

Effective Batching for Recurrent Neural Network Grammars

Hiroshi Noji

Artificial Intelligence Research Center
AIST

hiroshi.noji@aist.go.jp

Yohei Oseki

Graduate School of Arts and Sciences
University of Tokyo

oseki@g.ecc.u-tokyo.ac.jp

Abstract

As a language model that integrates traditional symbolic operations and flexible neural representations, recurrent neural network grammars (RNNGs) have attracted great attention from both scientific and engineering perspectives. However, RNNGs are known to be harder to scale due to the difficulty of batched training. In this paper, we propose effective batching for RNNGs, where every operation is computed in parallel with tensors across multiple sentences. Our PyTorch implementation effectively employs a GPU and achieves x6 speedup compared to the existing C++ DyNet implementation with model-independent auto-batching. Moreover, our batched RNNG also accelerates inference and achieves x20-150 speedup for beam search depending on beam sizes. Finally, we evaluate syntactic generalization performance of the scaled RNNG against the LSTM baseline, based on the large training data of 100M tokens from English Wikipedia and the broad-coverage targeted syntactic evaluation benchmark.¹

1 Introduction

Neural language models have an excellent word prediction ability, which motivates researchers to develop several analysis methods for fine-grained evaluation, aiming at understanding which linguistic abilities the models have acquired during training (Linzen et al., 2016; Wilcox et al., 2018; Marvin and Linzen, 2018; Warstadt et al., 2020). So far, many efforts have been made on the evaluation of *syntactic* performance of models, including the abilities to resolve distant subject-verb number agreement in English. Since neural language models are the foundation of contemporary NLP systems, building a language model having robust sentence processing abilities like humans is an

¹Our RNNG implementation is available at <https://github.com/aistairc/rnng-pytorch/>.

important goal, especially toward a system with human-like syntactic generalization abilities, not relying on the data-specific superficial cues found in the training data (McCoy et al., 2019; Linzen, 2020).

Past work has revealed that while sequential and unstructured models, such as LSTM and Transformer language models (Hochreiter and Schmidhuber, 1997; Vaswani et al., 2017), can induce several interesting syntactic behaviors, there is also a notable advantage in explicitly modeling syntax with specific architectures (Kuncoro et al., 2018; Wilcox et al., 2019; Hu et al., 2020). The representative of such models is the recurrent neural network grammars (RNNGs; Dyer et al., 2016), the top-down, left-to-right generative models of a parse tree and sentence.

While these results may suggest that RNNGs are a better modeling choice for language, unfortunately, they have a practical drawback in terms of scalability, due to their structure-sensitive computation mechanism (Kuncoro et al., 2019). Since the computational graphs of RNNGs depend on the tree structures of the sentences, training cannot be mini-batched easily. This is in contrast to LSTMs and Transformers, for which token-wise operations can be batched across sentences, allowing efficient computation on GPUs, which is the key to the data scalability. Although RNNGs are claimed to be a fascinating language model, in practice, they still do not replace the unstructured, computationally favorable models like LSTMs.

In this paper, we directly address the data scalability issue of RNNGs by showing that most computations during training can be batched across sentences. At the computational core of RNNGs is stack LSTMs (Dyer et al., 2015). In past work, Ding and Koehn (2019) have already shown that stack LSTM update operations can be reduced to a tensor operation by implementing the stack as a sin-

gle tensor with predefined maximum stack depth. Our work is built on this idea but with a few additional techniques to bridge the gap between the simple stack LSTMs and RNNGs. Importantly, we devise the efficient batching method for *composition* operations on the arbitrary number of stack items, which is unsolved in previous work.

The existing RNNG implementation is based on DyNet that supports the mechanism called Auto-batch (Neubig et al., 2017), which automatically finds mini-batch units from the independent computational graphs over multiple sentences with lazy computation. While this mechanism is model-independent and allows intuitive implementation, the utility of this method rapidly plateaus as we increase the batch size. On the other hand, our present method allows effective parallel computation, increasing the training speed almost linearly as we increase the batch size.

In addition to this new batching mechanism for improved scalability, we also provide a new analysis on the role of the strong syntactic inductive bias for models that can access the larger amount of data. For syntactic generalization abilities, while Hu et al. (2020) suggest that the model inductive bias plays a more important role than data scale, they also report that LSTMs or Transformers such as the off-the-shelf large-scale models (e.g., GPT-2 or JRNN) perform much better than their scale-controlled LSTMs. Does an RNNG, which already works relatively well on a modest amount of data, still benefit from the data scale to further strengthen its syntactic ability? We train a new RNNG on about 100M tokens in Wikipedia and evaluate its syntactic performance on SyntaxGym test circuits (Gauthier et al., 2020), finding that the data scale generally brings further performance gains, while the model tends to lose some heuristics on surface patterns that LSTMs seem to find. Our result suggests that RNNGs’ reliance on structures will be strengthened with more data, motivating future research on developing better syntactic representation itself as supervision to structured language models.

A related approach to our work is adding the syntactic bias into *sequential* language models, such as LSTMs, with knowledge distillation from RNNGs (Kuncoro et al., 2019, 2020). While motivations are similar, we provide a rather direct solution to resolve the scalability issue of RNNGs, opening up a new possibility of directly using them as an alternative to LSTMs.

From another perspective, our work can be complementary to this work, because knowledge distillation requires a teacher RNNG model, which itself is costly to obtain. For example, Kuncoro et al. (2020) trained an RNNG on a relatively large dataset of 3.6M sentences, which is approximately similar to the training data we use. While the detail is missing, they report that training takes three weeks on a GPU. On the other hand, our models get almost converged in three days. This direct improvement in training time greatly expands the applicability of RNNGs including a teacher of sequential models, and more direct use in computational psycholinguistics (Hu et al., 2020) and NLP applications such as syntactic neural machine translation (Eriguchi et al., 2017).

2 Preliminaries

2.1 Recurrent neural network grammars

RNNGs are joint generative models of a sentence and constituency tree. While RNN language models assign a next token probability, RNNGs assign a probability to next *action*, by which the parse state (stack LSTM) changes dynamically. In this work, we focus on the stack-only RNNG (Kuncoro et al., 2017), which has some resemblance to RNNs in that a single state vector \mathbf{h}_t defines next action probability a_t :

$$a_t \sim \text{softmax}(\mathbf{W}_a \text{MLP}(\mathbf{h}_t) + \mathbf{b}_a)$$

At each step, \mathbf{h}_t is obtained from the top element of stack LSTM, which preserves intermediate LSTM states up to \mathbf{h}_t . As a preparation for our batched RNNGs (Section 3), we try to formalize how this stack LSTM states change with each action. An RNNG internally preserves two different stacks: S_h and S_e . S_h is a stack LSTM, keeping the LSTM hidden states $\mathbf{h}_0 \cdots \mathbf{h}_t$.² S_e keeps stack elements, each of which is a word embedding \mathbf{e}_w , an open nonterminal embedding \mathbf{e}_x , or a closed constituent embedding \mathbf{e}_c obtained by REDUCE action.

At each step, the number of candidate actions is $|\mathcal{N}| + 2$ given the set of nonterminal symbols \mathcal{N} . Each action changes S_h and S_e as follows:

- NT(x): Push open nonterminal embedding \mathbf{e}_x onto S_e , getting a new LSTM state by $\mathbf{h}_{\text{new}} = \text{LSTM}(\text{top}(S_h), \mathbf{e}_x)$, and then push \mathbf{h}_{new} onto

²Precisely, we also have to keep LSTM cell states. We omit this part for brevity.

S_h . This action corresponds to generating an open nonterminal, e.g. “(VP” (when $x=VP$), which will be closed later.

- GEN: First, generate a next token by sampling from $w \sim \text{softmax}(\mathbf{W}_w \text{MLP}(\mathbf{h}_t) + \mathbf{b}_w)$. Then, as in NT, push \mathbf{e}_w onto S_e , getting a new LSTM state $\mathbf{h}_{\text{new}} = \text{LSTM}(\text{top}(S_h), \mathbf{e}_w)$, and push \mathbf{h}_{new} onto S_h .
- REDUCE: First, repeatedly pop from S_e k -times until we find \mathbf{e}_x , an open nonterminal embedding. Letting $\mathbf{e}^{t-k} = \mathbf{e}_x$, then, apply a composition function, which is BiLSTM (Dyer et al., 2016) by default, to obtain a composed phrase representation \mathbf{e}_c :

$$\mathbf{e}_c = \text{BiLSTM}([\mathbf{e}^{t-k}, \dots, \mathbf{e}^t]).$$

\mathbf{e}_c is then pushed onto S_e . To synchronize two stacks, we also pop k -times from S_h and update the LSTM state by $\mathbf{h}_{\text{new}} = \text{LSTM}(\text{top}(S_h), \mathbf{e}_c)$, pushing it onto S_h .

By declaring the operations as above, we notice that the main reasons to prevent mini-batching are twofold: (1) the stacks have variable length, which varies at each step for each sentence; and more crucially, (2) internal operations in an action, especially in REDUCE and others, are largely different.

As we describe next, the issue regarding (1) has been largely solved in previous work. For (2), our strategy is essentially not joining different action types, but trying to improve the efficiency of each action as much as possible after grouping by action types. We find that in practice this strategy works quite well (Section 5.2), allowing models to benefit from a large batch size effectively.

2.2 Batched stack LSTMs

Ding and Koehn (2019) propose a sentence-level batched training algorithm for a restricted class of stack LSTMs designed for unlabeled dependency parsing without composition operations (Dyer et al., 2015). More specifically, Ding and Koehn (2019) deal with the parsing models defined by the following two operations only:³

- PUSH: Push $\text{LSTM}(\text{top}(S_h), \mathbf{e}_w)$ to S_h . \mathbf{e}_w is the embedding of the next token.

³A restricted model of an arc-eager system (Nivre, 2004), which just POPs when LEFT-ARC occurs, can be achieved with these operations. RIGHT-ARC is modeled by PUSH. Essentially, this stack LSTM can only model the right spine of a tree at each step.

- POP: Pop the top element from S_h .

At each step, the next action is either PUSH or POP for each sentence. This model still suffers from the problem (1) above. However, they show that by changing the data structure of stack, next PUSH and POP across sentences can be performed in batch. Given B sentences in a batch, let S_h^i be a stack for i -th sentence. What we need to do is to access all top elements of S_h^i ($i \in [0, \dots, B-1]$) jointly, and this is possible by summarizing all stacks into a single stack *tensor*, denoted by \mathbf{S}_h , for which $\mathbf{S}_h[i, p]$ denotes p -th element (LSTM state) on the stack for i -th sentence.

The core idea behind achieving PUSH and POP jointly is that we perform LSTM updates for *all* stack top elements in a batch, but *only* proceed top stack pointers for PUSH batches. Given next actions $\mathbf{a} = [\text{PUSH}, \text{PUSH}, \text{POP}, \dots]$ of length B , we get a vector $\mathbf{op} = [+1, +1, -1, \dots]$, denoting whether next stack pointer is +1 (PUSH) or -1 (POP). By keeping stack top pointer vector \mathbf{p}_h , each step can be batched as the following two operations:

$$\begin{aligned} \mathbf{S}_h[(0, \mathbf{p}_h[0] + 1) \cdots (B-1, \mathbf{p}_h[B-1] + 1)] &\leftarrow \\ \text{LSTM}(\mathbf{S}_h[(0, \mathbf{p}_h[0]) \cdots (B-1, \mathbf{p}_h[B-1])], \mathbf{E}_w), \\ \mathbf{p}_h &\leftarrow \mathbf{p}_h + \mathbf{op}, \end{aligned}$$

in which \mathbf{E}_w is the next token embeddings.

Unfortunately, this batching relies on a strong assumption about models that one action (PUSH) involves all operations (LSTM update and pointer move by \mathbf{op}). This is not the case for RNNs, for which any action cannot be reduced to a subset of other actions, necessitating a different strategy.

3 Batched RNNs

Our batching algorithm for RNNs is built on the following two observations:

- For all $a_t \in \{\text{NT}, \text{GEN}, \text{REDUCE}\}$, the last step is common and corresponds to PUSH operation for stack LSTM above with newly created embeddings $\{\mathbf{e}_x, \mathbf{e}_w, \mathbf{e}_c\}$. This final step can be batched if we get all new embeddings as a single tensor \mathbf{E}_{next} (with size of $(B, |\mathbf{e}|)$).
- Then, the main problem is reduced to getting \mathbf{E}_{next} efficiently. This is possible by separately filling \mathbf{E}_{next} for each action, using a few additional pointer vectors to keep track of a stack state for each sentence.

Algorithm 1 One training step for batched RNNG

Input Next action vector at \mathbf{a} ;
 index vector for each action: $\mathbf{i}_{\text{gen}}, \mathbf{i}_{\text{nt}}, \mathbf{i}_{\text{red}}$

- 1: $\mathbf{E}_{\text{next}} \leftarrow$ new tensor of size $(B, |\mathbf{e}_x|)$
- 2: $\mathbf{E}_{\text{next}}[\mathbf{i}_{\text{gen}}] \leftarrow \text{word_emb}(\mathbf{x}[(\mathbf{i}_{\text{gen}}, \mathbf{b}[\mathbf{i}_{\text{gen}}])])$
- 3: $\mathbf{b}[\mathbf{i}_{\text{gen}}] = \mathbf{b}[\mathbf{i}_{\text{gen}}] + 1$ ▷ Move to next word.
- 4: $\mathbf{E}_{\text{next}}[\mathbf{i}_{\text{nt}}] \leftarrow \text{nt_emb}(\mathbf{a}[\mathbf{i}_{\text{nt}}])$
- 5: $\mathbf{p}_q[\mathbf{i}_{\text{nt}}] \leftarrow \mathbf{p}_q[\mathbf{i}_{\text{nt}}] + 1$
- 6: $\mathbf{q}[(\mathbf{i}_{\text{nt}}, \mathbf{p}_q[\mathbf{i}_{\text{nt}}])] \leftarrow \mathbf{p}_h[\mathbf{i}_{\text{nt}}] + 1$ ▷ Keep new NT depth.
- 7: $\mathbf{p}_{\text{prev_nt}} \leftarrow \mathbf{q}[(\mathbf{i}_{\text{red}}, \mathbf{p}_q[\mathbf{i}_{\text{red}}])]$
- 8: $\mathbf{E}_{\text{red}} \leftarrow \text{gather_children}(\mathbf{p}_{\text{prev_nt}}, \mathbf{p}_h[\mathbf{i}_{\text{red}}], \mathbf{S}_e)$
- 9: $\mathbf{E}_{\text{next}}[\mathbf{i}_{\text{red}}] \leftarrow \text{BiLSTM}(\mathbf{E}_{\text{red}})$ ▷ Composition.
- 10: $\mathbf{p}_q[\mathbf{i}_{\text{red}}] \leftarrow \mathbf{p}_q[\mathbf{i}_{\text{red}}] - 1$ ▷ Forget about reduced nts.
- 11: $\mathbf{p}_h[\mathbf{i}_{\text{red}}] \leftarrow \mathbf{p}_{\text{prev_nt}} - 1$
- 12: $\mathbf{p}_h \leftarrow \mathbf{p}_h + 1$
- 13: $\mathbf{S}_h[\mathbf{p}_h] \leftarrow \text{LSTM}(\mathbf{S}_h[\mathbf{p}_h - 1], \mathbf{E}_{\text{next}})$
- 14: $\mathbf{S}_e[\mathbf{p}_h] \leftarrow \mathbf{E}_{\text{next}}$

To obtain \mathbf{E}_{next} , for NT and GEN, we just need to lookup embeddings for next words and nonterminal symbols. We need an additional effort to obtain multiple \mathbf{e}_c s at once. Assuming a stack *tensor* as in Ding and Koehn (2019), we wish to pop k elements, up to \mathbf{e}_c , for multiple stacks by a single operation. Now the stack top positions can be accessed by \mathbf{p}_h (Section 2.2), which will be the end indices. To obtain the last open nonterminal positions across a batch, just keeping the last nonterminal positions is insufficient because there are multiple open nonterminals in general. The following matrix and vector allow this operation:

- \mathbf{q} : A matrix of size (B, D) given a predefined stack depth bound D . $\mathbf{q}[b, d]$ denotes the position of d -th nonterminal on the b -th stack.
- \mathbf{p}_q : A B -dimensional vector, pointing to the last index of \mathbf{q} (similar to \mathbf{p}_h for \mathbf{S}_h).

For example, by $\mathbf{q}[(0, \mathbf{p}_q[0]) \cdots (B - 1, \mathbf{p}_q[B - 1])]$, we can retrieve all the top open nonterminal positions in a batch. Note that for each $\mathbf{q}[b]$, the index beyond $\mathbf{p}_q[b]$ will not be accessed, so we can signify the remove of top nonterminals just by a decrement of \mathbf{p}_q without updating \mathbf{q} .

We need a few additional tensors to achieve fully batched stack tensor operations. Figure 1 shows an example.

- \mathbf{S}_h : A tensor of size (B, D, L, H) when the stack LSTM has L layers with H hidden dimensions. The core of batched stack LSTMs.
- \mathbf{S}_e : A tensor of size $(B, D, |\mathbf{e}|)$, corresponding to \mathbf{S}_e in non-batched models (Section 2.1).
- \mathbf{b} : A B -dimensional vector, keeping the next token index in the sentence.

		Stack	Next action
0	(S ₁ So ₂ (NP it) ₃		(VP
0	(S ₁ (NP he) ₂ (VP ₃ said)
0	(S ₁ (NP ₂ (NP ₃ A ₄ branch ₅))
		$\mathbf{b} = [1, 1, 2]$	
		$\mathbf{p}_h = [3, 3, 5]$	
		$\mathbf{p}_q = [0, 1, 2]$	
		$\mathbf{q} = \begin{bmatrix} \underline{1} & 3 & 0 & \cdots \\ 1 & \underline{3} & 0 & \cdots \\ 1 & 2 & \underline{3} & \cdots \end{bmatrix}$	

Figure 1: Example batched stack configuration. $x|_d$ means that x is at depth d . For example, “(NP it)” in the first sentence is closed so constitutes a single item on the stack. \mathbf{p}_q points to the top positions of \mathbf{q} , which are underlined.

- \mathbf{p}_h : A B -dimensional vector, pointing to the top elements of \mathbf{S}_h .

Full algorithm Algorithm 1 describes operations in each step given next actions \mathbf{a} . Action index vector \mathbf{i}_a , keeps the indices of action a in \mathbf{a} ; in Figure 1, $\mathbf{i}_{\text{gen}} = [1]$ and $\mathbf{i}_{\text{nt}} = [0]$. The operations are mainly categorized into filling \mathbf{E}_{next} for each action (in red), pointer updates according to action definitions (3, 5, 6, 10, 12), and finally stack updates (13, 14), corresponding to the observed common operations (a).⁴ *gather_children* is a function that returns a tensor summarizing reduced children node embeddings. Since the number of reduced children differs across batch, we implement this to return a padded tensor, using *gather* function in PyTorch.

Deviated from Ding and Koehn (2019), we *separately* perform each action, as indicated by the use of \mathbf{i}_a . This can be seen as a deficiency of our algorithm; however, this separation is necessary beyond very simple models, which are practically less attractive. Rather, our strategy can be applied to broader classes of structured neural models, including dependency parsing with composition, and we believe that our empirical success (Section 5) encourages further exploration of the presented strategy to various models.

How to set D ? As in Ding and Koehn (2019), we need to specify the stack depth bound D for each batch. Increasing this value incurs more GPU memory. For training, we can precompute the minimum value of each sentence by simulating oracle transitions beforehand and use it. For inference, we fix $D = 100$, since we find that even for very long

⁴By $A[(\mathbf{x}, \mathbf{y})]$ for vectors \mathbf{x} and \mathbf{y} , both with length l , we mean $A[(x[0], y[0]) \cdots (x[l - 1], y[l - 1])]$, corresponding to advanced indexing in PyTorch. We regard that $\mathbf{S}_h[0]$ is fixed by initial hidden vectors while $\mathbf{S}_e[0]$ is kept empty. $\mathbf{p}_h[b] = 0$ means that b -th stack is empty.

sentences (more than 150 words), the stack depth will never exceed 80 for English sentences.

A note on extra memory with stack tensors At first sight, our approach seems to suffer from the limitation in scalability due to fixed stack tensors (\mathbf{S}_h and \mathbf{S}_e). The sizes of these tensors grow by model size, implying that we may not be able to employ a large batch size for a large model. In practice, however, this extra memory will not be a bottleneck in the total memory for training. This is because the main cause of required memory during training is rather a computational graph itself, which keeps all intermediate hidden states at each step. Our stack tensors can be seen as a “storage” to allow computing these intermediate values effectively with tensor operations. The extra memory for this storage is smaller than the total memory in a computational graph because the former depends on D while the latter depends on the total action length A , and $D \ll A$ in general.⁵

4 Other Improvements

Batched beam search For inference as a language model or as an incremental parser, RNNs typically employ word-synchronous beam search (Stern et al., 2017; Hale et al., 2018), which is although known to be very slow (Crabbé et al., 2019) because it often requires large beam sizes, such as 100 or 1000, and operations are not batched. As a by-product of our batched training, we succeed at implementing fully batched beam search for RNNs, excluding any *for* loops, by which we drastically improve the search speed (Section 5.3). This is possible by adding the “beam” dimension to all state tensors (\mathbf{S}_h , \mathbf{q} , etc.).

Subwords Given an increased amount of training data, the vocabulary size naturally increases. To suppress this effect, using subwords (Sennrich et al., 2016) now becomes a standard technique. We thus incorporate subword modeling into our

⁵Our preliminary experiment suggests that our RNN implementation can be scaled at least comparable model and data sizes to ELMo (Peters et al., 2018), a large-scale LSTM-based model, given a similar amount of computing resources. We examine the maximum allowable batch size for a model with 1,256 hidden dimensions, amounting to 94M parameters, which are comparable to ELMo (93M), and find that the batch size can be increased to 256, with the maximum action size in a batch of 16,000 (see Section 5.1) on a single V100 GPU (16GB). Transformer-level scalability (Devlin et al., 2019) would still be infeasible because of the RNN’s limited parallelism that is only on sentence-level, not token-level as in Transformers.

RNNs and employ it for our largest experiment in Section 6. Kuncoro et al. (2020) recently incorporate subwords in RNNs, in which, each word is regarded as a new constituent with WORD label, e.g., (WORD cu| r| ry). This means that models always need to perform additional NT(WORD) and REDUCE for each *token*, even for unsegmented ones, e.g., (WORD I), greatly increasing the average action sequence length, which in turn affect the training time. In this work, we model subwords by a simpler method of just segmenting each token. For example, an NP looks like (NP Th| ai cu| r| ry). While Kuncoro et al. (2020) note that this simple modeling is less effective, our experiments suggest that this is a good enough strategy, considering the added computational costs with NT(WORD) actions.⁶

5 Evaluating Efficiency of Batching

The main focus of this section is a comparative evaluation of our PyTorch RNN implementation with the existing DyNet implementation.⁷ We show that: (1) with a large batch size training speed drastically improves, and models will tend to find better parameters (Section 5.2); and (2) our batched beam search hugely speeds up inference (Section 5.3).

5.1 Setting

While Penn Treebank (PTB; Marcus et al., 1993) has often been used to train RNNs (Wilcox et al., 2019, 2020), it is too small and here we use a larger dataset of BLLIP corpus (Charniak et al., 2000), expecting that the effects of large batch size become clearer by this modestly sized dataset.

Preprocessing For preprocessing, we largely follow Hu et al. (2020), which also train an RNN on this dataset. We partition the data according to their LG size, the largest training setting, amounting to 42 million tokens for training and 1,500 sentences for development. One difference we make is the handling of unknown tokens. We limit the vocabulary by the top frequent 50,000 word types in the training data. Hu et al. (2020) use all word types

⁶We provide a pilot study about this method in Appendix B. Using BLLIP corpus and Penn Treebank, we explore the relationship between a suitable number of subword units and model sizes. The main result is that large subword units are effective for larger models, and also subword modeling almost always improves parsing accuracy.

⁷<https://github.com/cpllab/rnng-incremental>. This implementation supports word-synchronous beam search. For the training part, this is not implemented to use DyNet Autobatch so we modified it to enable that.

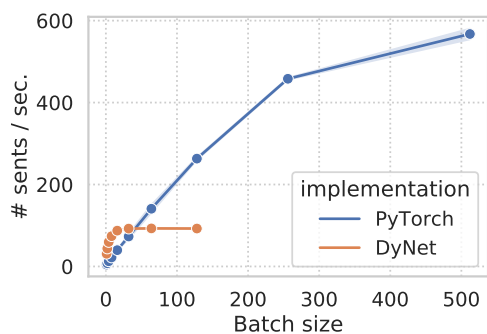


Figure 2: Training speed comparison (number of processed sentences / sec.) when increasing batch sizes to [1, 2, 4, \dots , 512]. Shade denotes standard deviation.

that appear at least twice; however, this method vastly increases the vocabulary size and hence the model size. Unknown tokens are created in the same way with the Berkeley parser’s surface feature rule (Petrov et al., 2006). The way to annotate parse trees is the same as well; we run Berkeley neural parser (Kitaev et al., 2019), a state-of-the-art constituency parser to assign accurate parses.

Model size and parameters We experiment with the most common model size of RNNG in the literature: 256 dimensions for input and LSTM hidden dimensions, with 2 layer LSTMs (Dyer et al., 2016; Hu et al., 2020). The total number of parameters is about 15M. The hyperparameters are summarized in Appendix A.

Other settings We employ some additional techniques to improve the efficiency of our batching mechanism. First, before training, we group sentences by their number of gold actions so that examples in each mini-batch have similar numbers of actions. Specifically, we first sort the sentences by action lengths, divide by every 4096 sentences, and then sample each batch from a single group.

Second, we predefine the maximum value for the total number of actions across sentences in a batch, which we set to 26,000. This is inspired by a similar mechanism in fairseq (Ott et al., 2019) for the maximum number of tokens. Using this means that the number of sentences in a batch will be adjusted to be smaller than the batch size when the action sequences (or sentences) are long, allowing us to interpret given batch size as the maximum that is fully exploited only for shorter sentences, which are in practice dominant in the data.⁸

⁸To reduce the memory further, we also employ mixed-precision training in PyTorch. Inference is performed with




Figure 3: Training wall clock time vs. total validation loss for different batch sizes. X-value of i -th point is an averaged duration time to i -th validation step across three random seeds. Shade denotes standard deviation.

Every experiment is run on a single V100 GPU with 16GB memory. Unless otherwise noted, we perform every experiment three times with different random seeds, reporting an average score with standard deviation.

5.2 Effects of batch sizes

Although our batched training involves action-specific operations (Section 3), to our surprise, the efficiency improvement for our RNNG with large batch sizes is almost linear up to 256 (Figure 2). The improvement is narrow at 512, though this is mainly due to the restriction of the maximum number of actions in a batch (Section 5.1), which reduces the actual batch size for longer inputs. DyNet’s Autobatch is quite effective up to 16, running much faster than ours due to the speed of C++, but further improvement is not obtained probably because of the increased overhead of finding mini-batch units themselves from a large computational graph.

Though this result clearly demonstrates the efficiency of our batching mechanism, it is *only* meaningful when the large batch size in fact leads to a faster model *convergence*. This is the case, as shown in Figure 3, where we compare the total validation losses as a function of actual wall clock time during training. The loss is calculated every 1000 batches. The model with batch size 512 converges fastest, and importantly, to better parameters. This result suggests that we can safely benefit from large batch size as long as memory permits. In the following experiments, we fix the batch size to 512.

half-precision (fp16). We find that this does not change the results of beam search at all.

5.3 Beam search speed improvement

As we discuss in Section 4, we have also improved the efficiency of word-synchronous beam search, a standard technique to calculate incremental prefix probabilities (Hale, 2001) and a parse tree for RNNs. Now, we evaluate the impact of this improvement. For PyTorch, we run it on V100 GPU; for DyNet, we find that it runs faster on CPUs so we instead use CPUs (Intel Xeon 6148, 20 cores x2), with Intel MKL. DyNet beam search is still too slow with this environment so we limit the number of tested sentences to 300 from the BLLIP development set. For PyTorch, we try two different batch sizes $\{1, 10\}$, with a restriction on the number of tokens in a batch, similarly to the total action size in training (Section 5.1). We fix this value to 250, with which the model can safely parse with the largest beam size of 1000.

Word-synchronous beam search employs two types of beam widths, action beam size (k) and word beam size (k_w), along with fast-tracked candidate size, denoted as k_s (see Stern et al. 2017). k is most akin to the standard beam size. Table 1 summarizes the results when increasing k (others are in the caption). The beam search of DyNet becomes prohibitively slow when $k \geq 50$. Strikingly, the increase in average runtime is more than linear against the beam size, especially for $10 \rightarrow 50$ and $50 \rightarrow 100$. The time increases, $0.5 \rightarrow 11.3$ (x22.6) and $11.3 \rightarrow 48.6$ (x4.3) are roughly quadratic to the increase of k (x5 and x2). This result is reasonable because in addition to the complexity of each step, which depends on k , the length of searched action sequence could also linearly grow by k .⁹ The naïve DyNet implementation directly suffers from this computational cost.

Our batched beam search largely resolves this issue and now the average runtime only gradually increases by k . We note that as a parser or a language model, this speed is still not very fast, considering that this is on a GPU.¹⁰ For the research purpose, including psycholinguistic assessments as

⁹For a sentence of length N , the runtime of beam search is $O(k \times N \times M_w)$, where M_w denotes the maximum number of actions between two tokens (until choosing next SHIFT). The expected number of actions between two tokens (bound by M_w) grows by k because at each step, with a large k the chance that non-shift beam items remain in the next beam increases; hence, the runtime becomes quadratic to k in the worst case. We conjecture that this inefficiency is bounded at some k (see $k = 200$), though is severe for smaller k s.

¹⁰For smaller k , we can increase the batch size and the maximum number of tokens in a batch to further speedup.

Action beam size k	10	50	100	200	400	1000
DyNet	0.5	11.3	48.6	100.6	201.3	NA
PyTorch (B=1)	1.7	2.1	2.3	2.5	2.9	4.1
PyTorch (B=10)	0.4	0.5	0.7	0.9	1.3	2.8

Table 1: Word-synchronous beam search speed (average seconds per sentence) comparison on the first 300 sentences in BLLIP development set. B denotes the batch size. Word beam sizes (k_w) / fast track sizes (k_s) are 10/1, 10/1, 10/1, 20/2, 40/4, and 100/10, respectively.

done in Section 6, however, this improvement is significant, making experiments much easier even with large beam sizes. We still need to work on improving efficiency further, possibly by modifying learning methods to replace word-synchronous search (Stanojević and Steedman, 2020).

6 Syntactic Generalization Ability of Scaled RNN

Finally, we evaluate the syntactic generalization abilities of the scaled RNN. For this purpose, we adopt the test circuits used in Hu et al. (2020) via SyntaxGym (Gauthier et al., 2020). Here, a test *circuit* is a collection of test *suites*; e.g., “Long-Distance Dependencies” circuit contains a suite on a specific type of “filler-gap dependencies” as well as a suite on (pseudo) “cleft”. For each example in a suite, a model succeeds if it can assign a higher likelihood on a grammatically critical position in the correct sentence. For example, given “*The farmer near the clerks knows/*know many people.*” in the “Agreement” circuit, a model is correct if it assigns $p(\text{knows}|h) > p(\text{know}|h)$. Note that for subword models the total likelihoods on subwords (not averaged) are compared.

In the previous literature, Hu et al. (2020) trained an RNN on BLLIP (42M tokens). Here, we train subword RNNs on 100M tokens from English Wikipedia, to which we assign parse trees with Berkeley neural parser. The model size is 35M with 30k subword units, following the experiment in Appendix B, which assesses the suitable number of subword units for different model sizes. We train this RNN for three days (with three different seeds), and at inference fix the beam size k to 100 ($k_w = 10, k_s = 1$). We also train an RNN with the subset of this data (42M tokens) to separate the effects of data size. Our LSTM baseline is the one used in Noji and Takamura (2020), which is basically an AWD-LSTM-LM (Merity et al., 2018)

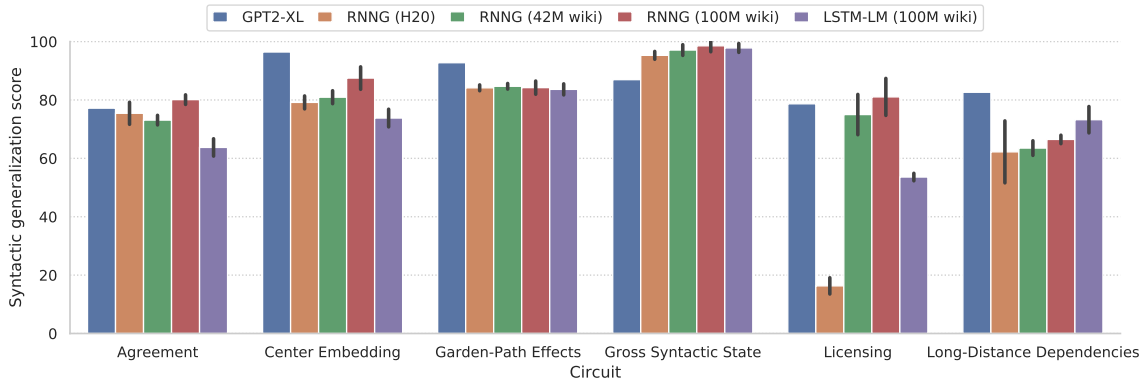


Figure 4: Circuit-level accuracies on SyntaxGym. For each circuit, suite-level accuracies are averaged across different test suites and random seeds to compute “Syntactic generalization score” of each model. Note that RNNG (H20) is a model trained on BLLIP (about 40M tokens) in Hu et al. (2020) but diverged from their results, because their suite-level accuracies are averaged across different models trained on various data sizes.

extended to be sentence-level and, on the Marvin and Linzen (2018) benchmark, shown to work better than GRNN (Gulordava et al., 2018), one of the models used in Hu et al. (2020).¹¹

The main result on circuit-level accuracies is summarized in Figure 4. On the effects of the data scale, we observe a consistent improvement from “RNNG (42M wiki)” to “RNNG (100M wiki)”. This result suggests that this amount of increase in training data is still beneficial for structural language models to strengthen their syntactic generalization ability. For some circuits, only RNNG (100M) outperforms GPT-2 (Radford et al., 2019) on average (“Agreement” and “Licensing”).

Comparing LSTM (100M) and RNNG (100M), RNNG generally outperforms LSTM, but with an exception on “Long-Distance Dependencies”. In order to inspect this, we break down this circuit into suites, (Figure 5), finding that this deficiency of RNNG is due to its poor performance on (pseudo) “cleft”, including the following example:

- (1) a. What he did was prepare the meal .
- b. *What he ate was prepare the meal .

On underlined tokens, models should assign a higher likelihood for (1a). Our LSTM performs nearly perfectly on these cases while our RNNGs perform badly. We conjecture that given more data and/or parameters, RNNGs will tend to strengthen

¹¹Our LSTM implementation is available at https://github.com/aistairc/lm_syntax_negative. We train this LSTM on our subword-segmented Wikipedia (30k units). The model size is adjusted so that the total number of parameters becomes 35M, the same size as our RNNGs (3 layer LSTMs with 1150 hidden and 450 input dimensions).




Figure 5: Suite-level accuracies on the “Long-Distance Dependencies” circuit. fgd-* is averaged across different test suites of filler-gap dependencies.

their commitment to provided syntactic supervision, and hence may lose some lexical heuristics which LSTMs can exploit from surface patterns (e.g., an association of *did* → *prepare*). In fact, the ability of LSTM on cleft is rather brittle, as shown in a huge drop on “cleft_modifier”, which include cleft constructions with intervening modifiers.

To rigorously handle these cases, models should notice that (1b) is a free relative clause and do not have an antecedent. However, the currently employed PTB annotation, which is limited to local structures, does not provide a distinction between these clause types, analyzing both as “(SBAR (WHNP What) (S (NP he) (VP did/ate)))”, which our RNNGs predict correctly. We also notice that Hu et al. (2020)’s RNNG (H20) performs rather similarly to our LSTM, while our RNNG (42M), trained on the comparable size of data to H20, is more similar to our RNNG (100M), suggesting that RNNG’s poor performance on cleft is not just due to the data scale. One possible explanation of the discrepancy between H20 and our 42M RNNG is

that our RNNG might be better optimized thanks to improved training, or due to the sizes of hidden layers (256 for H20 and 656 for ours).

This problem poses a new interesting challenge. While RNNGs have been compared to LSTMs several times, the provided syntactic structures are fixed and effects of different annotations (formalism, quantity, etc.) are not explored. For such investigation, the training cost of RNNGs has been a practical burden, but that problem largely goes away with the current study. We expect that our new implementation and batching strategy provide fruitful future research opportunities on structured neural language models.

7 Conclusion

A large computational cost of training structured neural language models was a main practical burden for employing these models in applications and analyses. With special focus on RNNGs, we have provided a direct solution to this problem by showing that batched effective training is in fact possible. On the large scale experiments with SyntaxGym test circuits, we found that the data quantity further strengthens the syntactic generalization abilities of RNNGs, while the annotation quality or quantity will also be of practical importance towards a language model with human-like strong syntactic performance.

Acknowledgement

This paper is based on results obtained from a project, JPNP20006, commissioned by the New Energy and Industrial Technology Development Organization (NEDO). This work was also supported by JSPS KAKENHI Grant Numbers 20K19877 and 19H04990, and the National Institute for Japanese Language and Linguistics (NINJAL) Collaborative Research Project “Computational Psycholinguistics of Language Processing with Large Corpora.”

References

Eugene Charniak, Don Blaheta, Niyu Ge, Keith Hall, John Hale, and Mark Johnson. 2000. [BLLIP 1987-89 WSJ Corpus Release 1 LDC2000T43](#). Linguistic Data Consortium.

Benoit Crabbé, Murielle Fabre, and Christophe Pallier. 2019. [Variable beam search for generative neural parsing and its relevance for the analysis of neuro-imaging signal](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language*

Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), pages 1150–1160, Hong Kong, China. Association for Computational Linguistics.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.

Shuoyang Ding and Philipp Koehn. 2019. [Parallelizable stack long short-term memory](#). In *Proceedings of the Third Workshop on Structured Prediction for NLP*, pages 1–6, Minneapolis, Minnesota. Association for Computational Linguistics.

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 (Volume 1: Long Papers)*, pages 334–343, Beijing, China. Association for Computational Linguistics.

Chris Dyer, Adhiguna Kuncoro, Miguel Ballesteros, and Noah A. Smith. 2016. [Recurrent neural network grammars](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 199–209, San Diego, California. Association for Computational Linguistics.

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 (Volume 2: Short Papers)*, pages 72–78, Vancouver, Canada. Association for Computational Linguistics.

Jon Gauthier, Jennifer Hu, Ethan Wilcox, Peng Qian, and Roger Levy. 2020. [SyntaxGym: An online platform for targeted evaluation of language models](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 70–76, Online. Association for Computational Linguistics.

Kristina Gulordava, Piotr Bojanowski, Edouard Grave, Tal Linzen, and Marco Baroni. 2018. [Colorless green recurrent networks dream hierarchically](#). 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 1195–1205, New Orleans, Louisiana. Association for Computational Linguistics.

- John Hale. 2001. [A probabilistic Earley parser as a psycholinguistic model](#). In *Second Meeting of the North American Chapter of the Association for Computational Linguistics*.
- John Hale, Chris Dyer, Adhiguna Kuncoro, and Jonathan Brennan. 2018. [Finding syntax in human encephalography with beam search](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2727–2736, Melbourne, Australia. Association for Computational Linguistics.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Jennifer Hu, Jon Gauthier, Peng Qian, Ethan Wilcox, and Roger Levy. 2020. [A systematic assessment of syntactic generalization in neural language models](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1725–1744, Online. Association for Computational Linguistics.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Nikita Kitaev, Steven Cao, and Dan Klein. 2019. [Multilingual constituency parsing with self-attention and pre-training](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3499–3505, Florence, Italy. Association for Computational Linguistics.
- Taku Kudo and John Richardson. 2018. [SentencePiece: A simple and language independent subword tokenizer and detokenizer for neural text processing](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 66–71, Brussels, Belgium. Association for Computational Linguistics.
- Adhiguna Kuncoro, Miguel Ballesteros, Lingpeng Kong, Chris Dyer, Graham Neubig, and Noah A. Smith. 2017. [What do recurrent neural network grammars learn about syntax?](#) In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 1249–1258, Valencia, Spain. Association for Computational Linguistics.
- Adhiguna Kuncoro, Chris Dyer, John Hale, Dani Yogatama, Stephen Clark, and Phil Blunsom. 2018. [LSTMs can learn syntax-sensitive dependencies well, but modeling structure makes them better](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1426–1436, Melbourne, Australia. Association for Computational Linguistics.
- Adhiguna Kuncoro, Chris Dyer, Laura Rimell, Stephen Clark, and Phil Blunsom. 2019. [Scalable syntax-aware language models using knowledge distillation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3472–3484, Florence, Italy. Association for Computational Linguistics.
- Adhiguna Kuncoro, Lingpeng Kong, Daniel Fried, Dani Yogatama, Laura Rimell, Chris Dyer, and Phil Blunsom. 2020. [Syntactic structure distillation pre-training for bidirectional encoders](#). *Transactions of the Association for Computational Linguistics*, 8:776–794.
- Tal Linzen. 2020. [How can we accelerate progress towards human-like linguistic generalization?](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5210–5217, Online. Association for Computational Linguistics.
- Tal Linzen, Emmanuel Dupoux, and Yoav Goldberg. 2016. [Assessing the ability of LSTMs to learn syntax-sensitive dependencies](#). *Transactions of the Association for Computational Linguistics*, 4:521–535.
- Mitchell P. Marcus, Beatrice Santorini, and Mary Ann Marcinkiewicz. 1993. [Building a large annotated corpus of English: The Penn Treebank](#). *Computational Linguistics*, 19(2):313–330.
- Rebecca Marvin and Tal Linzen. 2018. [Targeted syntactic evaluation of language models](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1192–1202, Brussels, Belgium. Association for Computational Linguistics.
- Tom McCoy, Ellie Pavlick, and Tal Linzen. 2019. [Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3428–3448, Florence, Italy. Association for Computational Linguistics.
- Stephen Merity, Nitish Shirish Keskar, and Richard Socher. 2018. [Regularizing and optimizing LSTM language models](#). In *International Conference on Learning Representations*.
- Sabrina J. Mielke. 2019. [Can you compare perplexity across different segmentations?](#)
- Sabrina J. Mielke, Ryan Cotterell, Kyle Gorman, Brian Roark, and Jason Eisner. 2019. [What kind of language is hard to language-model?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4975–4989, Florence, Italy. Association for Computational Linguistics.

- Graham Neubig, Yoav Goldberg, and Chris Dyer. 2017. [On-the-fly operation batching in dynamic computation graphs](#). In *Advances in Neural Information Processing Systems*, volume 30, pages 3971–3981. Curran Associates, Inc.
- Joakim Nivre. 2004. [Incrementality in deterministic dependency parsing](#). In *Proceedings of the Workshop on Incremental Parsing: Bringing Engineering and Cognition Together*, pages 50–57, Barcelona, Spain. Association for Computational Linguistics.
- Hiroshi Noji and Hiroya Takamura. 2020. [An analysis of the utility of explicit negative examples to improve the syntactic abilities of neural language models](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3375–3385, Online. Association for Computational Linguistics.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. [fairseq: A fast, extensible toolkit for sequence modeling](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 48–53, Minneapolis, Minnesota. Association for Computational Linguistics.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics.
- Slav Petrov, Leon Barrett, Romain Thibaux, and Dan Klein. 2006. [Learning accurate, compact, and interpretable tree annotation](#). In *Proceedings of the 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics*, pages 433–440, Sydney, Australia. Association for Computational Linguistics.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [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)*, pages 1715–1725, Berlin, Germany. Association for Computational Linguistics.
- Miloš Stanojević and Mark Steedman. 2020. [Max-margin incremental CCG parsing](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4111–4122, Online. Association for Computational Linguistics.
- Mitchell Stern, Daniel Fried, and Dan Klein. 2017. [Effective inference for generative neural parsing](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1695–1700, Copenhagen, Denmark. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30, pages 5998–6008. Curran Associates, Inc.
- Alex Warstadt, Alicia Parrish, Haokun Liu, Anhad Mohananey, Wei Peng, Sheng-Fu Wang, and Samuel R. Bowman. 2020. [BLiMP: The benchmark of linguistic minimal pairs for English](#). *Transactions of the Association for Computational Linguistics*, 8:377–392.
- Ethan Wilcox, Roger Levy, Takashi Morita, and Richard Futrell. 2018. [What do RNN language models learn about filler–gap dependencies?](#) In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 211–221, Brussels, Belgium. Association for Computational Linguistics.
- Ethan Wilcox, Peng Qian, Richard Futrell, Miguel Ballesteros, and Roger Levy. 2019. [Structural supervision improves learning of non-local grammatical dependencies](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 3302–3312, Minneapolis, Minnesota. Association for Computational Linguistics.
- Ethan Wilcox, Peng Qian, Richard Futrell, Ryosuke Kohita, Roger Levy, and Miguel Ballesteros. 2020. [Structural supervision improves few-shot learning and syntactic generalization in neural language models](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4640–4652, Online. Association for Computational Linguistics.

A Hyperparameters

We use the default parameter setting for the DyNet implementation. For our implementation, we use Adam optimizer (Kingma and Ba, 2015), which is found to be superior, while SGD has been used for DyNet implementation (Dyer et al., 2016; Wilcox et al., 2019). We set the learning rate and dropout rate to 0.001 and 0.1, respectively, which we find achieve lower validation loss robustly across different batch sizes.

B Effect of Number of Subword Units

We perform an experiment to understand the behaviors of our simple subword modeling (Section 4). We use the BLLIP corpus as preprocessed in Section 5.1 except the setting about vocabulary. We compare the fixed vocabulary models, which we train in the experiment of Section 5.2 (batch size 512), and several subword vocabulary models. The hyperparameters are the same as the fixed vocabulary models.

Subword units The number of subwords can be seen as a hyperparameter. To understand the effects of this size for RNNs, we prepare three different subword vocabularies: 10k (10,240), 20k (20,480), and 30k (30,720). We use byte-pair encoding (Sennrich et al., 2016) implemented in sentencepiece (Kudo and Richardson, 2018).

Model sizes We prepare two different model sizes, 15M and 35M, to see the interaction between the suitable size of subword units and model size, by adjusting the number of two dimensions so that the total number of parameters becomes comparable to these numbers. For 15M parameter models, the dimensions are 528 for 10k units; 432 for 20k; and 336 for 30k. For 35M parameter models, these are 864, 752, and 656, respectively. The number of LSTM layers is fixed to 2.

Results We investigate (1) the effectiveness of our simple subword modeling itself, and (2) whether the optimal number of subword units depends on model sizes. For (1), one way of evaluation is to compare the perplexities of subword models and fixed vocabulary models (see Section 5.1). However, they are not directly comparable because the fixed vocabulary models replace many tokens with unknown tokens, which are easy to predict and make the comparison unfair (Mielke et al., 2019).

V (# params.) / beam k	100	200	400	1000
$V_{fix} = 50k$ (15M)	92.34	93.70	94.26	94.59
$V_{sb} = 10k$ (15M)	92.74	93.85	94.37	94.64
$V_{sb} = 20k$ (15M)	92.77	93.95	94.48	94.80
$V_{sb} = 30k$ (15M)	92.33	93.76	94.38	94.72
$V_{sb} = 10k$ (35M)	92.67	93.87	94.37	94.59
$V_{sb} = 20k$ (35M)	92.84	93.93	94.50	94.79
$V_{sb} = 30k$ (35M)	92.92	93.84	94.44	94.72
Hale et al. (2018)	87.1	88.96	90.48	90.96

Table 2: PTB development set parsing accuracy (F1) when changing beam size, averaged on three models with different random seeds. V_{fix} is the vocabulary size for fixed vocabulary models while V_{sb} is that for subword models. Hale et al. (2018) is trained only on PTB training set and is not directly comparable. Word beam $k_w = k/10$ and $k_s = k/100$.

V (# params.) / beam k	100	200	400	1000
$V_{fix} = 50k$ (15M)	52.34	49.53	48.26	47.53
$V_{sb} = 10k$ (15M)	69.09	65.34	63.74	62.81
$V_{sb} = 20k$ (15M)	67.52	64.15	62.41	61.36
$V_{sb} = 30k$ (15M)	70.43	66.41	64.50	63.41
$V_{sb} = 10k$ (35M)	67.66	64.13	62.41	61.47
$V_{sb} = 20k$ (35M)	63.60	60.33	58.78	57.90
$V_{sb} = 30k$ (35M)	60.80	57.91	56.28	55.45

Table 3: Perplexity on BLLIP validation set for each setting described in Table 2, averaged on three models with different seeds. For V_{sb} , perplexity is not subword-level but token-level, by summing subword likelihoods for each token.

We instead validate the effectiveness of our subword modeling by not language modeling, but parsing performance. Note that the text in the BLLIP corpus is Wall Street Journal, the same as Penn Treebank (PTB). Thus, we expect that the quality of auto parses provided to our training data is high, allowing us to assume that a good model should parse the gold PTB data more accurately. We run beam search on the PTB development (section 22) for each model and the results are summarized in Table 2. We can see that F1 scores consistently improve by subword modeling compared to the fixed vocabulary setting. The effects of model size (15M vs. 35M) are negligible, suggesting that the upper bound parsing performance using the current silver quality data can be reached with smaller models.

For (2) above, while comparing perplexities across subword and fixed-vocabulary models is impossible, comparing different subword units is possible by casting the subword-level likelihoods

to token-level likelihoods (Mielke, 2019). Table 3 summarizes those values along with the results by fixed vocabulary models as reference. The effects of the number of subword units (V_{sb}) are clearer. For 15M models, the optimal V_{sb} is $20k$, while for larger 35M models, the optimal size is $30k$. This suggests that more parameters are needed to obtain better results for large models.