

# Conversational Semantic Parsing

Armen Aghajanyan Jean Maillard Akshat Shrivastava Keith Diedrick  
Mike Haeger Haoran Li Yashar Mehdad Ves Stoyanov  
Anuj Kumar Mike Lewis Sonal Gupta

Facebook

{armenag, jeanm, akshats, kdiedrick, mhaeger, aimeeli}@fb.com  
{mehdad, ves, anujk, mikelewis, sonalgupta}@fb.com

## Abstract

The structured representation for semantic parsing in task-oriented assistant systems is geared towards simple understanding of one-turn queries. Due to the limitations of the representation, the session-based properties such as co-reference resolution and context carry-over are processed downstream in a pipelined system. In this paper, we propose a semantic representation for such task-oriented conversational systems that can represent concepts such as co-reference and context carryover, enabling comprehensive understanding of queries in a session. We release a new session-based, compositional task-oriented parsing dataset of 20k sessions consisting of 60k utterances. Unlike Dialog State Tracking Challenges, the queries in the dataset have compositional forms. We propose a new family of Seq2Seq models for the session-based parsing above, which achieve better or comparable performance to the current state-of-the-art on ATIS, SNIPS, TOP and DSTC2. Notably, we improve the best known results on DSTC2 by up to 5 points for slot-carryover.

## 1 Introduction

At the core of conversational assistants lies the semantic representation, which provides a structured description of tasks supported by the assistant. Traditional dialog systems operate through a flat representation, usually composed of a single intent and a list of slots with non-overlapping content from the utterance

(Bapna et al., 2017; Gupta et al., 2018). Although flat representations are trivial to model with standard intent/slot tagging models, the semantic representation is fundamentally limiting. Gupta et al. (2018) explored the limitations of flat representations and proposed a compositional generalization which allowed slots to contain nested intents while allowing easy modeling through neural shift-reduce parsers such as RNNNG (Dyer et al., 2016).

Our contributions are the following:

- We explore the limitations of this compositional form and propose an extension which overcomes these limitations that we call *decoupled representation*.
- To parse this more complicated representation, we propose a family of Seq2Seq models based off the Pointer-Generator architecture that set state of the art in multiple semantic parsing and dialog tasks (See et al., 2017).
- To further advance session based task oriented semantic parsing, we release a publicly available set with 60k utterances constituting roughly 20k sessions.

## 2 Semantic Representation

The compositional extension proposed by Gupta et al. (2018) overcame the limitation of classical intent-slot frameworks by allowing nested intents in slots. But to maintain an easily model-able structure the following constraint was introduced: *the in-order traversal*

of the compositional semantic representation must reconstruct the utterance. Following this constraint it is possible to use discriminative neural shift reduce parsers such as RNNG to parse into this form (Dyer et al., 2016).

Although at face value this constraint seems reasonable, it has non-trivial implications for both the semantic parsing component (NLU) and downstream components in conversational assistants.

## 2.1 Surpassing Utterance Level Limitations with Decoupled Form

First we'll take a look at the space of utterances that can be covered by the compositional representation. One fundamental problem with the in-order constraint is that it disallows long-distance dependencies within the semantic representation. For example, the utterance *On Monday, set an alarm for 8am.* would optimally have a single date-time slot: [SL\_DATE\_TIME 8am on Monday]. But, because *8am* and *on Monday* are at opposite ends of the utterance, there is no way to construct a semantic parse tree with a single date-time slot. Gupta et al. (2018) mentioned this problem, but had some empirical data showing that utterances with long-distance dependencies are rare in English. Although this might be true, having fundamental limitations on what type of utterances can be supported even with a complete ontology is concerning.

In English, discontinuities are restricted in occurrence, despite emerging naturally within certain patterns, because English is a configurational language, which uses strongly marked word order to impart some level of semantic information (Chomsky, 1981). Beyond English, however, there are numerous world languages that are non-configurational and have much freer or potentially completely free word order. Non-configurational languages may often present the same semantic information through the use of Case Markers, Declensions, or other systems. The relatively free word order this allows creates much less emphasis on the collocation of a semantic unit's to-

kens. Therefore, as conversational assistants progress toward multiple languages it's important to consider that constraints that are acceptable if only English is considered will not analogously scale to other languages.

A simple solution is to convert a standard compositional intent-slot parse into a logical form containing two label types (slot and intent), with no constraints over intent spans. This is trivially accomplished by removing all text in the compositional semantic parse that does not appear in a leaf slot. We call this form of semantic parse the *decoupled* semantic representation, due to the semantic representation not being tightly coupled with the original utterance.

Figure 1 shows a side by side example of compositional and decoupled semantic representations for the utterance *Please remind me to call John.*

## 2.2 Session Based Limitations

Because traditional conversational systems historically have had a clear separation between utterance level semantic parsing and dialog systems (which stitch together utterance level information into sessions), semantic representations have not focused on session-based representations. Integrating session information into semantic parsers has been limited to refinement-based approaches.

Figure 2 shows an example of refinement and informationally complete based approaches to semantic parsing. The refinement approach delegates responsibility of session-based semantic parsing to a separate dialog component. Consequently, refinement approaches tend to have a very limited ontology due to the semantic parser operating over a fixed input (non-session utterances).

Predicting what slot to use for refining works for flat semantic representations, but it is non-trivial to extend to compositional or decoupled. The position of a slot in a flat semantic representation is not meaningful, thus it is sufficient to only predict the slot without specifying its position in the parse. But




Figure 1: Compositional and decoupled semantic representations for the single utterance “Please remind me to call John”.

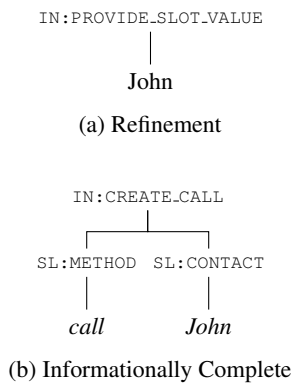


Figure 2: **Refinement** and **Complete** session based semantic representations for the utterance “call”.

both compositional and decoupled extensions to intent-slot parsing vary semantically by the position of the slot (or nested intent).

We present an example in Figure 3. Given the followup utterance *remind me to call*, a classical system would need to carry over the whole CONTACT slot, but the question is to where? The semantic parse is not flat. The slot could be carried over to the CREATE\_REMINDER intent or the nested GET\_CONTACT intent. So, if we were to extend classical slot carryover, we not only would need to predict what slot to carry over from the conversation, but what intent within the current semantic parse to place it under. We propose a new paradigm that does joint classical semantic parsing with co-reference resolution and slot-carryover.

### 2.3 Session Based Semantic Parsing

We present a simple extension to the decoupled paradigm of intent-slot semantic parsing by introduction of a new reference (REF) label type. The REF label type contains two elements in its set to represent co-references and slot-carryover as separate operations. Co-references can be seen as an explicit reference, namely a reference conditioned on an explicit word, while slot-carryover is treated as an implicit reference (conditioned by relevant contextual information).

As an example, refer to the sample session with decoupled semantic parses in Figure 4

## 3 Model

### 3.1 Sequence-to-Sequence Architecture

The decoupled semantic parsing model is an extension of the very common sequence-to-sequence learning approach (Sutskever et al., 2014), with the *source sequence* being the utterance and the *target sequence* being a linearized version of the target tree. Trees are linearized by bracketing them, using the same approach as Vinyals et al. (2015). The decoupled tree in Fig. 1b, for example, would be linearized to the following target sequence: [ IN:CREATE\_REMINDER, [ SL:PERSON\_REMINDED, me, ], [ SL:TODO, ..., ]. After tokenization, an encoder processes the source tokens  $w_i$  and produces corresponding encoder hidden states:

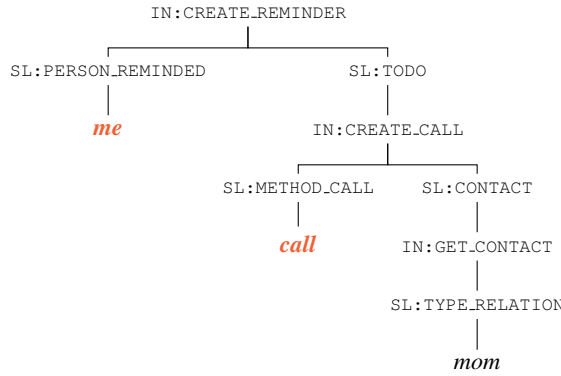


Figure 3: Sample session with complex slot-carryover: “Is mom available?” – “Remind me to call”

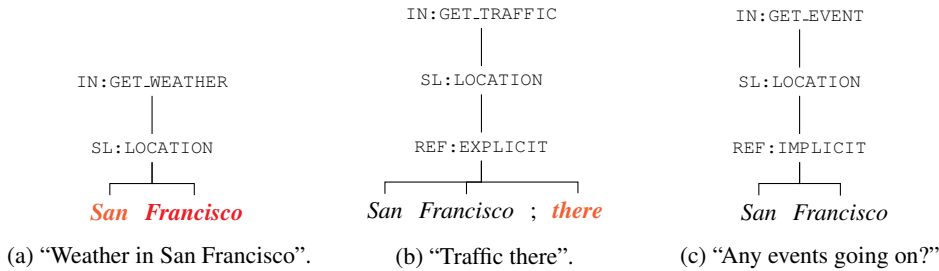


Figure 4: Decoupled semantic representations for a three-utterance session.

$$e_1, \dots, e_T = \text{Encoder}(w_1, \dots, w_T)$$

where the encoder, in our experiments, is either a standard bidirectional LSTM or a transformer.

In spite of its drawbacks, the rigid structure of the compositional semantic trees (Fig. 1a) has the advantage of readily mapping to the RNNG formalism and its inductive biases. The decoupled semantic representation, being more flexible, does not have such an easily exploitable form – but we can still exploit whatever structure exists. The tokens of the linearized decoupled representation (the *target sequence*) can always be divided into two classes: utterance tokens that are already present in the source sequence – which form the leaves of the tree – and ontology symbols. Taking again the example tree of Fig. 1b, *me*, *call*, and *John* are all tokens from the utterance, while [IN:CREATE\_REMINDER, [SL:PERSON\_REMINDED, ], etc., are ontology symbols. This partition is reflected in the

structure of the decoder: at every decoding step, the model can either generate an element from the ontology, or copy a token from the source sequence via a mechanism analogous to the pointer-generator network of See et al. (2017). At decoding time step  $t$ , the decoder is fed with the encoder’s outputs and produces a vector of features  $\mathbf{x}_t$ , which is used to compute an *ontology generation distribution*  $\mathbf{p}_t^g$ :

$$\begin{aligned} \mathbf{x}_t &= \text{Decoder}(e_1, \dots, e_t; \mathbf{d}_{t-1}; \mathbf{s}_{t-1}), \\ \mathbf{p}_t^g &= \text{softmax}(\text{Linear}_g[\mathbf{x}_t]), \end{aligned}$$

where  $\mathbf{d}_{t-1}$  is the previous output of the decoder,  $\mathbf{s}_{t-1}$  is the decoder’s incremental state, and  $\text{Linear}_\theta[\mathbf{x}]$  is short-hand for an affine transformation with parameters  $\theta$ , i.e.  $\mathbf{W}^\theta \mathbf{x} + \mathbf{b}^\theta$ . The decoder’s features are also used to calculate the attention distribution – using multi-head attention (Vaswani et al., 2017) – which then serves to produce the *utterance copy distribution*  $\mathbf{p}_t^c$ :

$$\begin{aligned} \mathbf{p}_t^c, \boldsymbol{\omega}_t &= \text{MhAttention}(e_1, \dots, e_t; \text{Linear}_c[\mathbf{x}_t]), \\ p_t^\alpha &= \sigma(\text{Linear}_\alpha[\mathbf{x}_t \parallel \boldsymbol{\omega}_t]), \end{aligned}$$

where  $\sigma(x) = \frac{1}{1+e^{-x}}$  is the standard sigmoid function,  $\parallel$  indicates concatenation, and MhAttention indicates multi-head attention which returns, respectively, the attention distribution and its weights. Finally, the extended probability distribution is computed as a mixture of the ontology generation and utterance copy distributions:

$$\mathbf{p}_t = p_t^\alpha \cdot \mathbf{p}_t^g + (1 - p_t^\alpha) \cdot \mathbf{p}_t^c.$$

### 3.2 Encoder and Decoder

We experiment with two main variants of the decoupled model: one based on recurrent neural networks, and one based on the transformer architecture (Vaswani et al., 2017).

**RNN** Our base model uses two distinct stacked bidirectional LSTMs as the encoder and stacked unidirectional LSTMs as the decoder. Both consist of two layers of size 512, with randomly initialized embeddings of size 300. The base model is optimized with LAMB while others are optimized with Adam, using parameters  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$ ,  $\epsilon = 10^{-8}$ , and L2 penalty  $10^{-5}$  (Kingma and Ba, 2014). The learning rate is found separately for each experiment via hyperparameter search. We also use stochastic weight averaging (Izmailov et al., 2018), and exponential learning rate decay. For an extended version of this model, we also try incorporating contextualized word vectors, by augmenting the input with ELMO embeddings (Peters et al., 2018).

**Transformer** We also experiment with two further variants of the model, that replace encoder and decoder with transformers. In the first variant, the encoder is initialized with RoBERTa (Liu et al., 2019), a pretrained language model. The decoder is a randomly initialized 3-layer transformer, with hidden size 512 and 4 attention heads. In the second variant, we initialise both encoder and decoder with BART (Lewis et al., 2019), a sequence-to-sequence pretrained model. Both encoder and decoder consist of 12 layers with hidden size 1024. We train these with stochastic weight averaging (Izmailov et al., 2018), and determine

optimal hyperparameters on the validation sets.

## 4 Experiments

### 4.1 Session Based Task Oriented Parsing

To incentivize further research into session based semantic parsing through the decoupled intent-slot paradigm we are releasing 20 thousand annotated sessions in 4 domains: calling, weather, music and reminder. We also allow for mixtures of domains within a session.

The data was collected in two stages. First we asked crowdsourced workers to write sessions (both from the users perspective as well as the Assistant’s output) tied to certain domains. Once we vetted the sessions, we asked a second group of annotators to annotate the user input per session. Each session was given to three separate annotators. We used majority voting to automatically resolve the correct parse when possible. In the cases where there was no agreement, we selected the maximum informative parse which abode by the labeling representations semantic constraints. The annotator agreement rate was 55%, while our final chosen semantic parses were correct 94% of the time. The large delta between the two numbers is due to multiple correct semantic parses existing for the same session.

We open source SB-TOP in the following link: <http://www.dl.fbaipublicfiles.com/sbtop/SBTOP.zip>. More information about the dataset can be found in the Table ?? in the Appendix.

### 4.2 Semantic Parsing


We evaluate the decoupled model on five semantic parsing datasets, four public and one internal. All but two are annotated with compositional semantic representations and the other with the standard flat intent-slot representation. In order to apply the decoupled models to them, we follow a mechanical procedure to transform the annotations to decoupled representations: all utterance tokens which are not part of a slot are stripped. This procedure effectively turns the tree of Fig. 1a into the tree

of Fig. 1b. We note that this procedure for all compositional and flat intent-slot data available is reversible, therefore we can convert from decoupled back to source representation.


The first public dataset is TOP (Gupta et al., 2018), which consists of over 31k training utterances covering the navigation, events, and navigation to events domains. The first internal dataset we use contains over 170k training utterances annotated with flat representations, covering over 140 distinct intents from a variety of domains including weather, communication, music, weather, and device control. The second internal dataset contains over 67k training utterances with fully hierarchical representations, and covers over 60 intents all in the communication domain.

The second and third public datasets are SNIPS Natural Language Understanding benchmark1 (SNIPS-NLU) and the Airline Travel Information Systems (ATIS) dataset (Hemphill et al., 1990). We follow the same procedure that was mentioned above for preparing the decoupled data for both of these datasets.

As can be seen from Table 1b, our proposed approach outperforms the previous state-of-the-art results on the ATIS, comparable to state-of-the-art on SNIPS, and TOP semantic parsing task, which had been obtained with the Seq2SeqPtr model by Rongali et al. (2020). Comparing the decoupled model to RNNs, we note that a single decoupled model, using either biLSTMs or transformers (with RoBERTa or BART pretraining) is able to outperform the RNN. In fact, the decoupled model even outperforms an ensemble of seven RNNs. The decoupled biLSTM extended with ELMo inputs is able to outperform the transformer model initialised with RoBERTa pretraining. However, the best performance is achieved by using the transformer model with BART-large pretraining, with the decoupled model fine-tuned jointly on top of it (Lewis et al., 2019). In order to understand how much of these gains are due to the semantic representation, we perform an ablation study by



(a) “i want a restaurant in the south part of town that serves moroccan food”.



(b) “how about modern european”

Figure 5: Example DSTC2 session, annotated for the decoupled model.

evaluating the biLSTM and RoBERTa-based models on TOP data using the standard logical form representation, and find a drop in frame accuracy of 0.32 and 0.55 respectively.

The TOP dataset contains to the order of 30k examples in its training set. In order to further tease out the differences between the biLSTM and transformer approaches, and to see how they compare when more training data is available, we also evaluate these models on our two larger internal datasets. Table 1c shows that the RoBERTa-based model does indeed benefit from the extra training data, being able to outperform the biLSTM-based model on the two datasets. In both cases, the decoupled model with BART pretraining achieves the top performance.

The same procedure was used over our SB-TOP dataset, with the only variant being we concatenated SB-TOP and TOP and jointly trained over both datasets. Table 2 shows the test results over

### 4.3 Slot carryover

To evaluate the ability of the decoupled models to work on session-based data, we evaluate them on a task which requires drawing information for multiple utterances. The DSTC2 dataset (Henderson et al., 2014) contains a number of dialogues annotated with dialogue state – slightly over 2k sessions in

Table 1: Frame accuracy of the decoupled models on semantic parsing tasks. † indicates results from Hakkani-Tür et al. (2016); ‡, from Goo et al. (2018); \*, from Zhang et al. (2018); ×, from Chen et al. (2019a).

(a) Accuracy on TOP.		(b) Accuracy on ATIS and SNIPS.			(c) Accuracy on internal datasets.	
Model	Acc.	Model	ATIS	SNIPS	Model	Acc.
RNNG	80.86	Joint biRNN <sup>†</sup>	80.7	73.2	Multi-domain (170k)	
RNNG + Ensembling	83.84	Slot gated <sup>‡</sup>	82.2	75.5	Decoupled ELMo	86.03
RNNG + ELMo	83.93	CapsuleNLU*	83.4	80.9	Decoupled RoBERTa	87.32
Decoupled biLSTM	79.51	Joint BERT <sup>×</sup>	88.2	<b>92.8</b>	Decoupled BART	<b>88.29</b>
Decoupled transformer	64.50	Joint BERT CRF <sup>×</sup>	88.6	92.6	Single-domain (67k)	
Decoupled ELMo	84.85	Decoupled BART	<b>89.25</b>	91.00	Decoupled ELMo	90.52
Decoupled RoBERTa	84.52	Best Seq2SeqPtr	87.12	87.14	Decoupled RoBERTa	91.51
Decoupled BART	<b>87.10</b>				Decoupled BART	<b>92.16</b>
Best Seq2SeqPtr	86.67					

Table 2: Decoupled model architecture results over the SB-TOP dataset. FA is exact match between canonicalized predicted and tree structures. Ref Only FA does not distinguish between implicit/explicit references. Intent accuracy is accuracy over top level intents while Inner Parse Accuracy is FA not considering top level intent.

Model	Oracle@Beam	FA	Ref-only FA	Intent Acc.	Inner Parse Acc.
Humans	1	55.04	57.4	84.32	60.12
Decoupled biLSTM	1	48.48	49.19	78.60	52.74
	5	60.24	69.88	93.71	72.01
Decoupled ELMo	1	51.22	52.03	80.93	55.07
	5	62.58	70.08	94.73	72.11
Decoupled BART	1	<b>53.45</b>	<b>54.18</b>	<b>82.46</b>	<b>56.84</b>
	5	<b>65.19</b>	<b>72.78</b>	<b>96.67</b>	<b>76.45</b>

the training set. They involve users searching for restaurants, by specifying constraints such as cuisine type and price range. Given that users will often take multiple turns to specify all constraints, determining the correct dialogue state requires the model to consider all past turns too. Consider the example of the two-turn DSTC2 session shown in Figure 5: the [SL:AREA south] slot, introduced in the first session, is said to *carry over* to the second session as it still applies to the dialogue state, despite not being explicitly mentioned.<sup>1</sup> To make previous utterances available to the model, we use a simple approach: all utterances are concatenated, with a separator token, and are fed to the encoder.

<sup>1</sup>The image shows the tree form to which we converted the DSTC2 native state tracking annotations, to make them easily linearizable and thus treatable by the decoupled models.

The decoupled models are evaluated on frame accuracy and slot carryover – the fraction of slots correctly carried over from one turn to the next. Carryover figures are split by slot distance: how many turns prior to the current one the slot under consideration first appeared. As shown in Table 3, the RoBERTa decoupled model outperforms the biLSTM model on frame accuracy, while the biLSTM model takes the lead in terms of raw slot carryover performance. BART outperforms both, achieving the best overall performance.

For informative purposes, we also include results from standard dialogue state tracking models. The results show that the decoupled models, despite not being specifically designed for the task of dialogue state tracking, compare favorably to other approaches in the literature. While our models outperform them on most metrics, it should be noted that they

Table 3: Performance of the decoupled models on a state tracking task (DSTC2).

Model	Accuracy	Slot distance			
		0	1	2	$\geq 3$
LSTM-based (Naik et al., 2018)	—	92.42	91.11	91.34	87.99
Pointer network decoder (Chen et al., 2019b)	—	92.70	92.04	92.90	91.39
Transformer decoder (Chen et al., 2019b)	—	93.00	92.69	92.80	89.49
GLAD (Zhong et al., 2018)	74.5	—	—	—	—
Decoupled biLSTM	88.3	93.34	94.73	95.28	95.73
Decoupled RoBERTa	89.8	91.98	92.94	93.58	94.28
Decoupled BART	<b>90.2</b>	94.21	95.47	95.90	97.05

are very different in nature: the decoupled models attend over all utterances leading up to and including the current turn, while state tracking models generally only have access to the current utterance and the previous system actions – in the case of Zhong et al. (2018) – or a fixed length dialogue representation. It is interesting to note that the decoupled models perform better on distant slots: this suggests that the models may be paying more attention to the beginning of the sentences, which may be an artifact of their pretraining.

## 5 Related Work

Traditional work on semantic parsing, either for the purposes of question answering or task-oriented request understanding, has focused on mapping utterances to logical form representations (Zelle and Mooney, 1996; Zettlemoyer and Collins, 2005; Kwiatkowski et al., 2010; Liang, 2016; van Noord et al., 2018). Logical forms, while very expressive, are also complex. Highly trained annotators are required for the creation of training data, and as a result there is a lack of large scale datasets that make use of these formalisms.

Intent-slot representations such as those used for the ATIS dataset (Price, 1990) or the datasets released as part of the DSTC challenges (Henderson et al., 2014; Rastogi et al., 2019) have less expressive power, but have the major advantage of being simple enough to enable the creation of large-scale datasets. Gupta et al. (2018) introduce a hierarchical

intent-slot representation, and show that it is expressive enough to capture the majority of user-generated queries in two domains.

Recent approaches to semantic parsing have focused on using techniques such as RNNs (Gupta et al., 2018), RNNs augmented with ensembling and re-ranking techniques or contextual embeddings (Einolghozati et al., 2018), sequence-to-sequence recurrent neural networks augmented with pointer mechanisms (Jia and Liang, 2016), capsule networks (Zhang et al., 2019), and Transformer-based architectures (Rongali et al., 2020).

## 6 Conclusions

We started this paper by exploring the limitations of compositional intent-slot representations for semantic parsing. Due to the constraints it imposes, it cannot represent certain utterances with long-term dependencies, and it is unsuitable for semantic parsing at the session (multi-utterance) level. To overcome these limitations we propose an extension of this representation, the *decoupled representation*. We propose a family of sequence-to-sequence models based on the pointer-generator architecture – using both recurrent neural network and transformer architectures – and show that they achieve top performance on several semantic parsing tasks. Further, to advance session-based task-oriented semantic parsing, we release to the public a new dataset of roughly 20k sessions (over 60k utterances).



## References

- Ankur Bapna, Gokhan Tür, Dilek Hakkani-Tür, and Larry Heck. 2017. [Towards zero-shot frame semantic parsing for domain scaling](#). In *Proc. Interspeech 2017*, pages 2476–2480.
- Qian Chen, Zhu Zhuo, and W. Wang. 2019a. Bert for joint intent classification and slot filling. *ArXiv*, abs/1902.10909.
- Tongfei Chen, Chetan Naik, Hua He, Pushpendre Rastogi, and Lambert Mathias. 2019b. [Improving long distance slot carryover in spoken dialogue systems](#). In *Proceedings of the First Workshop on NLP for Conversational AI*, pages 96–105, Florence, Italy. Association for Computational Linguistics.
- Noam Chomsky. 1981. Lectures on government and binding, foris, dordrecht. *Chomsky Lectures on Government and Binding 1981*.
- 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.
- Arash Einolghozati, Panupong Pasupat, Sonal Gupta, Rushin Shah, Mrinal Mohit, Mike Lewis, and Luke Zettlemoyer. 2018. Improving semantic parsing for task oriented dialog. In *Conversational AI Workshop at NeurIPS 2018*.
- Chih-Wen Goo, Guang Gao, Yun-Kai Hsu, Chih-Li Huo, Tsung-Chieh Chen, Keng-Wei Hsu, and Yun-Nung Chen. 2018. Slot-gated modeling for joint slot filling and intent prediction. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 753–757.
- Sonal Gupta, Rushin Shah, Mrinal Mohit, Anuj Kumar, and Mike Lewis. 2018. [Semantic parsing for task oriented dialog using hierarchical representations](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2787–2792, Brussels, Belgium. Association for Computational Linguistics.
- Dilek Hakkani-Tür, Gökhan Tür, Asli Celikyilmaz, Yun-Nung Chen, Jianfeng Gao, Li Deng, and Ye-Yi Wang. 2016. Multi-domain joint semantic frame parsing using bi-directional rnn-lstm. In *Interspeech*, pages 715–719.
- Charles T Hemphill, John J Godfrey, and George R Doddington. 1990. The atis spoken language systems pilot corpus. In *Speech and Natural Language: Proceedings of a Workshop Held at Hidden Valley, Pennsylvania, June 24-27, 1990*.
- Matthew Henderson, Blaise Thomson, and Jason D. Williams. 2014. [The second dialog state tracking challenge](#). In *Proceedings of the 15th Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL)*, pages 263–272, Philadelphia, PA, U.S.A. Association for Computational Linguistics.
- Pavel Izmailov, Dmitry Podoprikin, Timur Garipov, Dmitry Vetrov, and Andrew Gordon Wilson. 2018. Averaging weights leads to wider optima and better generalization. In *Conference on Uncertainty in Artificial Intelligence (UAI 2018)*.
- Robin Jia and Percy Liang. 2016. [Data recombination for neural semantic parsing](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12–22, Berlin, Germany. Association for Computational Linguistics.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Tom Kwiatkowski, Luke Zettlemoyer, Sharon Goldwater, and Mark Steedman. 2010. [Inducing probabilistic CCG grammars from logical form with higher-order unification](#). In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*, pages 1223–1233, Cambridge, MA. Association for Computational Linguistics.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. 2019. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *arXiv preprint arXiv:1910.13461*.
- Percy Liang. 2016. [Learning executable semantic parsers for natural language understanding](#). *Commun. ACM*, 59(9):68–76.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike

- Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Chetan Naik, Arpit Gupta, Hancheng Ge, Mathias Lambert, and Ruhi Sarikaya. 2018. [Contextual slot carryover for disparate schemas](#). In *Proc. Interspeech 2018*, pages 596–600.
- Rik van Noord, Lasha Abzianidze, Antonio Toral, and Johan Bos. 2018. [Exploring neural methods for parsing discourse representation structures](#). *Transactions of the Association for Computational Linguistics*, 6:619–633.
- 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.
- P. J. Price. 1990. [Evaluation of spoken language systems: the ATIS domain](#). In *Speech and Natural Language: Proceedings of a Workshop Held at Hidden Valley, Pennsylvania, June 24–27, 1990*.
- Abhinav Rastogi, Xiaoxue Zang, Srinivas Sunkara, Raghav Gupta, and Pranav Khaitan. 2019. Towards scalable multi-domain conversational agents: The schema-guided dialogue dataset. *arXiv preprint arXiv:1909.05855*.
- Subendhu Rongali, Luca Soldaini, Emilio Monti, and Wael Hamza. 2020. Don’t parse, generate! a sequence to sequence architecture for task-oriented semantic parsing. *arXiv preprint arXiv:2001.11458*.
- Abigail See, Peter J. Liu, and Christopher D. Manning. 2017. [Get to the point: Summarization with pointer-generator networks](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1073–1083, Vancouver, Canada. Association for Computational Linguistics.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. [Sequence to sequence learning with neural networks](#). In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 3104–3112. Curran Associates, Inc.
- 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 I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc.
- Oriol Vinyals, Łukasz Kaiser, Terry Koo, Slav Petrov, Ilya Sutskever, and Geoffrey Hinton. 2015. [Grammar as a foreign language](#). In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 2773–2781. Curran Associates, Inc.
- John M. Zelle and Raymond J. Mooney. 1996. Learning to parse database queries using inductive logic programming. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence - Volume 2, AAAI’96*, page 1050–1055. AAAI Press.
- Luke S. Zettlemoyer and Michael Collins. 2005. Learning to map sentences to logical form: Structured classification with probabilistic categorical grammars. In *Proceedings of the Twenty-First Conference on Uncertainty in Artificial Intelligence, UAI’05*, page 658–666, Arlington, Virginia, USA. AUAI Press.
- Chenwei Zhang, Yaliang Li, Nan Du, Wei Fan, and Philip Yu. 2019. [Joint slot filling and intent detection via capsule neural networks](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5259–5267, Florence, Italy. Association for Computational Linguistics.
- Chenwei Zhang, Yaliang Li, Nan Du, Wei Fan, and Philip S Yu. 2018. [Joint slot filling and intent detection via capsule neural networks](#). *arXiv preprint arXiv:1812.09471*.
- Victor Zhong, Caiming Xiong, and Richard Socher. 2018. [Global-locally self-attentive encoder for dialogue state tracking](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1458–1467, Melbourne, Australia. Association for Computational Linguistics.