

Semantic Mapping between Natural Language Questions and SQL Queries via Syntactic Pairing

Alessandra Giordani and Alessandro Moschitti

Department of Computer Science and Engineering
University of Trento
Via Sommarive 14, 38100 POVO (TN) - Italy
{agiordani,moschitti}@disi.unitn.it

Abstract. Automatically mapping natural language semantics into programming languages has always been a major and interesting challenge in Computer Science. In this paper, we approach such problem by carrying out mapping at syntactic level and then applying machine learning algorithms to derive an automatic translator of natural language questions into their associated SQL queries. To build the required training and test sets, we designed an algorithm, which, given an initial corpus of questions and their answers, semi-automatically generates the set of possible incorrect and correct pairs.

We encode such relational pairs in Support Vector Machines by means of kernel functions applied to the syntactic trees of questions and queries. The accurate results on automatic classification of the above pairs above, suggest that our approach captures the shared semantics between the two languages.

Keywords: Natural Language Processing; Kernel Methods; Support Vector Machines.

1 Introduction

The design of models for automatically mapping natural language semantics into programming languages has been always a major and interesting challenge in Computer Science since it would have a direct impact on industrial and social worlds. For example, accessing a database requires machine-readable instructions that common users are not supposed to know. Ideally, they should only pose a question in natural language without knowing either the underlying database schema or any complex machine language. The development of natural language interfaces to databases (NLIDBs) that translate the human intent into SQL instructions is indeed a classic problem, which is becoming of greater importance in today's world.

This could be addressed by finding a mapping between natural language and the database programming language. If we knew how to convert natural language questions into their associated SQL queries, it would be straightforward to obtain the answers by just executing a query. Unfortunately, previous work in

Natural Language Understanding has shown the inadequacy of logic and rule-based approaches to this problem; in contrast shallow and statistical methods appear to be promising.

In this paper, we exploit mapping at syntactic level between the two languages and apply machine learning models to derive the shared shallow semantics. Such approach requires the design of a dataset of relational pairs containing syntactic trees of questions and queries. For this purpose, we used syntactic parsers to obtain natural language and SQL trees and we designed an effective algorithm, which, given an initial corpus of correct question and query pairs, semi-automatically generates the labeled set of possible correct and incorrect instances.

We used the above dataset to train classifiers based on Support Vector Machines and kernel functions over pairs. Such functions are combinations of tree kernels applied to syntactic trees and linear kernels applied to bag-of-words. The cross-validation experiments on the task of selecting correct queries given a target set of questions show that our best kernel improves the baseline model of about 32%. The latter is the typical approach based on a linear kernel applied to the union of the bag-of-words from question and query texts. The most interesting finding is that the product between the two kernels representing questions and queries provides feature pairs, which can express the relational features between the syntactic/semantic representation of the two languages.

In the remainder, Section 2 introduces the problem of mapping questions into queries and illustrates the idea of our solution whereas Section 3 describes the technology to implement it, i.e. kernel methods. Section 4 shows our proposed algorithm to generate a training set of question and query pairs, Section 5 discusses our results and finally, Section 6 draws conclusions.

2 Automatic Mapping of Questions into SQL Queries

Studying the automatic mapping of questions into SQL queries is important for two main reasons: (a) it allows to design interesting applications based on databases and (b) it offers the possibility to understand the role of syntax in deriving a shared semantics between a natural language and an artificial language.

Given the complexity of theoretically modeling such relationship we use a statistical and shallow model. We consider a dataset of natural language questions \mathcal{N} and SQL queries \mathcal{S} related to a specific domain/database¹ and we automatically learn such mapping from the set of pairs $\mathcal{P} = \mathcal{N} \times \mathcal{S}$. More in detail, (a) we assume that pairs are annotated as correct when the SQL query answers to the question and incorrect otherwise and (b) we train a classifier on the pairs above for selecting the correct queries for a question. Then, to map new questions in the set of the available queries, (c) we rank the latter by means of the question classifier score and select the top one. In the following we provide the formal definition of our learning approach.

¹ We assume that for any database there is a core set \mathcal{S} of queries, which are frequently asked. \mathcal{S} should at least represent the syntactic structures of the most part of the frequent queries.




Fig. 1. Question/Query Syntactic trees

2.1 Pair Ranking

The problem of assigning a query (with its result) to a question, can be formally described as the following ranking problem: (i) given a question $n \in \mathcal{N}$ and a set of possible useful queries \mathcal{S} , we generate the set of possible pairs $P(n) = \{(n, s) : s \in \mathcal{S}\}$; (ii) we classify them with an automatic categorizer; (iii) we use the score/probability output by such model to rank $P(n)$; (vi) we select the top ranked pair.

For example, let us consider question n_1 : “Which states border Texas?” and the queries s_1 : `SELECT state_name FROM border_info WHERE border='texas'` and s_2 : `SELECT COUNT(state_name) FROM border_info WHERE border='texas'`. Since s_1 is a correct and s_2 is an incorrect interpretation of the question, the classifier should assign a higher score to the former, thus our ranker will output the $\langle n_1, s_1 \rangle$ pair. Note that both s_1 and s_2 share three terms, *state*, *border* and *texas*, with n_1 but $\langle n_1, s_2 \rangle$ is not correct. This suggests that we cannot only rely on the common terms but we should also take into account the syntax of both languages.

2.2 Pair Representation

The aim of our research is to derive the shared shallow semantics within pairs by means of syntax. Thus we represent questions and queries using their syntactic trees², as shown in Figure 1: for the question (a) we use the output of the Charniak’s syntactic parser [2], whereas for the query (b) we use a modification of the SQL derivation tree.

To build the SQL tree we implemented an ad-hoc parser that follows the syntactic derivation of a query according to our grammar. Since our database system embeds a MySQL server, we use the production rules of MySQL, shown at the top of Figure 2, slightly modified to manage punctuation, i.e. rules 5*, 6* and 20* related to comma and dot, as shown at the bottom.

More in detail, we change the non-terminals *Item* and *SelectItem* with the symbol • to have an uniform representation of the relationship between a table and its column in both the `SELECT` and `WHERE` clauses. This allows for matching between the subtrees containing table, column or both also when they appear in different clause types of two queries.

² Early work on the use of syntax for text categorization were based on part-of-speech tags, e.g. [1]. The efficiency of modern syntactic parsers allows us to use the complete parse tree.

```

1  SQL → SELECT ItemList FROM TableList |
      SELECT ItemList FROM TableList WHERE Cond
...
5  ItemList → ItemList , Item | Item
6  Item → Table . Column | Column
7  Column → * | ColumnName
...
15 WhereCondition → AndCondition AND AndCondition
16 AndCondition → Condition OR Condition
17 Condition → NOT Condition | Operand | ...
18 Operand → Factor | Operand + Factor | ...
19 Factor → Term | Factor * Term | ...
20 Term → Value | SelectItem | SQL
21 SelectItem → Table . ColumnName | ColumnName


---


5* ItemList → ItemList • | •
6*  • → Table | Column | Table Column
20* Term → Value | • | SQL

```

Fig. 2. Modified MySQL Grammar

It is worth noting that rule 20* still allows to parse nested queries and that the overall grammar, in general, is very expressive and powerful enough to express complex SQL queries involving nesting, aggregation, conjunctions and disjunctions in the `WHERE` clause.

Note that, although we eliminated comma and dot from the original SQL grammar, it is still possible to obtain the original SQL query by just performing a preorder traversal of the tree.

To represent the above structures in a learning algorithm we use tree kernels described in the following section.

3 Tree Kernels

Kernel Methods refer to a large class of learning algorithms based on inner product vector spaces, among which Support Vector Machines (SVMs) are one of the most well-known approaches. The main idea is that the parameter model vector \mathbf{w} generated by SVMs (or by other kernel-based machines) can be rewritten as

$$\sum_{i=1..l} y_i \alpha_i \mathbf{x}_i, \quad (1)$$

where y_i is equal to 1 for positive and -1 for negative examples, $\alpha_i \in \Re$ with $\alpha_i \geq 0$, $\forall i \in \{1, \dots, l\}$ \mathbf{x}_i are the training instances.

Therefore we can express the classification function as

$$\text{Sgn}\left(\sum_{i=1..l} y_i \alpha_i \mathbf{x}_i \cdot \mathbf{x} + b\right) = \text{Sgn}\left(\sum_{i=1..l} y_i \alpha_i \phi(o_i) \cdot \phi(o) + b\right), \quad (2)$$

where \mathbf{x} is a classifying object, b is a threshold and the product $K(o_i, o) = \langle \phi(o_i) \cdot \phi(o) \rangle$ is the kernel function associated with the mapping ϕ .

Note that it is not necessary to apply the mapping ϕ , we can use $K(o_i, o)$ directly. This allows, under the Mercer’s conditions [3], for defining abstract functions, which generate implicit feature spaces. The latter allow for an easier feature extraction and the use of huge feature spaces (possibly infinite), where the scalar product (i.e. $K(\cdot, \cdot)$) is implicitly evaluated.

In the remainder of this section, we illustrate some kernels for structured data: the Syntactic Tree Kernel (STK) [4], which computes the number of syntactic tree fragments and the Extended Syntactic Tree Kernel (STK_e) [5], which includes leaves in STK. In the last subsection, we show how to engineer new kernels from them.

3.1 Syntactic Tree Kernel (STK) and Its Extension (STK_e)

The main underlying idea of tree kernels is to compute the number of common substructures between two trees T_1 and T_2 without explicitly considering the whole fragment space. Let $\mathcal{F} = \{f_1, f_2, \dots, f_{|\mathcal{F}|}\}$ be the set of tree fragments and $\chi_i(n)$ an indicator function equal to 1 if the target f_i is rooted at node n and equal to 0 otherwise. A tree kernel function over T_1 and T_2 is defined as

$$TK(T_1, T_2) = \sum_{n_1 \in N_{T_1}} \sum_{n_2 \in N_{T_2}} \Delta(n_1, n_2), \quad (3)$$

where N_{T_1} and N_{T_2} are the sets of nodes in T_1 and T_2 , respectively, and $\Delta(n_1, n_2) = \sum_{i=1}^{|\mathcal{F}|} \chi_i(n_1)\chi_i(n_2)$.

The Δ function is equal to the number of common fragments rooted in nodes n_1 and n_2 , and thus, depends on the fragment type. We report its algorithm for the evaluation of the number of syntactic tree fragments (STFs) [4].

A syntactic tree fragment (STF) is a set of nodes and edges from the original tree, which is still a tree and with the constraint that any node must have all or none of its children. This is equivalent to state that the production rules contained in the STF cannot be partial.

To compute the number of common STFs rooted in n_1 and n_2 , the STK uses the following Δ function [4]:

1. if the productions at n_1 and n_2 are different then $\Delta(n_1, n_2) = 0$;
2. if the productions at n_1 and n_2 are the same, and n_1 and n_2 have only leaf children (i.e. they are pre-terminal symbols) then $\Delta(n_1, n_2) = \lambda$;
3. if the productions at n_1 and n_2 are the same, and n_1 and n_2 are not pre-terminals then $\Delta(n_1, n_2) = \lambda \prod_{j=1}^{l(n_1)} (1 + \Delta(c_{n_1}(j), c_{n_2}(j)))$, where $l(n_1)$ is the number of children of n_1 , $c_n(j)$ is the j -th child of node n and λ is a decay factor penalizing larger structures.

Figure 3.a shows some STFs of the left tree in Figure 1. STFs satisfy the constraint that grammatical rules cannot be broken.

STK does not include individual nodes as features. As shown in [5] we can include at least the leaves, (which in constituency trees correspond to words) by simply inserting the following step 0 in the algorithm above:

0. if n_1 and n_2 are leaf nodes and their labels are identical then $\Delta(n_1, n_2) = \lambda$;




Fig. 3. Feature spaces for the tree pair in Figure 1 a) joint space STK+STK b) Cartesian product STK×STK

3.2 Kernel Engineering

Kernel engineering [6,7,8] can be carried out by combining basic kernels with additive or multiplicative operators or by designing specific data objects, e.g. the tree representation for the SQL syntax, to which standard kernels are applied. Since our data is a set of pairs, we need to represent the members of a pair and their interdependencies. For this purpose, given two kernel functions, $k_1(.,.)$ and $k_2(.,.)$, and two pairs, $p_1 = \langle n_1, s_1 \rangle$ and $p_2 = \langle n_2, s_2 \rangle$, a first approximation is given by summing the kernels applied to the components: $K(p_1, p_2) = k_1(n_1, n_2) + k_2(s_1, s_2)$. This kernel will produce the union of the feature spaces of questions and queries (e.g. see [9,10]). For example, the explicit vector representation of the STK + STK space of the pair in Figure 1 is shown in Figure 3.a. The Syntactic Tree Fragments of the question will be in the same space of the Syntactic Tree Fragments of the query.

In theory a more effective kernel is the product $k(n_1, n_2) \times k(s_1, s_2)$ since it generates pairs of fragments as features, where the overall space is the Cartesian product of the used kernel spaces. For example Figure 3.b shows pairs of STF fragments, which are essential to capture the relational semantics between the syntactic tree subparts of the two languages (see [11]). In particular, the first fragment pair of the figure may suggest that a noun phrase composed by *state* expresses similar semantics of the syntactic construct `SELECT state_name`.

As additional feature and kernel engineering, we also exploit the ability of the polynomial kernel to add feature conjunctions. By simply applying the function $(1 + K(p_1, p_2))^d$, we can generate conjunction up to d features. Thus, we can obtain tree fragment conjunctions and conjunctions of pairs of tree fragments.

4 Dataset Generation

In the previous sections we have defined models to automatically learn the mapping between questions and queries. To apply such models we need training data, i.e. correct and incorrect pairs of questions and queries. Since acquiring such kind of data is the most costly aspect of our design, we should generate the learning set in a smart way. In this perspective, we assume that, in real world domains, we may find examples of questions and the associated queries answering to their information need. Such pairs may have been collected when users and operators of the database worked together for the accomplishment of

some tasks. In contrast, we cannot assume to have available pairs of incorrect examples, since (a) the operator tends to just provide the correct query and (b) both users and operators do not really understand the use of negative examples and the need to have unbiased distribution of them.

Therefore, we need techniques to generate negative examples from an initial set of correct pairs. Unfortunately, this is not a trivial task since when mixing a question and a query belonging to different pairs we cannot assume to only generate incorrect pairs, e.g. when swapping x with y in the two pairs $\langle \textit{Which states border Texas?}, x \rangle$ and $\langle \textit{What are the states bordering Texas?}, y \rangle$, we obtain other two correct pairs.

To generate a gold standard dataset we would need to manually check this aspect thus we design an algorithm to limit the human supervision. It consists of the following steps:

- Generalizing question and query instances: substitute the involved concepts in questions and their related field values in the SQL queries by means of variables (expressing the category of such values).
- Clustering the generalized pairs: intuitively each cluster represents the information need about a target semantic concept, e.g. “bordering state”, common to questions and queries. This requires a limited manual intervention.
- Pairing questions and queries of distinct clusters, i.e. the Cartesian product between the set of questions and the set of queries belonging to the pairs of a target cluster. This allows to find new positive examples that were not present in the initial corpus.
- Final dataset annotation: consider all possible pairs, i.e. Cartesian product between all the questions and queries of the dataset, and annotate them as negatives if they have not been annotated as positives in the previous step.

We use the GEOQUERIES250³ corpus translated by Popescu et al. [12] as our initial dataset. It consists of 250 pairs of NL questions and SQL queries over a small database about United States geography. In the following we describe in detail all the steps through which the final dataset is generated.

4.1 Pair Generalization

Our approach to automatically annotate pairs relies on automatically detecting if swapping the members of different pairs produces correct or incorrect examples. For this purpose, we detect similar syntactic structures of questions and queries by generalizing concept instances with variables.

In a database concepts are represented as tables’ fields. These are used in SQL to select data satisfying some conditions, i.e. concept constrained to a value.

Typically these values are natural language terms so we substitute them with variables if they appear in both questions and queries. For example, consider s_1 in Figure 1. The condition is `WHERE state_name = 'Texas'` and 'Texas' is

³ <http://www.cs.utexas.edu/users/ml/nldata.html>

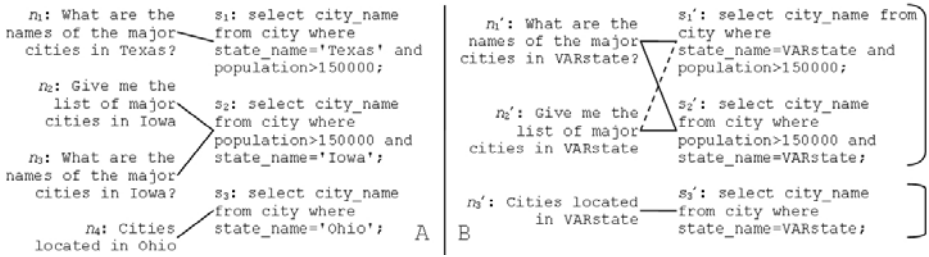


Fig. 4. Example of the initial corpus (*A*, on the left) and the generalized version (*B*, on the right). The latter is divided in two clusters (identified by the two brackets).

the value of the concept `state_name`. Since 'Texas' is also present in the related question we can substitute it with a variable *VARstate* (one variable for each different concept). Our assumption is that questions whose answer can be retrieved in a database tend to use the same terms stored in the database.

An example of the generalization phase is shown in Figure 4. On the left there is a set of four pairs containing four distinct questions and three related queries (connected by the lines) whereas on the right four generalized pairs are shown. We note that, after substituting instances with variables, both n_1 and n_3 are generalized into n'_1 , which is thus paired with two distinct SQL queries, i.e. s'_1 and s'_2 . This is not surprising since there can be multiple SQL queries that correctly retrieve an answer to a NL question. In this case we define them to be *semantically equivalent*, i.e. $s'_1 \equiv s'_2$. At the same time it is possible to write many NL questions that map to the same query.

It is worth noting that with the generalization process, we introduce redundancy that we eliminate by removing duplicated questions and queries. Thus, the output dataset is usually smaller than the initial one. However the number of training examples will be larger, not only because of the introduction of negatives but also due to the automatic discovering of new positives.

4.2 Pair Clustering and Final Dataset Annotation

Once the pairs have been generalized, we cluster them according to their semantic equivalence so that we can automatically derive new positive examples by swapping their members. We define semantically equivalent pairs those correct pairs with (a) equivalent NL questions or (b) equivalent SQL queries. Given that two equivalent queries should retrieve the same result set, we can automatically test their equivalence by simply executing them. Unfortunately, this is just a necessary condition (e.g. 'Texas' can be the answer of two different queries) therefore we manually evaluate new pairings satisfying this condition.

Note that automatically detecting semantic equivalence of natural language questions with perfect accuracy is a hard task, so we consider as semantically equivalent either identical questions or those associated with semantic equivalent queries. We also apply transitivity closure to both members of pairs to extend the set of equivalent pairs.


```

function CreateFinalDataset(LIST OF PAIRS  $I$ ) returns MATRIX  $M$ 
 $M[\mathcal{N}][\mathcal{S}] = -1$  // stores the cluster id for each pair, it is initialized to -1
 $k=0$  // at the beginning the number of cluster is zero
begin
 $\forall \langle n, s \rangle \in I$  s.t.  $M[n][s] = -1$  do // for all initial pairs not clustered yet
  begin
     $k:=k+1$  // add a new cluster
     $M[n][s]=k$  // the pair  $\langle n, s \rangle$  belongs to the new cluster
     $\forall \langle n', s' \rangle \in I$  do // find other pairs that belong to it
      if  $(n = n')$  OR  $(s = s')$  // if identical (indeed equivalent)
        then  $M[n'][s']=k$  // also  $\langle n', s' \rangle$  belongs to the cluster
      else // check equivalence if similar result set
        if  $(\text{fields}(s) = \text{fields}(s'))$  AND  $(\text{min}(s) = \text{min}(s'))$ 
          then if  $\text{res}(s) = \text{res}(s')$  // if  $s$  and  $s'$  retrieve the same result set
            then if  $n = n'$  // if the associated sentences are equivalent
              then  $M[n'][s']=k$  // label also the pair  $\langle n', s' \rangle$  with this cluster id
          end
        end
      end
    end
   $\forall k$  do // find all pairings within a cluster
     $\forall n, n', s, s'$  s.t.  $M[n][s]=k$  AND  $M[n'][s']=k$ 
      begin
         $M[n][s']=k$  // swap all queries
         $M[n'][s]=k$  // swap all questions
      end
    end
  return  $M$ 
end

```

Fig. 5. Clustering Algorithm

For example, in Figure 4.b s'_1 and s'_2 retrieve the same result set so we verify that they are semantically equivalent queries and we assign them to the same cluster (CL1), i.e. information need about the large cities of a state (with a population larger than 150,000 people). Alternatively, we can also consider that n'_1 and n'_2 are both paired with s'_2 to derive that they are equivalent, avoiding the human intervention. Concerning s'_3 , it retrieves a result set different from the previous one so we can automatically assign it to a different cluster (CL2), i.e. involving questions about any city of a state. Note that, once n'_2 is shown to be semantically equivalent to n'_1 we can pair them with s'_1 to create the new pair highlighted with the dashed relation $\langle n'_2, s'_1 \rangle$. Thus the negative instance set is $\langle n'_3, s'_1 \rangle, \langle n'_3, s'_2 \rangle, \langle n'_1, s'_3 \rangle, \langle n'_2, s'_3 \rangle$.

The above steps are formally described by the algorithm in Figure 5. It takes as input the generalized dataset as a list of correct pairs $I \subset \{ \langle n, s \rangle : n \in \mathcal{N}, s \in \mathcal{S} \}$ and returns a matrix M storing all positive and negative pairs $\mathcal{P} = \mathcal{N} \times \mathcal{S}$. M is obtained by (a) dividing I in k clusters of semantically related pairs and (b) applying the transitive closure to the semantic relationship between member pairs. More in detail, we first initialize its entries with a negative value, i.e. $M[n, s] = -1 \forall n \in \mathcal{N}, s \in \mathcal{S}$.

Second, we group together each $\langle n, s \rangle$ and $\langle n', s' \rangle \in I$, if at least two of their members are identical. If not we test if the two query members, s and s' retrieve the same result set. Since this may be time consuming we run this test only if

the selected columns in their `SELECT` clause are the same and if the two result sets share the same minimum.

Third, since the condition above is only necessary for semantic equivalence, in case we find the same result sets, we manually check if the natural language question members are semantically equivalent. This is faster and easier than checking the SQL queries.

Finally, once the initial clusters have been created, we apply the transitive closure to the cluster c_k to include all possible pairing between questions and queries belonging to c_k , i.e. $c_k = \{\langle n, s \rangle : n, s \in c_k\}$. We store in M the *id* of the clusters in the related pairs, i.e. $M[n][s] = k$ for each c_k . As a side effect all entries of M still containing -1 will be negative examples.

5 The Experiments

In these experiments, we study the effectiveness of our algorithm for automatically mapping questions into queries by testing the accuracy of selecting for each question of the test set its correct query. For this purpose, we learn a classifier of correct and incorrect pairs and we use it as a ranker for the possible queries as described in Section 2.

5.1 Setup

The query ranker⁴ consists in an SVM using advanced kernels for representing question and query pairs. We used SVM-Light-TK⁵, which extends the SVM-Light optimizer [14] with tree kernels. i.e. Syntactic Tree Kernel (STK) and its extension (STK_e) as described in Section 3.1. We model many different combination described in the next section. We used the default parameters, i.e. the cost and trade-off parameters = 1 (for normalized kernels) and $\lambda = 0.4$ (see Section 3.1).

As test set, we use our dataset obtained from GEOQUERIES250 by applying our algorithm described in Section 4. After the generalization process the initial 250 pairs of questions/queries are reduced to 155 pairs containing 154 different NL question and 80 different SQL queries. We found 76 clusters, from which we generated 165 positive and 12,001 negative examples for a total of 154×80 pairs. Since the number of negatives is much larger than the positives, we eliminate some of them to make the learning more efficient. For this purpose we only keep the pairs whose members share at least two stems. Our assumption is that a positive pair should share at least one variable and one concept (e.g. *VARstate cities*). We applied such heuristics to both training and test data by considering the false negatives in the final computation of the system accuracy⁶.

⁴ We simply used the score of the SVM-based classifiers. More effective approaches have been proposed [4,13].

⁵ <http://disi.unitn.it/~moschitt/Tree-Kernel.htm>

⁶ However, our results show that the above pre-processing in our data does not lead to any errors.

Table 1. Kernel combination accuracy

K1	K2	K1×K2	K1+K2
LIN	LIN	70.7±12.0	57.3±10.4
POLY	POLY	71.9±11.5	55.1±8.4
STK	STK	70.3±9.3	54.9±10.1
STK _e	STK _e	70.1±10.9	56.7±12.0
LIN	STK _e	75.6±13.1	56.6±12.4

Table 2. Kernel engineering results

Advanced Kernels	Accuracy
STK _e +POLY ²	73.2±11.4
(1+LIN ²) ²	73.6±9.4
(1+POLY ²) ²	73.2±10.9
(1+STK _e ²) ²	70.0±12.2
(1+LIN ²) ² +STK _e ²	75.0±10.8

We excluded 10.685 negative examples from our initial set, reducing it to 1.316 elements.

We evaluated our automatic mapping by applying a standard 10-fold cross validation and measuring the average accuracy and the Std. Dev. in selecting the correct query for each question of the test sets.

5.2 Results

We tested several models for ranking based on different kernel combinations whose results are reported in tables Table 1 and Table 2. The first two columns of Table 1 show the kernels used for the question and the query, respectively. More specifically, our basic kernels are: (1) linear kernel (LIN) built on the bag-of-words (BOW) of the questions or of the query, e.g. `SELECT` is considered a feature for the query; (2) a polynomial kernel of degree 3 on the above BOW (POLY); (3) the Syntactic Tree Kernel (STK) on the parse tree of the question or the query and (4) STK extended with leaf features (STK_e).

Columns 3 and 4 show the average accuracy (over 10 folds) ± Std. Dev. of two main kernel combinations by means of product and sum. Note that we can also sum or multiply different kernels, e.g. LIN×STK_e.

An examination of the reported tables suggests that: first, the basic traditional model based on linear kernel and BOW, i.e. LIN + LIN, provides an accuracy of only 57.3%, which is greatly improved by LIN×LIN=LIN², i.e. by 13.5 points. The explanation is that the kernel sum cannot express the relational feature pairs coming from questions and queries, thus LIN does not capture the underlying shared semantics between them. It should be noted that only kernel methods allow for an efficient and easy design of LIN², since the traditional approach would have required to build the Cartesian product of the question BOW by query BOW. This can be very large, e.g. 10K features for both spaces leads to a pair space of 100M features.

Second, the K₁+K₂ column confirms that the feature pair space is essential since the accuracy of all kernels implementing individual spaces (e.g. kernels which are sums of kernels) is much lower than the baseline model for feature pairs, i.e. LIN².

Third, if we include conjunctions in the BOW representation by using POLY, we improve the LIN model, when we use the feature pair space, i.e. 71.9% vs 70.7%.

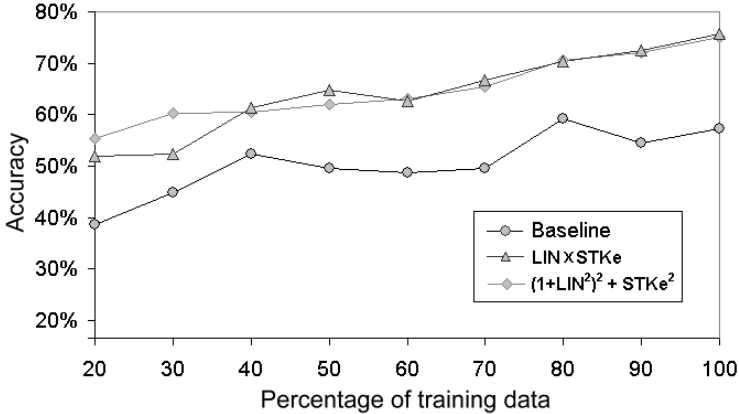


Fig. 6. Learning curves for GEOQUERIES250 corpora

Also, POLY² is better than STK² since it includes individual terms/words, which are not included by STK.

Next, the above point suggests that syntactic models can improve BOW although too many syntactic features (generated by STK) make the model unstable as suggested by the lower accuracy (70.1%) provided by $STK_e \times STK_e = STK_e^2$. This consideration leads us to experiment with the model $LIN \times STK_e$, which combines stems of the questions with syntactic constructs of SQL queries. This produces the highest result, i.e. 75.6%, suggesting that the syntactic parse tree of the SQL query is very reliable (it is obtained with 100% of accuracy) while the natural language parse tree, although accurate, introduces noise that degrades the overall feature representation. As a consequence it is more effective to use words only in the representation of the first member of the pairs.

Moreover, we experimented with very advanced kernels built on top of feature pair spaces as shown in Table 2. For example, we sum different pair spaces, STK_e^2 and POLY², and we apply the polynomial kernel on top of pair spaces by creating conjunctions, over feature pairs. This operation tends to increase too much the cardinality of the space and makes it ineffective. However, using the simplest initial space, i.e. LIN, to build pair conjunctions, i.e. $(1+LIN^2)^2$, we obtain a very interesting and high result, i.e. 73.6%⁷. This suggests that kernel methods have the potentiality to describe relational problems using simple building blocks although new theory describing the degradation of kernels when the space is too complex is required.

Finally, to study the stability of our complex kernels, we compared the learning curve of the baseline model, i.e. LIN+LIN, with the those of best models, i.e. $LIN \times STK_e$ and $STK^2 + (1+LIN^2)^2$. Figure 6 shows that complex kernels are

⁷ Although the Std. Dev. associated with the model accuracy is high, the one associated with the distribution of difference between the model accuracy is much lower. Considering that we used 10 folds, it is easy to verify that $(1+LIN^2)^2$ is better than LIN^2 at a 90% confidence limit.

not only more accurate but also more stable, i.e. their accuracy grows smoothly according to the increase of training data.

5.3 Related Work

As the literature suggest, NLIDBs can be classified according to the approach employed in deriving an SQL query that retrieves the answer of a given NL question against a database. In this section we review three systems based on different approaches and that were also tested on the GEOQUERIES250. For a complete review of many NLIDB refer to Chandra and Mihalcea [15].

Systems based on authoring interface rely on semantic grammar specified by an expert user to interpret question over the database. CatchPhrase [16] is an authoring tool where the author is asked to name database elements, tailor entries and define additional concepts. This tool achieves 80% Recall and 86% Precision.

Another approach is based on enriching the information contained in the pairs. An example is given by Precise system [12]. Reducing the problem of finding a semantic interpretation of ambiguous phrases to a graph matching problem, Precise achieves 100% Precision on a subset of semantically tractable questions (77,5% Recall).

The machine learning approach, that induces semantic grammar from a corpus of correct pairs of questions and queries, has been used in Krisp [17]. Krisp performs semantic parsing mapping sentences into their computer-executable meaning representations. For each production in the meaning representation language it trains an SVM classifier based on string subsequence kernels. Then it uses these classifiers to compositionally represent a natural language sentence in their meaning representations. Krisp achieves approximatively 94% Precision and 78% Recall. Our system, also based on the machine learning approach, does not decline to answer any questions and shows an accuracy of 76% when the SQL query of the pair with the highest rank is executed to retrieve the answer of the paired question.

Finally, another approach for mapping natural language sentences to artificial language using lambda calculus has been proposed in [18]. Such method showed 96.3% in Precision and 79.3% in Recall on the extended corpus GeoQueries880.

Regarding the use of tree kernels for natural language tasks several models have been proposed and experimented [4,19,20,21,22,23,24,25,26,27].

6 Conclusions

In this paper, we approach the problem of mapping natural language semantics into programming languages by automatically learning a model based on lexical and syntactic description of the training examples. In our study, these are pairs of NL questions and SQL queries, which we annotated by means of our semi-automatic algorithm based on the initial annotation available in the GEOQUERIES250 corpus.

To represent syntactic/semantic relationships expressed by the pairs above, we largely adopted kernel methods along with SVMs. We designed innovative combinations between different kernels for structured data applied to pairs of objects. To our knowledge, the functions that we propose for relational semantics description are novel. The experiments of the automatic question translation system show a satisfactory accuracy, i.e. 76%, although large improvements are still possible.

The main contributions of our study are: (i) we show that our automatic mapping between question and SQL queries is viable, (ii) in at least one task we have proved that kernel products are effective, (iii) syntax is important to map natural language into programming languages and (iv) we have generated a corpus for future studies, which we are going to make publically available.

In the future we would like to extend this research by focusing on advanced shallow semantic approaches such as predicate argument structures [28,29].

Acknowledgments

This research has been partially supported by the EC Project#231126, Living-Knowledge.

References

1. Basili, R., Moschitti, A., Pazienza, M.: A text classifier based on linguistic processing. In: Proceedings of IJCAI 1999, Machine Learning for Information Filtering (1999)
2. Charniak, E.: A maximum-entropy-inspired parser. In: Proceedings of NAACL 2000 (2000)
3. Shawe-Taylor, J., Cristianini, N.: Kernel Methods for Pattern Analysis. Cambridge University Press, Cambridge (2004)
4. Collins, M., Duffy, N.: New ranking algorithms for parsing and tagging: Kernels over discrete structures, and the voted perceptron. In: Proceedings of ACL 2002 (2002)
5. Zhang, D., Lee, W.S.: Question classification using support vector machines. In: Proceedings of the 26th annual international ACM SIGIR conference on Research and development in information retrieval, pp. 26–32. ACM Press, New York (2003)
6. Moschitti, A., Bejan, C.: A semantic kernel for predicate argument classification. In: CoNLL 2004, Boston, MA, USA (2004)
7. Moschitti, A., Coppola, B., Pighin, D., Basili, R.: Engineering of syntactic features for shallow semantic parsing. In: Proceedings of ACL 2005 Workshop on Feature Engineering for Machine Learning in NLP, USA (2005)
8. Moschitti, A., Pighin, D., Basili, R.: Tree kernels for semantic role labeling. Computational Linguistics 34(2), 193–224 (2008)
9. Moschitti, A., Quarteroni, S., Basili, R., Manandhar, S.: Exploiting syntactic and shallow semantic kernels for question/answer classification. In: Proceedings of ACL 2007, Prague, Czech Republic (2007)
10. Moschitti, A.: Kernel methods, syntax and semantics for relational text categorization. In: CIKM 2008: Proceeding of the 17th ACM conference on Information and knowledge management, pp. 253–262. ACM, New York (2008)

11. Moschitti, A., Zanzotto, F.: Fast and effective kernels for relational learning from texts. In: Ghahramani, Z. (ed.) *Proceedings of the 24th Annual International Conference on Machine Learning, ICML 2007* (2007)
12. Popescu, A.M., Etzioni, O., Kautz, H.: Towards a theory of natural language interfaces to databases. In: *Proceedings of the 2003 International Conference on Intelligent User Interfaces*, Miami, pp. 149–157. Association for Computational Linguistics (2003)
13. Moschitti, A., Pighin, D., Basili, R.: Semantic role labeling via tree kernel joint inference. In: *Proceedings of CoNLL-X*, New York City (2006)
14. Joachims, T.: Making large-scale SVM learning practical. In: Schölkopf, B., Burges, C., Smola, A. (eds.) *Advances in Kernel Methods* (1999)
15. Chandra, Y., Mihalcea, R.: *Natural language interfaces to databases*, University of North Texas, Thesis, M.S. (2006)
16. Minock, M., Olofsson, P., Näslund, A.: Towards building robust natural language interfaces to databases. In: Kapetanios, E., Sugumaran, V., Spiliopoulou, M. (eds.) *NLDB 2008*. LNCS, vol. 5039, pp. 187–198. Springer, Heidelberg (2008)
17. Kate, R.J., Mooney, R.J.: Using string-kernels for learning semantic parsers. In: *Proceedings of the 21st ICCL and 44th Annual Meeting of the ACL*, Sydney, Australia, July 2006, pp. 913–920. Association for Computational Linguistics (2006)
18. Zettlemoyer, L.S., Collins, M.: Learning to map sentences to logical form: Structured classification with probabilistic categorial grammars. In: *UAI*, pp. 658–666 (2005)
19. Kudo, T., Matsumoto, Y.: Fast Methods for Kernel-Based Text Analysis. In: Hinrichs, E., Roth, D. (eds.) *Proceedings of ACL*, pp. 24–31 (2003)
20. Cumby, C., Roth, D.: Kernel Methods for Relational Learning. In: *Proceedings of ICML 2003*, Washington, DC, USA, pp. 107–114 (2003)
21. Culotta, A., Sorensen, J.: Dependency Tree Kernels for Relation Extraction. In: *ACL 2004*, Barcelona, Spain, pp. 423–429 (2004)
22. Kudo, T., Suzuki, J., Isozaki, H.: Boosting-based parse reranking with subtree features. In: *Proceedings of ACL 2005*, US (2005)
23. Toutanova, K., Markova, P., Manning, C.: The Leaf Path Projection View of Parse Trees: Exploring String Kernels for HPSG Parse Selection. In: *Proceedings of EMNLP 2004*, Barcelona, Spain (2004)
24. Kazama, J., Torisawa, K.: Speeding up Training with Tree Kernels for Node Relation Labeling. In: *Proceedings of EMNLP 2005*, Toronto, Canada, pp. 137–144 (2005)
25. Shen, L., Sarkar, A., Joshi, A.k.: Using LTAG Based Features in Parse Reranking. In: *EMNLP*, Sapporo, Japan (2003)
26. Zhang, M., Zhang, J., Su, J.: Exploring Syntactic Features for Relation Extraction using a Convolution tree kernel. In: *Proceedings of NAACL*, New York City, USA, pp. 288–295 (2006)
27. Zhang, D., Lee, W.: Question classification using support vector machines. In: *Proceedings of SIGIR 2003*, Toronto, Canada. ACM, New York (2003)
28. Giuglea, A.M., Moschitti, A.: Knowledge discovery using framenet, verbnet and propbank. In: Meyers, A. (ed.) *Workshop on Ontology and Knowledge Discovering at ECML 2004*, Pisa, Italy (2004)
29. Giuglea, A.M., Moschitti, A.: Semantic Role Labeling via Framenet, Verbnet and Propbank. In: *Proceedings of ACL 2006*, Sydney, Australia (2006)