
Deep Residual Output Layers for Neural Language Generation

Nikolaos Pappas¹ James Henderson¹

Abstract

Many tasks, including language generation, benefit from learning the structure of the output space, particularly when the space of output labels is large and the data is sparse. State-of-the-art neural language models indirectly capture the output space structure in their classifier weights since they lack parameter sharing across output labels. Learning shared output label mappings helps, but existing methods have limited expressivity and are prone to overfitting. In this paper, we investigate the usefulness of more powerful shared mappings for output labels, and propose a deep residual output mapping with dropout between layers to better capture the structure of the output space and avoid overfitting. Evaluations on three language generation tasks show that our output label mapping can match or improve state-of-the-art recurrent and self-attention architectures, and suggest that the classifier does not necessarily need to be high-rank to better model natural language if it is better at capturing the structure of the output space.

1. Introduction

Learning the structure of the output space benefits a wide variety of tasks, such as object recognition and novelty detection in images (Weston et al., 2011; Socher et al., 2013; Frome et al., 2013; Zhang et al., 2016; Chen et al., 2018a), zero-shot prediction in texts (Dauphin et al., 2014; Yazdani & Henderson, 2015; Nam et al., 2016; Rios & Kavuluru, 2018), and structured prediction in either images or text (Srikumar & Manning, 2014a; Dyer et al., 2015; Belanger & McCallum, 2016; Graber et al., 2018). When the space of output labels is large or their data is sparse, treating labels as independent classes makes learning difficult, because identifying one label is not helped by data for other labels. This problem can be addressed by learning output label em-

beddings to capture the similarity structure of the output label space, so that data for similar labels can help classification, even to the extent of enabling few-shot or even zero-shot classification. This approach has been particularly successful in natural language generation tasks, where word embeddings give a useful similarity structure for next-word-prediction in tasks such as machine translation (Vaswani et al., 2017) and language modeling (Merity et al., 2017).

Existing neural language models typically use a log-linear classifier to predict words (Vaswani et al., 2017; Chen et al., 2018b). We can view the output label weights as a word embedding, and the input encoder as mapping the context to a vector in the same embedding space. Then the similarity between these two embeddings in this joint input-label space is measured with a dot product followed by the softmax function. We will refer to this part as the classifier, distinct from the input encoder which only depends on the input and the label encoder which only depends on the label. To improve performance and reduce model size, sometimes the output label weights are tied to the input word embedding vectors (Inan et al., 2016; Press & Wolf, 2017), but there is no parameter sharing taking place across different words, which limits the effective transfer between them.

Recent work has shown improvements over specific vanilla recurrent architectures by sharing parameters across outputs through a bilinear mapping on neural language modeling (Gulordava et al., 2018) or a dual nonlinear mapping on neural machine translation (Pappas et al., 2018), which can make the classifier more powerful. However, the shallow modeling constraints and the lack of regularization capabilities limit their applicability on arbitrary tasks and model architectures. Orthogonal to these studies, Yang et al. (2018) achieved state-of-the-art improvements on language modeling by increasing the power of the classifier using a mixture of softmax functions, albeit at the expense of computational efficiency. A natural question arises of whether one can make the classifier more powerful by simply increasing the power of the label mapping while using a single softmax function without modifying its dimensionality or rank.

In this paper, we attempt to answer this question by investigating alternative neural architectures for learning the embedding of an output label in the joint input-label space which address the aforementioned limitations. In particular,

¹Idiap Research Institute, Martigny, Switzerland. Correspondence to: Nikolaos Pappas <nikolaos.pappas@idiap.ch>.

we propose a deep residual nonlinear output mapping from word embeddings to the joint input-output space, which better captures the output structure while it avoids overfitting with two different dropout strategies between layers, and preserves useful information with residual connections to the word embeddings and, optionally, to the outputs of previous layers.¹ For the rest of the model, we keep the same input encoder architecture and still use the dot product and softmax function for output label prediction.

We demonstrate on language modeling and machine translation that we can match or improve state-of-the-art recurrent and self-attention architectures by simply increasing the power of the output mapping, while using a single softmax operation and without changing the dimensionality or rank of the classifier. The results suggest that the classifier does not necessarily need to be high rank to better model language if it better captures the output space structure. Further analysis reveals the significance of different model components and improvements on predicting low frequency words.

2. Background: Neural Language Generation

The output layer of neural models for language generation tasks such as language modeling (Bengio et al., 2003; Mikolov & Zweig, 2012; Merity et al., 2017), machine translation (Bahdanau et al., 2015; Luong et al., 2015; Johnson et al., 2017) and summarization (Rush et al., 2015; Paulus et al., 2018), typically consists of a linear unit with a weight matrix $\mathbf{W} \in \mathbb{R}^{d_h \times |\mathcal{V}|}$ and a bias vector $\mathbf{b} \in \mathbb{R}^{|\mathcal{V}|}$ followed by a softmax activation function, where \mathcal{V} is the vocabulary. Thus, at a given time t , the output probability distribution for the current output \mathbf{y}_t conditioned on the inputs i.e. the previous outputs, $\mathbf{y}_1^{t-1} = (\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_{t-1})$ with $\mathbf{y}_i \in \{0, 1\}^{|\mathcal{V}|} : \sum_j^{|\mathcal{V}|} \mathbf{y}_{ij} = 1 \forall i \in \mathcal{N}$, is defined as:

$$p(\mathbf{y}_t | \mathbf{y}_1^{t-1}) \propto \exp(\mathbf{W}^T \mathbf{h}_t + \mathbf{b}), \quad (1)$$

where \mathbf{h}_t is the input encoder’s hidden representation at time t with d_h dimensions. The parameterisation in Eq. 1 makes it difficult to learn the structure of the output space or to transfer this information from one label to another because the parameters for output label i , \mathbf{W}_i^T , are independent from the parameters for any other output label j , \mathbf{W}_j^T .

2.1. Weight Tying

Learning the structure of the output space can be helped by learning it jointly with the structure of input word embeddings, but this still does not support the transfer of learned information across output labels. In particular, since the output labels are words and thus the output parameters \mathbf{W}^T have one row per word, it is common to tie these parameters

¹Our code and settings are available at <http://github.com/idiap/drill>.

with those of the input word embeddings $\mathbf{E} \in \mathbb{R}^{|\mathcal{V}| \times d}$, by setting $\mathbf{W} = \mathbf{E}^T$ (Inan et al., 2016; Press & Wolf, 2017). Making this substitution in Eq 1, we obtain:

$$p(\mathbf{y}_t | \mathbf{y}_1^{t-1}) \propto \exp(\mathbf{E} \mathbf{h}_t + \mathbf{b}) \quad (2)$$

Although there is no explicit transfer across outputs, this parameterisation can implicitly learn the output structure, as can be seen if we assume an implicit factorization of the input embeddings, $\mathbf{E} \approx \mathbf{E}_1 \mathbf{W}_1$ as in (Mikolov et al., 2013).

2.2. Bilinear Mapping

The above bilinear form, excluding the bias, is similar to the form of joint input-output space learning models (Yazdani & Henderson, 2015; Nam et al., 2016) which have been proposed in the context of zero-shot text classification. This motivates the learning of explicit relationships across outputs and inputs through parameter sharing via \mathbf{W}_1 as above. By substituting this factorization in Eq 2, we obtain:

$$p(\mathbf{y}_t | \mathbf{y}_1^{t-1}) \propto \exp(\mathbf{E}_1 \mathbf{W}_1 \mathbf{h}_t + \mathbf{b}) \quad (3)$$

where $\mathbf{W}_1 \in \mathbb{R}^{d \times d_h}$ is the bilinear mapping and \mathbf{E} , \mathbf{h}_t are the output embeddings and the encoded input respectively, as above. This parametrization has been previously also proposed by Gulordava et al. (2018) for language modeling albeit with a different motivation, namely to decouple the hidden state from the word embedding prediction.

2.3. Dual Nonlinear Mapping

Another existing output layer parameterisation which explicitly learns the structure of the output is from (Pappas et al., 2018). Specifically, two nonlinear functions, $g_{out}(\cdot)$ and $g_{in}(\cdot)$, are introduced which aim to capture the output and context structure respectively:

$$p(\mathbf{y}_t | \mathbf{y}_1^{t-1}) \propto \exp(g_{out}(\mathbf{E}) g_{in}(\mathbf{h}_t) + \mathbf{b}), \quad (4)$$

$$\propto \exp(\sigma(\mathbf{E}\mathbf{U} + \mathbf{b}_u) \sigma(\mathbf{V}\mathbf{h}_t + \mathbf{b}_v) + \mathbf{b}) \quad (5)$$

where $\sigma(\cdot)$ is a nonlinear activation function such as ReLU or Tanh, the matrix $\mathbf{U} \in \mathbb{R}^{d \times d_j}$ and bias $\mathbf{b}_u \in \mathbb{R}^{d_j}$ are the linear projection of the encoded outputs, and the matrix $\mathbf{V} \in \mathbb{R}^{d_j \times d_h}$ and bias $\mathbf{b}_v \in \mathbb{R}^{d_j}$ are the linear projection of the context, and $\mathbf{b} \in \mathbb{R}^{|\mathcal{V}|}$ captures the biases of the target outputs in the vocabulary.

The parameterisation of Eq. 5 enables learning a more rich output structure than the bilinear mapping of Eq. 3 because it learns nonlinear relationships. Both, however, allow for controlling the capacity of the output layer independently of the dimensionality of the context \mathbf{h}_t and the word embedding \mathbf{E} , by increasing the breadth of the joint projection, e.g. the dimensionality of the \mathbf{U} and \mathbf{V} matrices in Eq. 5 above. This increased capacity can be seen in the inequalities below for the number of parameters of the output layers

discussed so far, assuming a fixed $|\mathcal{V}|, d, d_h$:

$$C_{tied} < C_{bilinear} \leq C_{dual} \leq C_{base}, \quad (6)$$

where C_{tied} , C_{base} , $C_{bilinear}$ and C_{dual} respectively correspond to the number of dedicated parameters of an output layer with (Eq. 2) and without (Eq. 1) weight tying, using the bilinear mapping (Eq. 3) and the dual nonlinear mapping (Eq. 5) which are assumed to be nonzero except C_{tied} .

Given this analysis, we identify and aim to address the following limitations of the previously proposed output layer parameterisations for language generation:

- (a) *Shallow modeling of the label space.* Output labels are mapped into the joint space with a single (possibly nonlinear) projection. Its power can only be increased by increasing the dimensionality of the joint space.
- (b) *Tendency to overfit.* Increasing the dimensionality of the joint space and thus the power of the output classifier can lead to undesirable effects such as overfitting in certain language generation tasks, which limits its applicability to arbitrary domains.

3. Deep Residual Output Layers

To address the aforementioned limitations we propose a deep residual output layer architecture for neural language generation which performs deep modeling of the structure of the output space while it preserves acquired information and avoids overfitting. Our formulation adopts the general form and the basic principles of previous output layer parametrizations which aim to capture the output structure explicitly in Section 2.3, namely (i) learning rich output structure, (ii) controlling the output layer capacity independently of the dimensionality of the vocabulary, the encoder and the word embedding, and, lastly, (iii) avoiding costly label-set-size dependent parameterisations.

3.1. Overview

A general overview of the proposed architecture for neural language generation is displayed in Fig. 1. We base our output layer formulation starting on the general form of the dual nonlinear mapping of Eq. 4:

$$p(\mathbf{y}_t | \mathbf{y}_1^{t-1}) \propto \exp(g_{out}(\mathbf{E})g_{in}(\mathbf{h}_t) + \mathbf{b}). \quad (7)$$

The input network $g_{in}(\cdot)$ takes as input a sequence of words represented by their input word embeddings \mathbf{E} which have been encoded in a context representation \mathbf{h}_t for the given time step t . The output or label network $g_{out}(\cdot)$ takes as input the word(s) describing each possible output label and encodes them in a label embedding $\mathbf{E}^{(k)}$ where k is the depth of the label encoder network. Next, we define these two proposed networks, and then we discuss how the model is trained and how it relates to previous output layers.

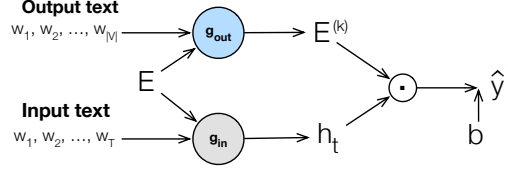


Figure 1. General overview of the proposed architecture.

3.2. Label Encoder Network

For language generation tasks, the output labels are each a word in the vocabulary \mathcal{V} . We assume that these labels are represented with their associated word embedding, which is a row in \mathbf{E} . In general, there may be additional information about each label, such as dictionary entries, cross-lingual resources, or contextual information, in which case we can add an initial encoder for these descriptions which outputs a label embedding matrix $\mathbf{E}' \in \mathbb{R}^{|\mathcal{V}| \times d}$. In this paper we make the simplifying assumption that $\mathbf{E}' = \mathbf{E}$ and leave the investigation of additional label information to future work.

3.2.1. LEARNING OUTPUT STRUCTURE

To obtain a label representation which is able to encode rich output space structure, we define the $g_{out}(\cdot)$ function to be a deep neural network with k layers which takes the label embedding E as input and outputs its deep label mapping at the last layer, $g_{out}(\mathbf{E}) = \mathbf{E}^{(k)}$, as follows:

$$\mathbf{E}^{(k)} = f_{out}^{(k)}(\mathbf{E}^{(k-1)}), \quad (8)$$

where k is the depth of the network and each function $f_{out}^{(i)}(\cdot)$ at the i_{th} layer is a nonlinear projection of the following form:

$$f_{out}^{(i)}(\mathbf{E}^{(i-1)}) = \sigma(\mathbf{E}^{(i-1)}\mathbf{U}^{(i)} + \mathbf{b}_u^{(i)}), \quad (9)$$

where $\sigma(\cdot)$ is a nonlinear activation function such as ReLU or Tanh, and the matrix $\mathbf{U}^{(i)} \in \mathbb{R}^{d \times d_i}$ and the bias $\mathbf{b}_u^{(i)} \in \mathbb{R}^{d_i}$ are the linear projection of the encoded outputs at the i_{th} layer. Note that when we restrict the above label network to have one layer depth the projection is equivalent to the label mapping from previous work in Eq. 5.

3.2.2. PRESERVING INFORMATION

The multiple layers of projections in Eq. 8 force the relationship between word embeddings \mathbf{E} and label embeddings $\mathbf{E}^{(k)}$ to be highly nonlinear. To preserve useful information from the original word embeddings and to facilitate the learning of the label network we add a skip connection directly to the input embedding. Optionally, for very deep label networks, we also add a residual connection to previous layers as in (He et al., 2016). With these additions the projection at the k_{th} layer becomes:

$$\mathbf{E}^{(k)} = f_{out}^{(k)}(\mathbf{E}^{(k-1)}) + \mathbf{E}^{(k-1)} + \mathbf{E} \quad (10)$$

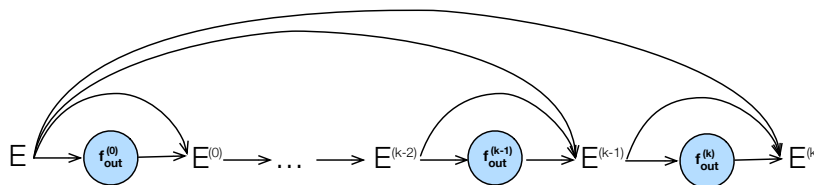


Figure 2. The proposed deep residual label network architecture for neural language generation. Straight lines represent the input to a function and curved lines represent shortcut or residual connections implying addition operations.

3.2.3. CONTROLLING OUTPUT NETWORK POWER

We can characterize the power of the proposed output network in terms of its number of parameters Θ_{drill} , including the proposed label encoder and output classifier:

$$C_{drill} \approx |\Theta_{drill}| = k \times (d \times d) + |\mathcal{V}|. \quad (11)$$

By controlling the depth of the label encoder we can make the number of parameters equal to that of each other output network. For weight tying, C_{tied} , this is $k = 0$, for full linear weights, C_{base} , this is $k = \frac{|\mathcal{V}|}{d}$, for a bilinear mapping, $C_{bilinear}$, this is $k = 1$, and for a dual nonlinear mapping, C_{joint} , this is $k = \frac{2d_j}{d}$. Hence, the power of the output network can be adjusted freely depending on the task at hand within the full spectrum of options defined by Ineq. 6.

3.2.4. AVOIDING OVERFITTING

The ability to increase power may be useful for high-resource data regimes, however it can lead to overfitting in the case we are in a low-resource data regime. To make sure that our network is robust to both data availability regimes, we choose to apply standard (Srivastava et al., 2014) or variational (Gal & Ghahramani, 2016b) dropout in between each of the k layers of the projection. Assuming $\delta(\cdot)$ to be the dropout mask sampling function, the above goal is achieved by modifying the function $f_{out}^{(i)}(\cdot)$ at the i th layer from Eq 9 as follows:

$$f_{out}^{(i)}(\mathbf{E}^{(i-1)}) = \delta(f_{out}^{(i)}(\mathbf{E}^{(i-1)})) \odot f_{out}^{(i)}(\mathbf{E}^{(i-1)}). \quad (12)$$

In standard dropout, a new binary dropout mask is sampled every time the dropout function is called. This means that new dropout masks are sampled independently for each dimension of each different label representation. In contrast, variational dropout samples a binary dropout mask only once upon the first call and then repeatedly uses that locked dropout mask for all label representations within the forward and backward pass.

3.3. Context Network

The context representation \mathbf{h}_t in most language generation tasks is typically the output of a deep neural network, and thus it can capture, in principle, the nonlinear structure of

the dual nonlinear mapping in Section 2.3. Eq. 2.3 has an additional nonlinearity $g_{in}(\cdot)$ in order to allow the dimensionality of the joint space to be larger than that of the context encoder’s output \mathbf{h}_t . However, in our proposed model we increase the power of the output network by increasing the depth of the label encoder, keeping the size of the joint space fixed. Thus, for our models, we make the simplifying assumption that there is no additional nonlinearity after the context encoder, setting $g_{in}(\cdot) = \mathbf{I}$.

3.4. Training Considerations

To perform maximum likelihood estimation of the model parameters, we use the negative log-likelihood of the data as our training objective. This involves computing the conditional likelihood of predicting the next word, as explained above. The normalized exponential function we use for converting the network scores to probability estimates is the typical softmax activation function.

In principle, our output layer parameterisation requires more computations than a typical softmax linear unit, with or without weight tying. Hence, it tends to get slower as the depth of the label encoder or the size of the vocabulary increases. In case either of them becomes extremely large, we can resort to recent sampling-based or hierarchical softmax approximation methods such as the ones proposed by Jean et al. (2015) and Grave et al. (2017).²

3.5. Relation to Previous Output Layer Forms

Our output layer parameterisation has the same general form as the one with the dual nonlinear mapping in Eq. 4. Hence, it preserves the property of being a generalization of output layers based on bilinear mapping and weight tying described in Section 2.3. The bilinear form in Eq. 3 can be simply derived from the general form of Eq 7 if we restrict the output mapping depth to be equal to one, set its bias equal to zero, and make the $\sigma(\cdot)$ activation function linear, so we have $\mathbf{U}^{(0)} = \mathbf{W}_1$. By further setting the matrix \mathbf{U} to be the identity matrix, we can also derive the output layer form based on weight tying in Eq. 2.

²Note that in practice, for our experiments with vocabularies up to 32K, we did not need to resort to a softmax approximation.

Model	#Param	Validation	Test
Mikolov & Zweig (2012) – RNN-LDA + KN-5 + cache	9M [‡]	-	92.0
Zaremba et al. (2014) – LSTM	20M	86.2	82.7
Gal & Ghahramani (2016a) – Variational LSTM (MC)	20M	-	78.6
Kim et al. (2016) – CharCNN	19M	-	78.9
Merity et al. (2016) – Pointer Sentinel-LSTM	21M	72.4	70.9
Grave et al. (2016) – LSTM + continuous cache pointer [†]	-	-	72.1
Inan et al. (2016) – Tied Variational LSTM + augmented loss	24M	75.7	73.2
Zilly et al. (2017) – Variational RHN	23M	67.9	65.4
Zoph & Le (2016) – NAS Cell	25M	-	64.0
Melis et al. (2017) – 2-layer skip connection LSTM	24M	60.9	58.3
Merity et al. (2017) – AWD-LSTM w/o finetune	24M	60.7	58.8
Merity et al. (2017) – AWD-LSTM	24M	60.0	57.3
Ours – AWD-LSTM-DRILL w/o finetune	24M	59.6	57.0
Ours – AWD-LSTM-DRILL	24M	58.2	55.7
Merity et al. (2017) – AWD-LSTM + continuous cache pointer [†]	24M	53.9	52.8
Krause et al. (2018) – AWD-LSTM + dynamic evaluation [†]	24M	51.6	51.1
Ours – AWD-LSTM-DRILL + dynamic evaluation [†]	24M	49.5	49.4
Yang et al. (2018) – AWD-LSTM-MoS	22M	56.54	54.44
Yang et al. (2018) – AWD-LSTM-MoS + dynamic evaluation [†]	22M	48.33	47.69

Table 1. Model perplexity with a single softmax (upper part) and multiple softmaxes (lower part) on validation and test sets on Penn Treebank. Baseline results are obtained from Merity et al. (2017) and Krause et al. (2018). † indicates the use of dynamic evaluation.

4. Experiments

We evaluate on three language generation tasks. The first two tasks are standard language modeling tasks, i.e. predicting the next word given the sequence of previous words. The third task is a conditional language modeling task, namely neural machine translation, i.e. predicting the next word in the target language given the source sentence and the previous words in the translation. To demonstrate the generality of the proposed output mapping we incorporate it in three different neural architectures which are considered state-of-the-art for their corresponding tasks.

4.1. Language Modeling

Datasets and Metrics. Following previous work in language modeling (Yang et al., 2018; Krause et al., 2018; Merity et al., 2017; Melis et al., 2017), we evaluate the proposed model in terms of perplexity on two widely used language modeling datasets, namely Penn Treebank (Mikolov et al., 2010) and WikiText-2 (Merity et al., 2017) which have vocabularies of 10,000 and 33,278 words, respectively. For fair comparison, we use the same regularization and optimization techniques with Merity et al. (2017).

Model Configuration. To compare with the state-of-the-art we use the proposed output layer within the best architecture by Merity et al. (2017), which is a highly regularized 3-layer LSTM with 400-dimensional embeddings and 1150-dimensional hidden states, noted as AWD-LSTM. Our hyper-parameters were optimized based on validation perplexity, as follows: 4-layer label encoder depth, 400-dimensional label embeddings, 0.6 dropout rate, residual

connection to \mathbf{E} , uniform weight initialization in the interval $[-0.1, 0.1]$, for both datasets, and, furthermore, *sigmoid* activation and *variational* dropout for PennTreebank, as well as *relu* activation and *standard* dropout for Wikitext-2. The rest of the hyper-parameters were set to the optimal ones found for each dataset by Merity et al. (2017).

For the implementation of the AWD-LSTM we used the language modeling toolkit in Pytorch provided by Merity et al. (2017),³ and for the dynamic evaluation the code in Pytorch provided by Krause et al. (2018).⁴

4.1.1. RESULTS

The results in terms of perplexity for our models, denoted by DRILL, and several competitive baselines, are displayed in Table 1 for PennTreebank and Table 2 for Wikitext-2. For the single-softmax models (above the double lines), for both datasets, our models improve over the state-of-the-art by +1.6 perplexity on PennTreebank and by +3.9 perplexity on Wikitext-2. Moreover, when our model is combined with the dynamic evaluation approach proposed by Krause et al. (2018), it improves even more over these models by +1.7 on PennTreebank and by +2.3 on Wikitext-2.

In contrast to other more complicated previous models, our model uses a standard LSTM architecture, following the work of Merity et al. (2017); Melis et al. (2017). For instance, Zilly et al. (2017) uses of a recurrent highway network which is an extension of an LSTM to allow multiple

³<http://github.com/salesforce/awd-lstm-lm>

⁴<http://github.com/benkrause/dynamic-evaluation>

Deep Residual Output Layers for Neural Language Generation

Model	#Param	Validation	Test
Inan et al. (2016) – Variational LSTM + augmented loss	28M	91.5	87.0
Grave et al. (2016) – LSTM + continuous cache pointer [†]	-	-	68.9
Melis et al. (2017) – 2-layer skip connection LSTM	24M	69.1	65.9
Merity et al. (2017) – AWD-LSTM w/o finetune	33M	69.1	66.0
Merity et al. (2017) – AWD-LSTM	33M	68.6	65.8
Ours – AWD-LSTM-DRILL w/o finetune	34M	65.7	62.8
Ours – AWD-LSTM-DRILL	34M	64.9	61.9
Merity et al. (2017) – AWD-LSTM + continuous cache pointer [†]	33M	53.8	52.0
Krause et al. (2018) – AWD-LSTM + dynamic evaluation [†]	33M	46.4	44.3
Ours – AWD-LSTM-DRILL + dynamic evaluation [†]	34M	43.9	42.0
Yang et al. (2018) – AWD-LSTM-MoS	35M	63.88	61.45
Yang et al. (2018) – AWD-LSTM-MoS + dynamical evaluation [†]	35M	42.41	40.68

Table 2. Model perplexity with a single softmax (upper part) and multiple softmaxes (lower part) on validation and test sets on WikiText-2. Baseline results are obtained from Merity et al. (2017) and Krause et al. (2018). † indicates the use of dynamic evaluation.

hidden state updates per time step, Zoph & Le (2016) uses reinforcement learning to generate an RNN cell which is even more complicated than an LSTM cell, and Merity et al. (2016) makes use of a probabilistic mixture model which combines a typical language model with a pointer network which reproduces words from the recent context.

Interestingly, our model also significantly reduces the performance gap against multiple softmax models. In particular, when our finetuned model is compared to the corresponding mixture-of-softmaxes (MoS) model, which makes use of 15 softmaxes in the classifier, it reduces the difference against AWD-LSTM from 2.8 to 1.2 points on PennTreebank and from 4.3 to 0.4 points on WikiText-2. When our model is compared to MoS with dynamic evaluation, the difference is reduced from 3.4 points to 1.7 points on PennTreebank and from 3.6 to 1.3 on WikiText-2. Note that the rank of the log-probability matrix for MoS on PennTreebank is 9,981, while for AWD-LSTM and our model the rank is only 400. This observation questions the high-rank hypothesis of MoS, which states that the log-probability matrix has to be high rank to better capture language. Our results suggest that the log-probability matrix does not need to be high rank if the classifier is better at capturing the output space structure.

Furthermore, as shown in Table 3, the MoS model is far slower than AWD-LSTM, even for these small datasets and reduced dimensionality settings,⁵ whereas adding our label encoder to AWD-LSTM results in only a small speed difference. In particular, on PennTreebank the MoS model takes about 139 seconds per epoch while AWD-LSTM about 47 seconds per epoch, which makes it slower by a factor of 3.0×, whereas our model is only 1.1× slower than this

⁵ Note that even though the MoS models have a comparable number of parameters to the other models, they use smaller values for several crucial hyper-parameters, such as word embedding size, hidden state size and batch size, likely to make the training speed more manageable and avoid overfitting.

Model	PennTreebank	Wikitext-2
AWD-LSTM	47 sec (1.0×)	89 sec (1.0×)
AWD-LSTM-DRILL	53 sec (1.1×)	106 sec (1.2×)
AWD-LSTM-MoS	139 sec (3.0×)	862 sec (9.7×)

Table 3. Average time taken per epoch on the two datasets: PennTreebank ($|\mathcal{V}| \approx 20K$) and Wikitext-2 ($|\mathcal{V}| \approx 33K$).

baseline. On Wikitext-2, the differences are even more pronounced due to the larger size of the vocabulary. The MoS model takes about 862 seconds per epoch while AWD-LSTM takes about 89 seconds per epoch, which makes it slower by a factor of 9.7×, whereas our model with 4-layers is only 1.2× slower than the baseline. We attempted to combine our label encoder with the MoS model, but its training speed exceeded our computation budget.

Overall, these results demonstrate that the proposed deep residual output mapping improves significantly the state-of-the-art single-softmax neural architecture for language modeling, namely AWD-LSTM, without hurting its efficiency. Hence, it could be a useful and practical addition to other existing architectures. In addition, our model remains competitive against models based on multiple softmaxes and could be combined with them in the future, since our work is orthogonal to using multiple softmaxes. To demonstrate that our model is also applicable to larger datasets as well, in Section 4.2 below we apply our method to neural machine translation. But before moving to that experiment, we first perform an ablation analysis of these results.

4.1.2. ABLATION ANALYSIS

To give further insights into the source of the improvement from our output layer parameterisation, in Table 4 we compare its ablated variants with previous output layer parameterisations. Each alternative is combined with the state-of-the-art encoder network AWD-LSTM (Merity et al., 2017). We observe that full softmax produces the highest perplexity

Output Layer	#Param	Validation	Test
Full softmax	43.8M	69.9	66.8
Weight tying [PW17]	24.2M	60.0	57.3
Bilinear map. [G18]	24.3M	60.7	58.5
Dual nonlinear map. [PH18]	24.5M	58.8	56.4
DRILL 1-layer	24.3M	58.8	56.2
DRILL 2-layers	24.5M	58.7	56.0
DRILL 3-layers	24.7M	58.5	55.9
DRILL 4-layers	24.8M	58.2	55.7
+ residuals between layers	24.8M	59.6	57.5
- no variational dropout	24.8M	63.4	60.7

Table 4. Ablation results and comparison with previous output layers when using AWD-LSTM (Merity et al., 2017) as an encoder network on PennTreebank.

scores, despite having almost 20M parameters more than the other models. This shows that the power of the output layer or classifier, as measured by number of parameters, is not indicative of generalization ability.

The output layer with weight tying (Press & Wolf, 2017), noted [PW17], has lower perplexity than the full softmax by 9.5 points. The bilinear mapping (Gulordava et al., 2018), noted [G18], has lower perplexity than the full softmax by 8.3 points, but it is still higher than weight tying by 1.2 points. The dual nonlinear mapping (Pappas et al., 2018), noted [PH18], has even lower perplexity compared to the full softmax by 10.4 points, and has lower perplexity than weight tying by 0.9 points.⁶ DRILL with only 1-layer depth is slightly better than [PH18], and with 2-layers depth outperforms all previous output mappings, improving over full softmax by 10.8 points, weight tying by 1.3 points, and dual non-linear mapping by 0.4 points. Increasing the depth even more provides further improvements of up to 0.3 points. This shows the benefits of learning deep output label mappings, as opposed to shallower ones. Lastly, DRILL with residual connections between layers has an increase of 1.8 perplexity points, likely because of an effective reduction in depth, and not using variational dropout has a significant increase in perplexity of namely 5 points, which highlights the importance of regularization between layers for this task.

To verify the hypothesis that our output layer facilitates information transfer across words, we also analyzed the loss for words in different frequency bands, created by computing statistics on the training set. Figure 3 displays the mean relative cross-entropy difference (%) between our output layer and the previous output layers for the different word frequency bands on the test set of PennTreebank. Overall, the graph shows that most of the improvements in perplexity between 5% to 17.5% brought by DRILL against baselines comes from predicting more accurately the words in lower word frequency bands (1 to 100 occurrences). The results

⁶For fair comparison, we also used dropout and residual connections to E and h_t when they lead to better validation performance.

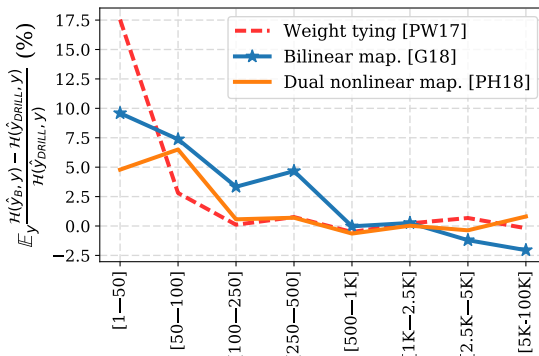


Figure 3. Mean relative cross-entropy loss difference (%) between each baseline output layer (B) and our output layer (DRILL) computed over different word frequency intervals on PennTreebank.

are consistent with Table 4, since the second best output layer is the one with the bilinear mapping followed by the bilinear mapping and weight tying baselines. One exception occurs in the highest frequency band, where DRILL has 2.5% higher perplexity than the bilinear mapping, but this difference is less significant because it is computed based on 16 unique words as opposed to the lowest frequency band which corresponds to 4116 unique words. These results validate our hypothesis that learning a deeper label encoder leads to better transfer of learned information across labels. More specifically, because low frequency words lack data to individually learn the complex structure of the output space, transfer of learned information from other words is crucial to improving performance, whereas this is not the case for higher frequency words. This analysis suggests that our model could also be useful for zero-resource scenarios, where labels need to be predicted without any training data, similarly to other joint input-output space models.

4.2. Neural Machine Translation

Dataset and Metrics. Following previous work in neural machine translation (Vaswani et al., 2017), we train on the WMT 2014 English-German dataset with ~ 4.5 M sentence pairs, using the Newstest2013 set for validation and the Newstest2014 set for testing. We pre-process the texts using the BPE algorithm (Sennrich et al., 2016) with 32K operations. Following the standard evaluation practices in the field (Bojar et al., 2017), the translation quality is measured using BLEU score (Papineni et al., 2002) on *tokenized* text.

Model configuration. We compare against the state-of-the-art Transformer (base) architecture from Vaswani et al. (2017) with a 6-layer encoder and decoder depth, 512-dimensional word embeddings, 2048-dimensional hidden feed-forward states and 8 heads.⁷ Our hyper-parameters were optimized based on validation accuracy, as follows: 2-layer label encoder depth, 512-dimensional label embed-

⁷We chose the base model because it can be trained much faster than the big model (12 hours vs 3.5 days), for efficiency reasons.

dings, 0.0 dropout rate, *sigmoid* activation function, residual connection to **E**, and uniform weight initialization in $[-0.1, 0.1]$. The rest of the hyper-parameters were set to the optimal ones in (Vaswani et al., 2017), except that we did not perform model averaging over last 5 for Transformer (base) model. To ensure fair comparison, we trained the Transformer (base) from scratch for the same number of training steps as ours, namely 350K, and thereby reproduced about the same score as in (Vaswani et al., 2017) with a slight difference of +0.1 point. For the implementation of the Transformer, we used OpenNMT (Klein et al., 2017).⁸

4.3. Results

The results displayed in Table 5 show that our model, namely Transformer-DRILL (base) with 79.9M parameters, outperforms the Transformer (base) model with 79.4M parameters by 0.8 points, and is only 0.3 points behind the Transformer (big) model which has 240 parameters due to its increased dimensionality. This result almost matches the single-model state-of-the-art, without resorting to very high capacity encoders or model averaging over different epochs. Transformer-DRILL also outperforms by 0.6 points our implementation of Transformer (base) model combined with the dual nonlinear mapping by Pappas et al. (2018), highlighting once more the importance of deeper label mappings. Note that our improvement is noticeable even when the vocabulary is based on sub-word units (Sennrich et al., 2016), instead of regular word units as in Section 4.1.1.

Lastly, our model even surpasses the performance of some ensemble models such as GNMT + RL and ConvS2S. The RNMT+ model is marginally better than Transformer (big) even though it has two layers deeper decoder and more powerful layers, namely bidirectional LSTMs instead of self-attention. RNMT+ cascaded and multicol are ensemble architectures which combine LSTMs with self-attention in different ways and increase the overall model complexity even more while providing marginal gains over simpler architectures. Combining our output layer with Transformer (big) should, in principle, make this difference even smaller.

5. Other Related Work

Several studies focus on learning the structure of the output space from texts for zero-shot classification (Dauphin et al., 2014; Nam et al., 2016; Rios & Kavuluru, 2018; Pappas & Henderson, 2019) and structured prediction (Srikumar & Manning, 2014b; Dyer et al., 2015; Yeh et al., 2018). Fewer such studies exist for neural language generation, for instance the ones described in Section 2. Their mappings can increase the power of the classifier by controlling its dimensionality or rank, but unlike ours, they have limited

⁸<http://github.com/OpenNMT/OpenNMT-py>

Model	BLEU
Bidirectional GRU (Sennrich et al., 2016)	22.8
ByteNet (Kalchbrenner et al., 2016)	23.7
GNMT + RL (Johnson et al., 2017)	24.6
ConvS2S (Gehring et al., 2017)	25.1
MoE (Shazeer et al., 2017)	26.0
GNMT + RL Ensemble (Johnson et al., 2017)	26.3
ConvS2S Ensemble (Gehring et al., 2017)	26.3
Transformer (base) (Vaswani et al., 2017)	27.3
Transformer-Dual (base) [PH18]	27.5
Ours – Transformer-DRILL (base)	28.1
Transformer (big) (Vaswani et al., 2017)	28.4
RNMT+ (Chen et al., 2018b)	28.5
RNMT+ cascaded (Chen et al., 2018b)	28.6
RNMT+ multicol (Chen et al., 2018b)	28.8

Table 5. Translation results in terms of BLEU on English to German with a 32K BPE vocabulary.

expressivity and a tendency to overfit. Yang et al. (2018) showed that the softmax layer which is low-rank creates a ‘bottleneck’ problem, i.e. limits model expressivity, and increased the classifier rank by using a mixture of softmaxes. Takase et al. (2018) improved MoS by computing the mixture based on the last and the middle recurrent layers. Two alternative ways to increase the classifier rank are obtained by multiplying the softmax with a non-parametric sigmoid function (Kanai et al., 2018), and by learning parametric monotonic functions on top of the logits (Ganea et al., 2019). Both of these methods have close to or higher perplexity than ours without using MoS, even though we keep the rank or power of the classifier the same. Instead, we specifically increase the power of the output label encoder, and the obtained results suggest that the classifier does not necessarily need to be high-rank to better capture language.

6. Conclusion

Typical log-linear classifiers for neural language modeling tasks can be significantly improved by learning a deep residual output label encoding, regardless of the input encoding architecture. Deeper representations of the output structure lead to better transfer across the output labels, especially the low-resource ones. The results on three tasks show that the proposed output layer parameterisation can match or improve state-of-the-art context encoding architectures and outperform previous output layer parameterisations based on a joint input-output space, while preserving their basic principles and generality. Our findings should apply on other conditional neural language modeling tasks, such as image captioning and summarization. As future work, it would be interesting to learn from more elaborate descriptions or contextualized representations of the output labels and investigate their transferability in different tasks.

Acknowledgements

This work was supported by the European Union through SUMMA project (n. 688139) and the Swiss National Science Foundation within INTERPID project (FNS-30106).

References

- Bahdanau, D., Cho, K., and Bengio, Y. Neural machine translation by jointly learning to align and translate. In *Proceedings of the 5th International Conference on Learning Representations*, San Diego, CA, USA, 2015. URL <https://arxiv.org/pdf/1409.0473.pdf>.
- Belanger, D. and McCallum, A. Structured prediction energy networks. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48, ICML'16*, pp. 983–992. JMLR.org, 2016. URL <http://dl.acm.org/citation.cfm?id=3045390.3045495>.
- Bengio, Y., Ducharme, R., Vincent, P., and Janvin, C. A neural probabilistic language model. *J. Mach. Learn. Res.*, 3:1137–1155, March 2003. ISSN 1532-4435. URL <http://dl.acm.org/citation.cfm?id=944919.944966>.
- Bojar, O., Chatterjee, R., Federmann, C., Graham, Y., Hadrow, B., Huang, S., Huck, M., Koehn, P., Liu, Q., Logacheva, V., Monz, C., Negri, M., Post, M., Rubino, R., Specia, L., and Turchi, M. Findings of the 2017 conference on machine translation (wmt17). In *Proceedings of the Second Conference on Machine Translation, Volume 2: Shared Task Papers*, pp. 169–214, Copenhagen, Denmark, September 2017. Association for Computational Linguistics. URL <http://www.aclweb.org/anthology/W17-4717>.
- Chen, L., Zhang, H., Xiao, J., Liu, W., and Chang, S.-F. Zero-shot visual recognition using semantics-preserving adversarial embedding networks. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1043–1052, 2018a. URL http://openaccess.thecvf.com/content_cvpr_2018/papers/Chen_Zero-Shot_Visual_Recognition_CVPR_2018_paper.pdf.
- Chen, M. X., Firat, O., Bapna, A., Johnson, M., Macherey, W., Foster, G., Jones, L., Schuster, M., Shazeer, N., Parmar, N., Vaswani, A., Uszkoreit, J., Kaiser, L., Chen, Z., Wu, Y., and Hughes, M. The best of both worlds: Combining recent advances in neural machine translation. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 76–86, Melbourne, Australia, July 2018b. Association for Computational Linguistics. URL <http://www.aclweb.org/anthology/P18-1008>.
- Dauphin, Y. N., Tür, G., Hakkani-Tür, D., and Heck, L. P. Zero-shot learning and clustering for semantic utterance classification. In *International Conference on Learning Representations*, Banff, Canada, 2014. URL <http://arxiv.org/abs/1401.0509>.
- Dyer, C., Ballesteros, M., Ling, W., Matthews, A., and Smith, N. A. Transition-based dependency parsing with stack long short-term memory. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pp. 334–343, Beijing, China, July 2015. Association for Computational Linguistics. doi: 10.3115/v1/P15-1033. URL <https://www.aclweb.org/anthology/P15-1033>.
- Frome, A., Corrado, G. S., Shlens, J., Bengio, S., Dean, J., Ranzato, M. A., and Mikolov, T. Devise: A deep visual-semantic embedding model. In Burges, C. J. C., Bottou, L., Welling, M., Ghahramani, Z., and Weinberger, K. Q. (eds.), *Advances in Neural Information Processing Systems 26*, pp. 2121–2129. Curran Associates, Inc., 2013. URL <http://papers.nips.cc/paper/5204-devise-a-deep-visual-semantic-embedding-model.pdf>.
- Gal, Y. and Ghahramani, Z. A theoretically grounded application of dropout in recurrent neural networks. In *Proceedings of the 30th International Conference on Neural Information Processing Systems, NIPS'16*, pp. 1027–1035, USA, 2016a. Curran Associates Inc. ISBN 978-1-5108-3881-9. URL <http://dl.acm.org/citation.cfm?id=3157096.3157211>.
- Gal, Y. and Ghahramani, Z. A theoretically grounded application of dropout in recurrent neural networks. In *Proceedings of the 30th International Conference on Neural Information Processing Systems, NIPS'16*, pp. 1027–1035, USA, 2016b. Curran Associates Inc. ISBN 978-1-5108-3881-9. URL <http://dl.acm.org/citation.cfm?id=3157096.3157211>.
- Ganea, O., Gelly, S., Bécigneul, G., and Severyn, A. Breaking the softmax bottleneck via learnable monotonic pointwise non-linearities. *CoRR*, abs/1902.08077, 2019. URL <http://arxiv.org/abs/1902.08077>.
- Gehring, J., Auli, M., Grangier, D., Yarats, D., and Dauphin, Y. N. Convolutional sequence to sequence learning. In Precup, D. and Teh, Y. W. (eds.), *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 1243–1252, International

- Convention Centre, Sydney, Australia, 06–11 Aug 2017. PMLR. URL <http://proceedings.mlr.press/v70/gehring17a.html>.
- Graber, C., Meshi, O., and Schwing, A. Deep structured prediction with nonlinear output transformations. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 31*, pp. 6320–6331. Curran Associates, Inc., 2018. URL <http://papers.nips.cc/paper/7869-deep-structured-prediction-with-nonlinear-output-transformations.pdf>.
- Grave, E., Joulin, A., and Usunier, N. Improving neural language models with a continuous cache. *CoRR*, abs/1612.04426, 2016. URL <http://arxiv.org/abs/1612.04426>.
- Grave, É., Joulin, A., Cissé, M., Grangier, D., and Jégou, H. Efficient softmax approximation for GPUs. In Precup, D. and Teh, Y. W. (eds.), *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 1302–1310, International Convention Centre, Sydney, Australia, 06–11 Aug 2017. PMLR. URL <http://proceedings.mlr.press/v70/grave17a.html>.
- Gulordava, K., Aina, L., and Boleda, G. How to represent a word and predict it, too: Improving tied architectures for language modelling. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 2936–2941, Brussels, Belgium, October–November 2018. Association for Computational Linguistics. URL <http://www.aclweb.org/anthology/D18-1323>.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778, June 2016. doi: 10.1109/CVPR.2016.90. URL <https://ieeexplore.ieee.org/document/7780459>.
- Inan, H., Khosravi, K., and Socher, R. Tying word vectors and word classifiers: A loss framework for language modeling. *CoRR*, abs/1611.01462, 2016. URL <http://arxiv.org/abs/1611.01462>.
- Jean, S., Cho, K., Memisevic, R., and Bengio, Y. On using very large target vocabulary for neural machine translation. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pp. 1–10, Beijing, China, July 2015. Association for Computational Linguistics. URL <http://www.aclweb.org/anthology/P15-1001>.
- Johnson, M., Schuster, M., Le, Q., Krikun, M., Wu, Y., Chen, Z., Thorat, N., Vidas, F. a., Wattenberg, M., Corrado, G., Hughes, M., and Dean, J. Google’s multilingual neural machine translation system: Enabling zero-shot translation. *Transactions of the Association for Computational Linguistics*, 5:339–351, 2017. ISSN 2307-387X. URL <https://transacl.org/ojs/index.php/tacl/article/view/1081>.
- Kalchbrenner, N., Espeholt, L., Simonyan, K., van den Oord, A., Graves, A., and Kavukcuoglu, K. Neural machine translation in linear time. *CoRR*, abs/1610.10099, 2016. URL <http://arxiv.org/abs/1610.10099>.
- Kanai, S., Fujiwara, Y., Yamanaka, Y., and Adachi, S. Sigsoftmax: Reanalysis of the softmax bottleneck. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 31*, pp. 286–296. Curran Associates, Inc., 2018. URL <http://papers.nips.cc/paper/7312-sigsoftmax-reanalysis-of-the-softmax-bottleneck.pdf>.
- Kim, Y., Jernite, Y., Sontag, D., and Rush, A. M. Character-aware neural language models. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, AAAI’16, pp. 2741–2749. AAAI Press, 2016. URL <http://dl.acm.org/citation.cfm?id=3016100.3016285>.
- Klein, G., Kim, Y., Deng, Y., Senellart, J., and Rush, A. Opennmt: Open-source toolkit for neural machine translation. In *Proceedings of ACL 2017, System Demonstrations*, pp. 67–72. Association for Computational Linguistics, 2017. URL <http://aclweb.org/anthology/P17-4012>.
- Krause, B., Kahembwe, E., Murray, I., and Renals, S. Dynamic evaluation of neural sequence models. In Dy, J. and Krause, A. (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 2766–2775, Stockholm, Sweden, 10–15 Jul 2018. PMLR. URL <http://proceedings.mlr.press/v80/krause18a.html>.
- Luong, T., Pham, H., and Manning, C. D. Effective approaches to attention-based neural machine translation. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pp. 1412–1421,

- Lisbon, Portugal, September 2015. Association for Computational Linguistics. URL <http://aclweb.org/anthology/D15-1166>.
- Melis, G., Dyer, C., and Blunsom, P. On the state of the art of evaluation in neural language models. *CoRR*, abs/1707.05589, 2017. URL <http://arxiv.org/abs/1707.05589>.
- Merity, S., Xiong, C., Bradbury, J., and Socher, R. Pointer sentinel mixture models. *CoRR*, abs/1609.07843, 2016. URL <http://arxiv.org/abs/1609.07843>.
- Merity, S., Keskar, N. S., and Socher, R. Regularizing and optimizing LSTM language models. *CoRR*, abs/1708.02182, 2017. URL <http://arxiv.org/abs/1708.02182>.
- Mikolov, T. and Zweig, G. Context dependent recurrent neural network language model. In *2012 IEEE Spoken Language Technology Workshop (SLT)*, pp. 234–239, Dec 2012. doi: 10.1109/SLT.2012.6424228. URL <https://ieeexplore.ieee.org/document/6424228>.
- Mikolov, T., Karafit, M., Burget, L., Cernock, J., and Khudanpur, S. Recurrent neural network based language model. In Kobayashi, T., Hirose, K., and Nakamura, S. (eds.), *INTERSPEECH*, pp. 1045–1048. ISCA, 2010. URL <http://dblp.uni-trier.de/db/conf/interspeech/interspeech2010.html#MikolovKBCK10>.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., and Dean, J. Distributed representations of words and phrases and their compositionality. In Burges, C. J. C., Bottou, L., Welling, M., Ghahramani, Z., and Weinberger, K. Q. (eds.), *Advances in Neural Information Processing Systems 26*, pp. 3111–3119. Curran Associates, Inc., 2013. URL <http://papers.nips.cc/paper/5021-distributed-representations-of-words-and-phrases-and-their-compositionality.pdf>.
- Nam, J., Mencía, E. L., and Fürnkranz, J. All-in text: Learning document, label, and word representations jointly. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence*, pp. 1948–1954, Phoenix, AR, USA, 2016. URL <https://www.aaai.org/ocs/index.php/AAAI/AAAI16/paper/view/12058>.
- Papineni, K., Roukos, S., Ward, T., and Zhu, W.-J. BLEU: a method for automatic evaluation of machine translation. In *Proceedings of 40th Annual Meeting of the Association for Computational Linguistics*, pp. 311–318, Philadelphia, Pennsylvania, USA, July 2002. Association for Computational Linguistics. doi: 10.3115/1073083.1073135. URL <http://www.aclweb.org/anthology/P02-1040>.
- Pappas, N. and Henderson, J. Gile: A generalized input-label embedding for text classification. *Transactions of the Association for Computational Linguistics*, 7:139–155, 2019. doi: 10.1162/tacl.a.00259. URL <https://doi.org/10.1162/tacl.a.00259>.
- Pappas, N., Miculicich, L., and Henderson, J. Beyond weight tying: Learning joint input-output embeddings for neural machine translation. In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pp. 73–83. Association for Computational Linguistics, 2018. URL <http://aclweb.org/anthology/W18-6308>.
- Paulus, R., Xiong, C., and Socher, R. A deep reinforced model for abstractive summarization. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=HkAClQgA->.
- Press, O. and Wolf, L. Using the output embedding to improve language models. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pp. 157–163, Valencia, Spain, April 2017. Association for Computational Linguistics. URL <http://www.aclweb.org/anthology/E17-2025>.
- Rios, A. and Kavuluru, R. Few-shot and zero-shot multi-label learning for structured label spaces. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 3132–3142, Brussels, Belgium, October–November 2018. Association for Computational Linguistics. URL <https://www.aclweb.org/anthology/D18-1352>.
- Rush, A. M., Chopra, S., and Weston, J. A neural attention model for abstractive sentence summarization. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pp. 379–389, Lisbon, Portugal, 2015. URL <https://aclanthology.coli.uni-saarland.de/papers/D15-1044/d15-1044>.
- Sennrich, R., Haddow, B., and Birch, A. Neural machine translation of rare words with subword units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 1715–1725, Berlin, Germany, 2016. URL <https://aclanthology.coli.uni-saarland.de/papers/P16-1162/p16-1162>.
- Shazeer, N., Mirhoseini, A., Maziarz, K., Davis, A., Le, Q. V., Hinton, G. E., and Dean, J. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. *CoRR*, abs/1701.06538, 2017. URL <http://arxiv.org/abs/1701.06538>.

- Socher, R., Ganjoo, M., Manning, C. D., and Ng, A. Y. Zero-shot learning through cross-modal transfer. In *Proceedings of the 26th International Conference on Neural Information Processing Systems, NIPS'13*, pp. 935–943, Lake Tahoe, Nevada, 2013. URL <https://papers.nips.cc/paper/5027-zero-shot-learning-through-cross-modal-transfer>.
- Srikumar, V. and Manning, C. D. Learning distributed representations for structured output prediction. In Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N. D., and Weinberger, K. Q. (eds.), *Advances in Neural Information Processing Systems 27*, pp. 3266–3274. Curran Associates, Inc., 2014a. URL <http://papers.nips.cc/paper/5323-learning-distributed-representations-for-structured-output-prediction.pdf>.
- Srikumar, V. and Manning, C. D. Learning distributed representations for structured output prediction. In *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2, NIPS'14*, pp. 3266–3274, Cambridge, MA, USA, 2014b. MIT Press. URL <http://dl.acm.org/citation.cfm?id=2969033.2969191>.
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15:1929–1958, 2014. URL <http://jmlr.org/papers/v15/srivastava14a.html>.
- Takase, S., Suzuki, J., and Nagata, M. Direct output connection for a high-rank language model. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 4599–4609, Brussels, Belgium, October–November 2018. Association for Computational Linguistics. URL <https://www.aclweb.org/anthology/D18-1489>.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L. u., and Polosukhin, I. Attention is all you need. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 30*, pp. 5998–6008. Curran Associates, Inc., 2017. URL <http://papers.nips.cc/paper/7181-attention-is-all-you-need.pdf>.
- Weston, J., Bengio, S., and Usunier, N. WSABIE: Scaling up to large vocabulary image annotation. In *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence (Volume 3)*, pp. 2764–2770, Barcelona, Spain, 2011. ISBN 978-1-57735-515-1. URL <https://ai.google/research/pubs/pub37180>.
- Yang, Z., Dai, Z., Salakhutdinov, R., and Cohen, W. W. Breaking the softmax bottleneck: A high-rank RNN language model. In *Proceedings of the 6th International Conference on Learning Representations*, volume abs/1711.03953, Vancouver, Canada, 2018. URL <http://arxiv.org/abs/1711.03953>.
- Yazdani, M. and Henderson, J. A model of zero-shot learning of spoken language understanding. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pp. 244–249. Association for Computational Linguistics, 2015. doi: 10.18653/v1/D15-1027. URL <http://aclweb.org/anthology/D15-1027>.
- Yeh, C., Wu, W., Ko, W., and Wang, Y. F. Learning deep latent spaces for multi-label classification. In *In Proceedings of the 32nd AAAI Conference on Artificial Intelligence*, New Orleans, USA, 2018. URL <https://arxiv.org/abs/1707.00418>.
- Zaremba, W., Sutskever, I., and Vinyals, O. Recurrent neural network regularization. *CoRR*, abs/1409.2329, 2014. URL <https://arxiv.org/abs/1409.2329>.
- Zhang, Y., Gong, B., and Shah, M. Fast zero-shot image tagging. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, Las Vegas, USA, 2016. URL <https://arxiv.org/abs/1605.09759>.
- Zilly, J. G., Srivastava, R. K., Koutník, J., and Schmidhuber, J. Recurrent highway networks. In Precup, D. and Teh, Y. W. (eds.), *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 4189–4198, International Convention Centre, Sydney, Australia, 06–11 Aug 2017. PMLR. URL <http://proceedings.mlr.press/v70/zilly17a.html>.
- Zoph, B. and Le, Q. V. Neural architecture search with reinforcement learning. *CoRR*, abs/1611.01578, 2016. URL <http://arxiv.org/abs/1611.01578>.