
Improving Semantic Parsing for Task Oriented Dialog

Arash Einolghozati*
Facebook Conversational AI
arashe@fb.com

Panupong Pasupat*
Stanford University
ppasupat@cs.stanford.edu

Sonal Gupta
Facebook Conversational AI
sonalgupta@fb.com

Rushin Shah
Facebook Conversational AI
rushinshah@fb.com

Mrinal Mohit
Facebook Conversational AI
mrinalmohit@fb.com

Mike Lewis
Facebook AI Research
mikelewis@fb.com

Luke Zettlemoyer
Facebook AI Research
lsz@fb.com

Abstract

Semantic parsing using hierarchical representations has recently been proposed for task oriented dialog with promising results. In this paper, we present three different improvements to the model: contextualized embeddings, ensembling, and pairwise re-ranking based on a language model. We taxonomize the errors possible for the hierarchical representation, such as wrong top intent, missing spans or split spans, and show that the three approaches correct different kinds of errors. The best model combines the three techniques and gives 6.4% better exact match accuracy than the state-of-the-art, with an error reduction of 33%, resulting in a new state-of-the-art result on the Task Oriented Parsing (TOP) dataset.

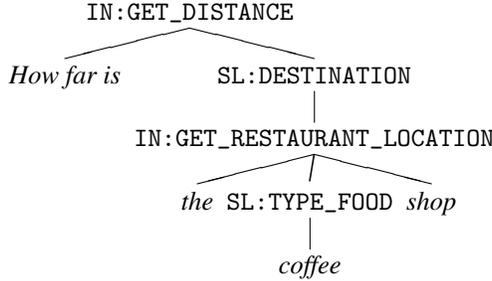
1 Introduction

Although intelligent personal assistants such as Alexa and Siri are now ubiquitous, modeling the semantics of compositional natural language queries remains challenging. The most common practice for natural language understanding in task oriented dialogs is to apply a slot-filling system [17, 16]. Given a simple query such as “*How far is San Francisco*”, a slot-filling system would classify the *intent* of the whole query (GET_DISTANCE) and tag the relevant *slots* (*San Francisco* being a DESTINATION) with a sequence tagging model. However, understanding more complex queries such as “*How far is the coffee shop*”, which requires locating the coffee shop (GET_RESTAURANT_LOCATION intent) before finding the distance (GET_DISTANCE intent), is not straightforward in traditional systems, since the compositionality is not captured when the task is posed as text classification or sequence tagging.

Recently, a Task Oriented Parsing (TOP) representation for intent-slot based dialog systems was introduced [9]. The representation, as illustrated in Figure 1, is expressive enough to capture the task-specific semantics of complex nested queries, but is easier to annotate and parse than full semantic representations such as logical forms [24] or abstract meaning representation [1].

A key advantage of the TOP representation is that its structure is similar to standard constituency parses, allowing us to adapt improvements in modeling techniques for syntactic parsers to this problem. In this paper, we propose three approaches to improve the parsing model:

*equal contribution



Action	Stack	Buffer
		<i>How far is the coffee shop</i>
IN:GD	IN:GD	<i>How far is the coffee shop</i>
SHIFT	IN:GD <i>How</i>	<i>far is the coffee shop</i>
SHIFT	IN:GD <i>How far</i>	<i>is the coffee shop</i>
SHIFT	IN:GD <i>How far is</i>	<i>the coffee shop</i>
SL:DEST	IN:GD <i>How far is</i> SL:DEST	<i>the coffee shop</i>
IN:GRL	IN:GD <i>How far is</i> SL:DEST IN:GRL	<i>the coffee shop</i>
SHIFT	IN:GD <i>How far is</i> SL:DEST IN:GRL <i>the</i>	<i>coffee shop</i>
SL:TF	IN:GD <i>How far is</i> SL:DEST IN:GRL <i>the</i> SL:TF	<i>coffee shop</i>
SHIFT	IN:GD <i>How far is</i> SL:DEST IN:GRL <i>the</i> SL:TF <i>coffee</i>	<i>shop</i>
REDUCE	IN:GD <i>How far is</i> SL:DEST IN:GRL <i>the</i> [SL:TF <i>coffee</i>]	<i>shop</i>
SHIFT	IN:GD <i>How far is</i> SL:DEST IN:GRL <i>the</i> [SL:TF <i>coffee</i>] <i>shop</i>	
REDUCE	IN:GD <i>How far is</i> SL:DEST [IN:GRL <i>the</i> [SL:TF <i>coffee</i>] <i>shop</i>]	
REDUCE	IN:GD <i>How far is</i> [SL:DEST [IN:GRL <i>the</i> [SL:TF <i>coffee</i>] <i>shop</i>]]	
REDUCE	[IN:GD <i>How far is</i> [SL:DEST [IN:GRL <i>the</i> [SL:TF <i>coffee</i>] <i>shop</i>]]]	

Figure 1: An example TOP annotation. Intents are prefixed with IN: and slots with SL:. In a traditional intent-slot system, the SL:DESTINATION could not have an intent nested inside it. The table below shows the parser actions used to generate the parse (with abbreviated labels). At each step, the candidate actions are scored based on the stack LSTM embeddings of three sequences: the previous actions, the current stack, and the current buffer.

- Ensembling with three strategies: majority voting, greedy action, and parser switch.
- Incorporating deep contextualized word embeddings, such as ELMo [18].
- Re-ranking the parses based on a language model (LM).

To compare the effectiveness of the different approaches, we also propose an error classification scheme for the TOP representation. Our analysis shows that the improvements from the three techniques are orthogonal. This allows us to construct the best model using an ensemble of seven models, of which three are trained with LM-based re-ranking.

2 Base Model and Data

Following [9], we use a shift-reduce parser based on Recurrent Neural Network Grammars (RNNG) [8] as our base model. The model constructs a parse tree by predicting a sequence of actions. The set of actions include SHIFT, REDUCE, and the generation of intent and slot labels. The SHIFT action consumes an input token, adding it as a child of most recent ‘open’ sub-tree node, while the REDUCE action closes a sub-tree. The final set of actions is to generate a non-terminal (an intent or a slot) as a new empty sub-tree node. Note that at each step, only a subset of actions will be valid. For examples, if all the input tokens have been added to the tree, the only valid action is REDUCE.

In an RNNG model, the action are scored based on the embeddings of three sequences: the list of actions performed so far, the partial construction of the tree (“stack”), and the list of remaining tokens (“buffer”). Each sequence is embedded using a stack LSTM, which dynamically recomputes the embedding as items are added to or removed from the sequence.

Table 1: Gains from ensembling.

Ensembling Strategy	Accuracy
Base	80.86
Oracle	90.81
Majority Vote	83.76
Greedy Action	83.84
Parser Switch	82.82

Table 2: Gains from ELMo embeddings.

ELMo layer	Accuracy
First	83.93
Last	83.08
Average	83.85
Learned	83.76
oncatenated	83.70

For training, we use the same hyperparameters as in [9]: 2-layer LSTMs of size 164, dropout rate of 0.34, and pre-trained word embeddings of size 200. We train with 16 workers using Hogwild [20] for 1 epoch. For optimization, we use Adam [15] with learning rate of 0.0004 and weight decay of 0.00004. The base model uses greedy decoding for inference at test time.

The open-sourced data released by [9] contains utterances of navigation and event domains. We remove the utterances where the top intent is IN:UNSUPPORTED as it is a noisy catch-all class for out-of-domain utterances. The final dataset contains 28,276 training, 4,014 evaluation, and 8,191 test utterances. Our evaluation metric is the exact match accuracy: the number of utterances whose full parse trees are correctly predicted.

3 Improvements

3.1 Ensembling

Ensembling is a powerful method to improve the performance of dependency and constituency parsers [12, 22]. In order to take advantage of parser diversity (e.g., sentence shuffling, dropout seeds, parameter initialization), we use different strategies to combine the output of individual parsers.

A parse tree has a one-to-one correspondence with the sequence of actions performed by the parser. In the approaches below, we propose and compare different ways of combining the list of actions from several base models. To ensure that the resulting parse is well-formed, we design the methods so that the output parse always matches one of the parses produced by ensembling parsers.

Majority Vote: We simply select the list which has been predicted by the majority of the individual parsers as the output of the ensemble.

Greedy Action: Starting from left (the top-level intent) to right, we pick the action that is predicted by the majority of the parsers, and then discard parsers that do not agree with the majority. Since the action lists can have different lengths, we pad the shorter lists with dummy actions which are eventually removed.

Parser Switch: We implement an approach similar to the one proposed in [12]. From each parse i , we first obtain the *constituent set* S_i , where each constituent is a triple (label, start index, end index) extracted from a sub-tree node. For example, the constituent corresponding to the top intent of the parse in Figure 1 is (IN:GET_DISTANCE, 0, 6). The score for each parser i is then computed as $\sum_j |S_i \cap S_j|$; i.e., the sum of intersection with the constituent sets from other parses. Finally, we choose the parser with the highest score. In contrast to the previous two methods, parser switch considers the whole tree to find the majority consensus, but it is computationally more expensive as it needs to construct trees from all parsers to obtain the constituent sets.

We also report the oracle ensemble parser which always picks the correct parser whenever possible in order to determine the upper bound for ensembling gains. For the rest of the paper, we use seven parsers for the ensemble model. The results in Table 1 show that the greedy action and majority vote strategies are almost equally good and better than the parser switch. On the other hand, all the methods fall short of the oracle parser which suggests that there might still be room for improving the ensemble parser. The effectiveness of the simple majority vote confirms the observation in [22], which showed that simple voting performs well for ensembling dependency parsers.

Table 3: Gains from LM-based re-ranking with beam size 5.

Method	Accuracy
Base (Greedy)	80.86
Oracle top-2 beam	89.99
Oracle top-5 beam	93.20
Top-1 beam	81.21
Naive LM ranker	82.80
SVM ranker	84.26

3.2 Contextual word embeddings

Following the substantial gains provided by contextual embeddings for multiple NLP tasks [18], we replace the regular pre-trained word embeddings with off-the-shelf pre-trained ELMo embeddings. It has 2 bi-LSTM layers with highway connections on top of a character n-gram convolutional layer, which after projection result in three vectors of size 1024 for each word.

The results for different ELMo strategies are shown in Table 2: using the first layer, last layer, average of 3 layers, concatenating the layers, and learned weighted average of the layers. The first and last layer weights result in the biggest and smallest gains, respectively. A justification for the first layer’s better performance on our task can be that the parser needs to identify syntactic properties (e.g., opening the non-terminal right after prepositions) to produce correct parses.

3.3 Language Model Re-ranking

The experiment in [9] showed that the top-5 accuracy of the model (i.e., the fraction of utterances where the correct parse is in any of the output from a beam of 5) is much higher than the top-1 accuracy. We train the base model as described in the previous section, but use beam search of size 5 instead of greedy decoding at inference. This improves the accuracy of the top-1 and top-5 hypotheses as shown in Table 3. This suggests that re-ranking the top hypotheses can be an effective way of increasing the accuracy.

Similar to [6], we propose to score the parses for re-ranking with a language model over the sequential serialization of the parses. Our hypothesis is that the generative nature of language model scoring could mitigate some of errors arising from the greedy decoding in RNNG. We serialize the parse tree as follows: the opening bracket and the non-terminal are considered as one token, and the closing bracket is mapped to the corresponding non-terminal and considered a different token. For example, [IN:GET_EVENT [SL:CAT_EVENT *Concerts*] by [SL:NAME_EVENT *Kendrick Lamar*]] would become "O_IN_GET_EVENT O_SL_CATEGORY_EVENT *Concerts* C_SL_CATEGORY_EVENT by O_SL_NAME_EVENT *Kendrick Lamar* C_SL_NAME_EVENT C_IN_GET_EVENT" (O stands for open and C stands for close).

We serialize the gold trees in the training data and train a neural language model on them. We explore two approaches to re-ranking using the trained language model:

- Use the LM score directly to rank the parses. We found that re-ranking only the top two hypotheses gives the best results.
- Train a SVM-based ranker [13] that takes the beam score and LM score as features.

At the end, the top decoded beams are scored and re-ranked based on one of the two strategies described above. The last two rows in Table 3 show that the LM re-ranking using the SVM ranker results in a substantial gain over the top-1 beam but there is still a huge difference compared to the oracle top-5 beam decoding. This suggests that more effective re-ranking approaches, such as including more features in the ranker, could be worth trying.

4 Error Analysis

In this section, we categorize the types of errors the parsing model makes in our task and analyze which ones can be mitigated via the methods described in the previous section. We classify the errors into seven major groups. Note that a single query can exhibit multiple types of errors. (In the list below, – and + denote the expected and predicted frames below, respectively.)

1. **Wrong Top intent (WT):** The first action taken is wrong.
 - [IN:GET_INFO_TRAFFIC *how many miles is* [SL:LOCATION [IN:GET_LOCATION [SL:CATEGORY_LOCATION *the interstate*]]] *backed up*]
 - + [IN:GET_DISTANCE *how many* [SL:UNIT_DISTANCE *miles*] *is* [SL:DESTINATION [IN:GET_LOCATION [SL:CATEGORY_LOCATION *the interstate*]]] *backed up*]
2. **Wrong Label (WL):** A predicted label for the sub-intents or slots is wrong.
 - [IN:GET_INFO_TRAFFIC *Should I avoid* [SL:PATH_AVOID *I - 26*] [SL:DATE_TIME *today*]]
 - + [IN:GET_INFO_TRAFFIC *Should I avoid* [SL:LOCATION *I - 26*] [SL:DATE_TIME *today*]]
3. **Spurious span (SS):** A Constituent is wrongly added.
 - [IN:GET_ESTIMATED_DEPARTURE *Do I need to leave earlier* [SL:DATE_TIME_DEPARTURE *today*] *due to traffic*]
 - + [IN:GET_ESTIMATED_DEPARTURE *Do I need to leave* [SL:SOURCE *earlier*] [SL:DATE_TIME_DEPARTURE *today*] *due to traffic*]
4. **Missing Spans (MS):** A Constituent is missing in the prediction.
 - [IN:GET_INFO_TRAFFIC *What is traffic like in* [SL:LOCATION *still water*] [SL:DATE_TIME *at the moment*]]
 - + [IN:GET_INFO_TRAFFIC *What is traffic like in* *still water* [SL:DATE_TIME *at the moment*]]
5. **Wrong Split (WS):** Constituents are wrongly split.
 - [IN:GET_EVENT *When is* [SL:CATEGORY_EVENT *Christmas in the Park*] [SL:DATE_TIME *this year*] *in* [SL:LOCATION *San Antonio TX*]]
 - + [IN:GET_EVENT *When is* [SL:DATE_TIME *Christmas*] *in* [SL:LOCATION [IN:GET_LOCATION [SL:CATEGORY_LOCATION *the Park*]]] [SL:DATE_TIME *this year*] *in* [SL:LOCATION *San Antonio TX*]]
6. **Wrong join (WJ):** Constituents are wrongly joined.
 - [IN:GET_EVENT [SL:LOCATION *Dayton*] [SL:CATEGORY_EVENT *parties*] [SL:DATE_TIME *for NYE*]]
 - + [IN:GET_EVENT [SL:CATEGORY_EVENT *Dayton parties*] [SL:DATE_TIME *for NYE*]]
7. **Bad Boundary (BB):** A constituent has wrong boundaries.
 - [IN:GET_EVENT *Whats* [SL:DATE_TIME *tomorrows*] *events for* [SL:LOCATION *Houston*]]
 - + [IN:GET_EVENT *Whats* [SL:CATEGORY_EVENT *tomorrows events*] *for* [SL:LOCATION *Houston*]]

The absolute number of errors for the base model and the relative change for each of the methods in the previous section are summarized in Table 4. We use the majority vote ensemble and the average ELMo setting for the experiments. We can see that ELMo is effective across all types of errors, while ensemble and LM ranking mitigate most types of errors. Ensembling is much more effective at top intent classification (a semantic task) whereas ELMo is the only method effective against missing span and wrong join (mostly syntactic tasks). This confirms the hypothesis that the RNN model uses the syntactic information in the first layer of ELMo more than the other layers. The above analysis suggests that combining the aforementioned methods can be effective in mitigating different types of errors together, which we explore in the next section.

Table 4: Effectiveness of methods on different types of errors (as percentage of error reduction).

Method	WT	WL	MS	SS	WS	WJ	BB
Base	222	850	246	333	110	88	330
ELMo	-9%	-20%	-11%	-25%	-22%	-28%	-21%
Ensemble	-25%	-21%	+2%	-32%	-29%	+3%	-18%
LM re-rank	-20%	-20%	+3%	-36%	-47%	-8%	-15%

Table 5: Combining ensembling and ELMo.

Ensemble Strategy	Accuracy
Oracle	92.43
Majority Vote	86.19
Greedy Action	85.95
Parser Switch	84.62

Table 6: LM re-ranking with ELMo.

Method	Accuracy
Oracle top-2 beam	91.53
Oracle top-5 beam	95.04
Base ELMo (Greedy)	83.85
Top-1 beam	84.05
LM re-ranked	86.30

5 Combining the Methods

Based on the analysis in the previous section, we combine the methods to seek further gains. We first try the ELMo + Ensemble combination. We use the average ELMo strategy and experiment with the three ensemble strategies. The results are shown in Table 5. The combined gain confirms that the ensemble and ELMo work almost orthogonally and the combination results in about 28% error reduction compared to the base parser.

We also experiment with adding the LM SVM re-ranking to the ELMo-enabled model. The results are shown in Table 6. We observe that combined with ELMo, the LM re-ranking of the top beams still keeps most of its gains compared to the case without ELMo.

Finally, we report the results of combining LM re-ranking and ensemble (with and without using ELMo). Here, we employ two strategies. In the first strategy, using the ranking SVM strategy as before, we re-rank the top five hypotheses for each parser inside the ensemble and then apply the ensemble strategy. In the second strategy, we use the number of times each hypothesis appears in the top-5 beam of the parsers as an additional feature to the ranking SVM. We use the average ELMo strategy alongside with the majority vote ensemble for these experiments.

Table 7 compares the results to the baseline using the top beam. We can see that LM re-ranking can be effective on top of ensemble and ELMo, and the biggest gains are achieved by using the voting decision inside the ranking SVM.

6 Related Work

Our work builds on top of two related but distinct directions of research. At one end, there has been a large literature on language understanding for task oriented dialog, such as the work that tackle the ATIS and DSTC datasets [17, 16]. Most work in this area assumes that the utterance is not

Table 7: The results from combining all techniques.

Method	Accuracy
Base Ensemble w/o ELMo	83.83
LM re-ranked Ensemble w/o ELMo	84.86
Extended SVM Ranking w/o ELMo	85.22
Base Ensemble with ELMo	86.26
LM re-ranked Ensemble with ELMo	86.67
Extended SVM Ranking with ELMo	87.25

compositional. The current state-of-the-art [25] frames the problem as one of non-recursive intent and slot tagging, and assumes that the NLU output is passed along to a dialog manager in order to be executed. There has also been work on end-to-end task oriented dialog [3], but there too, the problem is usually framed as one of selecting a single API call and its arguments, as opposed to compositional API calls.

At the other end of the spectrum, in the traditional semantic parsing literature, the problem is framed as predicting compositional semantic representations. These are mainly geared towards question-answering rather than task completion, and are usually directly executed against a knowledge base. Some of the standard datasets in this area include GeoQuery [24] and WebQuestions [2].

Within both of these areas, neural approaches have supplanted previous feature-engineering based approaches in recent years [10, 14]. In the context of tree-structured semantic parsing, some other interesting approaches include Seq2Tree [7] which modifies the standard Seq2Seq decoder to better output trees; SCANNER [5, 4] which extends the RNNG formulation specifically for semantic parsing such that the output is no longer coupled with the input; and TRANX [23] and Abstract Syntax Network [19] which generate code along a programming language schema. For graph-structured semantic parsing [1, 11], SLING [21] produces graph-structured parses by modeling semantic parsing as a neural transition parsing problem with a more expressive transition tag set. While graph structures can provide more detailed semantics, they are more difficult to parse and can be an overkill for understanding task oriented utterances.

7 Conclusion

In this paper, we propose three different techniques to improve the hierarchical representation based semantic parsing models using ensembling, contextualized word embeddings, and language model based re-ranking. We propose a categorization of errors for the TOP representation. Our results show that the three approaches improve the model on different types of errors and the best model uses a combination of them. Our best model reduces the error rate by 33% on the TOP dataset.

References

- [1] L. Banarescu, C. Bonial, S. Cai, M. Georgescu, K. Griffitt, U. Hermjakob, K. Knight, P. Koehn, M. Palmer, and N. Schneider. Abstract meaning representation for sembanking. In *7th Linguistic Annotation Workshop and Interoperability with Discourse*, 2013.
- [2] J. Berant, A. Chou, R. Frostig, and P. Liang. Semantic parsing on freebase from question-answer pairs. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1533–1544, 2013.
- [3] A. Bordes, Y.-L. Boureau, and J. Weston. Learning end-to-end goal-oriented dialog. In *Proc. of ICLR*, 2016.
- [4] J. Cheng, S. Reddy, V. Saraswat, and M. Lapata. Learning structured natural language representations for semantic parsing. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 44–55. Association for Computational Linguistics, 2017. doi: 10.18653/v1/P17-1005. URL <http://www.aclweb.org/anthology/P17-1005>.
- [5] J. Cheng, S. Reddy, V. Saraswat, and M. Lapata. Learning an executable neural semantic parser. *CoRR*, abs/1711.05066, 2017. URL <http://arxiv.org/abs/1711.05066>.
- [6] M. Collins and T. Koo. Discriminative reranking for natural language parsing. *Comput. Linguist.*, 31(1):25–70, Mar. 2005. ISSN 0891-2017. doi: 10.1162/0891201053630273. URL <http://dx.doi.org/10.1162/0891201053630273>.
- [7] L. Dong and M. Lapata. Language to logical form with neural attention. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 33–43, Berlin, Germany, August 2016. Association for Computational Linguistics. URL <http://www.aclweb.org/anthology/P16-1004>.

- [8] C. Dyer, A. Kuncoro, M. Ballesteros, and N. A. Smith. Recurrent neural network grammars. In *Proc. of NAACL*, 2016.
- [9] S. Gupta, R. Shah, M. Mohit, A. Kumar, and M. Lewis. Semantic parsing for task oriented dialog using hierarchical representations. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2018.
- [10] D. Hakkani-Tur, G. Tur, A. Celikyilmaz, Y.-N. Chen, J. Gao, L. Deng, and Y.-Y. Wang. Multi-domain joint semantic frame parsing using bi-directional RNN-LSTM. In *Interspeech 2016*, pages 715–719, 2016. doi: 10.21437/Interspeech.2016-402. URL <http://dx.doi.org/10.21437/Interspeech.2016-402>.
- [11] L. He, K. Lee, M. Lewis, and L. Zettlemoyer. Deep semantic role labeling: What works and what’s next. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, 2017.
- [12] J. C. Henderson and E. Brill. Exploiting diversity in natural language processing: Combining parsers. *CoRR*, cs.CL/0006003, 2000. URL <http://arxiv.org/abs/cs.CL/0006003>.
- [13] R. Herbrich, T. Graepel, and K. Obermayer. Support vector learning for ordinal regression. In *1999 Ninth International Conference on Artificial Neural Networks ICANN 99. (Conf. Publ. No. 470)*, volume 1, pages 97–102 vol.1, Sept 1999. doi: 10.1049/cp:19991091.
- [14] S. Iyer, I. Konstas, A. Cheung, J. Krishnamurthy, and L. Zettlemoyer. Learning a neural semantic parser from user feedback. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers*, pages 963–973, 2017. doi: 10.18653/v1/P17-1089. URL <https://doi.org/10.18653/v1/P17-1089>.
- [15] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *International Conference for Learning Representations*, 2015.
- [16] B. Liu and I. Lane. Attention-based recurrent neural network models for joint intent detection and slot filling. In *INTERSPEECH*, 2016.
- [17] G. Mesnil, X. He, L. Deng, and Y. Bengio. Investigation of recurrent-neural-network architectures and learning methods for spoken language understanding. In *Interspeech*, 2013.
- [18] M. E. Peters, M. Neumann, M. Iyyer, M. Gardner, C. Clark, K. Lee, and L. Zettlemoyer. Deep contextualized word representations. In *Proc. of NAACL*, 2018.
- [19] M. Rabinovich, M. Stern, and D. Klein. Abstract syntax networks for code generation and semantic parsing. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2017.
- [20] B. Recht, C. Re, S. Wright, and F. Niu. Hogwild: A lock-free approach to parallelizing stochastic gradient descent. In *Advances in Neural Information Processing Systems (NIPS)*, 2011.
- [21] M. Ringgaard, R. Gupta, and F. C. N. Pereira. SLING: A framework for frame semantic parsing. *CoRR*, abs/1710.07032, 2017. URL <http://arxiv.org/abs/1710.07032>.
- [22] M. Surdeanu and C. D. Manning. Ensemble models for dependency parsing: Cheap and good? In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics, HLT ’10*, pages 649–652, Stroudsburg, PA, USA, 2010. Association for Computational Linguistics.
- [23] P. Yin and G. Neubig. Tranx: A transition-based neural abstract syntax parser for semantic parsing and code generation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 7–12. Association for Computational Linguistics, 2018.
- [24] J. M. Zelle and R. J. Mooney. Learning to parse database queries using inductive logic programming. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence (AAAI)*, 1996.

- [25] S. Zhu and K. Yu. Encoder-decoder with focus-mechanism for sequence labelling based spoken language understanding. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5675–5679, March 2017. doi: 10.1109/ICASSP.2017.7953243.