SpatialNLI: A Spatial Domain Natural Language Interface to Databases Using Spatial Comprehension

Jingjing Li*
Auburn University
jingjingli@auburn.edu

Wenlu Wang*
Auburn University
wenluwang@auburn.edu

Wei-Shinn Ku
Auburn University
weishinn@auburn.edu

Yingtao Tian
Stony Brook University
yittian@cs.stonybrook.edu

Haixun Wang
WeWork Research
haixun.wang@wework.com

ABSTRACT

A natural language interface (NLI) to databases is an interface that translates a natural language question to a structured query that is executable by database management systems (DBMS). However, an NLI that is trained in the general domain is hard to apply in the spatial domain due to the idiosyncrasy and expressiveness of the spatial questions. Inspired by the machine comprehension model, we propose a spatial comprehension model that is able to recognize the meaning of spatial entities based on the semantics of the context. The spatial semantics learned from the spatial comprehension model is then injected to the natural language question to ease the burden of capturing the spatial-specific semantics. With our spatial comprehension model and information injection, our NLI for the spatial domain, named *SpatialNLI*, is able to capture the semantic structure of the question and translate it to the corresponding syntax of an executable query accurately. We also experimentally ascertain that *SpatialNLI* outperforms state-of-the-art methods.

CCS CONCEPTS

- **Human-centered computing** → **Natural language interfaces**;
- **Information systems** → **Geographic information systems**.

KEYWORDS

Natural Language Interface, Spatial Data Science

ACM Reference Format:

Jingjing Li, Wenlu Wang, Wei-Shinn Ku, Yingtao Tian, and Haixun Wang. 2019. SpatialNLI: A Spatial Domain Natural Language Interface to Databases Using Spatial Comprehension. In *27th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems (SIGSPATIAL '19)*, November 5–8, 2019, Chicago, IL, USA. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3347146.3359069>

*Both authors contributed equally to this research.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGSPATIAL '19, November 5–8, 2019, Chicago, IL, USA
© 2019 Association for Computing Machinery.
ACM ISBN 978-1-4503-6909-1/19/11...\$15.00
<https://doi.org/10.1145/3347146.3359069>

1 INTRODUCTION

Many business applications rely on relational databases. To facilitate the usage of database management systems to the public, NLI to databases has been extensively studied [1–10]. Spatial Domain NLI to databases has drawn great attention due to the popularity of spatial applications [11–14]. An intuitive solution is to adopt existing NLI in general databases to the spatial domain. However, due to the idiosyncrasy and expressiveness of the spatial semantics, it is unfeasible to adopt general NLI for the spatial domain directly. The challenge of adopting the existing general domain NLI to spatial domain lies to harnessing the expressiveness of spatial semantics. The expressiveness of spatial semantics can be justified based on the following observations [15]:

The meaning of spatial phrase “ <i>Mississippi</i> ”	
How many rivers does Mississippi have ?	state
How many cities does Mississippi run through ?	river
The meaning of spatial phrase “ <i>over</i> ”	
How many people walked over the bridge ?	on
How many birds flew over the bridge ?	above
The meaning of spatial phrase “ <i>at the back of</i> ”	
How many trees are at the back of the building ?	exterior
How many rooms are at the back of the building ?	interior

Figure 1: Three examples show that the spatial semantics is encyclopedic.

The examples as mentioned earlier show that the same spatial phrase in different questions embodies divergent senses expressing divergent query intentions. In the first two questions, “*Mississippi*” as a name can refer to either a state or a river, depending on the context where it is mentioned. In this example, the type of word “*Mississippi*” depends on the verb located after the name (“*have*” or “*run through*”). In the second two questions, the preposition “*over*” means either a superior position or on the surface. Its spatial meaning depends on the verb located before the preposition (“*walk*” or “*fly*”). In the last two questions, the prepositional phrase “*at the back of*” means either outside the building or inside the building, which depends on the noun before the prepositional phrase (“*tree*” or “*room*”). Such contextually dependent spatial semantics raises serious challenges for NLI to spatial domain databases. For instance, in the third example, if there are two spatial tables (one for the interior architecture of a building and one for the surroundings of a building), a wrongly comprehended spatial semantics would cause the NLI to query a wrong table. In general, spatial semantic understanding relies heavily on its contextual interpretation.

Existing works of NLI rely on conventional grammar-based methods or neural network-based methods. The former line of existing work uses predefined templates or manually designed features, which has the lower-transfer ability, thus confined in its specific dataset. The latter line of existing work uses grammar embedded neural networks. Embedding grammar into a model relies on converting the process of generating a sequence of tokens to the task of generating a sequence of actions that expands a syntax tree. Converting word space to action space will inevitably introduce transformation error, which can not guarantee overall accuracy. To the best of our knowledge, the state-of-the-art Syntax-based method TRANX [16] achieves an accuracy of 88.2% which is 2.1% lower than our accuracy.

The aforementioned observations and survey inspired us to propose a Spatial Domain NLI that is able to support the idiosyncrasy of spatial semantics. Inspired by the NLI in [9], we propose a strategy to address the ambiguity of spatial meaning (mentioned in Figure 1) and data sparsity problem by feeding necessary spatial semantics to the deep model. Here ambiguous spatial phrases are those that can not be uniquely identified by the schema. By feeding external spatial semantics, our NLI is able to support various spatial questions even when it has not seen similar semantics in the training set. The extra spatial semantics is recognized by our external spatial comprehension model, whose functionality is to recognize pre-defined spatial semantics.

We propose to capture spatial semantics using an external spatial comprehension model, where the interpretation of each word is based on the attentive combination of the context. We then complete our NLI model using a sequence-to-sequence (seq2seq) translation, which is not only able to achieve grammar correctness but also robust with data sparsity problem. Our fundamental strategy is to separate the tasks of NLI to (1) learning semantic structure of a natural language question, and (2) learning the spatial semantics of a spatial question.

The necessity of the external spatial comprehension model is due to the seq2seq translation model's failure to capture all the spatial semantics while learning the structure of the question. In our design, Task (1) is assigned to the seq2seq model, while an external spatial comprehension model is in charge of Task (2). We propose our spatial comprehension model as a bi-directional attentive workflow [17, 18], where the attentive spatial phrases of the input are enclosed with special symbols.

Our strategy is a general-purpose automatic solution that only relies on database content, an external model, a seq2seq model, and a minimum amount of human knowledge. To the best of our knowledge, we are the first to use an external spatial semantic understanding model to enhance the performance of the main seq2seq model. Our solution not only addresses the problem of data sparsity but also introduces minimum error since the spatial comprehension model achieves an accuracy of 98% and 100% for Geoquery and Restaurant datasets.

Our contribution is described as follows

- We propose a spatial comprehension model that is able to recognize the meaning (e.g., POI type) of an ambiguous spatial phrase (e.g., POI name) based on contextual interpretation.
- After injecting spatial semantics learned from spatial comprehension into the question, our model outperforms the state-of-the-art.
- We evaluate our strategies systematically and show that our spatial comprehension model and injection format perform well as expected.

2 RELATED WORK

In recent years, a line of works has been focusing on semantic parsing, which aims at converting natural language utterances to formal meaning representations. ZC05 [19], ZC07 [20], UBL [21] and FUBL [22] induce the specific grammars to make the translation, which defines the meanings of individual words and phrases. KCAZ13 [23] and [7] use ontologies to help form the grammar. [24] uses domain-independent facts to make the translation and ZH15 [25] builds the grammar based on the specific entity type of words. DCS+L [26] and [4] introduce tree structure for input natural language to solve this problem. However, most of the conventional methods rely on predefined templates or manually designed features to complete the translation, which is not comparable to ours, as we avoid using such lexicon mapping and predefined templates of prior knowledge in our system.

Also, some of the work focuses on the Natural Language Interface to Databases (NLIDB) for users to interact with the database without acknowledging the grammar of structured queries executed by the database engine. [1] first explores this issue with a specific database and concrete examples. [6, 27–30] also work on this issue depending on grammars and processes the semantic grammars manually for each individual database. [2] and [3] work on the NLIDB systems requiring large sets of natural language and SQL pairs. [5] and [8] present an interface with the help of the feedback from users and PEK03 [6] also defines the coverage of the NLIDB system, which is certainly not suitable for all databases. The problem is that the NLIDB study mentioned above is all designed for the general domain and is hard to apply to spatial natural language queries directly without loss of accuracy.

Now deep neural network models have been applied successfully to semantic parsing to exploit the sequential structure on both input and output side. One of them is the Encoder-Decoder model [31–33]. FKZ18 [34] works on translating the input to SQL queries based on the Encoder-Decoder model. TRANX [16] and ASN [35] construct Decoder-Encoder models with the tree structure. SQL2TREE [36] proposes a seq2seq model based on the Encoder-Decoder architecture and JL16 [37] enhances the performance of seq2seq by adding attention-based copying in the output and implementing data augmentation. [9, 38] work on an Encoder-Decoder based transfer learning for semantic parsing. [39] trains one model that is able to parse natural language sentences from multiple different languages and [40] exploits the Encoder-Decoder model in different domains. [41] introduces a framework with reinforcement learning to generate SQL queries. Here, we introduce on seq2seq model in our system. Compared with SQL2TREE [36] and JL16 [37], our model solves the spatial ambiguity problem for the input natural language more efficiently.

Natural language process for the spatial domain has been observed in literature. [42] annotates the spatial relation in natural

Figure 2: SpatialNLI overview.

language based on the specific annotation schema. [43] focuses on spatial ontologies to process the input spatial natural language queries. [44] maps the objects and spatial relations to formal linguistic terms, which disambiguate the spatial meanings of objects. [45] uses a form of symbolic expressions to extract spatial terms from natural language descriptions to represent spatial features and relations between them. All of them limit the query in the fixed form and have difficulty dealing with different kinds of spatial complex queries. For [46], it introduces a system that is capable of capturing the semantics of spatial relations in natural language using the neural network. But none of the above gives users an interface to interact with the database.

3 CHALLENGES OF SPATIAL NLI

- I. Sparse training data. Even though data augmentation is a feasible solution to sparse training data, it is likely that the deep model will be forced to handle unseen questions that are not covered by data augmentation, and required to support transfer learning.
- II. Spatial semantics ambiguity. A unique feature of spatial questions is its expressiveness in the spatial domain, and a spatial phrase often has an ambiguous meaning. For example, “Mississippi” could be either a state or a river, “New York” could be either a city or a state (taking Figure 3 as examples).

	Question
State	How many people live in <i>Mississippi</i> ?
River	How many states does the <i>Mississippi</i> run through?
City	Is <i>New York</i> or London bigger?
State	What is the capital of <i>New York</i> ?

Figure 3: Spatial POI ambiguity

Theoretically, a powerful data augmentation should be able to address the first challenge; however, such data augmentation strategy is rare in practice. Moreover, a seq2seq model is designed to translate a sentence, it is reasonable that it fails to capture the context precisely and infer the correct spatial semantics. For example, in the question “How many rivers does Mississippi have?” (shown in Figure 4), a seq2seq model should be able to understand the context and infer “Mississippi” as a state. However, since the word “rivers” appears in the question and precedes the word “Mississippi”, which means it has a major impact on the prediction when attentive on

Question	How many rivers does Mississippi have?
Ground Truth	answer(A, count(B, (river(B), const(C, <i>stateid</i> (Mississippi)), loc(B, C)), A))
Infer	answer(A, count(B, (river(B), const(C, <i>riverid</i> (Mississippi)), loc(B, C)), A))

Figure 4: POI type recognition without spatial comprehension model

“Mississippi”, in that case, it is highly possible that *stateid* will be inferred instead of *riverid*.

Therefore, we propose another deep model for the purpose of spatial semantic understanding; despite the fact that we use the same parse training data, an external model targeted on understanding the context is able to infer the correct spatial semantics (Challenge II) precisely. With the spatial semantics retrieved from the spatial comprehension model, we adopt symbol insertion strategy [9] to inject external information and help the seq2seq to infer an unseen sample correctly (Challenge I).

Overview To address the aforementioned challenges, we present our *SpatialNLI* overview shown in Figure 2. The workflow of our SpatialNLI involves the following steps:

1. Identify ambiguous spatial semantics in the NL query.
2. Build a spatial comprehension model that is able to understand a spatial-related question semantically.
3. Injecting spatial semantics retrieved from the spatial comprehension model into the question ($q \rightarrow q'$).
4. “Translating” the question into a structured query (Lambda expression in our example) ($q' \rightarrow l'$).
5. Replace the symbols injected to their original text ($l' \rightarrow l$).

4 SPATIALNLI

Since most of the keywords or data elements in spatial queries (e.g., lambda expression) are spatial-related, we propose a strategy to inject latent spatial semantics into the natural language question to help the seq2seq model to capture the semantic meaning of the question. For example, for the question “How many rivers does Mississippi have?”, its correspondence lambda expression is “count(B, (river(B), const(C, *stateid*(Mississippi)), loc(B, C)))”, which has five keywords “count”, “river”, “const”, “*stateid*”, “loc”, and three of them are spatial-related. We will illustrate the workflow of our SpatialNLI with this running example.

Our SpatialNLI model is composed of the following steps (corresponding to Algorithm 1)

Algorithm 1 SpatialNLI

-
- ```

1: function SPATIALNLI(q, D, E)
2: $P, V = \text{SPATIALMAPPER}(D, q, E)$;
3: $q', s2p = \text{SPATIALINJECTION}(q, V, P)$;
4: $l' = \text{SEQ2SEQ}(q')$;
5: $l = \text{RECOVER}(l', s2p)$;
6: Return l ;

```
- 
1. **Spatial Semantics Detection.** Having access to GeoSpatial databases, we detect potential keywords or data elements using 1) string match, 2) edit distance, and 3) cosine distance in semantic embedding space (e.g., Glove). In the aforementioned question, “*Mississippi*” can be detected by comparing against the data in the databases using string match, “*river*” can be detected by edit distance since “*rivers*” is in the table. We will detail semantic distance measurement later in Section 4.1.
  2. **Spatial Comprehension Model.** For ambiguous spatial phrases, we propose a spatial comprehension model to resolve the ambiguity. As we mentioned in Section 3, “*Mississippi*” is an ambiguous POI, and it will be identified as a river type using our spatial comprehension model.
  3. **Spatial Semantics Injection.** With identified keyword “*river*” and data element “*Mississippi*” (river name), we inject such information into the question by inserting pre-defined symbols “*How many*  $\langle k0 \rangle$  *rivers*  $\langle eok \rangle$  *does*  $\langle k1 \rangle$  *stateid*  $\langle eok \rangle$   $\langle v0 \rangle$  *Mississippi*  $\langle eov \rangle$  *have?*”
  4. **Seq2seq Translation.** We then feed the modified question to a seq2seq translation model. In the previous example, the predicted output sequence is “*answer*( $A$ , *count*( $B$ , ( $\langle k0 \rangle$ ( $B$ ), *const*( $C$ ,  $\langle k1 \rangle$ ( $\langle v0 \rangle$ ))), *loc*( $B$ ,  $C$ )),  $A$ )”.
  5. **Query Recovery.** The generated sequence of the seq2seq model is then recovered to an executable query. Following the previous example, we have “*answer*( $A$ , *count*( $B$ , (*river*( $B$ ), *const*( $C$ , *stateid*(*Mississippi*))), *loc*( $B$ ,  $C$ )),  $A$ )”.

**4.1 Spatial Semantics Detection**

Even though our major contribution is spatial comprehension, we formally define our strategy to detect keywords and data elements mentioned in the question (denoted as *SpatialMapper*) to keep our work self-contained.

$$P, V = \text{SPATIALMAPPER}(D, q, E)$$

The inputs are the GeoSpatial database  $D$ , a natural language question  $q$ , and an embedding function  $E$  (e.g., Glove).  $E$  will change a word to a high-dimensional vector, which represents its location in the embedding space. We collect the table names, column names, and column values from  $D$ , thus  $D$  refers to a collection of entities in our spatial mapper. The table names (e.g., *river*) and column names (e.g., *river length*) in  $D$  are potential keywords of executable queries (e.g., Lambda expression), and column values (e.g., *Mississippi*) are potentially data elements that might be mentioned in executable queries. The detail of the algorithm is presented in Algorithm 2, in the aforementioned example,  $P = [\langle \text{rivers}, \text{river} \rangle, \langle \text{Mississippi}, \text{Mississippi} \rangle]$ , since “*Mississippi*” is detected by exact string match, and “*rivers*” (in  $q$ ) and “*river*” (in  $D$ ) has a small edit distance. We define semantic distance measurement

**Algorithm 2** Spatial Semantics Mapper

- 
- ```

1: function SPATIALMAPPER( $D, q, E$ )
2:    $P = \emptyset$ ; ▷ Spatial semantics matching pairs.
3:    $V = \emptyset$ ; ▷ Spatial values with its semantic meaning.
4:   for  $k$  in  $K..1$  do ▷ Iterating from  $K$ -gram to 1-gram
5:     for all  $p_q \in k\text{-gram of } q$  do
6:       for all  $c \in D$  do
7:         if  $p_q == c$  or semantic_distance( $p_q, c$ )  $< \tau_{sem}$ 
           or edit_distance( $p_q, c$ )  $< \tau_{ed}$  then ▷  $\tau_{ed}$  is the threshold for edit distance.  $\tau_{sem}$  is the threshold for semantic distance.
8:            $P.add(\langle p_q, c \rangle)$ ;
9:           if  $c$  is a column value then
10:            for all table  $tb$  that has  $c$  do
11:               $p_q\text{-types}.add(tb)$ 
12:               $V.add(p_q)$ 
13:   Return  $P, V$ ;

```
-

as

$$\text{semantic_distance}(a, b) = 1 - \frac{E(a) \cdot E(b)}{\|E(a)\|_2 \|E(b)\|_2}$$

The semantic distance is also the spatial distance in the embedding space. If any operand is a phrase which comprises multiple tokens, for example, A is a list of tokens, we define $E(A) = \text{avg}_{a \in A} (E(a))$.

Taking question “*Where is the lowest spot in Iowa?*” as an example, its corresponding logic form is *answer*(A , *lowest*(A , (*place*(A), *loc*(A, B), *const*(B , *stateid*(*Iowa*))))), “*spot*” in the NL question is matched to keyword “*place*” since *semantic_distance*(*place*, *spot*) < 0.368 , which is relatively small.

We care about the spatial phrases that have semantic ambiguities (e.g., *Mississippi*). An intuitive solution is to use pre-collected human knowledge. However, to devise an automatic and intelligent approach, we propose using Geospatial database; for example, we discover that “*Mississippi*” is an ambiguous value by simply searching for this phrase in the database, and it appears in two tables *RIVER* and *STATE*. In Algorithm 2 Line 9-12, if a phrase appears in multiple tables, we collect all the ambiguous information in V . For example, in V , “*Mississippi*”.*types* = [*RIVER*, *STATE*], “*New York*”.*types* = [*CITY*, *STATE*] and “*Alabama*”.*types* = [*STATE*]. In other words, for a spatial phrase that is a value, we collect the tables it belongs to and stored in V . For most of the Geospatial databases, the table names are able to represent the meaning of the value. In the question “*How many rivers does Mississippi have?*”, $V = [\text{“Mississippi”}]$ and “*Mississippi*”.*types* = [*RIVER*, *STATE*].

It is worth noticing, we use minimum human knowledge to cover phrase mapping that is not covered by Glove. For example, in a question where “*population per km²*” refers to “*population density*”, such mapping is not easy to be covered by Glove or a deep model, thus human knowledge is necessary. However, such cases are rare in practice, and we only require minimum human knowledge.

4.2 Spatial Comprehension

A critical challenge in understanding the spatial question is the meaning of an ambiguous phrase, such as a point of interest (POI). For example in Figure 3, where “*Mississippi*” could be a river or a state, we have to differentiate its meaning by the context and

Figure 5: Spatial Comprehension Model

understanding the semantic meaning of the question. In the question “How many people live in *Mississippi*?”, we would interpret “*Mississippi*” as a state name, and in the question “How many states does the *Mississippi* run through?”, people would understand “*Mississippi*” is referring to a river.

With large training corps, a deep model might be able to capture that information; however, existing spatial question answer data sets are inadequate and sparse due to the difficulty in collecting ground truth. To address this challenge, we propose a principle method to enable semantic understanding of spatial questions using sparse training data, which relieves the burden of collecting large training sets. By our definition, *Spatial Comprehension* is spatial semantic understanding using machine comprehension. Our strategy is to exploit pre-trained GloVe embedding to understand spatial keywords in the question first, then use a seq2seq machine comprehension model to learn the semantic meaning of the question (context) without the burden of extracting the spatial relations.

Model Structure Our spatial comprehension model is designed to understand the “meaning” (e.g., type of POI) of an ambiguous spatial phrase mentioned in the question based on its context.

Inspired by the machine comprehension model using an attention flow [17], we propose our spatial comprehension model composed of two stacked LSTM layers on each input with another shared attentive LSTM layer. The design of the bi-directional attentive workflow [17] is to answer a question given a premise – i.e., locate the sentences in the premise that is most relevant to the answer of the question. The task of the machine comprehension is different from ours; however, they do share the same strategy: understanding one of them (question and premise) semantically based on the context of the other. So we use LSTM to pre-process our NL question and the possible meaning of the ambiguous spatial phrase separately, and conduct an attentive workflow over the hidden states of the question (shown in Figure 5). Also inspired by [9], we enclose the ambiguous spatial phrase in special symbols to indicate it has more influence than the other tokens in the question.

We denote a question as $q = [q_1, \dots, q_n]$ and the meaning of the ambiguous spatial phrase (mentioned in the question), such as a POI type, as $t = [t_1, \dots, t_m]$, both of which are fed to a word embedding layer ϕ (initialized with GloVe [47]). On top of that, we use an LSTM layer to capture the hidden states of each time step i .

We build the same structure for both q and t . We denote the hidden states of the top stacked LSTM layers as

$$H^q = \text{LSTM}(\phi(q)) = (h_1^q, \dots, h_n^q) \quad H^t = \text{LSTM}(\phi(t)) = (h_1^t, \dots, h_m^t)$$

Inspired by natural language understanding proposed in [10, 17], we build an extra LSTM layer on H^q with attention over H^t as the follows

$$d_i = \text{LSTM}([h_i^t, \beta_{i-1}], d_{i-1})$$

$$e_{ij} = v^T \text{Tanh}(W_0 H^t + W_1 h_j^q + W_2 d_i)$$

$$\alpha_{ij} = e_{ij} / \sum_j e_{ij}$$

$$\beta_i = \sum_{j=1}^n \alpha_{ij} h_j^q$$

$d_0 = 0$. W_0, W_1, W_2, v , and U are model parameters. Here i is the time step while enumerating t , and j enumerates each token in q . The final output d_m is fed to a multi-layer perceptron (MLP) and then resized to a binary prediction. If t involves a sequence of tokens, we use bi-directional attentive flow as in [17] and compute bi-directional output $d_i = [\vec{d}_i, \overleftarrow{d}_i]$.

With the attentive flow on type t while reading the question, our spatial comprehension model is able to make the prediction based on the memory of the context. However, with the observation that, given a question “How many rivers in *Mississippi*?” and a POI type “river”, the machine comprehension model is highly likely to produce a positive prediction (false prediction), since “river” is mentioned in the question by “rivers”, but the model would fail to capture our intention to categorize the type of “*Mississippi*”. In order to feed our intention into the model, we insert special symbols (e.g., $\langle @ \rangle$) to enclose the POI mentioned in the question.

Question	POI Type	Label
How many states does the $\langle @ \rangle$ Mississippi $\langle @ \rangle$ run through?	State	False
How many states does the $\langle @ \rangle$ Mississippi $\langle @ \rangle$ run through?	City	False
How many states does the $\langle @ \rangle$ Mississippi $\langle @ \rangle$ run through?	River	True

Figure 6: Spatial Comprehension Model Training Samples

The corresponding model structure is shown in Figure 5. For the question shown in Figure 5, to address the ambiguity of “*Mississippi*”, we feed three records shown in Figure 6. Our spatial comprehension function is defined as follows:

$$\text{SPATIALCOMPREHENSION}(q, p, t)$$

where q is the question, p is the ambiguous phrase, and t is the meaning of the phrase. If SPATIALCOMPREHENSION returns true, the semantic meaning of p is identified as t .

4.3 Spatial Semantics Injection

Now we are able to understand the context and recognize the meaning of each ambiguous spatial phrase correctly through the spatial comprehension model. The question is how to inject the external information to the main seq2seq model. We propose an injection strategy shown in Figure 7. The general idea is to insert symbols into the question at the locations before and after every spatial phrase to emphasize its semantics, then feed the inserted question into a seq2seq model.

The first step is to search for components in the question that need a spatial semantics injection—i.e., ambiguous spatial semantics. For example, in the question “What is the population of San Antonio?”, we do not feed extra information for every word, instead, we only focus on spatial information or tokens that could be shown in the corresponding logic form (e.g., keywords). In other words, we only care about the tokens in the NL question that contribute to its logic form. The tokens such as question word “what” and stop words “is” “the” “of” do not contain the question’s information, thus are not annotated.

For the ambiguity of spatial phrases, we believe it is necessary to feed the meaning of the phrase in the question. For example, we feed the **Type of POI** in the question to address the POI ambiguity. We present our *Information Injection Format With Type Feeding* in Figure 7. We will validate in the experiment section that our type feeding improves the accuracy dramatically.

We propose a general purpose automatic injection algorithm shown in Algorithm 3. For the input natural language question q , after we recognize the meaning (e.g., type) t_i of each phrase p_i (e.g., POI) correctly through spatial comprehension model, for each pair of $\langle p_i, t_i \rangle \in \langle P, T \rangle$, we will insert the type t_i before the p_i in the input question q . Also, for $p_i \in P$, we store the symbol sym of each phrase p_i in $s2p$, which will be used later in the query recovery. For example, as Figure 7 shows, the symbol for the phrase “san antonio” is $\langle v0 \rangle$. Then the $\langle \langle v0 \rangle, \text{san antonio} \rangle$ is stored in $s2p$ which will later be used for the recovery of $\langle v0 \rangle$ in the output logic form query.

Figure 7 presents our detailed injection format. For a phrase or token in the question that is identified as a keyword (e.g., population), that phrase or token will be enclosed with $\langle ki \rangle$ and $\langle eok \rangle$. For values such as “San Antonio” that appear in the question, we enclose them with $\langle vi \rangle$ and $\langle eov \rangle$ where $\langle eok \rangle$ represents “end of keyword” and $\langle eov \rangle$ represents “end of value”. Note that we use the spatial databases and the grammar of executable queries to identify keywords and values without referring to the ground truth. Here i indicates it is the i -th spatial semantics that is injected. For a value $\langle v \rangle$, if ambiguity exists, we predict its spatial meaning using spatial comprehension model and feed the spatial semantics into the question using the symbol $\langle k \rangle$.

As shown in Figure 7, since the output of the seq2seq model involves symbols that are inserted into the question which need to be transformed to its original literal form, we propose a query recovery model. The detailed algorithm will be presented in Section 5.2.

Algorithm 3 Spatial Semantics Injection

```

1: function SPATIALINJECTION( $q, V, P$ )
2:    $s2p = \emptyset$ ; ▷ Symbol phrase mapping.
3:    $index_v = 0$ ; ▷ Value Symbol index.
4:    $index_k = 0$ ; ▷ Keyword symbol index.
5:    $q' = q$ ;
6:   for all  $\langle p, c \rangle \in P$  do ▷ Iterate each matched pairs.
7:     if  $c$  is a keyword then
8:        $sym = \langle k' \rangle + index_k$ ;
9:     else if  $c$  is a value then
10:       $sym = \langle v' \rangle + index_v$ ;
11:      Search for  $c.types$  from  $V$ ;
12:      if  $|c.types| > 1$  then ▷  $c$  is an ambiguous spatial
13:        phrase
14:         $T = p.types$ 
15:        for all  $t \in T$  do
16:          if SPATIALCOMPREHENSION( $q, p, t$ ) is True
17:            then
18:               $c.type = t$ ;
19:              else if  $|c.types| == 1$  then ▷  $c$  is not an ambiguous
20:                spatial phrase
21:                 $c.type = c.types[0]$ ;
22:                Insert  $c.type$  to  $q'$  (using symbol  $\langle k' \rangle + index_k$ )
23:                 $index_k = index_k + 1$ ;
24:                 $index_v = index_v + 1$ ;
25:                 $s2p.add(\langle sym, c \rangle)$ ;
26:                Insert  $sym$  to  $q'$ ;
27:   Return  $q', s2p$ ;

```

Following the aforementioned example in spatial comprehension $q = \text{“How many states does the Mississippi run through?”}$, if we do not address the ambiguity problem and rely on the seq2seq model to infer the spatial meaning of “Mississippi”, q' will be “How many $\langle k0 \rangle$ states $\langle eok \rangle$ does the $\langle v0 \rangle$ Mississippi $\langle eov \rangle$ run through?”. After translated by the seq2seq model, the recovered query is likely to be

Question q	What is the population of San Antonio ?
Keyword Detection	What is the <i>population</i> of San Antonio ?
Symbol q' Injection	what is the $\langle k0 \rangle$ population $\langle eok \rangle$ of $\langle k1 \rangle$ cityid $\langle eok \rangle$ $\langle v0 \rangle$ San Antonio $\langle eov \rangle$?
Seq2seq Model	
Output l'	answer(A, $\langle k0 \rangle$ (B, A), const(B, $\langle k1 \rangle$ ($\langle v0 \rangle$)))
Recover l	answer(A, population(B, A), const(B, cityid(San Antonio)))

Figure 7: An example of Information Injection Format with Type Feeding

answer(A, count(B, (state(B), const(C, stateid(mississippi)), traverse(C, B)), A)) since “states” is mentioned in the question. Our model is fundamentally built upon a seq2seq translation model, where the context is transformed into a weighted sum of all the tokens, in which “states” will be embedded as part of the context, and the model is easy to be confused and outputs “Mississippi” as a state name. We will mention in Section 5.1 later that, since the output vocabulary size is much smaller than the input vocabulary size and most of the tokens in output appear in the input as well, we adopt Copying Mechanism [37], where the output token has a higher chance to be copied from the input sequence. Copying Mechanism makes the ambiguity harder to address, such as in the aforementioned question, “river” does not appear in the input question, the model has a higher probability of copying “state” from the input sequence. Thus we need to insert ‘riverid’ before the word “Mississippi” in the input sentence to help the model make the translation. The final input will be “How many <k0> states <eok> does the <k1> riverid <eok> <v0> Mississippi <eov> run through?”.

5 TRANSLATE & RECOVER

5.1 Translation Model

Since seq2seq models have been widely adopted in translation tasks, and our NLI task is simpler than a translation task due to small vocabulary size. We believe a seq2seq model is able to capture the logic and the spatial structure of the question as long as it is able to understand the entities that are mentioned in the question. So we adopt a seq2seq model with copying mechanism following [9].

$$l' = \text{SEQ2SEQ}(q')$$

5.2 Query Recovery Model

We detail our strategy of query recovery through Algorithm 4, whose inputs are l' , the output of translation model, and $s2p$, the symbol-phrase pairs detected by spatial semantics injection model (Algorithm 3). Just as Figure 7 shows, the output l' of the seq2seq translation model is a sequence of the symbol, for example, <k0>, <k1> and <v0> here. Then we need to recover the output logic form query. For each pair $\langle \text{sym}, c \rangle \in s2p$, every symbol sym l' needs to be replaced by the original phrase c . After replacing all symbols in l' , we finally get the output logic form. In Figure 7, for output l' , we replace the symbols <k0>, <k1> and <v0> by their corresponding phrase “population”, “cityid” and “San Antonio” based on the pairs in $s2p$. After recovery, we finally get the right output logic form.

6 DATA AUGMENTATION BY SHUFFLING

Just as mentioned before, one of the Challenges right now is the lack of training set. The sparsity of training data causes two problems: 1). The semantic structures of questions are sparse; 2). The data entities mentioned in the questions are inadequate. Problem 2 can

Algorithm 4 Symbol Recovery

```

1: function RECOVER( $l'$ ,  $s2p$ )
2:    $l = l'$ ;
3:   for all  $\langle \text{sym}, c \rangle \in s2p$  do
4:      $l.\text{replace}(\text{sym}, c)$ ;
5:   Return  $l$ ;
```

be simply addressed by replacing data entities. However, addressing problem 1 is non-trivial. So we propose to shuffle the prepositional phrases to augment the semantic structures of the training set.

We propose our unique augmentation strategy as follows: If a question has a prepositional phrase (PP as a POS tag), and the question can be decomposed as $q = q^{\text{prefix}} | q^{\text{PP}}$ or $q = q^{\text{PP}} | q^{\text{suffix}}$, where q^{prefix} are the words placed before the prepositional phrase and q^{suffix} are the words placed after the prepositional phrase, we will shuffle the position of the prepositional phrase. Note that we only consider the questions that start with a prepositional phrase or end with a prepositional phrase.

Consider the example in Figure 8, for the query “Which states does the Mississippi river run through”, the format of the question is $q = q^{\text{prefix}} | q^{\text{PP}}$, and $q^{\text{prefix}} =$ “Which states does the Mississippi river run” and $q^{\text{PP}} =$ “through”. By exchanging q^{prefix} and q^{PP} , we can get a new sentence “Through which states does the Mississippi river run” and the meaning of the new query remains the same. Also, for the other question “In what state is Mount Mckinley”, $q = q^{\text{PP}} | q^{\text{suffix}}$, $q^{\text{PP}} =$ “In what state”, and $q^{\text{suffix}} =$ “is Mount Mckinley”. we can shuffle the position of the prepositional phrase to get a new sentence.

q	Which states does the Mississippi river run through ?
	q^{prefix} q^{PP}
Augment	Through which states does the Mississippi river run ?
	q^{PP} q^{prefix}
q	In what state is Mount Mckinley ?
	q^{PP} q^{suffix}
Augment	Mount Mckinley is in what state ?
	q^{suffix} q^{PP}

Figure 8: Two examples of data augmentation

7 EXPERIMENTAL VALIDATION

7.1 Experimental Settings

Configuration All our experiments are conducted on a machine equipped with 2 Intel CPU E5-2670 v3 running at 2.3GHz with 256GB of RAM and 2 NVIDIA Tesla K80 GPUs.

Dataset To evaluate the effectiveness of our system, we performed an experimental evaluation on dataset Geoquery and Restaurants.¹

- *Geoquery* [27] is a collection of 880 natural language questions and corresponding executable database query pairs about U.S. geography. The answers in this dataset are defined in λ -calculus logical form. We follow the standard training-test split to that of [19], of which the dataset was divided into 600 training examples and 280 test examples respectively. As [37], [26] and [48], we determine its *Acc* based on the denotation match.
- *Restaurant* (Rest) [6, 28] is a dataset with 251 question-answer pairs about restaurants, their food types, and locations. The questions are all human natural language and the answers are in λ -calculus logical form.

7.2 Data Augmentation

We not only propose our new data augmentation strategy by shuffling in Section 6, but also adopt a data augmentation strategy that is based on the recombination [37] of a sentence itself.

¹Our code is publicly available at <https://github.com/VV123/SpatialNLI>

Type 1	Original	What is the highest point in Florida ?
	Augment	What is the highest point in Rhode Island ?
Type 2	Original	What is the highest point in Florida ?
	Augment	what state has the smallest population density ?
Type 3	Original	What is the highest point in state that has the smallest population density ?
	Augment	what state has the largest population ?
Type 3	Original	what state has no rivers ?
	Augment	what state has the largest population and has no rivers ?

Figure 9: Three examples of data augmentation

For example, as in Figure 9, for the query “What is the highest point in Florida?”, we can simply identify that the word “Florida” is the name of a state based on the spatial database. Given this example, we change it to new questions, in which the word “Florida” is replaced by the name of other states in the database. Here, the word “Florida” is replaced by “Rhode Island”. For the second example in Figure 9, for the query “What is the highest point in Florida?” and the query “What state has the smallest population density?”, we can infer that the entire expression of the second sentence could map to the word “Florida” in the first query since this query is asking about one state. Then we can generate one new question by replacing the word “Florida” with the second sentence. For the third example, the two queries, “What state has the largest population?” and “What state has no rivers?”, are both asking about one state, so we combine them together to generate a new query.

7.3 Spatial Comprehension Model

We preprocess the dataset for spatial comprehension model so that each record contains (1) A question with each POI phrase enclosed with symbols (e.g., $\langle @ \rangle$) indicating the attentive position; (2) A POI type (e.g., River, State, and City.).

For a question “How many states does the Mississippi run through?” with one ambiguous POI “Mississippi”, we have the three records as shown in Figure 6. To balance the positive and negative samples in the training set, we replicate positive samples. For samples in Figure 6, we replicate positive samples by 2 times.

We run experiments with 200 hidden units and 300-dimensional pre-trained Glove embedding. We minimize the cross entropy using Adam Optimization Algorithm. We evaluated the performance of our spatial comprehension model in Table 1. Acc_{rcd} represents the percentage of correctly predicted records. Acc_{qu} represents the percentage of correctly predicted questions where all POIs are recognized correctly. For the example in Figure 6, the total number of samples for Acc_{rcd} is 3, and the total number of samples for Acc_{qu} is 1. Even the training objective function is to optimize Acc_{rcd} . In fact Acc_{qu} is what we are trying to optimize, and we prove that Acc_{rcd} and Acc_{qu} are optimized simultaneously.

We evaluate on Geoquery and Rest datasets, respectively (shown in Table 1). Test Acc_{qu} is 98.1% for Geoquery, and 100.0% for Restaurant data, respectively. All the Acc_{qu} is not less than Acc_{rcd} . In other words, our spatial comprehension model is able to recognize the spatial semantics with high confidence.

Dataset		Train	Test
Geoquery	Acc_{rcd}	97.4%	91.9%
	Acc_{qu}	98.3%	98.1%
Rest(aurant)	Acc_{rcd}	100.0%	100.0%
	Acc_{qu}	100.0%	100.0%

Table 1: Spatial Comprehension Model evaluation.

7.4 Evaluation

For the encoder and the decoder of our seq2seq model, we use one layer of Gated Recurrent Unit (GRU) with a hidden size of 800 and $800 * 2$, respectively. The input and output of both encoder and decoder share the same embedding layer, which is initialized with 300-dimensional pre-trained Glove embedding. Special symbols inserted (e.g., k_1 and v_1) are treated as special tokens; they are represented by the concatenation of an embedding of the symbol type (e.g., k and v) and an index, where the embedding of the symbol type and the index are randomly initialized with 150-dimension (the concatenation has a dimension of 300). The other unknown token is initialized with a 300-dimension random vector. For training, we use gradient clipping with a threshold 5.0, and for inference, we use beam search with width 5.

Table 2 presents our experiment results for (a) Geoquery dataset and (b) Rest(aurant) dataset. For Geoquery, compared with the previous models, our method outperforms the state-of-the-art. The conventional methods are overdependent on predefined templates and manually designed features, which have lower accuracy on the test set. For neural network-based methods such as ASN [35] and TRANX [16], they convert word space to action space, which inevitably introduces transformation error. SQL2TREE [36] and JL16 [37] use seq2seq model as well, but fail to address spatial semantics ambiguity. For Rest, the state-of-the-art achieves 100%

Dataset	Geoquery	
Conventional	ZC05 [19]	79.3%
	ZC07 [20]	86.1%
	UBL [21]	87.9%
	DCS+L [26]	87.9%
	FUBL [22]	88.6%
	ZH15 [25]	88.9%
Deep Model	KCAZ13 [23]	89.0%
	ASN [35]	87.1%
	SQL2TREE [36]	87.1%
	TRANX [16]	88.2%
Ours	JL16 [37]	89.3%
	SpatialNLI	90.4%
	– Copy Mechanism	88.9 %
	– Spatial Comprehension	86.4 %
	– Type Feeding	85.0 %
	– Data Augmentation	83.2 %
	– Information Injection	82.9 %

(a)

Dataset	Restaurant	
Conventional	PEK03 [6]	97.0%
	TM00 [28]	99.6%
Deep Model	FKZ18 [34]	100.0%
Ours	SpatialNLI	100.0%
	– Spatial Comprehension	96.1 %
	– Copy Mechanism	94.1 %
	– Data Augmentation	92.2 %
	– Type Feeding	70.6 %
	– Information Injection	60.8 %

(b)

Table 2: “–” means the removal of each component. The accuracy is measured as denotation match [37] on test set.

	q	How many states does the <i>Mississippi</i> run through ?
+SC	q'	$\langle k0 \rangle$ How many $\langle eok \rangle$ $\langle k1 \rangle$ states $\langle eok \rangle$ does the $\langle k2 \rangle$ <i>riverid</i> $\langle eok \rangle$ $\langle v0 \rangle$ <i>Mississippi</i> $\langle eov \rangle$ run $\langle k3 \rangle$ through $\langle eok \rangle$?
	Infer	answer(A, $\langle k0 \rangle$ (B, ($\langle k1 \rangle$ (B), const(C, $\langle k2 \rangle$ ($\langle v0 \rangle$))), $\langle k3 \rangle$ (C, B)), A))
	Recover	answer(A, count(B, (state(B), const(C, <i>riverid</i> (<i>Mississippi</i>)), traverse(C, B)), A))
-SC	q'	$\langle k0 \rangle$ How many $\langle eok \rangle$ $\langle k1 \rangle$ states $\langle eok \rangle$ does the $\langle k2 \rangle$ <i>stateid</i> $\langle eok \rangle$ $\langle v0 \rangle$ <i>mississippi</i> $\langle eov \rangle$ run $\langle k3 \rangle$ through $\langle eok \rangle$?
	Infer	answer(A, $\langle k0 \rangle$ (B, ($\langle k1 \rangle$ (B), const(C, $\langle k2 \rangle$ ($\langle v0 \rangle$))), $\langle k3 \rangle$ (C, B)), A))
	Recover	answer(A, count(B, (state(B), const(C, <i>stateid</i> (<i>Mississippi</i>)), traverse(C, B)), A))

Figure 10: Spatial Comprehension Case Study. * +SC means using Spatial Comprehension, -SC means without.

accuracy, which states the Rest dataset is an easier task than Geoquery. Our model exhibits excellent downward compatibility by achieving 100% accuracy on Rest dataset.

To validate the performance of our system, several ablation experiments were conducted by the removal of (1) Copy Mechanism, (2) Spatial Comprehension Model, (3) Data Augmentation, (4) Type Feeding, and (5) Information Injection, respectively. By the removal of the spatial comprehension model, we random guess the meaning (type) of ambiguous POI and inject it to the question. For removing information injection, we feed the original content to the model without inserting any symbols. By the removal of type feeding, we conduct symbol injection but omit to inject the extra spatial information (e.g., $\langle k1 \rangle$ cityid $\langle eok \rangle$ in Figure 11).

First, we measure the contribution of the spatial comprehension mechanism to the overall performance of the model. We train and evaluate two models: one with the spatial comprehension model and one without. Training is done with data augmentation and information injection. In Table 2, for Geoquery and Restaurant, with the removal of spatial comprehension model, the denotation match accuracy drops 4% for Geo and drops 3.9% for Rest. Since only 19.3% of the test set for Geoquery and 4% of the test set for Restaurant has POI ambiguity problem, it is obvious that our machine comprehension model is able to resolve the majority of them.

As shown in Figure 10, by comparing against the spatial database, “*Mississippi*” appears in two tables: River table and State table. Without spatial comprehension, if we are using random guess, “*river*” has only 50% chance to be correctly categorized. The ‘+comprehension’ in the figure means we use the spatial comprehension model and ‘-comprehension’ is for the result without a right understanding of “*Mississippi*”. Without the spatial comprehension model, it is possible for the system to recognize the “*Mississippi*” as a state name. As the figure shows, once “*Mississippi*” is recognized as a state, it will insert “*stateid*” in the input question and finally get a wrong result after recovery. One interesting thing is that the infer for ‘+comprehension’ and ‘-comprehension’ are the same, both correct. This is because, for seq2seq model, it just outputs the result with $\langle ki \rangle$, not the specific word. Here the “*riverid*” and “*stateid*” are both replaced by $\langle k2 \rangle$. Thus we get the same infer result from seq2seq model. However, after recovery, the result without spatial comprehension model is wrong.

Table 2 shows that by removing type feeding, the accuracies drop 5.4% on Geoquery and 29.4% on Restaurant. The symbol injection significantly improves the accuracy of the Restaurant dataset since for most samples in the Restaurant dataset, one token in the input question always corresponds to multiple tokens/symbols in the output sequence, which relationships are hard for the seq2seq model to capture.

-SI	q	What is the population of San Antonio ?
	Infer	answer(A, population(B, A), const(B, cityid(<i>San Jose</i>)))
+SI	q	What is the population of San Antonio ?
	q'	what is the $\langle k0 \rangle$ population $\langle eok \rangle$ of $\langle k1 \rangle$ cityid $\langle eok \rangle$ $\langle v0 \rangle$ San Antonio $\langle eov \rangle$?
	Infer	answer(A, $\langle k0 \rangle$ (B, A), const(B, $\langle k1 \rangle$ ($\langle v0 \rangle$)))
	Recover	answer(A, population(B, A), const(B, cityid(<i>San Antonio</i>)))

Figure 11: A Case Study using Symbol Inject (Geoquery). * +/- SI means with/without Symbol Injection.

As shown in Table 2, the information injection component improves test accuracy by 7.5% on Geoquery and 39.2% on Restaurant. When we stop injecting information into the natural language question, the seq2seq is not able to capture all the necessary information to infer correctly and suffers from a large accuracy decrease. A case study of our symbol injection strategy is shown in Figure 11, where a seq2seq model generates outputs token by token and a large number of entities involve a sequence of tokens. Without symbol injection, the seq2seq model has to infer “*San Antonio*” token by token using two steps. On the other hand, with symbol injection, the seq2seq model generates $v0$ as a representation of “*San Antonio*”, which only requires one step. Our symbol injection format is able to replace a name entity composed of a sequence of tokens to a single symbol, which prevents wrong name entity caused by a long sequential generation.

We also jointly train both datasets in a shared model compared with separate training, shown in Table 3. Jointly training achieves an accuracy of 90.7%. Our experiment results show that a shared model performs better than two separate models.

Training	Geoquery	Restaurant
Separately	90.4%	100%
Jointly	90.7%	100%

Table 3: Evaluation of jointly training on denotation match.

8 CONCLUSION

In this work, we propose an NLIDB applied for the spatial domain to convert natural language queries to structured queries executable by database. The main contribution of our work is to recognize the meaning of the ambiguous spatial phrases based on contextual interpretation and capture the semantic structure of the question by the seq2seq model with injecting spatial information. Our extensive experimental analysis demonstrates the advantage of our approach over state-of-the-art methods.

9 ACKNOWLEDGEMENT

This research has been funded in part by the U.S. National Science Foundation grants IIS-1618669 (III) and ACI-1642133 (CICI).

REFERENCES

- [1] Ion Androutsopoulos, Graeme D Ritchie, and Peter Thanisch. Natural language interfaces to databases—an introduction. *Natural language engineering*, 1(1):29–81, 1995.
- [2] Florin Brad, Radu Cristian Alexandru Iacob, Ionel-Alexandru Hosu, and Traian Rebedea. Dataset for a neural natural language interface for databases (NNLIDB). In *Proceedings of the Eighth International Joint Conference on Natural Language Processing, IJCNLP*, pages 906–914, 2017.
- [3] Prasetya Utama, Nathaniel Weir, Fuat Basik, Carsten Binnig, Ugur Çetintemel, Benjamin Hättasch, Amir Ilkhechi, Shekar Ramaswamy, and Arif Usta. An end-to-end neural natural language interface for databases. *CoRR*, abs/1804.00401, 2018.
- [4] Fei Li and Hosagrahar Visvesvaraya Jagadish. Nalir: an interactive natural language interface for querying relational databases. In *International Conference on Management of Data, SIGMOD*, pages 709–712, 2014.
- [5] Yunyao Li, Huahai Yang, and H. V. Jagadish. Nalix: an interactive natural language interface for querying XML. In *Proceedings of the ACM SIGMOD International Conference on Management of Data*, pages 900–902, 2005.
- [6] Ana-Maria Popescu, Oren Etzioni, and Henry A. Kautz. Towards a theory of natural language interfaces to databases. In *Proceedings of the 8th International Conference on Intelligent User Interfaces*, 2003.
- [7] Diptikalyan Saha, Avrilia Floratou, Karthik Sankaranarayanan, Umar Farooq Minhas, Ashish R. Mittal, and Fatma Özcan. ATHENA: an ontology-driven system for natural language querying over relational data stores. *PVLDB*, 9(12):1209–1220, 2016.
- [8] Srinivasan Iyer, Ioannis Konstantas, Alvin Cheung, Jayant Krishnamurthy, and Luke Zettlemoyer. Learning a neural semantic parser from user feedback. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL*, pages 963–973, 2017.
- [9] Wenlu Wang, Yingtao Tian, Hongyu Xiong, Haixun Wang, and Wei-Shinn Ku. A transfer-learnable natural language interface for databases. *arXiv preprint arXiv:1809.02649*, 2018.
- [10] Wenlu Wang. A cross-domain natural language interface to databases using adversarial text method. *Database*, 1:q2, 2019.
- [11] Wenlu Wang, Ji Zhang, M-T Sun, and W-S Ku. Efficient parallel spatial skyline evaluation using mapreduce. In *Proceedings of the 20th international conference on extending database technology*, 2017.
- [12] Wenlu Wang, Ji Zhang, Min-Te Sun, and Wei-Shinn Ku. A scalable spatial skyline evaluation system utilizing parallel independent region groups. *The VLDB Journal The International Journal on Very Large Data Bases*, 28(1):73–98, 2019.
- [13] Wenlu Wang and Wei-Shinn Ku. Dynamic indoor navigation with bayesian filters. *SIGSPATIAL Special*, 8(3):9–10, 2017.
- [14] Wenlu Wang and Wei-Shinn Ku. Recommendation-based smart indoor navigation. In *Proceedings of the Second International Conference on Internet-of-Things Design and Implementation*, pages 311–312. ACM, 2017.
- [15] Jordan Zlatev. Spatial semantics. *The Oxford handbook of cognitive linguistics*, pages 318–350, 2007.
- [16] Pengcheng Yin and Graham Neubig. TRANX: A transition-based neural abstract syntax parser for semantic parsing and code generation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, EMNLP*, pages 7–12, 2018.
- [17] Min Joon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi. Bidirectional attention flow for machine comprehension. In *5th International Conference on Learning Representations, ICLR*, 2017.
- [18] Shuohang Wang and Jing Jiang. Machine comprehension using match-lstm and answer pointer. In *5th International Conference on Learning Representations, ICLR*, 2017.
- [19] Luke S. Zettlemoyer and Michael Collins. Learning to map sentences to logical form: Structured classification with probabilistic categorial grammars. In *UAI '05, Proceedings of the 21st Conference in Uncertainty in Artificial Intelligence*, 2005.
- [20] Luke Zettlemoyer and Michael Collins. Online learning of relaxed ccg grammars for parsing to logical form. In *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*, 2007.
- [21] Tom Kwiatkowski, Luke S. Zettlemoyer, Sharon Goldwater, and Mark Steedman. Inducing probabilistic CCG grammars from logical form with higher-order unification. In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing, EMNLP*, pages 1223–1233, 2010.
- [22] Tom Kwiatkowski, Luke S. Zettlemoyer, Sharon Goldwater, and Mark Steedman. Lexical generalization in CCG grammar induction for semantic parsing. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing, EMNLP*, pages 1512–1523, 2011.
- [23] Tom Kwiatkowski, Eunsol Choi, Yoav Artzi, and Luke Zettlemoyer. Scaling semantic parsers with on-the-fly ontology matching. In *Proceedings of the 2013 conference on empirical methods in natural language processing*, 2013.
- [24] Adrienne Wang, Tom Kwiatkowski, and Luke Zettlemoyer. Morpho-syntactic lexical generalization for ccg semantic parsing. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2014.
- [25] Kai Zhao and Liang Huang. Type-driven incremental semantic parsing with polymorphism. In *NAACL HLT 2015, The 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1416–1421, 2015.
- [26] Percy Liang, Michael I. Jordan, and Dan Klein. Learning dependency-based compositional semantics. *Computational Linguistics*, 39(2):389–446, 2013.
- [27] John M Zelle and Raymond J Mooney. Learning to parse database queries using inductive logic programming. In *Proceedings of the national conference on artificial intelligence*, pages 1050–1055, 1996.
- [28] Lappoon R Tang and Raymond J Mooney. Automated construction of database interfaces: Integrating statistical and relational learning for semantic parsing. In *Proceedings of the 2000 Joint SIGDAT conference on Empirical methods in natural language processing and very large corpora*, pages 133–141. Association for Computational Linguistics, 2000.
- [29] Lappoon R Tang and Raymond J Mooney. Using multiple clause constructors in inductive logic programming for semantic parsing. In *European Conference on Machine Learning*, pages 466–477. Springer, 2001.
- [30] Ruifang Ge and Raymond J Mooney. A statistical semantic parser that integrates syntax and semantics. In *Proceedings of the ninth conference on computational natural language learning*. Association for Computational Linguistics, 2005.
- [31] Nal Kalchbrenner and Phil Blunsom. Recurrent continuous translation models. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1700–1709, 2013.
- [32] Ilya Sutskever, Oriol Vinyals, and Quoc V Le. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pages 3104–3112, 2014.
- [33] Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014.
- [34] Catherine Finegan-Dollak, Jonathan K. Kummerfeld, Li Zhang, Karthik Ramamathan, Sesh Sadasivam, Rui Zhang, and Dragomir R. Radev. Improving text-to-sql evaluation methodology. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL*, pages 351–360, 2018.
- [35] Maxim Rabinovich, Mitchell Stern, and Dan Klein. Abstract syntax networks for code generation and semantic parsing. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL*, pages 1139–1149, 2017.
- [36] Li Dong and Mirella Lapata. Language to logical form with neural attention. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL*, 2016.
- [37] Robin Jia and Percy Liang. Data recombination for neural semantic parsing. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL*, 2016.
- [38] Xing Fan, Emilio Monti, Lambert Mathias, and Markus Dreyer. Transfer learning for neural semantic parsing. In *Proceedings of the 2nd Workshop on Representation Learning for NLP*, pages 48–56, 2017.
- [39] Raymond Hendy Susanto and Wei Lu. Neural architectures for multilingual semantic parsing. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, pages 38–44, 2017.
- [40] Jonathan Herzig and Jonathan Berant. Neural semantic parsing over multiple knowledge-bases. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL*, pages 623–628, 2017.
- [41] Victor Zhong, Caiming Xiong, and Richard Socher. Seq2sql: Generating structured queries from natural language using reinforcement learning. *CoRR*, abs/1709.00103, 2017.
- [42] Qijun Shen, Xueying Zhang, and Wenming Jiang. Annotation of spatial relations in natural language. In *2009 International Conference on Environmental Science and Information Application Technology*, volume 3, pages 418–421. IEEE, 2009.
- [43] John A Bateman, Joana Hois, Robert Ross, and Thora Tenbrink. A linguistic ontology of space for natural language processing. *Artificial Intelligence*, 174(14):1027–1071, 2010.
- [44] Parisa Kordjamshidi, Paolo Frasconi, Martijn Van Otterlo, Marie-Francine Moens, and Luc De Raedt. Relational learning for spatial relation extraction from natural language. In *International Conference on Inductive Logic Programming*, pages 204–220. Springer, 2011.
- [45] Arbaz Khan, Maria Vasardani, and Stephan Winter. Extracting spatial information from place descriptions. In *COMP@ SIGSPATIAL*, page 62, 2013.
- [46] Tiago Ramalho, Tomáš Kociský, Frederic Besse, SM Eslami, Gábor Melis, Fabio Viola, Phil Blunsom, and Karl Moritz Hermann. Encoding spatial relations from natural language. *arXiv preprint arXiv:1807.01670*, 2018.
- [47] Jeffrey Pennington, Richard Socher, and Christopher Manning. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543, 2014.
- [48] Yushi Wang, Jonathan Berant, and Percy Liang. Building a semantic parser overnight. 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, pages 1332–1342, 2015.