Interactive Factual Knowledge Learning in Dialogues

Anonymous Author(s)
Affiliation
Address
email

Abstract

Knowledge bases (KBs) used in applications need to be continuously expanded in order to serve the users well. This process is known as *knowledge base completion* (KBC). A piece of *knowledge* or a *fact* is often represented as a triple (s, r, t), meaning the entity s and the entity t have the relation r or are linked by r. KBC builds a model to infer missing facts from existing ones in a given KB. Existing KBC research typically makes the *closed-world assumption* that to infer a new fact (s, r, t), it assumes that s, r and t are already in the KB, but are not linked. Clearly, this assumption is a serious limitation. In this paper, we eliminate this assumption and allow s, r and/or t to be unknown to the KB, which we call *open-world knowledge base completion* (OKBC). We focus on solving OKBC via user interactions, which enables the proposed system to potentially serve as an engine for learning new knowledge during dialogues. Experimental results will show the effectiveness of the proposed approach.

1 Introduction

Knowledge bases (KBs) are instrumental to many AI applications, e.g., question answering [5], dialogues [17, 55, 27, 59]. Although many large KBs exist like Freebase [6] and WordNet [33], they miss a large percentage of facts about common entities [50], which is also true for numerous application specific KBs. In recent years, researchers have studied ways of expanding KBs automatically. This area is called *knowledge base completion* (KBC) [23, 15]. This paper studies this problem further, but unlike most existing studies, we study the problem in the *open-world setting* (to be defined shortly).

The type of knowledge that we focus on is the fact triples commonly used in *knowledge graphs*, (s, r, t), which means that the *source entity* s and the *target entity* t can be linked by the *relation* r. For example, (*Obama*, *CitizenOf*, *USA*), meaning that *Obama is a citizen of USA*. Traditional KBC infers new facts (s, r, t) from existing ones in the KB, which is often defined as a binary classification problem: Given a (query) triple (s, r?, t), it predicts whether s and t can be linked by r. We use ? to indicate a query. Existing work typically makes the *closed-world assumption* that s, r and t are all *known* to exist in the KB [23, 11, 9, 36]. This paper removes this assumption and allows s, r and/or t not to exist in the KB. We call this *open-world KBC* (OKBC). The main existing approach to solving OKBC is to extract facts from documents and add them to the KB [34, 20, 3, 52, 13]. Two recent works are [43, 42]. This paper proposes a different but complementary approach that involves user interactions and inference based on knowledge gained from the interactions.

The motivation for proposing an interactive approach to solving OKBC came from the need of *chatbots* like SIRI, Amazon Alexa, Google Assistant, not to mention numerous customer support bots used by large companies. Chatbots often use KBs to answer user questions. Many deep learning methods now use KBs too [17, 24, 55, 27, 59]. However, the incompleteness [50] of the KBs limits the scope of their applications. It is thus highly desirable for a chatbot to learn new knowledge in the dialogue process to expand its KB to make it more and more knowledgeable.

Submitted to 34th Conference on Neural Information Processing Systems (NeurIPS 2020). Do not distribute.

USER:	(*Obama*, *CitizenOf?*, *USA*) "Is Obama a citizen of USA?"	[**Query**]
IKAI:	I do not know what "**CitizenOf**" means? Can you provide me an example?	[***Ask for Clue***]
USER:	(*David Cameron*, *CitizenOf*, *UK*). "David Cameron is a citizen of UK."	[**SF1**]
IKAI:	Got it. Can you tell me how "**Obama**" and "**Honolulu**" are related?	[***CLQ***]
USER:	(*Obama*, *BornIn*, *Honolulu*). "Obama was born in Honolulu."	[**SF2**]
IKAI:	Got it. Can you tell also me how "**Honolulu**" and "**Hawaii**" are related?	[***MLQ***]
USER:	(*Honolulu*, *CapitalOfState*, *Hawaii*). "Honolulu is the state capital of Hawaii."	[**SF3**]
IKAI:	(*Obama*, *CitizenOf?*, *USA*) holds. "Yes, Obama is a US citizen."	[**Answer**]

Figure 1: An example of interactive learning and inference. Each triple above is assumed to be extracted from the sentence after it. ***Ask for Clue***, ***CLQ*** and ***MLQ*** are interaction query types, discussed in Sec 3.

There are many opportunities to learn during a dialogue. For example, we can extract knowledge from user utterances using information extraction methods [18]. We do not study it in this paper. This paper focuses on another opportunity, i.e., when the user asks the chatbot a *factual* verification (yes/no) question that the chatbot cannot answer. To learn from the user in this new setting is challenging because it not only involves extraction but also asking user questions and making inference.

KBC forms a suitable model for this dialogue-based knowledge learning. The KBC model is insufficient as in a conversation the user may ask or say anything, which may include entities and/or relations that are not in the KB. This paper proposes a novel interactive method to solve OKBC, called *Interactive Knowledge Acquisition and Inference* (IKAI), which involves two tasks: **(1)** ***Interactive acquisition of supporting facts***, formulating an inference strategy to ask the user suitable questions to convert an OKBC query to a KBC query. Those user answers, which are *supporting facts*, are also added to the KB. **(2)** ***Knowledge inference***, building a predictive model using the supporting facts and the knowledge in the KB to infer if the converted KBC query is true to answers the original user query. Fig. 1 shows an example. IKAI acquires *supporting facts* SF1, SF2, SF3 to accomplish the *interactive knowledge acquisition* and then utilizes these pieces of knowledge along with existing KB facts to *infer* over the query relation "*CitizenOf*" to answer the user query. The new knowledge (which are supporting facts) is also added to the KB for future use.

Note that, this paper only focuses on developing the core learning engine. It does not study fact or relation extraction from user utterances, entity linking, etc., which have been studied before. Thus, IKAI only works with structured queries (s, r?, t) (? indicates a query) and assumes that a relation extraction system (e.g., OpenIE [2]) is available to convert natural language queries or statements from the user into query triples (s, r?, t) or fact triples (s, r, t). We also assume that the user has good intentions, i.e., user answers are 100% correct, but we do not assume that the user can answer all questions as opposed to the teacher-student setup, where the teacher is assumed to know everything.

IKAI is implemented with a *Finite State Machine* (FSM) and a *Predictor*. The FSM dynamically formulates an inference strategy for a given OKBC query. While executing the strategy, IKAI processes the query and acquires the supporting facts by asking the user and thereby, transforms an OKBC problem to a KBC one (i.e., all s, r and t become *known* to the KB). The *Predictor* is invoked for knowledge inference to solve the *transformed* KBC problem. We evaluated IKAI using two real-world KBs: *Freebase* and *WordNet*. Experimental results demonstrate its effectiveness.

2 Related Work

KBC has been studied extensively [35, 12, 11, 8, 15]. The work in [35] partially deals with OKBC, but it cannot handle unknown entities. The work in [43, 42] solves OKBC using an external text corpus to perform inference on unknown entities. They cannot handle unknown relations. They or other existing methods do not perform KBC or OKBC through user interactions like IKAI.

IKAI is also related to interactive language learning (ILL) [49, 48], but ILL does not grow KB. Among dialogue systems [46, 25, 26, 10, 51, 56], the work in [26] allows the learner to ask questions. However, it learns only about *whether to ask*. The *what to ask* aspect is manually designed. More importantly, its learning is for building a dialogue system in the *Teacher-Student setup*, not for knowledge learning during conversations with users like IKAI. IKAI also formulates query-specific inference strategies. No existing dialogue systems [46, 25, 10, 51, 41, 56, 58, 28, 53] can train prediction models continually online using information retained in the past.

The most closely related works to IKAI are those in [30, 18]. [30] proposed a method to acquire knowledge when a WH-question is asked by the user. Its method for dealing with WH-questions by

trying everything in the KB is entirely different from dealing with yes/no questions in IKAI. [18] is based on extraction of relations directly from user utterances. IKAI is based on answering user queries and it formulates a strategy to ask and also does inference for knowledge acquisition.

Recently, many conversational systems have been proposed, e.g., Conversational Recommender Systems [44, 1, 57] and Information-seeking Conversation [54, 37]. However, they do not acquire new knowledge to deal with OKBC in conversation. There are also works on KB question-answering (KBQA) [21, 7, 4, 39, 38], open-vocabulary semantic parsing [14, 45], information extraction [40, 2], and KB updating using facts extracted from the Web [34]. We do not do KBQA, semantic parsing or Web fact extraction, which is complementary to our work. Rather, we focus on solving OKBC interactively, i.e., learning through interactions with the user.

3 IKAI: The Proposed System

Overview. As discussed in Sec 1, IKAI performs following two main tasks:

Task 1 : Converting OKBC to KBC. Given an OKBC query (s, r?, t) [e.g., (*Obama*, *CitizenOf*?, *USA*) - *Is Obama is citizen of USA?*], IKAI formulates some relevant questions to ask the user (see Fig. 1) to acquire supporting facts (SFs), which convert an OKBC query (s, r?, t) to a KBC query. The KBC query is also (s, r?, t) except that all components, s, r, and t, must already exist in the current KB. OKBC allows any of s, r and t to be *unknown* or not in the KB.

The conversion involves the following two cases- **(1) If r is unknown**, IKAI asks the user to provide a ***clue*** [an example triple r]. In Fig. 1, IKAI asks the user to provide an example for the *unknown* relation "*CitizenOf*" and the user gave the clue (*David Cameron*, *CitizenOf*, *UK*), meaning *David Cameron is a citizen of UK*. With this clue added to the KB, the relation "*CitizenOf*" becomes known. If "*David Cameron*" or "*UK*" is not in the KB, the clue may need to be further processed, i.e., using the case below. **(2) If s or t is unknown**, IKAI asks the user to provide a ***link*** (relation) to connect the unknown entity s or t with an automatically selected existing entity (discussed in "Actions" in Sec. 4) in the KB. In Fig. 1, IKAI asks the user to provide a relation to connect the *unknown* entity "*Obama*" with an existing KB entity "*Honolulu*" [Can you tell me how "*Obama*" and "*Honolulu*" are related?]. The acquired knowledge (e.g., the two user answers) thus reduces OKBC to KBC as all r, s and t are now in the KB. We deal with the case when the user is unable to answer later.

Task 2: Solving KBC by Building a Predictor. With the resulting KBC query, $(s, r?, t)$, IKAI uses a *path-ranking* (PR) approach [23, 22, 16, 35] to build the predictive model (Predictor) to predict if $(s, r?, t)$ is true. We choose the PR approach due to its better performance [47]. PR algorithms enumerate relation paths between two entities (s, t) in a KB (encoded as a multi-relation graph) and use those relation paths as features to train the predictive model for relation r. Given a sequence of entities and relations $\langle s, r_1, e_1, r_2, e_2.., r_l, t\rangle$ linking s to t in the KB, we use the term ***path feature*** to refer to the *relation path* $\langle r_1, r_2, .., r_l\rangle$ and the full sequence $\langle s, r_1, e_1, r_2, e_2.., r_l, t\rangle$ as simply ***path*** onwards. An entity e_i appearing in the path *excluding the source and target entities* is termed as an ***intermediate*** entity for the path. *A path feature is extracted from a path by just excluding the entities in the sequence.* In this work, we adopt the PR method, C-PR [29], which leverages the context relatedness (learned from text documents using neural word embedding [31, 32]) of the source s and target t and performs bidirectional random walk to enumerate path features to make more accurate prediction. More details can be found in Supplementary (Section E).

Note, in bidirectional path exploration, C-PR may get stuck with e_f (e_b) being the last forward (backward) node without being linked to form a full path. Unlike C-PR, in IKAI, we connect the forward and backward path sequence by linking e_f and e_b with a template (fake) relation "-?-" to create a full path (called a *incomplete path*). We call "-?-" the missing link marker. For an incomplete path, IKAI asks the user to provide a relation/link to connect (e_f, e_b) for path completion. We refer to such a query as a missing link query (MLQ) (see Sec. 3.1 for more details).

IKAI manages the aforementioned two tasks and also data collection for building the predictive model for task 2 (discussed in Sec 3.2) using a Finite State Machine (FSM) as discussed below.

3.1 Finite State Machine of IKAI

IKAI manages everything using a Finite State Machine (FSM). It formulates and simultaneously executes an *inference strategy*. That is, at each state S_i, FSM chooses an action a_j to execute. After an action is executed, FSM goes to a new state S_{i+1}. When FSM transits into its final state, the

Predictor is invoked to train to solve the reduced KBC problem to answer the original user query. Table 1 gives a description of the parameters of IKAI referred on-wards.

IKAI's FSM is $(\mathcal{S}, \mathcal{A}, \mathcal{S}_0, \mathcal{S}_F, \Delta)$, where $\mathcal{S}$ is the finite set of states, $\mathcal{A}$ the set of actions, $\mathcal{S}_0$ the set of initial states, $\mathcal{S}_F$ the set of terminal states and $\Delta : (\mathcal{S}, \mathcal{A}) \rightarrow \mathcal{S}$ the state transition function. We denote a triple (s, r, t) under processing as d on-wards. d can be the OKBC query (denoted as q) from the user or a clue triple (denoted as c^+) given by the user for an unknown relation (e.g., *CitizenOf* in Fig. 1), because one or both of the entities [e.g., "*David Cameron*" and/or "*UK*" in the clue (*David Cameron*, *CitizenOf*, *UK*)] may be unknown. Thus, IKAI also process clues just like OKBC queries.

Table 1: Parameters of IKAI.

Para.	Description
δ_p	Minimum number of complete paths in extracted feature set *per query* as a requirement for the feature set to be qualified as complete and used for training/inference.
δ_{IL}	Interaction limit, i.e., maximum number of questions that IKAI is allowed to ask the user per query.
η_p	Maximum path length for C-PR
η_w	Number of random walks per query for C-PR
l, h	Low and high contextual similarity threshold respectively
ρ	% of learned tasks selected for additional clue acquisition.

Table 2: State bits (SBs) and their meanings.

SB	Name	Description
QERS	Query entities and relation searched	Whether the query source (s) and target (t) entities and query relation (r) have been searched in KB.
SEF	Source Entity Found	Whether the source entity (s) has been found in KB.
TEF	Target Entity Found	Whether the target entity (t) has been found in KB.
QRF	Query Relation Found	Whether the query relation (r) has been found in KB.
CLUE	Clue bit set	Whether the triple (to be processed) is a clue from user.
ILO	Interaction Limit Over	Whether the interaction limit is over for the query.
PFE	Path Feature Extracted	Whether path feature extraction has been done.
CPF	Complete Path Found	Whether the extracted path features are complete.
INFI	Inference Invoked	Whether inference module has been invoked.

Whenever IKAI receives a clue from the user, it postpones the processing of the OKBC query and initiates the strategy formulation for each clue triple consecutively. Once all clue triples are processed (i.e., strategies are formulated and executed), the unknown query relation is now in the KB and IKAI can then infer the OKBC query. Thus, the strategy formulation for OKBC query resumes. A *Processing Stack* (PS) is used to manage the process, which basically holds the triples to be processed by IKAI along with FSM state information. Whenever IKAI receives d (a query or a clue) from the user, it pushes d and the initial state (S_0) corresponding to d [see "States" below] onto PS as a pair (d, S_0). FSM processes the stack top $PS[top]$ and $PS[top]$ gets modified on a state transition due to execution of some action a_j. After the processing of d finishes, i.e., $PS[top] = (d, S_F)$ with S_F being a terminal state, the pair (d, S_F) is popped out of PS and processing of new stack top resumes. Since clues being received after the query (q), they always become stack top and get processed first, postponing the processing of q.

States ($\mathcal{S}$): A state $S \in \mathcal{S}$ is defined by 9 binary variables (Table 2), each of which keeps track of the results of an action $a \in \mathcal{A}$ taken by IKAI and thus, records the progress of the inference. For example, $S_0 \in \mathcal{S}_0$ is the initial state with all state bits set as 0. If d is a clue, the CLUE bit of S_0 is set as 1.

Actions ($\mathcal{A}$): There are six actions in IKAI: $a_0, ..., a_5$. They function to convert an OKBC query to a KBC query and enumerate path features for training the Predictor. We discuss the actions one by one below. Fig. 2 shows the state transitions and actions executed at various FSM states for processing the example conversation in Fig. 1 along with the changes in PS. In Fig. 2, the state bits are shown next to each state ID S_i and the content of PS is shown below S_i.

• a_0 **- Search query entities and relation.** This action searches whether s, r, and t of d are already in the KB and then sets appropriate bits of the current state (Table 2) and the FSM moves to a new state. For example, in Fig. 2, IKAI executes a_0 at S_0 [at $PS[top] = (q, S_0)$] and detects that entity "*Obama*" and relation "*CitizenOf*" are *unknown* and the FSM transits to S_1 [$PS[top] = (q, S_1)$]. If d is a clue and $s, t \in$ KB, a_0 also updates the KB with d and its inverse triple (t, r^{-1}, s). E.g., (*UK*, *CitizenOf*$^{-1}$, *David Cameron*) is the inverse of triple (*David Cameron*, *CitizenOf*, *UK*). Adding inverse triple helps path enumeration to collect more path features [29].

• a_1 **- Ask for a clue.** This action asks the user to provide a clue (+ve example) c^+ for an unknown relation r. For example, in Fig. 1, a_1 is executed to acquire clue (SF1) for unknown relation "*CitizenOf*" from the user, who answers (*David Cameron*, *CitizenOf*, *UK*). For learning the predictive model for unseen query relation r, a_1 also corrupts s and t of the clue triple c^+ once at a time to generate -ve clue triple(s) c^- by randomly sampling a node from the KB and pairing it with either s or t, e.g., (*David Cameron*, *CitizenOf*, *Russia*) and (*Chicago*, *CitizenOf*, *UK*).

Considering the Fig. 1 example and state transitions in Fig. 2, after executing a_1 for query q, IKAI receives the +ve clue triple c^+ (*David Cameron*, *CitizenOf*, *UK*) and pushes the generated two -ve clues (c^-, S_0) pairs and then, (c^+, S_0) into PS. Now, c^+ being $PS[top]$, strategy formulation for q

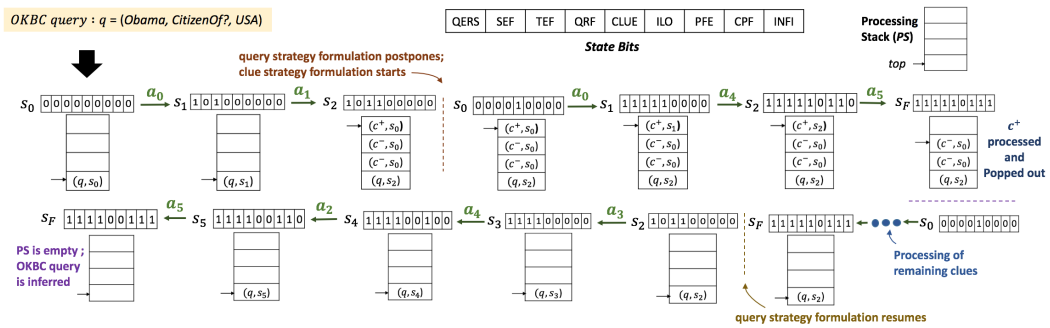

Figure 2: IKAI state transitions for the example in Fig. 1 [q is OKBC query and c^+ (c^-) is +ve (-ve) clue]

Table 3: Transition Conditions and corresponding Actions and their descriptions.

State Transition Conditions (for current state bits S_j [.])	Action Id : Operation
$QERS = 0$	a_0 : Search source (s), target (t) entities and query relation (r) in KB.
$ILO = 0 \wedge CLUE = 0 \wedge QERS = 1 \wedge QRF = 0$	a_1 : Ask user to provide an clue/example for query relation r.
$PFE = 1 \wedge ILO = 0 \wedge CPF = 0$	a_2 : Ask user to provide a missing link for path feature completion.
$QERS = 1 \wedge (SEF = 0 \vee TEF = 0) \wedge ILO = 0$	a_3 : Ask user to provide a connecting link to add a new entity to KB.
$QERS = 1 \wedge PFE = 0 \wedge SEF = 1 \wedge TFE = 1$	a_4 : Extract path features between source (s) and target (t) entities using C-PR.
$QRF = 1 \wedge CPF = 1$	a_5 : Invoke prediction model for inference.

postpones and that for c^+ begins. The strategy formulation for query q again resumes from "$\langle a_0, a_1 \rangle$" [at $PS[top]$ =(q, S_2)] when all clues finish their processing and consecutively popped out of the stack. Considering Fig. 1, the formulated strategy for each clue triple here is "$\langle a_0, a_4, a_5 \rangle$" (assuming both *David Cameron* and *UK* are already in the KB).

• a_2 **- Ask missing link query (MLQ).** This action formulates a *missing link query* (MLQ) to ask the user. That is, IKAI selects an incomplete path $\widetilde{p}$ (Sec. 3, Overview) from the extracted path set (by action a_4) with the highest context similarity (computed using the same embedding based method as used in C-PR, discussed in Sec. Sec. 3, Overview and Supplementary) $sim(e_f, e_b)$, where e_f and e_b are the two entities directly connected by -?- (missing link marker). A high contextual similarity indicates that e_f and e_b are more likely to possess a relation [29] and so, is a good candidate for MLQ. Note we don't want to ask the user too many questions by processing all incomplete paths.

For the Fig. 1 example, let IKAI choose an incomplete path "$Obama - BornIn \rightarrow Honolulu - ? - Hawaii - StateOf \rightarrow USA$" from the extracted path set (see action a_4) and with missing link between entity pair ($Honolulu$, $Hawaii$). Then, this missing link is used to formulate a MLQ: "*Can you also tell me how "Honolulu" and "Hawaii" are related?*". On receiving the answer from the user, i.e., (*Honolulu*, *CapitalOfState*, *Hawaii*) - "*Honolulu is the state capital of Hawaii*", IKAI fills the missing link in the path as "$Obama - BornIn \rightarrow Honolulu - CapitalOfState \rightarrow Hawaii - StateOf \rightarrow USA$". This complete path is then used to extract path feature as "$BornIn \rightarrow CapitalOfState \rightarrow StateOf$" to be used by the prediction model for inference. In Fig. 2, a_2 for query q is executed at state S_4 [i.e., when $PS[top] = (q, S_4)$]].

If the user does not respond to MLQ, a *guessing mechanism* is employed. Since entity pairs with high contextual similarity are likely to possess a relation, IKAI divides the similarity range [-1, 1] into three segments, using a low (l) and a high (h) similarity thresholds to replace the missing link in $\widetilde{p}$ with r_g to make it "complete": If $h \geq sim(s', t') \geq l$, r_g="-LooselyRelatedTo-"; else if $sim(s', t') \leq l$, r_g ="-NotRelatedTo-"; Otherwise, r_g="-RelatedTo-".

• a_3 **- Ask connecting link query (CLQ).** This action asks the user to connect unknown entities s and/or t with the KB by selecting the most contextually relevant node/entity (wrt s or t) from the KB and asking the user to provide a relation/link for s (or r) and the node. The contextual similarity of two entities is computed based on the same idea as in [29] using word embedding [32]In Fig. 1, IKAI asks CLQ for "*Obama*" with the chosen contextually relevant entity "*Honolulu*" and acquires SF2. In Fig. 2, a_3 for query q is executed at state S_2 [i.e., at $PS[top] = (q, S_2)$]] when the strategy formulation of q resumes after all clue triples finish processing.

• a_4 **- Enumerate path features.** This action extracts path feature set P between (s, t) involving both *complete* and *incomplete* paths for training the Predictor. IKAI always trains the Predictor with complete path features $P_c \subseteq P$. The incomplete paths (see Sec. 3, Overview) corresponding to the path features in $(P - P_c)$ are used as candidates for formulating and asking MLQs (see a_2) until $|P_c| = \delta_p$ (see Table 1) or δ_{IL} decrements to 0. δ_{IL} denotes the maximum number of questions that

IKAI can ask the user per OKBC query and is set to limit the number of interactions with user. In Fig. 1, a_4 for query q gets executed after a_0, a_1, a_3 [i.e., at $PS[top] = (q, S_3)$ in Fig. 2].

• a_5 **- Invoke Predictor.** This action first checks whether the OKBC query has been reduced to a KBC one. The Predictor can only work, *iff* (1) r, s and t exist in the KB (OKBC is reduced to KBC), and (2) path features (relation sequences corresponding to each path in P_c) are extracted between s and t for building a predictor. If the conditions are met, IKAI adds d (the KBC query triple or the clues) along with P_c in a data buffer $\mathcal{D}$ and invokes the Predictor for either training (on clue triples) or inference/prediction (on the reduced KBC query) by executing a_5. Note that, as clues involving query relation r get processed before the OKBC query, when the predictor model is invoked as part of the inference strategy for the query (involving r), the predictor is already trained to perform inference over query triples involving r. In Fig. 2, a_5 for query q gets executed at $PS[top] = (q, S_5)$.

State Transition Function (Δ): The transition function Δ is characterized by a set of unique transition conditions specified in terms of state bits. Table 3 gives all the state transition conditions. Given a state S_i, if S_i satisfies a transition condition C, an action corresponding to C is fired. For example, if in S_i, ILO=0, $CLUE$=0, $QERS$=1, QRF=0, action a_1 is invoked and on executing a_1, bits in S_i get updated and state transition occurs. Note that, following these conditions in Table 3, only a_2 can repeat up to δ_{IL} (see Table 1) number of times. It is not meaningful to execute any other action more than once for each d (OKBC query or clue from the user) in a given trail of state transitions from the initial state to a terminal state. Thus, if the current state of the FSM remains the same for δ_{IL} times, it means that inference is in-feasible for the current query (e.g., when no complete path feature is found) and the query remains unanswered.

3.2 IKAI Relation Prediction Model

Inferring the correctness of the resulting KBC query is performed by the relation Predictor (a neural network). It learns a single model ($\mathcal{F}$) with parameter Θ for all relations. The model is trained continuously with streaming clue triples. Note, the CLQs (acquired via action a_3) and MLQs (acquired via action a_2) in the KB are also used as training examples, besides the clue triples.

Given a relation r, IKAI uses the clue triples involving r (including the generated negative examples for them stored in $\mathcal{D}$ to train $\mathcal{F}$. When a query involving r is encountered at time j, IKAI selects k most recent examples of r from $\mathcal{D}$ (we select max. 500 examples per class label) to train $\mathcal{F}$ with a very small number of training epochs (to ensure real-time training of $\mathcal{F}$). For an already learned relation, this is like a quick recollection of the past learned knowledge before the query is answered. For a new r, only the acquired clues of r take part in training, where the learned weights Θ_{j-1} of $\mathcal{F}$ at time $j-1$ acts as the prior knowledge for learning Θ_j.

Often for a (new) relation r learned recently, the number of clue triples acquired in $\mathcal{D}$ may not be sufficient to learn r. Thus, to improve $\mathcal{F}$ continuously on *poorly learned tasks*, IKAI acquires additional clues for such r. Here, a *poorly learned task* denotes the task (relation) for which the validation performance of the *Predictor* at a given point in time is low compared to that on other tasks, measured by a evaluation metric. In our work, we use macro-F1 score to judge the overall learning performance of $\mathcal{F}$ on a given task. IKAI ranks all learned tasks based on validation macro-F1 scores and chooses ρ% (see Table 1) relations with lowest macro-F1 to acquire more clues. In this process, while executing action a_0 (see Sec. 4), besides the search operation, if the query relation r is evaluated as a poorly learned task, IKAI randomly sets $QRF = 0$ to fire a_1 for clue acquisition.

Our Predictor is based on the compositional vector space model [35, 12]. The model composes the implication of a path using RNN that takes as inputs the vector embeddings of binary relations in the path and outputs a vector in the semantic neighborhood of the test relation. For a training instance $i \in \mathcal{D}$, IKAI uses a LSTM [19] to compose the vector representation of each path $p \in P_c$ as v_p and vector representation of r as v_r. Next, IKAI computes the prediction value, $\mathbb{P}(r|s,t)$ as sigmoid of the mean cosine similarity of all v_p with v_r: $\mathbb{P}(r|s,t) = sigmoid(\frac{1}{|P_c|}\sum_{p \in P_c} cos(v_r, v_p))$.

$\mathcal{F}$ is trained by maximizing the log-likelihood. While predicting a query q of r, we infer q as +ve if $\mathbb{P}(r|h,t) \geq 0.5$ and -ve otherwise.

4 Experiments

We evaluate IKAI in terms of its predictive performance and strategy formulation ability. IKAI is designed for a *multi-user* chatting environment. Thus, it is natural for IKAI to observe many query

triples (and hence, accumulate more clues) for a relation from different users over time. Presently IKAI only adds the supporting facts into the KB. The predicted query triples are not added as they are unverified. In practice, IKAI can store these predicted triples in the KB as well after *cross-verification* by asking other users later. Similar, cross-verification strategy can also be used to deal with (intentional or unintentional) incorrect knowledge injection by user to the system's KB.

4.1 Evaluation Setup

Evaluating IKAI with real users in a crowd-source environment would be prohibitively time-consuming because both training and testing need a large number of real-time interactions between users and the system. Thus, we design a simulated interactive environment for the IKAI evaluation.

In the simulated environment, a **simulated user** (a program) is created to interact with **IKAI**. The (simulated) user has (1) a ***knowledge base*** ($\mathcal{K}_u$) for answering questions from IKAI, and (2) an ***OKBC query dataset*** (D_q) from which the user issues queries to IKAI. Here, D_q consists of a set of structured OKBC query triples q of the form (s, r?, t) readable by IKAI In practice, the user only issues OKBC queries to IKAI, but cannot evaluate the performance of the system unless the user knows the answer. To evaluate the performance of IKAI on D_q, we also label D_q while preparing the dataset (see Supplementary).

Table 4: Data statistics [*kwn*: known, *unk*: unknown, rel: relation]

KB Statistics	Freebase (FB)	WordNet (WN)
# Relations ($\mathcal{K}_{org}$ / $\mathcal{K}_b$)	1,345 / 1,273	18 / 12
# Entities ($\mathcal{K}_{org}$ / $\mathcal{K}_b$)	13, 871 / 13, 223	13, 595 / 13, 150
# Triples ($\mathcal{K}_{org}$ / $\mathcal{K}_b$)	854, 362 / 652, 790	107, 146 / 66, 338
# Test relations (*kwn* / *unk*)	25 (17 / 8)	18 (12 / 6)
# Train / valid / test instances	11,260 / 1223 / 7083	6628 / 711 / 3500
Entity statistics in Test query triples [s = source entity; t = target entity]		
Avg. % triples per rel with **only** s ***unk***	16.28	13.29
Avg. % triples per rel with **only** t ***unk***	16.29	10.94
Avg. % triples per rel **both** s **and** t ***unk***	4.69	4.17

As IKAI performs continuous online knowledge acquisition and learning, we evaluate its performance over time at different temporal checkpoints. We assume that, IKAI has already interacted with some users up to the evaluation start time t_{eval} and gathered a knowledge base ($\mathcal{K}_b$) (called the **base KB**), which includes its initial knowledge base and the knowledge obtained in its past interactions with users. Apart from $\mathcal{K}_b$, we also assume that at t_{eval}, IKAI's predictor model has already been trained to perform inference over a set of query relations seen so far. Recall our single predictor covers all past learned relations. We call these relations the *known* query relations. The query relations in dataset D_q not belonging to the set of *known* query relations are called *unknown* query relations. Note, D_q consists of query triples involving both *known* and *unknown* relations.

Evaluation KB Datasets. We create the simulated user's KB ($\mathcal{K}_u$), IKAI's base KB ($\mathcal{K}_b$) and the OKBC query dataset (D_q) along with the training and validation datasets (used in IKAI's *initial training phase*) from two standard KB datasets (see Table 4): (1) Freebase FB15k and (2) WordNet WN18 [9]. From each KB dataset, we build a large knowledge graph and use it as the original KB ($\mathcal{K}_{org}$). We also augment $\mathcal{K}_{org}$ with inverse triples (t, r^{-1}, s) for each (s, r, t) in $\mathcal{K}_{org}$ to increase graph connectivity following the existing KBC methods. Here, r^{-1} refers to the inverse relation of r that connects t to s (see "Actions" in Sec. 3.1). We provide details about the creation of simulated user's KB ($\mathcal{K}_u$), base KB ($\mathcal{K}_b$) and the OKBC query dataset (D_q) from $\mathcal{K}_{org}$, Hyper-parameter Settings and additional results in Supplementary Material.

Compared Models. Since no existing KBC method can solve the proposed OKBC problem, we compare various versions of IKAI: **(1) IKAI**: The proposed full version of IKAI.
(2) BG (Blind Guessing). To evaluate the guessing mechanism, BG fills all the missing or connecting links blindly with "-RelatedTo-" (the user has no answer). Note that, if no guessing is used and the user does not answer an OKBC query, IKAI will reject the query.
(3) w/o PTL (without Past Task Learning): IKAI does not ask for additional clues for poorly learned tasks and retrain $\mathcal{F}$ on past learned tasks. Only training of $\mathcal{F}$ on new tasks is enabled.

Note that although [43, 42] deal with OKBC, they infer facts using a text corpus, while we do through user-interactions. The two approaches are complementary and not comparable.

Evaluation Metrics. We use positive (+) F1 score and Macro F1 score for evaluation.

4.2 Results and Analysis

For evaluation on a given KB (Freebase or WordNet), we first randomly generate a chronological ordering of the corresponding training and validation instances and train IKAI so that we get initial

Table 5: Comparison of predictive performance of *various versions* of IKAI .

Dataset	Models	Rel - K / Ent -K		Rel - K / Ent -UNK		Rel - UNK / Ent - K		Rel - UNK / Ent -UNK		Overall	
		F1(+)	Macro-F1	F1(+)	Macro-F1	F1(+)	Macro-F1	F1(+)	Macro-F1	F1(+)	Macro-F1
Freebase	BG	**0.58**	0.62	0.49	0.56	0.43	0.53	0.38	0.50	0.50	0.57
	w/o PTL	0.55	0.65	**0.53**	0.62	**0.52**	0.42	**0.52**	0.41	0.53	0.58
	IKAI	**0.58**	**0.68**	0.49	0.59	**0.52**	**0.61**	0.44	**0.57**	**0.53**	**0.63**
WordNet	BG	0.54	0.47	0.53	0.52	0.48	0.47	0.49	0.48	0.52	0.48
	w/o PTL	**0.66**	**0.74**	**0.56**	**0.62**	0.46	0.28	0.48	0.32	0.55	0.59
	IKAI	**0.66**	0.71	0.55	0.60	**0.61**	**0.66**	**0.51**	**0.50**	**0.62**	**0.65**

relation prediction model for all known tasks (relations). Once this past training phase is over, we randomly generate a chronological ordering of all test instances in D_q, which are fed to IKAI in a streaming fashion, and then evaluate IKAI online. The **avg. test query processing time** of IKAI **for Freebase is 0.24 sec and for WordNet is 0.15 sec (on a Nvidia Titan RTX GPU).**

Table 5 shows the +ve F1 and Macro F1 scores of various versions of IKAI. We present the overall performance on the whole query dataset as well as performance on subsets of query datasets, denoted as (*Rel-X, Ent-Y*), where X and Y can be either *known* ('K') or *unknown* ('UNK') and '*Rel*' denotes query relation and '*Ent*' denotes query entities. So, here, (*Rel-K, Ent-UNK*) denotes the subset of the dataset having query triples involving known relation, but unknown entity (either s or t is unknown) with respect to $\mathcal{K}_b$. From Table 5, IKAI performs the best among all variants on both datasets. Due to the use of contextual similarity of entity-pairs, IKAI's guessing mechanism also works better than blind guessing (BG). The past task selection mechanism of IKAI also improves over w/o PTL, as it acquires more clues during testing for poorly performed tasks (evaluated on validation set).

Table 6: Performance (+ve-F1 scores) of IKAI on (filtered) test queries observed over time [% TTO = % of Test Triples Observed]. Number in [] shows the test triples count for TTO = 100%. "all" denotes overall performance.

% TTO	Freebase			WordNet		
	kwn [9]	*unk* [293]	*all* [302]	*kwn* [21]	*unk* [105]	*all* [126]
50%	0.0	0.49	0.50	0.94	0.79	0.82
100%	**0.54**	**0.58**	**0.58**	**0.95**	**0.87**	**0.88**

Table 7: Performance (+ve F1) of IKAI on user's responses to CLQs and MLQs evaluated on Full (100%) test data.

Response	Freebase			WordNet		
	kwn	*unk*	*Overall*	*kwn*	*unk*	*overall*
No	0.58	0.48	0.52	0.65	0.56	0.61
Yes	**0.58**	**0.50**	**0.53**	**0.66**	**0.57**	**0.62**

Table 6 shows the result of IKAI performance improvement over time at various temporal checkpoints for known (kwn) [corresponds to (*Rel-K, Ent-K*)], unknown (unk) [corresponds to queries for which either Rel-UNK or Ent-UNK] and overall queries. We noted in our experiments that, the simulated user's query satisfaction rate (proportions of queries that can be answered by the simulated user; 1% MLQs and 11% CLQs) is very low. Thus, to fairly evaluate the performance improvement of IKAI over time, we report the results (in Table 6) in *filtered setting*, where we only consider test queries that the user is able to answer at least one CLQ or MLQ (when asked). From Table 6, we see that the IKAI performance improves significantly as they see more test queries, and gain more knowledge (whenever user is able to respond) over time.

Table 7 shows the results of IKAI on user responses to MLQ's and CLQ's on the full test data. Answering MLQ's and CLQ's is hard for the simulated user as $\mathcal{K}_u$ often lacks the required triples. Thus, we analyze how the performance is affected if the user does not respond at all. The results show a clear trend in overall performance gain when the user responds (also, supported by Table 6). However, the improvement is not very large as the simulated user's query satisfaction rate (1% MLQs and 11% CLQs) is very small. But, the analysis shows the effectiveness of IKAI's *guessing mechanism* with such minimal user interaction (see performance gain of IKAI over BG in Table 5).

5 Conclusions

This paper proposed to solve the *open-world knowledge base completion* (OKBC) problem through user interactions and inference. This approach is motivated by the need to learn new knowledge in human-machine dialogues. IKAI can potentially serve as a chatbot knowledge learning engine. Our experimental results demonstrated both IKAI's predictive quality and strategy formulation ability. In our future work, we plan to learn other forms of knowledge to make IKAI more complete.

References

[1] Vito Walter Anelli, Pierpaolo Basile, Derek Bridge, Tommaso Di Noia, Pasquale Lops, Cataldo Musto, Fedelucio Narducci, and Markus Zanker. Knowledge-aware and conversational recommender systems. In *ACM RecSys*, 2018.

[2] Gabor Angeli, Melvin Jose Johnson Premkumar, and Christopher D Manning. Leveraging linguistic structure for open domain information extraction. In *ACL-IJCNLP*, 2015.

[3] Michele Banko, Michael J Cafarella, Stephen Soderland, Matthew Broadhead, and Oren Etzioni. Open information extraction from the web. In *IJCAI*, 2007.

[4] Junwei Bao, Nan Duan, Zhao Yan, Ming Zhou, and Tiejun Zhao. Constraint-based question answering with knowledge graph. In *COLING*, 2016.

[5] Jonathan Berant, Andrew Chou, Roy Frostig, and Percy Liang. Semantic parsing on freebase from question-answer pairs. In *EMNLP*. 2013.

[6] Kurt Bollacker, Colin Evans, Praveen Paritosh, Tim Sturge, and Jamie Taylor. Freebase: a collaboratively created graph database for structuring human knowledge. In *SIGMOD*, 2008.

[7] Antoine Bordes, Sumit Chopra, and Jason Weston. Question answering with subgraph embeddings. In *EMNLP*, 2014.

[8] Antoine Bordes, Nicolas Usunier, Sumit Chopra, and Jason Weston. Large-scale simple question answering with memory networks. *arXiv preprint arXiv:1506.02075*, 2015.

[9] Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. Translating embeddings for modeling multi-relational data. In *NIPS*, 2013.

[10] Antoine Bordes and Jason Weston. Learning end-to-end goal-oriented dialog. *arXiv preprint arXiv:1605.07683*, 2016.

[11] Antoine Bordes, Jason Weston, Ronan Collobert, and Yoshua Bengio. Learning structured embeddings of knowledge bases. In *AAAI*, 2011.

[12] Rajarshi Das, Arvind Neelakantan, David Belanger, and Andrew McCallum. Chains of reasoning over entities, relations, and text using recurrent neural networks. *arXiv preprint arXiv:1607.01426*, 2016.

[13] Anthony Fader, Stephen Soderland, and Oren Etzioni. Identifying relations for open information extraction. In *EMNLP*, 2011.

[14] Matt Gardner and Jayant Krishnamurthy. Open-vocabulary semantic parsing with both distributional statistics and formal knowledge. In *AAAI*, 2017.

[15] Matt Gardner and Tom M Mitchell. Efficient and expressive knowledge base completion using subgraph feature extraction. In *EMNLP*, 2015.

[16] Matt Gardner, Partha Pratim Talukdar, Jayant Krishnamurthy, and Tom Mitchell. Incorporating vector space similarity in random walk inference over knowledge bases. In *ACL*, 2014.

[17] Marjan Ghazvininejad, Chris Brockett, Ming-Wei Chang, Bill Dolan, Jianfeng Gao, Wen-Tau Yih, and Michel Galley. A knowledge-grounded neural conversation model. In *AAAI-2018*, 2018.

[18] Ben Hixon, Peter Clark, and Hannaneh Hajishirzi. Learning knowledge graphs for question answering through conversational dialog. In *The 2015 Annual Conference of the North American Chapter of the ACL (NAACL-2015)*, 2015.

[19] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 1997.

[20] Johannes Hoffart, Fabian M Suchanek, Klaus Berberich, and Gerhard Weikum. Yago2: A spatially and temporally enhanced knowledge base from wikipedia. *Artificial Intelligence*, 194:28–61, 2013.

[21] Yoji Kiyota, Sadao Kurohashi, and Fuyuko Kido. Dialog navigator: A question answering system based on large text knowledge base. In *COLING*, 2002.

[22] Ni Lao, Einat Minkov, and William W Cohen. Learning relational features with backward random walks. In *ACL*, 2015.

[23] Ni Lao, Tom Mitchell, and William W Cohen. Random walk inference and learning in a large scale knowledge base. In *EMNLP*, 2011.

[24] Phong Le, Marc Dymetman, and Jean-Michel Renders. Lstm-based mixture-of-experts for knowledge-aware dialogues. *arXiv preprint arXiv:1605.01652*, 2016.

[25] Jiwei Li, Alexander H Miller, Sumit Chopra, Marc'Aurelio Ranzato, and Jason Weston. Dialogue learning with human-in-the-loop. *arXiv preprint arXiv:1611.09823*, 2016.

[26] Jiwei Li, Alexander H Miller, Sumit Chopra, Marc'Aurelio Ranzato, and Jason Weston. Learning through dialogue interactions. *arXiv preprint arXiv:1612.04936*, 2016.

[27] Yinong Long, Jianan Wang, Zhen Xu, Zongsheng Wang, Baoxun Wang, and Zhuoran Wang. A knowledge enhanced generative conversational service agent. In *DSTC6 Workshop*, 2017.

[28] Andrea Madotto, Chien-Sheng Wu, and Pascale Fung. Mem2seq: Effectively incorporating knowledge bases into end-to-end task-oriented dialog systems. *arXiv preprint arXiv:1804.08217*, 2018.

[29] Sahisnu Mazumder and Bing Liu. Context-aware path ranking for knowledge base completion. In *IJCAI*, 2017.

[30] Sahisnu Mazumder, Bing Liu, Shuai Wang, and Nianzu Ma. Lifelong and interactive learning of factual knowledge in dialogues. In *SIGDIAL*, 2019.

[31] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. *In ICLR Workshop*, 2013.

[32] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In *NIPS*, 2013.

[33] George A Miller. Wordnet: a lexical database for english. *In Communications of the ACM*, 1995.

[34] Tom M Mitchell, William W Cohen, Partha Pratim Talukdar, Justin Betteridge, Andrew Carlson, Matthew Gardner, Bryan Kisiel, Jayant Krishnamurthy, et al. Never ending learning. In *AAAI*, 2015.

[35] Arvind Neelakantan, Benjamin Roth, and Andrew McCallum. Compositional vector space models for knowledge base completion. In *ACL*, 2015.

[36] Maximilian Nickel, Kevin Murphy, Volker Tresp, and Evgeniy Gabrilovich. A review of relational machine learning for knowledge graphs. *arXiv preprint arXiv:1503.00759*, 2015.

[37] Chen Qu, Liu Yang, W Bruce Croft, Johanne R Trippas, Yongfeng Zhang, and Minghui Qiu. Analyzing and characterizing user intent in information-seeking conversations. In *SIGIR*, 2018.

[38] Amrita Saha, Vardaan Pahuja, Mitesh M Khapra, Karthik Sankaranarayanan, and Sarath Chandar. Complex sequential question answering: Towards learning to converse over linked question answer pairs with a knowledge graph. In *AAAI*, 2018.

[39] Denis Savenkov and Eugene Agichtein. When a knowledge base is not enough: Question answering over knowledge bases with external text data. In *SIGIR*, 2016.

[40] Michael Schmitz, Robert Bart, Stephen Soderland, Oren Etzioni, et al. Open language learning for information extraction. In *EMNLP and CoNLL*, 2012.

[41] Iulian Vlad Serban, Alessandro Sordoni, Ryan Lowe, Laurent Charlin, Joelle Pineau, Aaron Courville, and Yoshua Bengio. A hierarchical latent variable encoder-decoder model for generating dialogues. In *AAAI*, 2017.

[42] Haseeb Shah, Johannes Villmow, Adrian Ulges, Ulrich Schwanecke, and Faisal Shafait. An open-world extension to knowledge graph completion models. In *AAAI*, 2019.

[43] Baoxu Shi and Tim Weninger. Open-world knowledge graph completion. In *Proceedings of AAAI-2018*, 2018.

[44] Y. Sun and Yi Zhang. Conversational recommender system. In *SIGIR*, 2018.

[45] Patrick Verga, Arvind Neelakantan, and Andrew McCallum. Generalizing to unseen entities and entity pairs with row-less universal schema. In *EACL*, 2017.

[46] Oriol Vinyals and Quoc Le. A neural conversational model. *arXiv preprint arXiv:1506.05869*, 2015.

[47] Quan Wang, Jing Liu, Yuanfei Luo, Bin Wang, and C Lin. Knowledge base completion via coupled path ranking. In *ACL*, 2016.

[48] Sida I Wang, Samuel Ginn, Percy Liang, and Christoper D Manning. Naturalizing a programming language via interactive learning. *arXiv preprint arXiv:1704.06956*, 2017.

[49] Sida I Wang, Percy Liang, and Christopher D Manning. Learning language games through interaction. *arXiv preprint arXiv:1606.02447*, 2016.

[50] Robert West, Evgeniy Gabrilovich, Kevin Murphy, Shaohua Sun, Rahul Gupta, and Dekang Lin. Knowledge base completion via search-based question answering. In *WWW*, 2014.

[51] Jason E Weston. Dialog-based language learning. In *NIPS*, 2016.

[52] Fei Wu and Daniel S Weld. Open information extraction using wikipedia. In *ACL*, 2010.

[53] Rui Yan and Dongyan Zhao. Coupled context modeling for deep chit-chat: towards conversations between human and computer. In *ACM SIGKDD*. ACM, 2018.

[54] Liu Yang, Minghui Qiu, Chen Qu, Jiafeng Guo, Yongfeng Zhang, W Bruce Croft, Jun Huang, and Haiqing Chen. Response ranking with deep matching networks and external knowledge in information-seeking conversation systems. In *SIGIR*, 2018.

[55] Tom Young, Erik Cambria, Iti Chaturvedi, Hao Zhou, Subham Biswas, and Minlie Huang. Augmenting end-to-end dialog systems with commonsense knowledge. In *AAAI*, 2018.

[56] Haichao Zhang, Haonan Yu, and Wei Xu. Listen, interact and talk: Learning to speak via interaction. *arXiv preprint arXiv:1705.09906*, 2017.

[57] Yongfeng Zhang, Xu Chen, Qingyao Ai, Liu Yang, and W Bruce Croft. Towards conversational search and recommendation: System ask, user respond. In *CIKM*. ACM, 2018.

[58] Hao Zhou, Huang Minlie, Tianyang Zhang, Xiaoyan Zhu, and Bing Liu. Emotional chatting machine: Emotional conversation generation with internal and external memory. In *Proceedings of AAAI-2018*, 2018.

[59] Hao Zhou, Tom Young, Minlie Huang, Haizhou Zhao, Jingfang Xu, and Xiaoyan Zhu. Commonsense knowledge aware conversation generation with graph attention. In *IJCAI*, 2018.

Supplementary Material

A Simulated User Creation and Test Dataset Generation

To simulate the inital interactions and knowledge learning process of IKAI up to time t_{eval}, we create a training and validation dataset (in addition to the test or query set D_q) involving triples of *known* query relations and entities. We assume that IKAI has already been trained on these training and validation dataset triples (in the past) before the evaluation on D_q starts. We call this training process the ***initial training phase*** of IKAI. Once the initial training phase is over, IKAI is ready to perform the online training (with clues to be acquired) evaluation process (with D_q) at t_{eval}.

Given the original KB graph ($\mathcal{K}_{org}$) (Freebase or WordNet), we create the simulated user's knowledge base ($\mathcal{K}_u$), the base KB ($\mathcal{K}_b$), the OKBC query or test dataset (D_q) and the training and validation datasets from $\mathcal{K}_{org}$ (Table 4 in main paper shows the results).

For Freebase, we found 86 relations with ≥ 1000 triples and randomly selected 25 from various domains. We randomly shuffle the list of 25 relations, select 34% of them as *unknown* relations and consider the rest (66%) as *known* relations. For each *known* relation r, we randomly shuffle the list of distinct triples for r, choose (maximum) 500 triples and split them into 60% training, 10% validation and 20% test. Rest 10% along with the leftover (not included in the list of 500) triples are added to the user knowledge $\mathcal{K}_u$. The training and validation triples are used in IKAI's past training phase and test triples constitute D_q.

For each *unknown* relation r, we remove all triples of r from $\mathcal{K}_{org}$, randomly choose 20% triples among these triples and reserve them as test instances for unknown r. Rest of the triples (80%) for *unknown* relation r are added to $\mathcal{K}_u$. In this process, we also make sure that the query instances involving *unknown* relation r are excluded from $\mathcal{K}_u$. Thus, the user cannot provide the query triple itself as a clue to IKAI (during inference) and also, to simulate the case that the user does not know the answer of its issued query. Note that, at this point, $\mathcal{K}_u$ has at least 10% of triples for each r (*known* and *unknown* to IKAI) and thus, the user is always able to provide clues for OKBC queries involving both *known* and *unknown* relations. Note, if the user cannot provide a clue for an unknown relation (not likely), the OKBC query cannot be converted to a KBC query and IKAI will not be able to answer the query.

To create queries involving unknown entities, we randomly choose 10% of the entities present in test triples, remove all triples involving those entities from $\mathcal{K}_{org}$ and add them to $\mathcal{K}_u$. At this point, $\mathcal{K}_{org}$ gets reduced to $\mathcal{K}_b$ (base KB) for IKAI. Note that, due to removal of a significant amount of triples, $\mathcal{K}_b$ becomes *sparser* (see "# triples" row in Table 4), making the inference task harder. The WordNet dataset being small, we use all its 18 relations for evaluation and create the OKBC query (test) dataset, training and validation dataset, $\mathcal{K}_u$, $\mathcal{K}_b$ following Freebase.

Although the simulated user may provide clues 100% of the time (otherwise, the query is not answerable), it often cannot respond to MLQs and CLQs due to lack of the required triples/facts. Thus, we further enrich $\mathcal{K}_u$ with external KB triples. Due to a fair amount of entity overlapping, we choose NELL[1] for enriching $\mathcal{K}_u$ in case of Freebase, and ConceptNet [2] for enriching $\mathcal{K}_u$ in case of WordNet.

B Dataset Labeling

We now discuss labeling of the query dataset (D_q) and training and validation datasets. Given a relation r and a triple (s, r, t) in the created training, validation or test query dataset so far, each pair (s, t) in them is regarded as a +ve instance for r. Thus, (s, r?, t) represents a +ve OKBC query (i.e., the true answer of the query is "YES"). For each +ve instance (s, t), we generate two negative ones, one by randomly corrupting the source s, and the other by corrupting the target t (same as in action a_1 in Sec. 3.1). The two corrupted triples corresponding to each +ve triple are treated as -ve OKBC query (i.e., the true answer of the query is "NO") and added to the corresponding dataset (training / validation / test). These datasets are thus ***expanded*** with negative instances. Note, the test triples are

[1] http://rtw.ml.cmu.edu/rtw/kbbrowser/

[2] http://conceptnet.io/

Table 8: IKAI's Inference strategies (ordered by frequency).

$\delta_{IL} = 0, \delta_p = 3$ **[C: 0.50]**	$\langle a_0, a_4, a_2, a_2, a_5 \rangle$	$\delta_{IL} = 5, \delta_p = 1$ **[C: 1.0]**
$\langle a_0, a_4, a_5 \rangle$	$\langle a_0, a_4, a_2, a_5 \rangle$	$\langle a_0, a_4, a_2, a_5 \rangle$
$\delta_{IL} = 1, \delta_p = 3$ **[C: 0.97]**	$\langle a_0, a_4, a_5 \rangle$	$\langle a_0, a_4, a_5 \rangle$
$\langle a_0, a_4, a_2, a_5 \rangle$	$\langle a_0, a_3, a_4, a_2, a_2, a_5 \rangle$	$\langle a_0, a_3, a_4, a_2, a_5 \rangle$
$\langle a_0, a_1, a_4, a_5 \rangle$	$\langle a_0, a_3, a_4, a_2, a_5 \rangle$	$\langle a_0, a_3, a_4, a_5 \rangle$
$\langle a_0, a_3, a_4, a_5 \rangle$	$\langle a_0, a_3, a_4, a_5 \rangle$	$\langle a_0, a_1, a_3, a_4, a_2, a_5 \rangle$
$\langle a_0, a_4, a_5 \rangle$	$\langle a_0, a_1, a_3, a_4, a_2, a_5 \rangle$	$\langle a_0, a_1, a_4, a_5 \rangle$
$\delta_{IL} = 3, \delta_p = 3$ **[C: 1.0]**	$\langle a_0, a_1, a_4, a_2, a_2, a_5 \rangle$	$\langle a_0, a_1, a_3, a_4, a_5 \rangle$
$\langle a_0, a_4, a_2, a_2, a_2, a_5 \rangle$	$\langle a_0, a_1, a_4, a_2, a_5 \rangle$	$\langle a_0, a_1, a_4, a_2, a_5 \rangle$

not in $\mathcal{K}_b$ or $\mathcal{K}_u$. Also, each triple/instance in training or testing is represented with its path features extracted while executing the inference strategy.

C Hyper-parameter Setting

Unless specified, the *empirically* set parameters (Table 1) of IKAI are: $\delta_{IL} = 5$ for test queries and $\delta_{IL} = 2$ for training/clue instances, $\delta_p = 3$, $\eta_\pi = 7$, $\eta_w = 20$, $l = 0.07$, $h = 0.2$, $\rho = 25\%$. For online training, we set batch-size as 32, training epoch as 2 (3) for learned (unlearned) task for each query processing, dropout as 0.2, hidden units and embedding size as 300 and learning rate as 5e-4. Adam optimizer is used.

D IKAI Inference Strategy Formulation

To evaluate the correctness of strategy formulation, we introduce *Coverage* (C), the fraction of total query triples answered by IKAI. Note, C depends on the parameter δ_{IL} and δ_p.

We use WordNet query dataset to illustrate various inference strategies formulated by IKAI (see Table 8) with varying δ_{IL} and δ_p values (column 2 follows collomn 1 in Table 8). When $\delta_{IL} = 0$, IKAI cannot interact with user and thus, can only solve KBC queries from user and rejects all queries involving unknown s, r or t. Thus, coverage C drops significantly (0.5). The only formulated strategy here is $\langle a_0, a_4, a_5 \rangle$, which involves only processing actions. None of the interactive actions a_1, a_2, a_3 was invoked by the FSM.

When $\delta_{IL} = 1$, IKAI can only interact once with the user per query for knowledge acquisition. In such a case, IKAI acquires knowledge well for instances where either of an entity or relation is unknown. However, as one unknown entity may appear in multiple test triples, once the entity becomes known, IKAI doesn't need to ask for it again (knowledge accumulation). Thus, C increases significantly (0.97). Each of the formulated strategies (4 strategies in Table 8) here has only one interactive action [one of $\{a_1, a_2, a_3\}$].

When $\delta_{IL} = 3$, IKAI is able to perform inference on all instances and coverage C becomes 1. For $\delta_p = 1$, IKAI uses a_2 only once (as only one MLQ satisfies δ_p) compared to at most 3 MLQs in formulated strategies for $\delta_p = 3$. In summary, IKAI's FSM module can correctly formulate inference strategies based on specified parameter values.

E Path Ranking: Overview

Path Ranking. PR algorithm is to enumerate relation paths between two entities (s, t) in a KB (encoded as a multi-relation graph) and use those relation paths as features to train the predictive model for relation r. To enumerate a set of path features for a given entity pair (s, t), PR performs random walks over KB to find paths starting from s (or t) and ending at t (or s). Once path features are extracted for a set of positively labeled (linked by r) and a set of negatively labeled (not linked by r) entity-pairs, the PR method trains a binary classifier (Predictor) for relation r using those features. The trained classifier for r is then used to predict if a given test entity pair (not linked by r in the KB) can be linked by r.

Context-aware Path Ranking (C-PR). Given s and t, C-PR performs context-aware bidirectional random walk to enumerate path features. In particular, C-PR enumerates a forward path starting from s and a backward path starting from t such that the intermediate entities appearing in a path (forward or backward) has high contextual similarity with respect to both s and t. The contextual similarity between any two entities is computed as the cosine similarity between the corresponding entity embedding vectors learned from text documents using neural word embedding. C-PR treats entities in the KB as words and learns and leverages their embeddings to guide the path exploration process. If a common entity appears in both forward and backward path at any time step during exploration, C-PR stops exploration, merges the forward and backward paths into a single path from s to t and returns the relation sequence in the merged path as a contextual path feature from s to t, which is then used in inference.