

Improving Multi-hop Knowledge Base Question Answering by Learning Intermediate Supervision Signals

Gaole He^{1,4†}, Yunshi Lan², Jing Jiang², Wayne Xin Zhao^{3,4*} and Ji-Rong Wen^{1,3,4}

¹School of Information, Renmin University of China

²School of Information System, Singapore Management University

³Gaoling School of Artificial Intelligence, Renmin University of China

⁴Beijing Key Laboratory of Big Data Management and Analysis Methods

{hegaole, jrwen}@ruc.edu.cn, batmanfly@gmail.com, {yslan, jingjiang}@smu.edu.sg

ABSTRACT

Multi-hop Knowledge Base Question Answering (KBQA) aims to find the answer entities that are multiple hops away in the Knowledge Base (KB) from the entities in the question. A major challenge is the lack of supervision signals at intermediate steps. Therefore, multi-hop KBQA algorithms can only receive the feedback from the final answer, which makes the learning unstable or ineffective.

To address this challenge, we propose a novel teacher-student approach for the multi-hop KBQA task. In our approach, the student network aims to find the correct answer to the query, while the teacher network tries to learn intermediate supervision signals for improving the reasoning capacity of the student network. The major novelty lies in the design of the teacher network, where we utilize both forward and backward reasoning to enhance the learning of intermediate entity distributions. By considering bidirectional reasoning, the teacher network can produce more reliable intermediate supervision signals, which can alleviate the issue of spurious reasoning. Extensive experiments on three benchmark datasets have demonstrated the effectiveness of our approach on the KBQA task.

CCS CONCEPTS

• **Computing methodologies** → **Reasoning about belief and knowledge; Search with partial observations.**

KEYWORDS

Knowledge Base Question Answering; Teacher-student Network; Intermediate Supervision Signals

ACM Reference Format:

Gaole He, Yunshi Lan, Jing Jiang, Wayne Xin Zhao and Ji-Rong Wen. 2021. Improving Multi-hop Knowledge Base Question Answering by Learning Intermediate Supervision Signals. In *Proceedings of the Fourteenth ACM*

*Corresponding author.

† This work is done when the first author visited SMU.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

WSDM '21, March 8–12, 2021, Virtual Event, Israel

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8297-7/21/03...\$15.00

<https://doi.org/10.1145/3437963.3441753>

International Conference on Web Search and Data Mining (WSDM '21), March 8–12, 2021, Virtual Event, Israel. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3437963.3441753>




Figure 1: A spurious case from MetaQA-3hop dataset. We use green, red, yellow and grey circles to denote the topic entities, correct answer, intermediate entities and irrelevant entities respectively.

1 INTRODUCTION

Knowledge Base Question Answering (KBQA) is a challenging task that aims at finding answers to questions expressed in natural language from a given knowledge base (KB). Traditional solutions [2, 5, 7, 36] usually develop a specialized pipeline consisting of multiple machine-learned or hand-crafted modules (e.g., named entity recognition, entity linking). Recently, end-to-end deep neural networks [21, 31] become the popular paradigm for this task by automatically learning data representations and network parameters.

For the KBQA task, there have been growing interests in solving complex questions that require a multi-hop reasoning procedure [20], called *multi-hop KBQA*. Besides the final answer, it is also important that a multi-hop KBQA algorithm can identify a reasonable relation path leading to the answer entities [6, 27]. In some cases, even if the answer was correctly found, the relation path might be spurious. We present an example of spurious multi-hop reasoning in Fig. 1. The question is “what types are the films starred by actors in the nine lives of fritz the cat?”. Besides the correct path (with red arrows), two spurious paths (with blue arrows) which include entities who are directors at the first step can also reach the correct answer. It is mainly due to the lack of supervision signals at the intermediate reasoning steps (which we call *intermediate supervision signals*). For the multi-hop KBQA task, training data is typically in the form of $\langle \text{question}, \text{answer} \rangle$ instead of the ideal form of $\langle \text{question}, \text{relation path} \rangle$. Therefore, multi-hop reasoning

algorithms can only receive the feedback at the final answer using such datasets.

To address this issue, several studies formulate multi-hop KBQA as a reinforcement learning (RL) task [3, 22, 27]. They set up a policy-based agent to sequentially extend its inference path until it reaches a target entity. Its states are usually defined as tuple of query and current entity, and action as traverse on KB through outgoing edges of current entity. RL-based approaches heavily rely on the terminal reward to bias the search. To prevent spurious paths in the search, reward shaping [22, 27] and action dropout [22] have been proposed to improve the model learning. However, these solutions either require expert experience or still lack effective supervision signals at intermediate steps.

Different from previous studies, our idea is to set up two models with different purposes for multi-hop KBQA. The main model aims to find the correct answer to the query, while the auxiliary model tries to learn intermediate supervision signals for improving the reasoning capacity of the main model. Specifically, the auxiliary model infers which entities at the intermediate steps are more relevant to the question, and these entities are considered as intermediate supervision signals. Although the idea is appealing, it is challenging to learn an effective auxiliary model, since we do not have such labeled data for training.

Our solution is inspired by the bidirectional search algorithms (e.g., bidirectional BFS [17]) on graphs, in which an ideal path connecting the source and the destination can be more effectively identified with bidirectional exploration. Indeed, for KBQA we also have two different views to consider the task setting: the forward reasoning that finds the path starting from the topic entities (i.e., entities in the queries) to the answer entities and the backward reasoning that returns from answer entities to the topic entities. Most existing methods only consider forward reasoning. However, it is possible to jointly model the two reasoning processes, since topic entities and answer entities are all known in the training data. Such a bidirectional reasoning mechanism is able to incorporate additional self-supervision signals at intermediate steps. As shown in Fig. 1, the entity distribution obtained by forward reasoning at the second step should be similar to that from backward reasoning at the first step. Irrelevant entities “*Devil’s Doorway*” and “*Coffin Rock*” are likely to be reached at the second reasoning step of forward reasoning but unreachable at the first step of backward reasoning. To maintain the correspondence between the two processes, we should avoid including the director “*Robert Taylor*” at the first step of forward reasoning. Such a potential correspondence is useful to improve the learning of each individual reasoning process at intermediate steps. That is the key point how we learn reliable intermediate supervision signals.

To this end, in this paper, we propose a novel teacher-student approach for the multi-hop KBQA task. Specifically, the student network (i.e., the main model), which aims to find the answer, is implemented by adapting the Neural State Machine (NSM) [14] from visual question answering. In our approach, the student network can improve itself according to intermediate entity distributions learned from the teacher network. The major novelty lies in the design of the teacher network (i.e., the auxiliary model), which provides intermediate supervision signals. We utilize the correspondence between the state information from the forward

and backward reasoning processes to enhance the learning of intermediate entity distributions. We further design two reasoning architectures that support the integration between forward and backward reasoning. By considering bidirectional reasoning, the teacher network can alleviate the issue of spurious reasoning, and produce more reliable intermediate supervision signals.

To evaluate our approach, we conduct extensive experiments on three benchmark datasets. Extensive experiments have demonstrated the effectiveness of our approach on the multi-hop KBQA task, especially for cases lacking training data. To the best of our knowledge, it is the first time that intermediate supervision signals have been explicitly learned with a teacher-student framework.

2 RELATED WORK

Our work is closely related to the studies on KBQA, multi-hop reasoning and teacher-student framework.

Knowledge Base Question Answering. For the KBQA task, various methods have been developed over the last decade. They can be categorized into two groups: semantic parsing based methods and retrieval based methods. Semantic parsing based methods [2, 19, 21, 36, 37] learn a semantic parser that converts natural language questions into intermediate logic forms, which can be executed against a KB. Retrieval-based methods [5, 24, 30, 31, 34] directly retrieve answers from the KB in light of the information conveyed in the questions.

Recently, researchers pay more attention to multi-hop based KBQA. Some work [24, 31, 38] employed classical methods (e.g., Variational Reasoning Network, Key-Value Memory Network and Graph Convolution Network) to conduct multi-hop reasoning within the KB. Moreover, Sun *et al.* [30] and Saxena *et al.* [29] leveraged extra corpus and enriched knowledge graph embeddings to boost the performance of multi-hop KBQA. However, these methods take the performance of final prediction as the only objective, which are vulnerable to the spurious examples.

Multi-hop Reasoning. In recent years, multi-hop reasoning becomes a hot research topic for both computer vision and natural language processing domains. Min *et al.* [25] proposed to decompose complex queries into several 1-hop queries and solved them by turn. Hudson *et al.* [13] designed a novel recurrent Memory, Attention, and Composition (MAC) cell, which splits complex reasoning into a series of attention-based reasoning steps. Das *et al.* [3, 22] conducted multi-hop reasoning on a graph under the reinforcement learning setting and treated every reasoning step as an edge transition on the graph. Besides, there are quite a few studies that adopt Graph Neural Network (GNN) [16, 33] to conduct explicit reasoning on graph structure [12, 31].

Teacher-student Framework. Knowledge distillation (KD) is introduced and generalized by early work [10]. They proposed a teacher-student framework, where a complicated high-performance model and a light-weight model are treated as teacher and student respectively. The predictions of the teacher model are treated as “soft labels” and the student model is trained to fit the soft labels. While knowledge distillation was initially proposed for model compression, recent work [9, 39] found that applying the soft labels as the training target can help the student achieve better performance.

Several studies also apply the teacher-student framework in question answering task. Yang *et al.* [35] designed a multi-teacher knowledge distillation paradigm in a Web Question Answering system. Do *et al.* [4] and Hu *et al.* [11] applied the teacher-student framework to visual question answering task and reading comprehension task, respectively. In this work, we try to address spurious reasoning caused by weak supervision in the multi-hop KBQA task with an elaborate teacher-student framework.

3 PRELIMINARY

In this section, we introduce the background and define the task.

Knowledge Base (KB). A knowledge base typically organizes factual information as a set of triples, denoted by $\mathcal{G} = \{\langle e, r, e' \rangle | e, e' \in \mathcal{E}, r \in \mathcal{R}\}$, where \mathcal{E} and \mathcal{R} denote the entity set and relation set, respectively. A triple $\langle e, r, e' \rangle$ denotes that relation r exists between head entity e and tail entity e' . Furthermore, we introduce *entity neighborhood* to denote the set of triples involving an entity e , denoted by $\mathcal{N}_e = \{\langle e, r, e' \rangle \in \mathcal{G}\} \cup \{\langle e', r, e \rangle \in \mathcal{G}\}$, containing both incoming and outgoing triples for e . For simplicity, we replace a triple $\langle e, r, e' \rangle$ with its reverse triple $\langle e', r^{-1}, e \rangle$, so that we can have $\mathcal{N}_e = \{\langle e', r, e \rangle \in \mathcal{G}\}$. For convenience, we further use italic bold fonts to denote the embeddings of entities or relations. Let $E \in \mathbb{R}^{d \times |\mathcal{E}|}$ and $R \in \mathbb{R}^{d \times |\mathcal{R}|}$ denote the embedding matrices for entities and relations in the KB, respectively, and each column vector $e \in \mathbb{R}^d$ or $r \in \mathbb{R}^d$ is a d -dimensional embedding for entity e or relation r .

Knowledge Base Question Answering (KBQA). We focus on factoid question answering over a knowledge base. We assume that a KB \mathcal{G} is given as the available resource and the answers will be the entities in \mathcal{G} . Formally, given a natural language question $q = \{w_1, w_2, \dots, w_l\}$ and a KB \mathcal{G} , the task of KBQA is to figure out the answer entity(ies), denoted by the set \mathcal{A}_q , to query q from the candidate entity set \mathcal{E} . The entities mentioned in a question are called *topic entities*. Specially, we consider solving complex questions where the answer entities are multiple hops away from the topic entities in the KB, called *multi-hop KBQA*.

4 THE PROPOSED APPROACH

In this section, we present the proposed approach for the multi-hop KBQA task under the teacher-student framework.

4.1 Overview

A major difficulty for multi-hop KBQA is that it usually lacks supervision signals at intermediate reasoning steps, since only the answer entities are given as ground-truth information. To tackle this issue, we adopt the recently proposed teacher-student learning framework [10, 28]. The main idea is to train a student network that focuses on the multi-hop KBQA task itself, while another teacher network is trained to provide (pseudo) supervision signals (*i.e.*, inferred entity distributions in our task) at intermediate reasoning steps for improving the student network.

In our approach, the student network is implemented based on Neural State Machine (NSM) [14], which was originally proposed for visual question answering on scene graph extracted from image data. We adapt it to the multi-hop KBQA task by considering KB

as a graph, and maintain a gradually learned entity distribution over entities during the multi-hop reasoning process. To develop the teacher network, we modify the architecture of NSM by incorporating a novel bidirectional reasoning mechanism, so that it can learn more reliable entity distributions at intermediate reasoning steps, which will be subsequently used by the student network as the supervision signals.

In what follows, we first describe the adapted architecture of NSM for multi-hop KBQA, and then present the teacher network and model learning.

4.2 Neural State Machine for Multi-hop KBQA

We present an overall sketch of NSM in Fig. 2. It mainly consists of an instruction component and a reasoning component. The instruction component sends instruction vectors to the reasoning component, while the reasoning component infers the entity distribution and learns the entity representations.

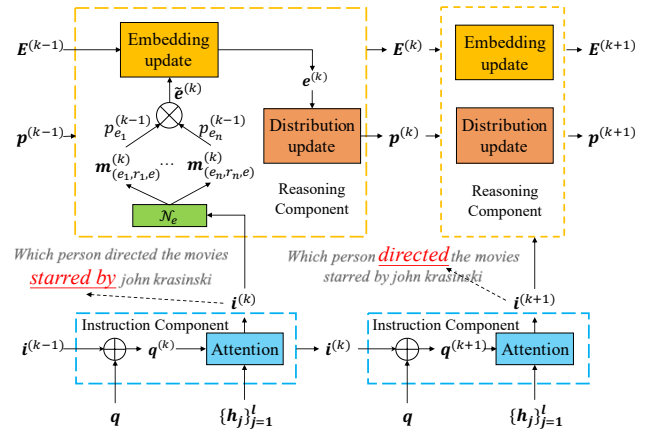


Figure 2: Illustration of the two reasoning steps for neural state machine on question “which person directed the movies starred by john krasinski?”. In different reasoning steps, the instruction vector focuses on different parts of the question.

4.2.1 Instruction Component. We first describe how to transform a given natural language question into a series of instruction vectors that control the reasoning process. The input of the instruction component consists of a query embedding and an instruction vector from the previous reasoning step. The initial instruction vector is set as zero vector. We utilize GloVe [26] to obtain the embeddings of the query words. Then we adopt a standard LSTM encoder to obtain a set of hidden states $\{h_j\}_{j=1}^l$, where $h_j \in \mathbb{R}^d$ and l is the length of the query. After that, the last hidden state is considered to be the question representation, *i.e.*, $q = h_l$. Let $i^{(k)} \in \mathbb{R}^d$ denote the instruction vector at the k -th reasoning step. We adopt the following method to learn the instruction vector $i^{(k)}$:

$$\begin{aligned} i^{(k)} &= \sum_{j=1}^l \alpha_j^{(k)} h_j, \\ \alpha_j^{(k)} &= \text{softmax}_j(W_\alpha(q^{(k)} \odot h_j) + b_\alpha), \\ q^{(k)} &= W^{(k)}[i^{(k-1)}; q] + b^{(k)}, \end{aligned} \quad (1)$$

where $W^{(k)} \in \mathbb{R}^{d \times 2d}$, $W_\alpha \in \mathbb{R}^{d \times d}$ and $b^{(k)}, b_\alpha \in \mathbb{R}^d$ are parameters to learn. The core idea is to attend to specific parts of a query when learning the instruction vectors at different time steps. In such a process, we also dynamically update the query representation, so that it can incorporate the information of previous instruction vectors. By repeating the process above, we can obtain a list of instruction vectors $\{i^{(k)}\}_{k=1}^n$ after n reasoning steps.

4.2.2 Reasoning Component. Once we obtain the instruction vector $i^{(k)}$, we can use it as a guiding signal for the reasoning component. The input of the reasoning component consists of the instruction vector of the current step, and the entity distribution and entity embeddings obtained from the previous reasoning step. The output of the reasoning component includes the entity distribution $p^{(k)}$ and the entity embeddings $\{e^{(k)}\}$. First, we set the initial entity embeddings by considering the relations involving e :

$$e^{(0)} = \sigma \left(\sum_{\langle e', r, e \rangle \in N_e} r \cdot W_T \right), \quad (2)$$

where $W_T \in \mathbb{R}^{d \times d}$ are the parameters to learn. Unlike previous studies [24, 31], we explicitly utilize the information of related relation types for encoding entities. In the multi-hop KBQA task, a reasoning path consisting of multiple relation types can reflect important semantics that lead to the answer entities. Besides, such a method is also useful to reduce the influence of noisy entities, and easy to apply to unseen entities of known context relations. Note that we do not use the original embedding of e when initializing $e^{(0)}$ because for intermediate entities along the reasoning path the identifiers of these entities are not important; it is the relations that these intermediate entities are involved in that matter the most.

Given a triple $\langle e', r, e \rangle$, a match vector $m_{\langle e', r, e \rangle}^{(k)}$ is learned by matching the current instruction $i^{(k)}$ with relation vector r :

$$m_{\langle e', r, e \rangle}^{(k)} = \sigma(i^{(k)} \odot W_R r), \quad (3)$$

where $W_R \in \mathbb{R}^{d \times d}$ are the parameters to learn. Furthermore, we aggregate the matching messages from neighboring triples and assign weights to them according to how much attention they receive at the last reasoning step:

$$\tilde{e}^{(k)} = \sum_{\langle e', r, e \rangle \in N_e} p_{e'}^{(k-1)} \cdot m_{\langle e', r, e \rangle}^{(k)}, \quad (4)$$

where $p_{e'}^{(k-1)}$ is the assigned probability of entity e' at the last reasoning step, which we will explain below. Such a representation is able to capture the relation semantics associated with an entity in the KB. Then, we update entity embeddings as follows:

$$e^{(k)} = \text{FFN}([e^{(k-1)}; \tilde{e}^{(k)}]), \quad (5)$$

where $\text{FFN}(\cdot)$ is a feed-forward layer taking as input of both previous embedding $e^{(k-1)}$ and relation-aggregated embedding $\tilde{e}^{(k)}$.

Through such a process, both the relation path (from topic entities to answer entities) and its matching degree with the question can be encoded into node embeddings. The probability distribution over intermediate entities derived at step k can be calculated as:

$$p^{(k)} = \text{softmax}(E^{(k)T} w), \quad (6)$$

where $E^{(k)}$ is a matrix where each column vector is the embedding of an entity at the k -th step, and $w \in \mathbb{R}^d$ are the parameters that derive the entity distribution $p^{(k)}$, and $E^{(k)}$ is the updated entity embedding matrix by Eq. 5.

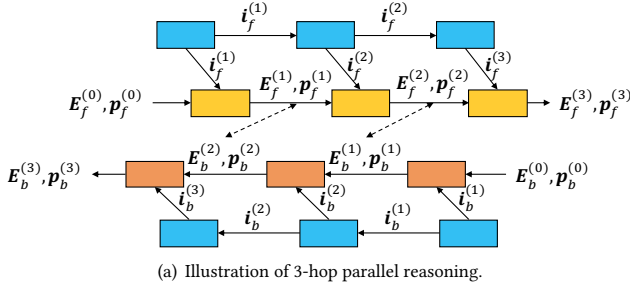
4.2.3 Discussion. For our task, the reason that we adopt the NSM model as the student network are twofold. First, our core idea is to utilize intermediate entity distributions derived from the teacher network as the supervision signals for the student network. In contrast, most previous multi-hop KBQA methods do not explicitly maintain and learn such an entity distribution at intermediate steps. Second, NSM can be considered as a special graph neural network, which has excellent reasoning capacity over the given knowledge graph. As shown in Section 4.2.2, the learning of entity distributions and entity embeddings can indeed correspond to the general “propagate-then-aggregate” update mechanism of graph neural networks. We would like to utilize such a powerful neural architecture to solve the current task.

The NSM [14] was proposed to conduct visual reasoning in an abstract latent space. We make two major adaptations for multi-hop KBQA. First, in Eq. 2, we initialize the node embeddings by aggregating the embeddings of those relations involving the entity. In our task, the given KB is usually very large. An entity is likely to be linked to a large number of other entities. Our initialization method is able to reduce the influence of noisy entities, focusing on the important relational semantics. Besides, it is also easy to generalize to new or unseen entities with known relations, which is especially important to incremental training. Second, in Eq. 5, we update entity embeddings by integrating previous embedding $e^{(k-1)}$ and relation-aggregated embedding $\tilde{e}^{(k)}$. For comparison, original NSM [14] separately modeled the two parts, whereas we combine the two factors in a unified update procedure, which is useful to derive more effective node embeddings.


4.3 The Teacher Network

Different from the student network, the teacher network aims to learn or infer reliable entity distributions at intermediate reasoning steps. Note that there are no such labeled entity distributions for training the teacher network. Instead, inspired by the bidirectional search algorithm (e.g., bidirectional BFS [17]), we incorporate the bidirectional reasoning mechanism for enhancing the learning of intermediate entity distributions in the teacher network.

4.3.1 Bidirectional Reasoning for Multi-hop KBQA. Given a knowledge base, the reasoning process for multi-hop KBQA can be considered to be an exploration and search problem on the graph. Most existing multi-hop KBQA methods start from the topic entities and then look for the possible answer entities, called *forward reasoning*. On the other hand, the opposite search from answer entities to topic entities (which we refer to as *backward reasoning*) has been neglected by previous studies. Our core idea is to consider the exploration in both directions and let the two reasoning processes synchronize with each other at intermediate steps. In this way, the derived intermediate entity distributions can be more reliable than those learned from a single direction. More specifically, given a n -hop reasoning path, let $p_f^{(k)}$ and $p_b^{(n-k)}$ denote the entity distributions from the forward reasoning at the k -th step and from the



(a) Illustration of 3-hop parallel reasoning.



(b) Illustration of 3-hop hybrid reasoning.

Figure 3: Illustration of the designed teacher architectures. We use blue, yellow and orange squares to denote the instruction component, forward reasoning component and backward reasoning component, respectively. The dotted arrows link the corresponding intermediate entity distributions of the two reasoning processes. We use f and b as subscripts to distinguish forward reasoning and backward reasoning, respectively.

backward reasoning at the $(n - k)$ -th step, respectively. The key point is that the two distributions should be similar or consistent if the two reasoning processes have been stable and accurate, *i.e.*, $\mathbf{p}_f^{(k)} \approx \mathbf{p}_b^{(n-k)}$. We will utilize such a correspondence as constraints in the following models.

4.3.2 Reasoning Architectures. Based on the idea above, we design two kinds of neural architectures for the teacher network, namely parallel reasoning and hybrid reasoning.

Parallel Reasoning. The first way is to set up two separate NSMs for both forward and backward reasoning, respectively. These two NSM networks are relatively isolated, and do not share any parameters. We only consider incorporating correspondence constraints on the intermediate entity distributions between them.

Hybrid Reasoning. In the second way, we share the same instruction component and arrange the two reasoning processes in a cycled pipeline. Besides the correspondence constraints, the two processes receive the same instruction vectors. Furthermore, the derived information at the final step of the forward reasoning is fed into the backward reasoning as initial values. Formally, the following equations hold in this case:

$$\begin{aligned} \mathbf{p}_b^{(0)} &= \mathbf{p}_f^{(n)}, \mathbf{E}_b^{(0)} = \mathbf{E}_f^{(n)}, \\ \mathbf{i}_b^{(k)} &= \mathbf{i}_f^{(n+1-k)}, k = 1, \dots, n. \end{aligned} \quad (7)$$

We present the illustrative examples of the parallel reasoning and hybrid reasoning in Fig. 3(a) and Fig. 3(b). Comparing the two reasoning architectures, it can be seen that parallel reasoning has a more loose integration, while hybrid reasoning requires a deeper fusion between the information from both reasoning processes. Unlike bidirectional BFS, in our task, backward reasoning might not be able to exactly mimic the inverse process of forward reasoning, since the two processes correspond to different semantics in multi-hop KBQA. Considering this issue, we share the instruction vectors and recycle the final state of the forward reasoning for initializing backward reasoning. In this way, backward reasoning receives more information about forward reasoning, so that it can better trace back the reasoning path of forward reasoning.

4.4 Learning with the Teacher-Student Framework

In this part, we present the details of model learning with our teacher-student framework.

4.4.1 Optimizing the Teacher Network. The two reasoning architectures of the teacher network can be optimized in the same way. We mainly consider two parts of loss, namely reasoning loss and correspondence loss.

The reasoning loss reflects the capacity of predicting the accurate entities, which can be decomposed into two directions:

$$\mathcal{L}_f = D_{KL}(\mathbf{p}_f^{(n)}, \mathbf{p}_f^*), \mathcal{L}_b = D_{KL}(\mathbf{p}_b^{(n)}, \mathbf{p}_b^*), \quad (8)$$

where $\mathbf{p}_f^{(n)}$ ($\mathbf{p}_b^{(n)}$) denotes the final entity distribution for forward (backward) reasoning process, \mathbf{p}_f^* (\mathbf{p}_b^*) denotes the ground-truth entity distribution, and $D_{KL}(\cdot, \cdot)$ is the Kullback-Leibler divergence [18], which measures the difference between the two distributions in an asymmetric way. To obtain \mathbf{p}_f^* and \mathbf{p}_b^* , we transform the occurrences of ground-truth entities into a frequency-normalized distribution. Specifically, if k entities in the graph are ground-truth entities, they are assigned a probability of $\frac{1}{k}$ in the final distribution.

The correspondence loss reflects the consistency degree between intermediate entity distributions from the two reasoning processes. It can be computed by summing the loss at each intermediate step:

$$\mathcal{L}_c = \sum_{k=1}^{n-1} D_{JS}(\mathbf{p}_f^{(k)}, \mathbf{p}_b^{(n-k)}), \quad (9)$$

where $D_{JS}(\cdot, \cdot)$ is the Jensen-Shannon divergence [8], which measures the difference between two distributions in a symmetric way.

To combine the above loss terms, we define the entire loss function of the teacher network \mathcal{L}_t as:

$$\mathcal{L}_t = \mathcal{L}_f + \lambda_b \mathcal{L}_b + \lambda_c \mathcal{L}_c, \quad (10)$$

where $\lambda_b \in (0, 1)$ and $\lambda_c \in (0, 1)$ are hyper-parameters to control the weights of the factors.

4.4.2 Optimizing the Student Network. After the teacher model is trained to convergence, we can obtain intermediate entity distributions in the two reasoning processes of the teacher network. We

take the average of the two distributions as the supervision signal:

$$\mathbf{p}_t^{(k)} = \frac{1}{2}(\mathbf{p}_f^{(k)} + \mathbf{p}_b^{(n-k)}), k = 1, \dots, n-1 \quad (11)$$

As described before, we adopt the NSM model as the student network to conduct forward reasoning. Besides the reasoning loss, we also incorporate the loss between the predictions of the student network and the supervision signal of the teacher network:

$$\begin{aligned} \mathcal{L}_1 &= D_{KL}(\mathbf{p}_s^{(n)}, \mathbf{p}_f^*), \\ \mathcal{L}_2 &= \sum_{k=1}^{n-1} D_{KL}(\mathbf{p}_s^{(k)}, \mathbf{p}_t^{(k)}), \\ \mathcal{L}_s &= \mathcal{L}_1 + \lambda \mathcal{L}_2. \end{aligned} \quad (12)$$

where $\mathbf{p}_t^{(k)}$ and $\mathbf{p}_s^{(k)}$ denote the intermediate entity distributions at the k -th step from the teacher network and student network, respectively, and λ is a hyperparameter to tune.

In practice, labeled data for intermediate reasoning steps is seldom available. Most existing methods only rely on the final answer to learn the entire model, which may not be well trained or form spurious reasoning paths. Our approach adopts the teacher network for improving the student network. The main novelty is to utilize both forward and backward reasoning in producing more reliable intermediate entity distributions. Note that we do not incorporate any additional labeled data for training intermediate reasoning steps in the teacher network. Instead, we try to learn such intermediate entity distributions by enforcing the correspondence in the bidirectional reasoning process. To our knowledge, backward reasoning has been seldom considered in multi-hop KBQA task, especially its correspondence with forward reasoning. Such an idea is indeed related to recent progress in self-supervised learning [15], in which we leverage internal supervision signal to learn the model.

5 EXPERIMENT

In this section, we perform the evaluation experiments for our approach on the KBQA task.

5.1 Datasets

We adopt three benchmark datasets for the multi-hop KBQA task:

MetaQA [38] contains more than 400k single and multi-hop (up to 3-hop) questions in the domain of movie, containing three datasets, namely MetaQA-1hop, MetaQA-2hop and MetaQA-3hop.

WebQuestionsSP (webqsp) [36] contains 4737 natural language questions that are answerable using Freebase as the knowledge base. The questions require up to 2-hop reasoning from knowledge base. We use the same train/dev/test splits as GraftNet [31].

Complex WebQuestions 1.1 (CWQ) [32] is generated from WebQuestionsSP by extending the question entities or adding constraints to answers. There are four types of question: composition (45%), conjunction (45%), comparative (5%), and superlative (5%). The questions require up to 4-hops of reasoning on the KB. Following [30, 31], we use the topic entities labeled in original datasets and adopt PageRank-Nibble algorithm (PRN) [1] to find KB entities close to them. With these entities, we can obtain a relatively small subgraph that is likely to contain the answer entity. For CWQ and webqsp datasets, we first obtain the neighborhood graph within

Table 1: Statistics of all datasets. “#entity” denotes average number of entities in subgraph, and “coverage” denotes the ratio of at least one answer in subgraph.

Datasets	Train	Dev	Test	#entity	coverage
MetaQA-1hop	96,106	9,992	9,947	487.6	100%
MetaQA-2hop	118,980	14,872	14,872	469.8	100%
MetaQA-3hop	114,196	14,274	14,274	497.9	99.0%
webqsp	2,848	250	1,639	1,429.8	94.9%
CWQ	27,639	3,519	3,531	1,305.8	79.3%

two hops of topic entities and then run PRN algorithm on it. We further expand one hop for CVT entities in Freebase to obtain the neighborhood subgraph. As shown in Table 1, 2-hop graphs are sufficient to cover most of the answer entities. While on MetaQA datasets, we run PRN algorithm on the entire KB. Specifically, we use the PRN algorithm [1] with $\epsilon = 1e^{-6}$ and then select the m top-scoring entities. We set $m = 500$ for the smaller MetaQA KB and $m = 2000$ for larger Freebase. For the reserved triples, both their head and tail entities are obtained from the top m entities identified by PRN algorithm. We summarize the statistics of the three datasets in Table 1.

5.2 Experimental Setting

5.2.1 Evaluation Protocol. We follow [30, 31] to cast the multi-hop KBQA task as a ranking task for evaluation. For each test question in a dataset, a list of answers are returned by a model according to their predictive probabilities. We adopt two evaluation metrics widely used in previous works, namely *Hits@1* and *F1*. Specifically, *Hits@1* refers to whether the top answer is correct. For all the methods, we learn them using the training set, and optimize the parameters using the validation set and compare their performance on the test set.

5.2.2 Methods to Compare. We consider the following methods for performance comparison:

- **KV-Mem** [24] maintains a memory table for retrieval, which stores KB facts encoded into key-value pairs.
- **GraftNet** [31] adopts a variant of graph convolution network to perform multi-hop reasoning on heterogeneous graph.
- **PullNet** [30] utilizes the shortest path as supervision to train graph retrieval module and conduct multi-hop reasoning with GraftNet on the retrieved sub-graph.
- **SRN** [27] is a multi-hop reasoning model under the RL setting, which solves multi-hop question answering through extending inference paths on knowledge base.
- **EmbedKGQA** [29] conducts multi-hop reasoning through matching pretrained entity embeddings with question embedding obtained from RoBERTa [23].
- **NSM**, **NSM_{+p}** and **NSM_{+h}** are three variants of our model, which (1) do not use the teacher network, (2) use the teacher network with parallel reasoning, and (3) use the teacher network with hybrid reasoning, respectively.

5.2.3 Implementation Details. Before training the student network, we pre-train the teacher network on multi-hop KBQA task. To avoid overfitting, we adopt early-stopping by evaluating *Hits@1* on the validation set every 5 epochs. We optimize all models with Adam

optimizer, where the batch size is set to 40. The learning rate is tuned amongst $\{0.01, 0.005, 0.001, 0.0005, 0.0001\}$. The reasoning steps is set to 4 for CWQ dataset, while 3 for other datasets. The coefficient λ (in Eq. 12) and λ_b, λ_c (in Eq. 10) are tuned amongst $\{0.01, 0.05, 0.1, 0.5, 1.0\}$.

5.3 Results

The results of different methods for KBQA are presented in Table 2. It can be observed that:

(1) Among the baselines, KV-Mem performs the worst. This is probably because it does not explicitly consider the complex reasoning steps. Most methods perform very well on the MetaQA-1hop and MetaQA-2hop datasets, which require only up to 2 hops of reasoning. On the other hand, the other datasets seem to be more difficult, especially the webqsp and CWQ datasets. Overall, EmbedKGQA and PullNet are better than the other baselines. PullNet trains an effective subgraph retrieval module based on the shortest path between topic entities and answer entities. Such a module is specially useful to reduce the subgraph size and produce high-quality candidate entities.

(2) Our base model (*i.e.*, the single student network) NSM performs better than the competitive baselines in most cases. It is developed based on a graph neural network with two novel extensions for this task (Sec. 4.2). The gains of teacher-student framework show variance on different datasets. Specifically, on the two most difficult datasets, namely Webqsp and CWQ, the variants of NSM_{+p} and NSM_{+h} are substantially better than NSM and other baselines. These results have shown the effectiveness of the teacher network in our approach, which largely improves the student network. Different from SRN and PullNet, our approach designs a novel bidirectional reasoning mechanism to learn more reliable intermediate supervision signals. Comparing NSM_{+p} and NSM_{+h} , we find that their results are similar. On Webqsp and CWQ datasets, the hybrid reasoning is slightly better to improve the student network than parallel reasoning.

Table 2: Performance comparison of different methods for KBQA (Hits@1 in percent). We copy the results for KV-Mem, GraftNet and PullNet from [30], and copy the results for SRN and EmbedKGQA from [27, 29]. Bold and underline fonts denote the best and the second best methods.

Models	Webqsp	MetaQA-1	MetaQA-2	MetaQA-3	CWQ
KV-Mem	46.7	96.2	82.7	48.9	21.1
GraftNet	66.4	97.0	94.8	77.7	32.8
PullNet	68.1	97.0	99.9	91.4	45.9
SRN	-	97.0	95.1	75.2	-
EmbedKGQA	66.6	97.5	<u>98.8</u>	<u>94.8</u>	-
NSM	68.7	97.1	99.9	98.9	47.6
NSM_{+p}	<u>73.9</u>	<u>97.3</u>	99.9	98.9	<u>48.3</u>
NSM_{+h}	74.3	97.2	99.9	98.9	48.8


5.4 Detailed Performance Analysis

Table 2 has shown that our approach overall has a better performance. Next, we perform a series of detailed analysis experiments. For clarity, we only incorporate the results of NSM as the reference, since it performs generally well among all the baselines.

5.4.1 Ablation Study. Previous experiments have indicated that the major improvement is from the contribution of the teacher network. Here, we compare the effect of different implementations of the teacher network. The compared variants include: (1) NSM_{+f} using only the forward reasoning (unidirectional); (2) NSM_{+b} using only the backward reasoning (unidirectional); (3) NSM_{+p} using the parallel reasoning (bidirectional); (4) NSM_{+h} using the hybrid reasoning (bidirectional); (5) $NSM_{+p,-c}$ removing the correspondence loss (Eq. 9) from NSM_{+p} ; and (6) $NSM_{+h,-c}$ removing the correspondence loss (Eq. 9) from NSM_{+h} . In Table 3, we can see that unidirectional reasoning is consistently worse than bidirectional reasoning: the variants of NSM_{+f} and NSM_{+b} have a lower performance than the other variants. Such an observation verifies our assumption that bidirectional reasoning can improve the learning of intermediate supervision signals. Besides, by removing the correspondence loss from the teacher network, the performance substantially drops, which indicates that forward and backward reasoning can mutually enhance each other.

Table 3: Ablation study of the teacher network (in percent).

Models	Webqsp		CWQ	
	Hits	F1	Hits	F1
NSM	68.7	62.8	47.6	42.4
NSM_{+f}	70.7	64.7	47.2	41.5
NSM_{+b}	71.1	65.4	47.1	42.7
$NSM_{+p,-c}$	72.5	66.5	47.7	42.7
$NSM_{+h,-c}$	73.0	66.9	47.5	42.1
NSM_{+p}	73.9	66.2	48.3	44.0
NSM_{+h}	74.3	67.4	48.8	44.0



(a) Varying λ on webqsp dataset.

(b) Varying λ on CWQ dataset.

Figure 4: Performance tuning of our approach.

5.4.2 Parameter Tuning. In our approach, we have several combination coefficients to tune, including λ in Eq. 12, and λ_b and λ_c in Eq. 10. We first tune λ amongst $\{0.01, 0.05, 0.1, 0.5, 1.0\}$, which controls the influence of the teacher network on the student network. As shown in Fig. 4, hybrid reasoning seems to work well with small λ (*e.g.*, 0.05), while parallel reasoning works better with relatively large λ (*e.g.*, 1.0). Similarly, we can tune the parameters of λ_b and λ_c . Overall, we find that $\lambda_c = 0.01$ and $\lambda_b = 0.1$ are good choices for our approach. Another parameter to tune is the embedding dimension d (which is set to 100), and we do not observe significant improvement when $d > 100$. The reasoning steps n should

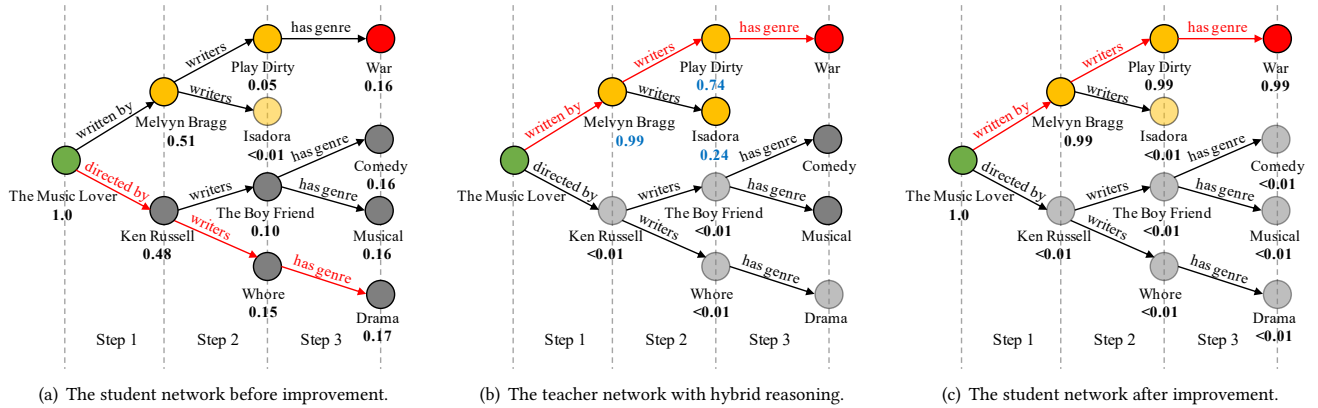


Figure 5: A case from the MetaQA-3hop dataset. We use green, red, yellow and grey circles to denote the topic entity, correct answer, intermediate entities and irrelevant entities respectively. The red colored edges denote the actual reasoning paths for different methods. The color darkness indicates the relevance degree of an entity by a method. For simplicity, we only visualize the entities with a probability equal to or above 0.01.

be adjusted for different datasets. We observe that our approach achieves the best performance on CWQ dataset with $n = 4$, while $n = 3$ for the other datasets with exhaustive search. Due to space limit, we omit these tuning results.

5.4.3 Evaluating Intermediate Entities. A major assumption we made is that our teacher network can obtain more reliable intermediate entities than the student network. Here, we compare the performance of the two networks in finding intermediate entities. Since the MetaQA-3hop dataset is created using pre-defined templates, we can recover the ground-truth entities at intermediate hops. We consider it a retrieval task and adopt the standard *Precision*, *Recall* and *F1* as evaluation metrics. From Table 4, we can see that the teacher network is much better than the student network in finding intermediate entities, but has slightly worse performance at the second hop. Note that the results of the third hop have been omitted, since it is the last hop. Since the student network only utilizes forward reasoning, the results of the first hop are more important than those of subsequent hops. These results also explain why our teacher-student approach is better than the single student model.

Table 4: Performance comparison w.r.t. different hops on MetaQA-3hop dataset (in percent).

Models	Hop 1			Hop 2		
	Pre	Rec	F1	Pre	Rec	F1
Student	61.0	60.6	60.4	99.9	70.2	80.8
Teacher _{+p}	80.0	59.0	66.3	95.0	68.9	78.8
Teacher _{+h}	99.9	56.0	70.9	99.7	63.0	75.4

5.4.4 One-Shot Evaluation. In Table 2, we have found that the improvement of our approach over the basic NSM model is very small on the MetaQA datasets. We suspect that this is because the amount of training data for MetaQA is more than sufficient: 100K training cases for no more than 300 templates in each dataset. To examine this, we randomly sample a single training case for every

question template from the original training set, which forms a one-shot training dataset. We evaluate the performance of our approach trained with this new training dataset. The results are shown in Table 5. As we can see, our approach still works very well, and the improvement over the basic NSM becomes more substantial.

Table 5: Results under one-shot setting (in percent).

Models	MetaQA-1		MetaQA-2		MetaQA-3	
	Hits	F1	Hits	F1	Hits	F1
NSM	93.3	92.6	97.7	96.0	90.6	74.5
NSM _{+p}	94.3	93.9	98.7	96.4	97.0	79.8
NSM _{+h}	93.9	93.7	98.4	95.8	95.6	81.6

5.5 Case Study

The major novelty of our approach lies in the teacher network. Next, we present a case study for illustrating how it helps the student network.

Given the question “*what types are the movies written by the screenwriter of the music lovers*”, the correct reasoning path is “*The Music Lovers*” (movie) $\rightarrow_{\text{written by}}$ “*Melvyn Bragg*” (screenwriter) $\rightarrow_{\text{write}}$ “*Play Dirty*” (movie) $\rightarrow_{\text{has genre}}$ “*War*” (genre). Note that “*Isadora*” is also qualified at the second step. However, its genre is missing in the KB. Fig. 5 presents a comparison between the learned results of the student before improvement (*i.e.*, without the teacher network), the teacher network and the student network after improvement.

As shown in Fig. 5(a), the original student network has selected a wrong path leading to an irrelevant entity. At the first hop, NSM mainly focuses on the two entities “*Ken Russell*” and “*Melvyn Bragg*” with probabilities of 0.48 and 0.51 respectively. Since it mistakenly includes “*Ken Russell*” (director of “*The Music Lovers*”) at the first reasoning step, it finally ranks “*Drama*” as the top entity and chooses an irrelevant entity as the answer. In comparison, the teacher network (Fig. 5(b)) is able to combine forward and backward reasoning to enhance the intermediate entity distributions. As we can see,

our teacher assigns a very high probability of 0.99 to the entity “Melvyn Bragg” at the first step. When the supervision signals of the teacher are incorporated into the student, it correctly finds the answer entity “War” with a high probability of 0.99 (Fig. 5(c)).

This example has shown that our teacher network indeed provides very useful supervision signals at intermediate steps to improve the student network.

6 CONCLUSION

In this paper, we developed an elaborate approach based on teacher-student framework for the multi-hop KBQA task. In our approach, the student network implemented by a generic neural state machine focuses on the task itself, while the teacher network aims to learn intermediate supervision signals to improve the student network. For the teacher network, we utilized the correspondence between state information from a forward and a backward reasoning process to enhance the learning of intermediate entity distributions. We further designed two reasoning architectures that support the integration between forward and backward reasoning. We conducted evaluation experiments with three benchmark datasets. The results show that our proposed model is superior to previous methods in terms of effectiveness for the multi-hop KBQA task.

Currently, we adopt the NSM model as the student network. It is flexible to extend our approach to other neural architectures or learning strategies on graphs. In the future, we will also consider enhancing the entity embeddings using KB embedding methods, and obtain better intermediate supervision signals.

ACKNOWLEDGEMENT

We thank Kun Zhou and Junyi Li for the helpful discussions. This work is partially supported by the National Research Foundation, Singapore under its International Research Centres in Singapore Funding Initiative, the National Natural Science Foundation of China under Grant No. 61872369 and 61832017, Beijing Academy of Artificial Intelligence (BAAI), and Beijing Outstanding Young Scientist Program under Grant No. BJJWZYJH012019100020098. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of National Research Foundation, Singapore.

REFERENCES

- [1] Reid Andersen, Fan R. K. Chung, and Kevin J. Lang. 2006. Local Graph Partitioning using PageRank Vectors. In *FOCS 2006*.
- [2] Jonathan Berant, Andrew Chou, Roy Frostig, and Percy Liang. 2013. Semantic Parsing on Freebase from Question-Answer Pairs. In *EMNLP 2013*. 1533–1544.
- [3] Rajarshi Das, Shehzaad Dhuliawala, Manzil Zaheer, Luke Vilnis, Ishan Durugkar, Akshay Krishnamurthy, Alex Smola, and Andrew McCallum. 2018. Go for a Walk and Arrive at the Answer: Reasoning Over Paths in Knowledge Bases using Reinforcement Learning. In *ICLR 2018*.
- [4] Tuong Do, Huy Tran, Thanh-Toan Do, Erman Tjiputra, and Quang D. Tran. 2019. Compact Trilinear Interaction for Visual Question Answering. In *ICCV 2019*.
- [5] Li Dong, Furu Wei, Ming Zhou, and Ke Xu. 2015. Question answering over freebase with multi-column convolutional neural networks. In *ACL 2015*.
- [6] Dheeru Dua, Sameer Singh, and Matt Gardner. 2020. Benefits of Intermediate Annotations in Reading Comprehension. In *ACL 2020*. 5627–5634.
- [7] David A. Ferrucci, Eric W. Brown, Jennifer Chu-Carroll, James Fan, David Gondek, Aditya Kalyanpur, Adam Lally, J. William Murdock, Eric Nyberg, John M. Prager, Nico Schlaefer, and Christopher A. Welty. 2010. Building Watson: An Overview of the DeepQA Project. *AI Magazine* 31, 3 (2010), 59–79.
- [8] Bent Fuglede and Flemming Topsøe. 2004. Jensen-Shannon divergence and Hilbert space embedding. In *ISIT 2004*. 31.
- [9] Tommaso Furlanello, Zachary Chase Lipton, Michael Tschannen, Laurent Itti, and Anima Anandkumar. 2018. Born-Again Neural Networks. In *ICML 2018*.
- [10] Geoffrey E. Hinton, Oriol Vinyals, and Jeffrey Dean. 2015. Distilling the Knowledge in a Neural Network. *CoRR abs/1503.02531* (2015).
- [11] Minghao Hu, Yuxing Peng, Furu Wei, Zhen Huang, Dongsheng Li, Nan Yang, and Ming Zhou. 2018. Attention-Guided Answer Distillation for Machine Reading Comprehension. In *EMNLP 2018*.
- [12] Ronghang Hu, Anna Rohrbach, Trevor Darrell, and Kate Saenko. 2019. Language-Conditioned Graph Networks for Relational Reasoning. In *ICCV 2019*.
- [13] Drew A. Hudson and Christopher D. Manning. 2018. Compositional Attention Networks for Machine Reasoning. In *ICLR 2018*.
- [14] Drew A. Hudson and Christopher D. Manning. 2019. Learning by Abstraction: The Neural State Machine. In *NeurIPS 2019*. 5901–5914.
- [15] Longlong Jing and Yingli Tian. 2020. Self-supervised Visual Feature Learning with Deep Neural Networks: A Survey. *TPAMI 2020* (2020).
- [16] Thomas N. Kipf and Max Welling. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *ICLR 2017*.
- [17] Richard E. Korf. 1999. A Divide and Conquer Bidirectional Search: First Results. In *IJCAI 99*. 1184–1191.
- [18] Solomon Kullback and Richard A. Leibler. 1951. On information and sufficiency. *The annals of mathematical statistics* 22, 1 (1951), 79–86.
- [19] Yunshi Lan and Jing Jiang. 2020. Query Graph Generation for Answering Multi-hop Complex Questions from Knowledge Bases. In *ACL 2020*. 969–974.
- [20] Yunshi Lan, Shuohang Wang, and Jing Jiang. 2019. Multi-hop Knowledge Base Question Answering with an Iterative Sequence Matching Model. In *ICDM 2019*.
- [21] Chen Liang, Jonathan Berant, Quoc V. Le, Kenneth D. Forbus, and Ni Lao. 2017. Neural Symbolic Machines: Learning Semantic Parsers on Freebase with Weak Supervision. In *ACL 2017*. 23–33.
- [22] Xi Victoria Lin, Richard Socher, and Caiming Xiong. 2018. Multi-Hop Knowledge Graph Reasoning with Reward Shaping. In *EMNLP 2018*. 3243–3253.
- [23] Yinhan Liu, Mylène Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. RoBERTa: A Robustly Optimized BERT Pretraining Approach. *CoRR abs/1907.11692* (2019).
- [24] Alexander H. Miller, Adam Fisch, Jesse Dodge, Amir-Hossein Karimi, Antoine Bordes, and Jason Weston. 2016. Key-Value Memory Networks for Directly Reading Documents. In *EMNLP 2016*.
- [25] Sewon Min, Victor Zhong, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2019. Multi-hop Reading Comprehension through Question Decomposition and Rescoring. In *ACL 2019*.
- [26] Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. Glove: Global Vectors for Word Representation. In *EMNLP 2014*. 1532–1543.
- [27] Yunqi Qiu, Yuanzhuo Wang, Xiaolong Jin, and Kun Zhang. 2020. Stepwise Reasoning for Multi-Relation Question Answering over Knowledge Graph with Weak Supervision. In *WSDM 2020*. 474–482.
- [28] Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. 2015. FitNets: Hints for Thin Deep Nets. In *ICLR 2015*.
- [29] Apoorv Saxena, Aditya Tripathi, and Partha P. Talukdar. 2020. Improving Multi-hop Question Answering over Knowledge Graphs using Knowledge Base Embeddings. In *ACL 2020*. 4498–4507.
- [30] Haitian Sun, Tania Bedrax-Weiss, and William W. Cohen. 2019. PullNet: Open Domain Question Answering with Iterative Retrieval on Knowledge Bases and Text. In *EMNLP 2019*.
- [31] Haitian Sun, Bhuvan Dhingra, Manzil Zaheer, Kathryn Mazaitis, Ruslan Salakhutdinov, and William W. Cohen. 2018. Open Domain Question Answering Using Early Fusion of Knowledge Bases and Text. In *EMNLP 2018*. 4231–4242.
- [32] Alon Talmor and Jonathan Berant. 2018. The Web as a Knowledge-Base for Answering Complex Questions. In *NAACL-HLT 2018*. 641–651.
- [33] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. Graph Attention Networks. *ICLR 2018* (2018).
- [34] Kun Xu, Siva Reddy, Yansong Feng, Songfang Huang, and Dongyan Zhao. 2016. Question Answering on Freebase via Relation Extraction and Textual Evidence. In *ACL 2016*.
- [35] Ze Yang, Linjun Shou, Ming Gong, Wutao Lin, and Daxin Jiang. 2020. Model Compression with Two-stage Multi-teacher Knowledge Distillation for Web Question Answering System. In *WSDM 2020*. 690–698.
- [36] Wen-tau Yih, Ming-Wei Chang, Xiaodong He, and Jianfeng Gao. 2015. Semantic Parsing via Staged Query Graph Generation: Question Answering with Knowledge Base. In *ACL 2015*.
- [37] Wen-tau Yih, Matthew Richardson, Christopher Meek, Ming-Wei Chang, and Jina Suh. 2016. The Value of Semantic Parse Labeling for Knowledge Base Question Answering. In *ACL 2016*.
- [38] Yuyu Zhang, Hanjun Dai, Zornitsa Kozareva, Alexander J. Smola, and Le Song. 2018. Variational Reasoning for Question Answering with Knowledge Graph. In *AAAI 2018*.
- [39] Ying Zhang, Tao Xiang, Timothy M. Hospedales, and Huchuan Lu. 2018. Deep Mutual Learning. In *CVPR 2018*. 4320–4328.