

# Parallelizable Stack Long Short-Term Memory

Shuoyang Ding Philipp Koehn

Center for Language and Speech Processing

Johns Hopkins University

{dings, phi}@jhu.edu

## Abstract

Stack Long Short-Term Memory (StackLSTM) is useful for various applications such as parsing and string-to-tree neural machine translation, but it is also known to be notoriously difficult to parallelize for GPU training due to the fact that the computations are dependent on discrete operations. In this paper, we tackle this problem by utilizing state access patterns of StackLSTM to homogenize computations with regard to different discrete operations. Our parsing experiments show that the method scales up almost linearly with increasing batch size, and our parallelized PyTorch implementation trains significantly faster compared to the Dynet C++ implementation.

## 1 Introduction

Tree-structured representation of language has been successfully applied to various applications including dependency parsing (Dyer et al., 2015), sentiment analysis (Socher et al., 2011) and neural machine translation (Eriguchi et al., 2017). However, most of the neural network architectures used to build tree-structured representations are not able to exploit full parallelism of GPUs by minibatched training, as the computation that happens for each instance is conditioned on the input/output structures, and hence cannot be naïvely grouped together as a batch. This lack of parallelism is one of the major hurdles that prevent these representations from wider adoption practically (e.g., neural machine translation), as many natural language processing tasks currently require the ability to scale up to very large training corpora in order to reach state-of-the-art performance.

We seek to advance the state-of-the-art of this problem by proposing a parallelization scheme for one such network architecture, the Stack Long Short-Term Memory (StackLSTM) proposed in Dyer et al. (2015). This architecture has been

successfully applied to dependency parsing (Dyer et al., 2015, 2016; Ballesteros et al., 2017) and syntax-aware neural machine translation (Eriguchi et al., 2017) in the previous research literature, but none of these research results were produced with minibatched training. We show that our parallelization scheme is feasible in practice by showing that it scales up near-linearly with increasing batch size, while reproducing a set of results reported in (Ballesteros et al., 2017).

## 2 StackLSTM

StackLSTM (Dyer et al., 2015) is an LSTM architecture (Hochreiter and Schmidhuber, 1997) augmented with a stack  $\mathcal{H}$  that stores some of the hidden states built in the past. Unlike traditional LSTMs that always build state  $h_t$  from  $h_{t-1}$ , the states of StackLSTM are built from the head of the state stack  $\mathcal{H}$ , maintained by a stack top pointer  $p(\mathcal{H})$ . At each time step, StackLSTM takes a real-valued input vector together with an additional discrete operation on the stack, which determines what computation needs to be conducted and how the stack top pointer should be updated. Throughout this section, we index the input vector (e.g. word embeddings)  $x_t$  using the time step  $t$  it is fed into the network, and hidden states in the stack  $h_j$  using their position  $j$  in the stack  $\mathcal{H}$ ,  $j$  being defined as the 0-base index starting from the stack bottom.

The set of input discrete actions typically contains at least `PUSH` and `POP` operations. When these operations are taken as input, the corresponding computations on the StackLSTM are listed below:<sup>1</sup>

- `PUSH`: read previous hidden state  $h_{p(\mathcal{H})}$ , perform LSTM forward computation with  $x_t$  and

<sup>1</sup>To simplify the presentation, we omitted the updates on cell states, because in practice the operations performed on cell states and hidden states are the same.

Transition Systems	Transition Op	Stack Op	Buffer Op	Composition Op
Arc-Standard	Shift	push	pop	none
	Left-Arc	pop, pop, push	hold	$S1 \leftarrow g(S0, S1)$
	Right-Arc	pop	hold	$S1 \leftarrow g(S1, S0)$
Arc-Eager	Shift	push	pop	none
	Reduce	pop	hold	none
	Left-Arc	pop	hold	$B0 \leftarrow g(B0, S0)$
	Right-Arc	push	pop	$B0 \leftarrow g(S0, B0)$
Arc-Hybrid	Shift	push	pop	none
	Left-Arc	pop	hold	$B0 \leftarrow g(B0, S0)$
	Right-Arc	pop	hold	$S1 \leftarrow g(S1, S0)$

Table 1: Correspondence between transition operations and stack/buffer operations for StackLSTM, where  $g$  denotes the composition function as proposed by (Dyer et al., 2015).  $S0$  and  $B0$  refers to the token-level representation corresponding to the top element of the stack and buffer, while  $S1$  and  $B1$  refers to those that are second to the top. We use a different notation here to avoid confusion with the states in StackLSTM, which represent non-local information beyond token-level.

$h_{p(\mathcal{H})}$ , write new hidden state to  $h_{p(\mathcal{H})+1}$ , update stack top pointer with  $p(\mathcal{H}) \leftarrow p(\mathcal{H}) + 1$ .

- **POP**: update stack top pointer with  $p(\mathcal{H}) \leftarrow p(\mathcal{H}) - 1$ .

Reflecting on the aforementioned discussion on parallelism, one should notice that StackLSTM falls into the category of neural network architectures that is difficult to perform minibatched training. This is caused by the fact that the computation performed by StackLSTM at each time step is dependent on the discrete input actions. The following section proposes a solution to this problem.

### 3 Parallelizable StackLSTM

Continuing the formulation in the previous section, we will start by discussing our proposed solution under the case where the set of discrete actions contains only **PUSH** and **POP** operations; we then move on to discussion of the applicability of our proposed solution to the transition systems that are used for building representations for dependency trees.

The first modification we perform to the **PUSH** and **POP** operations above is to unify the pointer update of these operations as  $p(\mathcal{H}) \leftarrow p(\mathcal{H}) + op$ , where  $op$  is the input discrete operation that could either take the value  $+1$  or  $-1$  for **PUSH** and **POP** operation. After this modification, we came to the following observations:

**Observation 1** *The computation performed for POP operation is a subset of Push operation.*

Now, what remains to homogenize **PUSH** and **POP** operations is conducting the extra computations needed for **PUSH** operation when **POP** is fed in as well, while guaranteeing the correctness of the resulting hidden state both in the current time step and in the future. The next observation points out a way for this guarantee:

**Observation 2** *In a StackLSTM, given the current stack top pointer position  $p(\mathcal{H})$ , any hidden state  $h_i$  where  $i > p(\mathcal{H})$  will not be read until it is overwritten by a Push operation.*

What follows from this observation is the guarantee that we can always safely overwrite hidden states  $h_i$  that are indexed higher than the current stack top pointer, because it is known that any read operation on these states will happen after another overwrite. This allows us to *do the extra computation anyway* when **POP** operation is fed, because the extra computation, especially updating  $h_{p(\mathcal{H})+1}$ , will not harm the validity of the hidden states at any time step.

Algorithm 1 gives the final forward computation for the Parallelizable StackLSTM. Note that this algorithm does not contain any if-statements that depends on stack operations and hence is homogeneous when grouped into batches that are consisted of multiple operations trajectories.

In transition systems (Nivre, 2008; Kuhlmann et al., 2011) used in real tasks (e.g., transition-based parsing) as shown in Table 1, it should be noted that more than **push** and **pop** operations are needed for the StackLSTM. Fortunately, for Arc-Eager and

---

**Algorithm 1:** Forward Computation for Parallelizable StackLSTM

---

**Input:** input vector  $x_t$   
discrete stack operation  $op$   
**Output:** current top hidden state  $h_{p(\mathcal{H})}$   
 $h_{prev} \leftarrow h_{p(\mathcal{H})};$   
 $h \leftarrow \text{LSTM}(x_t, h_{prev});$   
 $h_{p(\mathcal{H})+1} \leftarrow h;$   
 $p(\mathcal{H}) \leftarrow p(\mathcal{H}) + op;$   
**return**  $h_{p(\mathcal{H})};$

---

Arc-Hybrid transition systems, we can simply add a `hold` operation, which is denoted by value 0 for the discrete operation input. For that reason, we will focus on parallelization of these two transition systems for this paper. It should be noted that both observations discussed above are still valid after adding the `hold` operation.

## 4 Experiments

### 4.1 Setup

We implemented<sup>2</sup> the architecture described above in PyTorch (Paszke et al., 2017). We implemented the batched stack as a float tensor wrapped in a non-leaf variable, thus enabling in-place operations on that variable. At each time step, the batched stack is queried/updated with a batch of stack head positions represented by an integer vector, an operation made possible by `gather` operation and advanced indexing. Due to this implementation choice, the stack size has to be determined at initialization time and cannot be dynamically grown. Nonetheless, a fixed stack size of 150 works for all the experiments we conducted.

We use the dependency parsing task to evaluate the correctness and the scalability of our method. For comparison with previous work, we follow the architecture introduced in Dyer et al. (2015); Ballesteros et al. (2017) and chose the Arc-Hybrid transition system for comparison with previous work. We follow the data setup in Chen and Manning (2014); Dyer et al. (2015); Ballesteros et al. (2017) and use Stanford Dependency Treebank (de Marneffe et al., 2006) for dependency parsing, and we extract the Arc-Hybrid static oracle using the code associated with Qi and Manning (2017). The part-of-speech (POS) tags are generated with Stanford POS-tagger (Toutanova et al., 2003) with

<sup>2</sup><https://github.com/shuoyangd/hoolock>




Figure 1: Training speed at different batch size. Note that the  $x$ -axis is in log-scale in order to show all the data points properly.

a test set accuracy of 97.47%. We use exactly the same pre-trained English word embedding as Dyer et al. (2015).

We use Adam (Kingma and Ba, 2014) as the optimization algorithm. Following Goyal et al. (2017), we apply linear warmup to the learning rate with an initial value of  $\tau = 5 \times 10^{-4}$  and total epoch number of 5. The target learning rate is set by  $\tau$  multiplied by batch size, but capped at 0.02 because we find Adam to be unstable beyond that learning rate. After warmup, we reduce the learning rate by half every time there is no improvement for loss value on the development set (`ReduceLRonPlateau`). We clip all the gradient norms to 5.0 and apply a  $L_2$ -regularization with weight  $1 \times 10^{-6}$ .

We started with the hyper-parameter choices in Dyer et al. (2015) but made some modifications based on the performance on the development set: we use hidden dimension 200 for all the LSTM units, 200 for the parser state representation before the final softmax layer, and embedding dimension 48 for the action embedding.

We use Tesla K80 for all the experiments, in order to compare with Neubig et al. (2017b); Dyer et al. (2015). We also use the same hyper-parameter setting as Dyer et al. (2015) for speed comparison experiments. All the speeds are measured by running through one training epoch and averaging.

### 4.2 Results

Figure 1 shows the training speed at different batch sizes up to 256.<sup>3</sup> The speed-up of our model is

<sup>3</sup>At batch size of 512, the longest sentence in the training data cannot be fit onto the GPU.

$b$	dev		test	
	UAS	LAS	UAS	LAS
1*	92.50*	89.79*	92.10*	89.61*
8	92.93	90.42	92.54	90.11
16	92.62	90.19	92.53	90.13
32	92.43	89.89	92.31	89.94
64	92.53	90.04	92.22	89.73
128	92.39	89.73	92.55	90.02
256	92.15	89.46	91.99	89.43

Table 2: Dependency parsing result with various training batch size  $b$  and without composition function. The results marked with asterisks were reported in the [Ballesteros et al. \(2017\)](#).

close to linear, which means there is very little overhead associated with our batching scheme. Quantitatively, according to Amdahl’s Law ([Amdahl, 1967](#)), the proportion of parallelized computations is 99.92% at batch size 64. We also compared our implementation with the implementation that comes with [Dyer et al. \(2015\)](#), which is implemented in C++ with DyNet ([Neubig et al., 2017a](#)). DyNet is known to be very optimized for CPU computations and hence their implementation is reasonably fast even without batching and GPU acceleration, as shown in Figure 1.<sup>4</sup> But we would like to point out that we focus on the speed-up we are able to obtain rather than the absolute speed, and that our batching scheme is framework-universal and superior speed might be obtained by combining our scheme with alternative frameworks or languages (for example, the torch C++ interface).

The dependency parsing results are shown in Table 2. Our implementation is able to yield better test set performance than that reported in [Ballesteros et al. \(2017\)](#) for all batch size configurations except 256, where we observe a modest performance loss. Like [Goyal et al. \(2017\)](#); [Keskar et al. \(2016\)](#); [Masters and Luschi \(2018\)](#), we initially observed more significant test-time performance deterioration (around 1% absolute difference) for models trained without learning rate warmup, and concurring with the findings in [Goyal et al. \(2017\)](#), we find warmup very helpful for stabilizing large-batch training. We did not run experiments with batch size below 8 as they are too slow due to

<sup>4</sup>Measured on one core of an Intel Xeon E7-4830 CPU.

Python’s inherent performance issue.

## 5 Related Work

DyNet has support for automatic minibatching ([Neubig et al., 2017b](#)), which figures out what computation is able to be batched by traversing the computation graph to find homogeneous computations. While we cannot directly compare with that framework’s automatic batching solution for StackLSTM<sup>5</sup>, we can draw a loose comparison to the results reported in that paper for BiLSTM transition-based parsing ([Kiperwasser and Goldberg, 2016](#)). Comparing batch size of 64 to batch size of 1, they obtained a 3.64x speed-up on CPU and 2.73x speed-up on Tesla K80 GPU, while our architecture-specific manual batching scheme obtained 60.8x speed-up. The main reason for this difference is that their graph-traversing automatic batching scheme carries a much larger overhead compared to our manual batching approach.

Another toolkit that supports automatic minibatching is Matchbox<sup>6</sup>, which operates by analyzing the single-instance model definition and deterministically convert the operations into their minibatched counterparts. While such mechanism eliminated the need to traverse the whole computation graph, it cannot homogenize the operations in each branch of `if`. Instead, it needs to perform each operation separately and apply masking on the result, while our method does not require any masking. Unfortunately we are also not able to compare with the toolkit at the time of this work as it lacks support for several operations we need.

Similar to the spirit of our work, [Bowman et al. \(2016\)](#) attempted to parallelize StackLSTM by using *Thin-stack*, a data structure that reduces the space complexity by storing all the intermediate stack top elements in a tensor and use a queue to control element access. However, thanks to PyTorch, our implementation is not directly dependent on the notion of *Thin-stack*. Instead, when an element is popped from the stack, we simply shift the stack top pointer and potentially re-write the corresponding sub-tensor later. In other words, there is no need for us to directly maintain all the intermediate stack top elements, because in PyTorch, when the element in the stack is re-written, its underlying sub-tensor will not be destructed as there

<sup>5</sup>This is due to the fact that DyNet automatic batching cannot handle graph structures that depends on runtime input values, which is the case in StackLSTM.

<sup>6</sup><https://github.com/salesforce/matchbox>



are still nodes in the computation graph that point to it. Hence, when performing back-propagation, the gradient is still able to flow back to the elements that are previously popped from the stack and their respective precedents. Hence, we are also effectively storing all the intermediate stack top elements only once. Besides, Bowman et al. (2016) didn't attempt to eliminate the conditional branches in the StackLSTM algorithm, which is the main algorithmic contribution of this work.

## 6 Conclusion

We propose a parallelizable version of StackLSTM that is able to fully exploit the GPU parallelism by performing minibatched training. Empirical results show that our parallelization scheme yields comparable performance to previous work, and our method scales up very linearly with the increasing batch size.

Because our parallelization scheme is based on the observation made in section 1, we cannot incorporate batching for neither Arc-Standard transition system nor the token-level composition function proposed in Dyer et al. (2015) efficiently yet. We leave the parallelization of these architectures to future work.

Our parallelization scheme makes it feasible to run large-data experiments for various tasks that requires large training data to perform well, such as RNNG-based syntax-aware neural machine translation (Eriguchi et al., 2017).

## Acknowledgement

The authors would like to thank Peng Qi for helping with data preprocessing and James Bradbury for helpful technical discussions. This material is based upon work supported in part by the DARPA LORELEI and IARPA MATERIAL programs.

## References

Gene M. Amdahl. 1967. Validity of the single processor approach to achieving large scale computing capabilities. In *American Federation of Information Processing Societies: Proceedings of the AFIPS '67 Spring Joint Computer Conference, April 18-20, 1967, Atlantic City, New Jersey, USA*, pages 483–485.

Miguel Ballesteros, Chris Dyer, Yoav Goldberg, and Noah A. Smith. 2017. Greedy transition-based dependency parsing with stack lstms. *Computational Linguistics*, 43(2):311–347.

Samuel R. Bowman, Jon Gauthier, Abhinav Rastogi, Raghav Gupta, Christopher D. Manning, and Christopher Potts. 2016. A fast unified model for parsing and sentence understanding. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*.

Danqi Chen and Christopher D. Manning. 2014. A fast and accurate dependency parser using neural networks. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25-29, 2014, Doha, Qatar, A meeting of SIGDAT, a Special Interest Group of the ACL*, pages 740–750.

Chris Dyer, Miguel Ballesteros, Wang Ling, Austin Matthews, and Noah A. Smith. 2015. 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 of the Asian Federation of Natural Language Processing, ACL 2015, July 26-31, 2015, Beijing, China, Volume 1: Long Papers*, pages 334–343.

Chris Dyer, Adhiguna Kuncoro, Miguel Ballesteros, and Noah A. Smith. 2016. Recurrent neural network grammars. In *NAACL HLT 2016, The 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, San Diego California, USA, June 12-17, 2016*, pages 199–209.

Akiko Eriguchi, Yoshimasa Tsuruoka, and Kyunghyun Cho. 2017. Learning to parse and translate improves neural machine translation. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 2: Short Papers*, pages 72–78.

Priya Goyal, Piotr Dollár, Ross B. Girshick, Pieter Noordhuis, Lukasz Wesolowski, Aapo Kyrola, Andrew Tulloch, Yangqing Jia, and Kaiming He. 2017. Accurate, large minibatch SGD: training imagenet in 1 hour. *CoRR*, abs/1706.02677.

Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural Computation*, 9(8):1735–1780.

Nitish Shirish Keskar, Dheevatsa Mudigere, Jorge Nocedal, Mikhail Smelyanskiy, and Ping Tak Peter Tang. 2016. On large-batch training for deep learning: Generalization gap and sharp minima. *CoRR*, abs/1609.04836.

Diederik P. Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980.

Eliyahu Kiperwasser and Yoav Goldberg. 2016. Simple and accurate dependency parsing using bidirectional LSTM feature representations. *CoRR*, abs/1603.04351.

- Marco Kuhlmann, Carlos Gómez-Rodríguez, and Giorgio Satta. 2011. Dynamic programming algorithms for transition-based dependency parsers. In *The 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies, Proceedings of the Conference, 19-24 June, 2011, Portland, Oregon, USA*, pages 673–682.
- Marie-Catherine de Marneffe, Bill MacCartney, and Christopher D. Manning. 2006. Generating typed dependency parses from phrase structure parses. In *Proceedings of the Fifth International Conference on Language Resources and Evaluation, LREC 2006, Genoa, Italy, May 22-28, 2006.*, pages 449–454.
- Dominic Masters and Carlo Luschi. 2018. Revisiting small batch training for deep neural networks. *CoRR*, abs/1804.07612.
- Graham Neubig, Chris Dyer, Yoav Goldberg, Austin Matthews, Waleed Ammar, Antonios Anastasopoulos, Miguel Ballesteros, David Chiang, Daniel Clothiaux, Trevor Cohn, Kevin Duh, Manaal Faruqui, Cynthia Gan, Dan Garrette, Yangfeng Ji, Lingpeng Kong, Adhiguna Kuncoro, Gaurav Kumar, Chaitanya Malaviya, Paul Michel, Yusuke Oda, Matthew Richardson, Naomi Saphra, Swabha Swayamdipta, and Pengcheng Yin. 2017a. Dynet: The dynamic neural network toolkit. *CoRR*, abs/1701.03980.
- Graham Neubig, Yoav Goldberg, and Chris Dyer. 2017b. On-the-fly operation batching in dynamic computation graphs. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*, pages 3974–3984.
- Joakim Nivre. 2008. Algorithms for deterministic incremental dependency parsing. *Computational Linguistics*, 34(4):513–553.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. 2017. Automatic differentiation in pytorch. In *NIPS-W*.
- Peng Qi and Christopher D. Manning. 2017. Arc-swift: A novel transition system for dependency parsing. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 2: Short Papers*, pages 110–117.
- Richard Socher, Cliff Chiung-Yu Lin, Andrew Y. Ng, and Christopher D. Manning. 2011. Parsing natural scenes and natural language with recursive neural networks. In *Proceedings of the 28th International Conference on Machine Learning, ICML 2011, Bellevue, Washington, USA, June 28 - July 2, 2011*, pages 129–136.
- Kristina Toutanova, Dan Klein, Christopher D. Manning, and Yoram Singer. 2003. Feature-rich part-of-speech tagging with a cyclic dependency network. In *Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics, HLT-NAACL 2003, Edmonton, Canada, May 27 - June 1, 2003*.