

Don't Parse, Generate! A Sequence to Sequence Architecture for Task-Oriented Semantic Parsing

Subendhu Rongali*

University of Massachusetts Amherst
Amherst, MA, USA
srongali@cs.umass.edu

Emilio Monti

Amazon Alexa
Cambridge, UK
monti@amazon.co.uk

Luca Soldaini

Amazon Alexa Search
Manhattan Beach, CA, USA
lssoldai@amazon.com

Wael Hamza

Amazon Alexa AI
New York, NY, USA
waelhamz@amazon.com

ABSTRACT

Virtual assistants such as Amazon Alexa, Apple Siri, and Google Assistant often rely on a semantic parsing component to understand which action(s) to execute for an utterance spoken by its users. Traditionally, rule-based or statistical slot-filling systems have been used to parse “simple” queries; that is, queries that contain a single action and can be decomposed into a set of non-overlapping entities. More recently, shift-reduce parsers have been proposed to process more complex utterances. These methods, while powerful, impose specific limitations on the type of queries that can be parsed; namely, they require a query to be representable as a parse tree.

In this work, we propose a unified architecture based on Sequence to Sequence models and Pointer Generator Network to handle both simple and complex queries. Unlike other works, our approach does not impose any restriction on the semantic parse schema. Furthermore, experiments show that it achieves state of the art performance on three publicly available datasets (ATIS, SNIPS, Facebook TOP), relatively improving between 3.3% and 7.7% in exact match accuracy over previous systems. Finally, we show the effectiveness of our approach on two internal datasets.

KEYWORDS

Natural Language Understanding, Semantic Parsing, Voice Assistants, Sequence to Sequence models

ACM Reference Format:

Subendhu Rongali, Luca Soldaini, Emilio Monti, and Wael Hamza. 2020. Don't Parse, Generate! A Sequence to Sequence Architecture for Task-Oriented Semantic Parsing. In *Proceedings of The Web Conference 2020 (WWW '20)*, April 20–24, 2020, Taipei, Taiwan. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3366423.3380064>

*Work done when the author was at Amazon Alexa AI for a summer internship.

This paper is published under the Creative Commons Attribution 4.0 International (CC-BY 4.0) license. Authors reserve their rights to disseminate the work on their personal and corporate Web sites with the appropriate attribution.

WWW '20, April 20–24, 2020, Taipei, Taiwan

© 2020 IW3C2 (International World Wide Web Conference Committee), published under Creative Commons CC-BY 4.0 License.

ACM ISBN 978-1-4503-7023-3/20/04.

<https://doi.org/10.1145/3366423.3380064>

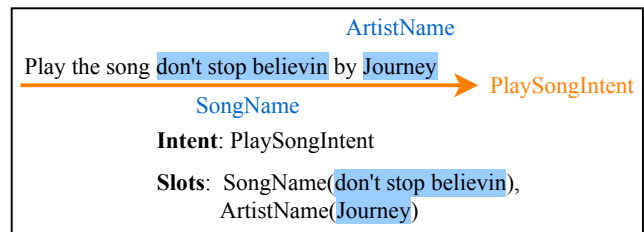


Figure 1: Semantic parsing of a “simple” query. Simple queries define single action (intent) and can be decomposed into a set of non-overlapping entities (slots).

1 INTRODUCTION

Adoption of intelligent voice assistants such as Amazon Alexa, Apple Siri, and Google Assistant has increased dramatically among consumers in the past few years: as of early 2019, it is estimated that 21% of U.S. adults own a smart speaker, a 78% year-over-year growth [20]. These systems are built to process user dialog and perform tasks such as media playback and online shopping.

A major part of any voice assistant is a semantic parsing component designed to understand the action requested by its users: given the transcription of an utterance, a voice assistant must identify the action requested by a user (play music, turn on lights, etc.), as well as parse any entities that further refine the action to perform (which song to play? which lights to turn on?) Despite huge advances in the field of Natural Language Processing (NLP), this task still remains challenging due to the sheer number of possible combinations a user can use to express a command.

Traditional approaches for task-oriented semantic dialog parsing frame the problem as a slot filling task. For example, given the query *Play the song don't stop believin by Journey*, a traditional slot filling system parses it in two independent steps: (i) It first classifies the *intent* of the user utterance as `PlaySongIntent`, and then (ii) identifies relevant *named entities* and tags those slots, such as *don't stop believin* as a `SongName` and *Journey* as an `ArtistName`. Traditional semantic parsing can therefore be reduced to a text classification and a sequence tagging problem, which is a standard architecture for many proposed approaches in literature [14, 16, 18]. This is shown in Figure 1.

With increasing expectations of users from virtual assistants, there is a need for the systems to handle more complex queries

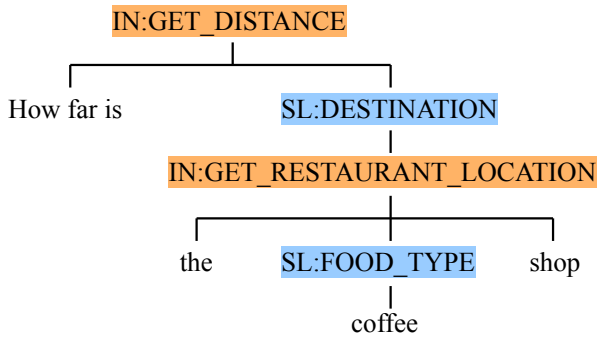


Figure 2: Semantic parse for a “complex” query in the Facebook TOP dataset. This complex query is represented as a tree containing two nested intents and slots.

– ones that are composed of multiple intents and nested slots or contain conditional logic. For example, the query *Are there any movie in the park events nearby?* involves first finding the location of parks that are nearby and then finding relevant movie events in them. This is not straightforward in traditional slot filling systems. Gupta et al. [6] and Einolghozati et al. [4] proposed multiple approaches for this using a Shift-reduce parser based on Recurrent Neural Network Grammars [3] that performs the tagging.

In this paper, we propose a unified approach to tackle semantic parsing for natural language understanding based on Transformer Sequence to Sequence models [33] and a Pointer Generator Network [29, 34]. Furthermore, we demonstrate how our approach can leverage pre-trained resources, such as neural language models, to achieve state of the art performance on several datasets. In particular, we obtain relative improvements between 3.3% and 7.7% over the best single systems on three public datasets (SNIPS [1], ATIS [26] and TOP [6], the last consisting of complex queries); on two internal datasets, we show relative improvements of up to 4.9%.

Furthermore, our architecture can be easily used to parse queries that do not conform to the grammar of either the slot filling or RNN systems. Some examples include semantic entities that correspond to overlapping spans in the query, and entities comprising of non-consecutive spans. We do not report any results on these kinds of datasets but we explain how to formulate the problems using our architecture.

In summary, our contributions are as follows.

- We propose a new architecture based on Sequence to Sequence models and a Pointer Generator Network to solve the task of semantic parsing for understanding user queries.
- We describe how to formulate different kinds of queries in our architecture. Our formulation is unified across queries with different kinds of tagging.
- We achieve state-of-the-art results on three public datasets and two internal datasets.

2 METHODOLOGY

We propose a unified architecture to solve the task of semantic parsing for both simple and complex queries. This architecture can also be adapted to handle queries containing slots with overlapping

spans. It consists of a Sequence to Sequence model and a Pointer Generator Network. We choose a pretrained BERT [2] model as our encoder. Our decoder is modeled after the transformer decoder described in Vaswani et al. [33] and is augmented with a Pointer Generator Network [9, 34] which allows us to learn to generate pointers to the source sequence in our target sequence. Figure 3 shows this architecture parsing an example query. We train the model using a cross-entropy loss function with label smoothing.

In this section, we first describe how we formulate queries and their semantic parses as sequences with pointers for our architecture. We then describe our encoder and decoder components.

2.1 Query Formulation

A Sequence to Sequence architecture is trained on samples with a source sequence and a target sequence. When some words in the target sequence are contained in the source sequence, they can be replaced with a separate *pointer* token that points to that word in the source to be able to apply the Pointer Generator Network.

Take the example query from Figure 1. In our architecture, we use the query as our source sequence. The target sequence is constructed by combining the intent with all the slots, in order, with each slot also containing its source words. The source and target sequences now look as follows.

```

Source: play the song don't stop believin by journey
Target: PlaySongIntent SongName( @ptr3 @ptr4 @ptr5
      )SongName ArtistName( @ptr7 )ArtistName
  
```

Here, each token $@ptr_i$ is a pointer to the i^{th} word in the source sequence. So $@ptr_3$, $@ptr_4$ and $@ptr_5$ point to the song words *don't stop believin*, and $@ptr_7$ points to the artist word *journey*. The slots have open and close tags since they are enclosing a consecutive span of source tokens. The intent is just represented as a single tag at the beginning of the target sequence. We can do this for simple queries since they consist of just one intent. The target vocabulary hence consists of all the available intents, two times the number of different slots, and the pointers.

Complex queries with multiple intents and nested slots can also be transformed easily into this formulation. Figure 2 shows an example from the Facebook TOP dataset along with its parse tree. This query *How far is the coffee shop* can be converted into our formulation as follows.

```

Source: How far is the coffee shop
Target: [IN:GET_DISTANCE @ptr0 @ptr1 @ptr2 [SL:DESTINATION
      [IN:GET_RESTAURANT_LOCATION @ptr3 [SL:TYPE_FOOD
      @ptr4 SL:TYPE_FOOD] @ptr5 IN:GET_RESTAURANT_LOCATION]
      SL:DESTINATION] IN:GET_DISTANCE]
  
```

We made a minor modification to the reference parses from the TOP dataset for our formulation. We replaced the end-brackets with custom end-brackets corresponding to the intent or slot they close. We found that this formulation helped our models perform better.

Finally, we show how we can express queries from datasets that don't conform to either the slot-filling or Shift-reduce systems. Take the following example from the healthcare domain, where the task is to extract a patient diagnosis and related information from a clinician's notes.

```

Source: The pt. was diagnosed with GI upper bleed today.
Annotations: Bleeding_Event (GI bleed),
      Anatomical_Site (upper)
  
```

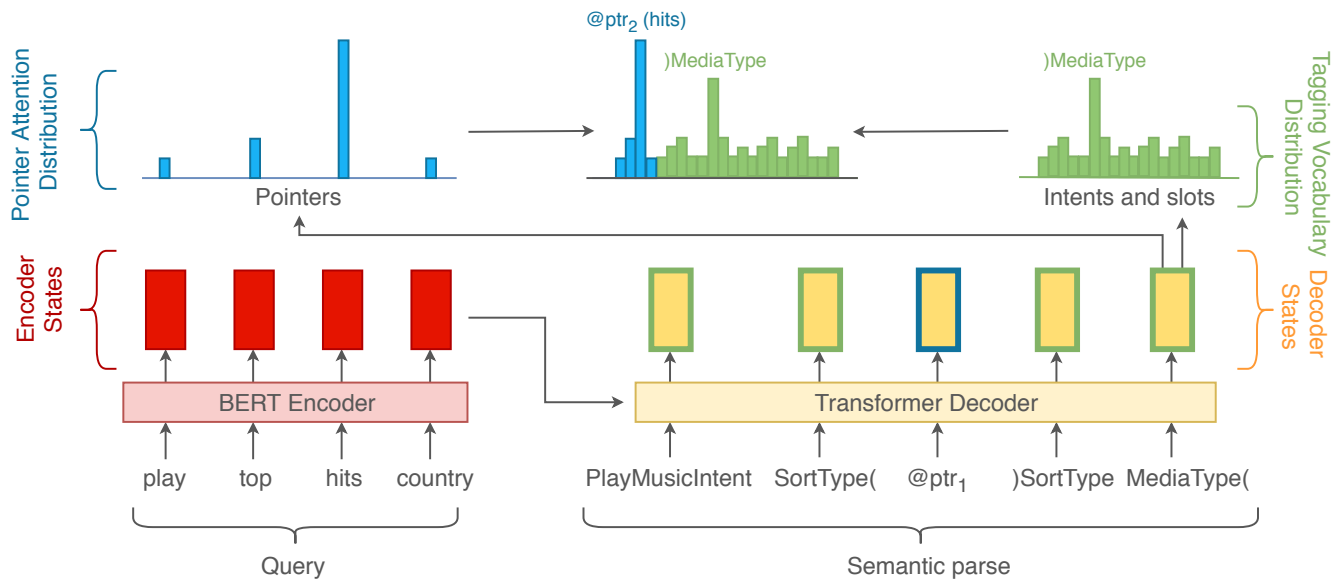


Figure 3: Our architecture - Sequence to Sequence model with Pointer Generator Network (SEQ2SEQ-Ptr). The model is currently decoding the symbol after `MediaType(` by looking at the scores over the tagging vocabulary and the attentions over the source pointers. It generates `@ptr2` since it has the highest overall score.

A traditional slot filling system wouldn't know which non consecutive slots to combine, while a shift-reduce parser cannot split the middle word into a separate tag. In our architecture, we simply formulate the target sequence as follows.

Target: `Bleeding_Event(@ptr5 @ptr7)Bleeding_Event`
`Anatomical_Site(@ptr6)Anatomical_Site`

2.2 BERT Encoder

Language model pretraining has been shown to improve the downstream performance on many NLP tasks [2, 24, 27]. The idea is to train a language model on a large amount of text using a next word prediction objective to learn good representations for each of the words. These representations can then be fine-tuned on a given NLP task to improve the performance of an existing model. Pretrained models improve the performance of task models since they already contain a lot of useful semantic information learned through the pretraining phase. This has even more significance when the task-specific dataset is fairly small. Some examples of pretrained models in literature include word embeddings such as Word2Vec [19] and Glove [23], and contextualized representations such as ELMo [24], OpenAI-GPT [27], and BERT [2].

We choose BERT to encode the source sequence in our architecture. BERT (Bidirectional Encoder Representations from Transformers) is a language representation model architecture based on Transformers [33]. The original publicly available model was pre-trained on a millions of lines of text from BooksCorpus and English Wikipedia. Unlike other language models (ELMo, OpenAI-GPT), which are trained to predict the next token given the previous sequence of words, BERT uses a composite objective that combines masked word prediction and next sentence prediction.

BERT's architecture is based on a multi-layer bidirectional Transformer, originally implemented in Vaswani et al. [33]. The detailed implementation of this architecture can be found in Devlin et al. [2]. For our experiments, we use three different variants of BERT.

For the three public datasets, we used the checkpoint released by Devlin et al. [2]. Experiments on the two internal datasets were carried out using a model we pretrained over a large sample of queries from the live traffic of Amazon Alexa. We also experimented with a publicly-available variant of BERT called RoBERTa [17]. RoBERTa (A Robustly Optimized BERT Pretraining Approach) uses the same architecture as BERT but changes the pretraining process. The next sentence prediction objective is removed and a dynamic masking scheme is used instead of a static one like the original BERT implementation. RoBERTa was also trained with longer sequences, higher batch-sizes, and for a longer time, and was reported to match or exceed the performance of BERT in several NLP benchmarks. The detailed implementation can be found in Liu et al. [17]. Finally as an ablation study, we experimented with an encoder with no pretrained weights.

2.3 Decoder with Pointer Generator Network

We use the transformer decoder proposed in Vaswani et al. [33] in our architecture. The self-attention mechanism in the decoder learns to attend to target words before the current step, as well as all the source words in the encoder.

We set up the decoder with different numbers of units, layers, and attention-heads for different tasks based on the size and complexity of the queries. These details are provided in the experiments section.

In a traditional Sequence to Sequence model, the target words are generated from the decoder hidden states through a feed-forward layer that obtains unnormalized scores over a target vocabulary

distribution. In our architecture, we use a Pointer Generator Network to generate two different kinds of target words: words from the target vocabulary consisting of parse symbols (the intent and slot delimiters), and words that are simply pointers to the source sequence. Our Pointer Generator Network is based on the models in Vinyals et al. [34] and See et al. [29], and is closest in implementation to Jia and Liang [9].

We now describe our decoding process. For each input source sequence $[x_1 \dots x_n]$, we use the BERT encoder to encode it into a sequence of encoder hidden states $[e_1 \dots e_n]$. Having generated the first $t - 1$ output tokens, the transformer decoder generates the token at step t as follows.

First, the decoder produces the decoder hidden state at time t , d_t by building multi-layer multi-head self-attention on the encoded output as well as the embeddings of the previously generated output sequence as described in Vaswani et al. [33]. We feed d_t through a dense layer to produce scores $[s_1, \dots, s_{|V|}]$ for each word in the vocabulary V . V contains all symbols necessary for the parse (intents, slots) but not regular words appearing on the source side.

We also use d_t as a query and compute unnormalized attention scores $[a_1 \dots a_n]$ with the encoded sequence. Concatenating the unnormalized attention scores (size n) and the output of the dense layer (size $|V|$), we obtain an unnormalized distribution over $|V| + n$ tokens, the first $|V|$ of which are the output parsing vocabulary and the last n of which are the $@ptr_i$ ($0 < i < n$) words pointing to the source tokens. We then feed this through a softmax layer to obtain the final probability distribution. This probability is used in the loss function during training and will be used to choose the next token to generate during inference. Since the transformer decoder uses embeddings of previously generated tokens, we use a set of special embeddings to represent $@ptr_i$ tokens.

In the example in Figure 3, we are trying to predict a target word after the token 'MediaTpe(' at step 5. As shown in the figure, we compute the scores $[a_1 \dots a_4]$ (blue, left) over each of the source tokens, and the scores $[s_1 \dots s_{|V|}]$ (green, right) over the parsing vocabulary. We expect the model to produce the highest score for a_3 , which corresponds to $@ptr_2$, representing the word *hits*.

3 DATASETS

We test our approach on five different datasets (three publicly available, two internal), which we describe in this section.

3.1 Facebook TOP

The Task Oriented Parsing (TOP) [6] dataset contains complex hierarchical and nested queries that make the task of semantic parsing more challenging. It contains around 45k annotations with 25 intents and 36 slots, randomly split into 31k training, 5k validation and 9k test utterances. The dataset mainly consists of user queries about navigation and various public events. An example from this dataset can be seen in Figure 2. The IN: prefix stands for intent while SL: is for slot. We can see how there are multiple intents and nested slots in the semantic interpretation. This makes the query much harder to interpret and parse using a simple slot tagging model that tags each word with a single slot.

3.2 SNIPS

The SNIPS dataset [1] is a public dataset that is used for training and testing semantic parsing models for voice assistants. It consists of utterances that belong to seven different intents: SearchCreativeWork, GetWeather, BookRestaurant, PlayMusic, AddToPlaylist, RateBook, and SearchScreeningEvent. Each intent contains around 2000 examples to train and 100 to test.

This dataset contains only simple queries with single intents and flat slots. An example is *Will there be fog in Tahquamenon Falls State Park*, where the intent is GetWeather and the slots are condition_description for *fog* and geographic_poi for *Tahquamenon Falls State Park*.

The dataset was originally used to evaluate models in the Snips Voice Platform. It has since been a widely used dataset to benchmark the performance of various task-oriented parsing models.

3.3 ATIS

The Airline Travel Information System (ATIS) [26] corpus is a widely used dataset in spoken language understanding. It was built by collecting and transcribing audio recordings of people making flight reservations in the early 90s. It consists of simple queries.

There are seventeen different goals or intents such as Flight or Aircraft capacity. This distribution is however skewed, with the Flight intent covering about 70% of the total queries. An example from this dataset consists of the query *How much is the cheapest flight from Boston to New York tomorrow morning?* The intent is Airfare, while the slots tag important information like the departure and arrival cities, and the departure times.

The ATIS corpus has supported research in the field of spoken language understanding for more than twenty years. Some researchers have performed extensive error analysis on the state of the art discriminative models for this dataset and reported that despite really low error rates, there exist many unseen categories and sequences in the dataset that can benefit from incorporating linguistically motivated features [31]. This supports the continued utility of ATIS as a research corpus.

3.4 Internal Datasets

Our internal datasets consist of millions of user utterances that are used to train and test Amazon Alexa. For our experiments, we sampled two datasets of utterances, one from the music domain and the other from video domain. Utterances in these domains naturally included a large amount of entities (e.g. artists and albums names, movie and video titles), and thus represent a good benchmark for the ability of any neural model to generalize over a diverse set of queries. The example in Figure 3 is from the music domain.

The sampled music domain dataset contains 6.2M training and 200k test utterances, with 23 intents and 100 slots. The video domain dataset contains 1M training and just 5k test utterances; parses in this dataset are comprised of 24 distinct intents and 59 slots.

4 BASELINE MODELS

We benchmark our performance on the internal datasets by comparing it to a well tuned RNN based model. The model learns to perform joint intent and slot tagging using a bidirectional LSTM and a Conditional Random Field (CRF) [8]. We further enhanced

Method	Accuracy	
	exact match	intent
Shift Reduce (SR) Parser [4]	80.86	–
SR with ELMo embeddings [4]	83.93	–
SR ensemble + ELMo + SVMRank [4]	87.25	–
SEQ2SEQ-PTR (no pretraining)	79.25	97.43
SEQ2SEQ-PTR (BERT encoder)	83.13	97.91
SEQ2SEQ-PTR (RoBERTa encoder)	86.67	98.13

Table 1: Results on TOP [6]. Our system (SEQ2SEQ-PTR) outperforms the best single method (SR + ELMo) by 3.3%. It is close to the best ensemble approach (SR + ELMo + SVRank).

Method	Accuracy	
	exact match	intent
Joint BiRNN [7]	73.20	96.90
Attention BiRNN [16]	74.10	96.70
Slot Gated Full Attention [5]	75.50	97.00
CAPSULENLU [36]	80.90	97.30
SEQ2SEQ-PTR (no pretraining)	85.43	97.00
SEQ2SEQ-PTR (BERT encoder)	86.29	98.29
SEQ2SEQ-PTR (RoBERTa encoder)	87.14	98.00

Table 2: Results on SNIPS [1]. Our system (SEQ2SEQ-PTR) outperforms the previous state of the art by 7.7%.

this baseline by replacing its embedding and encoder layers with a language model pretrained on a subset of Alexa’s live traffic. These components were fine-tuned on the two datasets described in Section 3.4.

For the ATIS and SNIPS datasets, we use the top four performing methods reported by Zhang et al. [36] as baselines. All these models perform joint intent and slot tagging. There are two variants that use RNNs: a simple RNN based model, and an RNN model augmented with attention. There is also a model that works completely with just attention, the slot gated full attention model. The final baseline, CapsuleNLU, uses Capsule Networks [28].

For the TOP dataset, we pick a model based on Recurrent Neural Network Grammars (RNNG) [3], the Shift Reduce Parser. We provide a brief overview of this model as described in Gupta et al. [6] - the parse tree is constructed using a sequence of transitions, or *actions*. The transitions are defined as a set of SHIFT, REDUCE, and the generation of intent and slot labels. SHIFT action consumes an input token (that is, adds the token as a child of the right most *open* sub-tree node) and REDUCE closes a sub-tree. The third set of actions is generating non-terminals: the slot and intent labels. The model learns to perform one of these actions at each step in time.

We report scores of three experimental setups with the shift reduce parser from Einolghozati et al. [4]: a simple shift reduce parser, a shift reduce parser augmented with ELMo embeddings, and an ensemble of these models augmented with ELMo and an SVM language model reranker.

5 EXPERIMENTAL SETUP

All our models were trained on a machine with 8 NVIDIA Tesla V100 GPUs, each with 16GB of memory. When using pretrained encoders, we leveraged gradual unfreezing to effectively tune the language

Method	Accuracy	
	exact match	intent
Joint BiRNN [7]	80.70	92.60
Attention BiRNN [16]	78.90	91.10
Slot Gated Full Attention [5]	82.20	93.60
CAPSULENLU [36]	83.40	95.00
SEQ2SEQ-PTR (no pretraining)	81.08	95.18
SEQ2SEQ-PTR (BERT encoder)	86.37	97.42
SEQ2SEQ-PTR (RoBERTa encoder)	87.12	97.42

Table 3: Results on ATIS [26]. Our system (SEQ2SEQ-PTR) outperforms the previous best method by 4.5%.

model layers on our datasets. We used the "BASE" variant of BERT and RoBERTa encoders, which uses 768-dimensional embeddings, 12 layers, 12 heads, and 3072 hidden units. When training from scratch, we used a smaller encoder consisting of 512-dimensional embeddings, 6 layers, 8 heads, and 1024 hidden units.

Depending on the dataset, we used either a 128 units, 4 layers, 3 heads, and 512 hidden units decoder (Facebook TOP, ATIS, SNIPS) or a larger 512 units, 6 layers, 8 heads, and 1024 hidden units decoder (internal Music and Video datasets). We used bi-linear product attention to score the source words in the Pointer Network.

While training, the cross entropy loss function was modified with label smoothing with $\epsilon = 0.1$. We used the Adam [12] optimizer with noam learning rate schedule [33], each adjusted differently for different datasets. At inference time, we used beam search decoding with a beam size of 4.

6 RESULTS AND DISCUSSION

We use exact match (EM) accuracy as the main metric to measure the performance of our models across all datasets. Under this metric, the entire semantic parse for a query has to match the reference parse to be counted as correct. Because EM is generally more challenging than slot-level precision and recall or semantic error rate [30], it is better suited to compare high performing systems like the ones studied in this work. For completeness, we also report the intent classification accuracy for our models.

The results from our experiments are documented in Tables 1-5. Our models match or beat the baselines across all datasets on both exact match and intent classification accuracies. We see significant improvements on both simple and complex datasets.

6.1 Complex Queries

We achieve an improvement of 2.7 (+3.3%) EM accuracy points on the TOP dataset over the state-of-the-art single model on this dataset (Table 1). Our SEQ2SEQ-PTR model with RoBERTa encoder is only surpassed by the ensemble model reported in Einolghozati et al. [4] (+0.6% EM accuracy points).

In addition, we find that even without specifying any hard requirements for the grammar of the parse trees in the complex queries, 98% of the generated parses are well formatted. For the simple query datasets, it was greater than 99% but difference is expected since the grammar is easier to learn there.

During error analysis, we found an interesting example in the TOP dataset where we believe our model generates a valid, more

Method	Accuracy	
	exact match	intent
BiLSTM-CRF (no pretraining)	<i>baseline</i>	
BiLSTM-CRF (pretrained LM)	+3.0%	+0.1%
SEQ2SEQ-PTR (no pretraining)	-0.3%	-0.7%
SEQ2SEQ-PTR (BERT encoder)	-2.2%	-0.8%
SEQ2SEQ-PTR (RoBERTa encoder)	-3.5%	-0.7%

Table 4: Results on an internal dataset (music). The best configuration of our method (SEQ2SEQ-PTR) is comparable to a BiLSTM-CRF pretrained on a large conversational dataset.

meaningful parse than the reference annotation. For the query *What time do I need to leave to get to Helen by 8pm*, our model parses *Helen* as [SL: DESTINATION [IN: GET_LOCATION_HOME [SL: CONTACT Helen]]], while it is annotated as [SL: DESTINATION Helen]. Our parse resolves the query as finding the estimated departure time to get to a location that is the home location of a contact named Helen, while the reference annotation suggests that the correct interpretation is to find the estimated departure time to get to a destination named *Helen*. We believe our parse is more likely to be correct given that Helen is most likely the name of a person.

6.2 Simple Queries

We report results of our sequence to sequence model (SEQ2SEQ-PTR) on four datasets (SNIPS, ATIS, internal music, internal video) that contain simple queries in Tables 2, 3, 4, and 5.

On the SNIPS and ATIS datasets, we note that the best version of our method (SEQ2SEQ-PTR with RoBERTa encoder) achieves a significant improvement in EM accuracy over existing baselines (+7.7% and +4.5% respectively.) Using a BERT encoder causes a slight decrease in performance, but still achieves a meaningful improvement over the previous state of the art [36]; this is consistent with what has been observed on other NLP tasks [17]. If no pretraining is used, performance is further reduced but it is notable that this variant still beats all the baselines on the SNIPS dataset.

For our internal Alexa datasets, we note that the proposed SEQ2SEQ-PTR method obtains comparable results to a BiLSTM-CRF tagger on the music domain, and slightly better EM accuracy (+1.9%) on the video domain. We believe our model wasn't able to outperform the baseline on the music domain because the entities in this domain are very diverse, especially song or album names. Sequence tagging methods therefore benefit from having to solve a simpler task of having to tag each word in the sequence, as opposed to our unconstrained model. We would however like to note that our from-scratch variants beat the from-scratch baselines on both domains. Also curiously, the performance of SEQ2SEQ-PTR with a BERT encoder fell behind that of a sequence to sequence model trained from scratch. Since the scratch model uses a smaller transformer encoder (6 layers with 8 heads per layer instead of 12/12), we believe it was able to converge more effectively than the BERT encoder.

7 RELATED WORK

The task of semantic parsing for intent and slot detection is well established in literature. Traditionally, this was done with slot filling systems that classify the query and then label each word in the

Method	Accuracy	
	exact match	intent
BiLSTM-CRF (no pretraining)	<i>baseline</i>	
BiLSTM-CRF (pretrained LM)	+3.0%	+0.1%
SEQ2SEQ-PTR (no pretraining)	+2.9%	-0.1%
SEQ2SEQ-PTR (BERT encoder)	+0.1%	-0.2%
SEQ2SEQ-PTR (RoBERTa encoder)	+4.9%	-0.2%

Table 5: Results on an internal dataset (video). The best configuration of our method (SEQ2SEQ-PTR) outperforms a pre-trained BiLSTM-CRF network by 1.9%.

query. There were a few approaches that followed this system, using Recurrent Neural Networks [16, 18]. Researchers have also experimented with Convolutional Neural Networks and showed good results [11] and more recently, Capsule Networks [28, 36].

Prior to the advent of deep learning models, the task of sequence labeling was tackled with the use of Conditional Random Fields (CRF) [10, 14, 22]. CRFs learn pairwise potentials on labeling subsequent words which allow models to find more probable label sequences for a given query.

Most of this work is valid for semantic parsing for simple queries which boils down to a sequence labeling task. To handle more complex cases with hierarchical slots such as the example in Figure 2, researchers have experimented with Sequence to Sequence models and models based on Recurrent Neural Network Grammars (RNNG) [3]. RNNGs were shown to perform better on complex queries than RNN or Transformer-based Sequence to Sequence models [6]. Researchers have also explored models involving logical forms and discourse for language representation [15, 32, 35].

The Pointer Generator Network in our architecture was introduced in Vinyals et al. [34]. It was used in NLP applications where some words from the source sequence reappeared in the target sequence such as text summarization and style transfer [21, 25, 29]. They were also used to copy out of vocabulary words from the source to target in machine translation [13]. Our implementation of the Pointer Network is closest to the architecture in Jia and Liang [9]. By using pointers to represent the source tokens and imposing no particular logical form over our target sequence, we can handle any kind of queries for parsing. This makes our architecture as expressive as logical forms, while also being able to learn as easily as simple slot tagging systems.

8 CONCLUSION

We propose a unified architecture for the task of semantic parsing for different kinds of queries. We show that our architecture matches or outperforms existing approaches across multiple datasets : internal music and video datasets, SNIPS, ATIS, and Facebook TOP. We significantly outperform the current state of the art models on the public datasets TOP (3.3%), SNIPS (7.7%), and ATIS (4.5%).

We describe how to apply this architecture to both simple queries and complex queries with hierarchical and nested slots. We also describe how to formulate any set of queries with non-conforming grammars to work with our architecture, making this model applicable to many different types of semantic parsing. We leave the non-conforming grammar task to future work.

REFERENCES

- [1] Alice Coucke, Alaa Saade, Adrien Ball, Théodore Bluche, Alexandre Caulier, David Leroy, Clément Doumouro, Thibault Gisselbrecht, Francesco Caltagirone, Thibaut Lavril, et al. 2018. Snips voice platform: an embedded spoken language understanding system for private-by-design voice interfaces. *arXiv preprint arXiv:1805.10190* (2018).
- [2] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [3] Chris Dyer, Adhiguna Kuncoro, Miguel Ballesteros, and Noah A Smith. 2016. Recurrent neural network grammars. *arXiv preprint arXiv:1602.07776* (2016).
- [4] Arash Einolghozati, Panupong Pasupat, Sonal Gupta, Rushin Shah, Mrinal Mohit, Mike Lewis, and Luke Zettlemoyer. 2019. Improving Semantic Parsing for Task Oriented Dialog. *arXiv preprint arXiv:1902.06000* (2019).
- [5] 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)*. Association for Computational Linguistics, New Orleans, Louisiana, 753–757. <https://doi.org/10.18653/v1/N18-2118>
- [6] 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*. 2787–2792.
- [7] Dilek Hakkani-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*.
- [8] Zhiheng Huang, Wei Xu, and Kai Yu. 2015. Bidirectional LSTM-CRF models for sequence tagging. *arXiv preprint arXiv:1508.01991* (2015).
- [9] Robin Jia and Percy Liang. 2016. Data recombination for neural semantic parsing. *arXiv preprint arXiv:1606.03622* (2016).
- [10] Feng Jiao, Shaojun Wang, Chi-Hoon Lee, Russell Greiner, and Dale Schuurmans. 2006. Semi-supervised conditional random fields for improved sequence segmentation and labeling. In *Proceedings of the 21st International Conference on Computational Linguistics and the 44th annual meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, 209–216.
- [11] Yoon Kim. 2014. Convolutional neural networks for sentence classification. *arXiv preprint arXiv:1408.5882* (2014).
- [12] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [13] Guillaume Klein, Yoon Kim, Yuntian Deng, Jean Senellart, and Alexander M Rush. 2017. Opennmt: Open-source toolkit for neural machine translation. *arXiv preprint arXiv:1701.02810* (2017).
- [14] John Lafferty, Andrew McCallum, and Fernando CN Pereira. 2001. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. (2001).
- [15] Percy Liang. 2016. Learning executable semantic parsers for natural language understanding. *arXiv preprint arXiv:1603.06677* (2016).
- [16] Bing Liu and Ian Lane. 2016. Attention-based recurrent neural network models for joint intent detection and slot filling. *arXiv preprint arXiv:1609.01454* (2016).
- [17] 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* (2019).
- [18] Grégoire Mesnil, Xiaodong He, Li Deng, and Yoshua Bengio. 2013. Investigation of recurrent-neural-network architectures and learning methods for spoken language understanding. In *Interspeech*. 3771–3775.
- [19] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*. 3111–3119.
- [20] National Public Radio. 2019. Smart Speakers See 78% Increase YOY. <https://www.npr.org/about-npr/682946406/npr-report-smart-speakers-see-78-increase-yoy>
- [21] Romain Paulus, Caiming Xiong, and Richard Socher. 2017. A deep reinforced model for abstractive summarization. *arXiv preprint arXiv:1705.04304* (2017).
- [22] Fuchun Peng, Fangfang Feng, and Andrew McCallum. 2004. Chinese segmentation and new word detection using conditional random fields. In *Proceedings of the 20th international conference on Computational Linguistics*. Association for Computational Linguistics, 562.
- [23] Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*. 1532–1543.
- [24] 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* (2018).
- [25] Shrimai Prabhumoye, Yulia Tsvetkov, Ruslan Salakhutdinov, and Alan W Black. 2018. Style transfer through back-translation. *arXiv preprint arXiv:1804.09000* (2018).
- [26] Patti 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*.
- [27] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI Blog* 1, 8 (2019).
- [28] Sara Sabour, Nicholas Frosst, and Geoffrey E Hinton. 2017. Dynamic routing between capsules. In *Advances in neural information processing systems*. 3856–3866.
- [29] Abigail See, Peter J Liu, and Christopher D Manning. 2017. Get to the point: Summarization with pointer-generator networks. *arXiv preprint arXiv:1704.04368* (2017).
- [30] B. Thomson, M. Gasic, M. Henderson, P. Tsiakoulis, and S. Young. 2012. N-best error simulation for training spoken dialogue systems. In *2012 IEEE Spoken Language Technology Workshop (SLT)*. 37–42. <https://doi.org/10.1109/SLT.2012.6424194>
- [31] Gokhan Tur, Dilek Hakkani-Tür, and Larry Heck. 2010. What is left to be understood in ATIS?. In *2010 IEEE Spoken Language Technology Workshop*. IEEE, 19–24.
- [32] 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 (2018), 619–633.
- [33] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*. 5998–6008.
- [34] Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. 2015. Pointer networks. In *Advances in Neural Information Processing Systems*. 2692–2700.
- [35] Luke S Zettlemoyer and Michael Collins. 2012. Learning to map sentences to logical form: Structured classification with probabilistic categorial grammars. *arXiv preprint arXiv:1207.1420* (2012).
- [36] 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*. Association for Computational Linguistics, Florence, Italy, 5259–5267. <https://doi.org/10.18653/v1/P19-1519>