

Context-based Virtual Adversarial Training for Text Classification with Noisy Labels

Do-Myoung Lee^{1*}, Yeachan Kim^{2*}, Chang-gyun Seo³

¹ShinhanCard, ²Deargen Inc., ³GC Company

¹domyoung89@shinhan.com, ²yeachan@deargen.me, ³ocean@gccompany.co.kr

Abstract

Deep neural networks (DNNs) have a high capacity to completely memorize noisy labels given sufficient training time, and its memorization unfortunately leads to performance degradation. Recently, virtual adversarial training (VAT) attracts attention as it could further improve the generalization of DNNs in semi-supervised learning. The driving force behind VAT is to prevent the models from overfitting to data points by enforcing consistency between the inputs and the perturbed inputs. These strategy could be helpful in learning from noisy labels if it prevents neural models from learning noisy samples while encouraging the models to generalize clean samples. In this paper, we propose context-based virtual adversarial training (ConVAT) to prevent a text classifier from overfitting to noisy labels. Unlike the previous works, the proposed method performs the adversarial training in the context level rather than the inputs. It makes the classifier not only learn its label but also its contextual neighbors, which alleviate the learning from noisy labels by preserving contextual semantics on each data point. We conduct extensive experiments on four text classification datasets with two types of label noises. Comprehensive experimental results clearly show that the proposed method works quite well even with extremely noisy settings.

Keywords: Text Classification, Learning with Noisy Labels

1. Introduction

Deep neural networks (DNNs) have shown human-level performance in various domains, such as image classification, machine translation, and speech recognition. To achieve such an ability, it is indisputably evident that we have to collect a large amount of training data. As labeling such data is laborious and expensive, previous works utilize search engine (Blum et al., 2003; Li et al., 2017) or crowdsourcing (Yan et al., 2014; Yu et al., 2018) to collect labeled dataset. Unfortunately, these low-cost approaches introduce low-quality annotations with *label noise*. It causes DNNs to completely memorize such label noises (i.e., nearly 100% training accuracy) (Zhang et al., 2016), deteriorating generalization capability (Frénay and Verleysen, 2013; Sukhbaatar et al., 2014).

To deal with label noises, recent works primarily propose loss correction approaches. These methods directly correct loss function (e.g., cross-entropy, mean squared error) or the probabilities used to compute it. For example, (Zhang et al., 2016) learns sample weighting scheme via an auxiliary network (called *MentorNet*) and applies it to noisy labels such that corrupt data could get nearly zero sample weights on the loss function. On the one hand, (Han et al., 2018) utilizes an intervention between two same networks with small-loss samples which are considered as clean. However, there exists only a single work to handle label noise on natural language. (Jindal et al., 2019) utilize a noise transition matrix for text classification. This method is useful at text classification on label noises, but it turned out that estimating real noise transition matrix is difficult (Jiang et al., 2018; Han et al., 2018)

especially when the number of classes is large.

Recently, adversarial training at attracts attention as it could prevent networks from misclassifying an image that contains a small perturbation (e.g., adversarial examples). The at enforces the classifier to make consistent predictions on synthetic inputs that have small, approximately worst-case perturbations. (Miyato et al., 2015) extends the idea of at to the semi-supervised regime by removing label dependency on perturbations, which is called a virtual adversarial training (VAT). Adversarial training indirectly have a label smoothing effect as it prevents models from predicting given labels with high confidence by anisotropically smoothing around each data point. Several recent works have shown that label smoothing is effective as a means of coping with label noise (Lukasik et al., 2020) and preventing memorization (Xie et al., 2016). Based on such observations, we further study whether the adversarial training methods, which is an indirect label smoothing method, are useful at training a robust classifier on label noise.

In this paper, we propose a *context-based virtual adversarial training* (ConVAT) to build a robust text classifier on label noise. We inherently follow a fundamental strategy of VAT. Unlike the previous work (Miyato et al., 2016), ConVAT performs the adversarial training on the context-level feature space not the word-level. To that end, we solve min-max optimization by following two steps: *formulating perturbation* and *smoothing*. In the first phase, we calculate the worst-case perturbation into an adversarial direction that could maximize the classification loss on the given samples. We then minimize the distributional distance between a normal sample and a perturbed sample to learn robust classifier on adversarial perturbations. This strategy al-

*Equal contribution.

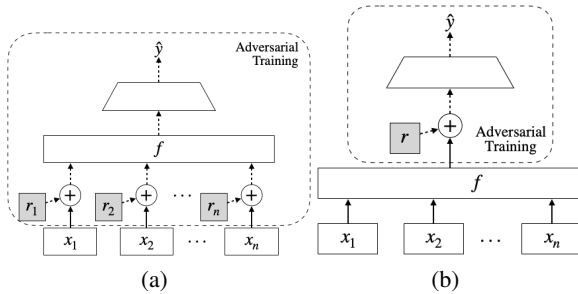


Figure 1: Overall training procedure in (a) virtual adversarial training (b) context-based virtual adversarial training (ConVAT). Dotted line indicates a duplicated propagation path to generate adversarial perturbation.

allows us to train a label noise-robust classifier without placing a burden in a network computation.

In order to show the strength of the proposed method, we conduct extensive experiments on four different datasets with the different kinds of label noise. Comprehensive evaluation results clearly show that ConVAT outperforms the state-of-the-art method in text classification with noisy labels. The in-depth analysis demonstrates that the proposed method have a strong advantage over previous adversarial methods in terms of time and memory complexity. Our code and dataset are publicly available¹.

2. Related Works

2.1. Adversarial Training

Adversarial examples are synthesized input that is created by making small perturbations to the input, which are designed to significantly increase the loss incurred by a machine learning model (Szegedy et al., 2013; Goodfellow et al., 2014). Adversarial training (Goodfellow et al., 2014) and virtual adversarial training (Miyato et al., 2015) are proposed to build a robust classifier against adversarial examples. The difference between the two works lies in how to make an adversarial perturbation. The former takes loss values from label information and makes a worst-case perturbation which prevents the classifier from correctly classifying an input image into a given label. Rather than using label information, the latter only uses the output distribution of around data points since nearby points should be similar in terms of their labels (Han et al., 2018). Recently, VAT attracts attention from various domains. For example, (Deng et al., 2019) applies this idea to graph convolutional neural networks for node classification by generating perturbations to all nodes in networks. (Chen et al., 2020) applies VAT into sequence labeling tasks by combining VAT with the conditional random field.

¹<https://github.com/domyounglee/baseline/tree/convat>

2.2. Noisy Label Learning

Several works try to train a robust classifier on label noise. Recently, these works popularly adopt loss correction approaches which directly correct loss function or the probabilities used to compute it. One common approach is modeling the noise transition matrix which defines the probability of one class flipped to another one (Natarajan et al., 2013; Patrini et al., 2017). However, it turns out that estimating real noise distribution is difficult (Jiang et al., 2018; Han et al., 2018). As a different line of work, several works utilize inference modules to differentiate clean and noisy samples using a neural network (Jiang et al., 2018; Lee et al., 2018), a graphical model (Xiao et al., 2015; Vahdat, 2017). These works are quite useful in learning a robust classifier, but they require the extra clean data or expensive noise detection, which is unpractical in real-world applications. Refined training strategies is another technique to handle label noise. These works typically use two same networks, which have different learning ability, and update parameters based on their predictions such as agreement (Wei et al., 2020), disagreement (Malach and Shalev-Shwartz, 2017; Yu et al., 2019) between two networks.

3. Methodology

In this section, we describe context-based virtual adversarial training, coined ConVAT, in details. We first formulate the problem definition for the noisy label learning (Section 3.1). Then, we take a contextual viewpoint for the classification (Section 3.2). Lastly, we elaborate ConVAT with the contextual viewpoint (Section 3.3). The training procedure is straightforward and described in Figure 1.

3.1. Problem Definition

As we stated before, collecting large-scale labeled datasets with low-cost approaches (e.g., search engine, crowdsourcing) results in label noise in the datasets. In other words, the label y is likely to be flipped into another class y' by human annotators.

Let the training dataset be denoted by $\mathcal{D} = \{(\mathbf{X}_1, y_1), (\mathbf{X}_2, y_2), \dots, (\mathbf{X}_N, y_N)\}$ where \mathbf{X}_i, y_i is the input and the class of i -th sample, respectively. Each input \mathbf{X}_i is the concatenated matrix of the word vectors x_j where subscript j indicates the j -th word in the given text. The noisy dataset can be represented as follows:

$$\mathcal{D}' = \{(\mathbf{X}_i, y'_i) | \mathbf{X}_i \in R^{T_i \times d}, y'_i \in \{1, 2, \dots, K\}, i = 1, \dots, N\} \quad (1)$$

where T_i is the i -th sentence length, d is the dimension of a word vector, and K is the number of the classes.

To corrupt dataset \mathcal{D} into \mathcal{D}' , we assume a Noise Transition Matrix $\Phi \in R^{K \times K}$ filled with probabilities of flipping from one class to another. This follows the same settings with (Jindal et al., 2019). We define

two types of Label Noise as follows and depicted each Noise Transition Matrix in Figure 2:

- **Uniform label noise:** the labels are flipped from one class to another with the same probability across all the classes.
- **Random label noise** the labels are flipped from one class to another based on a certain random distribution across all the classes.

62.50%	12.50%	12.50%	12.50%
12.50%	62.50%	12.50%	12.50%
12.50%	12.50%	62.50%	12.50%
12.50%	12.50%	12.50%	62.50%

58.33%	18.75%	10.41%	12.50%
12.50%	58.92%	14.28%	14.28%
21.05%	21.05%	50%	7.89%
15%	17.50%	17.50%	50%

Figure 2: Examples of the two types of noise transition matrix $\Phi \in R^{K \times K}$

Our goal is to train a robust text classifier from the noisy dataset D' . In order to achieve such objective, preventing the classifier from memorizing label noises is significant to avoid performance degradation.

3.2. Classification from a context viewpoint

Given a task-specific network (e.g., CNN, Transformer), we denote the term *context vector* as last layer activation before the softmax layer. The *context vector* contains rich contextual information which is central to various NLP tasks, such as word sense disambiguation, named entity recognition, coreference resolution (Melamud et al., 2016). Formally, given an input \mathbf{X}_i , a context vector \mathbf{c}_i can be written as $\mathbf{c}_i = f(\mathbf{X}_i)$ where f is the function of a task-specific network. We define the probability of the flipped class y'_i given the context vector \mathbf{c} and weights $\mathbf{W}_c \in R^{v \times K}$, v is the dimension of the context vector, as follows:

$$P(\mathbf{c}_i, y'_i; \theta_{sm}) = \frac{e^{\mathbf{c}_i^\top \mathbf{w}_{y'_i}}}{\sum_{k=1}^K e^{\mathbf{c}_i^\top \mathbf{w}_k}} \quad (2)$$

where θ_{sm} is the trainable parameter \mathbf{W}_c of the softmax layer. We also denote the parameter of the entire network including θ_{sm} as θ .

Based on the probability function P , we define the loss function l with cross entropy as follows :

$$l(\mathbf{c}_i, y'_i; \theta) = -\frac{1}{N} \sum_{i \in N} \log P(\mathbf{c}_i, y'_i; \theta) \quad (3)$$

However, several works have shown that cross-entropy, which is a *de-facto* standard loss function for the text classification, is vulnerable to label noises (Ghosh et al., 2017) due to a memorization effect. In the following subsection, we will describe how can we sidestep the memorization effect.

3.3. Context-based Virtual Adversarial Training

The goal of both AT and VAT is to train a robust classifier against adversarial inputs. To prevent corruption from an input, most works usually add a perturbation into an input (i.e., word embeddings) which is located in most-bottom architecture. However, we have to prevent corruption from an output layer due to label noises. We thus mainly focus on the output layer (i.e., softmax layer), which converts a context vector into a categorical distribution, and our strategy is to minimize the distributional distance between the context vector and the perturbed context vector.

$$\Delta_{KL}(\mathbf{r}, \mathbf{c}_i, \hat{\theta}_{sm}) \equiv \text{KL}[P(\mathbf{c}_i, \cdot; \hat{\theta}_{sm}) || P(\mathbf{c}_i + \mathbf{r}_i^{\text{c-adv}}, \cdot; \hat{\theta}_{sm})]$$

where

$$\mathbf{r}_i^{\text{c-adv}} \equiv \arg \max_{\mathbf{r}} \Delta_{KL}(\mathbf{r}, \mathbf{c}_i, \hat{\theta}_{sm}), \|\mathbf{r}\|_2 \leq \epsilon \quad (4)$$

where $\mathbf{r} \in R^v$ is a random *perturbation vector* and $\text{KL}[\mathbf{p} || \mathbf{q}]$ is a KL-divergence between distributions \mathbf{p} and \mathbf{q} . We refer to $\mathbf{r}_i^{\text{c-adv}}$ as *contextual adversarial perturbation vector* which has a direction of maximum perturbation called the adversarial direction. However, calculating the exact value of $\mathbf{r}_i^{\text{c-adv}}$ is intractable. With the linear approximation method of (Miyato et al., 2016), the resulting adversarial perturbation is re-defined as follows:

$$\mathbf{r}_i^{\text{c-adv}} \approx \epsilon \cdot \frac{\mathbf{g}}{\|\mathbf{g}\|_2} \text{ where } g = \nabla_{\mathbf{r}} \text{KL}[P(\mathbf{c}_i, \cdot; \hat{\theta}_{sm}) || P(\mathbf{c}_i + \mathbf{r}, \cdot; \hat{\theta}_{sm})]_{\mathbf{r}=\xi \mathbf{d}} \quad (5)$$

where $\mathbf{d} \in R^v$ is a normalized random unit vector and ϵ, ξ is a norm constant. This approximation coincides with a single iteration ($t = 1$) of the power method as in previous work (Miyato et al., 2016). Since the goal of the above process is to calculate the perturbation, *we do not update any parameters in this step*. With a view to maintaining the uniformity of context vector, we define the context-driven label smoothing (CLS) :

$$\text{CLS}(\mathbf{c}_i, \theta_{sm}) \equiv \Delta_{KL}(\mathbf{r}_i^{\text{c-adv}}, \mathbf{c}_i, \theta_{sm}) \quad (6)$$

Adding CLS term to the loss function, the label is smoothed by tying the context vector with its neighboring perturbed vectors. The model both learns its label and its most distant neighbor. Since neighbors have useful intermediate representations even with label noises (Bahri et al., 2020), CLS could make the classifier robust against label noises. Therefore, we re-define the loss function as:

$$l(\mathbf{c}_i, y'_i; \theta) = -\frac{1}{N} \sum_{i \in N} \log P(\mathbf{c}_i, y'_i; \theta) + \lambda \frac{1}{N} \sum_{i \in N} \text{CLS}(\mathbf{c}_i, \theta_{sm}) \quad (7)$$

where λ is a controlling factor of label smoothing and it adjusts the degree of how much the context vector count on its neighbors. We show our entire algorithm in Algorithm 1.

Algorithm 1 ConVAT loss function

```

1: Inputs :
2: context vector  $\mathbf{c} \in R^v$ , logit vector  $\mathbf{z} \in R^K$ 
3: , normalized random unit vector  $\mathbf{d} \in R^v$ 
4: Output :
5: KL-divergence  $\Delta_{KL}$  between  $\mathbf{z}$  and  $\hat{\mathbf{z}}$ 
6:  $P(\mathbf{x}) = \mathbf{x}^\top \mathbf{W}$ ,  $\mathbf{W} \in R^{v \times K}$ 
7:
8: Algorithm :
9: function CONVAT( $\mathbf{c}, \mathbf{d}, \mathbf{z}$ )
10:   for 1..t do
11:      $\hat{\mathbf{z}} = P(\mathbf{c} + \mathbf{r})$ 
12:      $\mathbf{g} \leftarrow \nabla_{\mathbf{r}} \Delta_{KL}(\mathbf{z}, \hat{\mathbf{z}})|_{\mathbf{r}=\mathbf{d}}$ 
13:      $\mathbf{r}^{c-adv} = \epsilon \cdot \mathbf{g} / \|\mathbf{g}\|^2$ 
14:      $\hat{\mathbf{z}} = P(\mathbf{c} + \mathbf{r}^{c-adv})$ 
15:   return  $\Delta_{KL}(\mathbf{z}, \hat{\mathbf{z}})$ 

```

4. Experiments

In this section, we conduct extensive experiments to verify the efficacy of our method compared to previous works. We describe the datasets and experimental settings in Section 4.1 and 4.2, respectively. We provide evaluation results compared to the state-of-the-art method in Section 4.3.

4.1. Datasets

In experiments, we use popular text classification datasets, which are SST-2 (Socher et al., 2011), TREC (Voorhees and Tice, 1999), Ag-News (Zhang et al., 2015) and DBpedia (Zhang et al., 2016), to evaluate the proposed method. We explain the datasets in below:

- **SST-2** : Stanford Sentiment Treebank dataset for predicting the sentiment of movie reviews. The dataset is composed of positive or negative reviews.
- **TREC**: The TREC dataset is a dataset for question classification consisting of open-domain, fact-based questions divided into broad semantic categories. We use the six-class (TREC-6) version of the dataset.
- **Ag-News**: The AG’s news topic classification dataset is constructed by choosing the 4 largest classes from the news corpus.
- **DBpedia**: DBpedia is a project aiming to extract structured content from the information created in Wikipedia. We used the classic 14 class version of the dataset.

We conduct experiments on two types of label noise (i.e., *Uniform* and *Random* noise) and evaluate the classification accuracy in the settings of different noise rates. For example, if the noise rate is 0.5, the correct class label is flipped with a 50% probability to other labels.

4.2. Experimental Settings

In this experiment, we use single-layered convolutional neural networks (Kim, 2014) for text classification. The pre-trained GloVe (Pennington et al., 2014) is used to initialize word vectors in a vocabulary. This backbone model is denoted as *CNN*. The comparison methods are equally applied to this model.

We mainly compare our method **ConVAT** with (Jindal et al., 2019) which is the first work that handles label noises in text classification and is denoted as **TransMAT**. When it comes to the regularization hyper-parameter for *TransMat*, we choose the best performing value for each task in uniform and random, respectively. We also compare the vanilla **VAT**, which performs the adversarial training in input-level feature space, to verify the strength of the context-level adversarial training. To fairly compare with our work, we equally set the network’s hyper-parameters (e.g., the number of filters, the size of the window) and these are tuned using the development set. We reproduce all the results in the same settings to compare the performance in the same noise datasets.

In our method, λ is set to 1.0 and epsilon ϵ differently for each dataset (details in Section 5.2). The rest of the hyper-parameters is the same as *TransMat* for a fair comparison. In training steps, we do early stopping based on development set accuracy¹. Entire models are implemented using Pytorch framework. We train the models on a single NVIDIA GeForce RTX 2080 Ti with 11GB of RAM. Our code and dataset are available at <https://github.com/domyounglee/baseline/tree/convat>.

4.3. Experimental Results

Table 1 shows the overall comparison results on the four datasets with different levels and types of label noises. We test all experiments in five times and report the average test accuracy. As can be seen from the table, the performance of the backbone model CNN is largely decreased as the noise rate increases both in two types of noise. However, TransMAT, ConVAT recover the performance loss even in extreme noise settings (i.e., 90% noisy dataset). Between these two works, ConVAT outperforms TransMAT in most settings and datasets. When the noise ratio is large, the performance gap becomes more evident. For example, ConVAT achieves as much as 3.4% and 5.8% improvement on accuracy in the uniform and random noise settings where the noise rate is set to highest. When it

¹Labels in the development set are also flipped in our experiment.

		Uniform (Symmetric)								Random (Asymmetric)							
Noise		0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.1	0.2	0.3	0.4	0.5	0.6	0.7	
AG-News	CNN	92.31	89.96	87.42	84.55	79.96	75.42	68.78	59.94	89.71	86.11	79.05	76.04	65.09	45.79	38.12	
	VAT	92.55	91.4	90.8	90.03	89.9	89.21	88.13	87.74	92.23	90.93	90.4	90.0	87.92	60.93	52.04	
	TransMat	92.53	92.36	91.8	91.13	90.17	89.85	88.9	88.52	91.98	91.5	90.56	89.75	83.22	60.22	49.31	
	ConVAT (Ours.)	92.44	92.02	91.85	91.24	91.01	90.70	89.85	88.92	91.66	91.52	91.06	89.86	88.28	62.93	53.94	
	Noise	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.1	0.2	0.3	0.4	0.5	0.6	0.7	
TREC	CNN	92.8	87.27	83.07	75	69.13	61.53	50.13	39.8	85.93	82.2	74	68.4	53.53	48.2	31.47	
	VAT	91.52	91.56	90.17	88.31	83.15	77.51	68.1	66.13	90	90.2	87.12	86.28	79.9	68.92	48.86	
	TransMat	92.6	91.53	90	86.1	84	82.4	78.67	74.4	90.78	89.6	85.8	84.5	74.4	73.4	48.2	
	ConVAT (Ours.)	92.9	92.06	91.6	89.58	87.27	85.48	79.47	78.01	90.12	89.53	87.73	89.2	82.93	76.93	49.4	
	Noise	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.1	0.2	0.3	0.4	0.5	0.6	0.7	
DBpedia	CNN	99.01	95.19	89.59	74.01	67.73	57.87	47.48	34.01	94.72	86.08	62.87	53.13	40.78	26.6	12.42	
	VAT	98.04	97.94	96.81	96.61	95.52	94.33	94.13	93.53	99.01	94.86	94.66	93.1	91.67	50.62	10.27	
	TransMat	98.78	95.10	98.01	97.46	96.95	96.24	95.23	92.185	98.18	97.24	91.39	88.28	76.73	42.76	9.29	
	ConVAT (Ours.)	98.86	98.65	98.39	97.97	97.67	97.69	96.77	95.36	98.58	98.17	96.1	95.08	90.55	54.01	10.12	
	Noise	0.0	0.3	0.5	0.7	0.75	0.8	0.85	0.9	0.3	0.5	0.7	0.75	0.8	0.85	0.9	
SST-2	CNN	87.27	83.29	79.08	73.42	64.03	58.1	54.73	49.7	81.44	75.58	71.88	63.39	57.12	55.81	52.32	
	VAT	87.82	86.35	84.21	82.89	79.73	69.76	63.4	48.74	85.17	85.35	79.9	77.8	69.96	68.56	46.6	
	TransMat	87.04	85.19	85.48	82.59	76.94	64.63	60.2	48.63	86.27	84.73	80.46	73.37	60.54	61.43	51.71	
	ConVAT (Ours.)	87.25	85.53	84.23	82.48	80.83	68.41	63.9	51.18	85.59	83.74	82.68	78.47	75.94	68.95	53.13	
	Noise	0.0	0.1	0.2	0.3	0.4	0.45	0.47	0.5	0.1	0.2	0.3	0.4	0.45	0.47	0.5	

Table 1: Performance evaluation results on entire datasets with different kinds of noise (Uniform, Random). The best results are highlighted in **boldface**.

comes to the comparison with VAT, ConVAT exhibits superior performance in almost every settings. These results clearly demonstrates that infusing perturbation into context-level is more effective than infusing it into input-level. That is to say, ConVAT has more resiliency on noisy labels, notably on extreme cases.

5. Analysis

In this section, we analyze the proposed method in depth. We first examine whether the proposed method could maintain accuracy without memorizing label noises in Section 5.1. We then describe the choice of the epsilon ϵ (in Eq. 5) at various noise levels in Section 5.2. Lastly, we confirm whether our method is indeed superior to the normal virtual adversarial training (Miayto et al., 2016) in terms of time and memory efficiency.

5.1. Denoising Effect of ConVAT

To analyze the denoising effects of ConVAT, we train ConVAT and CNN with the noisy datasets and recorded the accuracy of train, validation, and test accuracy for every epoch. We plot the results in Figure 3. In this analysis, we train each method on TREC, SST, and Ag-News where 50% of training samples are corrupted by uniform noises. Not surprisingly, we observe that train accuracies of CNN tend to keep increasing over training epochs and, however, the test accuracies drop rapidly after certain points where the memorization occurs. This performance trend can be widely observed in previous literature (Malach and Shalev-Shwartz, 2017; Han et al., 2018). However, ConVAT not only maintains its performance robustly, but it is also increasing little by little. This result indicates that the classifier with ConVAT does not overfit to noisy labels. This effect can be observed in the entire datasets.

To further analyze the effect, we visualize the context vectors of each sample from TREC datasets with 50% uniform noises through t-SNE (Van Der Maaten, 2014) and this is shown in Figure 6. The context vectors from ConVAT is represented in Figure 6 (a) and vectors from CNN is visualized in Figure 6 (b). We first found that vectors from ConVAT stay clustered and are arranged more distinctively as the training proceeds. We believe that such well-clustered results come from the smoothing effect between context vectors and near data points and this smoothing prevents the model from overfitting to noisy labels. In contrast, the vectors from CNN shows some clustered result on early training epochs, but it starts to lose each cluster as training proceeds. This shows that CNN overfits noisy labels and loses the ability to discriminate each data points into their classes.

5.2. Choice of Perturbation length ϵ

Epsilon in Eq. 5 plays a role in determining how far a data point considers the neighbors. If the epsilon is small, the data point only takes into account the close distance neighbor which has similar logits. If the epsilon is big, data points see more diverse contexts that lead the model to be more general and labels to be smoother. Therefore the optimal result of the dataset with a low noise rate can be driven with small epsilon. On the other hand, when training with a higher noise rate, the model must consider a longer range of neighbors.

We confirm whether such epsilon choice is indeed effective by finding optimal epsilon values through a grid search. The search space ranges from 0.0 to 3.0 with 0.1 steps. The best performing epsilon on datasets is shown in Figure 4. As can be seen from all datasets, the optimal epsilon increases when the noise rate in

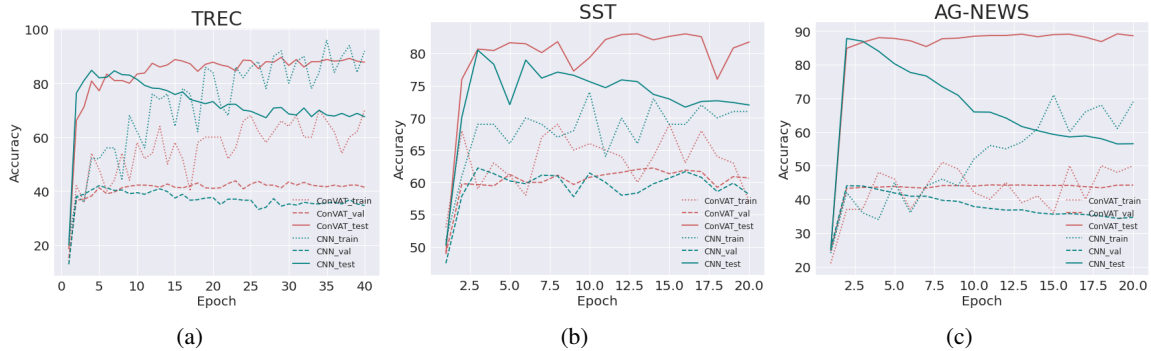


Figure 3: Train, validation and test accuracies versus training epochs on each datasets.

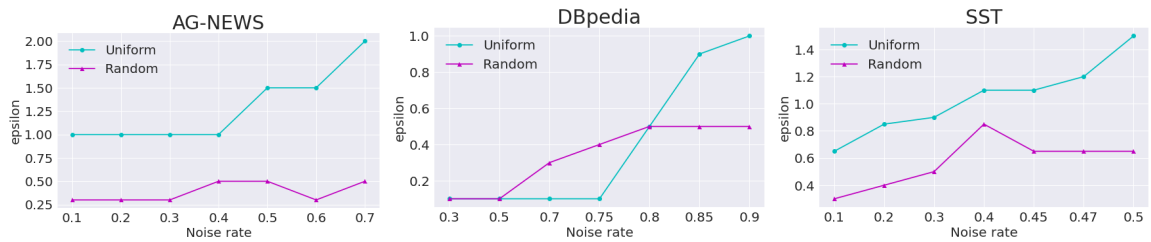


Figure 4: Optimal value of epsilon ϵ on the validation performance for supervised task on AG-NEWS, DBpedia, SST

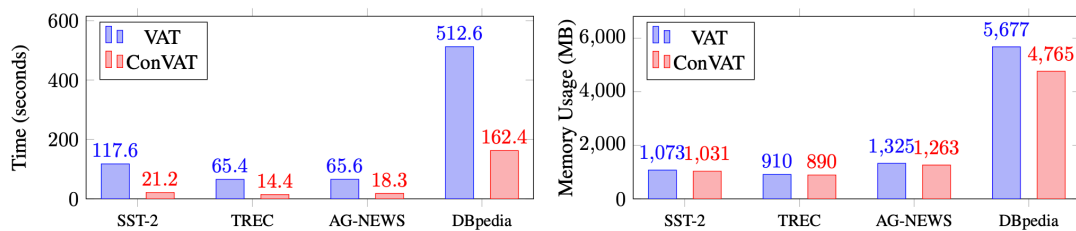


Figure 5: Time and memory usage comparison between ConVAT and VAT (Miyato et al., 2016).

datasets is large. This result supports our strategy of choosing proper epsilon. To find the best epsilon on datasets in practice, estimating the noise rate is required. Fortunately, noise rate can be easily estimated in practice (Liu and Tao, 2015; Patrini et al., 2017) and we, therefore, could find best performing hyper-parameters without an exhaustive grid search.

5.3. Time and Memory usage Analysis of ConVAT

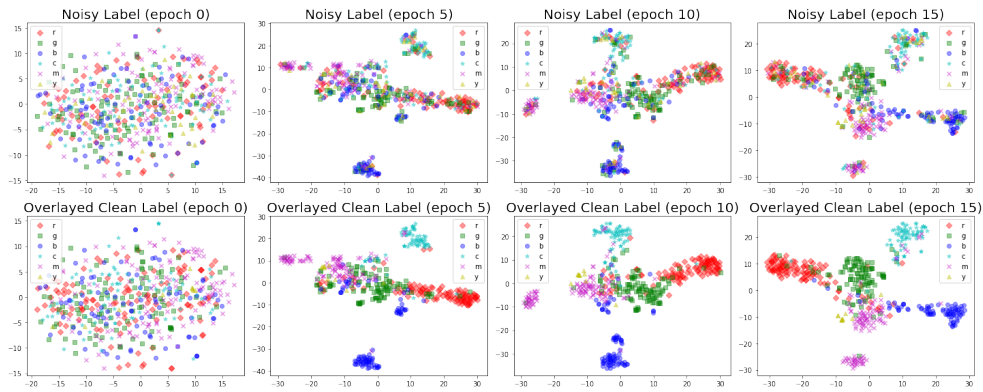
We design the proposed method to handle label noises instead of adversarial inputs. We thus have strong advantages in terms of both time and memory efficiency compared to the normal VAT (Miyato et al., 2016). Specifically, previous works perform a forward and backward propagation twice to calculate a perturbation, but ConVAT propagates the whole network once. It results in that the computational cost of ConVAT does not depend on network architecture unlike VAT. We quantitatively analyze the effect on entire datasets and Figure 5 shows the comparison results.

As can be seen from the figure (left), ConVAT requires 5x less time cost compared to VAT. The computational time gap becomes bigger as the size of the networks is larger.

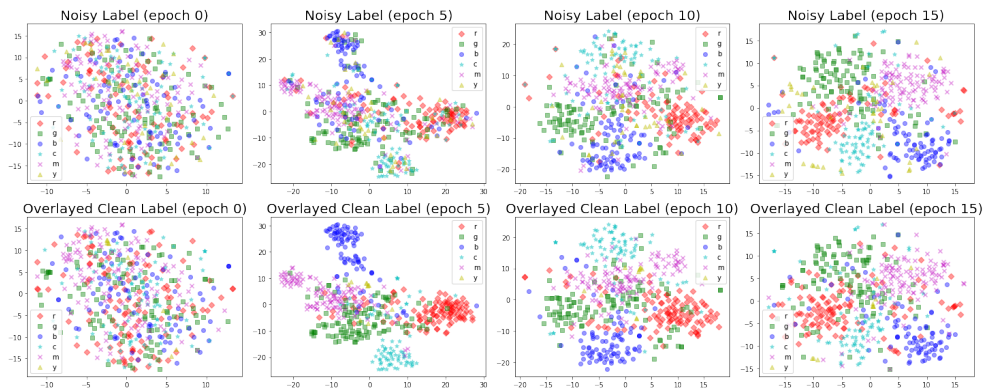
We also have found that our method requires fewer memory footprints than the VAT. The difference comes from the process of calculating worst-case perturbations (Eq. 5). Specifically, our method only requires the gradients of the softmax layer while VAT requires every gradient of entire networks (i.e., embedding layer to softmax layer).

6. Conclusion

In this paper, we have proposed context-based virtual adversarial training, coined ConVAT, that is a robust training method against label noises. Unlike previous works, ConVAT is designed as a network-agnostic manner and does not require additional training parameters. Comprehensive evaluation results have clearly shown that ConVAT is superior to previous works and has strong advantages in terms of the time complex-



(a) Context vectors from ConVAT



(b) Context vectors from CNN

Figure 6: tSNE visualization of the context vectors from ConVAT and CNN.

ity. Furthermore, we analyze our method in-depth and found that ConVAT robustly prevents networks from overfitting to label noises. As future work, we plan to apply our method to different natural language tasks such as a sequence labeling task where noise labels frequently occur.

7. References

- Bahri, D., Jiang, H., and Gupta, M. (2020). Deep k-nn for noisy labels. In *ICML*.
- Blum, A., Kalai, A., and Wasserman, H. (2003). Noise-tolerant learning, the parity problem, and the statistical query model. *ACM*.
- Chen, L., Ruan, W., Liu, X., and Lu, J. (2020). Seqvat: Virtual adversarial training for semi-supervised sequence labeling. In *ACL*.
- Deng, Z., Dong, Y., and Zhu, J. (2019). Batch virtual adversarial training for graph convolutional networks. *arXiv preprint arXiv:1902.09192*.
- Frénay, B. and Verleysen, M. (2013). Classification in the presence of label noise: a survey. *IEEE transactions on neural networks and learning systems*.
- Ghosh, A., Kumar, H., and Sastry, P. (2017). Robust loss functions under label noise for deep neural networks. *AAAI*.
- Goodfellow, I. J., Shlens, J., and Szegedy, C. (2014). Explaining and harnessing adversarial examples. *ICLR*.
- Han, B., Yao, Q., Yu, X., Niu, G., Xu, M., Hu, W., Tsang, I., and Sugiyama, M. (2018). Co-teaching: Robust training of deep neural networks with extremely noisy labels. In *NeurIPS*.
- Jiang, L., Zhou, Z., Leung, T., Li, L.-J., and Fei-Fei, L. (2018). Mentornet: Learning data-driven curriculum for very deep neural networks on corrupted labels. In *ICML*.
- Jindal, I., Pressel, D., Lester, B., and Nokleby, M. (2019). An effective label noise model for dnn text classification. In *NAACL*.
- Kim, Y. (2014). Convolutional neural networks for sentence classification. In *EMNLP*.
- Lee, K.-H., He, X., Zhang, L., and Yang, L. (2018). Cleannet: Transfer learning for scalable image classifier training with label noise. In *CVPR*.
- Li, W., Wang, L., Li, W., Agustsson, E., and Van Gool, L. (2017). Webvision database: Visual learning and understanding from web data. *arXiv preprint arXiv:1708.02862*.
- Liu, T. and Tao, D. (2015). Classification with noisy labels by importance reweighting. *IEEE Transactions on pattern analysis and machine intelligence*.
- Lukasik, M., Bhojanapalli, S., Menon, A. K., and Ku-

- mar, S. (2020). Does label smoothing mitigate label noise?
- Malach, E. and Shalev-Shwartz, S. (2017). Decoupling” when to update” from” how to update”. In *NeurIPS*.
- Melamud, O., Goldberger, J., and Dagan, I. (2016). context2vec: Learning generic context embedding with bidirectional lstm. In *Proceedings of the 20th SIGNLL conference on computational natural language learning*, pages 51–61.
- Miayto, T., Dai, A. M., and Goodfellow, I. (2016). Virtual adversarial training for semi-supervised text classification. *ICLR*.
- Miyato, T., Maeda, S.-i., Koyama, M., Nakae, K., and Ishii, S. (2015). Distributional smoothing with virtual adversarial training. *ICLR*.
- Natarajan, N., Dhillon, I. S., Ravikumar, P. K., and Tewari, A. (2013). Learning with noisy labels. In *NeurIPS*.
- Patrini, G., Rozza, A., Krishna Menon, A., Nock, R., and Qu, L. (2017). Making deep neural networks robust to label noise: A loss correction approach. In *CVPR*.
- Pennington, J., Socher, R., and Manning, C. D. (2014). Glove: Global vectors for word representation. In *EMNLP*.
- Socher, R., Pennington, J., Huang, E. H., Ng, A. Y., and Manning, C. D. (2011). Semi-supervised recursive autoencoders for predicting sentiment distributions. In *EMNLP*.
- Sukhbaatar, S., Bruna, J., Paluri, M., Bourdev, L., and Fergus, R. (2014). Training convolutional networks with noisy labels. *ICLR*.
- Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., and Fergus, R. (2013). Intriguing properties of neural networks. *ICLR*.
- Vahdat, A. (2017). Toward robustness against label noise in training deep discriminative neural networks. In *NeurIPS*.
- Van Der Maaten, L. (2014). Accelerating t-sne using tree-based algorithms. *The Journal of Machine Learning Research*.
- Voorhees, E. M. and Tice, D. M. (1999). The trec-8 question answering track evaluation. In *TREC*.
- Wei, H., Feng, L., Chen, X., and An, B. (2020). Combating noisy labels by agreement: A joint training method with co-regularization. In *CVPR*.
- Xiao, T., Xia, T., Yang, Y., Huang, C., and Wang, X. (2015). Learning from massive noisy labeled data for image classification. In *CVPR*.
- Xie, L., Wang, J., Wei, Z., Wang, M., and Tian, Q. (2016). Disturblabel: Regularizing cnn on the loss layer. In *CVPR*.
- Yan, Y., Rosales, R., Fung, G., Subramanian, R., and Dy, J. (2014). Learning from multiple annotators with varying expertise. *Machine learning*.
- Yu, X., Liu, T., Gong, M., and Tao, D. (2018). Learning with biased complementary labels. In *ECCV*.
- Yu, X., Han, B., Yao, J., Niu, G., Tsang, I. W., and Sugiyama, M. (2019). How does disagreement help generalization against label corruption? *ICML*.
- Zhang, X., Zhao, J., and LeCun, Y. (2015). Character-level convolutional networks for text classification. In *NeurIPS*.
- Zhang, C., Bengio, S., Hardt, M., Recht, B., and Vinyals, O. (2016). Understanding deep learning requires rethinking generalization. *ICLR*.