

Macro Grammars and Holistic Triggering for Efficient Semantic Parsing

Yuchen Zhang and Panupong Pasupat and Percy Liang
{zhangyuc, ppasupat, pliang}@cs.stanford.edu
Computer Science Department, Stanford University

Abstract

To learn a semantic parser from denotations, a learning algorithm must search over a combinatorially large space of logical forms for ones consistent with the annotated denotations. We propose a new online learning algorithm that searches faster as training progresses. The two key ideas are using *macro grammars* to cache the abstract patterns of useful logical forms found thus far, and *holistic triggering* to efficiently retrieve the most relevant patterns based on sentence similarity. On the WIKITABLEQUESTIONS dataset, we first expand the search space of an existing model to improve the state-of-the-art accuracy from 38.7% to 42.7%, and then use macro grammars and holistic triggering to achieve an 11x speedup and an accuracy of 43.7%.

1 Introduction

We consider the task of learning a semantic parser for question answering from question-answer pairs (Clarke et al., 2010; Liang et al., 2011; Berant et al., 2013; Artzi and Zettlemoyer, 2013; Pasupat and Liang, 2015). To train such a parser, the learning algorithm must somehow search for *consistent* logical forms (i.e., logical forms that execute to the correct answer denotation). Typically, the search space is defined by a compositional grammar over logical forms (e.g., a context-free grammar), which we will refer to as the *base grammar*.

To cover logical forms that answer complex questions, the base grammar must be quite general and compositional, leading to a huge search space that contains many useless logical forms. For example, the parser of Pasupat and Liang (2015) on

	Rank	Nation	Gold	Silver	Bronze
r_1 :	1	France	3	1	1
r_2 :	2	Ukraine	2	1	2
r_3 :	3	Turkey	2	0	1
r_4 :	4	Sweden	2	0	0
r_5 :	5	Iran	1	2	1

Table 1: A knowledge base for the question $x =$ “Who ranked right after Turkey?”. The target denotation is $y = \{\text{Sweden}\}$.

Wikipedia table questions (with beam size 100) generates and featurizes an average of 8,400 partial logical forms per example. Searching for consistent logical forms is thus a major computational bottleneck.

In this paper, we propose *macro grammars* to bias the search towards structurally sensible logical forms. To illustrate the key idea, suppose we managed to parse the utterance “Who ranked right after Turkey?” in the context of Table 1 into the following consistent logical form (in lambda DCS) (Section 2.1):

$$\mathbf{R}[\text{Nation}].\mathbf{R}[\text{Next}].\text{Nation.Turkey},$$

which identifies the cell under the Nation column in the row after Turkey. From this logical form, we can abstract out all relations and entities to produce the following *macro*:

$$\mathbf{R}\{\{Rel\#1\}\}.\mathbf{R}[\text{Next}].\{Rel\#1\}.\{Ent\#2\},$$

which represents the abstract computation: “identify the cell under the $\{Rel\#1\}$ column in the row after $\{Ent\#2\}$.” More generally, macros capture the overall shape of computations in a way that generalizes across different utterances and knowledge bases. Given the consistent logical forms of utterances parsed so far, we extract a set of macro rules. The resulting macro grammar consisting of these rules generates only logical forms conforming to these macros, which is a much smaller and

higher precision set compared to the base grammar.

Though the space of logical forms defined by the macro grammar is smaller, it is still expensive to parse with them as the number of macro rules grows with the number of training examples. To address this, we introduce *holistic triggering*: for a new utterance, we find the K most similar utterances and only use the macro rules induced from any of their consistent logical forms. Parsing now becomes efficient as only a small subset of macro rules are triggered for any utterance. Holistic triggering can be contrasted with the norm in semantic parsing, in which logical forms are either triggered by specific phrases (anchored) or can be triggered in any context (floating).

Based on the two ideas above, we propose an online algorithm for jointly inducing a macro grammar and learning the parameters of a semantic parser. For each training example, the algorithm first attempts to find consistent logical forms using holistic triggering on the current macro grammar. If it succeeds, the algorithm uses the consistent logical forms found to update model parameters. Otherwise, it applies the base grammar for a more exhaustive search to enrich the macro grammar. At test time, we only use the learned macro grammar.

We evaluate our approach on the WIKITABLE-QUESTIONS dataset (Pasupat and Liang, 2015), which features a semantic parsing task with open-domain knowledge bases and complex questions. We first extend the model in Pasupat and Liang (2015) to achieve a new state-of-the-art test accuracy of 42.7%, representing a 10% relative improvement over the best reported result (Haug et al., 2017). We then show that training with macro grammars yields an 11x speedup compared to training with only the base grammar. At test time, using the learned macro grammar achieves a slightly better accuracy of 43.7% with a 16x run time speedup over using the base grammar.

2 Background

We base our exposition on the task of question answering on a knowledge base. Given a natural language utterance x , a semantic parser maps the utterance to a logical form z . The logical form is executed on a knowledge base w to produce denotation $\llbracket z \rrbracket_w$. The goal is to train a semantic parser from a training set of utterance-denotation pairs.

2.1 Knowledge base and logical forms

A knowledge base refers to a collection of entities and relations. For the running example “*Who ranked right after Turkey?*”, we use Table 1 from Wikipedia as the knowledge base. Table cells (e.g., Turkey) and rows (e.g., r_3 = the 3rd row) are treated as entities. Relations connect entities: for example, the relation `Nation` maps r_3 to Turkey, and a special relation `Next` maps r_3 to r_4 .

A logical form z is a small program that can be executed on the knowledge base. We use lambda DCS (Liang, 2013) as the language of logical forms. The smallest units of lambda DCS are entities (e.g., Turkey) and relations (e.g., `Nation`). Larger logical forms are composed from smaller ones, and the denotation of the new logical form can be computed from denotations of its constituents. For example, applying the join operation on `Nation` and Turkey gives `Nation.Turkey`, whose denotation is $\llbracket \text{Nation.Turkey} \rrbracket_w = \{r_3\}$, which corresponds to the 3rd row of the table. The partial logical form `Nation.Turkey` can then be used to construct a larger logical form:

$$z = \mathbf{R}[\text{Nation}].\mathbf{R}[\text{Next}].\text{Nation.Turkey}, \quad (1)$$

where $\mathbf{R}[\cdot]$ represents the reverse of a relation. The denotation of the logical form z with respect to the knowledge base w is equal to $\llbracket z \rrbracket_w = \{\text{Sweden}\}$. See Liang (2013) for more details about the semantics of lambda DCS.

2.2 Grammar rules

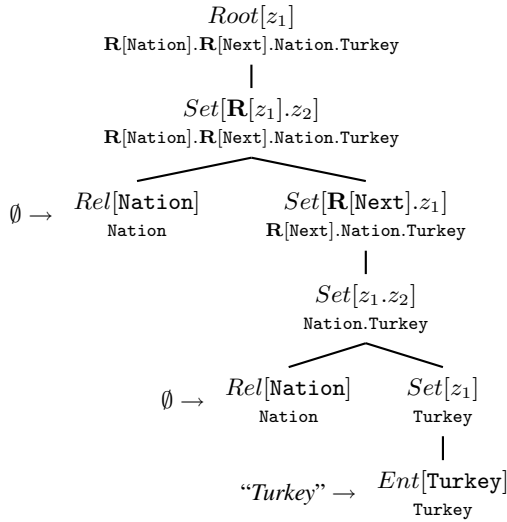
The space of logical forms is defined recursively by grammar rules. In this setting, each constructed logical form belongs to a category (e.g., *Entity*, *Rel*, *Set*), with a special category *Root* for complete logical forms. A rule specifies the categories of the arguments, category of the resulting logical form, and how the logical form is constructed from the arguments. For instance, the rule

$$\text{Rel}[z_1] + \text{Set}[z_2] \rightarrow \text{Set}[z_1.z_2] \quad (2)$$

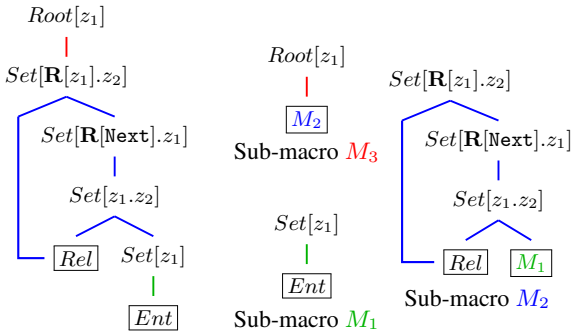
specifies that a partial logical form z_1 of category *Rel* and z_2 of category *Set* can be combined into $z_1.z_2$ of category *Set*. With this rule, we can construct `Nation.Turkey` if we have constructed `Nation` of type *Rel* and `Turkey` of type *Set*.

We consider the rules used by Pasupat and Liang (2015) for their floating parser.¹ The rules

¹Their grammar and our implementation use more fine-grained categories (*Atomic*, *Values*, *Records*) instead of *Set*. We use the coarser category here for simplicity.



(a) Derivation tree (z_i represents the i th child)



(b) Macro

(c) Atomic sub-macros

Figure 1: From the derivation tree (a), we extract a macro (b), which can be further decomposed into atomic sub-macros (c). Each sub-macro is converted into a macro rule.

are divided into *compositional rules* and *terminal rules*. Rule (2) above is an example of a compositional rule, which combines one or more partial logical forms together. A terminal rule has one of the following forms:

$$\text{TokenSpan}[span] \rightarrow c[f(span)] \quad (3)$$

$$\emptyset \rightarrow c[f(\emptyset)] \quad (4)$$

where c is a category. A rule with the form (3) converts an utterance token span (e.g., “Turkey”) into a partial logical form (e.g., Turkey). A rule with the form (4) generates a partial logical form without any trigger. This allows us to generate logical predicates that do not correspond to any part of the utterance (e.g., Nation).

A complete logical form is generated by recursively applying rules. We can represent the derivation process by a *derivation tree* such as in Fig-

ure 1a. Every node of the derivation tree corresponds to one rule. The leaf nodes correspond to terminal rules, and the intermediate nodes correspond to compositional rules.

2.3 Learning a semantic parser

Parameters of the semantic parser are learned from training data $\{(x_i, w_i, y_i)\}_{i=1}^n$. Given a training example with an utterance x , a knowledge base w , and a target denotation y , the learning algorithm constructs a set of candidate logical forms indicated by \mathcal{Z} . It then extracts a feature vector $\phi(x, w, z)$ for each $z \in \mathcal{Z}$, and defines a log-linear distribution over the candidates z :

$$p_\theta(z | x, w) \propto \exp(\theta^\top \phi(x, w, z)), \quad (5)$$

where θ is a parameter vector. The straightforward way to construct \mathcal{Z} is to enumerate all possible logical forms induced by the grammar. When the search space is prohibitively large, it is a common practice to use beam search. More precisely, the algorithm constructs partial logical forms recursively by the rules, but for each category and each search depth, it keeps only the B highest-scoring logical forms according to the model probability (5).

During training, the parameter θ is learned by maximizing the regularized log-likelihood of the correct denotations:

$$J(\theta) = \frac{1}{n} \sum_{i=1}^n \log p_\theta(y_i | x_i, w_i) - \lambda \|\theta\|_1, \quad (6)$$

where the probability $p_\theta(y_i | x_i, w_i)$ marginalizes over the space of candidate logical forms:

$$p_\theta(y_i | x_i, w_i) = \sum_{z \in \mathcal{Z}_i: \llbracket z \rrbracket_{w_i} = y_i} p_\theta(z | x_i, w_i).$$

The objective is optimized using AdaGrad (Duchi et al., 2010). At test time, the algorithm selects a logical form $z \in \mathcal{Z}$ with the highest model probability (5), and then executes it on the knowledge base w to predict the denotation $\llbracket z \rrbracket_w$.

3 Learning a macro grammar

The base grammar usually defines a large search space containing many irrelevant logical forms. For example, the grammar in Pasupat and Liang (2015) can generate long chains of join operations (e.g., $\mathbf{R}[\text{Silver}].\text{Rank}.\mathbf{R}[\text{Gold}].\text{Bronze}.2$) that rarely express meaningful computations.

Algorithm 1: Processing a training example

Data: example (x, w, y) , macro grammar, base grammar with terminal rules \mathcal{T}

- 1 Select a set \mathcal{R} of macro rules (Section 3.4);
- 2 Generate a set \mathcal{Z} of candidate logical forms from rules $\mathcal{R} \cup \mathcal{T}$ (Section 2.3);
- 3 **if** \mathcal{Z} contains consistent logical forms **then**
- 4 | Update model parameters (Section 3.5);
- 5 **else**
- 6 | Apply the base grammar to search for a consistent logical form (Section 2.3);
- 7 | Augment the macro grammar (Section 3.6);
- 8 **end**
- 9 Associate utterance x with the highest-scoring consistent logical form found;

The main contribution of this paper is a new algorithm to speed up the search based on previous searches. At a high-level, we incrementally build a *macro grammar* which encodes useful logical form macros discovered during training. Algorithm 1 describes how our learning algorithm processes each training example. It first tries to use an appropriate subset of rules in the macro grammar to search for logical forms. If the search succeeds, then the semantic parser parameters are updated as usual. Otherwise, it falls back to the base grammar, and then add new rules to the macro grammar based on the consistent logical form found. Only the macro grammar is used at test time.

We first describe macro rules and how they are generated from a consistent logical form. Then we explain the steps of the training algorithm in detail.

3.1 Logical form macros

A *macro* characterizes an abstract logical form structure. We define the macro for any given logical form z by transforming its derivation tree as illustrated in Figure 1b. First, for each terminal rule (leaf node), we substitute the rule by a placeholder, and name it with the category on the right-hand side of the rule. Then we merge leaf nodes that represent the same partial logical form. For example, the logical form (1) uses the relation `Nation` twice, so in Figure 1b, we merge the two leaf nodes to impose such a constraint.

While the resulting macro may not be tree-like, we call each node *root* or *leaf* if it is a root node or a leaf node of the associated derivation tree.

3.2 Constructing macro rules from macros

For any given macro M , we can construct a set of *macro rules* that, when combined with terminal rules from the base grammar, generates exactly the logical forms that satisfy the macro M . The straightforward approach is to associate a unique rule with each macro: assuming that its k leaf nodes contain categories c_1, \dots, c_k , we can define a rule:

$$c_1[z_1] + \dots + c_k[z_k] \rightarrow \text{Root}[f(z_1, \dots, z_k)], \quad (7)$$

where f substitutes z_1, \dots, z_k into the corresponding leaf nodes of macro M . For example, the rule for the macro in Figure 1b is

$$\text{Rel}[z_1] + \text{Ent}[z_2] \rightarrow \text{Root}[\mathbf{R}[z_1].\mathbf{R}[\text{Next}].z_1.z_2].$$

3.3 Decomposed macro rules

Defining a unique rule for each macro is computationally suboptimal since the common structures shared among macros are not being exploited. For example, while `max(R[Rank].Gold.Num.2)` and `R[Nation].argmin(Gold.Num.2, Index)` belong to different macros, the partial logical form `Gold.Num.2` is shared, and we wish to avoid generating and featurizing it more than once.

In order to reuse such shared parts, we decompose macros into *sub-macros* and define rules based on them. A subgraph M' of M is a sub-macro if (1) M' contains at least one non-leaf node; and (2) M' connects to the rest of the macro $M \setminus M'$ only through one node (the root of M'). A macro M is called *atomic* if the only sub-macro of M is itself.

Given a non-atomic macro M , we can find an atomic sub-macro M' of M . For example, from Figure 1b, we first find sub-macro $M' = M_1$. We detach M' from M and define a macro rule:

$$c'_1[z_1] + \dots + c'_k[z_k] \rightarrow c'_{\text{out}}[f(z_1, \dots, z_k)], \quad (8)$$

where c'_1, \dots, c'_k are categories of the leaf nodes of M' , and f substitutes z_1, \dots, z_k into the sub-macro M' . The category c'_{out} is computed by serializing M' as a string; this way, if the sub-macro M' appears in a different macro, the category name will be shared. Next, we substitute the subgraph M' in M by a placeholder node with name c'_{out} . The procedure is repeated on the new graph until the remaining macro is atomic. Finally, we define a single rule for the atomic macro. The

macro grammar uses the decomposed macro rules in replacement of Rule (7).

For example, the macro in Figure 1b is decomposed into three macro rules:

$$\begin{aligned} Ent[z_1] &\rightarrow M_1[z_1], \\ Rel[z_1] + M_1[z_2] &\rightarrow M_2[\mathbf{R}[z_1].\mathbf{R}[\text{Next}].z_1.z_2], \\ M_2[z_1] &\rightarrow Root[z_1]. \end{aligned}$$

These correspond to the three atomic sub-macros M_1 , M_2 and M_3 in Figure 1c. The first and the second macro rules can be reused by other macros.

Having defined macro rules, we now describe how Algorithm 1 uses and updates the macro grammar when processing each training example.

3.4 Triggering macro rules

Throughout training, we keep track of a set \mathcal{S} of training utterances that have been associated with a consistent logical form. (The set \mathcal{S} is updated by Step 9 of Algorithm 1.) Then, given a training utterance x , we compute its K -nearest neighbor utterances in \mathcal{S} , and select all macro rules that were extracted from their associated logical forms. These macro rules are used to parse utterance x .

We use token-level Levenshtein distance as the distance metric for computing nearest neighbors. More precisely, every utterance is written as a sequence of lemmatized tokens $x = (x^{(1)}, \dots, x^{(m)})$. After removing all determiners and infrequent nouns that appear in less than 2% of the training utterances, the distance between two utterances x and x' is defined as the Levenshtein distance between the two sequences. When computing the distance, we treat each word token as an atomic element. For example, the distance between “highest score” and “best score” is 1. Despite its simplicity, the Levenshtein distance does a good job in capturing the structural similarity between utterances. Table 2 shows that nearest neighbor utterances often map to consistent logical forms with the same macro.

In order to compute the nearest neighbors efficiently, we pre-compute a sorted list of $K_{\max} = 100$ nearest neighbors for every utterance before training starts. During training, calculating the intersection of this sorted list with the set \mathcal{S} gives the nearest neighbors required. For our experiments, the preprocessing time is negligible compared to the overall training time (less than 3%), but if computing nearest neighbors is expensive, then paral-

Who ranked right after Turkey?
Who took office right after Uriah Forrest?
How many more passengers flew to Los Angeles than to Saskatoon in 2013?
How many more Hungarians live in the Serbian Banat region than Romanians in 1910?
Which is deeper, Lake Tuz or Lake Palas Tuzla?
Which peak is higher, Mont Blanc or Monte Rosa?

Table 2: Examples of nearest neighbor utterances in the WIKITABLEQUESTIONS dataset.

lization or approximate algorithms (e.g., Indyk, 2004) could be used.

3.5 Updating model parameters

Having computed the triggered macro rules \mathcal{R} , we combine them with the terminal rules \mathcal{T} from the base grammar (e.g., for building Ent and Rel) to create a per-example grammar $\mathcal{R} \cup \mathcal{T}$ for the utterance x . We use this grammar to generate logical forms using standard beam search. We follow Section 2.3 to generate a set of candidate logical forms \mathcal{Z} and update model parameters.

However, we deviate from Section 2.3 in one way. Given a set \mathcal{Z} of candidate logical forms for some training example (x_i, w_i, y_i) , we pick the logical form z_i^+ with the highest model probability among consistent logical forms, and pick z_i^- with the highest model probability among inconsistent logical forms, then perform a gradient update on the objective function:

$$\begin{aligned} J(\theta) &= \frac{1}{n} \sum_{i=1}^n \left[\log \frac{p_i^+}{p_i^+ + p_i^-} \right] - \lambda \|\theta\|_1, \quad (9) \\ \text{where } p_i^+ &= p_\theta(z_i^+ | x_i, w_i) \\ p_i^- &= p_\theta(z_i^- | x_i, w_i). \end{aligned}$$

Compared to (6), this objective function only considers the top consistent and inconsistent logical forms for each example instead of all candidate logical forms. Empirically, we found that optimizing (9) gives a 2% gain in prediction accuracy compared to optimizing (6).

3.6 Updating the macro grammar

If the triggered macro rules fail to find a consistent logical form, we fall back to performing a beam search on the base grammar. For efficiency, we stop the search either when a consistent logical form is found, or when the total number of generated logical forms exceeds a threshold T . The two

stopping criteria prevent the search algorithm from spending too much time on a complex example. We might miss consistent logical forms on such examples, but because the base grammar is only used for generating macro rules, not for updating model parameters, we might be able to induce the same macro rules from other examples. For instance, if an example has an utterance phrase that matches too many knowledge base entries, it would be more efficient to skip the example; the macro that would have been extracted from this example can be extracted from less ambiguous examples with the same question type. Such omissions are not completely disastrous, and can speed up training significantly.

When the algorithm succeeds in finding a consistent logical form z using the base grammar, we derive its macro M following Section 3.1, then construct macro rules following Section 3.3. These macro rules are added to the macro grammar. We also associate the utterance x with the consistent logical form z , so that the macro rules that generate z can be triggered by other examples. Parameters of the semantic parser are not updated in this case.

3.7 Prediction

At test time, we follow Steps 1–2 of Algorithm 1 to generate a set \mathcal{Z} of candidate logical forms from the triggered macro rules, and then output the highest-scoring logical form in \mathcal{Z} . Since the base grammar is never used at test time, prediction is generally faster than training.

4 Experiments

We report experiments on the WIKITABLEQUESTIONS dataset (Pasupat and Liang, 2015). Our algorithm is compared with the parser trained only with the base grammar, the floating parser of Pasupat and Liang (2015) (PL15), the Neural Programmer parser (Neelakantan et al., 2016) and the Neural Multi-Step Reasoning parser (Haug et al., 2017). Our algorithm not only outperforms the others, but also achieves an order-of-magnitude speedup over the parser trained with the base grammar and the parser in PL15.

4.1 Setup

The dataset contains 22,033 complex questions on 2,108 Wikipedia tables. Each question comes with a table, and the tables during evaluation are dis-

“Which driver appears the most?” $\text{argmax}(\mathbf{R}[\text{Driver}].\text{Type}.\text{Row}, \mathbf{R}[\lambda x.\text{count}(\text{Driver}.x)])$
“What language was spoken more during the Olympic oath, English or French?” $\text{argmax}(\text{English} \sqcup \text{French}, \mathbf{R}[\lambda x.\text{count}(\text{Language}.x)])$
“Who is taller, Rose or Tim?” $\text{argmax}(\text{Rose} \sqcup \text{Tim}, \mathbf{R}[\lambda x.\mathbf{R}[\text{Num}].\mathbf{R}[\text{Height}].\text{Name}.x])$

Table 3: Several example logical forms our grammar can generate that are not covered by PL15.

joint from the ones during training. The training and test sets contain 14,152 and 4,344 examples respectively.² Following PL15, the development accuracy is averaged over the first three 80-20 training data splits given in the dataset package. The test accuracy is reported on the train-test data split.

We use the same features and logical form pruning strategies as PL15, but generalize their base grammar. To control the search space, the actual system in PL15 restricts the superlative operators argmax and argmin to be applied only on the set of table rows. We allow these operators to be applied on the set of tables cells as well, so that the grammar captures certain logical forms that are not covered by PL15 (see Table 3). Additionally, for terminal rule (3), we allow $f(\text{span})$ to produce entities that approximately match the token span in addition to exact matches. For example, the phrase “Greenville” can trigger both entities `Greenville_Ohio` and `Greensville`.

We chose hyperparameters using the first train-dev split. The beam size B of beam search is chosen to be $B = 100$. The K -nearest neighbor parameter is chosen as $K = 40$. Like PL15, our algorithm takes 3 passes over the dataset for training. The maximum number of logical forms generated in step 6 of Algorithm 1 is set to $T = 5,000$ for the first pass. For subsequent passes, we set $T = 0$ (i.e., never fall back to the base grammar) so that we stop augmenting the macro grammar. During the first pass, Algorithm 1 falls back to the base grammar on roughly 30% of the training examples.

For training the baseline parser that only relies on the base grammar, we use the same beam size $B = 100$, and take 3 passes over the dataset for training. There is no maximum constraint on the

²The remaining 3,537 examples were not included in the original data split.

	Dev	Test
Pasupat and Liang (2015)	37.0%	37.1%
Neelakantan et al. (2016)	37.5%	37.7%
Haug et al. (2017)	-	38.7%
This paper: base grammar	40.6%	42.7%
This paper: macro grammar	40.4%	43.7%

Table 4: Results on WIKITABLEQUESTIONS.

number of logical forms that can be generated for each example.

4.2 Coverage of the macro grammar

With the base grammar, our parser generates 13,700 partial logical forms on average for each training example, and hits consistent logical forms on 81.0% of the training examples. With the macro rules from holistic triggering, these numbers become 1,300 and 75.6%. The macro rules generate much fewer partial logical forms, but at the cost of slightly lower coverage.

However, these coverage numbers are computed based on finding any logical form that executes to the correct denotation. This includes spurious logical forms, which do not reflect the semantics of the question but are coincidentally consistent with the correct denotation. (For example, the question “*Who got the same number of silvers as France?*” on Table 1 might be spuriously parsed as $R[Nation].R[Next].Nation.France$, which represents the nation listed after France.) To evaluate the “true” coverage, we sample 300 training examples and manually label their logical forms. We find that on 48.7% of these examples, the top consistent logical form produced by the base grammar is semantically correct. For the macro grammar, this ratio is also 48.7%, meaning that the macro grammar’s effective coverage is as good as the base grammar.

The macro grammar extracts 123 macros in total. Among the 75.6% examples that were covered by the macro grammar, the top 34 macros cover 90% of consistent logical forms. By examining the top 34 macros, we discover explicit semantic meanings for 29 of them, which are described in detail in the supplementary material.

4.3 Accuracy and speedup

We report prediction accuracies in Table 4. With a more general base grammar (additional superlatives and approximate matching), and by optimiz-

	Acc.	Time (ms/ex)	
		Train	Pred
PL15	37.0%	619	645
Ours: base grammar	40.6%	1,117	1,150
Ours: macro grammar	40.4%	99	70
no holistic triggering	40.1%	361	369
no macro decomposition	40.3%	177	159

Table 5: Comparison and ablation study: the columns report averaged prediction accuracy, training time, and prediction time (milliseconds per example) on the three train-dev splits.

ing the objective function (9), our base parser outperforms PL15 (42.7% vs 37.1%). Learning a macro grammar slightly improves the accuracy to 43.7% on the test set. On the three train-dev splits, the averaged accuracy achieved by the base grammar and the macro grammar are close (40.6% vs 40.4%).

In Table 5, we compare the training and prediction time of PL15 as well as our parsers. For a fair comparison, we trained all parsers using the SEMPRES toolkit (Berant et al., 2013) on a machine with Xeon 2.6GHz CPU and 128GB memory without parallelization. The time for constructing the macro grammar is included as part of the training time. Table 5 shows that our parser with the base grammar is more expensive to train than PL15. However, training with the macro grammar is substantially more efficient than training with only the base grammar—it achieves 11x speedup for training and 16x speedup for test time prediction.

We run two ablations of our algorithm to evaluate the utility of holistic triggering and macro decomposition. The first ablation triggers all macro rules for parsing every utterance without holistic triggering, while the second ablation constructs Rule (7) for every macro without decomposing it into smaller rules. Table 5 shows that both variants result in decreased efficiency. This is because holistic triggering effectively prunes irrelevant macro rules, while macro decomposition is important for efficient beam search and featurization.

4.4 Influence of hyperparameters

Figure 2a shows that for all beam sizes, training with the macro grammar is more efficient than training with the base grammar, and the speedup rate grows with the beam size. The test time ac-

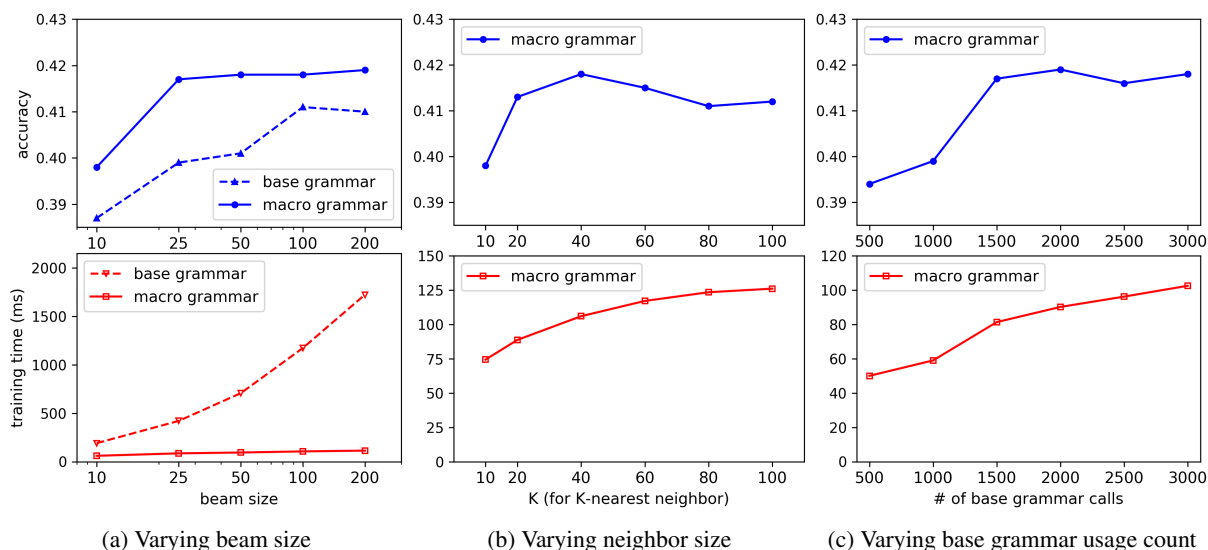


Figure 2: Prediction accuracy and training time (per example) with various hyperparameter choices, reported on the first train-dev split.

accuracy of the macro grammar is robust to varying beam sizes as long as $B \geq 25$.

Figure 2b shows the influence of the neighbor size K . A smaller neighborhood triggers fewer macro rules, leading to faster computation. The accuracy peaks at $K = 40$ then decreases slightly for large K . We conjecture that the smaller number of neighbors acts as a regularizer.

Figure 2c reports an experiment where we limit the number of fallback calls to the base grammar to m . After the limit is reached, subsequent training examples that require fallback calls are simply skipped. This limit means that the macro grammar will get augmented at most m times during training. We find that for small m , the prediction accuracy grows with m , implying that building a richer macro grammar improves the accuracy. For larger m , however, the accuracies hardly change. According to the plot, a competitive macro grammar can be built by calling the base grammar on less than 15% of the training data.

Based on Figure 2, we can trade accuracy for speed by choosing smaller values of (B, K, m) . With $B = 50$, $K = 40$ and $m = 2000$, the macro grammar achieves a slightly lower averaged development accuracy (40.2% rather than 40.4%), but with an increased speedup of 15x (versus 11x) for training and 20x (versus 16x) for prediction.

5 Related work and discussion

A traditional semantic parser maps natural language phrases into partial logical forms and

composes these partial logical forms into complete logical forms. Parsers define composition based on a grammar formalism such as Combinatory Categorical Grammar (CCG) (Zettlemoyer and Collins, 2007; Kwiatkowski et al., 2011, 2013; Kushman and Barzilay, 2013; Krishnamurthy and Kollar, 2013), Synchronous CFG (Wong and Mooney, 2007), and CFG (Kate and Mooney, 2006; Chen and Mooney, 2011; Berant et al., 2013; Desai et al., 2016), while others use the syntactic structure of the utterance to guide composition (Poon and Domingos, 2009; Reddy et al., 2016). Recent neural semantic parsers allow any sequence of logical tokens to be generated (Dong and Lapata, 2016; Jia and Liang, 2016; Kociský et al., 2016; Neelakantan et al., 2016; Liang et al., 2017; Guu et al., 2017). The flexibility of these composition methods allows arbitrary logical forms to be generated, but at the cost of a vastly increased search space.

Whether we have annotated logical forms or not has dramatic implications on what type of approach will work. When logical forms are available, one can perform grammar induction to mine grammar rules *without search* (Kwiatkowski et al., 2010). When only annotated denotations are available, as in our setting, one must use a base grammar to define the output space of logical forms. Usually these base grammars come with many restrictions to guard against combinatorial explosion (Pasupat and Liang, 2015).

Previous work on higher-order unification for

lexicon induction (Kwiatkowski et al., 2010) using factored lexicons (Kwiatkowski et al., 2011) also learns logical form macros with an online algorithm. The result is a lexicon where each entry contains a logical form template and a set of possible phrases for triggering the template. In contrast, we have avoided binding grammar rules to particular phrases in order to handle lexical variations. Instead, we use a more flexible mechanism—holistic triggering—to determine which rules to fire. This allows us to generate logical forms for utterances containing unseen lexical paraphrases or where the triggering is spread throughout the sentence. For example, the question “Who is X, John or Y” can still trigger the correct macro extracted from the last example in Table 3 even when X and Y are unknown words.

Our macro grammars bears some resemblance to adaptor grammars (Johnson et al., 2006) and fragment grammars (O’Donnell, 2011), which are also based on the idea of caching useful chunks of outputs. These generative approaches aim to solve the modeling problem of assigning higher probability mass to outputs that use reoccurring parts. In contrast, our learning algorithm uses caching as a way to constrain the search space for computational efficiency; the probabilities of the candidate outputs are assigned by a separate discriminative model. That said, the use of macro grammars does have a small positive modeling contribution, as it increases test accuracy from 42.7% to 43.7%.

An orthogonal approach for improving search efficiency is to adaptively choose which part of the search space to explore. For example, Berant and Liang (2015) uses imitation learning to strategically search for logical forms. Our holistic triggering method, which selects macro rules based on the similarity of input utterances, is related to the use of paraphrases (Berant and Liang, 2014; Fader et al., 2013) or string kernels (Kate and Mooney, 2006) to train semantic parsers. While the input similarity measure is critical for scoring logical forms in these previous works, we use the measure only to retrieve candidate rules, while scoring is done by a separate model. The retrieval bar means that our similarity metric can be quite crude.

6 Summary

We have presented a method for speeding up semantic parsing via macro grammars. The main source of efficiency is the decreased size of the

logical form space. By performing beam search on a few macro rules associated with the K -nearest neighbor utterances via holistic triggering, we have restricted the search space to semantically relevant logical forms. At the same time, we still maintain coverage over the base logical form space by occasionally falling back to the base grammar and using the consistent logical forms found to enrich the macro grammar. The higher efficiency allows us expand the base grammar without having to worry much about speed: our model achieves a state-of-the-art accuracy while also enjoying an order magnitude speedup.

Acknowledgements. We gratefully acknowledge Tencent for their support on this project.

Reproducibility. Code, data, and experiments for this paper are available on CodaLab platform: <https://worksheets.codalab.org/worksheets/0x4d6dbfc5ec7f44a6a4da4ca2a9334d6e/>.

References

- Y. Artzi and L. Zettlemoyer. 2013. UW SPF: The University of Washington semantic parsing framework. *arXiv preprint arXiv:1311.3011*.
- J. Berant, A. Chou, R. Frostig, and P. Liang. 2013. Semantic parsing on Freebase from question-answer pairs. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- J. Berant and P. Liang. 2014. Semantic parsing via paraphrasing. In *Association for Computational Linguistics (ACL)*.
- J. Berant and P. Liang. 2015. Imitation learning of agenda-based semantic parsers. *Transactions of the Association for Computational Linguistics (TACL)*, 3:545–558.
- D. L. Chen and R. J. Mooney. 2011. Learning to interpret natural language navigation instructions from observations. In *Association for the Advancement of Artificial Intelligence (AAAI)*, pages 859–865.
- J. Clarke, D. Goldwasser, M. Chang, and D. Roth. 2010. Driving semantic parsing from the world’s response. In *Computational Natural Language Learning (CoNLL)*, pages 18–27.
- A. Desai, S. Gulwani, V. Hingorani, N. Jain, A. Karkare, M. Marron, S. R., and S. Roy. 2016. Program synthesis using natural language. In *International Conference on Software Engineering (ICSE)*, pages 345–356.
- L. Dong and M. Lapata. 2016. Language to logical form with neural attention. In *Association for Computational Linguistics (ACL)*.

- J. Duchi, E. Hazan, and Y. Singer. 2010. Adaptive sub-gradient methods for online learning and stochastic optimization. In *Conference on Learning Theory (COLT)*.
- A. Fader, L. Zettlemoyer, and O. Etzioni. 2013. Paraphrase-driven learning for open question answering. In *Association for Computational Linguistics (ACL)*.
- K. Guu, P. Pasupat, E. Z. Liu, and P. Liang. 2017. From language to programs: Bridging reinforcement learning and maximum marginal likelihood. In *Association for Computational Linguistics (ACL)*.
- T. Haug, O. Ganea, and P. Grnarova. 2017. Neural multi-step reasoning for question answering on semi-structured tables. *arXiv preprint arXiv:1702.06589*.
- P. Indyk. 2004. Approximate nearest neighbor under edit distance via product metrics. In *Symposium on Discrete Algorithms (SODA)*, pages 646–650.
- R. Jia and P. Liang. 2016. Data recombination for neural semantic parsing. In *Association for Computational Linguistics (ACL)*.
- M. Johnson, T. Griffiths, and S. Goldwater. 2006. Adaptor grammars: A framework for specifying compositional nonparametric Bayesian models. In *Advances in Neural Information Processing Systems (NIPS)*, pages 641–648.
- R. J. Kate and R. J. Mooney. 2006. Using string-kernels for learning semantic parsers. In *International Conference on Computational Linguistics and Association for Computational Linguistics (COLING/ACL)*, pages 913–920.
- T. Kociský, G. Melis, E. Grefenstette, C. Dyer, W. Ling, P. Blunsom, and K. M. Hermann. 2016. Semantic parsing with semi-supervised sequential autoencoders. In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 1078–1087.
- J. Krishnamurthy and T. Kollar. 2013. Jointly learning to parse and perceive: Connecting natural language to the physical world. *Transactions of the Association for Computational Linguistics (TACL)*, 1:193–206.
- N. Kushman and R. Barzilay. 2013. Using semantic unification to generate regular expressions from natural language. In *Human Language Technology and North American Association for Computational Linguistics (HLT/NAACL)*, pages 826–836.
- T. Kwiatkowski, E. Choi, Y. Artzi, and L. Zettlemoyer. 2013. Scaling semantic parsers with on-the-fly ontology matching. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- T. Kwiatkowski, L. Zettlemoyer, S. Goldwater, and M. Steedman. 2010. Inducing probabilistic CCG grammars from logical form with higher-order unification. In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 1223–1233.
- T. Kwiatkowski, L. Zettlemoyer, S. Goldwater, and M. Steedman. 2011. Lexical generalization in CCG grammar induction for semantic parsing. In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 1512–1523.
- C. Liang, J. Berant, Q. Le, and K. D. F. N. Lao. 2017. Neural symbolic machines: Learning semantic parsers on Freebase with weak supervision. In *Association for Computational Linguistics (ACL)*.
- P. Liang. 2013. Lambda dependency-based compositional semantics. *arXiv preprint arXiv:1309.4408*.
- P. Liang, M. I. Jordan, and D. Klein. 2011. Learning dependency-based compositional semantics. In *Association for Computational Linguistics (ACL)*, pages 590–599.
- A. Neelakantan, Q. V. Le, and I. Sutskever. 2016. Neural programmer: Inducing latent programs with gradient descent. In *International Conference on Learning Representations (ICLR)*.
- T. J. O’Donnell. 2011. *Productivity and Reuse in Language*. Ph.D. thesis, Massachusetts Institute of Technology.
- P. Pasupat and P. Liang. 2015. Compositional semantic parsing on semi-structured tables. In *Association for Computational Linguistics (ACL)*.
- H. Poon and P. Domingos. 2009. Unsupervised semantic parsing. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- S. Reddy, O. Täckström, M. Collins, T. Kwiatkowski, D. Das, M. Steedman, and M. Lapata. 2016. Transforming dependency structures to logical forms for semantic parsing. In *Association for Computational Linguistics (ACL)*, pages 127–140.
- Y. W. Wong and R. J. Mooney. 2007. Learning synchronous grammars for semantic parsing with lambda calculus. In *Association for Computational Linguistics (ACL)*, pages 960–967.
- L. S. Zettlemoyer and M. Collins. 2007. Online learning of relaxed CCG grammars for parsing to logical form. In *Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP/CoNLL)*, pages 678–687.