

---

# Bayesian Sparsification of Gated Recurrent Neural Networks

---

Ekaterina Lobacheva<sup>1\*</sup>, Nadezhda Chirkova<sup>1\*</sup>, Dmitry Vetrov<sup>1,2</sup>

<sup>1</sup>Samsung-HSE Laboratory, National Research University Higher School of Economics

<sup>2</sup>Samsung AI Center Moscow, Russia  
{elobacheva, nchirkova, dvetrov}@hse.ru

## Abstract

Bayesian methods have been successfully applied to sparsify weights of neural networks and to remove structure units from the networks, e. g. neurons. We apply and further develop this approach for gated recurrent architectures. Specifically, in addition to sparsification of individual weights and neurons, we propose to sparsify preactivations of gates and information flow in LSTM. It makes some gates and information flow components constant, speeds up forward pass and improves compression. Moreover, the resulting structure of gate sparsity is interpretable and depends on the task.

## 1 Introduction

Recurrent neural networks (RNNs) yield high-quality results in many applications [1, 4, 18, 21] but often overfit due to overparametrization. In many practical problems, RNNs can be compressed orders of times with only slight quality drop or even with quality improvement [2, 15, 20]. Methods for RNN compression can be divided into three groups: based on matrix factorization [6, 19], quantization [7] or sparsification [2, 15, 20].

We focus on RNNs sparsification. Two main groups of approaches for sparsification are pruning and Bayesian sparsification. In pruning [15, 20], weights with absolute values less than a predefined threshold are set to zero. Such methods imply a lot of hyperparameters (thresholds, pruning schedule etc). Bayesian sparsification techniques [14, 16, 8, 9, 2] treat weights of an RNN as random variables and approximate posterior distribution over them given sparsity-inducing prior distribution. After training weights with low signal-to-noise ratio are set to zero. This allows eliminating the majority of weights from the model without time-consuming hyperparameters tuning. Also, Bayesian sparsification techniques can be easily extended to permanently set to zero intermediate variables in the network's computational graph [16, 8] (e.g. neurons in fully-connected networks or filters in convolutional networks). It is achieved by multiplying such a variable on a learnable weight, finding posterior over it and setting the weight to zero if the corresponding signal-to-noise ratio is small.

In this work, we investigate the last mentioned property for gated architectures, particularly for LSTM. Following [2, 14], we sparsify individual weights of the RNN. Following [8], we eliminate neurons from the RNN by introducing multiplicative variables on activations of neurons. Our main contribution is the introduction of multiplicative variables on preactivations of the gates and information flow in LSTM. This leads to several positive effects. Firstly, when some component of preactivations is permanently set to zero, the corresponding gate becomes constant. It simplifies LSTM structure and speeds up computations. Secondly, we obtain a three-level hierarchy of sparsification: sparsification of individual weights helps to sparsify gates and information flow (make their components constant), and sparsification of gates and information flow helps to sparsify neurons (remove them from the model). As a result, the overall compression of the model is higher.

---

\*Equal contribution

## 2 Preliminaries

Consider a dataset of  $N$  objects  $(x_i, y_i)$  and a model  $p(y|x, W)$  parametrized by a neural network with weights  $W$ . In [14], the authors propose a Bayesian technique called Sparse variational dropout (SparseVD) for neural networks sparsification. This model comprises log-uniform prior over weights:  $p(|w_{ij}|) \propto \frac{1}{|w_{ij}|}$  and fully factorized normal approximate posterior:  $q(w_{ij}) = \mathcal{N}(w_{ij}|m_{ij}, \sigma_{ij}^2)$ . To find parameters of the approximate posterior distribution, evidence lower bound (ELBO) is optimized:

$$\sum_{i=1}^N \mathbb{E}_{q(W)} \log p(y^i|x^i, W) - KL(q(W)||p(W)) \rightarrow \max_{m, \sigma} \quad (1)$$

Because of the log-uniform prior, for the majority of weights signal-to-noise ratio  $m_{ij}^2/\sigma_{ij}^2 \rightarrow 0$  and these weights do not affect network’s output. In [2] SparseVD is adapted to RNNs.

In [8] the authors propose to multiply activations of neurons on group variables  $z$  and to learn and sparsify group variables along with  $W$ . They put standard normal prior on  $W$  and log-uniform prior on  $z$ . The first prior moves mean values of  $W$  to 0, and it helps to set to zero  $z$  and to remove neurons from the model. This model is equivalent to multiplying rows of weight matrices on group variables.

## 3 Proposed method

To sparsify individual weights, we apply SparseVD [14] to all weights of the RNN, taking into account recurrent specifics underlined in [2]. To compress layers and remove neurons, we follow [8] and introduce group variables for the neurons of all layers (excluding output predictions), and specifically,  $z^x$  and  $z^h$  for input and hidden neurons of LSTM.

The key component of our model is introducing groups variables  $z^i, z^f, z^g, z^o$  on preactivations of gates and information flow. The resulting LSTM layer looks as follows:

$$f = \sigma \left( (W_f^h (h_{t-1} \odot z^h) + W_f^x (x_t \odot z^x)) \odot z^f + b_f \right) \quad \{ \text{same for } i, o \text{ and } g \} \quad (2)$$

$$c_t = f \odot c_{t-1} + i \odot g \quad h_t = o \odot \tanh(c_t) \quad (3)$$

Described model is equivalent to multiplying rows and columns of weight matrices on group variables:

$$w_{f,ij}^h = \hat{w}_{f,ij}^h \cdot z_i^h \cdot z_j^f \quad \{ \text{same for } i, o \text{ and } g \}$$

We learn group variables  $z$  in the same way as weights  $W$ : approximate posterior with fully factorized normal distribution given fully factorized log-uniform prior distribution<sup>2</sup>. To find approximate posterior distribution, we maximize ELBO (1). After learning, we set all weights and group variables with signal-to-noise ratio less than 0.05 to 0.

If some component of  $z^i, z^f, z^o, z^g$  is set to 0, the corresponding gate or information flow component becomes constant (equal to activation function of bias). It means that we don’t need to compute this component, and the forward pass through LSTM is accelerated.

**Related work.** In [20] the authors propose a pruning-based method that removes neurons from LSTM and argue that independent removing of  $i, f, g, o$  components may lead to invalid LSTM units. In our model, we do not remove these components but make them constant, gaining compression and acceleration with correct LSTM structure.

## 4 Experiments

We perform experiments with LSTM architecture on two types of problems: text classification (datasets IMDb [10] and AGNews [22]) and language modeling (dataset PTB [11], character and word level tasks). For text classification, we use networks with an embedding layer, one recurrent layer and an output dense layer at the last timestep. For language modeling, we use networks with one

<sup>2</sup>Our experiments show that log-uniform prior on individual weights gives better results than standard normal prior used in [8].

Task	Method	Quality	Compression	Neurons $x - h$	Gates
IMDb Accuracy %	Original	84.1	1x	300 – 128	512
	SparseVD W	<b>84.47</b>	1135x	7 – 9	22
	SparseVD W+N	83.98	17874x	1 – 5	12
	SparseVD W+G+N	83.98	<b>19747x</b>	<b>1 – 4</b>	<b>6</b>
AGNews Accuracy %	Original	<b>90.6</b>	1x	300 – 512	2048
	SparseVD W	89.01	350x	193 – 65	72
	SparseVD W+N	88.55	645x	43 – 17	62
	SparseVD W+G+N	88.41	<b>647x</b>	<b>43 – 14</b>	<b>39</b>
Char PTB Bits-per-char Valid-Test	Original	1.498 – 1.456	1x	50 – 1000	4000
	SparseVD W	1.472 – 1.429	7.9x	50 – 431	1718
	SparseVD W+N	1.478 – 1.430	<b>10.0x</b>	<b>50 – 390</b>	<b>1560</b>
	SparseVD W+G+N	<b>1.467 – 1.425</b>	9.8x	50 – 404	1563
Word PTB Perplexity Valid-Test	Original	135.6 – 129.5	1x	10000 – 256	1024
	SparseVD W	<b>115.0 – 109.0</b>	22.0x	9985 – 153	281
	SparseVD W+N	116.2 – 111.0	23.4x	9993 – 134	335
	SparseVD W+G+N	122.2 – 116.5	<b>24.4x</b>	<b>9973 – 114</b>	<b>220</b>

Table 1: Quantitative results. Compression is equal to  $|w|/|w \neq 0|$ . In last two columns number of remaining neurons and non-constant gates in the recurrent layer are reported.




Figure 1: Structure of gate sparsity. Non-constant gates are dark blue, constant ones are light blue. Numbers correspond to activation values of constant gates. For language modeling only 15 randomly chosen active neurons are presented.

recurrent layer and an output dense layer. More details about network architectures and experiment setup are presented in Appendix A.

We compare four models in terms of quality and sparsity: baseline model without any regularization, standard SparseVD model for weights sparsification only (W), SparseVD model with group variables for neurons sparsification (W+N) and SparseVD model with group variables for gates and neurons sparsification (W+G+N). In all SparseVD models, we sparsify weights matrices of all layers. Since in text classification tasks usually only a small number of input words are important, we use additional multiplicative weights to sparsify the input vocabulary in case of group sparsification (W+N, W+G+N) following [2]. On the contrary, in language modeling tasks all input characters or words are usually important, therefore we do not use  $z^x$  for this task. Additional sparsification of input neurons in this case noticeably damage models quality and sparsity level of hidden neurons. To measure the sparsity level of our models we calculate the compression rate of individual weights as follows:  $|w|/|w \neq 0|$ . To compute the number of remaining neurons or non-constant gates we use corresponding rows/columns of  $W$  and corresponding weights  $z$  if applicable.

Quantitative results are shown in Table 1. Multiplicative variables for neurons boost group sparsity level without a significant quality drop. Additional variables for gates and information flow not only make some gates constant but also increase group sparsity level even further. Moreover, for a lot of constant gates bias values tend to be very large or small making corresponding gates either always open or close.

Proposed gate sparsification technique also reveals an interesting work-flow structure of LSTM networks for different tasks. Figure 1 shows typical examples of gates of remaining hidden neurons. For language modeling tasks output gates are very important because models need both store all the information about the input in the memory and output only the current prediction at each timestep. On the contrary, for text classification tasks models need to output the answer only once at the end of the sequence, hence they do not really use output gates. Also, the character level language modeling task is more challenging than the word level one: the model uses the whole gating mechanism to solve it. We think this is the main reason why gate sparsification does not help here.

## Acknowledgments

The study has been supported by Russian Science Foundation (grant 17-71-20072) and Samsung Research, Samsung Electronics.

## References

- [1] Chan, W., Jaitly, N., Le, Q. V., and Vinyals, O. (2016). Listen, attend and spell: A neural network for large vocabulary conversational speech recognition. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*.
- [2] Chirkova, N., Lobacheva, E., and Vetrov, D. (2018). Bayesian compression for natural language processing. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- [3] Glorot, X. and Bengio, Y. (2010). Understanding the difficulty of training deep feedforward neural networks. In Teh, Y. W. and Titterton, M., editors, *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, volume 9 of *Proceedings of Machine Learning Research*, pages 249–256, Chia Laguna Resort, Sardinia, Italy. Proceedings of Machine Learning Research.
- [4] Ha, D., Dai, A., and Le, Q. V. (2017). Hypernetworks. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- [5] Kingma, D. P. and Ba, J. (2015). Adam: A method for stochastic optimization. In *Proceedings of the 3rd International Conference for Learning Representations (ICLR)*.
- [6] Le, Q. V., Jaitly, N., and Hinton, G. E. (2015). A simple way to initialize recurrent networks of rectified linear units. *Computing Research Repository*, arXiv:1504.00941.
- [7] Liu, X., Cao, D., and Yu, K. (2018). Binarized lstm language model. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2113–2121. Association for Computational Linguistics.
- [8] Louizos, C., Ullrich, K., and Welling, M. (2017). Bayesian compression for deep learning. In *Advances in Neural Information Processing Systems 30*, pages 3288–3298.
- [9] Louizos, C., Welling, M., and Kingma, D. P. (2018). Learning sparse neural networks through l0 regularization. In *International Conference on Learning Representations*.
- [10] Maas, A. L., Daly, R. E., Pham, P. T., Huang, D., Ng, A. Y., and Potts, C. (2011). Learning word vectors for sentiment analysis. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies - Volume 1, HLT '11*, pages 142–150, Stroudsburg, PA, USA. Association for Computational Linguistics.
- [11] Marcus, M. P., Marcinkiewicz, M. A., and Santorini, B. (1993). Building a large annotated corpus of english: The penn treebank. *Comput. Linguist.*, 19(2):313–330.
- [12] Mikolov, T., Kombrink, S., Burget, L., Cernocky, J., and Khudanpur, S. (2011). Extensions of recurrent neural network language model. In *2011 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5528–5531.
- [13] Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., and Dean, J. (2013). Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems 26*, pages 3111–3119.
- [14] Molchanov, D., Ashukha, A., and Vetrov, D. (2017). Variational dropout sparsifies deep neural networks. In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017*.
- [15] Narang, S., Diamos, G. F., Sengupta, S., and Elsen, E. (2017). Exploring sparsity in recurrent neural networks. In *Proceedings of the International Conference for Learning Representations (ICLR)*.

- [16] Neklyudov, K., Molchanov, D., Ashukha, A., and Vetrov, D. P. (2017). Structured bayesian pruning via log-normal multiplicative noise. In *Advances in Neural Information Processing Systems 30*, pages 6778–6787.
- [17] Pennington, J., Socher, R., and Manning, C. D. (2014). Glove: Global vectors for word representation. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, volume 14, pages 1532–1543.
- [18] Ren, M., Kiros, R., and Zemel, R. S. (2015). Exploring models and data for image question answering. In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 2953–2961.
- [19] Tjandra, A., Sakti, S., and Nakamura, S. (2017). Compressing recurrent neural network with tensor train. *Computing Research Repository*, arXiv:1705.08052.
- [20] Wen, W., He, Y., Rajbhandari, S., Zhang, M., Wang, W., Liu, F., Hu, B., Chen, Y., and Li, H. (2018). Learning intrinsic sparse structures within long short-term memory. In *International Conference on Learning Representations*.
- [21] Wu, Y., Schuster, M., Chen, Z., Le, Q. V., Norouzi, M., Macherey, W., Krikun, M., Cao, Y., Gao, Q., Macherey, K., Klingner, J., Shah, A., Johnson, M., Liu, X., Kaiser, L., Gouws, S., Kato, Y., Kudo, T., Kazawa, H., Stevens, K., Kurian, G., Patil, N., Wang, W., Young, C., Smith, J., Riesa, J., Rudnick, A., Vinyals, O., Corrado, G., Hughes, M., and Dean, J. (2016). Google’s neural machine translation system: Bridging the gap between human and machine translation. *Computing Research Repository*, arXiv:1609.08144.
- [22] Zhang, X., Zhao, J., and LeCun, Y. (2015). Character-level convolutional networks for text classification. In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems (NIPS)*.

## A Experimental setup

**Datasets.** To evaluate our approach on text classification task we use two standard datasets: IMDb dataset [10] for binary classification and AGNews dataset [22] for four-class classification. We set aside 15% and 5% of training data for validation purposes respectively. For both datasets, we use the vocabulary of 20,000 most frequent words. To evaluate our approach on language modeling task we use the Penn Treebank corpus [11] with the train/valid/test partition of Mikolov et al. mikolov11. The dataset has a vocabulary of 50 characters or 10,000 words.

**Architectures for text classification.** We use networks with one embedding layer of 300 units, one LSTM layer of 128 / 512 hidden units for IMDb / AGNews, and finally, a fully connected layer applied to the last output of the LSTM. Embedding layer is initialized with word2vec [13] / GloVe [17] and SparseVD models are trained for 800 / 150 epochs on IMDb / AGNews. Hidden-to-hidden weight matrices  $W^h$  are initialized orthogonally and all other matrices are initialized uniformly using the method by Glorot and Bengio pmlr-v9-glorot10a. We train our networks using Adam [5] with batches of size 128 and a learning rate of 0.0005. Baseline networks overfit for all our tasks, therefore, we present results for them with early stopping.

**Architectures for language modeling.** To solve character / word-level tasks we use networks with one LSTM layer of 1000 / 256 hidden units and fully-connected layer with softmax activation to predict next character or word. We train SparseVD models for 250 / 150 epochs on character-level / word-level tasks. All weight matrices of the networks are initialized orthogonally and all biases are initialized with zeros. Initial values of hidden and cell elements are not trainable and equal to zero. For the character-level task, we train our networks on non-overlapping sequences of 100 characters in mini-batches of 64 using a learning rate of 0.002 and clip gradients with threshold 1. For the word-level task, networks are unrolled for 35 steps. We use the final hidden states of the current mini-batch as the initial hidden state of the subsequent mini-batch (successive mini batches sequentially traverse the training set). The size of each mini-batch is 32. We train models using a learning rate of 0.002 and clip gradients with threshold 10. Baseline networks overfit for all our tasks, therefore, we present results for them with early stopping.

**Sparsification.** For all weights that we sparsify, we initialize  $\log \sigma$  with -3. We eliminate weights with signal-to-noise ratio less than  $\tau = 0.05$ .