

Explicitly Modeling Syntax in Language Models with Incremental Parsing and a Dynamic Oracle

Yikang Shen

Mila/Université de Montréal
yi-kang.shen@umontreal.ca

Shawn Tan

Mila/Université de Montréal
tanjings@mila.quebec

Alessandro Sordoni

Microsoft Research Montréal

Siva Reddy

Mila/McGill University

Aaron Courville

Mila/Université de Montréal

Abstract

Syntax is fundamental to our thinking about language. Failing to capture the structure of input language could lead to generalization problems and over-parametrization. In the present work, we propose a new syntax-aware language model: Syntactic Ordered Memory (SOM). The model explicitly models the structure with an incremental parser and maintains the conditional probability setting of a standard language model (left-to-right). To train the incremental parser and avoid exposure bias, we also propose a novel dynamic oracle, so that SOM is more robust to wrong parsing decisions. Experiments show that SOM can achieve strong results in language modeling, incremental parsing and syntactic generalization tests, while using fewer parameters than other models.

1 Introduction

Several recent works have systematically studied the linguistic abilities of modern language models, particularly syntax (Linzen et al., 2016; Marvin and Linzen, 2018; Gulordava et al., 2018). They find that most language models are good at capturing frequent syntactic structures but do not generalize well to those in the long tail. Moreover, although some excel at having low perplexity scores, this is less due to their syntactic ability but more due to capturing collocations (frequently co-occurring words). Recently, Hu et al. (2020) show that RNNs underperform on a syntactic generalization (SG) test set, whereas models that have an explicit notion of syntax, such as RNNG (Dyer et al., 2016), fare well on SG but at the cost of generally poorer language modeling (higher perplexity). Transformer-based models achieve strong performance when trained with large datasets, but are worse than random when trained on a small dataset.

These works showed that building language models with an explicit internal model of syntax




Figure 1: The mechanism of SOM. “Context” is a distributed representation of previous sentences. It could also represent the source sentence in a sequence to sequence task. It incrementally build subtrees given the input sentences. A RNN will takes the context representation and the representations of subtrees in the current sentence to predict next token.

helps in achieving better performance in SG tasks and is also thought to help learn more efficiently in low data settings. However, building syntax-aware models that also obtain strong language modeling performance, when compared with recent transformer-based models, has until now seemed elusive. In this work, we propose a new syntax-aware language model dubbed Syntactic Ordered Memory (SOM; Fig. 1), which jointly acts as a language model and an *incremental* parser. SOM inherits the syntax representation used in Ordered Memory (OM; Shen et al. 2019) in which syntax trees are embedded in a grid-like memory representation. Whereas OM was trained as an unsupervised parser, SOM is explicitly trained to predict both ground-truth syntax trees incrementally and, using the predicted partial syntactic structure, to predict the next token. Fig.1 shows the mechanism of SOM.

SOM factorizes the next-token prediction process into two steps: first, we predict the attachment position for the next token with a zero-step look-ahead parser, trained in a supervised fashion; then, we predict the next token distribution conditioned on the partially predicted structure. One way of training the incremental parser is to use teacher-forcing. However, this can lead to *exposure bias*, due to the fact that the model was never exposed to its own predictions during training. To avoid

this, we introduce a dynamic oracle (Goldberg and Nivre, 2012) for our model, so that our model can learn to recover from previous parsing mistakes during inference. We found this to be crucial to obtain good performance.

We compare SOM with existing methods that integrate syntax into language models. RNNs (Dyer et al., 2016) and Ordered Neurons (Shen et al., 2018) are particularly related. RNNs are generative models of language which define a joint distribution on syntactic structures and sequence of words. Ordered Neurons attempt to model the hierarchical structure of language by defining an ordering to the hidden states and the gates that impose that structure. We show that our proposed SOM model can achieve strong language modeling, parsing and SG performance even when trained on small amounts of data.

In summary, our contributions are threefold:

- We introduce SOM, a new syntax-augmented language model that learns an incremental parser and use its predictions to improve language modeling.
- We propose a novel dynamic oracle that allows to reduce the exposure bias and is instrumental to achieving good downstream performance.
- We report high SG score, language modeling and incremental parsing performance for various dataset sizes. We also find that jointly learning both language modelling and parsing improves both these capabilities in the model.

2 Related Work

Syntax-aware models There has been work to integrate syntax into our current models of language. Socher et al. (2013) used parse trees for composing sentences in order to predict sentiment over movie reviews. However, having an external parser and restriction of batched computations in that early model made the method unwieldy. Bowman et al. (2016) introduced the SPINN model, which alleviated those issues, turning sentences into a sequence of actions to be executed by a shift-reduce parser. Our SOM model is based on shift-reduce as well, because of the incremental nature of the parsing we want to achieve. RNN (Dyer et al., 2016; Kuncoro et al., 2016) was an example of integrating syntax information for language modelling.

There is also work that attempts to learn these syntactic structures without supervision. Kim et al. (2019) later devised an unsupervised version of the RNN, a method which produced good parsing performance. DIORA (Drozdov et al., 2019, 2020) was a method that leveraged the Inside-Outside algorithm to construct sentence embeddings for downstream tasks, with the benefit of being able to read off parse trees in the encoding process.

Swayamdipta et al. (2019) finds that there are no improvements over using ELMo (Peters et al., 2018) embeddings when shallow syntactic information is included, concluding that ELMo-style pretraining has learned the syntactic information. However, Kuncoro et al. (2019) investigated the importance of the learnt syntactic knowledge RNN in a large pre-trained model like BERT, they found that syntax information helps with downstream tasks. In our experiments, we find that explicitly training OM with syntax (with our dynamic oracle scheme) improves performance on syntactic generalization tasks.

Incremental Parsing & Language Modelling

In SOM, we specifically focus on incremental parsing. Ghezzi and Mandrioli (1979) discusses incremental parsing in the context of programming languages, with shift-reduce parsers being a specific type of incremental parsing. OM, RNN, and SPINN are models that were designed with shift-reduce in mind.

Incremental parsing lends itself well to the task of autoregressive language modelling. Since the parser only sees the prefix of a sentence, the model can use the partial parse to make a prediction about upcoming words. Demberg et al. (2013) summarises several empirical results that provide evidence for incremental and predictive parsing in humans, and makes several connections between *incrementality* (that comprehenders do not wait to the end of the sentence before building a representation) and *prediction* about future words coming in the sentence.

Given that an incremental parser processes a sentence from left to right, there are naturally some limitations. Hassan et al. (2009) show why either a beam or delay is necessary if performing incremental parsing with monotonic extensions: They experiment with a parser based on Combinatory Categorical Grammar (Steedman, 2000). They find that without the look-ahead, there is a 30 % point reduction in the parsing results. One of our con-

tributions in this paper is the one-step lookahead while performing parsing, but zero-step lookahead when performing next-word prediction, allowing the model to be trained jointly as an incremental parser and language model.

Despite the left-to-right nature of incremental parsing, this setting may aid language modelling too. Shieber (1983) suggests the biases may correspond to the way humans parse English, and use a modified shift-reduce parser to disambiguate between different parses of a sentence. There has been work that shows that incremental parsing can improve language modelling. Köhn and Baumann (2016) demonstrate that combining an incremental dependency parser with a language model yields improvements in perplexity. Roark (2001) presents a top-down phrase structure parser that performs beam-search to generate connected intermediate structures for every sentence prefix. This model can be used for language modeling and beats trigram models on the Penn Treebank (Marcus et al., 1994)

Dynamic Oracles Since incremental parsing requires that we break down the problem of structure prediction into sequential decisions, we are prone to *exposure bias*. There are techniques to address this by allowing the model to make mistakes and supervising future actions based on the state arrived at (Daumé et al., 2009). Goldberg and Nivre (2012) introduces the concept of dynamic oracles for dependency parsing. Coavoux and Crabbé (2016) uses this technique for incremental constituency parsing, but uses morphological features, and does not perform language modelling. Fried and Klein (2018) cover in further detail the related work relating to dynamic oracles and parsing. We find that using dynamic oracles for training is crucial in seeing benefits in both language modelling and incremental parsing.

Evaluating Syntactic Generalization Recent tests have been developed that attempt to probe the linguistic abilities of language models. Gulordava et al. (2018) explores the extent to which RNNs are able to model grammar, independent of the semantics of the sentence. Marvin and Linzen (2018) evaluate language models on their ability to score sentences with and without the proper subject-verb agreements over a variety of different settings.

Hu et al. (2020) expands on these ideas, and propose a suite of syntactic generalization tests

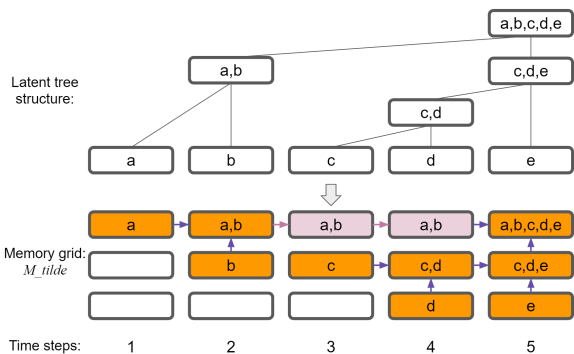


Figure 2: The grid view of a tree structure. Blue arrows represent composing children into parent. Gray arrows represent copying from previous time step. Orange slots are memories generated at the current time step. Gray slots are memories copied from previous time step.


for language models over a series of different sized datasets. They find that while GPT-2 performs well, their performance is highly dependent on the scale of the language modeling training dataset, while other models remain more robust. In this paper, we use this test suite for the evaluation.

3 Ordered Memory

We first provide useful background on Ordered Memory. Ordered Memory (OM, Shen et al. 2019) is a recurrent neural network that explicitly models recursive structure through memory writing and erasing operations. OM maps the latent syntax into a $T \times N$ memory grid \tilde{M} , where T is the length of input sequence and N is the maximum number of memory slots. Figure 2 gives an intuition of what the grid contains. Empty blocks in the figure represent memory slots that can be discarded during inference. Ideally, the memory network should generate the t -th column of the grid \tilde{M}_t at time step t . But generating \tilde{M}_t requires the model to have access about the tree structure which is usually latent. For this reason, OM induces the latent structure through inductive biases of its reading and writing operations.

As a recurrent model, OM performs one-step look-ahead incremental parsing through maintaining three states:

- **Memory M_t** : a matrix of dimension $N \times D$, where each occupied slot is a distributed representation for a node spanning an subsequence in x_1, \dots, x_{t-1} conditioned on x_t , i.e. M_t represents a one-step look-ahead parser stack. It's represented by gray blocks in Figure 3.



(a) The transition from time step 4 to 5. ① The one-step look-ahead parser combines \hat{M}_{t-1} and M_{t-1} considering on the current input x_t , in this example, the split point of \hat{M}_{t-1} and M_{t-1} is $i = 2$. ② Current input x_t is written into the lower slot of new candidate memory \hat{M}_t^{i-1} . ③ The rest of new candidate memories $\hat{M}_t^{\geq i}$ are generated with bottom-up recurrent composition.

(b) Predicting the next token at time step 4. ① The zero-step look-ahead parser combines M_t and \hat{M}_t at time step t . ② The recurrent network takes the combined memory M_t^{out} as input and output a hidden state $h_t = f(w_{\leq t})$. ③ h_t is then fed into a linear layer to compute $p(x_{t+1}|x_{\leq t})$.

Figure 3: The recurrent transition (left) and prediction network (right) of SOM. In the recurrent transition, a one-step look-ahead parser predict the syntax once e is observed and can be seen as a posterior over the syntax given the current word. The prediction network uses a zero-step look-ahead parser to predict the location of the next phrase and acts as a prior on the syntactic structure.

- **Candidate memory** \hat{M}_t : a matrix of dimension $N \times D$ contains representations for all possible new nodes at time step t . At next time step $t + 1$, the model will decide whether or not to write these candidates into memory M_{t+1} conditioned on x_{t+1} . They are represented by orange blocks in Figure 3. if the model is making correct parsing decisions, then $M_t = \hat{M}_{t-1}$.
- **Memory mask** $\vec{\pi}_t$: $\vec{\pi}_t \in \{0, 1\}^N$, where each entry indicates whether the respective slot in \hat{M}_t is occupied by a candidate, e.g., if $\vec{\pi}_t = (0, 1, 1)$, then the occupied slots are $\hat{M}_t^{\geq 2}$. At next time step, the model can only choose a candidate from masked slots to write into the memory M_{t+1} .

At each time step, the model takes $[M_{t-1}, \hat{M}_{t-1}, \vec{\pi}_{t-1}]$ and word embedding x_t as inputs, returning the outputs $[M_t, \hat{M}_t, \vec{\pi}_t]$.

To generate the new memory M_t , we combine M_{t-1} and \hat{M}_{t-1} to match \hat{M}_{t-1} . The model uses x_t as its query to attend on previous candidates \hat{M}_{t-1} . The attention distribution is p_t , which models the split point of gray blocks and orange blocks in Figure 2. Suppose p_t is a one-hot distribution and $p_t^i = 1$. The candidates $\hat{M}_{t-1}^{\leq i}$ are written into the respective memory slot $M_t^{\leq i}$, while $M_{t-1}^{\geq i}$ are copied to $M_t^{\geq i}$:

$$M_t^{\leq i} = \hat{M}_{t-1}^{\leq i}, \quad M_t^{\geq i} = M_{t-1}^{\geq i} \quad (1)$$

We will refer to the process of generating M_t as a one-step look-ahead parser, since the model is using the current input x_t as extra information to build the partial parse for time step $t - 1$. To generate new candidates \hat{M}_t , the input embedding x_t is written into \hat{M}_t^{i-1} , and $\hat{M}_t^{\geq i}$ are computed recurrently with eq.3:

$$\hat{M}_t^{\leq i-1} = \emptyset, \quad \hat{M}_t^{i-1} = x_t \quad (2)$$

$$\hat{M}_t^j = \text{cell}(M_t^j, \hat{M}_t^{j-1}), \quad \forall j \geq i \quad (3)$$

where $\text{cell}()$ is the composition function that takes its childrens' representations as input and output the parent's representation. The non-empty slots in candidate memory are then $\hat{M}_t^{\geq i-1}$, and they can be masked by:

$$\vec{\pi}_t^{\leq i-1} = 0, \quad \vec{\pi}_t^{\geq i-1} = 1 \quad (4)$$

In other words, $\vec{\pi}_t^i = \sum_{j \leq i+1} p_t^j$, and $\vec{\pi}_t^i$ is monotonically increasing. More details of the OM can be found in Shen et al. (2019).

4 Syntactic Ordered Memory

We propose two augmentations to OM in order to better perform language modelling and incremental parsing: a prediction network and the dynamic oracle. a) Previous language models mostly focus on predicting the next token or a missing token. In our case, we are explicitly modeling the latent structure. By predicting the structure for the next token, we exploit this latent structure for word prediction.

This helps the model better organize information for predicting next word, allowing shortcuts to be created for long-term dependencies, as shown in Fig.1. b) If the model only observes states resulting from correct past decisions at training time, it will not be prepared to recover from its own mistakes during prediction, suffering from exposure bias (Schmidt, 2019; Fried and Klein, 2018). In the experiment section, we demonstrate how this phenomenon will significantly hurt the language model performance and, to a lesser extent, also hurt the parsing performance.

4.1 Prediction Network

At time step t , the prediction network takes $[M_t, \hat{M}_t, \vec{\pi}_t]$ as input, and produces a probability distribution over the next token $p(w_{t+1}|w_{\leq t})$. To do this, we need to have a temporary estimate of the local structure. We therefore need to approximate p_{t+1} with a zero-step look-ahead prediction p'_t :

$$\alpha_t^i = \frac{\mathbf{w}_2^{Att} \text{ReLU}(\mathbf{W}_1^{Att} \hat{M}_t^i + b_1) + b_2}{\sqrt{N}} \quad (5)$$

$$p'_t = \text{masked_softmax}(\alpha_t, \text{mask} = \vec{\pi}_t) \quad (6)$$

where \mathbf{W}_1^{Att} is $N \times N$ weight matrix, \mathbf{w}_2^{Att} is a N dimension weight vector, and α_t^i is a scalar. We then sample the slot at index i from the distribution p'_t . i is the zero-step look-ahead parsing decision, which means that the next phrase will be a sibling of node \hat{M}_t^i . We therefore need to predict the next token conditioned on \hat{M}_t^i and its previous contexts. So we feed memory slots $[M_t^N, M_t^{N-1}, \dots, M_t^{i+1}, \hat{M}_t^i]$ into a recurrent neural network:

$$h_t = \text{RNN}(M_t^N, M_t^{N-1}, \dots, M_t^{i+1}, \hat{M}_t^i) \quad (7)$$

where h_t is the final hidden state of the RNN. As shown in Figure 3b, the input sequence are representations of non-overlapping subtrees spanning from x_1 to x_t . h_t can therefore be seen as a distributed representation of the sequence $w_{\leq t}$. In the RNN, we use the same architecture as the cell function in OM to model the recurrent transition

function:

$$\begin{bmatrix} f_j \\ i_j \\ c_j \\ u_j \end{bmatrix} = \mathbf{W}_2^{Cell} \text{ReLU} \left(\mathbf{W}_1^{Cell} \begin{bmatrix} h_t^{j+1} \\ M_j \end{bmatrix} + b_1 \right) + b_2 \quad (8)$$

$$h_t^j = \text{LN}(\sigma(f_j) \odot h_t^{j+1} + \sigma(i_j) \odot M_j + \sigma(c_j) \odot u_j) \quad (9)$$

where σ is the sigmoid function, LN is layer normalization function, f_j, i_j, c_j are controlling gates, c_j is cell state, and h_t^{N+1} is a zero vector. After obtaining h_t , we can compute the distribution over the next token and the language modelling loss:

$$p(w_{t+1}|w_{\leq t}) = \text{softmax}(\mathbf{W}_{emb} h_t + b) \quad (10)$$

$$L_{\text{LM}} = - \sum_t \log(p(w_{t+1}|w_{\leq t})) \quad (11)$$

4.2 Dynamic Oracle for SOM

Data: $\theta_1, \dots, \theta_T, \Gamma$

Result: ξ_1, \dots, ξ_T

initialize $\xi_1 = N$;

for $i \leftarrow 2$ **to** T **do**

$j = \text{first_sibling}_{\Gamma}(i)$;

$\mu_i = \max(\theta_{j+1}, \dots, \theta_{i-1})$;

$\xi_i = \max(\xi_j - 1, \mu_i)$;

end

Algorithm 1: The structure label generation algorithm, where Γ is the ground-truth tree and θ_i is the structural decisions made by our model. This algorithm produces a parse close to the original given the errors already made, and that new gold parse is converted into grid decisions. Given Γ , the function $\text{first_sibling}_{\Gamma}(i)$ returns the index of the first token in the smallest clause that contains w_i , and where w_i is not the first token. Ideally, w_i should be written into the slot $(\xi_j - 1)$. For example, in Figure 2, c is written into the slot 2, then d, e should be written into the slot 1. However, the model could make a wrong decision between w_j and w_i . If the model has merged information from w_j into a higher slot μ_i , x_i should be written into slot μ_i as well.

One way to provide a supervision signal for p_t and p'_t is to train the parser with static oracle: feed the gold tree to the model, and have the model predict future decisions. However, static oracle makes the language model overfit on the gold tree, resulting in bad perplexity scores (Table 2). Inspired

Type	Max	Median	Mean
Constituency	29	7	7.7
Dependency	16	4	4.2

Table 1: Statistics of tree depth for Penn Treebank. Dependency trees are converted from constituency tree with Stanford CoreNlp toolkit.




Figure 4: The universal dependency tree is converted into a constituency tree Γ through merging the head and its children into one single constituent. Since the grid view only works with binary trees, we binarize n-ary nodes with a left branching bias.

by the dynamic oracles proposed in (Goldberg and Nivre, 2012; Coavoux and Crabbé, 2016), we propose a dynamic oracle for ordered memory, which dynamically changes the reference structure based on mistakes made by our model on previous steps. To do this, we build the structure label for each time step based on the gold tree and previous decisions made by the model. During training, we sample the model’s decision from p_t :

$$\theta_t = \text{Multinomial}(p_t) \quad (12)$$

and we make greedy decisions during evaluation:

$$\theta_t = \text{argmax}(p_t) \quad (13)$$

The same operations are applied to p'_t as well.

We use the Algorithm.1 to convert the gold tree Γ into labels ξ_t for p_t . Since the zero-step look-ahead distribution p'_t should match the one-step look-ahead distribution p_{t+1} at next time step $t+1$, we use ξ_{t+1} as label for p'_t . The structure loss is the negative log-likelihood:

$$L_S = - \sum_t (\log(p_t(\xi_t | w_{\leq t})) + \log(p'_t(\xi_{t+1} | w_{\leq t})))$$

For our model, the depth of Γ has a linear relation to the computational complexity and GPU memory consumption. To maximize the model’s efficiency, the gold tree Γ is constructed from *universal dependency trees*.¹ There are two reasons

we chose universal dependency trees instead of constituency trees: 1) In Table 1, the dependency trees are on average shallower than constituency trees; this means faster computation time and less memory consumption for our model. 2) Universal dependency trees can be applied to many more languages than Penn Treebank-style constituency grammar. Additionally, Penn Treebank-style trees can easily be converted to universal dependency trees. As shown in Figure 4, we convert the universal dependency tree into Γ by merging the head and its children into one single constituent.

5 Experiments

We present the results of SOM on language modeling, syntactic generalization, and incremental parsing. Details of hyperparameters and experiment settings can be found in Appendix B.

5.1 Language Modeling

Penn Treebank has one million words of 1989 Wall Street Journal corpus annotated with constituency trees. Since SOM primarily focuses on sentence-level structure and language modeling, we use the same preprocessing schema as RNNG² (Dyer et al., 2016). Sentences are modeled separately, punctuation is retained, and singleton words are replaced with the Berkeley parser’s mapping rules³, resulting in 23,815-word types. Orthographic case distinction is preserved, and numbers (beyond singletons) are not normalized.

BLLIP is a large Penn Treebank-style parsed corpus of approximately 24 million sentences. We train and evaluate SOM on three splits of BLLIP: BLLIP-XS (40k sentences, 1M tokens), BLLIP-SM (200K sentences, 5M tokens), and BLLIP-MD (600K sentences, 14M tokens). They are obtained by randomly sampling sections from BLLIP 1987-89 Corpus Release 1. All models are tested on a shared held-out tested set.

Following the settings provided in (Hu et al., 2020), datasets are preprocessed into two different versions. The first setting is similar to the PTB dataset. Singleton words are mapped to UNK classes that preserve fine-grained information, such as orthographic case distinctions and morphological suffixes (e.g. UNK-ed, UNK-ly). The second setting use subword-level vocabulary extracted

²2-21 for training, 24 for validation, 23 for evaluation.

³<http://github.com/slavpetrov/berkeleyparser>

¹<https://universaldependencies.org/>

Model	# parameters	ppl	p acc	UF1	p' acc
SOM	17.7M	77.68	0.927	87.96	0.870
SOM – Prediction network	13.0M	83.63	0.923	87.09	–
SOM – Prediction network – Language Modeling Loss	13.0M	–	0.925	86.26	–
SOM – Dynamic Oracle + Static Oracle	17.7M	129.27	0.913	86.58	0.849
SOM – Dynamic Oracle + Left-branching Oracle	17.7M	82.01	–	–	–
<i>Inference with External Trees</i>					
SOM – Predicted tree + Gold tree	17.7M	60.87	0.947	100.00	0.884

Table 2: Ablation tests on the PTB dataset. “ p acc” and “ p' acc” are the prediction accuracies of the one-step look-ahead and zero-step look-ahead parsers respectively. “UF1” is the parsing performance with respect to the converted constituency tree Γ . “– Prediction network”: this model uses the last candidate memory slot \hat{M}_t^N to predict the next token, instead of using the h_t from the prediction network. “– Predicted tree + Gold tree”: the model’s parsing decisions were replaced with ground truth decisions; these results can be considered as the performance upper bound of SOM.

Model	PTB
<i>Without annotations</i>	
RNNLM	93.2
PRPN (Shen et al., 2017)	96.7
URNNG (Kim et al., 2019)	90.6
<i>With annotations</i>	
RNNG (Dyer et al., 2016)	88.7
RNNG \rightarrow URNNG (Kim et al., 2019)	85.9
SOM	77.7

Table 3: Perplexities on Penn Treebank datasets. *With annotations* are models that use the gold tree as supervision signal during training. Baseline results are from Kim et al. (2019)

Model	XS	SM	MD
n-gram	240.21	157.60	106.09
RNNG	122.46	86.72	69.57
LSTM	98.19	65.52	59.05
ON-LSTM	71.76	54.00	56.37
GPT-2	529.90*	183.10*	37.04*
SOM	70.41	51.47	31.95*

Table 4: Perplexities on BLLIP datasets achieved by different models. Perplexity scores across training dataset sizes are not strictly comparable for models that use word-level vocabulary. * results are using GPT-2’s subword vocabulary.

from the GPT-2 pretrained model rather than the BLLIP training corpora.

Results of language modeling are given in Table 3 and Table 4. SOM consistently outperforms both the annotated model and non-annotated models. While GPT-2 seems to fail to learn on smaller datasets, SOM still outperforms GPT-2 on the BLLIP-MD dataset with far fewer parameters (34.8M vs 124.4M), and achieves comparable results with the GPT-2 that is trained on a 3 times larger dataset BLLIP-LG (Hu et al., 2020).

The ablation test results are shown in Table 2.

The biggest performance drop comes from replacing the dynamic oracle with static oracle. We believe that this is due to the model overfitting on the gold tree, and suffering from exposure bias as a result. Another big performance drop happens after removing the prediction network. This suggests that predicting the attaching nodes of the next phrase with the zero-step look-ahead parsers helps to predict the next token. Replacing the gold tree labels with trivial left-branching tree labels also hurts the perplexity. This suggests that learning syntactic structure helps language modeling.

5.2 Syntactic Generalization

Syntactic Generalization (SG) test suites evaluate the syntactic knowledge of neural language models. Hu et al. (2020) proposed a set of 34 test suites to evaluation 6 different aspects of syntax: 1) agreement, 2) licensing, 3) garden-path effects, 4) gross syntactic expectation, 5) center embedding, 6) long-distance dependencies.

Following their settings, we evaluate our language models trained on the BLLIP datasets. Language models are presented with a group of sentences with minor differences. To pass each test, the model needs to assign higher conditional probabilities to designated phrases in the sentence that are more grammatical.

Figure 6 shows the average accuracy over all model on the complete set of SG test suites. SOM achieves the best average accuracy, outperforms models with hierarchical structure bias (RNNG, ON-LSTM), and transformer-based model (GPT-2). However, according to Figure 8a in Appendix C.1, GPT-2 trained on BLLIP-LG and BLLIP-MD still outperform SOM. This could due to that the number of parameters in SOM is largely falling behind GPT-2.




Figure 5: Evaluation results on all models, split across test suite circuits.




Figure 6: Average SG accuracy by model class.

Figure 5 provides fine-grained results on six SG classes. SOM achieves strong performance on licensing, gross syntactic state, center embedding, and long-distance embeddings. These classes require the model to keep track of syntactic features across large syntactic chunks (e.g., relative or subordination clauses). SOM can effectively keep this long-term information in higher-level memory slots, and revisit the information after the clause in the middle is ended. More detailed results can be found in Appendix C.1.

5.3 Incremental Parsing

Model	UF1
PRPN*	41.2
ONLSTM*	47.7
ONLSTM-SYD (Du et al., 2020)	61.3
Incremental Shift-reduce Parser	56.82
Shift-reduce + LM + Dynamic Oracle	58.04
SOM	67.27
Oracle Binary Trees	82.5

Table 5: Incremental parsing results on the standard PTB constituency trees. “*” means that the model is doing unsupervised grammar induction. Since we compare UF1 against the standard, nonbinarized trees (per convention), UF1 scores is upper bounded by the oracle binary trees score.

To evaluate SOM’s performance on incremental parsing, we trained and evaluated our models on the standard PTB constituency trees. Baseline models include: a) a standard incremental shift-reduce parser with one-step look-ahead; b) a incremental shift-reduce parser that equipped with our predic-

tion network and trained on same dynamic oracle and language model loss as our model; c) a recently proposed ONLSTM-SYD model (Du et al., 2020) that is also trained on both language model and parsing loss; d) unsupervised ONLSTM; e) unsupervised PRPN. As shown in Table 5, SOMs outperform all baseline models, including the shift-reduce parser that has the same extra components as SOMs. For language modelling performance, original constituency tree based models achieve similar perplexity as dependency tree based counterparts. But constituency tree based models require $2\times$ GPU time and memory to train and evaluate.

For ablation test, we also compare parsing results given by SOM with binary constituency trees Γ converted from universal dependency trees.⁴ These results are shown in Table 2. We observe that using static oracle instead of dynamic oracle results in the worst parsing performance. This suggests that our dynamic oracle helps the model to learn a better parser. After removing the language model loss, the UF1 drops 1.7 points. This suggests that the language model loss helps the model to learn better representations for syntax.

6 Conclusion

In this work, we propose a new language model with an integrated incremental parser. This was done by augmenting the Ordered Memory model with a prediction network, and by using a dynamic oracle for training it to perform incremental parsing. The resulting model models the joint distribution of syntactic structure and sequence words. We find that by using the dynamic oracle and explicitly modeling the syntax, we can achieve strong performance on language modelling and syntactic generalization and both these techniques are crucial in the model’s performance.

⁴UF1 scores are computed by EVALB <https://nlp.cs.nyu.edu/evalb/>

References

- 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. *arXiv preprint arXiv:1603.06021*.
- Maximin Coavoux and Benoit Crabbé. 2016. Neural greedy constituent parsing with dynamic oracles.
- Hal Daumé, John Langford, and Daniel Marcu. 2009. Search-based structured prediction. *Machine learning*, 75(3):297–325.
- Vera Demberg, Frank Keller, and Alexander Koller. 2013. Incremental, predictive parsing with psycholinguistically motivated tree-adjointing grammar. *Computational Linguistics*, 39(4):1025–1066.
- Andrew Drozdov, Subendhu Rongali, Yi-Pei Chen, Tim O’Gorman, Mohit Iyyer, and Andrew McCallum. 2020. Unsupervised parsing with s-diora: Single tree encoding for deep inside-outside recursive autoencoders. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4832–4845.
- Andrew Drozdov, Pat Verga, Mohit Yadav, Mohit Iyyer, and Andrew McCallum. 2019. Unsupervised latent tree induction with deep inside-outside recursive autoencoders. *arXiv preprint arXiv:1904.02142*.
- Wenyu Du, Zhouhan Lin, Yikang Shen, Timothy J O’Donnell, Yoshua Bengio, and Yue Zhang. 2020. Exploiting syntactic structure for better language modeling: A syntactic distance approach. *arXiv preprint arXiv:2005.05864*.
- Chris Dyer, Adhiguna Kuncoro, Miguel Ballesteros, and Noah A Smith. 2016. Recurrent neural network grammars. *arXiv preprint arXiv:1602.07776*.
- Daniel Fried and Dan Klein. 2018. Policy gradient as a proxy for dynamic oracles in constituency parsing. *arXiv preprint arXiv:1806.03290*.
- Carlo Ghezzi and Dino Mandrioli. 1979. Incremental parsing. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 1(1):58–70.
- Yoav Goldberg and Joakim Nivre. 2012. A dynamic oracle for arc-eager dependency parsing. In *Proceedings of COLING 2012*, pages 959–976.
- Kristina Gulordava, Piotr Bojanowski, Edouard Grave, Tal Linzen, and Marco Baroni. 2018. Colorless green recurrent networks dream hierarchically. In *Proceedings of NAACL-HLT*, pages 1195–1205.
- Hany Hassan, Khalil Sima’an, and Andy Way. 2009. Lexicalized semi-incremental dependency parsing. In *Proceedings of the International Conference RANLP-2009*, pages 128–134.
- Jennifer Hu, Jon Gauthier, Peng Qian, Ethan Wilcox, and Roger P Levy. 2020. A systematic assessment of syntactic generalization in neural language models. *arXiv preprint arXiv:2005.03692*.
- Yoon Kim, Alexander M Rush, Lei Yu, Adhiguna Kuncoro, Chris Dyer, and Gábor Melis. 2019. Unsupervised recurrent neural network grammars. *arXiv preprint arXiv:1904.03746*.
- Arne Köhn and Timo Baumann. 2016. Predictive incremental parsing helps language modeling.
- Adhiguna Kuncoro, Miguel Ballesteros, Lingpeng Kong, Chris Dyer, Graham Neubig, and Noah A Smith. 2016. What do recurrent neural network grammars learn about syntax? *arXiv preprint arXiv:1611.05774*.
- Adhiguna Kuncoro, Chris Dyer, Laura Rimell, Stephen Clark, and Phil Blunsom. 2019. Scalable syntax-aware language models using knowledge distillation. *arXiv preprint arXiv:1906.06438*.
- 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 Marcus, Grace Kim, Mary Ann Marcinkiewicz, Robert MacIntyre, Ann Bies, Mark Ferguson, Karen Katz, and Britta Schasberger. 1994. The penn treebank: annotating predicate argument structure. In *Proceedings of the workshop on Human Language Technology*, pages 114–119. Association for Computational Linguistics.
- Rebecca Marvin and Tal Linzen. 2018. Targeted syntactic evaluation of language models. *arXiv preprint arXiv:1808.09031*.
- Matthew E Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. *arXiv preprint arXiv:1802.05365*.
- Brian Roark. 2001. Probabilistic top-down parsing and language modeling. *Computational linguistics*, 27(2):249–276.
- Florian Schmidt. 2019. Generalization in generation: A closer look at exposure bias. *arXiv preprint arXiv:1910.00292*.
- Yikang Shen, Zhouhan Lin, Chin-Wei Huang, and Aaron Courville. 2017. Neural language modeling by jointly learning syntax and lexicon. *arXiv preprint arXiv:1711.02013*.
- Yikang Shen, Shawn Tan, Arian Hosseini, Zhouhan Lin, Alessandro Sordani, and Aaron C Courville. 2019. Ordered memory. In *Advances in Neural Information Processing Systems*, pages 5038–5049.

- Yikang Shen, Shawn Tan, Alessandro Sordoni, and Aaron Courville. 2018. Ordered neurons: Integrating tree structures into recurrent neural networks. *arXiv preprint arXiv:1810.09536*.
- Stuart M Shieber. 1983. Sentence disambiguation by a shift-reduce parsing technique. In *Proceedings of the 21st annual meeting on Association for Computational Linguistics*, pages 113–118. Association for Computational Linguistics.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Y Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 conference on empirical methods in natural language processing*, pages 1631–1642.
- Mark Steedman. 2000. *The syntactic process*, volume 24. MIT press Cambridge, MA.
- Swabha Swayamdipta, Matthew Peters, Brendan Roof, Chris Dyer, and Noah A Smith. 2019. Shallow syntax in deep water. *arXiv preprint arXiv:1908.11047*.

A Disentangling Semantic and Syntactic representations

Given the architecture of our model, we can easily disentangle the language model information flow and parsing information flow. Figure 7 illustrates the disentangled information and gradient flows in our model. The language model depends on both prior context and structural inputs, and derivatives are computed with respect to both of these inputs and backpropagated. However, while the structure also depends on both inputs, we limit backpropagation so that it can only update with respect to the syntactic input. This is because we want the parsing component to function independently of the language modelling component, but still leverage the semantic information to deal with syntactic ambiguity.




Figure 7: The schema of disentangling syntax from the language model. Solid lines represent dependency during inference, and gradients flow back during backpropagation. The dashed line represents the dependency during inference, but detached so that the gradients do not flow back during backpropagation.

It is possible that existing model architectures could implicitly learn to split these representations, even without the explicit disentanglement that we proposed here. Yet, Table 2 shows that entangled model can actually achieve stronger in-domain performance, thanks to the liberty to allocate capacity to the two different functionalities based on the training set.

To do so, we propose splitting word embeddings, memory slots, and intermediate hidden states into two segments: semantic segment and syntactic segment. We then replace linear layers in our cell functions with the following function:

$$\begin{bmatrix} y_{sem} \\ y_{syn} \end{bmatrix} = \begin{bmatrix} \mathbf{W}_{sem2sem} & \mathbf{W}_{syn2sem} \\ 0 & \mathbf{W}_{syn2syn} \end{bmatrix} \begin{bmatrix} x_{sem} \\ x_{syn} \end{bmatrix}$$

where y_{sem} and x_{sem} are the semantic segment which is optimized to minimize language modeling loss, y_{syn} and x_{syn} are the syntactic segment which is optimized to minimize both parsing and language modeling loss. This architecture results in the solid lines in Figure 7. Additionally, layer normalization functions are replaced with two separate functions for the two segments respectively.

Meanwhile, p_t still depends on both semantic and syntactic segment, but the structural loss does not backpropagate into the semantic segment:

$$p_t = f(x_{t,sem}, x_{t,syn}, \hat{M}_{t,sem}, \hat{M}_{t,syn}) \quad (14)$$

$$\frac{\partial p_t}{\partial x_{t,sem}} = 0, \quad \frac{\partial p_t}{\partial \hat{M}_{t,sem}} = 0 \quad (15)$$

and the same for p'_t :

$$p'_t = f(\hat{M}_{t,sem}, \hat{M}_{t,syn}) \quad (16)$$

$$\frac{\partial p'_t}{\partial \hat{M}_{t,sem}} = 0 \quad (17)$$

This gradient detachment is represented by the dash line in Figure 7. In the experiment section, the disentangled models are denoted as dSOM, and entangled models are denoted as SOM. For dSOM, the dimension of semantic and syntactic segments for memory slots are denoted as D_{sem} and D_{syn} respectively. Among the proposed models, the eSOM has the best performance on the in-domain test sets. Appendix C.2 shows that the dSOM slightly outperforms eSOM in perplexity on out-of-domain test sets.

B Hyperparameters

Model	XS	SM	MD
RNG	22.8M	48.4M	81.1M
LSTM	13.4M	30.5M	52.2M
ONLSTM+AWD	30.8M	44.2M	61.2M
GPT-2	124.4M	124.4M	124.4M
dSOM	16.4M	39.5M	34.8M
eSOM	17.8M	41.4M	37.9M

Table 6: Parameter counts for different models

Dropout is applied before all linear layers in our model. They all share the same dropout rate, except the dropout before language model output layer has a different rate. We also applied embedding dropout which randomly set some embedding vectors to 0. Hyperparameters are chosen based on the perplexity on validation set.

C More Experiment Results

C.1 Syntactic Generalization results

C.2 Out of Domain Evaluation

Out-of-domain Test set contains testsets from other English universal dependencies treebanks. It contains corpora of different genres, including academic, email, blog, fiction, legal, news, etc. We use these datasets to test the generalization ability of models that are trained on PTB.

Dataset	D_{sem}	D_{syn}	#slots	embedding dropout	dropout	output dropout
PTB	300	100	15	0.1	0.3	0.5
BLLIP-XS	300	100	15	0.1	0.3	0.5
BLLIP-SM	400	100	15	0.1	0.2	0.2
BLLIP-MD-BPE	400	100	15	0	0.1	0.1

Table 7: Hyperparameters. The hidden size of eSOM models are always the sum of D_{sem} and D_{syn}

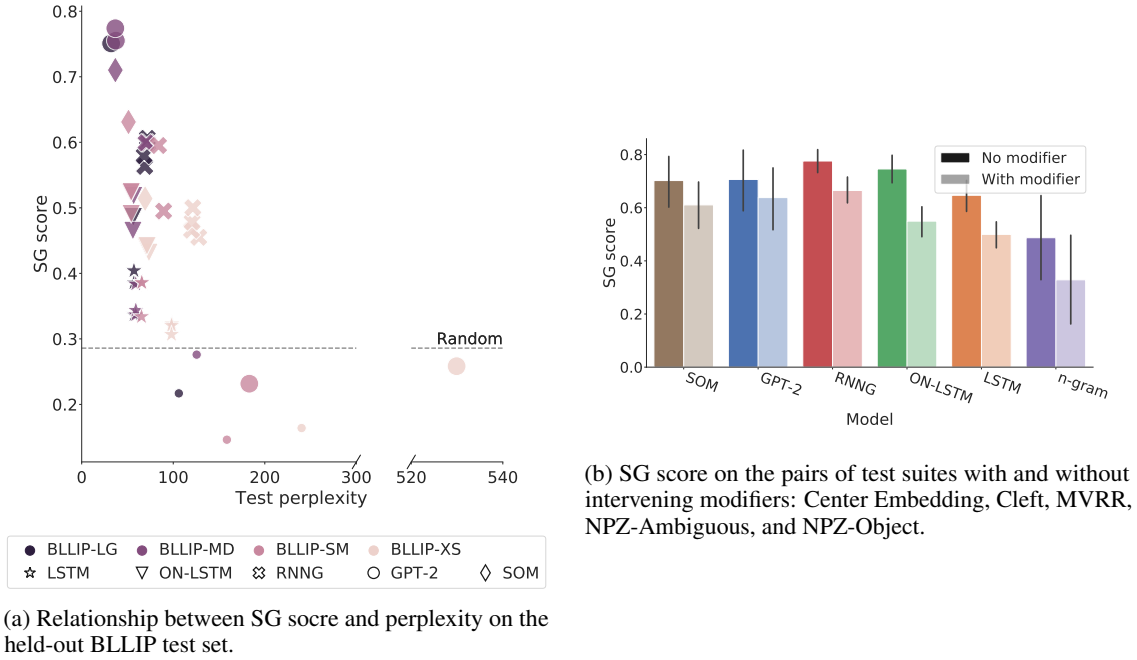


Figure 8

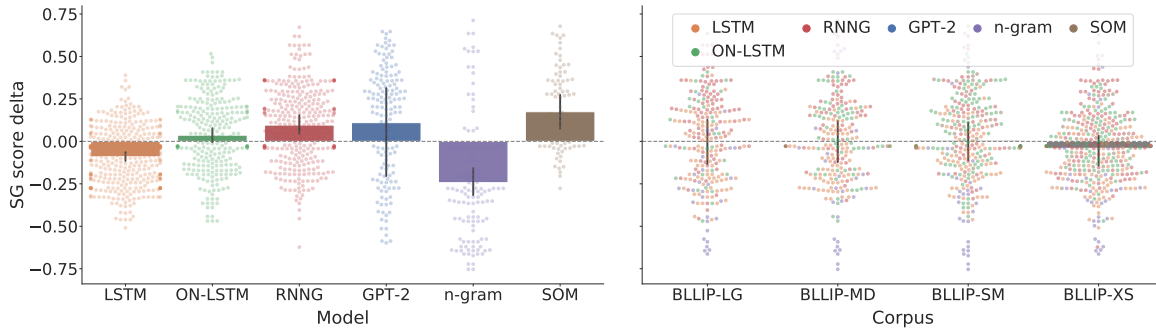


Figure 9: Left: Model class has a strong effect on SG scores. Right: Data scale has little effect on SG scores

Dataset	GUM		EWT		ParTUT		LinES		Pronouns		PUD		SG acc
	ppl	UF1	ppl	UF1	ppl	UF1	ppl	UF1	ppl	UF1	ppl	UF1	
eSOM	351.5	67.0	400.7	73.1	282.6	76.5	253.3	67.2	552.1	87.1	281.0	75.6	0.614
dSOM	350.3	66.4	403.0	72.3	269.8	74.3	252.1	66.6	565.8	87.3	280.3	75.1	0.581
LB	375.5	–	436.5	–	300.9	–	267.5	–	620.4	–	300.7	–	0.513

Table 8: Out of domain test results. Models are trained on PTB. The test sets are obtained from English universal dependencies treebank. “LB” stands for left-branching tree labels. Thanks to the structure information, our models generalize much better than the left-branching baseline.




Figure 10: Left: differences in model class induce significant differences in SG scores for several circuits. Right: differences in training data size do not reliably account for most of circuits.