

Decoupling Structure and Lexicon for Zero-Shot Semantic Parsing

Jonathan Herzig¹ and Jonathan Berant^{1,2}

¹Tel-Aviv University

²Allen Institute for Artificial Intelligence

jonathan.herzig@cs.tau.ac.il, joberant@cs.tau.ac.il

Abstract

Building a semantic parser quickly in a new domain is a fundamental challenge for conversational interfaces, as current semantic parsers require expensive supervision and lack the ability to generalize to new domains. In this paper, we introduce a zero-shot approach to semantic parsing that can parse utterances in unseen domains while only being trained on examples in other source domains. First, we map an utterance to an abstract, domain-independent, logical form that represents the structure of the logical form, but contains slots instead of KB constants. Then, we replace slots with KB constants via lexical alignment scores and global inference. Our model reaches an average accuracy of 53.4% on 7 domains in the OVERNIGHT dataset, substantially better than other zero-shot baselines, and performs as good as a parser trained on over 30% of the target domain examples.

1 Introduction

Semantic parsing, the task of mapping natural language utterances into executable logical forms, is a key paradigm in developing conversational interfaces (Zelle and Mooney, 1996; Zettlemoyer and Collins, 2005; Kwiatkowski et al., 2011; Berant and Liang, 2015). The recent success of conversational interfaces such as Amazon Alexa, Google Assistant, Apple Siri, and Microsoft Cortana has led to soaring interest in developing methodologies for training semantic parsers quickly in any new domain and from little data.

Prior work focused on alleviating data collection by training from weak supervision (Clarke et al., 2010; Liang et al., 2011; Kwiatkowski et al., 2013; Artzi and Zettlemoyer, 2013), or developing protocols for fast data collection through paraphrasing (Berant and Liang, 2014; Wang et al., 2015) or a human-in-the-loop (Iyer et al., 2017).




Figure 1: A test utterance is delexicalized (1) and mapped to its abstract logical form (2). Slots (“\$” variables) are then aligned to the abstract utterance (3), and are filled with the top assignment in terms of local and global scores (4). Logical forms throughout this paper are in λ -DCS (Liang, 2013).

However, all these approaches rely on supervised training data in the target domain and ignore data collected previously for other domains.

In this paper, we propose an alternative, zero-shot approach to semantic parsing, where no labeled or unlabeled examples are provided in the target domain, but annotated examples from other domains are available. This is a challenging setup as in semantic parsing each dataset is associated with its own knowledge-base (KB) and thus all target domain KB constants (relations and entities) are unobserved at training time. Moreover, this is a natural use-case as more and more conversational interfaces are developed in multiple domains.

Our approach is motivated by recent work (Herzig and Berant, 2017; Su and Yan, 2017; Fan et al., 2017; Richardson et al., 2018) that showed that while the lexicon and KB constants in dif-

ferent domains vary, the structure of language composition repeats across domains. Therefore, we propose that by abstracting away the domain-specific lexical items of an utterance, we can learn to map the structure of an abstract utterance to an abstract logical form that does not include any domain-specific KB constants, using data from other domains only.

Figure 1 illustrates this approach. A test utterance in the target domain is delexicalized and mapped to an abstract, domain-independent representation, where some content words are replaced by abstract tokens (step 1). Then, a structure-mapping model maps this representation into an abstract logical form that contains slots instead of KB constants (step 2). A major technical challenge at this point is to replace slots in the abstract logical form with KB constants from the target domain. We show that it is possible to learn a domain-independent lexical alignment model that aligns each slot to a word in the original utterance (step 3). This alignment, combined with a global inference procedure (step 4) allows one to find the best assignment of KB constants and produce a final logical form. Importantly, both of our models are trained from data in *other* domains only.

We show that our zero-shot framework parses 7 different unseen domains from the OVERNIGHT dataset with an average denotation accuracy of 53.4%. This result dramatically outperforms several natural baselines, and achieves the same result as training a parser on over 30% of the fully supervised target domain examples. To our knowledge, this work is the first to train a zero-shot semantic parser that can handle unseen domains. All our code is available at <https://github.com/jonathanherzig/zero-shot-semantic-parsing>.

2 Background

Neural Semantic Parsing Sequence-to-sequence models (Sutskever et al., 2014) were recently proposed for semantic parsing (Jia and Liang, 2016; Dong and Lapata, 2016). In this setting, a sequence of input language tokens x_1, \dots, x_m is mapped to a sequence of output logical tokens z_1, \dots, z_n . We briefly review the model by Jia and Liang (2016), which we use as part of our framework, and also as a baseline.

The **encoder** is a BiLSTM (Hochreiter and Schmidhuber, 1997) that converts x_1, \dots, x_m into

a sequence of context sensitive states. The attention-based **decoder** (Bahdanau et al., 2015; Luong et al., 2015) is an LSTM language model additionally conditioned on the encoder states. Formally, the decoder is defined by:

$$p(z_j = w \mid x, z_{1:j-1}) \propto \exp(U[s_j, c_j]),$$

$$s_{j+1} = LSTM([\phi^{(out)}(z_j), c_j], s_j),$$

where s_j are decoder states, U and the embedding function $\phi^{(out)}$ are the decoder parameters, and the context vector, c_j , is the result of global attention (Luong et al., 2015). We also employ attention-based copying (Jia and Liang, 2016), but omit details for brevity.

Semantic Parsing over Multiple KBs Recently, Herzig and Berant (2017), Su and Yan (2017) and Fan et al. (2017) proposed to exploit structural regularities in language across different domains. These works pooled together examples from multiple datasets in different domains, each corresponding to a separate KB, and trained a single sequence-to-sequence model over all examples, sharing parameters across domains. They showed that this substantially improves parsing accuracy. While these works implicitly capture linguistic regularities across domains, they rely on annotated data in the target domain. We, conversely, explicitly decouple structure mapping from the assignment of KB constants, and thus can tackle the zero-shot setting where no target domain examples are available. This is the focus of the next section.

3 Zero-Shot Semantic Parsing

3.1 Overview

Following the empirical success of sharing structural information between different semantic parsing domains, we propose in this paper to take a more radical approach and to explicitly decouple semantic parsing into a structure mapping model and a lexicon mapping model. We now provide an overview of our approach and explain how this decoupling facilitates zero-shot semantic parsing.

We assume access to D different source domains, where for every domain d we receive a KB \mathcal{K}_d , and a training set of pairs of utterances and logical forms $\{(x_i, z_i)\}_{i=1}^{N_d}$. We further assume a lexicon L that maps each KB constant in \mathcal{K}_d to a short phrase that describes it (e.g., $L(\text{PubYear}) \rightarrow \text{“publication year”}$), as in Wang

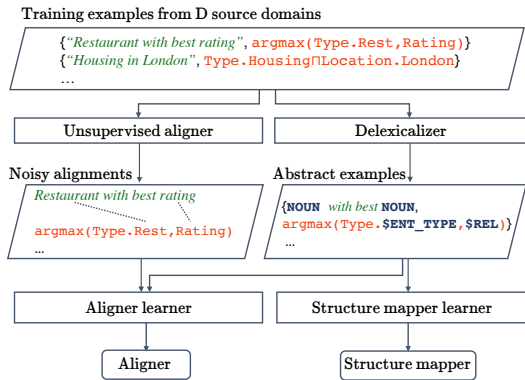


Figure 2: Training flow. Examples in source domains are delexicalized. Abstract examples are used to train both the structure mapper learner and aligner learner, where the aligner learner uses noisy alignments as labels.

et al. (2015). Finally, we assume a pre-trained, static, embedding function $\phi(w) \in \mathbb{R}^f$ for every word w , used to measure cross-domain lexical semantic similarity. Our goal is to train a semantic parser that maps a new utterance x to the correct logical form z from a new domain d_{new} given $\mathcal{K}_{d_{\text{new}}}$.

Figure 2 describes the flow of our training procedure: we first employ a simple rule-based method to transform training examples to an abstract representation, where content words (in utterances) and KB constants (in logical forms) are delexicalized. We then train the following two models that decouple structure from lexicon (a) The structure mapper that maps abstract utterances to abstract logical forms. (b) The aligner that provides an alignment from abstract logical form tokens to abstract utterance tokens. Training the aligner is challenging because no gold alignments between the abstract utterance and abstract logical form are available. To overcome this challenge we propose a distillation strategy: we obtain noisy supervision by training a state-of-the-art unsupervised alignment model on the D source domains. Then, we train a second supervised alignment model that receives abstract utterances, abstract logical forms, and target noisy alignments as input and learns to predict the noisy alignments.

Once the two models are trained, we can tackle a new domain without training examples (Figure 1). Given an utterance from the target domain, we first abstract it using the delexicalizer, and then predict its abstract structure using the structure

Lexical representation

“What meetings have no more than 3 attendees?”
 $\text{Type.Meeting} \sqcap \mathbf{R}[\lambda x.\text{count}(\text{Attendee}.x)]. \leq .3$
 “Which recipe needs no more than two ingredients?”
 $\text{Type.Recipe} \sqcap \mathbf{R}[\lambda x.\text{count}(\text{IngredientOf}.x)]. \leq .2$

Abstract representation

“What NOUN have no more than NUM NOUN?”
 $\text{Type}.\$SENT_TYPE \sqcap \mathbf{R}[\lambda x.\text{count}(\$REL.x)]. \leq .\$NUM$
 “Which NOUN VERB no more than NUM NOUN.”
 $\text{Type}.\$SENT_TYPE \sqcap \mathbf{R}[\lambda x.\text{count}(\$REL.x)]. \leq .\$NUM$

Figure 3: Examples in different domains (CALENDAR and RECIPES) in their original and abstract representations. A similar structural regularity (a comparative structure) maps to an identical abstract logical form.

mapper. We treat delexicalized logical form tokens as slots to be filled with KB constants. Candidate assignments are then scored locally according to the semantic similarity of a KB constant (represented by its entry value in the lexicon L) to words the slot aligns to according to the aligner. For this we use the pre-trained embedding function $\phi(\cdot)$ as the only cross-domain information. Finally, we choose a final assignment of KB constants by exactly maximizing a global scoring function, which takes into account both local alignment scores as well as global constraints.

We next describe in detail the four components of our framework: the delexicalizer, structure mapper, aligner, and inference procedure.

3.2 Delexicalizer

The goal of the delexicalizer is to strip utterances and logical forms from their domain-specific components and preserve domain-independent parts. We note that it is possible that some words contain both domain-specific and domain-general aspects (“cheapest”). However, we conjecture that it is possible to decompose examples in a manner that enables zero-shot semantic parsing.

The output of the delexicalizer is an abstract representation that should manifest structural linguistic regularities across domains (Figure 3). For example, a comparative structure will correspond to the same abstract logical form in different domains. In this representation, used as input to our models, content words and KB constants are transformed to an abstract type. This rule-based pre-processing step is applied to all D source domain training examples (utterances and logical forms), and to target domain utterances at test time. We now describe the process of delexicalization.

Source	Category	Abstract Type	Examples
Utterance	Noun	NOUN	"cuisines", "housing", "time"
	Verb	VERB	"published", "born", "posted"
	Adjective	ADJ	"high", "cooking", "monthly"
	Number	NUM	"4", "three"
	Date	DATE	"2018", "january 2nd"
	Entity	ENT	"midtown", "alice", "dinner"
Logical Form	Number	\$NUM	1, 2, 3
	Date	\$DATE	1.6.2018
	Entity	\$ENT	MidtownWest, CentralOffice
	Entity type	\$SENT_TYPE	Person, Location, Recipe
	Numerical entity	\$SENT_NUM	Rent, Size, CookingTime
	Binary relation	\$REL	Author, Attendee
	Unary relation	\$REL_UNARY	AllowsCats, WonAward
	Numerical relation	\$REL_NUM	Height, Length, StarRating
	Date relation	\$REL_DATE	PostingDate, PubYear

Table 1: Categories of content words and KB constants, and their corresponding abstract type notation.

Utterances Table 1 describes the full list of abstraction rules. We delexicalize several categories of content words and keep function words, which describe the utterance structure, in their lexicalized form. Specifically, any verb¹ whose lemma is not “be” or “do” is delexicalized. All nouns are delexicalized, except for a small vocabulary of three words (“average”, “total”, and “number”), which denote a domain-general operation. Adjectives tend to distribute more evenly between domain-specific words and domain-general words, thus discriminating them is harder (e.g., “outdoor”, “wide” and “cooking” are domain-specific words while “minimum”, “same” and “many” are domain-general words). Thus, we take a statistical approach and only delexicalize adjectives that are unique to the domain (i.e., did not appear in the training set of any other source domain). We also delexicalize dates and numbers, and identify entities in the utterance by string matching against the entities in the KB. These are then delexicalized to their corresponding abstract type (Table 1).

Logical Forms We delexicalize all KB constants to their abstract type, which is given as part of the KB schema (Table 1).

3.3 Structure Mapper

As a first step towards predicting the lexical logical form, we map an abstract utterance, to an abstract logical form. The model is the neural semantic parser described in Section 2, only here the input and output are the abstract examples in all D domains, which the delexicalizer outputs. The model utilizes a single encoder-decoder pair shared across all domains. As Figure 3 suggests, the model should learn, e.g., that a noun modified

¹Numbers, dates and part-of-speech tags are extracted using Stanford CoreNLP (Manning et al., 2014).

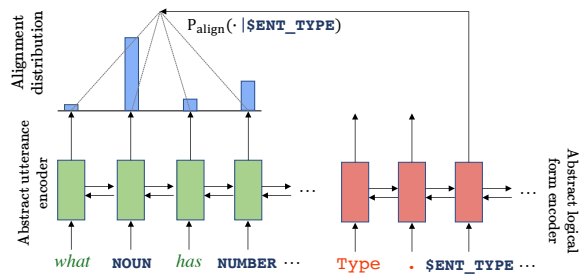


Figure 4: The aligner model. Alignments are derived by comparing the slot hidden state against all utterance hidden states.

by a wh-question often maps to $\$SENT_TYPE$, and that “no more than” maps to the \leq operator.

3.4 Aligner

The output of the structure mapper is an abstract logical form that contains slots instead of KB constants. To predict a complete logical form, we must assign a KB constant to each slot.

We observe that the description of a KB constant that appears in the logical form (`Article`) is often semantically similar to some word in the utterance (“paper”). Thus, we can obtain signal for the identity of a KB constant by solving an alignment problem: each slot can be aligned to words in the utterance that have similar meaning to that of the gold KB constant. Naturally, in some cases a KB constant is not semantically similar to any utterance word (e.g., the relation `Field` in Figure 1), which we will mediate by using a global inference procedure (Section 3.5).

Thus, our goal is to learn a model that given an abstract utterance-logical form pair (x^{abs}, z^{abs}) produces an alignment matrix A , where A_{ij} corresponds to the alignment probability $p(x_j^{abs} | z_i^{abs})$. A central challenge is that no gold alignments are provided in any domain. Therefore, we adopt a “distillation approach”, where we train a supervised model over abstract examples to mimic the predictions of an unsupervised model that has access to the full lexicalized examples.

Specifically, we use a standard unsupervised word aligner (Dyer et al., 2013), which takes all lexicalized examples $\{(x_i, z_i)\}_{i=1}^{N_d}$ in all D domains and produces an Alignment matrix A^* for every example, where $A_{ij}^* = 1$ iff token i in the logical form is aligned to token j in the utterance. Then, we treat A^* as gold alignments and gener-

ate examples $(x^{\text{abs}}, z^{\text{abs}}, A^*)$ to train the aligner. Learning alignments over abstract representations is possible, as a slot in a specific context tends to align to specific types of abstract words (e.g., Figure 3 suggests that a relation that is aggregated, often aligns to the NOUN that appears after the NUM in the abstract utterance).

We now present our alignment model, depicted in Figure 4. The model uses two different BiLSTMs to encode x^{abs} and z^{abs} to their context sensitive states b_1, \dots, b_m and s_1, \dots, s_n respectively. We model the alignment probability $p_{\text{align}}(x_j^{\text{abs}} | z_i^{\text{abs}})$ with a bi-linear form similar to attention (Luong et al., 2015):

$$e_{ij} = s_i^T W b_j,$$

$$p_{\text{align}}(x_j^{\text{abs}} | z_i^{\text{abs}}) = \frac{\exp(e_{ij})}{\sum_{j'=1}^m \exp(e_{ij'})},$$

where the parameters W are learned during training. We train the model to minimize the negative log-likelihood of gold alignments while considering only alignments of slots (since we only align slots at test time). The cross-entropy loss for a training example $(x^{\text{abs}}, z^{\text{abs}}, A^*)$ is then given by:

$$- \sum_{i:i \in S_z} \sum_{j=1}^m A_{ij}^* \log p_{\text{align}}(x_j^{\text{abs}} | z_i^{\text{abs}}),$$

where S_z are the slot indices in z^{abs} .

Our model can be viewed as an attention model, dedicated to aligning logical form tokens to utterance tokens. Using a separate alignment model rather than the attention weights of the structure mapper has two advantages: First, alignments are generated given the entire generated sequence z^{abs} rather than just a prefix. Second, our model focuses its capacity on the alignment task without worrying about generation of z^{abs} . In Section 4, we will demonstrate that training a dedicated aligner substantially improves performance.

3.5 Inference

The aligner provides a distribution over utterance tokens for every slot in the abstract logical form. To compute the final logical form, we must replace each slot with a KB constant. Formally, let $(z_{j_1}^{\text{abs}}, \dots, z_{j_l}^{\text{abs}})$ be the sequence of slots in z^{abs} and denote them for simplicity as $y = (y_1, \dots, y_l)$. Our goal is to predict a sequence of KB constants $c = (c_1, \dots, c_l)$, where each c_i is chosen from a

candidate set $\mathcal{C}(y_i)$ that is determined by the abstract token y_i according to Table 1 (e.g., if y_i is `$REL`, then $\mathcal{C}(y_i)$ is the set of binary relations).

Our scoring function depends on alignments computed by the aligner. However, because slots are independent in the aligner, we introduce a few global constraints that capture the dependence between different slots. Formally, we wish to find c^* that maximizes the following scoring function, which depends on the utterance x , the slot sequence y , the abstract logical form z^{abs} , the alignment matrix A and the embedding function ϕ :

$$\arg \max_c \sum_{i=1}^l (s_{\text{local}}(c_i, y_i, x, A, \phi)) + s_{\text{global}}(c, z^{\text{abs}}).$$

We now describe our scoring functions in detail.

Local Score Because inference is applied only at test time, we have access to the lexicalized utterance and not only the abstract one. Thus, the aligner outputs a distribution over words for a slot y (e.g., in Figure 1, `$REL_DATE` aligns with high probability to `VERB`, which corresponds to the word “*published*”). Each word, in turn, has different semantic similarity to each KB constant in $\mathcal{C}(y)$. Intuitively, we would like to assign a KB constant that has high similarity with words the slot is aligned to. Thus, we define s_{local} of a KB constant c_i for every slot y_i to be its expected semantic similarity under the alignment distribution:

$$s_{\text{local}}(c_i, y_i, x, A, \phi) = E_{x \sim p_{\text{align}}(x^{\text{abs}} | y_i)} [\text{sim}_{\phi}(x, c_i)]$$

$$= \sum_{j=1}^m p_{\text{align}}(x_j^{\text{abs}} | y_i) \cdot \text{sim}_{\phi}(x_j, c_i).$$

We define $\text{sim}_{\phi}(x_j, c_i)$ to be the cosine similarity between the embedding $\phi(x_j)$ and the embedding $\phi(c_i)$ (scaled to the range $[0,1]$), where $\phi(c_i)$ is defined to be the average embedding of all words in $L(c_i)$, that is, $\text{sim}(x_j, c_i) = \frac{1 + \cos(\phi(x_j), \phi(c_i))}{2}$.

Global Score Utilizing only a local scoring function raises several concerns. First, slots are treated independently and dependencies between slots are ignored, which might result in a final logical form that is globally inconsistent. For example, we could generate the logical form `Birthplace.ComputerScience`, which is semantically dubious. Second, some KB constants do not align to any word in the utterance and appear

in the logical form only implicitly. For example, the logical form in Figure 1 contains the `Field` relation, however “*field*” is implicit in the utterance. Therefore, we define $exe_{\mathcal{K}}(z)$ to be true iff z executes against \mathcal{K} without errors, and define a global score that prevents assignments c that result in a logical form z such that $exe_{\mathcal{K}}(z)$ is false.

Moreover, we can use similar constraints to prevent logical forms that are highly unlikely according to our prior knowledge. Specifically, we define $once(z)$ to be true iff each date, named entity, and number in the logical form z appear exactly once. We then define a global score that prevents logical forms in which $once(z)$ is false. Empirically, we find such assignments to be mostly wrong (e.g., `Type.Article \sqcap (Field.QA \sqcup Field.QA)`).

Formally, our scoring function is defined as:

$$s_{\text{global}}(c, z^{\text{abs}}) = \begin{cases} 0 & exe_{\mathcal{K}}(z^{\text{abs}}|_c), once(z^{\text{abs}}|_c) \\ -\infty & \text{otherwise,} \end{cases}$$

where $z^{\text{abs}}|_c$ is the result of assigning the KB constants c to the slots in z^{abs} .

Inference Algorithm. While each local scoring function can be efficiently maximized independently, the global constraints that depend on the entire assignment c make inference more complicated. However, because the global scoring function introduces hard constraints, an exact and efficient inference algorithm is still possible. Our inference algorithm generates solutions one-by-one sorted by the local scoring function only. Then, it checks for each one whether it satisfies the global constraints defined by s_{global} , and stops once a satisfying solution is found, which is guaranteed to maximize our scoring function. While in the worst case, this procedure is exponential in the size of c , in practice solutions are found after only a few steps. We also always halt after T steps if a solution has not been found.

Algorithm 1 describes the details of our inference procedure. We define $cands$ to be a data structure that contains l lists of candidate KB constants (a list for each slot), sorted according to the local scoring function s_{local} in descending order. Additionally, $getAssign(cands, a)$ is a function that accesses $cands$, and retrieves the assignment with indices a . For example, $getAssign(cands, \{0\}^l)$ retrieves the top scoring local assignment. Last, we define $a_{\text{inc}(i)}$ to be the indices a , where a_i is incremented by 1.

Algorithm 1 Exact inference algorithm

Input: $cands, T$

Output: c^* - the top scoring assignment

```

1:  $horizon \leftarrow \emptyset$  ▷ Max heap
2:  $a_{\text{init}} \leftarrow \{0\}^l$ 
3:  $push(horizon, a_{\text{init}})$ 
4: for  $t \leftarrow 1$  to  $T$  do
5:    $a \leftarrow pop(horizon)$ 
6:    $c \leftarrow getAssign(cands, a)$ 
7:   if  $s_{\text{global}}(c, z^{\text{abs}}) = 0$  then
8:     return  $c$ 
9:   for  $i \leftarrow 1$  to  $l$  do
10:     $push(horizon, a_{\text{inc}(i)})$ 
11: return  $NULL$ 

```

The algorithm proceeds as follows. First we initialize a maximum heap $horizon$ into which we will dynamically push candidate assignments. Then, we iteratively pop the best current assignment from the heap, and check if it satisfies the global constraints. If it does, we return this assignment and stop. Otherwise, we generate the next possible candidates, one from each list (there is no need to add more than one because candidates are sorted). If no satisfying assignment is found after T steps, we return $NULL$. It is easy to show that when the algorithm returns an assignment it is guaranteed to be the one that maximizes our global scoring function.

4 Experiments

4.1 Experimental Setup

Data We evaluated our method on the OVERNIGHT semantic parsing dataset, which contains 13,682 examples of language utterances paired with logical forms across eight domains, which were chosen to explore diverse types of language phenomena. As described, our approach depends on having linguistic regularities repeat across domains. However two domains contain logical forms that are based on neo-davidsonian semantics for treating events with multiple arguments. Since such logical forms are completely absent in six domains, it is not possible for our method to generalize to those in our zero-shot approach. Therefore, we do not evaluate on the BASKETBALL domain, in which 98% of the examples contain such logical forms, and omit all examples (68%) that contain such logical forms in the SOCIAL domain. We evaluated on the same train/test split as Wang et al. (2015), using the same accuracy metric, i.e., the proportion of questions for which the denotations of the

Model	Blocks	Calendar	Housing	Publications	Recipes	Restaurants	Social	Avg.
INLEX	59.9	73.8	72.0	79.5	79.2	76.2	83.4	74.8
INABSTRACT	39.6	57.7	51.9	59.6	66.2	68.4	66.1	58.5
CROSSLX	0.0	0.0	0.5	0.0	1.4	1.2	1.0	0.6
CROSSLXREP	6.5	1.8	2.6	1.9	16.7	6.6	1.7	5.4
ZEROSHOT	28.3	53.6	52.4	55.3	60.2	61.7	62.4	53.4

Table 2: Test accuracy for all models on all domains.

	In-domain	Cross-domain
Lexical	INLEX	CROSSLX
Abstract	INABSTRACT	ZEROSHOT

Table 3: Evaluated models.

predicted and gold logical forms are equal. We additionally used the lexicon L they provided with descriptions for KB constants.

Evaluated Models We evaluated different models (Table 3) according to the following two attributes. Firstly, whether the model is trained on target domain data (*in-domain*) or on source domains data only (*cross-domain*). Secondly, we trained the neural semantic parser described in Section 2 over the lexical data representation (*lexical*), or in comparison trained our model over the abstract representation (*abstract*).

As CROSSLX can not generate KB constants unseen during training, we additionally implemented CROSSLXREP. In this model, we added an additional step that modifies the output of CROSSLX: we replaced a generated KB constant with its most similar KB constant from the target KB that also shares its abstract type.

Implementation Details In all experiments, for our embedding function $\phi(\cdot)$, we used pre-trained GloVe (Pennington et al., 2014) vectors with dimension 300. In a single experiment we considered one domain as the target domain, while other domains were the source domains (and repeated for all domains). For INLEX, CROSSLX and CROSSLXREP we used exactly the same experimental setup as Jia and Liang (2016). For our zero-shot model, we used 20% of the training data as a development set for tuning hyper-parameters. We first tuned parameters for the structure mapper, and used the best setting for tuning the aligner.

We provide the list of hyper-parameters and their values for our zero-shot framework. **Structure mapper:** number of epochs (22, using early stopping), hidden unit dimension (300), word vector dimension (100), learning rate (0.1 with SGD

optimizer), $L2$ regularization (0.001). At test time, we used beam search with beam size 5, and then picked the highest-scoring logical form that we could infer an assignment for. **Aligner:** number of epochs (30, using early stopping), hidden unit dimension (250), word vector dimension (100), learning rate (0.0002 with Adam optimizer), dropout rate over hidden states (0.4). For both models, word vectors are updated during training. **Inference:** we used $T = 500$ steps, after which we halted.

4.2 Results

We trained all models above and evaluated on the test set for all seven domains. Results show (Table 2) that ZEROSHOT substantially outperforms other zero-shot baselines. CROSSLX performs poorly, as it can only generate KB constants seen during training. CROSSLXREP performs better, as it can generate KB constants from the target domain, however, generating the correct constant usually fails. This highlights the challenge in the zero-shot semantic parsing setting.

For baselines trained on target domain data, INLEX (re-implementation of (Jia and Liang, 2016)) achieved average accuracy of 74.8, which is comparable to the 74.4 average accuracy they report on our seven domains. Training on the target domain with our method INABSTRACT achieved 58.5% average accuracy, which shows that while the abstract representation in our framework loses some valuable information, it is still successful. Importantly, the performance of ZEROSHOT (53.4%) is only slightly lower than INABSTRACT, showing that our model degrades gracefully and generalizes well across domains compared to CROSSLX.

Model Ablations We now measure the effect of different components of our framework on denotation accuracy. We examined the effect of removing components completely, or replacing them with simpler ones. Thus, the following ablated models can be viewed as additional baselines.

Model	Blocks	Calendar	Housing	Publications	Recipes	Restaurants	Social	Avg.
ZEROSHOT	29.2	60.4	57.3	47.7	59.5	63.0	64.0	54.5
-GLOBALHEUR	29.2	60.4	56.7	46.9	60.1	46.0	62.2	51.6
-ALIGNER	22.6	57.5	52.7	46.9	58.4	43.0	60.4	48.8
-INFERENCE	28.2	51.5	44.7	43.0	42.2	44.5	45.8	42.8
-ALIGNER,INFERENCE	20.4	38.8	32.7	37.5	37.0	15.5	36.9	31.2

Table 4: Development accuracy for all ablations.

1. -ALIGNER: Replacing the alignment distribution from the aligner with alignment distribution from the decoder of the structure mapper.
2. -INFERENCE: Discarding the global scoring function and maximizing each slot independently.
3. -ALIGNER,INFERENCE: Discarding both of our main technical contributions.
4. -GLOBALHEUR: Discarding the global inference heuristics (denoted as $once(z)$).

Table 4 shows that ablating each of the components hurts performance. Discarding our two main technical contributions results in 31.2% accuracy compared to 54.5% in the full model. Performing inference with global constraints dramatically improve performance, showing that using the alignment model alone results often in incoherent logical forms. Our dedicated aligner also improves performance compared to alignments learned by the decoder of the structure mapper. This is pronounced without global constraints (a drop from 42.8% to 31.2%), but is less severe when global inference is used (a drop from 54.5% to 48.8%).

Intrinsic Analysis While we evaluated performance above via denotation accuracy, we now evaluate our framework’s modules with different metrics (on the development set). We evaluated the structure mapper by measuring the exact match of the top candidate in the beam to the gold abstract logical form (49.1%). We further evaluated the aligner by measuring alignment accuracy for top candidate alignments, in comparison to the unsupervised aligner output (72.9%).

Finally, we measured inference performance in the following ways. The fraction of cases where inference succeeded within T steps is 70% (as some predicted abstract logical forms are not valid in terms of their syntax), and the average number of steps in case of success (3.67 steps). In addition, the fraction of correct global assignments given an abstract logical form that exactly matches the gold one is 77.0%. To conclude, results show that the structure mapping problem is harder than




Figure 5: Learning curve for INLEX, compared to ZEROSHOT average performance.

slot filling, for which we learned good alignments and performed fast and mostly accurate inference.

Valuation To estimate the value of our zero-shot framework in terms of target domain examples, we plot a learning curve (Figure 5) that shows development set average accuracy for INLEX (trained on target domain data). In comparison, ZEROSHOT utilizes no target domain data, thus it is fixed. As Figure 5 shows, our framework’s value is equal to 30% of the target domain training data. In our setting this equals to 400 examples manually-annotated with full logical forms. Note that this value is gained every time a semantic parser for a new domain is needed. Moreover, our parser can be used as an initial system, deployed to begin training from user interaction directly.

Limitations We now outline some of the limitations of our approach for zero-shot semantic parsing. We hypothesized that language regularities repeat across domains, however as mentioned above, neo-davidsonian semantics occurs mostly in one domain in the OVERNIGHT dataset and thus we were not able to generalize to it. Our parser also obtained low accuracy in BLOCKS. This domain contains mostly spatial language, different from other domains in OVERNIGHT. Specifically, prepositions, which we did not lexicalize map to

relations in the KB (e.g., “below” and “above” map to the relations `Below` and `Above`). This shows the challenge involved in decomposing the structure from the lexicon with rules. In addition, since some spatial relations in this domain are semantically similar (`Length`, `Width` and `Height`), we found it hard to rank them correctly during inference. This stresses that in our framework, we assume KB constants to be sufficiently distinguishable in the pre-trained embedding space, which is not always the case.

5 Related Work

While zero-shot executable semantic parsing is still under-explored, some works focused on the open-vocabulary setting which handles unseen relations by replacing a formal KB with a probabilistic database learned from a text corpus (Choi et al., 2015; Gardner and Krishnamurthy, 2017).

Our abstract utterance representation is related to other attempts to generate intermediate representations that improve generalization such as dependency trees (Reddy et al., 2016), syntactic CCG parses (Krishnamurthy and Mitchell, 2015), abstract templates (Abujabal et al., 2017; Goldman et al., 2018) or masked entities (Dong and Lapata, 2016). Our abstract logical form representation is similar to that Dong and Lapata (2018) used in to guide the decoding of the full logical form. The main difference with our work is that we focus on a comprehensive abstract representation tailored for zero-shot semantic parsing.

It is worth mentioning other work that inspected various aspects of zero-shot parsing. Bapna et al. (2017) focused on frame semantic parsing, and assumed that relations appear across different domains to learn a better mapping in the target domain. Also in frame semantic parsing, Ferreira et al. (2015) utilized word embeddings to map words to unseen KB relations. Finally, Lake and Baroni (2017) inspected whether neural semantic parsers can handle types of compositionality that were unseen during training. The main difference between their work and ours is that we focus on a scenario where a compositional logical form is generated, but the target KB constants do not appear in any of the source domains.

6 Conclusion

In this paper we address the challenge of zero-shot semantic parsing. We introduce a model that

can parse utterances in unseen domains by decoupling structure mapping from lexicon mapping, and demonstrate its success on 7 domains from the OVERNIGHT dataset.

In future work, we would like to automatically learn a delexicalizer from data, tackle zero-shot parsing when the structure distribution in the target domain is very different from the source domains, and apply our framework to datasets where only denotations are provided.

Acknowledgments

We thank Kyle Richardson, Vivek Srikumar and the anonymous reviewers for their constructive feedback. This work was completed in partial fulfillment for the PhD degree of the first author. Herzig was supported by a Google PhD fellowship. This research was partially supported by The Israel Science Foundation grant 942/16 and The Blavatnik Computer Science Research Fund.

References

- Abdalghani Abujabal, Mohamed Yahya, Mirek Riedewald, and Gerhard Weikum. 2017. Automated template generation for question answering over knowledge graphs. In *Proceedings of the 26th international conference on world wide web*, pages 1191–1200. International World Wide Web Conferences Steering Committee.
- Y. Artzi and L. Zettlemoyer. 2013. Weakly supervised learning of semantic parsers for mapping instructions to actions. *Transactions of the Association for Computational Linguistics (TACL)*, 1:49–62.
- D. Bahdanau, K. Cho, and Y. Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *International Conference on Learning Representations (ICLR)*.
- Ankur Bapna, Gokhan Tur, Dilek Hakkani-Tur, and Larry Heck. 2017. Towards zero shot frame semantic parsing for domain scaling. In *Interspeech 2017*.
- 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.
- Eunsol Choi, Tom Kwiatkowski, and Luke Zettlemoyer. 2015. Scalable semantic parsing with partial ontologies. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference*

- on *Natural Language Processing (Volume 1: Long Papers)*, volume 1, pages 1311–1320.
- 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.
- L. Dong and M. Lapata. 2016. Language to logical form with neural attention. In *Association for Computational Linguistics (ACL)*.
- Li Dong and Mirella Lapata. 2018. Coarse-to-fine decoding for neural semantic parsing. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 731–742. Association for Computational Linguistics.
- Chris Dyer, Victor Chahuneau, and Noah A. Smith. 2013. A simple, fast, and effective reparameterization of IBM model 2. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 644–648, Atlanta, Georgia. Association for Computational Linguistics.
- Xing Fan, Emilio Monti, Lambert Mathias, and Markus Dreyer. 2017. Transfer learning for neural semantic parsing. In *Proceedings of the 2nd Workshop on Representation Learning for NLP*, pages 48–56, Vancouver, Canada. Association for Computational Linguistics.
- Emmanuel Ferreira, Bassam Jabaian, and Fabrice Lefèvre. 2015. Zero-shot semantic parser for spoken language understanding. In *Sixteenth Annual Conference of the International Speech Communication Association*.
- Matt Gardner and Jayant Krishnamurthy. 2017. Open-vocabulary semantic parsing with both distributional statistics and formal knowledge. In *AAAI*, pages 3195–3201.
- Omer Goldman, Veronica Latcinnik, Ehud Nave, Amir Globerson, and Jonathan Berant. 2018. Weakly supervised semantic parsing with abstract examples. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1809–1819. Association for Computational Linguistics.
- Jonathan Herzig and Jonathan Berant. 2017. Neural semantic parsing over multiple knowledge-bases. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 623–628. Association for Computational Linguistics.
- S. Hochreiter and J. Schmidhuber. 1997. Long short-term memory. *Neural Computation*, 9(8):1735–1780.
- Srinivasan Iyer, Ioannis Konstas, Alvin Cheung, Jayant Krishnamurthy, and Luke Zettlemoyer. 2017. Learning a neural semantic parser from user feedback. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 963–973, Vancouver, Canada. Association for Computational Linguistics.
- R. Jia and P. Liang. 2016. Data recombination for neural semantic parsing. In *Association for Computational Linguistics (ACL)*.
- Jayant Krishnamurthy and Tom M Mitchell. 2015. Learning a compositional semantics for freebase with an open predicate vocabulary. *Transactions of the Association for Computational Linguistics*, 3:257–270.
- 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. 2011. Lexical generalization in CCG grammar induction for semantic parsing. In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 1512–1523.
- Brenden M Lake and Marco Baroni. 2017. Still not systematic after all these years: On the compositional skills of sequence-to-sequence recurrent networks. *arXiv preprint arXiv:1711.00350*.
- P. Liang. 2013. Lambda dependency-based compositional semantics. *arXiv*.
- P. Liang, M. I. Jordan, and D. Klein. 2011. Learning dependency-based compositional semantics. In *Association for Computational Linguistics (ACL)*, pages 590–599.
- M. Luong, H. Pham, and C. D. Manning. 2015. Effective approaches to attention-based neural machine translation. In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 1412–1421.
- C. D. Manning, M. Surdeanu, J. Bauer, J. Finkel, S. J. Bethard, and D. McClosky. 2014. The Stanford CoreNLP natural language processing toolkit. In *ACL system demonstrations*.
- J. Pennington, R. Socher, and C. D. Manning. 2014. Glove: Global vectors for word representation. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- Siva Reddy, Oscar Täckström, Michael Collins, Tom Kwiatkowski, Dipanjan Das, Mark Steedman, and Mirella Lapata. 2016. Transforming dependency structures to logical forms for semantic parsing. *Transactions of the Association for Computational Linguistics*, 4:127–140.

- Kyle Richardson, Jonathan Berant, and Jonas Kuhn. 2018. Polyglot semantic parsing in apis. 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 720–730. Association for Computational Linguistics.
- Yu Su and Xifeng Yan. 2017. Cross-domain semantic parsing via paraphrasing. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1235–1246, Copenhagen, Denmark. Association for Computational Linguistics.
- I. Sutskever, O. Vinyals, and Q. V. Le. 2014. Sequence to sequence learning with neural networks. In *Advances in Neural Information Processing Systems (NIPS)*, pages 3104–3112.
- Y. Wang, J. Berant, and P. Liang. 2015. Building a semantic parser overnight. In *Association for Computational Linguistics (ACL)*.
- M. Zelle and R. J. Mooney. 1996. Learning to parse database queries using inductive logic programming. In *Association for the Advancement of Artificial Intelligence (AAAI)*, pages 1050–1055.
- L. S. Zettlemoyer and M. Collins. 2005. Learning to map sentences to logical form: Structured classification with probabilistic categorial grammars. In *Uncertainty in Artificial Intelligence (UAI)*, pages 658–666.