

# NEURAL ENQUIRER: Learning to Query Tables with Natural Language

Pengcheng Yin<sup>†\*</sup> Zhengdong Lu<sup>‡</sup> Hang Li<sup>‡</sup> Ben Kao<sup>†</sup>

<sup>†</sup>Dept. of Computer Science  
The University of Hong Kong  
{pcyin, kao}@cs.hku.hk

<sup>‡</sup>Noah’s Ark Lab, Huawei Technologies  
{Lu.Zhengdong, HangLi.HL}@huawei.com

## Abstract

We proposed NEURAL ENQUIRER as a neural network architecture to execute a natural language (NL) query on a knowledge-base (KB) for answers. Basically, NEURAL ENQUIRER finds the distributed representation of a query and then executes it on knowledge-base tables to obtain the answer as one of the values in the tables. Unlike similar efforts in end-to-end training of semantic parsers [11, 9], NEURAL ENQUIRER is fully “neuralized”: it not only gives distributional representation of the query and the knowledge-base, but also realizes the execution of compositional queries as a series of differentiable operations, with intermediate results (consisting of annotations of the tables at different levels) saved on multiple layers of memory. NEURAL ENQUIRER can be trained with gradient descent, with which not only the parameters of the controlling components and semantic parsing component, but also the embeddings of the tables and query words can be learned from scratch. The training can be done in an end-to-end fashion, but it can take stronger guidance, e.g., the step-by-step supervision for complicated queries, and benefit from it. NEURAL ENQUIRER is one step towards building neural network systems which seek to understand language by executing it on real-world. Our experiments show that NEURAL ENQUIRER can learn to execute fairly complicated NL queries on tables with rich structures.

## 1 Introduction

In models for natural language dialogue and question answering, there is ubiquitous need for querying a knowledge-base [13, 11]. The traditional pipeline is to put the query through a semantic parser to obtain some “executable” representations, typically logical forms, and then apply this representation to a knowledge-base for the answer. Both the semantic parsing and the query execution part can get quite messy for complicated queries like  $\tilde{Q}$ : “*Which city hosted the longest game before the game in Beijing?*” in Figure 1, and need carefully devised systems with hand-crafted features or rules to derive the correct logical form  $\tilde{F}$  (written in SQL-like style). Partially to overcome this difficulty, there has been effort [11] to “backpropagate” the result of query execution to revise the semantic representation of the query, which actually falls into the thread of work on learning from grounding [5]. One drawback of these

---

\*Work done when the first author worked as an intern at Noah’s Ark Lab, Huawei Technologies.

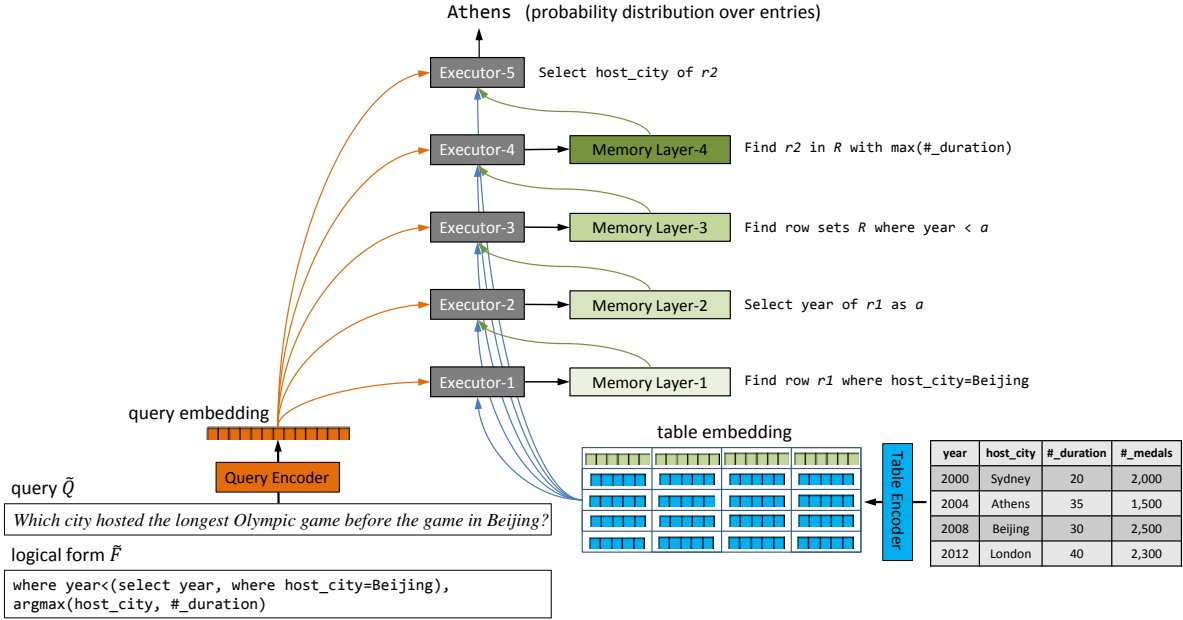


Figure 1: An overview of NEURAL ENQUIRER with five executors

semantic parsing models is rather symbolic with rule-based features, leaving only a handful of tunable parameters to cater to the supervision signal from the execution result.

On the other hand, neural network-based models are previously successful mostly on tasks with direct and strong supervision in natural language processing or related domain, with examples including machine translation and syntactic parsing. The recent work on learning to execute simple Python code with LSTM [15] pioneers in the direction on learning to parse structured objects through executing it in a purely neural way, while the later work on Neural Turing Machine (NTM) [6] introduces more modeling flexibility by equipping the LSTM with external memory and various means of interacting with it.

Our work, inspired by above-mentioned threads of research, aims to design a neural network system that can learn to understand the query and execute it on a knowledge-base table from examples of queries and answers. Our proposed NEURAL ENQUIRER encodes queries and KBs into distributed representations, and executes compositional queries against the KB through a series of differentiable operations. It can be trained using Query-Answer pairs, where the distributed representations of queries and the KB are optimized together with the query execution logic in an end-to-end fashion. We then demonstrates using a synthetic question-answering task that our proposed NEURAL ENQUIRER is capable of learning to execute compositional natural language queries with complex structures.

## 2 Overview of Neural Enquirer

Given an NL query  $Q$  and a KB table  $\mathcal{T}$ , NEURAL ENQUIRER executes the query against the table and outputs a ranked list of query answers. The execution is done by first using *Encoders* to encode the query and table into distributed representations, which are then sent to a cascaded pipeline of *Executors* to derive the answer. Figure 1 gives an illustrative example (with five executors) of various types of components involved:

**Query Encoder** (Section 3.1), which encodes the query into a distributed representation that carries the semantic information of the original query. The encoded query embedding will be sent to various executors to compute its execution result.

**Table Encoder** (Section 3.2), which encodes entries in the table into distributed vectors. Table Encoder outputs an embedding vector for each table entry, which retains the two-dimensional structure of the table.

**Executor** (Section 3.3), which executes the query against the table and outputs *annotations* that encode intermediate execution results, which are stored in the memory of each layer to be accessed by subsequent executor. Our basic assumption is that complex, compositional queries can be answered through multiple steps of computation, where each executor models a specific type of operation conditioned on the query. Figure 1 illustrates the operation each executor is supposed to perform in answering  $\hat{Q}$ . Different from classical semantic parsing approaches which require a predefined set of all possible logical operations, NEURAL ENQUIRER is capable of *learning* the logic of executors via end-to-end training using Query-Answer pairs. By stacking several executors, NEURAL ENQUIRER is able to answer complex queries involving multiple steps of computation.

### 3 Model

In this section we give a more detailed exposition of different types of components in the NEURAL ENQUIRER model.


#### 3.1 Query Encoder

Given an NL query  $Q$  composed of a sequence of words  $\{w_1, w_2, \dots, w_T\}$ , Query Encoder parses  $Q$  into a  $d_Q$ -dimensional vectorial representation  $\mathbf{q}$ :  $Q \xrightarrow{\text{encode}} \mathbf{q} \in \mathbb{R}^{d_Q}$ . In our implementation of NEURAL ENQUIRER, we employ a bidirectional RNN for this mission<sup>1</sup>. More specifically, the RNN summarizes the sequence of word embeddings of  $Q$ ,  $\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T\}$ , into a vector  $\mathbf{q}$  as the representation of  $Q$ , where  $\mathbf{x}_t = \mathbf{L}[w_t]$ ,  $\mathbf{x}_t \in \mathbb{R}^{d_w}$  and  $\mathbf{L}$  is the embedding matrix. See Appendix A for details.

It is worth noting that our Query Encoder can find the representation of rather general class of symbol sequences, agnostic to the actual representation of queries (e.g., natural language, SQL-like, etc). NEURAL ENQUIRER is capable of learning the execution logic expressed in the input query through end-to-end training, making it a generic model for query execution.

#### 3.2 Table Encoder

Table Encoder converts a knowledge-base table  $\mathcal{T}$  into its distributional representation as input to NEURAL ENQUIRER. Suppose the table has  $M$  rows and  $N$  columns, where each column comes with a *field name* (e.g., `host_city`), and the value of each table entry is a word (e.g., `Beijing`) in our vocabulary, Table Encoder first finds



<sup>1</sup>Other choices of sentence encoder such as LSTM or even convolutional neural networks are possible too

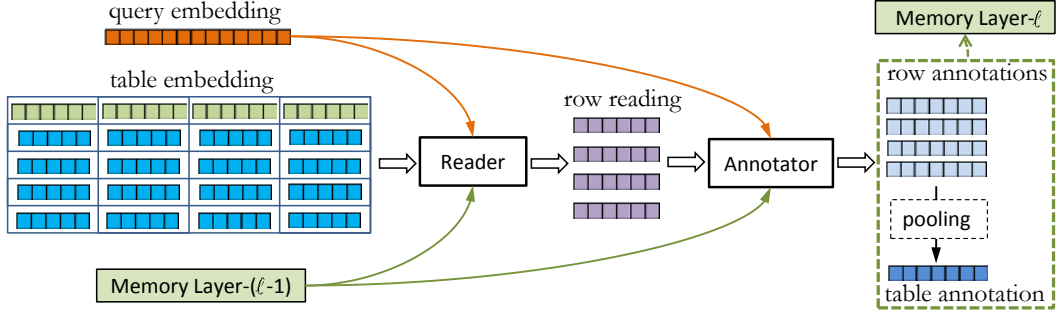


Figure 2: Overview of an Executor- $\ell$

the embedding for field names and values of table, and then it computes the (field, value) composite embedding for each of the  $M \times N$  entries in the table. More specifically, for the entry in the  $m$ -th row and  $n$ -th column with a value of  $w_{mn}$ , Table Encoder computes a  $d_{\mathcal{E}}$ -dimensional embedding vector  $\mathbf{e}_{mn}$  by fusing the embedding of the entry value with the embedding of its corresponding field name as follows:

$$\mathbf{e}_{mn} = \text{DNN}_0([\mathbf{L}[w_{mn}]; \mathbf{f}_n]) = \tanh(\mathbf{W} \cdot [\mathbf{L}[w_{mn}]; \mathbf{f}_n] + \mathbf{b})$$

where  $\mathbf{f}_n$  is the embedding of the field name (of the  $n$ -th column).  $\mathbf{W}$  and  $\mathbf{b}$  denote the weight matrices, and  $[\cdot; \cdot]$  the concatenation of vectors. The output of Table Encoder is a tensor of shape  $M \times N \times d_{\mathcal{E}}$ , consisting of  $M \times N$  embeddings of length  $d_{\mathcal{E}}$  for all entries.

Our Table Encoder functions differently from classical knowledge embedding models (e.g., TransE [4]), where embeddings of entities (entry values) and relations (field names) are learned in a unsupervised fashion via minimizing certain reconstruction errors. Embeddings in NEURAL ENQUIRER are optimized via supervised learning towards end-to-end QA tasks. Additionally, as will shown in the experiments, those embeddings function in a way as indices, which not necessarily encode the exact semantic meaning of their corresponding words.

### 3.3 Executor

NEURAL ENQUIRER executes an input query on a KB table through layers of execution. Each layer of executor captures a certain type of operation (e.g., `select`, `where`, `max`, etc.) relevant to the input query<sup>2</sup>, and returns intermediate execution results, referred to as *annotations*, saved in an external memory of the same layer. A query is executed step-by-step through a sequence of stacked executors. Such a cascaded architecture enables NEURAL ENQUIRER to answer complex, compositional queries. An illustrative example is given in Figure 1, with each executor annotated with the operation it is assumed to perform. We will demonstrate in Section 5 that NEURAL ENQUIRER is capable of learning the operation logic of each executor via end-to-end training.

As illustrated in Figure 2, an executor at Layer- $\ell$  (denoted as Executor- $\ell$ ) has two major neural network components: Reader and Annotator. An executor processes a table row-by-row. For the  $m$ -th row, with  $N$  (field, value) composite embeddings  $\mathcal{R}_m = \{\mathbf{e}_{m1}, \mathbf{e}_{m2}, \dots, \mathbf{e}_{mN}\}$ , the Reader fetches a read vector  $\mathbf{r}_m^\ell$  from  $\mathcal{R}_m$ , which is sent to the Annotator to compute a

<sup>2</sup>Depending on the query, an executor may perform different operations.

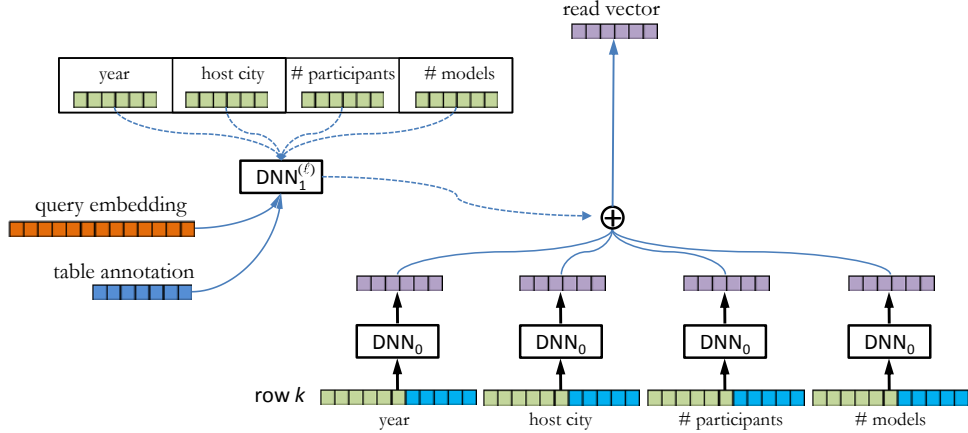


Figure 3: Illustration of the Reader for Executor- $\ell$ .

row annotation  $\mathbf{a}_m^\ell \in \mathbb{R}^{d_A}$ :

$$\text{Read Vector: } \mathbf{r}_m^\ell = f_R^\ell(\mathcal{R}_m, \mathcal{F}_T, \mathbf{q}, \mathcal{M}^{\ell-1}) \quad (1)$$

$$\text{Row Annotation: } \mathbf{a}_m^\ell = f_A^\ell(\mathbf{r}_m^\ell, \mathbf{q}, \mathcal{M}^{\ell-1}) \quad (2)$$

where  $\mathcal{M}^{\ell-1}$  denotes the content in memory Layer- $(\ell-1)$ , and  $\mathcal{F}_T = \{\mathbf{f}_1, \mathbf{f}_2, \dots, \mathbf{f}_N\}$  is the set of field name embeddings. Once all row annotations are obtained, Executor- $\ell$  then generates the *table annotation* through the following pooling process:

$$\text{Table Annotation: } \mathbf{g}^\ell = f_{\text{POOL}}(\mathbf{a}_1^\ell, \mathbf{a}_2^\ell, \dots, \mathbf{a}_M^\ell).$$

A row annotation captures the local execution result on each row, while a table annotation, derived from all row annotations, summarizes the global computational result on the whole table. Both row annotations  $\{\mathbf{a}_1^\ell, \mathbf{a}_2^\ell, \dots, \mathbf{a}_M^\ell\}$  and table annotation  $\mathbf{g}^\ell$  are saved in memory Layer- $\ell$ :  $\mathcal{M}^\ell = \{\mathbf{a}_1^\ell, \mathbf{a}_2^\ell, \dots, \mathbf{a}_M^\ell, \mathbf{g}^\ell\}$ .

Our design of executor is inspired by Neural Turing Machines [6], where data is fetched from an external memory using a read head, and subsequently processed by a controller, whose outputs are flushed back in to memories. An executor functions similarly by reading data from each row of the table, using a Reader, and then calling an Annotator to calculate intermediate computational results as annotations, which are stored in the executor’s memory. We assume that row annotations are able to handle operations which require only row-wise, local information (e.g., **select**, **where**), while table annotations can model superlative operations (e.g., **max**, **min**) by aggregating table-wise, global execution results. Therefore, a combination of row and table annotations enables NEURAL ENQUIRER to capture a variety of real-world query operations.

### 3.3.1 Reader

As illustrated in Figure 3, an executor at Layer- $\ell$  reads in a vector  $\mathbf{r}_m^\ell$  for each row  $m$ , defined as the weighted sum of composite embeddings for entries in this row:

$$\mathbf{r}_m^\ell = f_R^\ell(\mathcal{R}_m, \mathcal{F}_T, \mathbf{q}, \mathcal{M}^{\ell-1}) = \sum_{n=1}^N \tilde{\omega}(\mathbf{f}_n, \mathbf{q}, \mathbf{g}^{\ell-1}) \mathbf{e}_{mn}$$

where  $\tilde{\omega}(\cdot)$  is the normalized attention weights given by:

$$\tilde{\omega}(\mathbf{f}_n, \mathbf{q}, \mathbf{g}^{\ell-1}) = \frac{\exp(\omega(\mathbf{f}_n, \mathbf{q}, \mathbf{g}^{\ell-1}))}{\sum_{n'=1}^N \exp(\omega(\mathbf{f}_{n'}, \mathbf{q}, \mathbf{g}^{\ell-1}))} \quad (3)$$

and  $\omega(\cdot)$  is modeled as a DNN (denoted as  $\text{DNN}_1^{(\ell)}$ ).

Note that the  $\tilde{\omega}(\cdot)$  is agnostic to the values of entries in the row, i.e., in an executor all rows share the same set of weights  $\tilde{\omega}(\cdot)$ . Since each executor models a specific type of computation, it should only attend to a subset of entries pertain to its execution, which is modeled by the Reader. This is related to the content-based addressing of Neural Turing Machines [6] and the attention mechanism in neural machine translation models [2].

### 3.3.2 Annotator

In *Executor- $\ell$* , the Annotator computes row and table annotations based on the fetched read vector  $\mathbf{r}_m^\ell$  of the Reader, which are then stored in the  $\ell$ -th memory layer  $\mathcal{M}^\ell$  accessible to *Executor- $(\ell+1)$* . This process is repeated in intermediate layers, until the executor in the last layer to finally generate the answer.

**Row annotations** A row annotation encodes the local computational result on a specific row. As illustrated in Figure 4, a row annotation for row  $m$  in *Executor- $\ell$* , given by

$$\mathbf{a}_m^\ell = f_A^\ell(\mathbf{r}_m^\ell, \mathbf{q}, \mathcal{M}^{\ell-1}) = \text{DNN}_2^{(\ell)}([\mathbf{r}_m^\ell; \mathbf{q}; \mathbf{a}_m^{\ell-1}; \mathbf{g}^{\ell-1}]). \quad (4)$$

fuses the corresponding read vector  $\mathbf{r}_m^\ell$ , the results saved in previous memory layer (row and table annotations  $\mathbf{a}_m^{\ell-1}$ ,  $\mathbf{g}^{\ell-1}$ ), and the query embedding  $\mathbf{q}$ . Basically,

- row annotation  $\mathbf{a}_m^{\ell-1}$  represents the local status of the execution before Layer- $\ell$ ;
- table annotation  $\mathbf{g}^{\ell-1}$  summarizes the global status of the execution before Layer- $\ell$ ;
- read vector  $\mathbf{r}_m^\ell$  stores the value of current attention;
- query embedding  $\mathbf{q}$  encodes the overall execution agenda,

all of which are combined through  $\text{DNN}_2^{(\ell)}$  to form the annotation of row  $m$  in the current layer.




Figure 4: Illustration of Annotator for *Executor- $\ell$* .

**Table annotations** Capturing the global execution state, a table annotation is summarized from all row annotations via a global pooling operation. In our implementation of NEURAL ENQUIRER we employ max pooling:

$$\mathbf{g}^\ell = f_{\text{POOL}}(\mathbf{a}_1^\ell, \mathbf{a}_2^\ell, \dots, \mathbf{a}_M^\ell) = [g_1, g_2, \dots, g_{d_g}]^\top \quad (5)$$

where  $g_k = \max(\{\mathbf{a}_1^\ell(k), \mathbf{a}_2^\ell(k), \dots, \mathbf{a}_M^\ell(k)\})$  is the maximum value among the  $k$ -th elements of all row annotations. It is possible to use other pooling operations (e.g., gated pooling), but we find max pooling yields the best results.

### 3.3.3 Last Layer of Executor

Instead of computing annotations based on read vectors, the last executor in NEURAL ENQUIRER directly outputs the probability of the value of each entry in  $\mathcal{T}$  being the answer:

$$p(w_{mn}|Q, \mathcal{T}) = \frac{\exp(f_{\text{ANS}}^\ell(\mathbf{e}_{mn}, \mathbf{q}, \mathbf{a}_m^{\ell-1}, \mathbf{g}^{\ell-1}))}{\sum_{m'=1}^M \sum_{n'=1}^N \exp(f_{\text{ANS}}^\ell(\mathbf{e}_{m'n'}, \mathbf{q}, \mathbf{a}_{m'}^{\ell-1}, \mathbf{g}^{\ell-1}))} \quad (6)$$

where  $f_{\text{ANS}}^\ell(\cdot)$  is modeled as a DNN. Note that the last executor, which is devoted to returning answers, carries out a specific kind of execution using  $f_{\text{ANS}}^\ell(\cdot)$  based on the entry value, the query, and annotations from previous layer.

## 3.4 Handling Multiple Tables

Real-world KBs are often modeled by a schema involving various tables, where each table stores a specific type of factual information. We present NEURAL ENQUIRER-M, adapted for simultaneously operating on multiple KB tables. A key challenge in this scenario is that the multiplicity of tables requires modeling interaction between them. For example, NEURAL ENQUIRER-M needs to serve *join* queries, whose answer is derived by joining fields in different tables. Details of the modeling and experiments of NEURAL ENQUIRER-M are given in Appendix C.

## 4 Learning

NEURAL ENQUIRER can be trained in an *end-to-end* (N2N) fashion in Question Answering tasks. During training, both the representations of queries and table entries, as well as the execution logic captured by weights of executors are learned. More specifically, given a set of  $N_{\mathcal{D}}$  query-table-answer triples  $\mathcal{D} = \{(Q^{(k)}, \mathcal{T}^{(i)}, y^{(i)})\}$ , we optimize the model parameters by maximizing the log-likelihood of gold-standard answers:

$$\mathcal{L}_{\text{N2N}}(\mathcal{D}) = \sum_{i=1}^{N_{\mathcal{D}}} \log p(y^{(i)} = w_{mn}|Q^{(i)}, \mathcal{T}^{(i)}) \quad (7)$$

In end-to-end training, each executor discovers its operation logic from training data in a purely data-driven manner, which could be difficult for complicated queries requiring four or five sequential operations.

This can be alleviated by softly guiding the learning process via controlling the attention weights  $\tilde{w}(\cdot)$  in Eq. (3). By enforcing  $\tilde{w}(\cdot)$  to bias towards a field pertain to a specific operation,

we can “coerce” the executor to figure out the logic of this particular operation relative to the field. As an example, for Executor-1 in Figure 1, by biasing the weight of the `host_city` field towards 1.0, only the value of `host_city` field will be fetched and sent for computing annotations, in this way we can force the executor to learn to find the row whose `host_city` is *Beijing*. This setting will be referred to as *step-by-step* (SbS) training. Formally, this is done by introducing additional supervision signal to Eq. (7):

$$\mathcal{L}_{\text{Sbs}}(\mathcal{D}) = \sum_{i=1}^{N_{\mathcal{D}}} [\log p(y^{(i)} = w_{mn} | Q^{(i)}, \mathcal{T}^{(i)}) + \alpha \sum_{\ell=1}^L \log \tilde{w}(\mathbf{f}_{k,\ell}^*, \cdot, \cdot)] \quad (8)$$

where  $\alpha$  is a scalar and  $\mathbf{f}_{k,\ell}^*$  is the embedding of the field name known *a priori* to be relevant to the executor at Layer- $\ell$  in the  $k$ -th example.

## 5 Experiments

In this section we evaluate NEURAL ENQUIRER on synthetic QA tasks with queries with varying compositional depth. We will first briefly describe our synthetic QA task for benchmark and experimental setup, and then discuss the results under different settings.

### 5.1 Synthetic QA Task

We present a synthetic QA task to evaluate the performance of NEURAL ENQUIRER, where a large amount of QA examples at various levels of complexity are generated to evaluate the single table and multiple tables cases of the model. Starting with “artificial” tasks eases the process of developing novel deep models [14], and has gained increasing popularity in recent advances of the research on modeling symbolic computation using DNNs [6, 15].

Our synthetic dataset consists of query-table-answer triples  $\{(Q^{(i)}, \mathcal{T}^{(i)}, y^{(i)})\}$ . To generate such a triple, we first randomly sample a table  $\mathcal{T}^{(i)}$  of size  $10 \times 10$  from a synthetic schema of Olympic Games, which has 10 fields, whose values are drawn from a vocabulary of size 240, with 120 country and city names, and 120 numbers. Figure 5 gives an example table with one row. Next, we generate a query  $Q^{(i)}$  using predefined templates associated with its gold-standard answer  $y^{(i)}$  on  $\mathcal{T}^{(i)}$ . Our task consists of four types of natural language queries as summarized in Table 1, with annotated SQL-like logical forms for easy interpretation. We generate NL queries at various levels of compositionality to mimic the real world scenario. The complexity of those queries ranges from simple SELECT\_WHERE queries to more complicated NEST ones involving multiple steps of computation. Those queries are flexible enough to involve complex matching between NL phrases and logical constituents, which makes query understanding and execution nontrivial: (1) the same field is described by different NL phrases (e.g., “*How big is the country ...*” and “*What is the size of the country ...*” for `country_size` field); (2) different fields may be referred to by the same NL pattern (e.g., “*in China*” for `country` and “*in Beijing*” for `host_city`); (3) simple NL constituents may be grounded to complex logical operations (e.g., “*after the game in Beijing*” implies comparing between `year` fields). In our experiments we use the above procedure to generate benchmark datasets consisting of different types of queries. To make the artificial task harder, we enforce that all queries in the testing set do not appear in the training set.

To simulate the read-world scenario where queries of various types are issued to the model, we constructed two MIXTURED datasets for our main experiments, with 25K and 100K



year	host_city	#_participants	#_medals	#_duration	#_audience	host_country	GDP	country_size	population
2008	Beijing	4,200	2,500	30	67,000	China	2,300	960	130

Figure 5: An example table in the synthetic QA task (only show one row)

Query Type	Example Queries with Annotated SQL-like Logical Forms
SELECT_WHERE	<p>▷ <math>Q_1</math>: How many people participated in the game in Beijing?  <math>F_1</math>: <code>select #_participants, where host_city = Beijing</code></p> <p>▷ <math>Q_2</math>: In which country was the game hosted in 2012?  <math>F_2</math>: <code>select host_country, where year = 2012</code></p>
SUPERLATIVE	<p>▷ <math>Q_3</math>: When was the latest game hosted?  <math>F_3</math>: <code>argmax(host_city, year)</code></p> <p>▷ <math>Q_4</math>: How big is the country which hosted the shortest game?  <math>F_4</math>: <code>argmin(country_size, #_duration)</code></p>
WHERE_SUPERLATIVE	<p>▷ <math>Q_5</math>: How long is the game with the most medals that has fewer than 3,000 participants?  <math>F_5</math>: <code>where #_participants &lt; 3,000, argmax(#_duration, #_medals)</code></p> <p>▷ <math>Q_6</math>: How many medals are in the first game after 2008?  <math>F_6</math>: <code>where #_year &gt; 2008, argmin(#_medals, #_year)</code></p>
NEST	<p>▷ <math>Q_7</math>: Which country hosted the longest game before the game in Athens?  <math>F_7</math>: <code>where year &lt; (select year, where host_city=Athens), argmax(host_country, #_duration)</code></p> <p>▷ <math>Q_8</math>: How many people watched the earliest game that lasts for more days than the game in 1956?  <math>F_8</math>: <code>where #_duration &lt; (select #_duration, where year=1956), argmin(#_audience, #_year)</code></p>

Table 1: Example queries in our synthetic QA task

training examples respectively, where four types of queries are sampled with the ratio 1 : 1 : 1 : 2. Both datasets share the same testing set of 20K examples, 5K for each type of query.

## 5.2 Setup

**[Tuning]** We use a NEURAL ENQUIRER with five executors. The number of layers for  $DNN_1^{(\ell)}$  and  $DNN_2^{(\ell)}$  are set to 2 and 3, respectively. We set the dimensionality of word/entity embeddings and row/table annotations to 20, hidden layers to 50, and the hidden states of the GRU in query encoder to 150.  $\alpha$  in Eq. (8) is set to 0.2. We pad the beginning of all input queries to a fixed size.

NEURAL ENQUIRER is trained via standard back-propagation. Objective functions are optimized using SGD in a mini-batch of size 100 with adaptive learning rates (ADADelta [16]). The model converges fast within 100 epochs.

**[Baseline]** We compare our model with SEMPRES [11], a state-of-the-art semantic parser.

**[Metric]** We evaluate the performance of NEURAL ENQUIRER and SEMPRES (baseline) in terms of accuracy, defined as the fraction of correctly answered queries.

## 5.3 Main Results

Table 2 summarizes the results of SEMPRES (baseline) and our NEURAL ENQUIRER under *end-to-end* (N2N) and *step-by-step* (SbS) settings. We show both the individual performance for each type of query and the overall accuracy. We evaluate SEMPRES only on MIXTURED-25K because of its long training time even on the smaller MIXTURED-25K (> 3 days). We give discussion of efficiency issues in Appendix B.

We first discuss NEURAL ENQUIRER’s performance under end-to-end (N2N) training setting (the 3rd and 6th column in Table 2), and defer the discussion for SbS setting to Sec-

	MIXTURED-25K				MIXTURED-100K		
	SEMPRE	N2N	SbS	N2N - OOV	N2N	SbS	N2N - OOV
SELECT_WHERE	93.8%	96.2%	99.7%	90.3%	99.3%	100.0%	97.6%
SUPERLATIVE	97.8%	98.9%	99.5%	98.2%	99.9%	100.0%	99.7%
WHERE_SUPERLATIVE	34.8%	80.4%	94.3%	79.1%	98.5%	99.8%	98.0%
NEST	34.4%	60.5%	92.1%	57.7%	64.7%	99.7%	63.9%
Overall Acc.	65.2%	84.0%	96.4%	81.3%	90.6%	99.9%	89.8%

Table 2: Accuracies on MIXTURED datasets




Figure 6: Weights visualization of query  $Q_5$

tion 5.4. On MIXTURED-25K, our model outperforms SEMPRES on all types of queries, with a marginal gain on *simple* queries (SELECT\_WHERE, SUPERLATIVE), and significant improvement on *complex* ones (WHERE\_SUPERLATIVE, NEST). When the size of training set grows (MIXTURED-100K), NEURAL ENQUIRER achieves near 100% accuracy for the first three types of queries, while registering a decent overall accuracy of 90.6%. These results suggest that our model is very effective in answering compositional natural language queries, especially those with complex semantic structures compared with the state-of-the-art system.

To further understand why our model is capable of handling compositional queries, we study the attention weights  $\tilde{w}(\cdot)$  of Readers (Eq. 3) for executors in intermediate layers, and the answer probability (Eq. 6) the last executor outputs for each entry in the table. Those statistics are obtained from the model trained on MIXTURED-100K. We sampled two queries ( $Q_5$  and  $Q_7$  in Table 1) in the dataset that our model answers correctly and visualized their corresponding values, as illustrated in Figure 6 and 7, respectively. We find that each executor actually *learns* its execution logic from just the correct answers in end-to-end training, which corresponds with our assumption. For  $Q_5$ , the model executes the query in three steps, with each of the last three executors performs a specific type of operation. For each row, Executor-3 takes the value of the #\_participants field as input and computes intermediate annotations, while Executor-4 focuses on the #\_medals field. Finally, the last executor outputs high probability for the #\_duration field (the 5-th column) in the 3-rd row. The attention weights for Executor-1 and Executor-2 appear to be meaningless because  $Q_5$  requires only three steps of execution, and the model learns to defer the meaningful execution to the last three executors. We can guess confidently that in executing  $Q_5$ , Executor-3 performs the conditional filtering operation (**where** clause in  $F_5$ ), and Executor-4 performs the first part of **argmax** (finding the maximum value of #\_medals), while the last executor finishes the execution by assigning high probability for the #\_duration field of the row with the maximum value of #\_medals.

Compared with the relatively simple  $Q_5$ ,  $Q_7$  is more complicated, whose logical form  $F_7$  involves a nest sub-query, and requires five steps of execution. From the weights visualized in figure 7, we can find that the last three executors function similarly as the case in answering

$Q_7$ : Which country hosted the longest game before the game in Athens?



Figure 7: Weights visualization of query  $Q_7$

$Q_8$ : How many people watched the earliest game that lasts for more days than the game in 1956?




Figure 8: Weights visualization of query  $Q_8$  (an incorrectly answered query)

$Q_5$ , yet the execution logic for the first two executors is a bit obscure. We posit that this is because during end-to-end training, the supervision signal propagated from the top layer has decayed along the long path down to the first two executors, which causes vanishing gradient problem.

We also investigate the case where our model fails to deliver the correct answer for complicated queries. Figure 8 gives such a query  $Q_8$  in Table 1 together with its visualized weights. Similar as  $Q_7$ ,  $Q_8$  requires five steps of execution. Besides messing up the weights in the first two executors, the last executor, Executor-5, predicts a wrong entry as the query answer, instead of the highlighted (in red rectangle) correct entry.

## 5.4 With Additional Step-by-Step Supervision

To alleviate the vanishing gradient problem when training on complex queries as described in Section 5.3, in our next set of experiments we trained our NEURAL ENQUIRER model using step-by-step (SbS) training (Eq. 8), where we encourage each executor to attend to a specific field that is known *a priori* to be relevant to its execution logic. The results are shown in the 4-rd and 7-th columns of Table 2. With stronger supervision signal, the model significantly outperforms the results in end-to-end setting, and achieves near 100% accuracy on all types of queries, which shows that our proposed NEURAL ENQUIRER is capable of leveraging the additional supervision signal given to intermediate layers in SbS training setting, and answering complex and compositional queries with perfect accuracy.

Let us revisit the query  $Q_7$  in SbS setting with the weights visualization in Figure 9. In contrast to the result in N2N setting (Figure 7) where the attention weights for the first two executors are obscure, the weights in every executor are perfectly skewed towards the actual field pertain to each layer of execution (with a weight 1.0). Quite interestingly, the attention weights for Executor-3 and Executor-4 are exactly the same with the result in N2N setting, while the weights for Executor-1 and Executor-2 are significantly different, suggesting NEURAL ENQUIRER learned a different execution logic in the SbS setting.



Figure 9: Weights visualization of query  $Q_7$  in step-by-step training setting

## 5.5 Dealing with Out-Of-Vocabulary Words



Figure 10: Weights visualization of query  $Q_9$

One of the major challenges for applying neural network models to NLP applications is to deal with Out-Of-Vocabulary (OOV) words, which is particularly severe for QA. It is hard to cover existing tail entities, while at the same time new entities appear in user-issued queries and back-end KBs everyday. Quite interestingly, we find that a simple variation of NEURAL ENQUIRER is able to handle unseen entities almost without any loss of accuracy.

Basically, we divide words in the vocabulary into *entity* words and *operation* words. Embeddings of entity words (e.g., *Beijing*, *China*) function in a way as index to facilitate the matching between entities in queries and tables during the layer-by-layer execution, and do not need to be updated once initialized; while those of operation words, i.e., all non-entity words (e.g., numbers, *longest*, *before*, etc) carry semantic meanings relevant to execution and should be optimized during training. Therefore, after randomly initializing the embedding matrix  $\mathbf{L}$ , we only update the embeddings of operation words in training, while keeping those of entity words unchanged.

To test the model’s performance with OOV words, we modify queries in the testing portion of the MIXTURED dataset to replace all entity words (i.e., all country and city names) with OOV ones<sup>3</sup> unseen in the training set. Results obtained using N2N training are summarized in the 5th and 8th columns of Table 2. As it shows NEURAL ENQUIRER training in this OOV setting yields performance comparable to that in the non-OOV setting, indicating that operation words and entity words play different roles in query execution.

An interesting question to investigate in this OOV setting is how NEURAL ENQUIRER distinguishes between different types of entity words (i.e., cities and countries) in queries, since their embeddings are randomly initialized and fixed thereafter. An example query is  $Q_9$ : “How many people watched the game in Macau?”, where *Macau* is an OOV entity. To

<sup>3</sup>They also have embeddings in  $\mathbf{L}$ .

Query Type	SELECT_WHERE	SUPERLATIVE	WHERE_SUPERLATIVE	Overall
Accuracy	68.2%	84.8%	80.2%	77.7%

Table 3: Accuracies for large knowledge source simulation

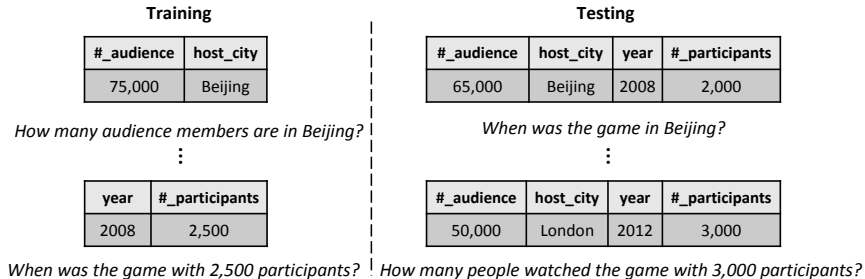


Figure 11: Large knowledge source simulation

help understand how the model knows *Macau* is a city, we give its weights visualization in Figure 10. Interestingly, the model first checks the `host_city` field in Executor-3, and then `host_country` in Executor-4, which suggests that the model learns to scan all possible fields where the OOV entity may belong to.

## 5.6 Simulating Large Knowledge Source

An important direction in semantic parsing research is to scale to large knowledge source [3, 11]. In this set of experiments we simulate a test case to evaluate NEURAL ENQUIRER’s ability to generalize to large knowledge source. We train a model on tables whose field sets are either  $\mathcal{F}_1, \mathcal{F}_2, \dots, \mathcal{F}_5$ , where  $\mathcal{F}_i$  (with  $|\mathcal{F}_i| = 5$ ) is a subset of the entire set  $\mathcal{F}_{\mathcal{T}}$ . We then test the model on tables with all fields  $\mathcal{F}_{\mathcal{T}}$  and queries whose fields span multiple subsets  $\mathcal{F}_i$ . Figure 11 illustrates the setting. Note that all testing queries exhibit field combinations unseen in the training data, to mimic the difficulty the system often encounter when scaling to large knowledge source, which usually poses great challenge on model’s generalization ability. We then train and test the model only on a new dataset of the first three types of relatively simple queries (namely SELECT\_WHERE, SUPERLATIVE and WHERE\_SUPERLATIVE). The sizes of training/testing splits are 75,000 and 30,000, with equal numbers for different query types. Table 3 lists the results. NEURAL ENQUIRER still maintains a reasonable performance even when the compositionality of testing queries is previously unseen, showing the model’s generalization ability in tackling unseen query patterns through composition of familiar ones, and hence the potential to scale to larger and unseen knowledge sources.

## 6 Related Work

Our work falls into the research area of Semantic Parsing, where the key problem is to parse Natural Language queries into logical forms executable on KBs. Classical approaches for Semantic Parsing can be broadly divided into two categories. The first line of research resorts to the power of grammatical formalisms (e.g., Combinatory Categorical Grammar) to parse NL queries and generate corresponding logical forms, which requires curated/learned lexicons defining the correspondence between NL phrases and symbolic constituents [17, 7, 1, 18]. The

model is tuned with annotated logical forms, and is capable of recovering complex semantics from data, but often constrained on a specific domain due to scalability issues brought by the crisp grammars and the lack of annotated training data. Another line of research takes a semi-supervised learning approach, and adopts the results of query execution (i.e., answers) as supervision signal [5, 3, 10, 11, 12]. The parsers, designed towards this new learning paradigm, take different types of forms, ranging from generic chart parsers [3, 11] to more specifically engineered, task-oriented ones [12, 8]. Semantic parsers in this category often scale to open domain knowledge sources, but lack the ability of understanding compositional queries because of the intractable search space incurred by the flexibility of parsing algorithms. Our work follows this line of research in using query answers as indirect supervision to facilitate end-to-end training using QA tasks, but performs semantic parsing in distributional spaces, where logical forms are “neuralized” to an executable distributed representation.

Our work is also related to the recent advances of modeling symbolic computation using Deep Neural Networks. Pioneered by the development of Neural Turing Machines (NTMs) [6], this line of research studies the problem of using differentiable neural networks to perform “hard” symbolic execution. As an independent line of research with similar flavor, Zaremba et al. [15] designed a LSTM-RNN to execute simple Python programs, where the parameters are learned by comparing the neural network output and the correct answer. Our work is related to both lines of work, in that like NTM, we heavily use external memory and flexible way of processing (e.g., the attention-based reading in the operations in `Reader`) and like [15], `NEURAL ENQUIRER` learns to execute a sequence with complicated structure, and the model is tuned from the executing them. As a highlight and difference from the previous work, we have a deep architecture with multiple layer of external memory, with the neural network operations highly customized to querying KB tables.

Perhaps the most related work to date is the recently published `NEURAL PROGRAMMER` proposed by Neelakantan et al. [9], which studies the same task of executing queries on tables with Deep Neural Networks. `NEURAL PROGRAMMER` uses a neural network model to select operations during query processing. While the query planning (i.e., which operation to execute at each time step) phase is modeled softly using neural networks, the symbolic operations are predefined by users. In contrast `NEURAL ENQUIRER` is fully distributional: it models both the query planning and the operations with neural networks, which are jointly optimized via end-to-end training. Our `NEURAL ENQUIRER` model learns symbolic operations using data-driven approach, and demonstrates that a fully neural, end-to-end differentiable system is capable of modeling and executing compositional arithmetic and logic operations upto certain level of complexity.

## 7 Conclusion and Future Work

In this paper we propose `NEURAL ENQUIRER`, a fully neural, end-to-end differentiable network that learns to execute queries on tables. We present results on a set of synthetic QA tasks to demonstrate the ability of `NEURAL ENQUIRER` to answer fairly complicated compositional queries across multiple tables. In the future we plan to advance this work in the following directions. First we will apply `NEURAL ENQUIRER` to natural language questions and natural language answers, where both the input query and the output supervision are noisier and less informative. Second, we are going to scale to real world QA task as in [11], for which we have to deal with a large vocabulary and novel predicates. Third, we are going to work

on the computational efficiency issue in query execution by heavily borrowing the symbolic operation.

## References

- [1] Y. Artzi, K. Lee, and L. Zettlemoyer. Broad-coverage ccg semantic parsing with amr. In *EMNLP*, pages 1699–1710, 2015.
- [2] D. Bahdanau, K. Cho, and Y. Bengio. Neural machine translation by jointly learning to align and translate. *ICLR*, 2015.
- [3] J. Berant, A. Chou, R. Frostig, and P. Liang. Semantic parsing on freebase from question-answer pairs. In *EMNLP*, pages 1533–1544, 2013.
- [4] A. Bordes, N. Usunier, A. Garca-Durn, J. Weston, and O. Yakhnenko. Translating embeddings for modeling multi-relational data. In *NIPS*, pages 2787–2795, 2013.
- [5] D. L. Chen and R. J. Mooney. Learning to sportscast: a test of grounded language acquisition. In *ICML*, pages 128–135, 2008.
- [6] A. Graves, G. Wayne, and I. Danihelka. Neural turing machines. *CoRR*, abs/1410.5401, 2014.
- [7] T. Kwiatkowski, E. Choi, Y. Artzi, and L. S. Zettlemoyer. Scaling semantic parsers with on-the-fly ontology matching. In *EMNLP*, pages 1545–1556, 2013.
- [8] D. K. Misra, K. Tao, P. Liang, and A. Saxena. Environment-driven lexicon induction for high-level instructions. In *ACL (1)*, pages 992–1002, 2015.
- [9] A. Neelakantan, Q. V. Le, and I. Sutskever. Neural Programmer: Inducing Latent Programs with Gradient Descent. *ArXiv e-prints*, Nov. 2015.
- [10] P. Pasupat and P. Liang. Zero-shot entity extraction from web pages. In *ACL (1)*, pages 391–401, 2014.
- [11] P. Pasupat and P. Liang. Compositional semantic parsing on semi-structured tables. In *ACL (1)*, pages 1470–1480, 2015.
- [12] W. tau Yih, M.-W. Chang, X. He, and J. Gao. Semantic parsing via staged query graph generation: Question answering with knowledge base. In *ACL (1)*, pages 1321–1331, 2015.
- [13] T.-H. Wen, M. Gasic, N. Mrksic, P. hao Su, D. Vandyke, and S. J. Young. Semantically conditioned lstm-based natural language generation for spoken dialogue systems. In *EMNLP*, pages 1711–1721, 2015.
- [14] J. Weston, A. Bordes, S. Chopra, and T. Mikolov. Towards ai-complete question answering: A set of prerequisite toy tasks. *CoRR*, abs/1502.05698, 2015.
- [15] W. Zaremba and I. Sutskever. Learning to execute. *CoRR*, abs/1410.4615, 2014.
- [16] M. D. Zeiler. ADADELTA: an adaptive learning rate method. *CoRR*, abs/1212.5701, 2012.
- [17] L. S. Zettlemoyer and M. Collins. Learning to map sentences to logical form: Structured classification with probabilistic categorial grammars. In *UAI*, pages 658–666, 2005.
- [18] L. S. Zettlemoyer and M. Collins. Online learning of relaxed ccg grammars for parsing to logical form. In *EMNLP-CoNLL*, pages 678–687, 2007.

## A Computation of Query Encoder

We use a bidirectional RNN as the Query Encoder, which consists of a *forward* GRU and a *backward* GRU. Given the sequence of word embeddings of  $Q$ :  $\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T\}$ , at each time step  $t$ , the forward GRU computes the hidden state  $\mathbf{h}_t$  as follows:

$$\begin{aligned} \mathbf{h}_t &= \mathbf{z}_t \mathbf{h}_{t-1} + (\mathbf{1} - \mathbf{z}_t) \tilde{\mathbf{h}}_t \\ \tilde{\mathbf{h}}_t &= \tanh(\mathbf{W} \mathbf{x}_t + \mathbf{U}(\mathbf{r}_t \circ \mathbf{h}_{t-1})) \\ \mathbf{z}_t &= \sigma(\mathbf{W}_z \mathbf{x}_t + \mathbf{U}_z \mathbf{h}_{t-1}) \\ \mathbf{r}_t &= \sigma(\mathbf{W}_r \mathbf{x}_t + \mathbf{U}_r \mathbf{h}_{t-1}) \end{aligned}$$

where  $\mathbf{W}$ ,  $\mathbf{W}_z$ ,  $\mathbf{W}_r$ ,  $\mathbf{U}$ ,  $\mathbf{U}_z$ ,  $\mathbf{U}_r$  are parametric matrices,  $\mathbf{1}$  the column vector of all ones, and  $\circ$  element-wise multiplication. The backward GRU reads the sequence in reverse order. We concatenate the last hidden states given by the two GRUs as the vectorial representation  $\mathbf{q}$  of the query.

## B Efficiency of Model Learning

We compared the efficiency of NEURAL ENQUIRER and SEMPRES (baseline) in training by plotting the accuracy on testing data by training time. Figure 12 illustrates the results. We train NEURAL ENQUIRER-CPU and SEMPRES on a machine with Intel Core i7-3770@3.40GHz and 16GB memory, while NEURAL ENQUIRER-GPU is tuned on Nvidia Tesla K40. NEURAL ENQUIRER-CPU is 10 times faster than SEMPRES, and NEURAL ENQUIRER-GPU is 100 times faster.




Figure 12: Accuracy on Testing by Training Time

## C Handling Multiple Tables

### C.1 NEURAL ENQUIRER-M Model

Basically, NEURAL ENQUIRER-M assigns an executor to each table  $\mathcal{T}_k$  in every execution layer  $\ell$ , denoted as  $\text{Executor}(\ell, k)$ . Figure 13 pictorially illustrates  $\text{Executor}(\ell, 1)$  and  $\text{Executor}(\ell, K)$  working on Table-1 and Table-K respectively. Within each executor, the Reader is designed the same way as single table case, while we modify the Annotator to let in the information from other tables. More specifically, for  $\text{Executor}(\ell, k)$ , we modify its Annotator by extending Eq. (4) to leverage computational



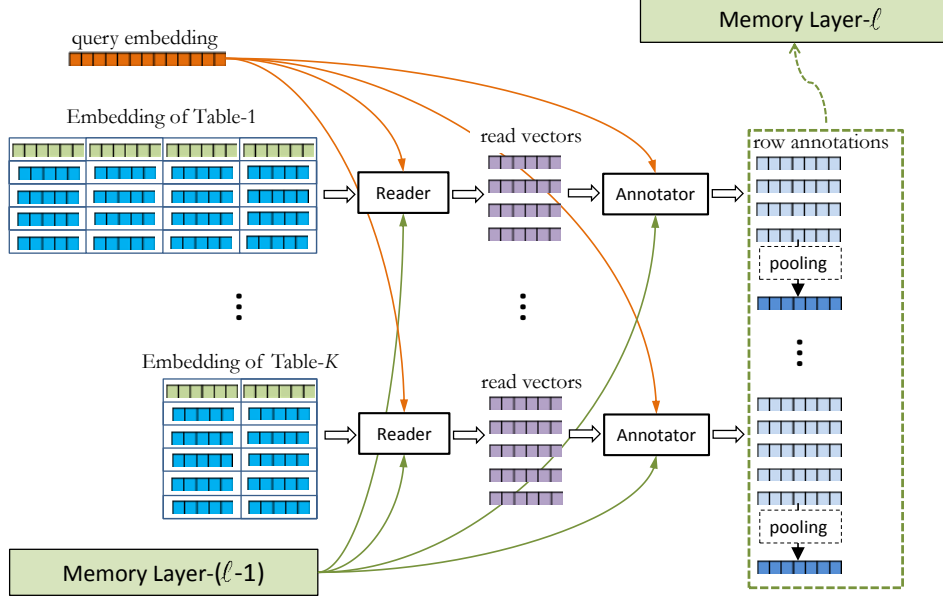


Figure 13: Executor- $(\ell, 1)$  and Executor- $(\ell, K)$  in multiple tables case

results from other tables when computing the annotation for the  $m$ -th row:

$$\begin{aligned} \mathbf{a}_{k,m}^\ell &= f_A^\ell(\mathbf{r}_{k,m}^\ell, \mathbf{q}, \mathbf{a}_{k,m}^{\ell-1}, \mathbf{g}_k^{\ell-1}, \hat{\mathbf{a}}_{k,m}^{\ell-1}, \hat{\mathbf{g}}_k^{\ell-1}) \\ &= \text{DNN}_2^{(\ell)}([\mathbf{r}_{k,m}^\ell; \mathbf{q}; \mathbf{a}_{k,m}^{\ell-1}; \mathbf{g}_k^{\ell-1}; \hat{\mathbf{a}}_{k,m}^{\ell-1}; \hat{\mathbf{g}}_k^{\ell-1}]) \end{aligned}$$

This process is illustrated in Figure 14. Note that we add subscripts  $k \in [1, K]$  to the notations to index tables. To model the interaction between tables, the Annotator incorporates the “relevant” row annotation,  $\hat{\mathbf{a}}_{k,m}^{\ell-1}$  and the “relevant” table annotation,  $\hat{\mathbf{g}}_k^{\ell-1}$  derived from the previous execution results of other tables when computing the current row annotation.

A relevant row annotation stores the data fetched from row annotations of other tables, while a relevant table annotation summarizes the table-wise execution results from other tables. We now describe how to compute those annotations. First, for each table  $\mathcal{T}_{k'}$ ,  $k' \neq k$ , we fetch a relevant row annotation  $\bar{\mathbf{a}}_{k,k',m}^{\ell-1}$  from all row annotations  $\{\mathbf{a}_{k',m'}^{\ell-1}\}$  of  $\mathcal{T}_{k'}$  via attentive reading:

$$\bar{\mathbf{a}}_{k,k',m}^{\ell-1} = \sum_{m'=1}^{M_{k'}} \frac{\exp(\gamma(\mathbf{r}_{k,m}^\ell, \mathbf{q}, \mathbf{a}_{k,m}^{\ell-1}, \mathbf{a}_{k',m'}^{\ell-1}, \mathbf{g}_k^{\ell-1}, \mathbf{g}_{k'}^{\ell-1}))}{\sum_{m''=1}^{M_{k'}} \exp(\gamma(\mathbf{r}_{k,m}^\ell, \mathbf{q}, \mathbf{a}_{k,m}^{\ell-1}, \mathbf{a}_{k',m''}^{\ell-1}, \mathbf{g}_k^{\ell-1}, \mathbf{g}_{k'}^{\ell-1}))} \mathbf{a}_{k',m'}^{\ell-1}$$

Intuitively, the attention weight  $\gamma(\cdot)$  (modeled by a DNN) captures how important the  $m'$ -th row annotation from table  $\mathcal{T}_{k'}$ ,  $\mathbf{a}_{k',m'}^{\ell-1}$  is with respect to the current step of execution. After getting the set of row annotations fetched from all other tables,  $\{\bar{\mathbf{a}}_{k,k',m}^{\ell-1}\}_{k'=1, k' \neq k}^K$ , we then compute  $\hat{\mathbf{a}}_{k,m}^{\ell-1}$  and  $\hat{\mathbf{g}}_k^{\ell-1}$  via a pooling operation<sup>4</sup> on  $\{\bar{\mathbf{a}}_{k,k',m}^{\ell-1}\}_{k'=1, k' \neq k}^K$  and  $\{\mathbf{g}_{k'}^{\ell-1}\}_{k'=1, k' \neq k}^K$ :

$$\langle \hat{\mathbf{a}}_{k,m}^{\ell-1}, \hat{\mathbf{g}}_k^{\ell-1} \rangle = \hat{f}_{\text{POOL}}(\{\bar{\mathbf{a}}_{k,1,m}^{\ell-1}, \bar{\mathbf{a}}_{k,2,m}^{\ell-1}, \dots, \bar{\mathbf{a}}_{k,K,m}^{\ell-1}\}; \{\mathbf{g}_1^{\ell-1}, \mathbf{g}_2^{\ell-1}, \dots, \mathbf{g}_K^{\ell-1}\}).$$

In summary, relevant row and table annotations encode the local and global computational results from other tables. By incorporating them into calculating row annotations, NEURAL ENQUIRER-M is capable of answering queries that involve interaction between multiple tables.

Finally, NEURAL ENQUIRER-M outputs the ranked list of answer probabilities by normalizing the value of  $g(\cdot)$  in Eq. (6) over each entry for very table.

<sup>4</sup>This operation is trivial in our experiments on two tables.

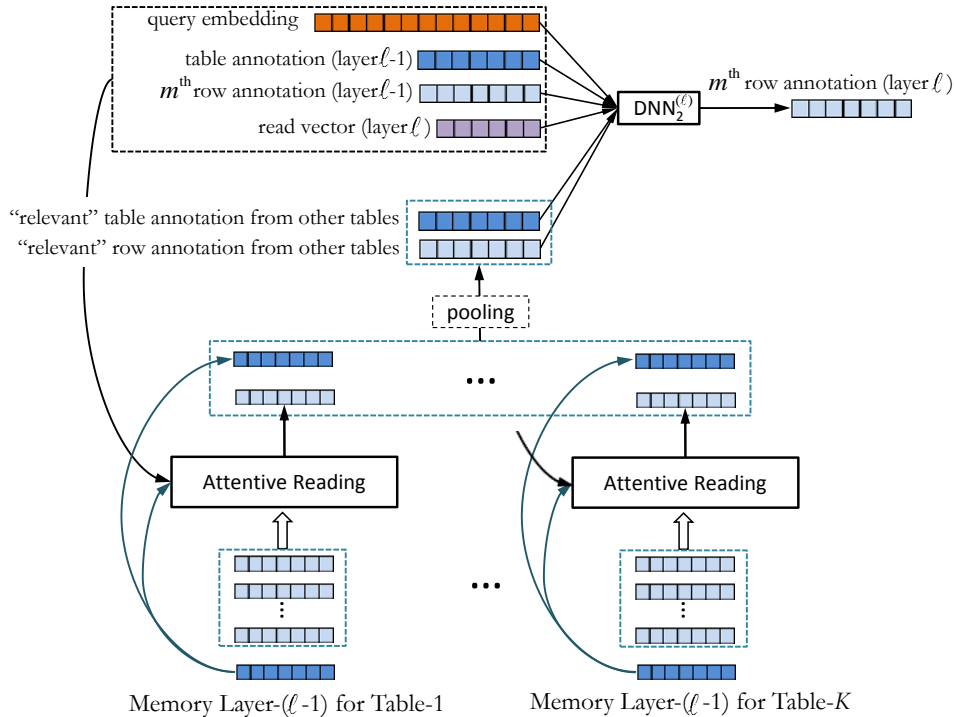


Figure 14: Illustration of Annotator for multiple tables case

country	year	host_city	#_participants	#_medals	#_duration	#_audience
China	2008	Beijing	3,500	4,200	30	67,000

country	continent	population	country_size
China	Asia	130	960

Figure 15: Multiple tables example in the synthetic QA task (only show one row)

## C.2 Experimental Results

We present preliminarily results for NEURAL ENQUIRER-M, which we evaluated on SQL-like SELECT\_WHERE logical forms (like  $F_1$ ,  $F_2$  in Table 1). We sampled a dataset of 100K examples, with each example having two tables as in Figure 15. Out of all SELECT\_WHERE queries, roughly half of the queries (denoted as “Join”) require joining the two tables to derive answers. We tested on a model with three executors. Table 4 lists the results. The accuracy of join queries is lower than that of non-join queries, which is caused by additional interaction between the two tables involved in answering join queries.

Query Type	Non-Join	Join	Overall
<b>Accuracy</b>	99.7%	81.5%	91.3%

Table 4: Accuracies of SELECT\_WHERE queries on two tables

We find that NEURAL ENQUIRER-M is capable of identifying that the *country* field is the foreign key linking the two tables. Figure 16 illustrates the attention weights for a correctly answered join




Figure 16: Weights visualization of query  $Q_9$

query  $Q_9$ . Although the query does not contain any hints for the foreign key (*country* field), **Executor-(1,1)** (the executor at Layer-1 on **Table-1**) operates on an ensemble of embeddings of the *country* and *year* fields, whose outputting row annotations (contain information of both the key *country* and the value *year*) are sent to **Executor-(2,2)** to compare with the *country* field in **Table-2**. We posit that the result of comparison is stored in the row annotations of **Executor-(2,2)** and subsequently sent to the executors at Layer-3 for computing the answer probability for each entry in the two tables.