Learning-to-Ask: Knowledge Acquisition via 20 Questions

Yihong Chen†, Bei Chen‡, Xuguang Duan†, Jian-Guang Lou‡, Yue Wang†, Wenwu Zhu†, Yong Cao*

† Tsinghua University, ‡ Microsoft Research, * Alibaba AI Labs

20 Questions

In the online interactive *20 Questions*, the game agent plays the role of the guesser and tries to figure out what is in the mind of the human player.

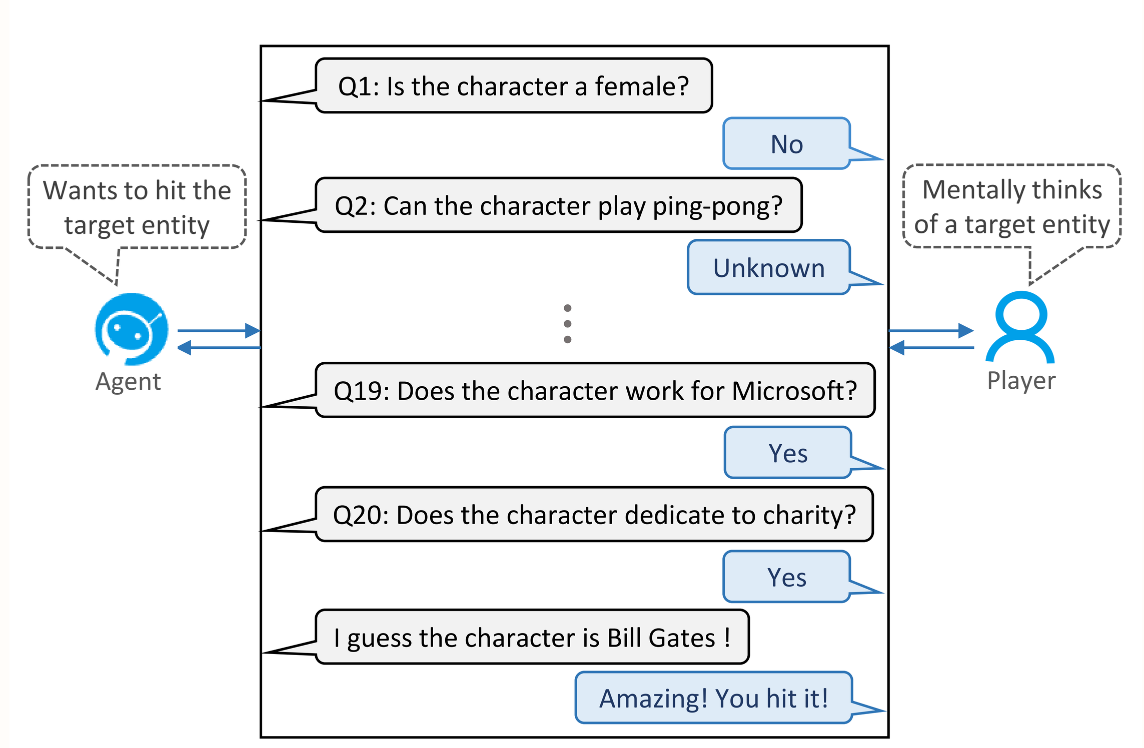

Figure 1: An example episode of *20 Questions*.

Information Seeking & Knowledge Acquisition ?

Challenges

(1) The agent needs efficient and robust *information seeking* (IS) strategies to work with noisy user responses and hit the target with as few questions as possible.

(2) To acquire valuable facts, the agent needs effective *knowledge acquisition* (KA) strategies which can identify important questions for a given entity.

Our goal is to learn both IS and KA strategies which boost each other, allowing the agent to start with a relatively small knowledge base and quickly improves in the absence of constant human supervision.

Contributions

(1) We are the first to formally integrate IS strategies and KA strategies into a unified framework, named Learning-to-Ask.

(2) Novel methods combining reinforcement learning (RL) with Bayesian treatment are proposed to learn IS strategies in LA.

(3) KA strategies learning in LA is achieved with a generalized matrix factorization (GMF) based method.

(4) We conduct experiments against simulators built on real world data.

Learning-to-Ask Framework

The knowledge base is represented as a $M \times N$ entity-question matrix D, with each row and each column corresponding to an entity in $\mathcal{E} = \{e_1, e_2, ..., e_M\}$ and a question in $\mathcal{Q} = \{q_1, q_2, ..., q_N\}$ respectively.

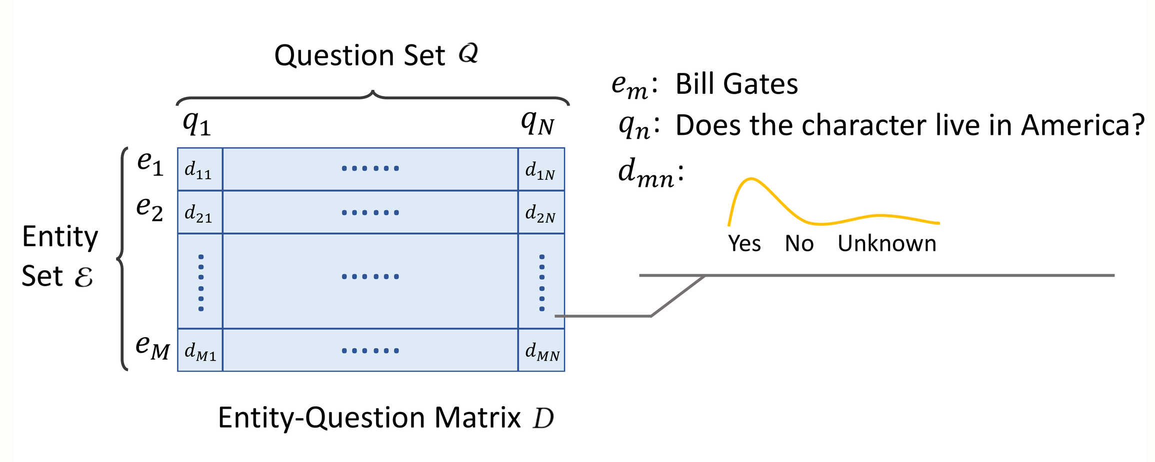

Figure 2: The knowledge base: entity-question matrix.

Backed by the knowledge base described above, the agent learns to ask questions within the LA framework.

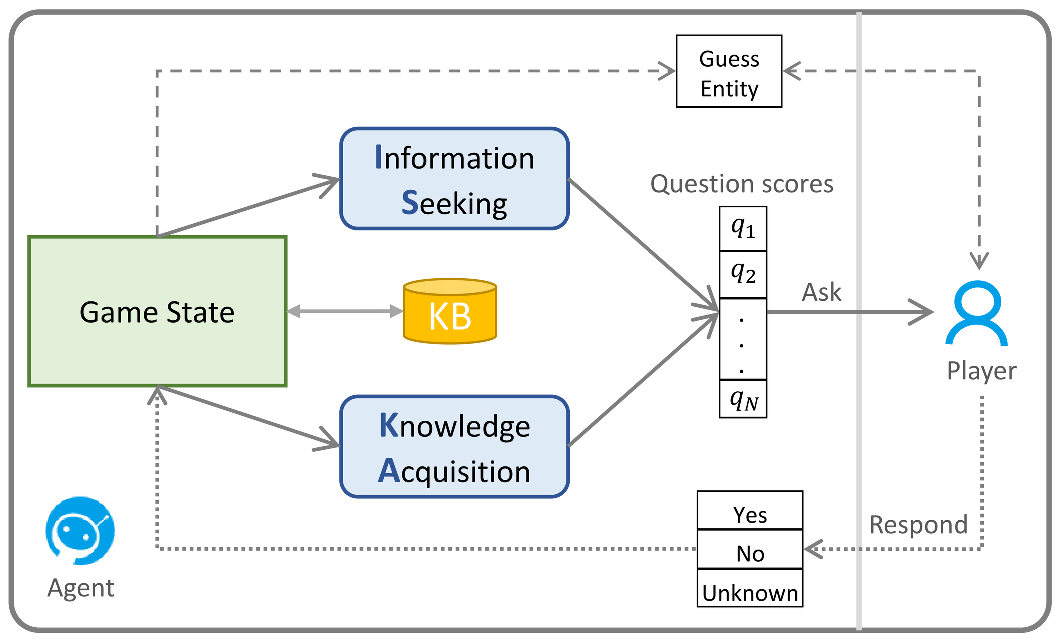

Figure 3: Learning-to-Ask framework.

IS Module

This module generates questioning strategies to collect necessary information to hit the target entity at the final guess. At each game step t, IS module detects the ambiguity in the collected information $(a_1, x_1, a_2, x_2, ..., a_{t-1}, x_{t-1})$, where a denotes a question and x denotes a response. And it scores the candidate questions according to how much uncertainty they can resolve about the target entity. The agent asks the next question based on the scores. Once adequate information collected, the agent makes a guess g.

KA Module

This module generates questioning strategies to improve the agent's knowledge base i.e. acquire knowledge about the missing entries in D. We assume that the ideal knowledge base for the agent should be consistent with the one human players possesses. Hence, the most valuable missing entries that KA module aims to acquire are those known to the human players but unknown to the agent. KA module scores the candidate question based on the values of entries pertaining to the guess g. The agent asks the question according to the scores and accumulates the response into its knowledge buffer.

Methods

Learning-to-Ask for Information Seeking

RL-based methods allow the agent to learn questioning strategies solely through experience. At the beginning, the agent has no idea what effective questioning strategies should be. It gradually learns to ask "good" questions by trial and error. Effective questioning strategies are reinforced by the reward signal that indicates whether the agent wins the game, namely whether it collects enough information to hit the target entity.

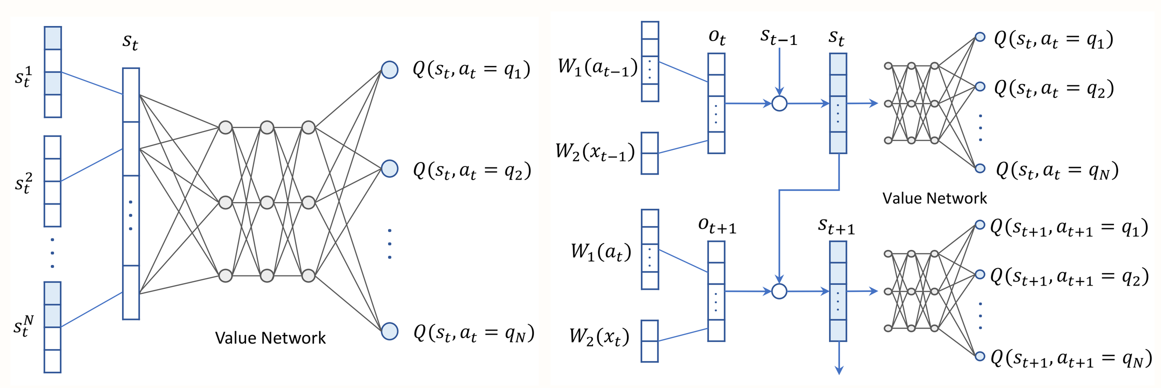

Figure 4: Q-network architecture for (a) LA-DQN and (b) LA-DRQN.

Learning-to-Ask for Knowledge Acquisition

Instead of asking questions blindly, we propose LA-GMF, a *generalized matrix factorization* (GMF) based method that takes into account two aspects to make KA more efficient: Uncertainty and Value. let $Y \in R^{M \times N}$ be the indicator matrix constructed from D: if the entry d_{mn} is known ($N_{mn} \neq 3$), then $y_{mn} = 1$; otherwise, $y_{mn} = 0$. Y can be factorized as $Y = U \cdot V$. GMF scores each entry d_{mn} and estimates the value as:

$$\hat{y}_{mn} = a_{out}(h^{\top}(U_m \odot V_n)). \quad (1)$$

Algorithm 1 One cycle of LA-GMF for KA module.

Input: knowledge base D, buffer size N_k
construct Y from D; factorize $Y = U \cdot V$; $Count = 0$
while $Count < N_k$ **do**
  $e_m = g$ is the guess of the current episode
  $Q_h = \{a_1, a_2, ..., a_{T_1}\}$ is the set of history asked questions
  $\theta_m = (\frac{1}{\sqrt{N_{m1}}}, \frac{1}{\sqrt{N_{m2}}}, ..., \frac{1}{\sqrt{N_{mN}}})$ is the uncertainty parameter
  for $step = 1, 2, \ldots, T_2$ **do**
    $Q_c = \varnothing$
    for $i = 1, 2, \ldots, N_c$ **do**
      $q_i \sim \text{Multinomial}(\theta_m)$ and $q_i \notin Q_c$ and $q_i \notin Q_h$
      $Q_c = Q_c + q_i$
    end for
    rank questions in Q_c by scores using Equation (1)
    ask the top question q in Q_c, collect the response in buffer
    $Q_h = Q_h + q$; $Count = Count + 1$
  end for
end while
update D using all the knowledge in buffer

Results

Q.1 Equipped with the learned IS strategies in LA, is the agent smart enough to hit the target entity?

Q.2 Equipped with the learned KA strategies in LA, can the agent acquire knowledge rapidly to improve its knowledge base?

Q.3 Is the acquired knowledge helpful to the agent in terms of increasing the efficiency in identifying the target entity?

Q.4 Starting with no prior, can the logical connections between questions learned by the agent over episodes?

Q.5 Can the agent engage players in contributing their knowledge unknowingly by asking appropriate questions?

Table 1: Statistics of datasets.

Dataset	# Entities	# Questions	Missing Ratio
Person	10440	1819	0.9597
PersonSub	500	500	0.5836

Table 2: Winning rate of the IS module.

Dataset	Entropy	LA-LIN	LA-DQN	LA-DRQN
PersonSub	0.5161	0.4020	0.7229	0.8535
Person	0.2873	–	–	0.6237

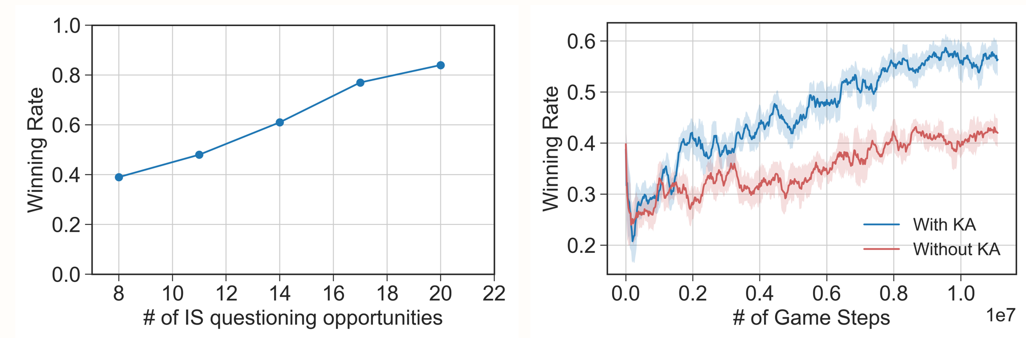

Figure 5: (a) LA-DRQN performance with varied IS questioning opportunities on PersonSub; (b) LA-DRQN performance of IS with KA on PersonSub.

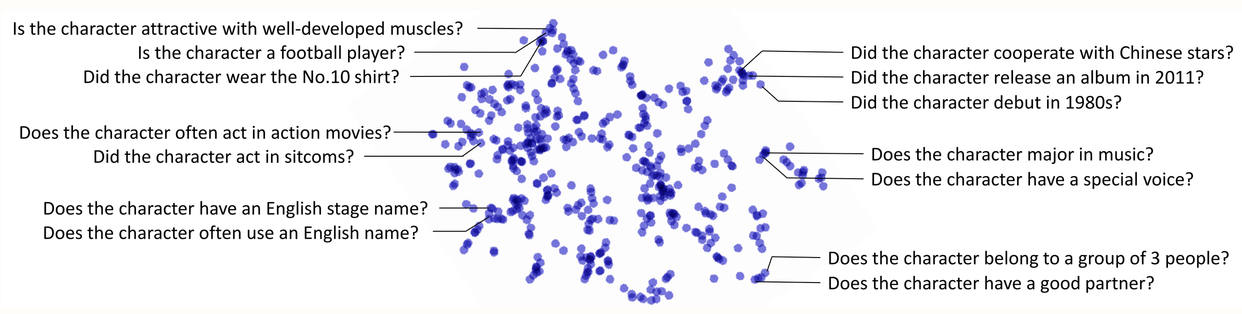

Figure 6: Visualization of question embeddings with LA-DRQN on PersonSub.

Table 3: Contrasting examples of KA questions.

Target Entity	KA questions with LA-GMF		KA questions with Uncertainty-Only	
3* Jay Chou[1]	Does the character mainly sing folk songs?	No	Did the character die unnaturally?	Unknown
	Is the character handsome?	Yes	Does the character like human beings?	Unknown
	Is the character born in Japan?	No	Is the character from a novel?	No
3* Jimmy Kudo[2]	Is the character brave?	Yes	Does the character only act in movies (no TV shows)?	Unknown
	Is the character from Japanese animation?	Yes	Has the character ever been traitorous?	Unknown
	Was the character dead?	No	Is the character a monster?	Unknown

24th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, London, United Kingdom