



The Importance of the Current Input in Sequence Modeling

Christian Oliva^(✉)  and Luis F. Lago-Fernández 

Universidad Autónoma de Madrid, Madrid 28049, Spain
{christian.oliva,luis.lago}@uam.es

Abstract. The last advances in sequence modeling are mainly based on deep learning approaches. The current state of the art involves the use of variations of the standard LSTM architecture, combined with several adjustments that improve the final prediction rates of the trained neural networks. However, in some cases, these adaptations might be too much tuned to the particular problems being addressed. In this article, we show that a very simple idea, to add a direct connection between the input and the output, skipping the recurrent module, leads to an increase of the prediction accuracy in sequence modeling problems related to natural language processing. Experiments carried out on different problems show that the addition of this kind of connection to a recurrent network always improves the results, regardless of the architecture and training-specific details. When this idea is introduced into the models that lead the field, the resulting networks achieve a new state-of-the-art perplexity in language modeling problems.

1 Introduction

Deep learning models constitute the current state of the art in most artificial intelligence applications, from computer vision to robotics or medicine. When dealing with sequential data, Recurrent Neural Networks (RNNs), specially those architectures with gating mechanisms such as the LSTM [7], the GRU [3] and other variants, are usually the default choice. One of the most interesting applications of RNNs is related to the field of Natural Language Processing, where most tasks, such as machine translation, document summarization or language modeling, involve the manipulation of sequences of textual data. Of these, language modeling has been extensively used to test different innovations in recurrent architectures, mainly due to the ease of obtaining very large datasets that can be used to train neural networks with millions of parameters.

Sequence modeling consists of predicting the next element in a sequence given the past history. In language modeling, the sequence is a text, and hence the task is to predict the next word or the next character. In this context, some of the best performing architectures include the Mogrifier LSTM [12] and different variations of the Averaged SGD Weight-Drop (AWD) LSTM [13], usually combined with dynamic evaluation and Mixture of Softmaxes (MoS) [6, 20]. These models

obtain the best state-of-the-art performance with moderate size datasets, such as the Penn Treebank [15] or the Wikitext-2 [14] corpora, when no additional data are used during training. When larger datasets are considered, or when external data are used to pre-train the networks, attention-based architectures usually outperform other models [2, 18].

In this work we use moderate-scale language modeling datasets to explore the effect of a mechanism recently proposed by [16], when combined with different LSTM-based models in the language modeling context. The idea consists of modifying a recurrent architecture by introducing a direct connection between the input and the output of the recurrent module. This has been shown to improve both the model’s generalization results and its readability in simple tasks related to the recognition of regular languages.

In a standard RNN, the output depends only on the network’s hidden state, h_t , which in turn depends on both the input, x_t , and the recent past, h_{t-1} . But there is no explicit dependence of the network’s output on its input. In some cases this could be a shortcoming, since the transformation of x_t needed to compute the network’s internal state is not necessarily the most appropriate to compute the output. However, an explicit dependence of the output on x_t can be forced by adding a *dual* connection that skips the recurrent layers. We claim that this strategy may be of general application in RNN models.

To test our hypothesis we perform a thorough comparison of several state-of-the-art RNN architectures, with and without the *dual* connection, on the Penn Treebank (PTB) and the Wikitext-2 (WT2) datasets. Our results show that, under all experimental conditions, the *dual* architectures outperform their non-dual counterparts. In addition, the Mogrifier-LSTM enhanced with a *dual* connection establishes a new state-of-the-art word-level perplexity for the Penn Treebank dataset when no additional data are used to train the models.

The remainder of the article is organized as follows. First, in Sect. 2, we present the different models we have used and the two possible architectures, the standard recurrent architecture and the *dual* architecture. In Sect. 3, we describe the datasets and the experimental setup. In Sect. 4, we present our results. And finally, in Sect. 5, we extract some conclusions and discuss further lines of research.

2 Models

We start by presenting the standard recurrent architecture which is common to all the models. In absence of a *dual* connection, the basic architecture involves an embedding layer, a recurrent layer and a fully-connected layer with *softmax* activation:

$$e_t = W^{ex} x_t \tag{1}$$

$$h_t = REC(e_t, S_{t-1}) \tag{2}$$

$$y_t = softmax(W^{yh} h_t + b^y), \tag{3}$$

where W^{**} and b^* are weight matrices and biases, respectively, and x_t is the input vector at time t . The *REC* module represents an arbitrary recurrent layer, with S_{t-1} being a set of vectors describing its internal state at the previous time step. In the most general case, this module will simply be an LSTM cell, but we consider other possibilities as well, as described below.

The *dual* architecture introduces an additional layer, with ReLU activation, which is fed with both the output of the embedding layer and the output of the recurrent module:

$$e_t = W^{ex}x_t \quad (4)$$

$$h_t = REC(e_t, S_{t-1}) \quad (5)$$

$$d_t = ReLU(W^{de}e_t + W^{dh}h_t + b^d) \quad (6)$$

$$y_t = softmax(W^{yd}d_t + b^y). \quad (7)$$

This way the network's input can reach the softmax layer following two different paths, through the recurrent layer and through the *dual* connection. In the following we consider different forms for the recurrent module in Eq. 2 and 5.

2.1 The LSTM Module

In the simplest approach the recurrent module consists of an LSTM cell, where the internal state includes both the output and the memory, $S_t = \{h_t; c_t\}$, which are computed as follows:

$$f_t = \sigma(W^{fe}e_t + W^{fh}h_{t-1} + b^f) \quad (8)$$

$$i_t = \sigma(W^{ie}e_t + W^{ih}h_{t-1} + b^i) \quad (9)$$

$$o_t = \sigma(W^{oe}e_t + W^{oh}h_{t-1} + b^o) \quad (10)$$

$$z_t = \tanh(W^{ze}e_t + W^{zh}h_{t-1} + b^z) \quad (11)$$

$$c_t = f_t \odot c_{t-1} + i_t \odot z_t \quad (12)$$

$$h_t = o_t \odot \tanh(c_t), \quad (13)$$

where, as before, W^{**} are weight matrices and b^* are bias vectors. The \odot operator denotes an element-wise product, and σ is the logistic sigmoid function. For convenience, we summarize the joint effect of Eqs. 8–13 as:

$$h_t = LSTM(e_t, \{h_{t-1}; c_{t-1}\}). \quad (14)$$

In the literature it is quite common to stack several LSTM layers. Here we consider a double-layer LSTM, where the output h_t of the recurrent module is obtained by the concatenated application of two LSTM layers:

$$h'_t = LSTM_1(e_t, \{h'_{t-1}; c'_{t-1}\}) \quad (15)$$

$$h_t = LSTM_2(h'_t, \{h_{t-1}; c_{t-1}\}). \quad (16)$$

We refer to this double LSTM module as *dLSTM*:

$$h_t = dLSTM(e_t, \{h_{t-1}; c_{t-1}; h'_{t-1}; c'_{t-1}\}) \quad (17)$$

$$= LSTM_2(LSTM_1(e_t, \{h'_{t-1}; c'_{t-1}\}), \{h_{t-1}; c_{t-1}\}). \quad (18)$$

2.2 The Mogrifier-LSTM Module

The Mogrifier-LSTM [12] is one of the state-of-the-art variations of the standard LSTM architecture achieving the lowest perplexity scores in language modeling tasks. It basically consists of a standard LSTM block, but the input e_t and the hidden state h_{t-1} are transformed before entering Eqs. 8–13. The mogrifier transformation involves several steps where e_t and h_{t-1} modulate each other:

$$e_t^i = 2\sigma(Q^i h_{t-1}^{i-1}) \odot e_t^{i-2}, \quad \text{for odd } i \in \{1, 2, \dots, r\} \quad (19)$$

$$h_{t-1}^i = 2\sigma(R^i e_t^{i-1}) \odot h_{t-1}^{i-2}, \quad \text{for even } i \in \{1, 2, \dots, r\}, \quad (20)$$

where Q^i and R^i are weight matrices and we have $e_t^{-1} = e_t$ and $h_{t-1}^0 = h_{t-1}$. The linear transformations $Q^i h_{t-1}^{i-1}$ and $R^i e_t^{i-1}$ can also include the addition of a bias vector, which has been omitted for the sake of clarity. The constant r is a hyperparameter whose value defines the number of rounds of the transformation. We refer to this recurrent module, including the mogrifier transformation and the subsequent application of the LSTM layer, as:

$$h_t = mLSTM(e_t, \{h_{t-1}; c_{t-1}\}) = LSTM(e_t^*, \{h_{t-1}^*; c_{t-1}\}), \quad (21)$$

where e_t^* and h_{t-1}^* are the highest indexed e_t^i and h_{t-1}^i in Eq. 19 and 20. Note that the choice $r = 0$ recovers the standard LSTM model.

The original work also explored the use of a double-layer LSTM enhanced with the mogrifier transformation. This strategy can be summarized as follows:

$$h_t = mdLSTM(e_t, \{h_{t-1}; c_{t-1}; h'_{t-1}; c'_{t-1}\}) \quad (22)$$

$$= mLSTM_2(mLSTM_1(e_t, \{h'_{t-1}; c'_{t-1}\}), \{h_{t-1}; c_{t-1}\}). \quad (23)$$

3 Experiments

3.1 Datasets

We perform experiments on two datasets: the Penn Treebank (PTB) corpus [11], as preprocessed by [15], and the WikiText-2 (WT2) dataset [14]. In both cases, the data are used without any additional preprocessing.

The Penn Treebank dataset has been widely used in the literature to experiment with language modeling. The standard data preprocessing is due to [15], and includes transformation of all letters to lower case, elimination of

punctuation symbols, and replacement of all numbers with a special token. The vocabulary is limited to the 10,000 most frequent words. The data is split into a training set which contains almost 930,000 tokens, and validation and test sets with around 80,000 words each.

The WikiText-2 dataset, introduced by [14], is a more realistic benchmark for language modeling tasks. It consists of more than 2 million words extracted from Wikipedia articles. The training, validation and test sets contain around 2,125,000, 220,000, and 250,000 words, respectively. The vocabulary includes over 30,000 words, and the data retain capitalization, punctuation, and numbers.

3.2 Experimental Setup

All the considered models follow one of the two architectures discussed in Sect. 2, either the Embedding-Recurrent-Softmax (ERS) architecture (Eqs. 1–3) or the *dual* architecture (Eqs. 4–7). In either case, the recurrent module can be any of *LSTM*, *dLSTM*, or *mdLSTM*. Weight tying [8, 17] is used to couple the weight matrices of the embedding and the output layers. This reduces the number of parameters and prevents the model from learning a one-to-one correspondence between the input and the output [13].

We run two different sets of experiments. First, we analyze the effect of the *dual* connection by comparing the performances of the two architectures (ERS vs Dual), using each of the recurrent modules, on both the PTB and the WT2 datasets. In this setting the hyperparameters are tuned for the ERS architecture, and then transferred to the *dual* case. Second, we search for the best hyperparameters for the *dual* architecture using the *mdLSTM* recurrence, and compare the perplexity score with current state-of-the-art values. All the experiments have been performed using the Keras library [4].

The networks are trained using the Nadam optimizer [5], a variation of Adam [9] where Nesterov momentum is applied. The number of training epochs is different for each experimental condition. On one hand, when the objective is to perform a pairwise comparison between *dual* and non-dual architectures, we train the models for 100 epochs. On the other hand, when the goal is to compare the *dual* network with state of the art approaches, we let the models run for 300 epochs. We use batch sizes of 32 and 128 for the PTB and the WT2 problems, respectively, and set the sequence length to 25 in all cases. The remaining hyperparameters are searched in the ranges described in Table 1.

Finally, all the models are run twice, both with and without dynamic evaluation [10]. Dynamic evaluation is a standard method commonly used to adapt the model parameters, learned during training, using also the validation data. This allows the networks to get adapted to the new evaluation conditions, which in general improves their performance. In order to keep the models as simple as possible, no additional modifications have been considered.

Table 1. List of all the hyperparameters and the search range associated with each of them. Those marked with an asterisk (*) refer to the *dual* architectures only.

Name	Description	Values
<i>Num epochs</i>	Number of training epochs	{100, 300}
<i>Learning rate</i>	Learning rate	[10^{-6} , 10^{-3}]
<i>Batch size</i>	Batch size	{32, 128}
<i>Seq len</i>	Sequence length	{10, 25, 50}
<i>Embedding units</i>	Size of the embedding layer	{400, 850}
<i>Recurrent units</i>	Size of the recurrent layers	{400, 850, 1150}
<i>LSTM layers</i>	Number of recurrent layers	{1, 2, 3}
<i>Dual units*</i>	Size of the <i>dual</i> layer	{400, 850}
<i>Embedding L2reg</i>	L2 regularization applied to the Embedding and output layers	{0, 10^{-6} , 10^{-5} }
<i>Rec. input L2reg</i>	L2 regularization applied to the input weights of the recurrent layer	{0, 10^{-6} , 10^{-5} }
<i>Rec. L2reg</i>	L2 regularization applied to the recurrent weights of the recurrent layer	{0, 10^{-6} , 10^{-5} }
<i>Activation L2reg</i>	L2 regularization applied to the recurrent layers output	{0, 10^{-6} , 10^{-5} }
<i>Dual L2reg*</i>	L2 regularization applied to <i>dual</i> layer	{0, 10^{-6} , 10^{-5} }
<i>Rec. input Dropout</i>	Dropout before the first recurrent layer	[0.0, 0.5]
<i>Rec. Dropout</i>	Dropout for the linear transformation of the recurrent state	[0.0, 0.5]
<i>Rec. internal Dropout</i>	Dropout between the recurrent layers	[0.0, 0.5]
<i>Rec. output Dropout</i>	Dropout after the last recurrent layer	[0.0, 0.5]
<i>Dual input Dropout*</i>	Dropout before the <i>dual</i> layer	[0.0, 0.5]
<i>Dual output Dropout*</i>	Dropout after the <i>dual</i> layer	[0.0, 0.5]
<i>Mogrifier deep</i>	Mogrifier rounds	{0, 2, 3, 4, 5, 6}
<i>Mogrifier L2reg</i>	L2 regularization applied to Mogrifier weights	{0, 10^{-6} , 10^{-5} }
<i>Mogrifier rank</i>	Weight factorization. $Q^i \in R^{m \times n} = Q_l^i Q_r^i$ with $Q_l^i \in R^{m \times k}$, $Q_r^i \in R^{k \times n}$	{0, 50, 100, 200}
<i>Mogrifier Dropout</i>	Dropout between the Mogrifier weights	[0.0, 0.2]
<i>Learning rate eval</i>	Learning rate when Dynamic evaluation	[10^{-6} , 10^{-3}]
<i>Seq len eval</i>	Sequence length when Dynamic evaluation	[5, 50]
<i>Clipnorm eval</i>	Gradients clipping to a maximum norm	[0.0, 1.0]

4 Results

We first show the results of the comparative analysis ERS vs Dual, then we focus on the search of the optimal hyperparameters for the *dual* architecture with the *mdLSTM* recurrence.

4.1 Dual vs Non-dual Architectures

Table 2 displays the validation and test perplexity scores obtained for each of the experimental configurations on the PTB and the WT2 problems, both with and without dynamic evaluation. To facilitate the comparison, each pair of rows contain the results for one of the recurrent modules (*LSTM*, *dLSTM* or *mdLSTM*) using the two architectures ERS and Dual, with the best values shown in bold. In each case, the hyperparameters are tuned for the standard ERS architecture and then used within the *dual* networks without any additional adaptation. The exceptions are hyperparameters, such as the *dual* dropout, which do not exist in the ERS configuration (those marked with an asterisk in Table 1). To give a measure of the model complexity, Table 2 contains also the approximate number of trainable parameters for each configuration.

Table 2. Validation and test word-level perplexity obtained for each of the experimental configurations on the PTB (top) and the WT2 (bottom) datasets.

Penn Treebank Dataset					
Model	No. PARAMS	No. Dyneval		Dyneval	
		Val.	Test	Val.	Test
<i>LSTM</i>	8.88 M	67.37	64.91	62.31	61.17
<i>Dual LSTM</i>	9.60 M	61.22	59.39	55.26	54.69
<i>dLSTM</i>	13.62 M	63.44	61.03	57.18	56.01
<i>Dual dLSTM</i>	13.94 M	60.99	59.56	56.11	54.87
<i>mdLSTM</i>	21.43 M	57.42	55.48	51.16	50.27
<i>Dual mdLSTM</i>	22.88 M	56.08	54.12	48.82	48.00
<i>mdLSTM+</i>	22.16 M	57.77	56.29	50.42	49.83
WikiText-2 Dataset					
Model	No. PARAMS	No. Dyneval		Dyneval	
		Val.	Test	Val.	Test
<i>LSTM</i>	20.23 M	92.84	88.28	74.98	69.42
<i>Dual LSTM</i>	20.95 M	85.88	82.48	61.94	57.61
<i>dLSTM</i>	29.60 M	78.65	75.60	63.26	59.42
<i>Dual dLSTM</i>	30.32 M	77.01	73.90	61.10	57.10
<i>mdLSTM</i>	37.51 M	72.05	69.06	57.42	53.93
<i>Dual mdLSTM</i>	38.95 M	71.78	70.83	53.48	50.71

As expected, dynamic evaluation improves the results regardless of the model or the dataset. The main observation, however, is that networks enhanced with the *dual* connection display lower perplexity scores for almost all the training conditions on both the PTB and the WT2 datasets. The advantage of the Dual vs the ERS architecture is larger for less complex models, and narrows as the model complexity increases. Nevertheless, even for networks with *mdLSTM* recurrence, the *dual* architectures outperform their non-dual counterparts in more than 2 perplexity points on the test set, when dynamic evaluation is used.

In order to test that this improvement is due to the *dual* connection and not to the presence of an extra processing layer, we performed an additional experiment with a *Dual mdLSTM* model, but removing the term $W^{de}e_t$ from Eq. 6. The results for the PTB dataset are shown in Table 2 as *mdLSTM+*. Note that, in spite of slightly improving the baseline, this enhanced mogrifier model is still well below the result obtained with the full *dual* architecture.

Finally, it is worth noting that all the results presented correspond to our own implementation of the models, and that in most cases we are not including some of the several training or validation adaptations frequently used in the literature (such as AWD or MoS, for example). This can explain the difference with respect to the results reported by [12] for the Mogrifier-LSTM model. We would expect a further improvement of the results if these additional mechanisms were implemented.

4.2 Dual Mogrifier Fine Tuning

The second part of the experiments consists of searching for the best hyperparameters in the configuration that provided the smallest perplexity in the previous setup, that is the *Dual mdLSTM* architecture. We carry out this experiment with the PTB problem. After an extensive search (see Table 1), the best performance is obtained with a model with 850 units in the embedding layer, 850 units in each of the mogrifier LSTM layers, and 850 units also in the *dual* layer. The input, recurrent, internal, and output dropout rates are all set to 0.5, the *dual* input and output dropout rates are set to 0.5 and 0.4, respectively, and the mogrifier dropout rate is set to 0.15. Both the embedding and the *dual* L2 regularization parameters are set to 10^{-5} . The mogrifier number of rounds is set to 4, and the rank to 100. All the remaining hyperparameters are set to 0.

After the training phase, we continue with a fine tuning of some additional hyperparameters, using the validation data. First, we look for the best sequence length in the range [5, 70], and then we fine-tune the softmax temperature in the range [0.9, 1.3]. When using dynamic evaluation, we also look for the best gradient clipping value (in the range [0.0, 1.0]) and, following [12], we repeat the whole procedure with the β_1 parameter of the Nadam optimizer set to 0, which resembles the RMSProp optimizer without momentum. The results are shown in Table 3, together with the top perplexity scores reported in the literature for the same problem.

The state-of-the-art is dominated by several variations of the AWD-LSTM network [13], the most common being the inclusion of a Mixture of Softmaxes

Table 3. Best validation and test word-level perplexity scores reported in the literature for the Penn Treebank dataset, with and without dynamic evaluation. Missing values in the last two columns correspond to works where the dynamic evaluation approach was not considered. The last row in the table displays the results obtained with our *Dual mdLSTM* network.

Model		No Dyneval		Dyneval	
		Val.	Test	Val.	Test
<i>AWD-LSTM</i> [13]	24 M	60.00	57.30	–	–
<i>AWD-LSTM-DOC</i> [19]	23 M	54.12	52.38	–	–
<i>AWD-LSTM</i> [10]	24 M	59.80	57.70	51.60	51.10
<i>mdLSTM - ours</i>	22 M	57.42	55.48	51.16	50.27
<i>AWD-LSTM +PDR</i> [1]	24 M	57.90	55.60	50.10	49.30
<i>AWD-LSTM +MoS</i> [21]	22 M	56.54	54.44	48.33	47.69
<i>AWD-LSTM +MoS +PDR</i> [1]	22 M	56.20	53.80	48.00	47.30
<i>AWD-LSTM-DOC x5</i> [19]	185 M	48.63	47.17	–	–
<i>AWD-LSTM +MoS +FRAGE</i> [6]	24 M	55.52	53.51	47.38	46.54
<i>AWD-LSTM +MoS +Adv</i> [20]	22 M	54.98	52.87	47.15	46.52
<i>AWD-LSTM +MoS +Adv +PS</i> [20]	22 M	54.10	52.20	46.63	46.01
<i>Mogrifier-LSTM</i> [12]	24 M	51.40	50.10	44.90	44.80
<i>Dual mdLSTM - ours</i>	23 M	52.87	51.19	45.13	44.61

(MoS) [21]. Other add-ons include Direct Output Connection (DOC) [19], which is a generalization of MoS, Frequency Agnostic word Embedding (FRAGE) [6], Past Decode Regularization (PDR) [1], or Partial Shuffling (PS) with Adversarial Training (Adv) [20]. The *mogrifier-LSTM* described in Sect. 2.2 combines many of these ideas with a mutual gating between the input and the hidden state vectors to obtain the best results reported in the literature for the PTB problem, among those obtained by networks that do not use additional data during the training phase. Compared with all these models, our current approach leads the ranking with a perplexity score of 44.61, even though most of the aforementioned adjustments have not been considered.

Finally, it is important to mention that the last two rows in the table do not correspond to comparable models. While the penultimate row shows the results reported by Melis et al. [12] with their *Mogrifier-LSTM* model, the last table row contains the results of our *Dual mdLSTM* model, which uses the *mogrifier* transformation but lacks many of the additional characteristics of the Melis et al. model. Hence, regarding the improvement associated with fine-tuning the *Dual mdLSTM* model, the fair comparison would be with the results shown in Table 2, which have been also included in Table 3 for the sake of clarity. In this case the dual model outperforms its non-dual equivalent in more than 5 perplexity points (50.27 versus 44.61) on the test set.

5 Discussion

In this work, we have presented a new network design for the Language Modeling task based on the *dual* network proposed by [16]. This network adds a direct connection between the input and the output, skipping the recurrent module, and can be adapted to any of the traditional Embedding-Recurrent-Softmax (ERS) models, opening the way to new approaches for this task. We have based our work on the Penn Treebank [15] and the WikiText-2 [14] datasets, comparing the ERS approach and its *dual* alternative. Regardless of the configuration, the *dual* version performs always better, even though it faces a slight disadvantage, since most of the hyperparameters are tuned using the ERS model. We can expect a much better performance if the complete set of hyperparameters is properly tuned for the *dual* network.

This is in fact the case for the second experiment, where a *Dual mdLSTM*, which includes a simplified version of the mogrifier LSTM [12] within a *dual* architecture, is fine tuned for the Penn Treebank dataset. After a thorough search of the hyperparameter space, we have found a network configuration that establishes a new state-of-the-art score for this problem. Interestingly, this new record has been obtained in spite of leaving aside many of the standard features used in most state-of-the-art approaches, such as *AWD* [13] or *MoS* [21]. The incorporation of these features into the *dual* architecture can be expected to further increase the model performance.

The *dual* architecture was firstly proposed as an alternative that reduces the computational load on the recurrent layer, letting it concentrate on modeling the temporal dependencies only. From a more abstract point of view, it has been argued that the dual architecture can be understood as a sort of Mealy machine, where the output explicitly depends on both the hidden state and the input [16]. Our results show that this explicit dependence on the input can indeed lead to better performance on language modeling tasks. This emphasizes the importance of the current input in RNN models.

Finally, although the new approach has not been tested with large-scale language corpora, we expect that our results scale well to larger datasets. Work in progress contemplates this extension. The *dual* architecture also needs further research concerning the deepness of the specific variations of Language Modeling and other families of problems not necessarily related to Natural Language Processing. This work opens a new line of research to be considered when processing any sequence or time series. The utility of this approach in more general problems will be addressed as future work.

Acknowledgments. This work has been partially funded by Spanish project PID2020-114867RB-I00, (MCIN/AEI and ERDF- “A way of making Europe”).

References

1. Brahma, S.: Improved language modeling by decoding the past. In: Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, pp. 1468–1476. Association for Computational Linguistics, Florence, Italy (Jul 2019)

2. Brown, T., et al.: Language models are few-shot learners. In: Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M.F., Lin, H. (eds.) *Advances in Neural Information Processing Systems*, vol. 33, pp. 1877–1901. Curran Associates, Inc. (2020)
3. Cho, K., et al.: Learning phrase representations using RNN encoder-decoder for statistical machine translation. In: Moschitti, A., Pang, B., Daelemans, W. (eds.) *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, 25–29 October 2014, Doha, Qatar, A meeting of SIG-DAT, a Special Interest Group of the ACL*, pp. 1724–1734. ACL (2014)
4. Chollet, F., et al.: Keras (2015). <https://keras.io>
5. Dozat, T.: Incorporating nesterov momentum into adam. In: *Proceedings of 4th International Conference on Learning Representations, Workshop Track* (2016)
6. Gong, C., He, D., Tan, X., Qin, T., Wang, L., Liu, T.Y.: Frage: Frequency-agnostic word representation. In: Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., Garnett, R. (eds.) *Advances in Neural Information Processing Systems*, vol. 31. Curran Associates, Inc. (2018)
7. Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural Comput.* **9**(8), 1735–1780 (1997)
8. Inan, H., Khosravi, K., Socher, R.: Tying word vectors and word classifiers: A loss framework for language modeling. In: *Proceedings of the 5th International Conference on Learning Representations* (2017)
9. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: Bengio, Y., LeCun, Y. (eds.) *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, 7–9 May 2015, Conference Track Proceedings* (2015)
10. Krause, B., Kahembwe, E., Murray, I., Renals, S.: Dynamic evaluation of neural sequence models. In: Dy, J., Krause, A. (eds.) *Proceedings of the 35th International Conference on Machine Learning. Proceedings of Machine Learning Research*, 10–15 Jul, vol. 80, pp. 2766–2775. PMLR (2018)
11. Marcus, M.P., Marcinkiewicz, M.A., Santorini, B.: Building a large annotated corpus of english: The penn treebank. *Comput. Linguist.* **19**(2), 313–330 (1993)
12. Melis, G., Kočický, T., Blunsom, P.: Mogrifier LSTM. In: *International Conference on Learning Representations* (2020)
13. Merity, S., Keskar, N.S., Socher, R.: Regularizing and optimizing LSTM language models. In: *International Conference on Learning Representations* (2018)
14. Merity, S., Xiong, C., Bradbury, J., Socher, R.: Pointer sentinel mixture models. In: *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, 24–26 April 2017, Conference Track Proceedings. OpenReview.net* (2017)
15. Mikolov, T., Karafiát, M., Burget, L., Cernocký, J., Khudanpur, S.: Recurrent neural network based language model. In: Kobayashi, T., Hirose, K., Nakamura, S. (eds.) *INTERSPEECH*, pp. 1045–1048. ISCA (2010)
16. Oliva, C., Lago-Fernández, L.F.: Separation of Memory and Processing in Dual Recurrent Neural Networks. In: Farkaš, I., Masulli, P., Otte, S., Wermter, S. (eds.) *ICANN 2021. LNCS*, vol. 12894, pp. 360–371. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-86380-7_29
17. Press, O., 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. Association for Computational Linguistics, Valencia, Spain (Apr 2017)
18. Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I.: Language models are unsupervised multitask learners (2019)

19. Takase, S., Suzuki, J., 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. Association for Computational Linguistics, Brussels, Belgium (Oct-Nov 2018)
20. Wang, D., Gong, C., Liu, Q.: Improving neural language modeling via adversarial training. In: Chaudhuri, K., Salakhutdinov, R. (eds.) Proceedings of the 36th International Conference on Machine Learning. Proceedings of Machine Learning Research, 09–15 Jun, vol. 97, pp. 6555–6565. PMLR (2019)
21. Yang, Z., Dai, Z., Salakhutdinov, R., Cohen, W.W.: Breaking the softmax bottleneck: A high-rank RNN language model. In: International Conference on Learning Representations (2018)