SAGPool
NCII	RTX 2080Ti * 1	0.7G	260s	190	4110	GCN+SAGPool
Fr-En translation	RTX 2080Ti * 1	4G	4h	-	233K	Transformer
De-En translation	RTX 2080Ti * 1	5G	4h	-	160K	Transformer
iNaturalist 2018	TITAN RTX * 4	15G	48h	200	438K	ResNet50
COCO Detection	RTX 2080Ti * 4	9G	7h	-	118K	ResNet50 +RetinaNet

Table 6: Training costs of each setting. Samples are sentences pairs in translation, images in image classification tasks, and graphs in graph classification. The training steps are calculated in the level of updates on Fr-En translation, De-EN translation, and COCO detection. For Fr-En translation, we train the model for $100K$ steps and each batch has $4K$ tokens. For De-En translation, we train the model for $20K$ steps and each batch has $16K$ tokens. For COCO detection, we train the model for $90K$ steps and each batch has 32 images.

Epsilon	FGSM(Linf)		PGD(Linf)		PGD (L2)		
	CE	EL	CE	EL	Epsilon	CE	EL
No_attack	95.43	97.21	94.3	96.88	No_attack	98.97	98.98
0.05	92.59	95.81	92.32	95.77	0.3	98.38	98.4
0.1	91.01	95.18	89.89	94.48	1	93.88	94.36
0.15	89.79	95.16	87.07	93.03	3	10.83	76.22
0.2	89.3	95.33	84.51	91.79	10	0.01	76.22
0.25	89.09	95.09	82.17	90.84	30	0.01	76.22
0.3	87.87	94.15	79.72	90.03	100	0.01	76.22

Table 7: Adversarial robustness after adversarial training under the attack of FGSM and PGD (Linf and L2). For generating adversarial training examples, the epsilon is 0.3 for two Linf norm attack methods and 3 for the PGD with L2 norm attacks.

On imbalanced classification, we find that the correct predictions matter for improving accuracy. Results are shown in Table 9 in terms of experiments on iNaturalist 2018 and in Table 10 on COCO detection.

C.2 Other Results

EL Reduces Excess Risk EL reduces the gap between train acc and test acc by 2 points, 3 points, and 2.5 points for training ResNet50 on CIFAR-10, CIFAR-100, ImageNet, respectively.

EL with Conservative Bonus also Reduces Energy on the Data and Enlarges Margins For a conservative bonus, the margin is from -60 to 60, centered by 17. And the energy for the target class "cat" is -17, which is -5 lower than that of CE loss. We plot margin distribution on CIFAR-10 and CIFAR-100 in Figure 10

Adversarial Robustness under AutoAttack Beyond the canonical attack methods PGD, we also conduct experiments with a recently proposed method, AutoAttack (Croce and Hein 2020). The results in Table 11 show that models trained with EL are also more robust to AutoAttack.

Maximum Conditional Energy is a Good Indicator for OOD Detection We do experiments on a large-scale OOD

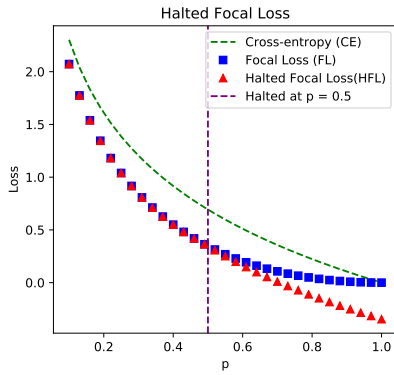
detection task, in which In-D data is ImageNet (It contains 1.2M images, including 50K images for the test), and OOD data is iNaturalist2018 (It contains 438 K samples, including 24.4K images for the test). At the same time, MNIST only has 70k low-resolution images in total. ImageNet contains 1K common objects, while iNaturalist2018 contains 8,142 species. These datasets are large so that the OOD detection is challenging. We show the results in Figure 11. Our indicator "minimum conditional energy among all classes, The "free energy" (Liu et al. 2020) method uses logarithm sum exp of energy among all classes. The difference is that "free energy" considers overall energy. The similarity between us and them is that we all consider energy, but the difference is that we consider the maximum value among classes, which is similar to "max prob."

We can also learn from figures that with all these three indicators, EL (LE=0.75) does better for OOD detection than CE, we reduce the FPR95 from 75.13% (Best for CE) to 71.86%, and EL (LE=1) does better than CE with two energy indicator. These results demonstrate that the EL improves the energy distribution in the case of OOD detection.

Label Smoothing Helps reduce ECE in EL We in Table 13 that Label Smoothing (LS) can reduce ECE for EL. Although label smoothing reduces the energy values for both

Table 8: Well-classified examples also help improve CE loss with label smoothing, and the improvements are on par with it.

Name	MNIST	Fr-En Translation	De-En Translation
CE loss	99.42 ± 0.06	36.66 ± 0.14	34.03 ± 0.06
+ label smoothing	99.54 ± 0.04	37.14 ± 0.15	35.09 ± 0.07
Encouraging Loss	99.56 ± 0.05	37.36 ± 0.08	34.60 ± 0.12
+ label smoothing	99.63 ± 0.04	37.80 ± 0.17	35.50 ± 0.11



Method	iNaturalist2018			
	Overall	Many	Medium	Few
Focal Loss (Lin et al. 2017)	62.9	72.9	64.1	58.8
Halted Focal Loss	63.6(+0.7)	73.1(+0.2)	64.6(+0.5)	59.7(+0.9)
Cross Entropy Loss	64.3	74.1	65.9	59.8
Encouraging Loss (conservative)	65.3	74.3	66.6	61.2
Encouraging Loss	65.8	74.4	66.6	62.4
Encouraging Loss (aggressive)	66.3 (+2.0)	75.1 (+1.0)	67.4(+1.6)	62.6(+2.8)

Table 9: Comparison between Focal Loss, Halted Focal Loss, CE Loss and Encouraging Loss on long-tailed classification dataset iNaturalist 2018, std is about 0.4.

Figure 9: Illustration of Halted Focal Loss which reverts Focal Loss to CE Loss when the correctness of a example is high.

Method	AP	AP ₅₀	AP ₇₅
FL	34.45 ± 0.13	54.23 ± 0.18	36.66 ± 0.13
HFL	34.60 ± 0.09	54.24 ± 0.15	36.73 ± 0.18

Table 10: Results of HFL on the COCO detection dataset, AP denotes average precision. As we can see, increasing the importance of the well-classified examples with halted focal loss achieves better results.

MNIST Epilson	AutoAttack (Linf)		AutoAttack (L2)		
	CE	EL	Epilson	CE	EL
No_attack	98.69	99.04	No_attack	98.69	99.04
0.05	93.51	94.99	0.3	96.76	97.54
0.1	74.89	79.31	1	80.76	84.03
0.15	32.98	37.74	3	0.05	0.27
0.2	5.04	7.59	10	0	0
0.25	0.01	0.26	30	0	0
0.3	0	0	100	0	0

Table 11: Adversarial robustness under the attack of AutoAttack.

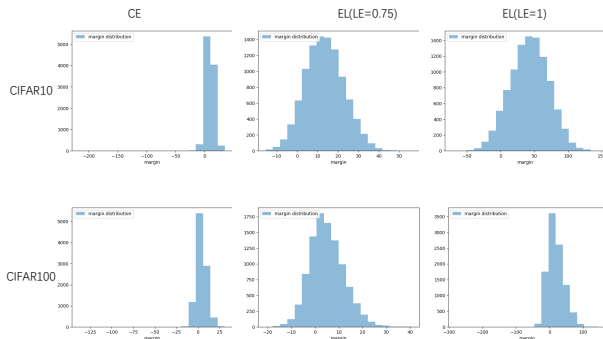


Figure 10: EL enlarges margins on CIFAR-10 and CIFAR-100.

Setting Metric	MNIST vs. F. MNIST	
	AUROC↑	FPR95↓
CE	95.68	20.92
CE+LS 97.10	8.10	
EL	98.04	8.69
EL+LS	99.55	0.98

Table 12: EL is additive with Label smoothing (LS) for improving OOD detection performance. We use F. MNIST to denote Fashion MNIST. Higher AUROC and lower FPR95 are better.

CIFAR-10	w/o LS			w LS		
	ECE	Energy on the data	Energy lowest out-of-data	ECE	Energy on the data	Energy lowest out-of-data
CE	0.031	-15.1	-6.2	0.088	-5	-2.3
EL(LE=0.75)	0.042	-20.8	-7.4	0.034	-4.8	-1.2
EL(LE=1)	0.062	-54.5	-11.73	0.056	-12.3	-0.8

MNIST	w/o LS			w LS		
	ECE	Energy on the data	Energy lowest out-of-data	ECE	Energy on the data	Energy lowest out-of-data
CE	0.005	-27	-5.85	0.009	-4	-0.08
EL(LE=1)	0.004	-131.6	-23.9	0.003	-12.6	-0.07

Table 13: We compared the ECE and Energy of EL and CE with and without label smoothing. Label smoothing reduces ECE when using EL.

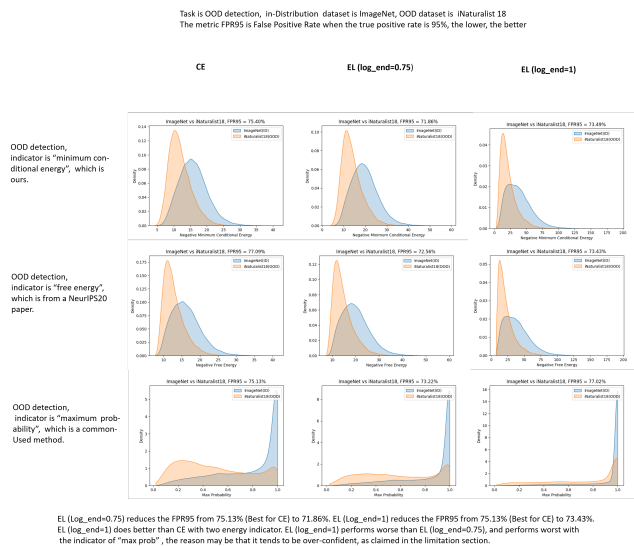


Figure 11: OOD detection on ImageNet vs. iNaturalist2018

CE and EL, we can see that in EL, the difference in energy value between non-data points and data points is still more significant than in CE. For example, on MNIST, after applying label smoothing, the average energy of EL at data points is lower than CE (EL: -12.6, CE:-4), while the average energy of non-data points is slightly higher than CE (EL: -0.07, CE:-0.08), on CIFAR-10, we can push down the energy on the data form -5 to -12.3 while pushing up the lowest energy out of the data from -2.3 to -0.8.

The combination of EL and label smoothing can signific
We in Table 13 that EL is much better than CE, even with LS. For example, we reduce the metric FPR95 from 8.10% to 0.95 on MNIST vs. Fashion MNIST. This is because we get a more discriminative energy distribution.

D Other Related Work

Rewarding the model The idea of encouraging loss that rewarding correct predictions is partly inspired by Reinforcement Learning (RL) (Sutton, Barto et al. 1998; Williams

1992). However, we are different since we take the likelihood into consideration when designing loss, and our goal is to make correct predictions take a more important role than before during learning. considering likelihood in design rewards and let correct predictions take a more important role in learning.

Discussion About Mixup Since the augmentation method Mixup (Zhang et al. 2018) that interpolates both input and output of two examples, unexpectedly pays attention to well-classified examples. So the idea cannot further improve it. However, the application of Mixup is limited because of its bad understandable semantics. including Aug-Mix (Hendrycks et al. 2019) and Token labeling (Jiang et al. 2021) can be combined with this technique. Moreover, the SoTA parameter-efficient model Efficientnet-B0 does not use the Mixup but uses the RandAugment (Cubuk et al. 2020), our idea improves it. Mixup improves the Resnet50 by 0.2 with the same number of epochs (90 epochs) and improves by 1.5 with more than 2x epochs, while we can improve the ResNet-50 by 0.9 without using extra epochs.