
Coupling Distributed and Symbolic Execution for Natural Language Queries

Lili Mou¹ Zhengdong Lu² Hang Li³ Zhi Jin¹

Abstract

Building neural networks to query a knowledge base (a table) with natural language is an emerging research topic in deep learning. An executor for table querying typically requires multiple steps of execution because queries may have complicated structures. In previous studies, researchers have developed either fully distributed executors or symbolic executors for table querying. A distributed executor can be trained in an end-to-end fashion, but is weak in terms of execution efficiency and explicit interpretability. A symbolic executor is efficient in execution, but is very difficult to train especially at initial stages. In this paper, we propose to couple distributed and symbolic execution for natural language queries, where the symbolic executor is pretrained with the distributed executor’s intermediate execution results in a step-by-step fashion. Experiments show that our approach significantly outperforms both distributed and symbolic executors, exhibiting high accuracy, high learning efficiency, high execution efficiency, and high interpretability.

1. Introduction

Using natural language to query a knowledge base is an important task in NLP and has wide applications in question answering (QA) (Yin et al., 2016a), human-computer conversation (Wen et al., 2017), etc. Table 1 illustrates an example of a knowledge base (a table) and a query “How long is the game with the largest host country size?” To answer the question, we should first find a row with the

¹Key Laboratory of High Confidence Software Technologies (Peking University), MoE; Software Institute, Peking University, China ²DeeplyCurious.ai ³Noah’s Ark Lab, Huawei Technologies. Work done when the first author was an intern at Huawei. Correspondence to: L.M. <double.power.mou@gmail.com>, Z.L. <luz@DeeplyCurious.ai>, H.L. <HangLi.HL@huawei.com>, Z.J. <zhijin@sei.pku.edu.cn>.

Query:

How long is the game with the largest host country size?

Knowledge base (table):

Year	City	...	Area	...	Duration
...
2000	Sydney	...	200	...	30
2004	Athens	...	250	...	20
2008	Beijing	...	350	...	25
2012	London	...	300	...	35
2016	Rio de Janeiro	...	200	...	40
...

Table 1. An example of a natural language query and a knowledge base (table).

largest value in the column *Area*, and then select the value of the chosen row with the column being *Duration*.

A typical approach to table querying is to convert a natural language sentence to an “executable” logic form, known as *semantic parsing*. Traditionally, building a semantic parser requires extensive human engineering of explicit features (Berant et al., 2013; Pasupat & Liang, 2015).

With the fast development of deep learning, an increasingly number of studies use neural networks for semantic parsing. Dong & Lapata (2016) and Xiao et al. (2016) apply sequence-to-sequence (*seq2seq*) neural models to generate a logic form conditioned on an input sentence, but the training requires groundtruth logic forms, which are costly to obtain and specific to a certain dataset. In realistic settings, we only assume groundtruth denotations¹ are available, and that we do not know execution sequences or intermediate execution results. Liang et al. (2017) train a *seq2seq* network by REINFORCE policy gradient. But it is known that the REINFORCE algorithm is sensitive to the initial policy; also, it could be very difficult to get started at early stages.

Yin et al. (2016b) propose a fully distributed neural enquirer, comprising several neuralized execution layers of field attention, row annotation, etc. The model can be trained in an end-to-end fashion because all components are differentiable. However, it lacks explicit interpretation and is not efficient in execution due to intensive matrix/vector operation during neural processing.

Neelakantan et al. (2016) propose a neural programmer

¹A *denotation* refers to an execution result.

by defining a set of symbolic operators (e.g., `argmax`, `greater_than`); at each step, all possible execution results are fused by a softmax layer, which predicts the probability of each operator at the current step. The step-by-step fusion is accomplished by weighted average and the model is trained with mean square error. Hence, such approaches work with numeric tables, but may not be suited for other operations like string matching. It also suffers from the problem of “exponential numbers of combinatorial states,” as the model explores the entire space at a time by step-by-step weighted average.

In this paper, we propose to couple distributed and symbolic execution for natural language queries. By “symbolic execution,” we mean that we define symbolic operators and keep discrete intermediate execution results.² Our intuition rises from the observation that a fully distributed/neuralized executor also exhibits some (imperfect) symbolic interpretation. For example, the field attention gadget in Yin et al. (2016b) generally aligns with column selection. We therefore use the distributed model’s intermediate execution results as supervision signals to pre-train a symbolic executor. Guided by such imperfect step-by-step supervision, the symbolic executor learns a fairly meaningful initial policy, which largely alleviates the cold start problem of REINFORCE. Moreover, the improved policy can be fed back to the distributed executor to improve the neural network’s performance.

We evaluated the proposed approach on the QA dataset in Yin et al. (2016b). Our experimental results show that the REINFORCE algorithm alone takes long to get started. Even if it does, it is stuck in poor local optima. Once pretrained by imperfect supervision signals, the symbolic executor can recover most execution sequences, and also achieves the highest denotation accuracy. It should be emphasized that, in our experiment, neither the distributed executor nor the symbolic executor is aware of groundtruth execution sequences, and that the entire model is trained with weak supervision of denotations only.

To the best of our knowledge, we are the first to couple distributed and symbolic execution for semantic parsing. Our study also sheds light on neural sequence prediction in general.

2. Approach

In Subsection 2.1, we introduce the fully distributed neural executor, which is mostly based on Yin et al. (2016b). For symbolic execution, we design a set of operators that

²The sense of *symbolic execution* here should not be confused with “symbolic execution of a program” (see https://en.wikipedia.org/wiki/Symbolic_execution for example).

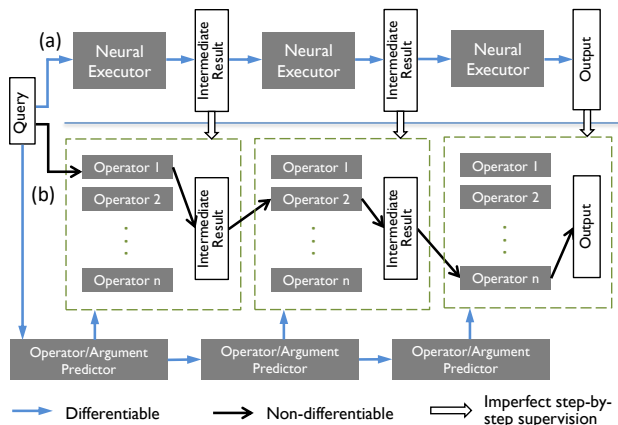


Figure 1. An overview of the coupled distributed and symbolic executors.

are complete to the task at hand; at each execution step, a neural network predicts a particular operator and possibly arguments (Subsection 2.2).

Subsection 2.3 provides a unified view of distributed and symbolic execution (Figure 1). We explain how the symbolic executor is pretrained by the distributed one’s intermediate execution results, and then further trained with the REINFORCE algorithm.

2.1. Distributed Executor

The distributed executor makes full use of neural networks for table querying. By “distributed,” we mean that all semantic units (including words in the query, entries in the table, and execution results) are represented as distributed, real-valued vectors and processed by neural networks. One of the most notable studies of distributed semantics is word embeddings, which map discrete words to vectors as meaning representations (Mikolov et al., 2013).

The distributed executor consists of the following main components.

- **Query encoder.** Words are mapped to word embeddings and a bi-directional recurrent neural network (RNN) aggregates information over the sentence. RNNs’ last states in both directions are concatenated as the query representation (denoted as q).
- **Table encoder.** All table cells are also represented as embeddings. For a cell c (e.g., *Beijing* in Table 1) with its column/field name being f (e.g., *City*), the cell vector is the concatenation of the embeddings of c and f , further processed by a feed-forward neural network (also known as multi-layer perceptron). We denote the representation of a cell as c .
- **Executor.** As shown in Figure 1a, the neural network comprises several steps of execution. In each execution step, the neural network selects a column by soft-

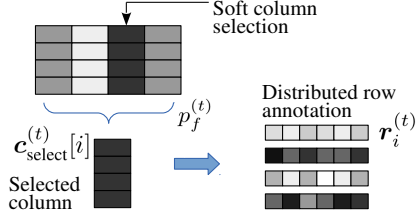


Figure 2. A single step of distributed execution.

max attention, and annotates each row with a vector, i.e., an embedding (Figure 2). The row vector can be intuitively thought of as row selection in query execution, but is represented by distributed semantics here.³ In the last step of execution, a softmax classifier is applied to the entire table to select a cell as the answer. Details are further explained as below.

Let $r_i^{(t-1)}$ be the previous step’s row annotation result, where the subscript i indexes a particular row. We summarize global execution information (denoted as $g^{(t-1)}$) by max-pooling the row annotation $r^{(t-1)}$, i.e.,

$$g^{(t-1)} = \text{MaxPool}_i \{ r_i^{(t-1)} \} \quad (1)$$

In the current execution step, we first compute a distribution $p_f^{(t)}$ over all fields as “soft” field selection. The computation is based on the query q , the previous global information $g^{(t-1)}$, and the field name embeddings f , i.e.,

$$p_{f_j}^{(t)} = \text{softmax} \left(\text{MLP} ([q; f_j; g^{(t-1)}]) \right) \quad (2)$$

$$= \frac{\exp \{ \text{MLP} ([q; f_j; g^{(t-1)}]) \}}{\sum_{j'} \exp \{ \text{MLP} ([q; f_{j'}; g^{(t-1)}]) \}} \quad (3)$$

where $[\cdot; \cdot; \dots]$ denotes vector concatenation; MLP refers to a multi-layer perceptron.

Here, the weights of softmax are field embeddings (rather than parameters indexed by positions). In this way, there is no difference if one shuffles table columns. Besides, for a same field name, its embedding is shared among all training samples containing this field (but different tables may have different fields).

We represent the selected cell in each row as the sum of all cells in that row, weighted by soft field selection $p_f^{(t)}$. Formally, for the i -th row, we have

$$c_{\text{select}}^{(t)}[i] = \sum_j p_{f_j}^{(t)} c_{ij} \quad (4)$$

³In a pilot experiment, we tried a gating mechanism to indicate the results of row selection in hopes of aligning symbolic table execution. However, our preliminary experiments show that such gates do not exhibit much interpretation, but results in performance degradation. The distributed semantics provide more information than a 1-bit gate for a row.

where $[i]$ indexes a row of the selected column.

The current row annotation is computed by another MLP, based on the query q , previous execution results $r_i^{(t-1)}$, previous global information $g^{(t-1)}$, as well as the selected row in the current step $c_{\text{select}}^{(t)}[i]$ (the selection is in a soft manner), i.e.,

$$r_i^{(t)} = \text{MLP} \left([q, g^{(t-1)}, r_i^{(t-1)}, c_{\text{select}}^{(t)}[i]] \right) \quad (5)$$

As said, the last execution layer applies a softmax classifier over all cells to select an answer. Similar to Equation 3, the weights of softmax are not associated with positions, but the cell embeddings. In other words, the probability of choosing the i -th row, j -th column is

$$p_{ij} = \frac{\exp \{ \text{MLP} (q, g^{(t-1)}, r_i^{(t-1)}, c_{ij}) \}}{\sum_{i'} \sum_{j'} \exp \{ \text{MLP} (q, g^{(t-1)}, r_{i'}^{(t-1)}, c_{i'j'}) \}} \quad (6)$$

In this way, the neural executor is invariant with respect to the order of rows and columns. While order-sensitive architectures (e.g., convolutional/recurrent neural networks) might also model a table by implicitly ignoring such order information, the current treatment is more reasonable, which also better aligns with symbolic interpretation.

2.2. Symbolic Executor

The methodology of designing a symbolic executor is to define a set of primitive operators for the task, and then to use a machine learning model to predict the operator sequence and its arguments.

Our symbolic executor is different from the neural programmer (Neelakantan et al., 2016) in that we keep discrete/symbolic operators as well as execution results, whereas Neelakantan et al. (2016) fuse execution results by weighted average.

2.2.1. PRIMITIVE OPERATORS

We design six operators for symbolic execution, which are complete as they cover all types of queries in our scenario. Similar to the distributed executor, the result of one-step symbolic execution is some information for a row; here, we use a 0-1 boolean scalar, indicating whether a row is selected or not after a particular step of execution. Then a symbolic execution step takes previous results as input, with a column/field being the argument. The green boxes in Figure 1b illustrate the process and Table 2 summarizes our primitive operator set.

In Table 1, for example, the first step of execution is `argmax` over the column `Area`, with previous (initial) row selection being all ones. This step yields a single

Operator	Explanation
select_row	Choose a row whose value of a particular column is mentioned in the query
argmin	Choose the row from previously selected candidate rows with the minimum value in a particular column
argmax	Choose the row from previously selected candidate rows with the maximum value in a particular column
greater_than	Choose rows whose value in a particular column is greater than a previously selected row
less_than	Choose rows whose value in a particular column is less than a previously selected row
select_value	Choose the value of a particular column and of the previously selected row
EOE	Terminate, indicating the end of execution

Table 2. Primitive operators for symbolic execution.

row $\langle 2008, \text{Beijing}, \dots, 350, \dots, 25 \rangle$. The second execution operator is `select_value`, with an argument `column=Duration`, yielding the result 25. Then the executor terminates (EOE).

Stacked with multiple steps of primitive operators, the executor can answer fairly complicated questions like “How long is the last game which has smaller country size than the game whose host country GDP is 250?” In this example, the execution sequence is

1. `select_row`: select the row where the column is *GDP* and the value is mentioned in the query.
2. `less_than`: select rows whose country size is less than that of the previously selected row.
3. `argmax`: select the row whose year is the largest among previously selected rows.
4. `select_value`: choose the value of the previously selected row with the column being *Duration*.

Then the execution terminates. In our scenario, the execution is limited to four steps (EOE excluded) as such queries are already very complicated in terms of logical depth.

2.2.2. OPERATOR AND ARGUMENT PREDICTORS

We also leverage neural models, in particular RNNs, to predict the operator and its argument (a selected field/column).

Let $\mathbf{h}^{(t-1)}$ be the previous state’s hidden vectors. The current hidden state is

$$\mathbf{h}_{\text{op}}^{(t)} = \text{sigmoid}(W_{\text{op}}^{(\text{rec})} \mathbf{h}_{\text{op}}^{(t-1)}) \quad (7)$$

where $W_{\text{op}}^{(\text{rec})}$ is weight parameters. (A bias term is omitted in the equation for simplicity.) The initial hidden state is the query embedding, i.e., $\mathbf{h}_{\text{op}}^{(0)} = \mathbf{q}$.

The predicted probability of an operator i is given by

$$p_{\text{op}_i}^{(t)} = \text{softmax} \left\{ \mathbf{w}_{\text{op}_i}^{(\text{out})\top} \mathbf{h}_{\text{op}}^{(t)} \right\} \quad (8)$$

The operator with the largest predicted probability is selected for execution.

Our RNN here does not have input, because the execution sequence is not dependent on the result of the previous execution step. Such architecture is known as a Jordan-type RNN (Jordan, 1997; Mesnil et al., 2013).

Likewise, another Jordan-type RNN selects a field. The only difference lies in the weight of the output softmax, i.e., \mathbf{w}_i in Equation 8 is substituted with the embedding of a field/column name \mathbf{f} , given by

$$\mathbf{h}_{\text{field}}^{(t)} = \text{sigmoid}(W_{\text{field}}^{(\text{rec})} \mathbf{h}_{\text{field}}^{(t-1)}) \quad (9)$$

$$p_{f_j}^{(t)} = \text{softmax} \left\{ \mathbf{f}_j^\top \mathbf{h}_{\text{field}}^{(t)} \right\} \quad (10)$$

Training a symbolic executor without step-by-step supervision signals is non-trivial. A typical training method is reinforcement learning in a trial-and-error fashion. However, for a random initial policy, the probability of recovering an accurate execution sequence is extremely low. Given a 10×10 table, for example, the probability is $1/(6^4 \cdot 10^4) \approx 7.7 \times 10^{-8}$; the probability of obtaining an accurate denotation is 1%, which is also very low. Therefore, symbolic executors are not efficient in learning.

2.3. A Unified View

We now have two worlds of execution:

- The distributed executor is end-to-end learnable, but it is of low execution efficiency because of intensive matrix/vector multiplication during neural information processing. The fully neuralized execution also lacks explicit interpretation.
- The symbolic executor has high execution efficiency and explicit interpretation. However, it cannot be trained in an end-to-end manner, and suffers from the cold start problem of reinforcement learning.

We propose to combine the two worlds by using the distributed executor’s intermediate execution results to pre-train the symbolic executor for an initial policy; we then use the REINFORCE algorithm to improve the policy. The well-trained symbolic executor’s intermediate results can also be fed back to the distributed executor to improve performance.

2.3.1. DISTRIBUTED \rightarrow SYMBOLIC

We observe that the field attention in Equation 3 generally aligns with column selection in Equation 10. We therefore pre-train the column selector in the symbolic executor with labels predicted by a fully neuralized/distributed executor.

Such pretraining can obtain up to 70% accurate field selection and largely reduce the search space during reinforcement learning.

Formally, the operator predictor (Equation 8) and argument predictor (Equation 10) in each execution step are the *actions* (denoted as \mathcal{A}) in reinforcement learning terminologies. If we would like to pretrain m actions $a_1, a_2, \dots, a_m \in \mathcal{A}$, the cost function of a particular data sample is

$$J = - \sum_{i=1}^m \sum_{j=1}^{n_{\text{label}}^{(i)}} \hat{t}_j^{(i)} \log p_j^{(i)} \quad (11)$$

where $n_{\text{label}}^{(j)}$ is the number of labels (possible candidates) for the j -th action. $\mathbf{p}^{(i)} \in \mathbb{R}^{n_{\text{label}}^{(i)}}$ is the predicted probability by the operator/argument predictors in Figure 1b. $\hat{\mathbf{t}}^{(i)} \in \mathbb{R}^{n_{\text{label}}^{(i)}}$ is the induced action from the fully distributed model in Figure 1a. In our scenario, we only pretrain column predictors.

After obtaining a meaningful, albeit imperfect, initial policy, we apply REINFORCE (Williams, 1992) to improve the policy.

We define a binary reward R indicating whether the final result of symbolic execution matches the groundtruth denotation. The loss function of a policy is the negative expected reward where actions are sampled from the current predicted probabilities

$$J = -\mathbb{E}_{a_1, a_2, \dots, a_n \sim \theta} [R(a_1, a_2, \dots, a_n)] \quad (12)$$

The partial derivative for a particular sampled action is

$$\frac{\partial J}{\partial \mathbf{o}_i} = \tilde{R} \cdot (\mathbf{p}_i - \mathbf{1}_{a_i}) \quad (13)$$

where \mathbf{p}_i is the predicted probability of all possible actions at the time step i , $\mathbf{1}_{a_i}$ is a onehot representation of a sampled action a_i , and \mathbf{o}_i is the input (also known as *logit*) of softmax. \tilde{R} is the adjusted reward, which will be described shortly.

To help the training of REINFORCE, we have two tricks:

- We balance exploration and exploitation with a small probability ϵ . In other words, we sample an action from the predicted action distribution with probability $1 - \epsilon$ and from a uniform distribution over all possible actions with probability ϵ . The small fraction of uniform sampling helps the model to escape from poor local optima, as it continues to explore the entire action space during training.
- We adjust the reward by subtracting the mean reward, averaged over sampled actions for a certain data point. This is a common practice for REINFORCE (Ranzato et al., 2016). We also truncate negative rewards as

zero to prevent gradient from being messed up by incorrect execution. This follows the idea of “reward-inaction,” where unsuccessful trials are ignored (Section 2.4 in Sutton & Barto (1998)). The adjusted reward is denoted as \tilde{R} in Equation 13.

Notice that these tricks are applied to both coupled training and baselines for fairness.

2.3.2. DISTRIBUTED \rightarrow SYMBOLIC \rightarrow DISTRIBUTED

After policy improvement by REINFORCE, we could further feed back the symbolic executor’s intermediate results to the distributed one, akin to the step-by-step supervision setting in Yin et al. (2016b). The loss is a combination of denotation cross entropy loss $J_{\text{denotation}}$ and field attention cross entropy loss J_{fields} . (J_{fields} is similar to Equation 11, and details are not repeated). The overall training objective is $J = J_{\text{denotation}} + \lambda J_{\text{fields}}$, where λ is a hyperparameter balancing the two factors.

As will be seen in Section 3.3.5, feeding back intermediate results improves the distributed executor’s performance. This shows the distributed and symbolic worlds can indeed be coupled well.

3. Experiments

3.1. Dataset

We evaluated our approach on a QA dataset in Yin et al. (2016b). The dataset comprises 25k different tables and queries for training; validation and test sets contain 10k samples, respectively, and do not overlap with the training data. Each table is of size 10×10 , but different samples have different tables; the queries can be divided into four types: *SelectWhere*, *Superlative*, *WhereSuperlative*, and *NestQuery*, requiring 2–4 execution steps (EOE excluded).

We have both groundtruth denotation and execution actions (including operators and fields), as the dataset is synthesized by complicated rules and templates for research purposes. However, only denotations are used as labels during training, which is a realistic setting; execution sequences are only used during testing. For the sake of simplicity, we presume the number of execution steps is known *a priori* during training (but not during testing). Although we have such (little) knowledge of execution, it is not a limitation of our approach and out of our current focus. One can easily design a dummy operator to fill an unnecessary step or one can also train a discriminative sentence model to predict the number of execution steps if a small number of labels are available.

We chose the synthetic datasets because it is magnitudes larger than existing resources (e.g., WEBQUESTIONS).

Query type	SEMPRE [†]	Denotation			Execution		
		Distributed [†]	Symbolic	Coupled	Distributed	Symbolic	Coupled
SelectWhere	93.8	96.2	99.2	99.6	–	99.1	99.6
Superlative	97.8	98.9	100.0	100.0	–	100.0	100.0
WhereSuperlative	34.8	80.4	51.9	99.9	–	0.0	91.0
NestQuery	34.4	60.5	52.5	100.0	–	0.0	100.0
Overall	65.2	84.0	75.8	99.8	–	49.5	97.6

Table 3. Accuracies (in percentage) of the SEMPRES toolkit, the distributed neural enquirer, the symbolic executor, and our coupled approach. [†]Results reported in Yin et al. (2016b).

The process of data synthesizing also provides intermediate execution results for in-depth analysis. Like BABI for machine comprehension, our dataset and setting are a prerequisite for general semantic parsing. The data are available at our project website⁴; the code for data generation can also be downloaded to facilitate further development of the dataset.

3.2. Settings

The symbolic executor’s settings were generally derived from Yin et al. (2016b) so that we can have a fair comparison.⁵ The dimensions of all layers were in the range of 20–50; the learning algorithm was AdaDelta with default hyperparameters.

For the pretraining of the symbolic executor, we applied maximum likelihood estimation for 40 epochs to column selection with labels predicted by the distributed executor. We then used the REINFORCE algorithm to improve the policy, where we generated 10 action samples for each data point with the exploration probability ϵ being 0.1.

When feeding back to the distributed model, we chose λ from $\{0.1, 0.5, 1\}$ by validation to balance denotation error and field attention error. 0.5 outperforms the rest.

Besides neural networks, we also included the SEMPRES system as a baseline for comparison. The results are reported in Yin et al. (2016b), where a SEMPRES version that is specially optimized for table query is adopted (Pasupat & Liang, 2015). Thus it is suited in our scenario.

3.3. Results

3.3.1. OVERALL PERFORMANCE

Table 3 presents the experimental results of our coupled approach as well as baselines. Because reinforcement learning is much more noisy to train, we report the test accuracy

⁴<https://sites.google.com/site/copleneuralsymbolic/>

⁵One exception is the query RNN’s hidden states. Yin et al. (2016b) used 300d BiRNN, but we found it more likely to overfit in the symbolic setting, and hence we used 50d. This results in slower training, more rugged error surfaces, but higher peak performance.

corresponding to highest validation accuracy among three different random initializations (called *trajectories*). This is also known as a *restart* strategy for non-convex optimization.

As we see, both distributed and symbolic executors outperform the SEMPRES system, showing that neural networks can capture query information more effectively than human-engineered features. Further, the coupled approach also significantly outperforms both of them.

If trained solely by REINFORCE, the symbolic executor can recover the execution sequences for simple questions (*SelectWhere* and *Superlative*). However, for more complicated queries, it only learns last one or two steps of execution and has trouble in recovering early steps, even with the tricks in Section 2.3.1. This results in low execution accuracy but near 50% denotation accuracy because, in our scenario, we still have half chance to obtain an accurate denotation even if the nested (early) execution is wrong—the ultimate result is either in the candidate list or not, given a wrong where-clause execution.

By contrast, the coupled training largely improves the symbolic executor’s performance in terms of all query types.

3.3.2. INTERPRETABILITY

The accuracy of execution is crucial to the interpretability of a model. We say an execution is *accurate*, if all actions (operators and arguments) are correct. As shown above, an accurate denotation does not necessarily imply an accurate execution.

We find that coupled training recovers most correct execution sequences, that the symbolic executor alone cannot recover complicated cases, and that a fully distributed enquirer does not have explicit interpretations of execution.

The results demonstrate high interpretability of our approach, which is helpful for human understanding of execution processes.

3.3.3. LEARNING EFFICIENCY

We plot in Figure 3 the validation learning curves of the symbolic executor, trained by either reinforcement learning

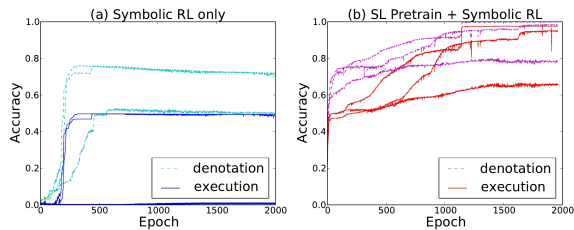


Figure 3. Validation learning curves. (a) Symbolic executor trained by REINFORCE (RL) only. (b) Symbolic executor with 40-epoch pretraining using a distributed executor in a supervised learning (SL) fashion. Both settings have three trajectories with different random initializations. Dotted lines: Denotation accuracy. Solid lines: Execution accuracy.

alone or our coupled approach.

Figure 3a shows that the symbolic executor is hard to train by REINFORCE alone: one trajectory obtains near-zero execution accuracy after 2000 epochs; the other two take 200 epochs to get started to escape from initial plateaus. Even if they achieve $\sim 50\%$ execution accuracy (for simple query types) and $\sim 75\%$ denotation accuracy, they are stuck in the poor local optima.

Figure 3b presents the learning curves of the symbolic executor pretrained with intermediate field attention of the distributed executor. Since the operation predictors are still hanging after pretraining, the denotation accuracy is near 0 before reinforcement learning. However, after only a few epochs of REINFORCE training, the performance increases sharply and achieves high accuracy gradually.

Notice that we have 40 epochs of (imperfectly) supervised pretraining. However, its time is negligible compared with reinforcement learning because in our experiments REINFORCE generates 10 samples and hence is theoretically 10 times slower. The results show that our coupled approach has much higher learning efficiency than a pure symbolic executor.

3.3.4. EXECUTION EFFICIENCY

Table 4 compares the execution efficiency of a distributed executor and our coupled approach. All neural networks are implemented in Theano with a TITAN Black GPU and Xeon e7-4820v2 (8-core) CPU; symbolic execution is assessed in C++ implementation. The comparison makes sense because the Theano platform is not specialized in symbolic execution, and fortunately, execution results do not affect actions in our experiment. Hence they can be easily disentangled.

As shown in the table, the execution efficiency of our approach is 2–5 times higher than the distributed executor, depending on the implementation. The distributed executor is when predicting because it maps every token to a

	Fully Distributed	Our approach		
		Op/Arg Pred.	Symbolic Exe. [†]	Total
CPU	13.86	2.65	0.002	2.65
GPU	1.05	0.44		0.44

Table 4. Execution efficiency. We present the running time (in seconds) of the test set, containing 10k samples. [†]The symbolic execution is assessed in C++ implementation. Others are implemented in Theano, including the fully distributed model as well as the operator and argument predictors.

Training Method	Accuracy (%)
End-to-end (w/ denotation labels) [†]	84.0
Step-by-step (w/ execution labels) [†]	96.4
Feeding back	96.5

Table 5. The accuracy of a fully distributed model, trained by different methods. In the last row, we first train a distributed executor and feed its intermediate execution results to the symbolic one; then the symbolic executor’s intermediate results are fed back to the distributed one. [†]Reported in Yin et al. (2016b).

distributed real-valued vector, resulting in intensive matrix-vector operations. The symbolic executor only needs a neural network to predict actions (operators and arguments), and thus is more lightweight. Further, we observe the execution itself is blazingly fast, implying that, compared with distributed models, our approach could achieve even more efficiency boost with a larger table or more complicated operation.

3.3.5. FEEDING BACK

We now feed back the well-trained symbolic executor’s intermediate results to the fully distributed one. As our well-trained symbolic executor has achieved high execution accuracy, this setting is analogous to strong supervision with step-by-step groundtruth signals, and thus it also achieves similar performance,⁶ shown in Table 5.

We showcase the distributed executor’s field attention⁷ in Figure 4. If trained in an end-to-end fashion, the neural network exhibits interpretation in the last three steps of this example, but in the early step, the field attention is incorrect (also more uncertain as it scatters a broader range). After feeding back the symbolic executor’s intermediate results as step-by-step supervision, the distributed executor exhibits near-perfect field attention.

This experiment further confirms that the distributed and

⁶We even have 0.1% performance boost compared with the step-by-step setting, but we think it should be better explained as variance of execution.

⁷The last neural executor is a softmax layer over all cells. We marginalize over rows to obtain the field probability.

Query: How many people watched the earliest game whose host country GDP is larger than the game in Cape Town?

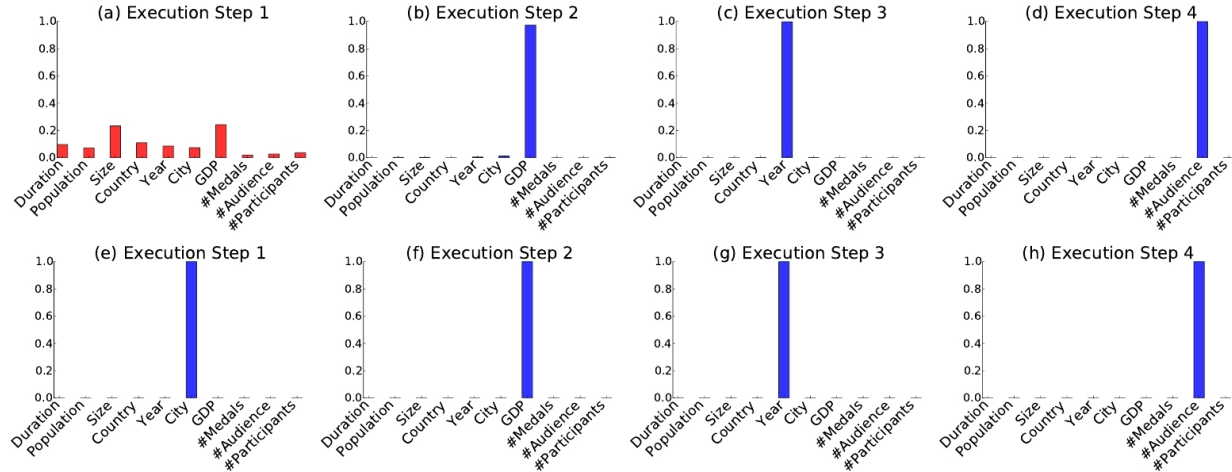


Figure 4. Distributed executor’s intermediate results of field attention. **Top:** Trained in an end-to-end fashion (a–d). **Bottom:** One-round co-training of distributed and symbolic executors (e–h). The red plot indicates incorrect field attention.

symbolic worlds can indeed be coupled well. In more complicated applications, there could also be possibilities in iteratively training one model by leveraging the other in a co-training fashion.

4. Related Work and Discussions

Neural execution has recently aroused much interest in the deep learning community. Besides SQL-like execution as has been extensively discussed in previous sections, neural Turing machines (Graves et al., 2014) and neural programmer-interpreters (Reed & De Freitas, 2016) are aimed at more general “programs.” The former is a “distributed analog” to Turing machines with soft operators (e.g., read, write, and address); its semantics, however, cannot be grounded to actual operations. The latter learns to generate an execution trace in a fully supervised, step-by-step manner.

Another related topic is incorporating neural networks with external (hard) mechanisms. Hu et al. (2016) propose to better train a neural network by leveraging the classification distribution induced from a rule-based system. Lei et al. (2016) propose to induce a sparse code by REINFORCE to make neural networks focus on relevant information. In machine translation, Mi et al. (2016) use alignment heuristics to train the attention signal of neural networks in a supervised manner. In these studies, researchers typically leverage external hard mechanisms to improve neural networks’ performance.

The uniqueness of our work is to train a fully neuralized/distributed model first, which takes advantage of its differentiability, and then to guide a symbolic model to achieve a meaningful initial policy. Further trained by reinforcement learning, the symbolic model’s knowledge can

improve neural networks’ performance by feeding back step-by-step supervision. Our work sheds light on neural sequence prediction in general, for example, exploring word alignment (Mi et al., 2016) or chunking information (Zhou et al., 2017) in machine translation by coupling neural and external mechanisms.

5. Conclusion and Future Work

In this paper, we have proposed a coupled view of distributed and symbolic execution for natural language queries. By pretraining with intermediate execution results of a distributed executor, we manage to accelerate the symbolic model’s REINFORCE training to a large extent. The well-trained symbolic executor could also guide a distributed executor to achieve better performance. Our proposed approach takes advantage of both distributed and symbolic worlds, achieving high interpretability, high execution efficiency, high learning efficiency, as well as high accuracy.

As a pilot study, our paper raises several key open questions: When do neural networks exhibit symbolic interpretations? How can we better transfer knowledge between distributed and symbolic worlds?

In future work, we would like to design interpretable operators in the distributed model to better couple the two worlds and to further ease the training with REINFORCE. We would also like to explore different ways of transferring knowledge, e.g., distilling knowledge from the action distributions (rather than using the max *a posteriori* action), or sampling actions by following the distributed model’s predicted distribution during symbolic one’s Monte Carlo policy gradient training (REINFORCE).

Acknowledgments

We thank Pengcheng Yin and Jiatao Gu for helpful discussions; we also thank the reviewers for insightful comments. This research is partially supported by the National Basic Research Program of China (the 973 Program) under Grant Nos. 2014CB340301 and 2015CB352201, and the National Natural Science Foundation of China under Grant Nos. 614201091, 61232015 and 61620106007.

References

- Berant, Jonathan, Chou, Andrew, Frostig, Roy, and Liang, Percy. Semantic parsing on Freebase from question-answer pairs. In *EMNLP*, pp. 1533–1544, 2013.
- Dong, Li and Lapata, Mirella. Language to logical form with neural attention. In *ACL*, pp. 33–43, 2016.
- Graves, Alex, Wayne, Greg, and Danihelka, Ivo. Neural Turing machines. *arXiv preprint arXiv:1410.5401*, 2014.
- Hu, Zhiting, Ma, Xuezhe, Liu, Zhengzhong, Hovy, Eduard, and Xing, Eric. Harnessing deep neural networks with logic rules. In *ACL*, pp. 2410–2420, 2016.
- Jordan, Michael I. Serial order: A parallel distributed processing approach. *Advances in Psychology*, 121:471–495, 1997.
- Lei, Tao, Barzilay, Regina, and Jaakkola, Tommi. Rationalizing neural predictions. In *EMNLP*, pp. 107–117, 2016.
- Liang, Chen, Berant, Jonathan, Le, Quoc, Forbus, Kenneth D, and Lao, Ni. Neural symbolic machines: Learning semantic parsers on Freebase with weak supervision. In *ACL (to appear)*, 2017.
- Mesnil, Grégoire, He, Xiaodong, Deng, Li, and Bengio, Yoshua. Investigation of recurrent-neural-network architectures and learning methods for spoken language understanding. In *INTERSPEECH*, pp. 3771–3775, 2013.
- Mi, Haitao, Sankaran, Baskaran, Wang, Zhiguo, and Ittycheriah, Abe. Coverage embedding models for neural machine translation. In *EMNLP*, pp. 955–960, 2016.
- Mikolov, Tomas, Sutskever, Ilya, Chen, Kai, Corrado, Greg S, and Dean, Jeff. Distributed representations of words and phrases and their compositionality. In *NIPS*, pp. 3111–3119, 2013.
- Neelakantan, Arvind, Le, Quoc V, and Sutskever, Ilya. Neural programmer: Inducing latent programs with gradient descent. In *ICLR*, 2016.
- Pasupat, Panupong and Liang, Percy. Compositional semantic parsing on semi-structured tables. In *ACL-IJCNLP*, pp. 1470–1480, 2015.
- Ranzato, MarcAurelio, Chopra, Sumit, Auli, Michael, and Zaremba, Wojciech. Sequence level training with recurrent neural networks. In *ICLR*, 2016.
- Reed, Scott and De Freitas, Nando. Neural programmer-interpreters. In *ICLR*, 2016.
- Sutton, Richard S and Barto, Andrew G. *Reinforcement Learning: An Introduction*. MIT Press Cambridge, 1998.
- Wen, Tsung-Hsien, Vandyke, David, Mrkšić, Nikola, Gasic, Milica, Rojas Barahona, Lina M., Su, Pei-Hao, Ultes, Stefan, and Young, Steve. A network-based end-to-end trainable task-oriented dialogue system. In *EACL*, pp. 438–449, 2017.
- Williams, Ronald J. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8(3):229–256, 1992.
- Xiao, Chunyang, Dymetman, Marc, and Gardent, Claire. Sequence-based structured prediction for semantic parsing. In *ACL*, pp. 1341–1350, 2016.
- Yin, Jun, Jiang, Xin, Lu, Zhengdong, Shang, Lifeng, Li, Hang, and Li, Xiaoming. Neural generative question answering. In *IJCAI*, pp. 2972–2978, 2016a.
- Yin, Pengcheng, Lu, Zhengdong, Li, Hang, and Kao, Ben. Neural enquirer: Learning to query tables with natural language. In *IJCAI*, pp. 2308–2314, 2016b.
- Zhou, Hao, Tu, Zhaopeng, Huang, Shujian, Liu, Xiaohua, Li, Hang, and Chen, Jiajun. Chunk-based bi-scale decoder for neural machine translation. In *ACL*, 2017.