

Boosting Co-teaching with Compression Regularization for Label Noise

Yingyi Chen¹ Xi Shen² Shell Xu Hu³ Johan A.K. Suykens¹

¹ESAT-STADIUS, KU Leuven ²LIGM (UMR 8049), Ecole des Ponts, UPE ³Upload AI LLC

{yingyi.chen, johan.suykens}@esat.kuleuven.be, xi.shen@enpc.fr, shell@uploadai.com

Abstract

In this paper, we study the problem of learning image classification models in the presence of label noise. We revisit a simple compression regularization named Nested Dropout [22]. We find that Nested Dropout [22], though originally proposed to perform fast information retrieval and adaptive data compression, can properly regularize a neural network to combat label noise. Moreover, owing to its simplicity, it can be easily combined with Co-teaching [5] to further boost the performance.

Our final model remains simple yet effective: it achieves comparable or even better performance than the state-of-the-art approaches on two real-world datasets with label noise which are Clothing1M [28] and ANIMAL-10N [24]. On Clothing1M [28], our approach obtains 74.9% accuracy which is slightly better than that of DivideMix [12]. On ANIMAL-10N [24], we achieve 84.1% accuracy while the best public result by PLC [30] is 83.4%. We hope that our simple approach can be served as a strong baseline for learning with label noise. Our implementation is available at <https://github.com/yingyichen-cyy/Nested-Co-teaching>.

1. Introduction

The availability of large-scale datasets with clean annotations has made indispensable contributions to the prosperity of deep learning. However, collecting these extensive high-quality data has always been a major challenge since the procedure is both expensive and time-consuming. Appealed by the inexpensive and convenient accesses to large but defective data, such as querying commercial search engines [14], downloading images from social media [16] and other web crawling strategies [19], efforts have been made in literature to learn with imperfect data, among which learning with noisy labels has been attached great importance.

The problem of learning with noisy labels dates back to [1, 21] and the mainstream solutions include adding regularization [15, 18], estimating the label transition matrix [20], training on selected or reweighted samples [5, 8, 17, 27, 29],

label correction [11, 30] and other strategies categorized into the semi-supervised learning genre [3, 9, 12]. In addition to works that provide concise and theoretically sound frameworks to combat label noise, there are other existing works devoted to achieving the state-of-the-art performances on benchmark datasets which require an appropriate tuning on multiple hyper-parameters [12, 30].

With the above in mind, we propose a simple method that combines the regularization and training on selected samples paradigms together to improve state-of-the-art performance on two real-world datasets. To be specific, we revisit a compression regularization called Nested Dropout [22], which was originally proposed to learn ordered feature representations, such that it can be used to perform fast information retrieval and adaptive data compression. It has also been shown in the paper that the ordered feature representation has a strong connection to the PCA solution, that is, the feature channels can be associated to the eigenvectors of the covariance of input data. We find this property important in combating label noise since we are then able to conduct a signal-to-noise separation on the learned feature channels. We verify this intuition empirically, which shows that Nested Dropout gives rise to a strong baseline for learning with noisy labels.

To further take full advantage of this compression regularization, we combine it with a widely acknowledged method called Co-teaching [5]. This is another strong baseline for learning with noisy labels. The basic idea is that two networks can be trained simultaneously where each network updates itself based on the small-loss mini-batch samples selected by its peer. The success of Co-teaching requires the two networks to be reliable enough to select clean samples for each other where we assume the smaller the loss, the cleaner the data [5, 8, 10, 26, 29]. In this regard, we propose a two-stage solution:

- In stage one, two Nested Dropout networks are trained separately to provide reliable base networks for the subsequent stage;
- In stage two, the two trained networks are further fine-tuned with Co-teaching.

As such, we are able to boost the classical strategy with a simple compression regularization.

The rest of the paper is organized as follows: Section 2 gives the architecture of the proposed method. Section 3 presents the experiments on an illustrative toy dataset and two real-world datasets, namely, Clothing1M [28] and ANIMAL-10N [24]. Empirical results demonstrate the effectiveness of our two-stage method given its superior performance comparing to several state-of-the-art approaches such as DivideMix [12] and PLC [30].

2. Method

In this section, we present our approach. We find that the compression regularization is a simple yet effective technique to combat label noise. Specifically, we first focus on one compression regularization named Nested Dropout [22] in Section 2.1, which may serve as a stronger substitute of Dropout [25]. To take full advantage of Nested Dropout, we further combine it with one commonly accepted method called Co-teaching [5] in Section 2.2.

2.1. Nested Dropout

Nested Dropout [22] is one regularization able to learn ordered representations where different dimensions have different degrees of importance. While applied, meaningless representations can be dropped, leading to a compressed network [4]. Taking above into consideration, these ordered representations could be particularly adapted to learning with noisy labels since representations learned from noisy data are expected to be meaningless.

Let $h \in \mathbb{R}^{K \times H \times W}$ be the hidden feature representation obtained by the feature network f , i.e. $h = f(x)$ with x being the input. To obtain an ordered feature representation, in each training iteration, we only keep the first k dimensional feature of h and mask the rest to zeros, that is,

$$z = [h_{1:k}, 0, \dots, 0] \in \mathbb{R}^{K \times \dots}$$

where k is sampled from a categorized distribution with parameters

$$\left\{ p_k \propto \exp\left(-\frac{k^2}{2\sigma_{\text{nest}}^2}\right), \quad \forall k = 1, \dots, K \right\}. \quad (1)$$

where σ_{nest} is one major hyper-parameter in our method. In this manner, smaller k 's are preferred when σ_{nest} is small.

2.2. Co-teaching

Co-teaching [5] is a standard baseline for learning with label noise. The idea is to train two deep networks f_1 and f_2 simultaneously where each network selects its $(1 - \lambda_{\text{forget}})$ percent small-loss instances, i.e. \mathcal{D}_1 and \mathcal{D}_2 , respectively, where λ_{forget} is the forget rate and it is an important hyper-parameter in the Co-teaching architecture. Networks update themselves based on the sample set selected by their peers.

Considering that small-loss instances are more likely to be clean [5, 8, 10, 26, 29], we could obtain classifiers resistant to noisy labels by training them on these instances. However, the above comes with one premise that the classifier should be reliable enough so that the small-loss instances are indeed clean. In the original Co-teaching [5], it proposes to keep all the instances in the mini-batch at the beginning, and then gradually decrease the number of instances in \mathcal{D}_1 and \mathcal{D}_2 until the N -th epoch, after which the number of samples used to train the models becomes fixed. Compared to tuning the hyper-parameter N , we find it more stable to conduct standard training for each model until convergence then fine-tune both models with Co-teaching.

To combine with Nested Dropout, the training procedure comes with two stages: (i) train two Nested Dropout networks separately; (ii) fine-tune these two networks with Co-teaching.

In the first stage, we set a learning rate warm-up to cope with the difficulty of training with Nested Dropout in early epochs resulting from the high probability of dropping most of the channels in the feature layer (i.e. large σ_{nest}). In the second stage, Nested Dropout is maintained during the training of each model except for selecting small-loss instances \mathcal{D}_1 and \mathcal{D}_2 . The final performance is the accuracy of the ensembled model.

3. Experiments

In this section, we present our experimental results. We first show how Nested Dropout deals with regression noise in a toy example in Section 3.1. In Section 3.2, we compare our method with state-of-the-art approaches on two real-world datasets: Clothing1M [28] and ANIMAL-10N [24]. Finally, an ablation study on ANIMAL-10N [24] is given in Section 3.3.

3.1. Toy example: a simple regression with noise

To gain an intuitive understanding on the reason why Nested Dropout [22] is able to resist label noise, we present a simulated regression experiment. We generate a dataset of noisy observations from $y_i = x_i + \epsilon_i$ for $i = 1, \dots, 64$ where x_i is the evenly spaced value between $[0, 10]$ and $\epsilon_i \sim \mathcal{N}(0, 1)$ are independently sampled. We adopt a multilayer perceptron (MLP) composed of three linear layers with input and output dimensions being $1 \rightarrow 64 \rightarrow 128 \rightarrow 1$. Each layer is followed by a ReLU activation except the last one. For the model with Nested Dropout [22], we apply Nested Dropout [22] to the last layer of this MLP and denote it by MLP+Nested. We set $\sigma_{\text{nest}} = 200$ and use the (1) to sample the number of features for training. Results after 100k epochs are shown in Figure 1. It can be seen that MLP overfits to the label noise while MLP+Nested with the first $k \in \{1, 10\}$ channels recovers the ground-truth $y = x$ better. However, with the number of channels increasing,




Figure 1: Comparisons of regression between MLP and MLP incorporated with Nested Dropout [22] on a synthetic noisy label dataset. (a) Training the regression with standard MLP; (b-d) Learning regression with MLP+Nested and plot the prediction results using only the first k channels where (b) $k = 1$, (c) $k = 10$, (d) $k = 100$.

Table 1: Test set accuracy (%) on Clothing1M [28]. We report average accuracy as well as the standard deviation for three runs. Results with “*” are either using a balanced subset or a balanced loss.

Methods	Acc. (%)
CE [27]	67.2
F-correction [27, 20]	68.9
Decoupling [27, 17]	68.5
Co-teaching [27, 5]	69.2
Co-teaching+ [27, 29]	59.3
JoCoR [27]	70.3
JO [26]	72.2
Dropout* [25]	72.8
PENCIL* [11]	73.5
MLNT [13]	73.5
PLC* [30]	74.0
DivideMix* [12]	74.8
Ours	
Nested*	73.1 ± 0.3
Nested + Co-teaching*	74.9 ± 0.2

MLP+Nested gradually overfits to the label noise due to over parameterization. This demonstrates that the first few channels in MLP+Nested contain the main data structure information, while channels towards the end are more likely to encode misled information that would overfit to the noise.

3.2. Comparison with state-of-the-art methods on real datasets

Datasets We conduct experiments on two real datasets: Clothing1M [28] and ANIMAL-10N [24]. The Clothing1M [28] dataset contains 1 million clothing images obtained from online shopping websites with 14 categories. The labels in this dataset are quite noisy with an unknown underlying structure. This dataset also provides 50k, 14k and 10k manually verified clean data for training, validation and testing, respectively. Note that the clean training set is not used during the training. Following [11, 30], in our experiment, we randomly sample a balanced subset which

Table 2: Test set accuracy (%) on ANIMAL-10N [24]. We report average accuracy as well as the standard deviation for three runs.

Methods	Acc. (%)
CE [24]	79.4 ± 0.1
Dropout [25]	81.3 ± 0.3
SELFIE [24]	81.8 ± 0.1
PLC [30]	83.4 ± 0.4
Ours	
Nested	81.3 ± 0.6
Nested + Co-teaching	84.1 ± 0.1

includes 260k images (18.5k images per category) from the noisy training set and use it as our training set and report the classification accuracy on the 10k clean test data. We adopt the standard data augmentation procedures to train ImageNet [2, 6], including random horizontal flip, and resizing the image with a short edge of 256 and then randomly cropping a 224×224 patch from the resized image. The ANIMAL-10N is recently proposed by [24]. It contains 10 animals with confusing appearance. The estimated label noise rate is 8%. There are 50k training and 5k testing images. We did not apply any data augmentation so that the setting is the same with [24].

Implementation details We implement our approach on Pytorch. Experiments on Clothing1M [28] are with ResNet-18 [6] pre-trained on ImageNet [2] following [27]. Note that Nested Dropout or Dropout is applied right before the linear classifier in the network. Models in stage one are optimised with SGD optimizer with a momentum of 0.9, a weight decay of $5e-4$, an initial learning rate of $2e-2$, and batch size of 448. During training, we run learning rate warm-up for 6000 iterations, then train the model for 30 epochs with the learning rate decayed by 0.1 after the 5th epoch. In stage two, we apply Co-teaching to fine-tune two well-trained models. SGD optimizer is utilized with the same settings only with an initial learning rate changing to $2e-3$. Moreover, we set forget rate λ_{forget} in the Co-

Table 3: Average test accuracy (%) with standard deviation (three runs) of different σ_{nest} on ANIMAL-10N [24]. The corresponding optimal number of channels k^* for each model is also provided (entry “ k^* ”). We report test accuracy of single model (entry “Acc.”) as well as the accuracy with the combination of Co-teaching (entry “Co-teaching Acc.”)

σ_{nest}	k^*	Acc. (%)	k^*	Co-teaching Acc. (%)
CE	4096	79.4 ± 0.1	4096	82.2 ± 1.1
25	17.7 ± 9.7	81.0 ± 0.6	16.3 ± 6.9	83.7 ± 0.1
50	18.8 ± 6.9	81.3 ± 0.6	13.4 ± 4.1	84.1 ± 0.2
100	13.6 ± 5.6	81.0 ± 0.5	16.8 ± 7.1	84.1 ± 0.1
150	16.0 ± 3.6	81.1 ± 0.5	18.8 ± 7.4	83.8 ± 0.2
250	13.2 ± 3.1	81.1 ± 0.2	21.0 ± 10.4	83.8 ± 0.1

teaching [5] to be 0.3, freeze batch norm and no warm-up is applied. Models are again trained for 30 epochs with the learning rate decayed by 0.1 after the 5th epoch.

For ANIMAL-10N, we use VGG-19 [23] with batch normalization [7] as in [24]. The two Dropout layers in the original VGG-19 architecture are changed to Nested Dropouts when Nested is applied. The SGD optimizer is employed. Following [24], we train the network for 100 epochs and use an initial learning rate of 0.1, which is divided by 5 at 50% and 75% of the total number of epochs. In stage one, models are trained with learning rate warm-up for 6000 iterations. In stage two, no warm-up is applied, batch norms are frozen, forget rate λ_{forget} is 0.2, initial learning rate is $4e-3$ and decayed by 0.2 after the 5th epoch with 30 epochs in total.

Results on the Clothing1M [28] We now compare our method to state-of-the-art approaches on Clothing1M [28] in Table 1. It is worth noting that Table 1 also includes very recent approaches such as DivideMix [12] and PLC [30]. Surprisingly, our single model with Nested Dropout [22] not only surpasses the standard Dropout [25], but also achieves comparable performance to PENCIL [11] and MLNT [11]. Incorporating Co-teaching [5] further boosts the performance to 74.9% and outperforms the state-of-the-art DivideMix [12].

Results on the ANIMAL-10N [28] Experimental results on ANIMAL-10N [28] are given in Table 2. The dataset is recently proposed, and we compare with two approaches that report performance on this dataset: SELFIE [24] and PLC [30]. It can be seen that our single model achieves comparable performance to Dropout as well as SELFIE [24] and Co-teaching provides a consistent performance boost, which is similar to the results on the Clothing1M [28]. Note that, our best performance by using Nested Dropout [22] and Co-teaching [5] achieves 84.1% accuracy outperforms recent approach PLC [30] by 0.7%.

3.3. Ablation study

In this section, we provide the ablation study of σ_{nest} on ANIMAL-10N [24]. The results are given in Table 3. As

we can see, the Nested Dropout [22] provides consistent improvement compared to training with standard cross entropy loss (entry “CE”) and the performance gain is also robust to the choices of the hyper-parameter σ_{nest} . Moreover, fine-tuning with Co-teaching [5] provides clear boost for all the models. We also show the optimal number of channels of each model (entry “ k^* ”) in the table. Note that though two layers of Nested Dropout have been applied to the classifier of VGG-19, the optimal number of channels k^* is recorded with regard to the last Nested Dropout layer for simplicity. Interestingly, the models trained with Nested Dropout [22] achieve better performance but with only using less than 1% of channels compare to the models trained with standard cross entropy (entry “CE”).

4. Conclusion

In this paper, we investigated the problem of image classification in the presence of noisy labels. Specifically, we first demonstrated that a simple compression regularization called Nested Dropout [22] can be used to combat label noise. Moreover, due to its simplicity, Nested Dropout [22] can be easily combined with Co-teaching [5] to further boost the performance. We validated our approach on two real-world noisy datasets and achieved state-of-the-art performance on both datasets. The proposed approach is simple comparing to many existing methods. Therefore, we hope that our approach can be served as a strong baseline for future research on learning with noisy label.

Acknowledgements The research leading to these results has received funding from the European Research Council under the European Union’s Horizon 2020 research and innovation program / ERC Advanced Grant E-DUALITY (787960). This paper reflects only the authors’ views and the Union is not liable for any use that may be made of the contained information. This work was also supported in part by EU H2020 ICT-48 Network TAILOR (Foundations of Trustworthy AI - Integrating Reasoning, Learning and Optimization), Leuven.AI Institute.

References

- [1] Dana Angluin and Philip Laird. Learning from noisy examples. *Machine Learning*, 1988.
- [2] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009.
- [3] Yifan Ding, Liqiang Wang, Deliang Fan, and Boqing Gong. A semi-supervised two-stage approach to learning from noisy labels. In *WACV*, 2018.
- [4] Aidan N Gomez, Ivan Zhang, Siddhartha Rao Kamalakara, Divyam Madaan, Kevin Swersky, Yarín Gal, and Geoffrey E Hinton. Learning sparse networks using targeted dropout. *arXiv*, 2019.
- [5] Bo Han, Quanming Yao, Xingrui Yu, Gang Niu, Miao Xu, Weihua Hu, Ivor Tsang, and Masashi Sugiyama. Co-teaching: Robust training of deep neural networks with extremely noisy labels. In *NeurIPS*, 2018.
- [6] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [7] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, 2015.
- [8] 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.
- [9] Kyeongbo Kong, Junggi Lee, Youngchul Kwak, Minsung Kang, Seong Gyun Kim, and Woo-Jin Song. Recycling: Semi-supervised learning with noisy labels in deep neural networks. *IEEE Access*, 2019.
- [10] M Pawan Kumar, Benjamin Packer, and Daphne Koller. Self-paced learning for latent variable models. In *NeurIPS*, 2010.
- [11] Yi Kun and Wu Jianxin. Probabilistic End-to-end Noise Correction for Learning with Noisy Labels. In *CVPR*, 2019.
- [12] Junnan Li, Richard Socher, and Steven CH Hoi. Dividemix: Learning with noisy labels as semi-supervised learning. In *ICLR*, 2020.
- [13] Junnan Li, Yongkang Wong, Qi Zhao, and Mohan S Kankanhalli. Learning to learn from noisy labeled data. In *CVPR*, 2019.
- [14] Wen Li, Limin Wang, Wei Li, Eirikur Agustsson, and Luc Van Gool. Webvision database: Visual learning and understanding from web data. *arXiv*, 2017.
- [15] Xingjun Ma, Yisen Wang, Michael E Houle, Shuo Zhou, Sarah Erfani, Shutao Xia, Sudanthi Wijewickrema, and James Bailey. Dimensionality-driven learning with noisy labels. In *ICML*, 2018.
- [16] Dhruv Mahajan, Ross Girshick, Vignesh Ramanathan, Kaiming He, Manohar Paluri, Yixuan Li, Ashwin Bharambe, and Laurens Van Der Maaten. Exploring the limits of weakly supervised pretraining. In *ECCV*, 2018.
- [17] Eran Malach and Shai Shalev-Shwartz. Decoupling” when to update” from” how to update”. In *NeurIPS*, 2017.
- [18] Nagarajan Natarajan, Inderjit S Dhillon, Pradeep Ravikumar, and Ambuj Tewari. Learning with noisy labels. In *NeurIPS*, 2013.
- [19] Christopher Olston and Marc Najork. *Web crawling*. Now Publishers Inc, 2010.
- [20] 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.
- [21] J. Ross Quinlan. Induction of decision trees. *Machine learning*, 1986.
- [22] Oren Rippel, Michael Gelbart, and Ryan Adams. Learning ordered representations with nested dropout. In *ICML*, 2014.
- [23] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. 2015.
- [24] Hwanjun Song, Minseok Kim, and Jae-Gil Lee. Selfie: Refurbishing unclean samples for robust deep learning. In *ICML*, 2019.
- [25] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *JMLR*, 2014.
- [26] Daiki Tanaka, Daiki Ikami, Toshihiko Yamasaki, and Kiyoharu Aizawa. Joint optimization framework for learning with noisy labels. In *CVPR*, 2018.
- [27] Hongxin Wei, Lei Feng, Xiangyu Chen, and Bo An. Combating noisy labels by agreement: A joint training method with co-regularization. In *CVPR*, 2020.
- [28] Tong Xiao, Tian Xia, Yi Yang, Chang Huang, and Xiaogang Wang. Learning from massive noisy labeled data for image classification. In *CVPR*, 2015.
- [29] Xingrui Yu, Bo Han, Jiangchao Yao, Gang Niu, Ivor Tsang, and Masashi Sugiyama. How does disagreement help generalization against label corruption? In *ICML*, 2019.
- [30] Yikai Zhang, Songzhu Zheng, Pengxiang Wu, Mayank Goswami, and Chao Chen. Learning with feature-dependent label noise: A progressive approach. In *ICLR*, 2021.