

# Direct Output Connection for a High-Rank Language Model

Sho Takase<sup>†</sup> Jun Suzuki<sup>†‡</sup> Masaaki Nagata<sup>†</sup>

<sup>†</sup>NTT Communication Science Laboratories

<sup>‡</sup>Tohoku University

{takase.sho, nagata.masaaki}@lab.ntt.co.jp

jun.suzuki@ecei.tohoku.ac.jp

## Abstract

This paper proposes a state-of-the-art recurrent neural network (RNN) language model that combines probability distributions computed not only from a final RNN layer but also from middle layers. Our proposed method raises the expressive power of a language model based on the matrix factorization interpretation of language modeling introduced by Yang et al. (2018). The proposed method improves the current state-of-the-art language model and achieves the best score on the Penn Treebank and WikiText-2, which are the standard benchmark datasets. Moreover, we indicate our proposed method contributes to two application tasks: machine translation and headline generation. Our code is publicly available at: [https://github.com/nttcs-lab-nlp/doc\\_lm](https://github.com/nttcs-lab-nlp/doc_lm).

## 1 Introduction

Neural network language models have played a central role in recent natural language processing (NLP) advances. For example, neural encoder-decoder models, which were successfully applied to various natural language generation tasks including machine translation (Sutskever et al., 2014), summarization (Rush et al., 2015), and dialogue (Wen et al., 2015), can be interpreted as conditional neural language models. Neural language models also positively influence syntactic parsing (Dyer et al., 2016; Choe and Charniak, 2016). Moreover, such word embedding methods as Skip-gram (Mikolov et al., 2013) and vLBL (Mnih and Kavukcuoglu, 2013) originated from neural language models designed to handle much larger vocabulary and data sizes. Neural language models can also be used as contextualized word representations (Peters et al., 2018). Thus, language modeling is a good benchmark task for investigating the general frameworks of neural methods in NLP field.

In language modeling, we compute joint probability using the product of conditional probabilities. Let  $w_{1:T}$  be a word sequence with length  $T$ :  $w_1, \dots, w_T$ . We obtain the joint probability of word sequence  $w_{1:T}$  as follows:

$$p(w_{1:T}) = p(w_1) \prod_{t=1}^{T-1} p(w_{t+1}|w_{1:t}). \quad (1)$$

$p(w_1)$  is generally assumed to be 1 in this literature, that is,  $p(w_1) = 1$ , and thus we can ignore its calculation. See the implementation of Zaremba et al. (2014)<sup>1</sup>, for an example. RNN language models obtain conditional probability  $p(w_{t+1}|w_{1:t})$  from the probability distribution of each word. To compute the probability distribution, RNN language models encode sequence  $w_{1:t}$  into a fixed-length vector and apply a transformation matrix and the softmax function.

Previous researches demonstrated that RNN language models achieve high performance by using several regularizations and selecting appropriate hyperparameters (Melis et al., 2018; Merity et al., 2018). However, Yang et al. (2018) proved that existing RNN language models have low expressive power due to the *Softmax bottleneck*, which means the output matrix of RNN language models is low rank when we interpret the training of RNN language models as a matrix factorization problem. To solve the *Softmax bottleneck*, Yang et al. (2018) proposed *Mixture of Softmaxes* (MoS), which increases the rank of the matrix by combining multiple probability distributions computed from the encoded fixed-length vector.

In this study, we propose *Direct Output Connection* (DOC) as a generalization of MoS. For stacked RNNs, DOC computes the probability distributions from the middle layers including input embeddings. In addition to raising the rank, the

<sup>1</sup><https://github.com/wojzaremba/lstm>

proposed method helps weaken the vanishing gradient problem in backpropagation because DOC provides a shortcut connection to the output.

We conduct experiments on standard benchmark datasets for language modeling: the Penn Treebank and WikiText-2. Our experiments demonstrate that DOC outperforms MoS and achieves state-of-the-art perplexities on each dataset. Moreover, we investigate the effect of DOC on two applications: machine translation and headline generation. We indicate that DOC can improve the performance of an encoder-decoder with an attention mechanism, which is a strong baseline for such applications. In addition, we conduct an experiment on the Penn Treebank constituency parsing task to investigate the effectiveness of DOC.

## 2 RNN Language Model

In this section, we briefly overview RNN language models. Let  $V$  be the vocabulary size and let  $P_t \in \mathbb{R}^V$  be the probability distribution of the vocabulary at timestep  $t$ . Moreover, let  $D_{h^n}$  be the dimension of the hidden state of the  $n$ -th RNN, and let  $D_e$  be the dimensions of the embedding vectors. Then the RNN language models predict probability distribution  $P_{t+1}$  by the following equation:

$$P_{t+1} = \text{softmax}(Wh_t^N), \quad (2)$$

$$h_t^n = f(h_t^{n-1}, h_{t-1}^n), \quad (3)$$

$$h_t^0 = Ex_t, \quad (4)$$

where  $W \in \mathbb{R}^{V \times D_{h^N}}$  is a weight matrix<sup>2</sup>,  $E \in \mathbb{R}^{D_e \times V}$  is a word embedding matrix,  $x_t \in \{0, 1\}^V$  is a one-hot vector of input word  $w_t$  at timestep  $t$ , and  $h_t^n \in \mathbb{R}^{D_{h^n}}$  is the hidden state of the  $n$ -th RNN at timestep  $t$ . We define  $h_t^n$  at timestep  $t = 0$  as a zero vector:  $h_0^n = \mathbf{0}$ . Let  $f(\cdot)$  represent an abstract function of an RNN, which might be the Elman network (Elman, 1990), the Long Short-Term Memory (LSTM) (Hochreiter and Schmidhuber, 1997), the Recurrent Highway Network (RHN) (Zilly et al., 2017), or any other RNN variant. In this research, we stack three LSTM layers based on Merity et al. (2018) because they achieved high performance.

## 3 Language Modeling as Matrix Factorization

Yang et al. (2018) indicated that the training of language models can be interpreted as a matrix

<sup>2</sup>Actually, we apply a bias term in addition to the weight matrix but we omit it to simplify the following discussion.

factorization problem. In this section, we briefly introduce their description. Let word sequence  $w_{1:t}$  be context  $c_t$ . Then we can regard a natural language as a finite set of the pairs of a context and its conditional probability distribution:  $\mathcal{L} = \{(c_1, P^*(X|c_1)), \dots, (c_U, P^*(X|c_U))\}$ , where  $U$  is the number of possible contexts and  $X \in \{0, 1\}^V$  is a variable representing a one-hot vector of a word. Here, we consider matrix  $A \in \mathbb{R}^{U \times V}$  that represents the true log probability distributions and matrix  $H \in \mathbb{R}^{U \times D_{h^N}}$  that contains the hidden states of the final RNN layer for each context  $c_t$ :

$$A = \begin{bmatrix} \log P^*(X|c_1) \\ \log P^*(X|c_2) \\ \dots \\ \log P^*(X|c_U) \end{bmatrix}; H = \begin{bmatrix} h_{c_1}^N \\ h_{c_2}^N \\ \dots \\ h_{c_U}^N \end{bmatrix}. \quad (5)$$

Then we obtain set of matrices  $F(A) = \{A + \Lambda S\}$ , where  $S \in \mathbb{R}^{U \times V}$  is an all-ones matrix, and  $\Lambda \in \mathbb{R}^{U \times U}$  is a diagonal matrix.  $F(A)$  contains matrices that shifted each row of  $A$  by an arbitrary real number. In other words, if we take a matrix from  $F(A)$  and apply the softmax function to each of its rows, we obtain a matrix that consists of true probability distributions. Therefore, for some  $A' \in F(A)$ , training RNN language models is to find the parameters satisfying the following equation:

$$HW^\top = A'. \quad (6)$$

Equation 6 indicates that training RNN language models can also be interpreted as a matrix factorization problem. In most cases, the rank of matrix  $HW^\top$  is  $D_{h^N}$  because  $D_{h^N}$  is smaller than  $V$  and  $U$  in common RNN language models. Thus, an RNN language model cannot express true distributions if  $D_{h^N}$  is much smaller than  $\text{rank}(A')$ .

Yang et al. (2018) also argued that  $\text{rank}(A')$  is as high as vocabulary size  $V$  based on the following two assumptions:

1. Natural language is highly context-dependent. In addition, since we can imagine many kinds of contexts, it is difficult to assume a basis that represents a conditional probability distribution for any contexts. In other words, compressing  $U$  is difficult.
2. Since we also have many kinds of semantic meanings, it is difficult to assume basic meanings that can create all other semantic meanings by such simple operations as addition and subtraction; compressing  $V$  is difficult.

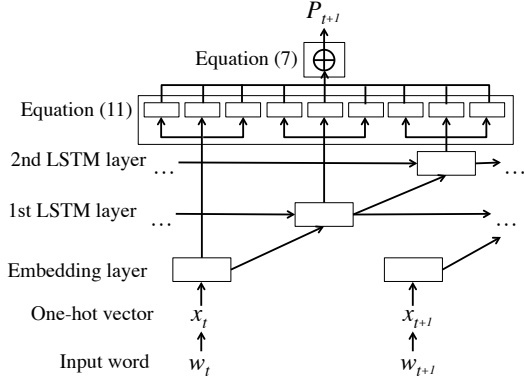


Figure 1: Overview of the proposed method: DOC. This figure represents the example of  $N = 2$  and  $i_0 = i_1 = i_2 = 3$ .

In summary, Yang et al. (2018) indicated that  $D_{h^N}$  is much smaller than  $\text{rank}(A)$  because its scale is usually  $10^2$  and vocabulary size  $V$  is at least  $10^4$ .

#### 4 Proposed Method: Direct Output Connection

To construct a high-rank matrix, Yang et al. (2018) proposed Mixture of Softmaxes (MoS). MoS computes multiple probability distributions from the hidden state of final RNN layer  $h^N$  and regards the weighted average of the probability distributions as the final distribution. In this study, we propose Direct Output Connection (DOC), which is a generalization method of MoS. DOC computes probability distributions from the middle layers in addition to the final layer. In other words, DOC directly connects the middle layers to the output.

Figure 1 shows an overview of DOC, that uses the middle layers (including word embeddings) to compute the probability distributions. Figure 1 computes three probability distributions from all the layers, but we can vary the number of probability distributions for each layer and select some layers to avoid. In our experiments, we search for the appropriate number of probability distributions for each layer.

Formally, instead of Equation 2, DOC computes the output probability distribution at timestep  $t + 1$  by the following equation:

$$P_{t+1} = \sum_{j=1}^J \pi_{j,c_t} \text{softmax}(\tilde{W} k_{j,c_t}), \quad (7)$$

$$s.t. \sum_{j=1}^J \pi_{j,c_t} = 1, \quad (8)$$

where  $\pi_{j,c_t}$  is a weight for each probability distribution,  $k_{j,c_t} \in \mathbb{R}^d$  is a vector computed from each hidden state  $h^n$ , and  $\tilde{W} \in \mathbb{R}^{V \times d}$  is a weight matrix. Thus,  $P_{t+1}$  is the weighted average of  $J$  probability distributions. We define the  $U \times U$  diagonal matrix whose elements are weight  $\pi_{j,c}$  for each context  $c$  as  $\Phi$ . Then we obtain matrix  $\tilde{A} \in \mathbb{R}^{U \times V}$ :

$$\tilde{A} = \log \sum_{j=1}^J \Phi \text{softmax}(K_j \tilde{W}^\top), \quad (9)$$

where  $K_j \in \mathbb{R}^{U \times d}$  is a matrix whose rows are vector  $k_{j,c}$ .  $\tilde{A}$  can be an arbitrary high rank because the righthand side of Equation 9 computes not only the matrix multiplication but also a nonlinear function. Therefore, an RNN language model with DOC can output a distribution matrix whose rank is identical to one of the true distributions. In other words,  $\tilde{A}$  is a better approximation of  $A'$  than the output of a standard RNN language model.

Next we describe how to acquire weight  $\pi_{j,c_t}$  and vector  $k_{j,c_t}$ . Let  $\pi_{c_t} \in \mathbb{R}^J$  be a vector whose elements are weight  $\pi_{j,c_t}$ . Then we compute  $\pi_{c_t}$  from the hidden state of the final RNN layer:

$$\pi_{c_t} = \text{softmax}(W_\pi h_t^N), \quad (10)$$

where  $W_\pi \in \mathbb{R}^{J \times D_{h^N}}$  is a weight matrix. We next compute  $k_{j,c_t}$  from the hidden state of the  $n$ -th RNN layer:

$$k_{j,c_t} = W_j h_t^n, \quad (11)$$

where  $W_j \in \mathbb{R}^{d \times D_{h^n}}$  is a weight matrix. In addition, let  $i_n$  be the number of  $k_{j,c_t}$  from  $h_t^n$ . Then we define the sum of  $i_n$  for all  $n$  as  $J$ ; that is,  $\sum_{n=0}^N i_n = J$ . In short, DOC computes  $J$  probability distributions from all the layers, including the input embedding ( $h^0$ ). For  $i_N = J$ , DOC becomes identical to MoS. In addition to increasing the rank, we expect that DOC weakens the vanishing gradient problem during backpropagation because a middle layer is directly connected to the output, such as with the auxiliary classifiers described in Szegedy et al. (2015).

For a network that computes the weights for several vectors, such as Equation 10, Shazeer et al. (2017) indicated that it often converges to a state where it always produces large weights for few vectors. In fact, we observed that DOC tends to assign large weights to shallow layers. To prevent this phenomenon, we compute the coefficient of

variation of Equation 10 in each mini-batch as a regularization term following Shazeer et al. (2017). In other words, we try to adjust the sum of the weights for each probability distribution with identical values in each mini-batch. Formally, we compute the following equation for a mini-batch consisting of  $w_b, w_{b+1}, \dots, w_{\tilde{b}}$ :

$$B = \sum_{t=b}^{\tilde{b}} \pi_{c_t} \quad (12)$$

$$\beta = \left( \frac{\text{std}(B)}{\text{avg}(B)} \right)^2, \quad (13)$$

where functions  $\text{std}(\cdot)$  and  $\text{avg}(\cdot)$  are functions that respectively return an input’s standard deviation and its average. In the training step, we add  $\lambda_\beta$  multiplied by weight coefficient  $\beta$  to the loss function.

## 5 Experiments on Language Modeling

We investigate the effect of DOC on the language modeling task. In detail, we conduct word-level prediction experiments and show that DOC improves the performance of MoS, which only uses the final layer to compute the probability distributions. Moreover, we evaluate various combinations of layers to explore which combination achieves the best score.

### 5.1 Datasets

We used the Penn Treebank (PTB) (Marcus et al., 1993) and WikiText-2 (Merity et al., 2017) datasets, which are the standard benchmark datasets for the word-level language modeling task. Mikolov et al. (2010) and Merity et al. (2017) respectively published preprocessed PTB<sup>3</sup> and WikiText-2<sup>4</sup> datasets. Table 1 describes their statistics. We used these preprocessed datasets for fair comparisons with previous studies.

### 5.2 Hyperparameters

Our implementation is based on the averaged stochastic gradient descent Weight-Dropped LSTM (AWD-LSTM)<sup>5</sup> proposed by Merity et al. (2018). AWD-LSTM consists of three LSTMs with various regularizations. For the hyperparameters, we used the same values as Yang et al. (2018) except for the

<sup>3</sup><http://www.fit.vutbr.cz/~imikolov/rnnlm/>

<sup>4</sup><https://einstein.ai/research/the-wikitext-long-term-dependency-language-modeling-dataset>

<sup>5</sup><https://github.com/salesforce/awd-lstm-lm>

		PTB	WikiText-2
Vocab		10,000	33,278
#Token	Train	929,590	2,088,628
	Valid	73,761	217,646
	Test	82,431	245,569

Table 1: Statistics of PTB and WikiText-2.

Hyperparameter	PTB	WikiText-2
Learning rate	20	15
Batch size	12	15
Non-monotone interval	60	60
$D_e$	280	300
$D_{h^1}$	960	1150
$D_{h^2}$	960	1150
$D_{h^3}$	620	650
Dropout rate for $x_t$	0.1	0.1
Dropout rate for $h_t^0$	0.4	0.65
Dropout rate for $h_t^1, h_t^2$	0.225	0.2
Dropout rate for $h_t^3$	0.4	0.4
Dropout rate for $k_{j,c_t}$	0.6	0.6
Recurrent weight dropout	0.50	0.50

Table 2: Hyperparameters used for training DOC.

#DOC				$\lambda_\beta$	Valid	Test
$i_3$	$i_2$	$i_1$	$i_0$			
15	0	0	0	0	56.54†	54.44†
20	0	0	0	0	56.88‡	54.79‡
15	0	0	5	0	56.21	54.28
15	0	5	0	0	55.26	53.52
15	5	0	0	0	54.87	53.15
15	5	0	0	0.0001	54.95	53.16
15	5	0	0	0.001	<b>54.62</b>	<b>52.87</b>
15	5	0	0	0.01	55.13	53.39
10	5	0	5	0	56.46	54.18
10	5	5	0	0	56.00	54.37

Table 3: Perplexities of AWD-LSTM with DOC on the PTB dataset. We varied the number of probability distributions from each layer in situation  $J = 20$  except for the top row. The top row (†) represents MoS scores reported in Yang et al. (2018) as a baseline. ‡ represents the perplexity obtained by the implementation of Yang et al. (2018)<sup>6</sup> with identical hyperparameters except for  $i_3$ .

dropout rate for vector  $k_{j,c_t}$  and the non-monotone interval. Since we found that the dropout rate for vector  $k_{j,c_t}$  greatly influences  $\beta$  in Equation 13, we varied it from 0.3 to 0.6 with 0.1 intervals. We selected 0.6 because this value achieved the best score on the PTB validation dataset. For the non-monotone interval, we adopted the same value as Zolna et al. (2018). Table 2 summarizes the hyperparameters of our experiments.

### 5.3 Results

Table 3 shows the perplexities of AWD-LSTM with DOC on the PTB dataset. Each value of columns  $i_n$

<sup>6</sup><https://github.com/zihangdai/mos>

$\lambda_\beta$	Valid	Test
0	0.276	0.279
0.0001	0.254	0.252
0.001	0.217	0.213
0.01	0.092	0.086

Table 4: Coefficient of variation of Equation 10:  $\sqrt{\beta}$  in validation and test sets of PTB.

Model	Valid	Test
AWD-LSTM	401	401
AWD-LSTM-MoS	10000	10000
AWD-LSTM-DOC	10000	10000

Table 5: Rank of output matrix ( $\tilde{A}$  in Equation 9) on the PTB dataset.  $D_3$  of AWD-LSTM is 400.

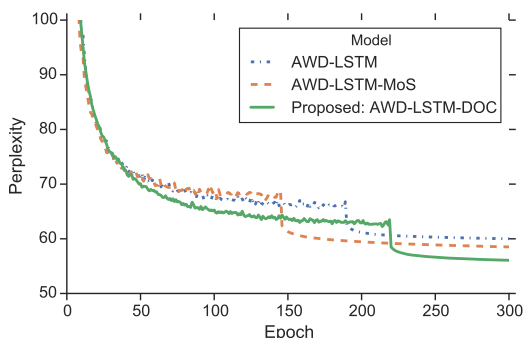


Figure 2: Perplexities of each method on the PTB validation set.

represents the number of probability distributions from hidden state  $h_t^n$ . To find the best combination, we varied the number of probability distributions from each layer by fixing their total to 20:  $J = 20$ . Moreover, the top row of Table 3 shows the perplexity of AWD-LSTM with MoS reported in Yang et al. (2018) for comparison. Table 3 indicates that language models using middle layers outperformed one using only the final layer. In addition, Table 3 shows that increasing the distributions from the final layer ( $i_3 = 20$ ) degraded the score from the language model with  $i_3 = 15$  (the top row of Table 3). Thus, to obtain a superior language model, we should not increase the number of distributions from the final layer; we should instead use the middle layers, as with our proposed DOC.

Table 3 shows that the  $i_3 = 15, i_2 = 5$  setting achieved the best performance and the other settings with shallow layers have a little effect. This result implies that we need some layers to output accurate distributions. In fact, most previous studies adopted two LSTM layers for language modeling. This suggests that we need at least two layers to obtain high-quality distributions.

Model	Valid	Test
AWD-LSTM <sup>†</sup>	58.88	56.36
AWD-LSTM-MoS <sup>†</sup>	56.36	54.26
AWD-LSTM-MoS <sup>‡</sup>	55.67	53.75
AWD-LSTM-DOC	<b>54.62</b>	<b>52.87</b>
AWD-LSTM-DOC (fin)	<b>54.12</b>	<b>52.38</b>

Table 6: Perplexities of our implementations and re-runs on the PTB dataset. We set the non-monotone interval to 60. <sup>†</sup> represents results obtained by original implementations with identical hyperparameters except for non-monotone interval. <sup>‡</sup> indicates the result obtained by our AWD-LSTM-MoS implementation with identical dropout rates as AWD-LSTM-DOC. For (fin), we repeated fine-tuning until convergence.

For the  $i_3 = 15, i_2 = 5$  setting, we explored the effect of  $\lambda_\beta$  in  $\{0, 0.01, 0.001, 0.0001\}$ . Although Table 3 shows that  $\lambda_\beta = 0.001$  achieved the best perplexity, the effect is not consistent. Table 4 shows the coefficient of variation of Equation 10, i.e.,  $\sqrt{\beta}$  in the PTB dataset. This table demonstrates that the coefficient of variation decreases with growth in  $\lambda_\beta$ . In other words, the model trained with a large  $\lambda_\beta$  assigns balanced weights to each probability distribution. These results indicate that it is not always necessary to equally use each probability distribution, but we can acquire a better model in some  $\lambda_\beta$ . Hereafter, we refer to the setting that achieved the best score ( $i_3 = 15, i_2 = 5, \lambda_\beta = 0.001$ ) as AWD-LSTM-DOC.

Table 5 shows the ranks of matrices containing log probability distributions from each method. In other words, Table 5 describes  $\tilde{A}$  in Equation 9 for each method. As shown by this table, the output of AWD-LSTM is restricted to  $D_3$ <sup>7</sup>. In contrast, AWD-LSTM-MoS (Yang et al., 2018) and AWD-LSTM-DOC outputted matrices whose ranks equal the vocabulary size. This fact indicates that DOC (including MoS) can output the same matrix as the true distributions in view of a rank.

Figure 2 illustrates the learning curves of each method on PTB. This figure contains the validation scores of AWD-LSTM, AWD-LSTM-MoS, and AWD-LSTM-DOC at each training epoch. We trained AWD-LSTM and AWD-LSTM-MoS by setting the non-monotone interval to 60, as with AWD-LSTM-DOC. In other words, we used hyperparameters identical to the original ones to train AWD-LSTM and AWD-LSTM-MoS, except for the non-monotone interval. We note that the opti-

<sup>7</sup>Actually, the maximum rank size of an ordinary RNN language model is  $D_N + 1$  when we use a bias term.

Model	#Param	Valid	Test
LSTM (medium) (Zaremba et al., 2014)	20M	86.2	82.7
LSTM (large) (Zaremba et al., 2014)	66M	82.2	78.4
Variational LSTM (medium) (Gal and Ghahramani, 2016)	20M	81.9 ± 0.2	79.7 ± 0.1
Variational LSTM (large) (Gal and Ghahramani, 2016)	66M	77.9 ± 0.3	75.2 ± 0.2
Variational RHN (Zilly et al., 2017)	32M	71.2	68.5
Variational RHN + WT (Zilly et al., 2017)	23M	67.9	65.4
Variational RHN + WT + IOG (Takase et al., 2017)	29M	67.0	64.4
Neural Architecture Search (Zoph and Le, 2017)	54M	-	62.4
LSTM with skip connections (Melis et al., 2018)	24M	60.9	58.3
AWD-LSTM (Merity et al., 2018)	24M	60.0	57.3
AWD-LSTM + Fraternal Dropout (Zolna et al., 2018)	24M	58.9	56.8
AWD-LSTM-MoS (Yang et al., 2018)	22M	56.54	54.44
Proposed method: AWD-LSTM-DOC	23M	<b>54.62</b>	<b>52.87</b>
Proposed method: AWD-LSTM-DOC (fin)	23M	<b>54.12</b>	<b>52.38</b>
Proposed method (ensemble): AWD-LSTM-DOC × 5	114M	<b>49.99</b>	<b>48.44</b>
Proposed method (ensemble): AWD-LSTM-DOC (fin) × 5	114M	<b>48.63</b>	<b>47.17</b>

Table 7: Perplexities of each method on the PTB dataset.

Model	#Param	Valid	Test
Variational LSTM + IOG (Takase et al., 2017)	70M	95.9	91.0
Variational LSTM + WT + AL (Inan et al., 2017)	28M	91.5	87.0
LSTM with skip connections (Melis et al., 2018)	24M	69.1	65.9
AWD-LSTM (Merity et al., 2018)	33M	68.6	65.8
AWD-LSTM + Fraternal Dropout (Zolna et al., 2018)	34M	66.8	64.1
AWD-LSTM-MoS (Yang et al., 2018)	35M	63.88	61.45
Proposed method: AWD-LSTM-DOC	37M	<b>60.97</b>	<b>58.55</b>
Proposed method: AWD-LSTM-DOC (fin)	37M	<b>60.29</b>	<b>58.03</b>
Proposed method (ensemble): AWD-LSTM-DOC × 5	185M	<b>56.14</b>	<b>54.23</b>
Proposed method (ensemble): AWD-LSTM-DOC (fin) × 5	185M	<b>54.91</b>	<b>53.09</b>

Table 8: Perplexities of each method on the WikiText-2 dataset.

mization method converts the ordinary stochastic gradient descent (SGD) into the averaged SGD at the point where convergence almost occurs. In Figure 2, the turning point is the epoch when each method drastically decreases the perplexity. Figure 2 shows that each method similarly reduces the perplexity at the beginning. AWD-LSTM and AWD-LSTM-MoS were slow to decrease the perplexity from 50 epochs. In contrast, AWD-LSTM-DOC constantly decreased the perplexity and achieved a lower value than the other methods with ordinary SGD. Therefore, we conclude that DOC positively affects the training of language modeling.

Table 6 shows the AWD-LSTM, AWD-LSTM-MoS, and AWD-LSTM-DOC results in our configurations. For AWD-LSTM-MoS, we trained our implementation with the same dropout rates as AWD-LSTM-DOC for a fair comparison. AWD-LSTM-DOC outperformed both the original AWD-LSTM-MoS and our implementation. In other words, DOC outperformed MoS.

Since the averaged SGD uses the averaged parameters from each update step, the parameters of the early steps are harmful to the final parameters. Therefore, when the model converges, recent

studies and ours eliminate the history of and then retrains the model. Merity et al. (2018) referred to this retraining process as fine-tuning. Although most previous studies only conducted fine-tuning once, Zolna et al. (2018) argued that two fine-tunings provided additional improvement. Thus, we repeated fine-tuning until we achieved no more improvements in the validation data. We refer to the model as AWD-LSTM-DOC (fin) in Table 6, which shows that repeated fine-tunings improved the perplexity by about 0.5.

Tables 7 and 8 respectively show the perplexities of AWD-LSTM-DOC and previous studies on PTB and WikiText-2<sup>8</sup>. These tables show that AWD-LSTM-DOC achieved the best perplexity. AWD-LSTM-DOC improved the perplexity by almost 2.0 on PTB and 3.5 on WikiText-2 from the state-of-the-art scores. The ensemble technique provided further improvement, as described in pre-

<sup>8</sup>We exclude models that use the statistics of the test data (Grave et al., 2017; Krause et al., 2017) from these tables because we regard neural language models as the basis of NLP applications and consider it unreasonable to know correct outputs during applications, e.g., machine translation. In other words, we focus on neural language models as the foundation of applications although we can combine the method using the statistics of test data with our AWD-LSTM-DOC.

vious studies (Zaremba et al., 2014; Takase et al., 2017), and improved the perplexity by at least 4 points on both datasets. Finally, the ensemble of the repeated finetuning models achieved 47.17 on the PTB test and 53.09 on the WikiText-2 test.

## 6 Experiments on Application Tasks

As described in Section 1, a neural encoder-decoder model can be interpreted as a conditional language model. To investigate the effect of DOC on an encoder-decoder model, we incorporate DOC into the decoder and examine its performance.

### 6.1 Dataset

We conducted experiments on machine translation and headline generation tasks. For machine translation, we used two kinds of sentence pairs (English-German and English-French) in the IWSLT 2016 dataset<sup>9</sup>. The training set respectively contains about 189K and 208K sentence pairs of English-German and English-French. We experimented in four settings: from English to German (En-De), its reverse (De-En), from English to French (En-Fr), and its reverse (Fr-En).

Headline generation is a task that creates a short summarization of an input sentence (Rush et al., 2015). Rush et al. (2015) constructed a headline generation dataset by extracting pairs of first sentences of news articles and their headlines from the annotated English Gigaword corpus (Napoles et al., 2012). They also divided the extracted sentence-headline pairs into three parts: training, validation, and test sets. The training set contains about 3.8M sentence-headline pairs. For our evaluation, we used the test set constructed by Zhou et al. (2017) because the one constructed by Rush et al. (2015) contains some invalid instances, as reported in Zhou et al. (2017).

### 6.2 Encoder-Decoder Model

For the base model, we adopted an encoder-decoder with an attention mechanism described in Kiyono et al. (2017). The encoder consists of a 2-layer bidirectional LSTM, and the decoder consists of a 2-layer LSTM with attention proposed by Luong et al. (2015). We interpreted the layer after computing the attention as the 3rd layer of the decoder. We refer to this encoder-decoder as EncDec. For the hyperparameters, we followed the setting of Kiyono et al. (2017) except for the sizes of hidden

<sup>9</sup><https://wit3.fbk.eu/>

Model	En-De	De-En	En-Fr	Fr-En
EncDec	23.05	28.18	34.37	34.07
EncDec+DOC ( $i_3 = 2$ )	23.62	29.12	36.09	34.41
EncDec+DOC ( $i_3 = i_2 = 2$ )	<b>23.97</b>	<b>29.33</b>	<b>36.11</b>	<b>34.72</b>

Table 9: BLEU scores on test sets in the IWSLT 2016 dataset. We report averages of three runs.

Model	RG-1	RG-2	RG-L
EncDec	46.77	24.87	43.58
EncDec+DOC ( $i_3 = 2$ )	46.91	24.91	43.73
EncDec+DOC ( $i_3 = i_2 = 2$ )	<b>46.99</b>	<b>25.29</b>	<b>43.83</b>
ABS (Rush et al., 2015)	37.41	15.87	34.70
SEASS (Zhou et al., 2017)	46.86	24.58	43.53
Kiyono et al. (2017)	46.34	24.85	43.49

Table 10: ROUGE F1 scores in headline generation test data provided by Zhou et al. (2017). RG in table denotes ROUGE. For our implementations (the upper part), we report averages of three runs.

states and embeddings. We used 500 for machine translation and 400 for headline generation. We constructed a vocabulary set by using Byte-Pair-Encoding<sup>10</sup> (BPE) (Sennrich et al., 2016). We set the number of BPE merge operations at 16K for the machine translation and 5K for the headline generation.

In this experiment, we compare DOC to the base EncDec. We prepared two DOC settings: using only the final layer, that is, a setting that is identical to MoS, and using both the final and middle layers. We used the 2nd and 3rd layers in the latter setting because this case achieved the best performance on the language modeling task in Section 5.3. We set  $i_3 = 2$  and  $i_2 = 2, i_3 = 2$ . For this experiment, we modified a publicly available encode-decoder implementation<sup>11</sup>.

### 6.3 Results

Table 9 shows the BLEU scores of each method. Since an initial value often drastically varies the result of a neural encoder-decoder, we reported the average of three models trained from different initial values and random seeds. Table 9 indicates that EncDec+DOC outperformed EncDec.

Table 10 shows the ROUGE F1 scores of each method. In addition to the results of our implementations (the upper part), the lower part represents the published scores reported in previous studies. For the upper part, we reported the average of three models (as in Table 9). EncDec+DOC outperformed EncDec on all scores. Moreover, EncDec

<sup>10</sup><https://github.com/rsennrich/subword-nmt>

<sup>11</sup><https://github.com/mlpnlp/mlpnlp-nmt/>

outperformed the state-of-the-art method (Zhou et al., 2017) on the ROUGE-2 and ROUGE-L F1 scores. In other words, our baseline is already very strong. We believe that this is because we adopted a larger embedding size than Zhou et al. (2017). It is noteworthy that DOC improved the performance of EncDec even though EncDec is very strong.

These results indicate that DOC positively influences a neural encoder-decoder model. Using the middle layer also yields further improvement because EncDec+DOC ( $i_3 = i_2 = 2$ ) outperformed EncDec+DOC ( $i_3 = 2$ ).

## 7 Experiments on Constituency Parsing

Choe and Charniak (2016) achieved high F1 scores on the Penn Treebank constituency parsing task by transforming candidate trees into a symbol sequence (S-expression) and reranking them based on the perplexity obtained by a neural language model. To investigate the effectiveness of DOC, we evaluate our language models following their configurations.

### 7.1 Dataset

We used the Wall Street Journal of the Penn Treebank dataset. We used the section 2-21 for training, 22 for validation, and 23 for testing. We applied the preprocessing codes of Choe and Charniak (2016)<sup>12</sup> to the dataset and converted a token that appears fewer than ten times in the training dataset into a special token *unk*. For reranking, we prepared 500 candidates obtained by the Charniak parser (Charniak, 2000).

### 7.2 Models

We compare AWD-LSTM-DOC with AWD-LSTM (Merity et al., 2018) and AWD-LSTM-MoS (Yang et al., 2018). We trained each model with the same hyperparameters from our language modeling experiments (Section 5). We selected the model that achieved the best perplexity on the validation set during the training.

### 7.3 Results

Table 11 shows the bracketing F1 scores on the PTB test set. This table is divided into three parts by horizontal lines; the upper part describes the scores by single language modeling based rerankers, the middle part shows the results by ensembling five rerankers, and the lower part repre-

<sup>12</sup><https://github.com/cdg720/emnlp2016>

Model	F1	
	Base	Rerank
Reranking with single model		
Choe and Charniak (2016)	89.7	92.6
AWD-LSTM	89.7	93.2
AWD-LSTM-MoS	89.7	93.2
AWD-LSTM-DOC	89.7	<b>93.3</b>
Reranking with model ensemble		
AWD-LSTM × 5 (ensemble)	89.7	93.4
AWD-LSTM-MoS × 5 (ensemble)	89.7	93.4
AWD-LSTM-DOC × 5 (ensemble)	89.7	<b>93.5</b>
AWD-LSTM-DOC × 5 (ensemble)	91.2	<b>94.29</b>
AWD-LSTM-DOC × 5 (ensemble)	93.12	<b>94.47</b>
State-of-the-art results		
Dyer et al. (2016)	91.7	93.3
Fried et al. (2017) (ensemble)	92.72	94.25
Suzuki et al. (2018) (ensemble)	92.74	94.32
Kitaev and Klein (2018)	95.13	-

Table 11: Bracketing F1 scores on the PTB test set (Section 23). This table includes reranking models trained on the PTB without external data.

sents the current state-of-the-art scores in the setting without external data. The upper part also contains the score reported in Choe and Charniak (2016) that reranked candidates by the simple LSTM language model. This part indicates that our implemented rerankers outperformed the simple LSTM language model based reranker, which achieved 92.6 F1 score (Choe and Charniak, 2016). Moreover, AWD-LSTM-DOC outperformed AWD-LSTM and AWD-LSTM-MoS. These results correspond to the language modeling task.

The middle part shows that AWD-LSTM-DOC also outperformed AWD-LSTM and AWD-LSTM-MoS in the ensemble setting. In addition, we can improve the performance by exchanging the base parser with a stronger one. In fact, we achieved 94.29 F1 score by reranking the candidates from retrained Recurrent Neural Network Grammars (RNNG) (Dyer et al., 2016)<sup>13</sup>, that achieved 91.2 F1 score in our configuration. Moreover, the lowest row of the middle part indicates the result by reranking the candidates from the retrained neural encoder-decoder based parser (Suzuki et al., 2018). Our base parser has two different parts from Suzuki et al. (2018). First, we used the sum of the hidden states of the forward and backward RNNs as the hidden layer for each RNN<sup>14</sup>. Second, we tied the embedding matrix to the weight matrix to compute

<sup>13</sup>The output of RNNG is not in descending order because it samples candidates based on their scores. Thus, we prepared more candidates (i.e., 700) to be able to obtain correct instances as candidates.

<sup>14</sup>We used the deep bidirectional encoder described at <http://opennmt.net/OpenNMT/training/models/> instead of a basic bidirectional encoder.



the probability distributions in the decoder. The retrained parser achieved 93.12 F1 score. Finally, we achieved 94.47 F1 score by reranking its candidates with AWD-LSTM-DOC. We expect that we can achieve even better score by replacing the base parser with the current state-of-the-art one (Kitaev and Klein, 2018).

## 8 Related Work

Bengio et al. (2003) are pioneers of neural language models. To address the curse of dimensionality in language modeling, they proposed a method using word embeddings and a feed-forward neural network (FFNN). They demonstrated that their approach outperformed n-gram language models, but FFNN can only handle fixed-length contexts. Instead of FFNN, Mikolov et al. (2010) applied RNN (Elman, 1990) to language modeling to address the entire given sequence as a context. Their method outperformed the Kneser-Ney smoothed 5-gram language model (Kneser and Ney, 1995; Chen and Goodman, 1996).

Researchers continue to try to improve the performance of RNN language models. Zaremba et al. (2014) used LSTM (Hochreiter and Schmidhuber, 1997) instead of a simple RNN for language modeling and significantly improved an RNN language model by applying dropout (Srivastava et al., 2014) to all the connections except for the recurrent connections. To regularize the recurrent connections, Gal and Ghahramani (2016) proposed variational inference-based dropout. Their method uses the same dropout mask at each timestep. Zolna et al. (2018) proposed fraternal dropout, which minimizes the differences between outputs from different dropout masks to be invariant to the dropout mask. Melis et al. (2018) used black-box optimization to find appropriate hyperparameters for RNN language models and demonstrated that the standard LSTM with proper regularizations can outperform other architectures.

Apart from dropout techniques, Inan et al. (2017) and Press and Wolf (2017) proposed the word tying method (WT), which unifies word embeddings ( $E$  in Equation 4) with the weight matrix to compute probability distributions ( $W$  in Equation 2). In addition to quantitative evaluation, Inan et al. (2017) provided a theoretical justification for WT and proposed the augmented loss technique (AL), which computes an objective probability based on word embeddings. In addition to these regularization

techniques, Merity et al. (2018) used DropConnect (Wan et al., 2013) and averaged SGD (Polyak and Juditsky, 1992) for an LSTM language model. Their AWD-LSTM achieved lower perplexity than Melis et al. (2018) on PTB and WikiText-2.

Previous studies also explored superior architecture for language modeling. Zilly et al. (2017) proposed recurrent highway networks that use highway layers (Srivastava et al., 2015) to deepen recurrent connections. Zoph and Le (2017) adopted reinforcement learning to construct the best RNN structure. However, as mentioned, Melis et al. (2018) established that the standard LSTM is superior to these architectures. Apart from RNN architecture, Takase et al. (2017) proposed the input-to-output gate (IOG), which boosts the performance of trained language models.

As described in Section 3, Yang et al. (2018) interpreted training language modeling as matrix factorization and improved performance by computing multiple probability distributions. In this study, we generalized their approach to use the middle layers of RNNs. Finally, our proposed method, DOC, achieved the state-of-the-art score on the standard benchmark datasets.

Some studies provided methods that boost performance by using statistics obtained from test data. Grave et al. (2017) extended a cache model (Kuhn and De Mori, 1990) for RNN language models. Krause et al. (2017) proposed dynamic evaluation that updates parameters based on a recent sequence during testing. Although these methods might also improve the performance of DOC, we omitted such investigation to focus on comparisons among methods trained only on the training set.

## 9 Conclusion

We proposed *Direct Output Connection* (DOC), a generalization method of MoS introduced by Yang et al. (2018). DOC raises the expressive power of RNN language models and improves quality of the model. DOC outperformed MoS and achieved the best perplexities on the standard benchmark datasets of language modeling: PTB and WikiText-2. Moreover, we investigated its effectiveness on machine translation and headline generation. Our results show that DOC also improved the performance of EncDec and using a middle layer positively affected such application tasks.

## References

- Yoshua Bengio, Réjean Ducharme, Pascal Vincent, and Christian Janvin. 2003. A neural probabilistic language model. *Journal of Machine Learning Research*, 3:1137–1155.
- Eugene Charniak. 2000. A maximum-entropy-inspired parser. In *1st Meeting of the North American Chapter of the Association for Computational Linguistics (NAACL 2000)*, pages 132–139.
- Stanley F. Chen and Joshua Goodman. 1996. An empirical study of smoothing techniques for language modeling. In *Proceedings of the 34th Annual Meeting on Association for Computational Linguistics (ACL 1996)*, pages 310–318.
- Do Kook Choe and Eugene Charniak. 2016. Parsing as language modeling. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing (EMNLP 2016)*, pages 2331–2336.
- Chris Dyer, Adhiguna Kuncoro, Miguel Ballesteros, and Noah A. Smith. 2016. Recurrent neural network grammars. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT 2016)*, pages 199–209.
- Jeffrey L Elman. 1990. Finding Structure in Time. *Cognitive science*, 14(2):179–211.
- Daniel Fried, Mitchell Stern, and Dan Klein. 2017. Improving neural parsing by disentangling model combination and reranking effects. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (ACL 2017)*, pages 161–166.
- Yarin Gal and Zoubin Ghahramani. 2016. A Theoretically Grounded Application of Dropout in Recurrent Neural Networks. In *Advances in Neural Information Processing Systems 29 (NIPS 2016)*.
- Edouard Grave, Armand Joulin, and Nicolas Usunier. 2017. Improving Neural Language Models with a Continuous Cache. In *Proceedings of the 5th International Conference on Learning Representations (ICLR 2017)*.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long Short-Term Memory. *Neural Computation*, 9(8):1735–1780.
- Hakan Inan, Khashayar Khosravi, and Richard Socher. 2017. Tying Word Vectors and Word Classifiers: A Loss Framework for Language Modeling. In *Proceedings of the 5th International Conference on Learning Representations (ICLR 2017)*.
- Nikita Kitaev and Dan Klein. 2018. Constituency parsing with a self-attentive encoder. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (ACL 2018)*, pages 2676–2686.
- Shun Kiyono, Sho Takase, Jun Suzuki, Naoaki Okazaki, Kentaro Inui, and Masaaki Nagata. 2017. Source-side prediction for neural headline generation. *CoRR*.
- Reinhard Kneser and Hermann Ney. 1995. Improved backing-off for m-gram language modeling. In *In Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP 1995)*, pages 181–184.
- Ben Krause, Emmanuel Kahembwe, Iain Murray, and Steve Renals. 2017. Dynamic evaluation of neural sequence models. *CoRR*.
- Roland Kuhn and Renato De Mori. 1990. A cache-based natural language model for speech recognition. 12:570–583.
- Thang Luong, Hieu Pham, and Christopher D. Manning. 2015. Effective approaches to attention-based neural machine translation. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing (EMNLP 2015)*, pages 1412–1421.
- Mitchell P. Marcus, Mary Ann Marcinkiewicz, and Beatrice Santorini. 1993. Building a Large Annotated Corpus of English: The Penn Treebank. *Computational Linguistics*, 19(2):313–330.
- Gábor Melis, Chris Dyer, and Phil Blunsom. 2018. On the state of the art of evaluation in neural language models. *Proceedings of the 6th International Conference on Learning Representations (ICLR 2018)*.
- Stephen Merity, Nitish Shirish Keskar, and Richard Socher. 2018. Regularizing and Optimizing LSTM Language Models. In *Proceedings of the 6th International Conference on Learning Representations (ICLR 2018)*.
- Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. 2017. Pointer Sentinel Mixture Models. In *Proceedings of the 5th International Conference on Learning Representations (ICLR 2017)*.
- Tomas Mikolov, Martin Karafiát, Lukás Burget, Jan Cernocký, and Sanjeev Khudanpur. 2010. Recurrent neural network based language model. In *Proceedings of the 11th Annual Conference of the International Speech Communication Association (INTER-SPEECH 2010)*, pages 1045–1048.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed Representations of Words and Phrases and their Compositionality. In *Advances in Neural Information Processing Systems 26 (NIPS 2013)*, pages 3111–3119.
- Andriy Mnih and Koray Kavukcuoglu. 2013. Learning Word Embeddings Efficiently with Noise-Contrastive Estimation. In *Advances in Neural Information Processing Systems 26 (NIPS 2013)*, pages 2265–2273.

- Courtney Napoles, Matthew Gormley, and Benjamin Van Durme. 2012. Annotated gigaword. In *Proceedings of the Joint Workshop on Automatic Knowledge Base Construction and Web-scale Knowledge Extraction*, pages 95–100.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL 2018)*, pages 2227–2237.
- Boris T Polyak and Anatoli B Juditsky. 1992. Acceleration of Stochastic Approximation by Averaging. *SIAM Journal on Control and Optimization*, 30(4):838–855.
- Ofir Press and Lior Wolf. 2017. Using the Output Embedding to Improve Language Models. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics (EACL 2017)*, pages 157–163.
- Alexander M. Rush, Sumit Chopra, and Jason Weston. 2015. A Neural Attention Model for Abstractive Sentence Summarization. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing (EMNLP 2015)*, pages 379–389.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Improving neural machine translation models with monolingual data. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (ACL 2016)*, pages 86–96.
- Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarz, Andy Davis, Quoc V. Le, Geoffrey E. Hinton, and Jeff Dean. 2017. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. In *Proceedings of the 5th International Conference on Learning Representations (ICLR 2017)*.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(1):1929–1958.
- Rupesh Kumar Srivastava, Klaus Greff, and Jürgen Schmidhuber. 2015. Highway networks. In *Proceedings of the Deep Learning Workshop in ICML 15*.
- Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. 2014. Sequence to Sequence Learning with Neural Networks. In *Advances in Neural Information Processing Systems 27 (NIPS 2014)*, pages 3104–3112.
- Jun Suzuki, Sho Takase, Hidetaka Kamigaito, Makoto Morishita, and Masaaki Nagata. 2018. An empirical study of building a strong baseline for constituency parsing. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (ACL 2018)*, pages 612–618.
- Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. 2015. Going deeper with convolutions. In *Proceedings of the 28th IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2015)*, pages 1–9.
- Sho Takase, Jun Suzuki, and Masaaki Nagata. 2017. Input-to-output gate to improve rnn language models. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (IJCNLP 2017)*, pages 43–48.
- Li Wan, Matthew Zeiler, Sixin Zhang, Yann L Cun, and Rob Fergus. 2013. Regularization of Neural Networks using DropConnect. In *Proceedings of the 30th International Conference on Machine Learning (ICML 2013)*, pages 1058–1066.
- Tsung-Hsien Wen, Milica Gasic, Nikola Mrkšić, Pei-Hao Su, David Vandyke, and Steve Young. 2015. Semantically Conditioned LSTM-based Natural Language Generation for Spoken Dialogue Systems. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing (EMNLP 2015)*, pages 1711–1721.
- Zhilin Yang, Zihang Dai, Ruslan Salakhutdinov, and William W. Cohen. 2018. Breaking the softmax bottleneck: A high-rank RNN language model. In *Proceedings of the 6th International Conference on Learning Representations (ICLR 2018)*.
- Wojciech Zaremba, Ilya Sutskever, and Oriol Vinyals. 2014. Recurrent neural network regularization. In *Proceedings of the 2nd International Conference on Learning Representations (ICLR 2014)*.
- Qingyu Zhou, Nan Yang, Furu Wei, and Ming Zhou. 2017. Selective encoding for abstractive sentence summarization. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (ACL 2017)*, pages 1095–1104.
- Julian Georg Zilly, Rupesh Kumar Srivastava, Jan Koutník, and Jürgen Schmidhuber. 2017. Recurrent Highway Networks. *Proceedings of the 34th International Conference on Machine Learning (ICML 2017)*, pages 4189–4198.
- Konrad Zolna, Devansh Arpit, Dendi Suhubdy, and Yoshua Bengio. 2018. Fraternal dropout. In *Proceedings of the 6th International Conference on Learning Representations (ICLR 2018)*.
- Barret Zoph and Quoc V. Le. 2017. Neural Architecture Search with Reinforcement Learning. In *Proceedings of the 5th International Conference on Learning Representations (ICLR 2017)*.