

Adversarial Dropout for Recurrent Neural Networks

Sungrae Park,^{1*} Kyungwoo Song,² Mingi Ji,² Wonsung Lee,^{3*} Il-Chul Moon²

¹Clova AI Research, NAVER Corp., Korea

²Industrial & Systems Engineering, KAIST, Korea

³AI Center, SK Telecom, Korea

Abstract

Successful application processing sequential data, such as text and speech, requires an improved generalization performance of recurrent neural networks (RNNs). Dropout techniques for RNNs were introduced to respond to these demands, but we conjecture that the dropout on RNNs could have been improved by adopting the *adversarial* concept. This paper investigates ways to improve the dropout for RNNs by utilizing intentionally generated dropout masks. Specifically, the guided dropout used in this research is called as *adversarial dropout*, which adversarially disconnects neurons that are dominantly used to predict correct targets over time. Our analysis showed that our regularizer, which consists of a gap between the original and the reconfigured RNNs, was the upper bound of the gap between the training and the inference phases of the random dropout. We demonstrated that minimizing our regularizer improved the effectiveness of the dropout for RNNs on sequential MNIST tasks, semi-supervised text classification tasks, and language modeling tasks.

Introduction

Many effective regularization methods have been introduced to address the issue of large-scale neural networks predisposed to overfitting. Among several regularization techniques (Bishop 1995; Ioffe and Szegedy 2015; Salimans and Kingma 2016), dropout (Srivastava et al. 2014; Kingma, Salimans, and Welling 2015) has become a common methodology because of its simplicity and effectiveness. The dropout randomly disconnects neural units during training to prevent the feature co-adaptation. Srivastava et al. (Srivastava et al. 2014) interpreted the dropout as an extreme form of a model ensemble by sharing the extensive parameters of a neural network. However, the naive application of the dropout to hidden states of recurrent neural networks (RNNs) failed to prove performance gains (Zaremba, Sutskever, and Vinyals 2015) because it interferes with abstracting long-term information. This issue is caused by the structural differences between feed-forward neural networks (FFNNs) and RNNs.

Recent works have investigated the applications of the dropout on the recurrent connections of RNNs. Gal and Ghahramani (Gal and Ghahramani 2016) suggested the application of the dropout on the same neural units through time

steps, as opposed to dropping different neural units at each time step. Their approach is similar to the L2 regularization on the weight parameters in the recurrent connection. After this suggestion, Semeniuta et al. (Semeniuta, Severyn, and Barth 2016) and Moon et al. (Moon et al. 2015) explored the dropout applications within long-short term memory (LSTM) cells by regulating the gating mechanism. Recently, Merity et al. (Merity, Keskar, and Socher 2018) proposed the use of *DropConnect* (Wan et al. 2013) on the recurrent hidden-to-hidden weight matrices. This approach also allows the recurrent units to share the same dropout mask. These dropout techniques for RNNs have shown that recurrent units should have the same transition metric to process the sequential information. These results are used in extending the dropout techniques on RNNs.

Recently, dropout-based ensemble (DE) regularization (Bachman, Alsharif, and Precup 2015; Ma et al. 2017; Laine and Aila 2017; Tarvainen and Valpola 2017) has been established to improve the dropout techniques. DE regularization conceptually consists of two phases: the generation of the dropout masks and the comparison of the original and the perturbed networks by the dropout masks. Bachman et al. (Bachman, Alsharif, and Precup 2015) and Ma et al. (Ma et al. 2017) tried to minimize the distance between the output distributions of the original network and its randomly perturbed network. Laine and Aila (Laine and Aila 2017) suggested self-ensembling models, or the Π model, containing a distance term between two randomly perturbed networks. Park et al. (Park et al. 2018) proposed *adversarial dropout* that intentionally deactivates neurons that are dominantly used to predict the correct target. By utilizing the adversarial dropout, the regularization consists of the distance between the original and the adversarially perturbed networks. Additionally, they proved further improvements for supervised and semi-supervised image classification tasks. These DE regularizers are fundamentally developed for all types of neural networks, but these DE regularizations are not compatible to RNNs because of the sequential information abstraction.

To investigate the effectiveness of the DE regularization on RNNs, Zolna et al. (Zolna et al. 2018) proposed *fraternal dropout* (FD), an RNN version of *self-ensembling* (Laine and Aila 2017). Specifically, two randomly perturbed RNNs by *DropConnect* are simultaneously supervised from the true labels while maintaining the distance between their predic-




Figure 1: A recurrent neural network with dropout noise on recurrent connections. Each recurrent connection shares the time-invariant dropout mask, ϵ , to alleviate the destructive loss.

tions. They showed further generalization improvements on the language modeling, considered to be one of the most demanding applications of RNNs. Their results motivated us to extend the *adversarial dropout* to RNNs.

This paper proposes *adversarial dropout* on recurrent connections to adaptively regularize RNNs. Similar to the adversarial dropout on FFNNs, adversarial dropout for RNNs also deactivates dominating hidden states to predict the correct target sequence. In order to adapt the concept of the adversarial dropout to the RNN structure, first, we analytically showed that the regularizer, which consists of a gap between the original and the perturbed RNNs, is the upper bound of the gap between the training and the inference phases of the random dropout on an RNN (Bachman, Alsharif, and Precup 2015; Ma et al. 2017). Second, we investigated ways to measure which recurrent neurons are mainly used to predict target sequences in using a time-invariant dropout mask. Finally, we improved the algorithm to find a better adversarial dropout condition compared to the previous algorithm finding the adversarial dropout. According to our experiments, adversarial dropout for RNNs showed the advanced performances on the sequential versions of MNIST, the semi-supervised text classification, and the language modeling tasks.

Preliminary

Recurrent neural networks with dropout noise

We denoted simple RNN models for brevity of notation because derivations for LSTM and GRU follow similarly. Given input sequence $\mathbf{x} = \{x_1, \dots, x_T\}$ of length T , a simple RNN model is formed by the repeated application of function f_h . This generates a hidden state \mathbf{h}_t at time step t by taking x_t and \mathbf{h}_{t-1} as input. For descriptions of the dropout for RNNs, we denoted $d(\mathbf{h}, \epsilon) = \mathbf{h} \odot \epsilon / (1 - p)$ as the random dropout function where \mathbf{h} is the dropout target and ϵ is the dropout mask, sampled from the Bernoulli distribution with success probability $(1 - p)$. Gal et al. (Gal and Ghahramani 2016) proposed variational dropout to drop the previous hidden state \mathbf{h}_{t-1} with the same dropout mask for every time steps. The following is the RNN equation with the variational dropout:

$$\mathbf{h}_t = f_h(x_t, \mathbf{h}_{t-1}) = \sigma(x_t \mathbf{W}_h + d(\mathbf{h}_{t-1}, \epsilon) \mathbf{U}_h + \mathbf{b}_h) \quad (1)$$

where \mathbf{W}_h and \mathbf{U}_h are parameter matrices that model input-to-hidden and hidden-to-hidden (recurrent) connections, respectively; \mathbf{b}_h is a vector of bias terms; and σ is the nonlinear activation function. We notice that all hidden states share the time-invariant dropout mask, ϵ . Figure 1 describes the

dropout noise space on the recurrent connections sharing the same mask across every time steps. The goal of this paper is to investigate the effectiveness of adversarial dropout on RNNs. When we applied the adversarial dropout perturbation to the model, we replaced the dropout mask, ϵ , with adversarial dropout mask, ϵ^{adv} .

Inference gap of dropout for recurrent neural networks

Recent results have shown that there is a gap between objectives for training and testing when using a dropout technique. This is called the *inference gap of dropout* (Ma et al. 2017; Bulò, Porzi, and Kotschieder 2016). Dropout training indicates learning an ensemble of neural networks, but the output of each network in the ensemble should be averaged to provide the final prediction. Unfortunately, this averaging over an exponential number of subnetworks is intractable, and standard dropout typically implements an approximation by applying the expectation of the dropout variable to compute outputs with a deterministic network. However, the approximation causes the gap between the training and the inference phase of the dropout techniques.

Expectation-linearization (EL) regularization (Ma et al. 2017), which is similar to a pseudo-ensemble agreement (Bachman, Alsharif, and Precup 2015), tries to reduce the gap by adding a penalty to the objective function. Let ϵ be the dropout mask applied to the recurrent connection of RNNs, and $\mathbf{p}_t(\mathbf{x}, \epsilon; \theta)$ be the prediction of the model for the input sequence \mathbf{x} at time t . The EL regularizer can be represented as a generalized form as shown below:

$$\mathcal{R}^{\text{EL}}(\mathbf{x}; \theta) := \mathbb{E}_{\epsilon} \left[\sum_{t=1}^T \lambda_t D[\mathbf{p}_t(\mathbf{x}, \mathbb{E}_{\epsilon}[\epsilon]; \theta) \| \mathbf{p}_t(\mathbf{x}, \epsilon; \theta)] \right] \quad (2)$$

where $D[\cdot \| \cdot]$ indicates a non-negative function that represents the distance between two output vectors such as cross-entropy, ϵ is a Bernoulli random vector with the probability p , and λ_t is a hyperparameter controlling the intensity of the loss at time t . The goal of the EL regularizer is to minimize network loss as well as the expected difference between the prediction from the random dropout mask and the prediction from the expected dropout mask, which is fully connected. However, the EL regularizer requires a Monte Carlo (MC) approximation because calculating the penalty is still intractable due to the expectation.

FD regularization (Zolna et al. 2018), which is an RNN version of II models (Laine and Aila 2017), indirectly reduces the inference gap by minimizing the variance of output distributions caused by random dropout masks. Its regularizer consists of the distance between two outputs perturbed by two sampled dropout masks as shown in the following:

$$\mathcal{R}^{\text{FD}}(\mathbf{x}; \theta) := \mathbb{E}_{\epsilon^1, \epsilon^2} \left[\sum_{t=1}^T \lambda_t D[\mathbf{p}_t(\mathbf{x}, \epsilon^1; \theta) \| \mathbf{p}_t(\mathbf{x}, \epsilon^2; \theta)] \right], \quad (3)$$

where ϵ^1 and ϵ^2 are Bernoulli random vectors with the probability p . Zolna et al. (Zolna et al. 2018) proved that the FD regularizer is related to the lower bound of the EL regularizer,

$\frac{1}{4}\mathcal{R}^{\text{FD}}(\mathbf{x}; \theta) \leq \mathcal{R}^{\text{EL}}(\mathbf{x}; \theta)$. To apply the FD regularizer, the MC approximation is required due to the expectation over the dropout masks.

We should note that the adversarial dropout for RNNs consists of the distance between the output distributions perturbed by a base dropout mask and an adversarial dropout mask, ϵ^{adv} . In section 3, we prove that our regularizer is the upper bound of the EL regularizer without any further MC approximation.

Adversarial training

Adversarial training is a training process that uses adversarial examples crafted to include a small perturbation in inputs so as to not classify the perturbed inputs correctly (Goodfellow, Shlens, and Szegedy 2015; Papernot et al. 2016; Miyato et al. 2017). The perturbation is a real-valued vector and is mainly used to generate intentional noisy data. Recently, the concepts of adversarial training has been extended to sequential modeling tasks. Samanta et al. (Samanta and Mehta 2017) proposed an algorithm that crafts adversarial text examples by using gradient-based saliency maps. They proved that the retraining with the crafted text improves generalization performance on the text classification task. In addition to directly crafting the adversarial text examples, Miyato et al. (Miyato, Dai, and Goodfellow 2016) injected adversarial noise in the word-embedding layer and built a regularizer including adversarial perturbation. They evaluated adversarial training on semi-supervised text classification tasks and they achieved state-of-the-art performances. In contrast to the adversarial perturbation on input space, adversarial dropout identifies subnetworks misclassifying data even though a few neurons are deactivated. The successful applications of adversarial training on RNNs become additional motivation to extend adversarial dropout to RNNs. Additionally, our results showed that the model, when combining adversarial training and adversarial dropout, achieved an advanced level of performance on Internet Movie Database (IMDB).

Adversarial dropout for recurrent neural networks

Adversarial dropout (Park et al. 2018) is a novel DE regularization to reduce the distance between two output distributions of the original network and an intentionally perturbed network. Specifically, the adversarial dropout generate a perturbed network by disconnecting neural units that are dominantly used to predict correct targets. The perturbed network reconfigured by the adversarial dropout might not classify the correct target even though few neurons are deactivated from the full network. Therefore, its learning stimulates the useless or the incorrectly learned neurons to better contribute to more accurate predictions. The following is an RNN version of the adversarial dropout regularizer:

$$\mathcal{R}^{\text{AdD}}(\mathbf{x}, \epsilon^0; \theta) := \sum_{t=1}^T \lambda_t \text{D}[\mathbf{p}_t(\mathbf{x}, \epsilon^0; \theta) \parallel \mathbf{p}_t(\mathbf{x}, \epsilon^{\text{adv}}; \theta)] \quad (4)$$

$$\text{where } \epsilon^{\text{adv}} := \underset{\epsilon; \|\epsilon - \epsilon^0\|_2 \leq \delta}{\text{argmax}} \sum_{t=1}^T \lambda_t \text{D}[\mathbf{p}_t(\mathbf{x}, \epsilon^0; \theta) \parallel \mathbf{p}_t(\mathbf{x}, \epsilon; \theta)]. \quad (5)$$

where δ is the hyperparameter controlling the intensity of the noise. In this equation, ϵ^0 is the base dropout mask, which

represents a target network that supervises an adversarially dropped network. For a example, ϵ^0 can be set as $\mathbf{1}$ vector that indicates the original network without any dropped neurons. At each training step, we identified the worst case dropout condition, ϵ^{adv} , against the current model, $\mathbf{p}_t(\mathbf{x}, \epsilon^0; \theta)$, and trained the model to be robust to such dropout perturbations by minimizing the regularization term.

We note that the regularization term is equivalent to three statistical relationships between two output distributions by a dropout mask, ϵ^0 , and its adversarial dropout mask, ϵ^{adv} , as shown below (proof in the appendix).

Remark 1. Let ϵ be an i.i.d. dropout mask; $p_{t,i}(\epsilon)$ be a i^{th} element of $\mathbf{p}_t(\mathbf{x}, \epsilon; \theta) \in \mathbb{R}^M$ where M is the size of the output dimension; $\mathcal{R}_t^{\text{AdD}}(\mathbf{x}, \epsilon; \theta)$ be the distance term at time t in Eq.4; and ϵ^{adv} be an adversarial dropout mask by setting the base dropout mask as ϵ when the distance metric is L2 norm. Then,

$$\begin{aligned} \mathbb{E}_{\epsilon}[\mathcal{R}_t^{\text{AdD}}(\mathbf{x}, \epsilon; \theta)] &= \overbrace{\sum_i V_{\epsilon}[p_{t,i}(\epsilon)] + \sum_i V_{\epsilon}[p_{t,i}(\epsilon^{\text{adv}})]}^{(1)} \\ &\quad - \overbrace{2 \sum_i \text{Cov}_{\epsilon}[p_{t,i}(\epsilon), p_{t,i}(\epsilon^{\text{adv}})]}^{(2)} \\ &\quad + \overbrace{\sum_i (\mathbb{E}_{\epsilon}[p_{t,i}(\epsilon)] - \mathbb{E}_{\epsilon}[p_{t,i}(\epsilon^{\text{adv}})])^2}^{(3)}, \end{aligned} \quad (6)$$

where V_{ϵ} is the variance and Cov_{ϵ} is the covariance. We note that ϵ^{adv} depends on the base dropout mask, ϵ , so its output, $\mathbf{p}_t(\mathbf{x}, \epsilon^{\text{adv}}; \theta)$, contains randomness derived from ϵ . Minimizing the regularization in Eq.5 pursues (1) invariant outputs over random dropout masks and adversarial dropout masks, (2) positive relationship between the two outputs, and (3) minimized distance between the means of the two outputs. The first terms (1) are consistent with the goal of the FD regularization that minimizes the variance of the output distributions over dropout masks (Zolna et al. 2018). The second term (2) indicates that outputs move in the same direction even though dominantly used neurons for the correct predictions are deactivated in the adversarially perturbed network. If the output of sub-network A is higher than the output of sub-network B, the output of adversarial sub-network from A becomes higher than the one from B. This means that a paired gap between a sub-network and its adversarial sub-network becomes similar to the other paired gaps caused by dropping dominantly used neurons. This interpretation provides the connection between one paired case to the others. The third term (3) minimizes the mean of the gaps. This might increase training error of the base network because the adversarially perturbed networks have poor performance. On the other hand, it can improve generalization performance by preventing the base network from overfitting to training data. Figure 2 shows a graphical description of the three statistical relationships.

When comparing our regularizer with the other EL regularizer, our regularization is the upper bound of the EL regularization when $\epsilon^0 = \mathbb{E}_{\epsilon}[\epsilon]$ without the MC approxi-




Figure 2: Regularization term of adversarial dropout indicates three statistical relationships between the outputs of randomly dropped networks and of adversarially dropped networks: (1) the variances of two outputs, (2) the covariance between two outputs, and (3) the distance between the means.

mation under the constrained domain space of the dropout masks, ϵ and ϵ^0 (proof in the appendix).

Proposition 1. $\frac{1}{4}\mathcal{R}^{\text{FD}}(\mathbf{x}; \theta) \leq \mathcal{R}^{\text{EL}}(\mathbf{x}; \theta) \leq \mathcal{R}^{\text{AdD}}(\mathbf{x}, E_{\epsilon}[\epsilon]; \theta)$ when $\epsilon \in \{\epsilon \mid \|E_{\epsilon}[\epsilon] - \epsilon\|_2 \leq \delta\}$.

We note that our regularization is also positioned in the upper bound of the FD regularization. This result shows that our regularization has relevance with others. The adversarial dropout has a largest upper bound, and this suggests the robustness, yet slow learning. Given a set of regularizers, one cannot tell which one is a right structure for his/her domain. Often, researchers just try multiple cases and choose one experimentally.

Finding adversarial dropout masks for recurrent neural networks

In order to identify ϵ^{adv} , we needed to minimize the distance term in Eq.4 with respect to ϵ . However, exact minimization is intractable for modern neural networks because of non-linear activations of the network and the discrete domain of ϵ . This challenge led us to approximate the worst case of the dropout masks by applying the following two steps: (1) identifying the influences of the recurrent neurons and (2) applying a greedy algorithm based on the influence scores (Park et al. 2018).

The influences of the recurrent neurons should be time-invariant because the recurrent neurons are repeatedly applied over every time step. Fortunately, we can derive the influences by using a gradient of the distance term with respect to the dropout mask because the recurrent neurons have a time-invariant dropout mask. To get the influences, we first relax the discrete domain of $\epsilon \in \{0, 1\}^D$ to the continuous domain, $\tilde{\epsilon} \in [0, 1]^D$, then we calculate the gradients with respect to $\tilde{\epsilon}$. We named the gradient values as influence map (IM) because they indicated the influences of the recurrent neurons.

Let $\mathbf{h}_t(\tilde{\epsilon})$ and $\hat{\mathbf{h}}_t(\tilde{\epsilon})$ be a hidden state and a dropped hidden state by $\tilde{\epsilon}$ at time t , and $D_t(\tilde{\epsilon}) = D[\mathbf{p}_t(\mathbf{x}, \epsilon^0; \theta) \parallel \mathbf{p}_t(\mathbf{x}, \tilde{\epsilon}; \theta)]$. The IM can be calculated as follows:

$$\begin{aligned} \text{IM}_i(\epsilon^s) &:= \left. \frac{\partial \sum_{t=1}^T \lambda_t D_t(\tilde{\epsilon})}{\partial \tilde{\epsilon}_i} \right|_{\tilde{\epsilon}=\epsilon^s} \\ &= \sum_{t=1}^T \lambda_t \frac{\partial D_t(\epsilon^s)}{\partial \hat{\mathbf{h}}_t(\epsilon^s)} \left\{ \sum_{u=1}^{t-1} \frac{\partial \hat{\mathbf{h}}_t(\epsilon^s)}{\partial \hat{h}_{u,i}(\epsilon^s)} h_{u,i}(\epsilon^s) \right\}, \end{aligned} \quad (7)$$

where ϵ^s is the initial dropout mask of $\tilde{\epsilon}$. In the equation, $h_{u,i}(\tilde{\epsilon})$ and $\hat{h}_{u,i}(\tilde{\epsilon})$ are the i^{th} components of $\mathbf{h}_t(\tilde{\epsilon})$ and $\hat{\mathbf{h}}_t(\tilde{\epsilon})$ respectively. The equation on the right shows an alternative view of the IM without any Jacobian matrices with respect to $\tilde{\epsilon}$. The equation contains the gradient propagation, $\partial \hat{\mathbf{h}}_t(\epsilon^s) / \partial \hat{h}_{u,i}(\epsilon^s)$, which is calculated by the backpropagation of the recurrent directions from time u to time t . This alternative equation shows that the IM provide the degree of the influence of the recurrent units over time. Additionally, the IM depends on the initial values of the dropout mask, ϵ^s . It is important to note that the initial dropout mask, ϵ^s , should be different from ϵ^0 because the same conditions of dropout mask cause the zero values of the gradient if there is no another stochastic process in the model. In this paper, we initialized ϵ^s by randomly flipping a element of the base dropout mask ϵ^0 to set a initial dropout mask ϵ^s . Calculated IM values indicate the relations between the recurrent neurons and the distance term.

After calculating the IM, the adversarial dropout mask, ϵ^{adv} , should be identified under the constraint, $\|\epsilon^{\text{adv}} - \epsilon^0\|_2 \leq \delta$. For the worst case dropout conditions, the neuron with a positive IM value should be activated and the neuron with a negative IM value should be deactivated. Park et al. (Park et al. 2018) proposed a greedy algorithm to find the worst case of the dropout mask utilizing the IM. The algorithm estimates IM values once and iteratively flips a dropout element shown to be highly effective to the distance term in view of the IM. The adversarial dropout with the greedy algorithm is shown below:

$$\epsilon^{\text{adv}} = \text{Flip}_{\|\epsilon^s - \epsilon^0\|_2 \leq \delta}(\epsilon^s, \text{IM}(\epsilon^s)), \quad (8)$$

where we introduce a flip function, which changes elements of the dropout mask, ϵ^s , in the descending order of $(1 - 2\epsilon^s) \odot \text{IM}(\epsilon^s)$ until the constraint $\|\epsilon^s - \epsilon^0\|_2 \leq \delta$ is satisfied (detail algorithm in appendix). In the flip function, $(1 - 2\epsilon^s) \odot \text{IM}(\epsilon^s)$ indicates the influence of flipping a dropout element because the relation estimated by the IM depends on the state of the current dropout mask, ϵ^s . This approach is simple but can be improved because IM values are unstable when an element is flipped.

We constructed a straightforward way to extend the greedy algorithm. Specifically, we modified the greedy algorithm to update IM values multiple times, as shown here:

$$\begin{aligned} \epsilon^{\text{adv},(0)} &= \epsilon^s, \\ \epsilon^{\text{adv},(k+1)} &= \text{Flip}_{\|\epsilon^{\text{adv},(k+1)} - \epsilon^0\|_2 \leq \frac{k+1}{K} \delta}(\epsilon^{\text{adv},(k)}, \text{IM}(\epsilon^{\text{adv},(k)})), \end{aligned} \quad (9)$$

where K is the maximum number of the iteration. For a better approximation of the worst-case dropout mask, the IM values are updated using the current dropout mask and the bound of the constraint is increased through the iterations. There is a trade-off between the better approximation and the computational cost. The large number of the iteration, K , leads to a better approximation, but it also increases the computational cost proportionately. In our experiment, we tested cases of $K=1$ and $K=2$ to investigate performance improvement obtained by sacrificing the computational cost.

Table 1: Test error rates of supervised learning on *s*MNIST and *p*MNIST. Each setting is repeated ten times.

| Model | <i>s</i> MNIST | <i>p</i> MNIST |
|---------------|------------------------------|------------------------------|
| Unregularized | 1.018 (± 0.165) | 9.924 (± 0.307) |
| VD (2016) | 0.721 (± 0.111) | 5.218 (± 0.132) |
| FD (2018) | 0.720 (± 0.061) | 5.121 (± 0.121) |
| AD ($K=1$) | 0.705 (± 0.047) | 5.046 (± 0.067) |
| AD ($K=2$) | 0.644 (± 0.046) | 5.030 (± 0.065) |

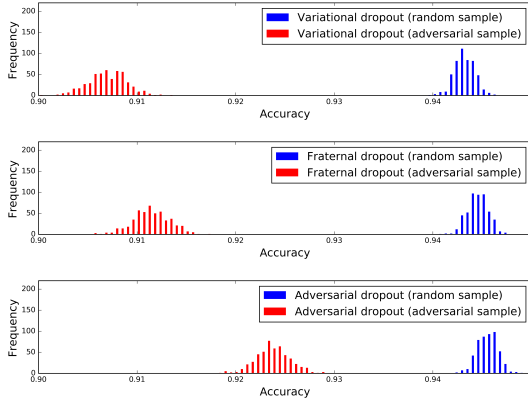


Figure 3: Test accuracies of the perturbed RNNs by random dropout and adversarial dropout masks.

Experiments

Sequential MNIST

Sequential MNIST tasks, also known as pixel-by-pixel MNIST, process each image one pixel at a time and predicts the label of the image. The tasks can be categorized by the order of the pixels in a sequence: *s*MNIST in scanline order and *p*MNIST in a fixed random order. Because the size of an MNIST image is 28×28 , the length of the sequence becomes 784.

Our baseline model included a single LSTM layer of 100 units with a softmax classifier to produce a prediction from the final hidden state. For the settings for the dropout, we set the dropout probability as 0.1 for the baseline models. In the case of the adversarial dropout, we adapted $\epsilon^0 = E_{\epsilon}[\epsilon]$, which indicates the expectation of the dropout mask, and $\delta = 0.03$, which represents the maximum changes from the base dropout mask as 3%. These hyperparameters of the baseline models as well as our models were retrieved in the validation phase. All models were trained with the same optimizer (detail settings in appendix).

Table 1 shows the test performances of the dropout-based regularizations. When applying variational dropout on the recurrent connections, the performance was improved from the unregularized LSTM. By adding regularization terms, the performance on the test dataset was improved in the following order: FD, and adversarial dropout ($K=1, 2$).

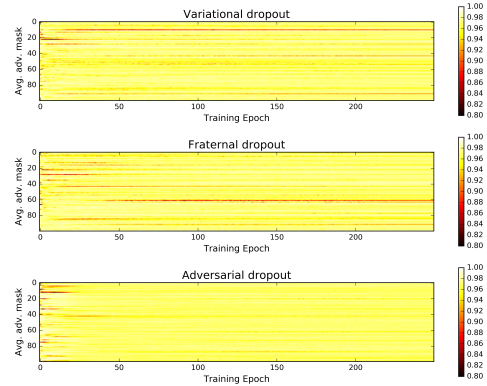


Figure 4: Averages of the adversarial dropout masks over the test dataset through the training iterations.

In order to investigate how adversarial dropout training affects the accuracy distribution of the dropped subnetworks, we sampled random dropout masks and adversarial dropout masks respectively 500 times and tested the performance of the subnetwork reconfigured by the sampled dropout masks. Figure 3 (a) shows the histogram of the results. In this analysis, we applied the dropout probability, $p=0.03$, and the hyperparameters, $\delta=0.03$ and $K=2$, to the adversarial dropout. In the case of the variational dropout, we can see that the test performances of adversarially dropped networks were worse than the test performance of randomly dropped networks. By adding the regularization term of the FD, the performances of adversarially dropped networks were improved and caused the performance distribution of randomly dropped networks to move the right and its variance to reduce. That is, regularization of the FD indirectly improved the performance of the subnetwork in its worst case scenario. On the other hand, adversarial dropout directly controlled the sub-network in the worst case scenario. As a result, the performances of the base subnetwork and the adversarially perturbed networks were improved more than the performances in the case of FD regularization.

We additionally investigated the adversarial dropout condition changes throughout the training phase. We calculated $E_{\text{test}}[\epsilon^{\text{adv}}]$ on every training epoch. Figure 3 (b) shows the visualized adversarial dropout masks. As can be seen, the adversarial dropout condition is equally spread over the test dataset in the case of adversarial dropout training, whereas the other cases show that a certain number of elements are selected more frequently for the adversarial dropout condition. These results show that adversarial dropout training stimulates the useless or the incorrectly learned neurons to better contribute to a more accurate prediction.

Semi-supervised text classification

The text classification task is one of the most important tasks utilizing RNN architecture. We evaluated our method on two text datasets, IMDB and Elec. IMDB is a standard benchmark

Table 2: Test error rates (%) on the IMDB and Elec classification tasks. All models were pre-trained by neural language models. Our experiments were repeated five times.

| | Method | IMDB | Elec |
|-----------------------|---|------------------------------|-----------------------|
| Reported | One-hot CNN (Johnson and Zhang 2015; 2016) | 6.05 | 5.87 |
| | One-hot bi-LSTM (Johnson and Zhang 2016) | 5.94 | 5.55 |
| | Adversarial training (Miyato, Dai, and Goodfellow 2016) | 6.21 | 5.61 |
| | Virtual adversarial training (Miyato, Dai, and Goodfellow 2016) | 5.91 | 5.54 |
| | Adversarial + virtual adversarial training (Miyato, Dai, and Goodfellow 2016) | 6.09 | 5.40 |
| Our base model | | | |
| | LSTM | 7.164 (± 0.070) | 6.155 (± 0.092) |
| Our experiment | Input perturbation on embeddings | | |
| | Adversarial training | 6.296 (± 0.115) | 5.652 (± 0.064) |
| | Virtual adversarial training | 5.914 (± 0.051) | 5.610 (± 0.045) |
| | Dropout perturbation on hidden states | | |
| | Variational dropout | 7.042 (± 0.059) | 6.114 (± 0.074) |
| | Expectation-Linearization dropout | 6.358 (± 0.065) | 5.734 (± 0.052) |
| | Fraternal dropout | 6.349 (± 0.060) | 5.785 (± 0.085) |
| | Adversarial dropout on the last hidden state | 6.471 (± 0.058) | 5.784 (± 0.064) |
| | Adversarial dropout on the recurrent connection ($K=1$) | 6.155 (± 0.053) | 5.622 (± 0.051) |
| | Adversarial dropout on the recurrent connection ($K=2$) | 6.005 (± 0.043) | 5.603 (± 0.049) |
| | Input and dropout perturbations | | |
| | Adversarial dropout ($K=1$) + virtual adversarial training | 5.715 (± 0.081) | 5.638 (± 0.044) |
| | Adversarial dropout ($K=2$) + virtual adversarial training | 5.687 (± 0.068) | 5.621 (± 0.058) |

movie review dataset for sentiment classification (Maas et al. 2011). Elec is a dataset on electronic product reviews from Amazon (Johnson and Zhang 2015). Both datasets consist of labeled text data and unlabeled text data. In Elec, we removed some duplicated examples in labeled and unlabeled training dataset. We describe the detail summarization in the appendix.

Following previous studies on these datasets, we implemented a pre-training phase through neural language modeling and a training phase using the classification model. Most experiment settings are similar to (Miyato, Dai, and Goodfellow 2016) (detail settings in the appendix). Our implementation code will be available at https://github.com/sungraepark/adversarial_dropout_text_classification.

Table 2 shows the test performance on IMDB and Elec with each training method. The base LSTM achieved a 7.164% error rate on IMDB and 6.155% error rate on Elec with only embedding dropout. By adding variational dropout on the hidden states with the time-invariant dropout mask, the error rate was slightly reduced to 7.042% and 6.114%, respectively. When we tested EL regularizers, that are marked as the dropout perturbation in the table, the performance was improved in the following order: EL, FD, and adversarial dropout ($K=1, 2$). In the results, we found that adversarial dropout improved performance more than the FD. This might be caused by the fact that the regularization term of adversarial dropout is the upper bound of the regularization term of the FD. When comparing from adversarial dropout on the last hidden state, aka. original adversarial dropout, we found that applying adversarial dropout at the recurrent connections

was more effective. Additionally, our iterative algorithm with $K=2$ provided the further performance gains. Furthermore, we evaluated a jointed method of adversarial dropout, which is dropout perturbation on hidden states, and virtual adversarial training, which is a linear perturbation on embedding space. As a result, we found that the jointed method achieved an impressive performance of 5.687% on IMDB.

Word-level language models

We investigated the performance of our regularization model in language modeling for two benchmark datasets, the Penn Treebank (PTB) (Mikolov et al. 2010) and WikiText-2 (WT2) dataset (Merity et al. 2017). We preprocessed both datasets: for PTB as specified by Mikolov et al. (Mikolov et al. 2010) and for WT2 as specified by Koehn et al. (Koehn et al. 2007). In this experiment, we used AWD-LSTM 3-layer architecture that was introduced by Merity et al. (Merity, Keskar, and Socher 2018). Specifically, the architecture contains one 400 dimensional embedding layer and three LSTM layers whose dimensions are 1150 for first two layers and 400 for the third layer. The embedding matrix and the weight matrix of the last layer for predictions were tied (Merity, McCann, and Socher 2017). In the training phase, Merity et al. (Merity, Keskar, and Socher 2018) applied DropConnect to improve model performances. In this experiment, we left a vast majority of hyperparameters used in the baseline model, i.e. embedding and hidden sizes, learning rate, and so on. Instead, we controlled the hyperparameters, δ and K , related to the adversarial dropout. We did not use the additional regularization terms except for the L2 norm for weight

Table 3: Perplexity on word-level language modeling on Penn Treebank and WikiText-2.

| Model | PTB | | | WikiText-2 | | |
|---|--------|-----------------------|-----------------------|------------|-----------------------|-----------------------|
| | Param. | Val. | Test | Param. | Val. | Test |
| Variational dropout (Gal and Ghahramani 2016) | 51M | 71.1 | 68.5 | 28M | 91.5 | 87.0 |
| 5-layer RHN (Melis, Dyer, and Blunsom 2018) | 24M | 64.8 | 62.2 | 24M | 78.1 | 75.6 |
| AWD-LSTM (Merity, Keskar, and Socher 2018) | 24M | 60.0 | 58.3 | 34M | 68.6 | 65.8 |
| Fraternal dropout (Zolna et al. 2018) | 24M | 58.9 | 56.8 | 34M | 66.8 | 64.1 |
| Adversarial dropout | 24M | 58.7 ± 0.3 | 56.4 ± 0.2 | 34M | 66.5 ± 0.1 | 63.4 ± 0.2 |

parameters. Our implementation code will be available at https://github.com/sungraepark/adversarial_dropout_lm.

Previous results with LSTM showed that fine-tuning is important to achieve state-of-the-art performances (Merity, Keskar, and Socher 2018; Li et al. 2018), so we fine-tuned model parameters once without regularization after learning was over with our regularization term. Table 3 shows the perplexity on both the PTB and WikiText-2 validation and test datasets. Our approach showed an advanced performance compared to existing benchmarks.

Conclusion

The improvement of generalization performance in RNNs is required for the successful application of sequential data processing. The existing methods utilizing the dropout depend on the random dropout mask without considering a guided sampling on the dropout mask. In contrast, we developed an RNN version of adversarial dropout, which is a deterministic dropout technique to find a subnetwork inferior to target predictions. Specifically, we proved that the regularizer with the adversarial dropout is the upper bound of the EL and FD regularizers. Additionally, we found a way to measure which recurrent neurons are mainly used for target predictions. Furthermore, we improved the algorithm to find the adversarial dropout condition. In our experiments, we showed that the adversarial dropout for RNNs improved generalization performance on the sequential MNIST tasks, the semi-supervised text classification tasks, and word-level language modeling tasks. More importantly, we achieved a highly advanced performance of 5.687% on IMDB when applying an adversarial perturbation on word embeddings and an adversarial dropout perturbation, together.

Acknowledgement

This research was supported by Basic Science Research Program through the National Research Foundation of Korea(NRF) funded by the Ministry of Education(NRF-2018R1C1B6008652)

References

Bachman, P.; Alsharif, O.; and Precup, D. 2015. Learning with pseudo-ensembles. In *Advances in Neural Information Processing Systems*, 3365–3373.

Bishop, C. M. 1995. Training with noise is equivalent to tikhonov regularization. *Neural computation* 7(1):108–116.

Bulò, S. R.; Porzi, L.; and Kotschieder, P. 2016. Dropout distillation. In *International Conference on Machine Learning*, 99–107.

Gal, Y., and Ghahramani, Z. 2016. A theoretically grounded application of dropout in recurrent neural networks. In *Advances in Neural Information Processing Systems*, 1019–1027.

Goodfellow, I.; Shlens, J.; and Szegedy, C. 2015. Explaining and harnessing adversarial examples. In *International Conference on Learning Representations*.

Ioffe, S., and Szegedy, C. 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, 448–456.

Johnson, R., and Zhang, T. 2015. Semi-supervised convolutional neural networks for text categorization via region embedding. In *Advances in neural information processing systems*, 919–927.

Johnson, R., and Zhang, T. 2016. Supervised and semi-supervised text categorization using lstm for region embeddings. In *International Conference on Machine Learning*, 526–534.

Kingma, D. P.; Salimans, T.; and Welling, M. 2015. Variational dropout and the local reparameterization trick. In *Advances in Neural Information Processing Systems*, 2575–2583.

Koehn, P.; Hoang, H.; Birch, A.; Callison-Burch, C.; Federico, M.; Bertoldi, N.; Cowan, B.; Shen, W.; Moran, C.; Zens, R.; et al. 2007. Moses: Open source toolkit for statistical machine translation. In *Proceedings of the 45th annual meeting of the ACL on interactive poster and demonstration sessions*, 177–180. Association for Computational Linguistics.

Laine, S., and Aila, T. 2017. Temporal ensembling for semi-supervised learning. In *Proc. International Conference on Learning Representations (ICLR)*.

Li, Z.; He, D.; Tian, F.; Chen, W.; Qin, T.; Wang, L.; and Liu, T.-Y. 2018. Towards binary-valued gates for robust lstm training. In *International Conference on Machine Learning*.

Ma, X.; Gao, Y.; Hu, Z.; Yu, Y.; Deng, Y.; and Hovy, E. 2017. Dropout with expectation-linear regularization. In *International Conference on Learning Representations*.

Maas, A. L.; Daly, R. E.; Pham, P. T.; Huang, D.; Ng, A. Y.; and Potts, C. 2011. Learning word vectors for sentiment analysis. In *Proceedings of the 49th annual meeting of the*

- association for computational linguistics: Human language technologies-volume 1*, 142–150. Association for Computational Linguistics.
- Melis, G.; Dyer, C.; and Blunsom, P. 2018. On the state of the art of evaluation in neural language models. In *International Conference on Learning Representations*.
- Merity, S.; Xiong, C.; Bradbury, J.; and Socher, R. 2017. Pointer sentinel mixture models. In *International Conference on Learning Representations*.
- Merity, S.; Keskar, N. S.; and Socher, R. 2018. Regularizing and optimizing lstm language models. In *International Conference on Learning Representations*.
- Merity, S.; McCann, B.; and Socher, R. 2017. Revisiting activation regularization for language rnns. *arXiv preprint arXiv:1708.01009*.
- Mikolov, T.; Karafiát, M.; Burget, L.; Černocký, J.; and Khudanpur, S. 2010. Recurrent neural network based language model. In *Eleventh Annual Conference of the International Speech Communication Association*.
- Miyato, T.; Maeda, S.-i.; Koyama, M.; and Ishii, S. 2017. Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *arXiv preprint arXiv:1704.03976*.
- Miyato, T.; Dai, A. M.; and Goodfellow, I. 2016. Adversarial training methods for semi-supervised text classification. *arXiv preprint arXiv:1605.07725*.
- Moon, T.; Choi, H.; Lee, H.; and Song, I. 2015. Rnndrop: A novel dropout for rnns in asr. In *Automatic Speech Recognition and Understanding (ASRU), 2015 IEEE Workshop on*, 65–70. IEEE.
- Papernot, N.; McDaniel, P.; Jha, S.; Fredrikson, M.; Celik, Z. B.; and Swami, A. 2016. The limitations of deep learning in adversarial settings. In *Security and Privacy (EuroS&P), 2016 IEEE European Symposium on*, 372–387. IEEE.
- Park, S.; Park, J.-K.; Shin, S.-J.; and Moon, I.-C. 2018. Adversarial dropout for supervised and semi-supervised learning. In *AAAI*, volume 4, 12.
- Salimans, T., and Kingma, D. P. 2016. Weight normalization: A simple reparameterization to accelerate training of deep neural networks. In *Advances in Neural Information Processing Systems*, 901–901.
- Samanta, S., and Mehta, S. 2017. Towards crafting text adversarial samples. *arXiv preprint arXiv:1707.02812*.
- Semeniuta, S.; Severyn, A.; and Barth, E. 2016. Recurrent dropout without memory loss. Technical report.
- Srivastava, N.; Hinton, G. E.; Krizhevsky, A.; Sutskever, I.; and Salakhutdinov, R. 2014. Dropout: a simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research* 15(1):1929–1958.
- Tarvainen, A., and Valpola, H. 2017. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *Advances in neural information processing systems*, 1195–1204.
- Wan, L.; Zeiler, M.; Zhang, S.; Le Cun, Y.; and Fergus, R. 2013. Regularization of neural networks using dropconnect. In *International Conference on Machine Learning*, 1058–1066.
- Zaremba, W.; Sutskever, I.; and Vinyals, O. 2015. Recurrent neural network regularization. *ICLR*.
- Zolna, K.; Arpit, D.; Suhubdy, D.; and Bengio, Y. 2018. Fraternal dropout. In *International Conference on Learning Representations*.