



Efficient relation extraction via quantum reinforcement learning

Xianchao Zhu^{1,2,3} · Yashuang Mu³ · Xuetao Wang³ · William Zhu⁴

Received: 28 December 2023 / Accepted: 3 February 2024
© The Author(s) 2024

Abstract

Most existing relation extraction methods only determine the relation type after identifying all entities, thus not fully modeling the interaction between relation-type recognition and entity mention detection. This article introduces a novel paradigm for relation extraction by treating relevant entities as parameters of relations and harnessing the strong expressive capabilities and acceleration advantages of quantum computing to address the relation extraction task. In this article, we develop a quantum hierarchical reinforcement learning approach to enhance the interaction between relation-type recognition and entity mention detection. The entire relation extraction process is broken down into a hierarchical structure of two layers of quantum reinforcement learning strategies dedicated to relation detection and entity extraction, demonstrating greater feasibility and expressiveness, especially when dealing with superimposed relations. Our proposed method outperforms existing approaches through experimental evaluations on commonly used public datasets, mainly showcasing its significant advantages in extracting superimposed relationships.

Keywords Relation extraction · Hierarchical reinforcement learning · Quantum computing · Quantum reinforcement learning

Introduction

Extracting entities, relationships, or events from a vast amount of unstructured text is critically important for building large-scale, reusable knowledge [1–5]. It can promote many real-world tasks, including construction of knowledge base [6–9], automatic question and answering system [10, 11], and biomedical text mining [12–14]. The input consists of unstructured text, while the output comprises triples containing source entities, target entities, and their corresponding entity relationships.

Traditional methods for relation extraction, known as pipelined approaches, treat extraction as two independent subtasks: first, entity identification, and then, extraction of the relationships between them [15–17]. This method is flexible and straightforward, but its phased execution does not exploit deeper interactions between the subtasks. Consequently, the upstream and downstream subtasks cannot improve their extraction strategies through interactive execution. In contrast, joint entity and relationship extraction frameworks use a single model to extract both entities and relationships, achieving better performance by leveraging the relationship between the two subtasks [18, 19, 19, 20, 20–22]. Most notably, Takanobu and colleagues develops a hierarchical reinforcement learning relationship extraction framework called HRL-RE. By decomposing the total extraction procedure into a hierarchical structure with two reinforcement learning strategies dedicated to relation detection and entity extraction, this framework enhances the interaction between entity identification and relation-type extraction. It makes handling superimposed relationships more feasible and natural [23]. Nevertheless, this approach has not achieved satisfying results in figuring out overlapped entities and sentence relations. One of the leading factors is that the learning procedure is sluggish and has many futile attempts, leading to low efficiency of strategy learning.

✉ Xianchao Zhu
xczhuiffs@163.com

¹ Key Laboratory of Grain Information Processing and Control, Ministry of Education, Henan University of Technology, Zhengzhou 450001, China

² Henan Key Laboratory of Grain Photoelectric Detection and Control, Henan University of Technology, Zhengzhou 450001, China

³ School of Artificial Intelligence and Big Data, Henan University of Technology, Zhengzhou 450001, China

⁴ Institute of Fundamental and Frontier Sciences, University of Electronic Science and Technology of China, Chengdu 610054, China

To address the issues mentioned above, we present a novel relation extraction approach known as quantum hierarchical reinforcement learning for relationship extraction (QHRL-RE), which incorporates the quantum computing advantages of quantum entanglement and superposition into a hierarchical reinforcement learning relation extraction model. Specifically, drawing inspiration from the breakthroughs of quantum reinforcement learning in speech recognition and control domains [24–27], we employ quantum long short-term memory (QLSTM) network models [28] for encoding and decoding representations in relation extraction tasks. These QLSTM models can better capture long-term dependencies in unstructured text data than traditional methods. Then, we utilize a hybrid quantum-classical algorithm, which iteratively optimizes tasks applicable to relation extraction while harnessing the enhanced expressive power conferred by quantum superposition. As a result, our presented approach is more efficient in discovering superimposed entities and relationships from structureless text than traditional methods. The experiment results on two classical relationship extraction datasets, NTY10 and NTY11, demonstrate that our proposed method outperforms classical relation extraction methods with similar architectures and model parameter counts, showcasing improved performance.

The organization of this article is as follows: in the second section, a brief overview of previous relationship extraction methods is provided, along with introductions to hierarchical reinforcement learning, quantum reinforcement learning, and a brief overview of the HRL-RE approach. The third section introduces the technique, quantum hierarchical reinforcement learning for relationship extraction (QHRL-RE) presented in this paper. This section explains how this method addresses the challenges of entity and relationship extraction in cases of overlap, employing quantum hierarchical reinforcement learning techniques. Subsequently, the fourth section presents experimental results, including experiments conducted on two publicly accessible New York Times corpora to show the superiority of the proposed algorithm in relationship extraction tasks, particularly in the context of superimposed entity relationships. Finally, the fifth section summarizes the main findings and contributions of the entire paper.

Related work

This section presents a concise review of relation extraction methods, a brief review of quantum reinforcement learning methods and hierarchical reinforcement learning, and an introduction to hierarchical reinforcement learning for relation extraction.

Relations extraction

Relationship extraction plays a powerful role in information extraction applications [1, 29–33]. Javeed proposes a distant supervised relation extraction model based on the attention mechanism of a new relation representation [34]. Various joint extraction methods have been proposed [18–20, 35]. For example, Zheng et al. treat entity and relationship extraction as sequence tagging tasks. They use bidirectional LSTM and unidirectional LSTM for encoding and decoding, with the output layer simultaneously labeling entities and relationships, achieving entity relationship joint extraction [20]. Bjerne et al. introduces the concept of relation identifiers, explicitly representing phrases that indicate the presence of relationships in the sentences and then selecting their parameters to reduce the intrinsic complexity of tasks [36]. More recently, reinforcement learning has been effective in relationship extraction tasks [37–40]. Feng et al. uses reinforcement learning to discover entities and relationship types jointly [41]. Qin et al. proposes a deep reinforcement learning method for relationship extraction [40]. Feng et al. also suggests a relationship extraction method comprising an RL entity trigger and a CNN relationship identifier [42]. These approaches aim to improve the accuracy and robustness of relationship extraction by considering entity recognition and relationship identification as interconnected tasks, in contrast to the traditional pipeline methods.

Quantum reinforcement learning

Quantum reinforcement learning (QRL) can be dated to reference [43]. Nevertheless, this method requires quantifying the environment, which may not be feasible in most real environment scenarios. This paper focuses on the latest developments of variational quantum circuit (VQC)-based QRL for traditional domains. The first VQC-based QRL [44] is a quantum version of deep Q-learning (DQN) and adopts discrete state spaces and action spaces in experimental domains such as Frozen Lake. Subsequent advanced work in quantum deep Q-learning has been deemed continuous observation spaces, for example, in the Cart-Pole problem [45–49]. The work in [50] extends the VQC framework further to improve DQN into double-deep Q-learning (DDQN) and adopts QRL to address robot operation tasks. In addition to learning Q functions as value functions, recent developments have introduced QRL methods to learn policy functions. For example, [51] describes quantum policy gradient reinforcement learning using the REINFORCE algorithm. Subsequently, work [52] considers an enhanced policy gradient algorithm known as proximal policy optimization (PPO) with VQC and demonstrates that quantum models with few parameters can outperform classical models.

Hierarchical reinforcement learning

Hierarchical reinforcement learning (HRL) is a significant branch of reinforcement learning (RL) that distinguishes itself from classical RL methods [53–55]. HRL leverages hierarchical abstraction techniques to improve RL structurally, focusing on addressing challenges RL struggles with, such as sparse rewards, sequential decision-making, and weak transferability. This approach enhances exploration and transfer capabilities. Most representatively, the options framework may be the most common formalism that allows agents to reason regarding extended actions [56–60]. This framework models courses of action as options, which can accelerate learning in different ways, allowing, for example, faster credit assignment, planning, transfer learning, and better exploration.

Hierarchical reinforcement learning for relationship extraction

The HRL-RE approach breaks down the total entity and relationship extraction mission into two component tasks [23]. It first identifies sentence relationships and then discovers a pair of entities corresponding to that relation type. HRL-RE calculates a policy based on the states processed by a Bi-LSTM and obtains the relationship type in the high-level relationship detection subtask. Once the relationship type is received, this high-level strategy delegates the low-level component task of entity extraction. HRL-RE computes a strategy using the Monte Carlo (MC) gradient estimation approach to get the entity pair associated with that relationship in the low-level subtask. After the present low-level component task is completed, the high-level reinforcement learning component task searches for the subsequent relationship in the sentence. This hierarchical method aims to ameliorate the effectiveness of relationship extraction by breaking down the task into more manageable subtasks, with each level focusing on a specific aspect of the extraction process [23].

The HRL-RE method enhances the accuracy of entity and relationship extraction and, to some extent, addresses the issue of superimposed entities and relationships. However, this approach only achieves satisfactory results when handling superimposed entities and sentence relationships. The main reason behind this is that the learning process is cumbersome, with many ineffective attempts leading to inefficient policy learning. In cases involving superimposed entities and complex sentence structures, the learning process may need help to navigate the search space effectively. This inefficiency can hinder the method’s ability to accurately extract relationships and entities from such sentences. Improving the efficiency of the strategy learning process by optimizing the reinforcement learning algorithm or explor-

ing alternative approaches could be a potential avenue for addressing this limitation.

Quantum hierarchical reinforcement learning for relationship extraction

This section introduces a novel approach to jointly extract superimposed entities and relationships, called quantum hierarchical reinforcement learning for relationship extraction (QHRL-RE, as shown in Figure 3). This method leverages the advantages of quantum computing, specifically quantum entanglement, and superposition, in combination with hierarchical reinforcement learning to address the problems outlined in “[Hierarchical reinforcement learning for relationship extraction](#)”. Specifically, drawing inspiration from the breakthroughs of quantum reinforcement learning in the speech recognition domain, we employ quantum long short-term memory (QLSTM) network models [28] for encoding and decoding representations in relation extraction tasks. These QLSTM models can better capture long-term dependencies in unstructured text data. In our proposed method, we utilize a hybrid quantum-classical method, which iteratively optimizes tasks applicable to relation extraction while harnessing the enhanced expressive power conferred by quantum superposition (as shown in Algorithm 1).

High-level QRL model for relationship detection

We employ the perspective of the HRL-RE approach to accomplish the high-level relation recognition mission [23]. In the high-level relations recognition component task, the sentence is scanned progressively, and the current high-level strategy \mathcal{O} ($\mathcal{O} \in NR \cup \mathcal{R}$) for the current time step is computed based on the state. Here, \mathcal{R} represents all the relationship types in the current dataset, and NR stands for “no relationship.”

State: $s_t^h \in \mathcal{S}$ of the high-level task at current step t is calculated as Eq. 1. It is calculated from the present hidden state h_t of the current time step t , the relationship type vector v_t^r of the latest $non - NR$ high-level strategy o' and the state s_{t-1} of the previous time step $t - 1$.

$$s_t^h = f^h \left(\mathbf{W}_s^h [\mathbf{q}h_t; v_t^r; \mathbf{s}_{t-1}] \right), \tag{1}$$

where $f^h(\cdot)$ denotes a non-linear transfer function, and \mathbf{W}_s^h denotes a weighting matrix. To get the hidden layer state h_t , we adopt a quantum Bi-LSTM over the present input word vectoring w_t :

$$\begin{aligned} \overrightarrow{\mathbf{q}h}_t &= \overrightarrow{QLSTM}(\overrightarrow{\mathbf{q}h}_{t-1}, \mathbf{w}_t) \\ \overleftarrow{\mathbf{q}h}_t &= \overleftarrow{QLSTM}(\overleftarrow{\mathbf{q}h}_{t+1}, \mathbf{w}_t) \end{aligned}$$

$$\mathbf{q}\mathbf{h}_t = [\overrightarrow{\mathbf{q}\mathbf{h}}_t, \overleftarrow{\mathbf{q}\mathbf{h}}_t]. \tag{2}$$

The mathematical expression of QLSTM is as follows:

$$\begin{aligned} g_t &= \sigma(VQC_1(\mathbf{w}_t)) \\ \tau_t &= \sigma(VQC_2(\mathbf{w}_t)) \\ \tilde{L}_t &= \tanh(VQC_3(\mathbf{w}_t)) \\ v_t &= g_t * v_{t-1} + \tau_t * \tilde{L}_t \\ k_t &= \sigma(VQC_4(s_t)) \\ z_t &= VQC_5(k_t * \tanh(v_t)) \\ y_t &= VQC_6(k_t * \tanh(v_t)), \end{aligned} \tag{3}$$

where \mathbf{w}_t represents the input at time t , z_t represents the hidden layer state, v_t represents the cell state, and y_t represents the output, $*$ represents element-wise multiplication (as shown in Fig. 1).

A random policy $\mu : S \rightarrow \mathcal{O}$ is employed, computed based on the current time step state \mathbf{s}_t^h through a softmax layer. The output of the softmax layer is stochastically sampled to get the o_t behavior of the present time step:

$$o_t \sim \mu(o_t | \mathbf{s}_t^h) = \text{softmax}(\mathbf{W}_\mu \mathbf{s}_t^h), \tag{4}$$

where \mathbf{W}_μ denotes a weighting matrix.

Reward: The experiment domain provides a return signal r_t^h to estimate the reward for performing policy o_t :

$$r_t^h = \begin{cases} -1, & \text{if } o_t \text{ not in } S \\ 0, & \text{if } o_t = NR \\ 1, & \text{if } o_t \text{ in } S. \end{cases} \tag{5}$$

Finally, the ultimate reward r_{fin}^h is gained to evaluate the sentence-level extraction manifestation that μ discovers:

$$r_{fin}^h = F_\beta(\mathcal{R}) = \frac{(1 + \beta^2)Prec \cdot Rec}{\beta^2 Prec \cdot Rec}, \tag{6}$$

where F_β denotes the weighted harmonic mean (WHM) of accuracy and recall rate in terms of the relationships in \mathcal{R} .

Low-level QRL model for entity extraction

We employ the perspective of the HRL-RE algorithm to construct the low-level entity identification mission [23]. In the low-level mission, the sentence is scanned line by line, and the action of the current time step is computed based on the state s_t^l and strategy π . If the high-level RL policy forecasts the NR (*non - NR*) relationship type, the low-level RL will extract entity information within the relationship. The high-level strategy o_t from high-level RL is an additional input parameter to the low-level RL mission.

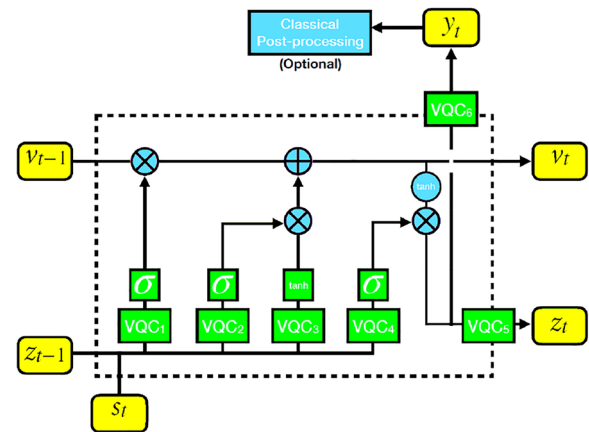


Fig. 1 The framework of quantum long short-term memory (QLSTM)

Action: this action assigns an entity tag to each word at every time step. The entity tags are represented as $\mathcal{A} = (\{\mathcal{S}, \mathcal{T}, \mathcal{O}\} \times \{\mathcal{B}, \mathcal{I}\} \cup \{\mathcal{U}\})$, where \mathcal{S} indicates the original entity, \mathcal{T} is the subjective entity, \mathcal{O} indicates the unrelated entity, \mathcal{N} denotes non-entity words, \mathcal{B} denotes the beginning of an entity and \mathcal{I} indicates internal parts of an entity.

State: the normative expression of the state s_t^l for the low-level task is as follows:

$$\begin{aligned} \mathbf{c}_t &= g(\mathbf{W}_h^l \mathbf{s}_t^h), \\ \mathbf{s}_t^l &= f^l(\mathbf{W}_s^l [\mathbf{q}\mathbf{h}_t; \mathbf{v}_t^e; \mathbf{s}_{t-1}; \mathbf{c}_t^l]), \end{aligned} \tag{7}$$

where $\mathbf{q}\mathbf{h}_t$ is the hidden state obtained from the Quantum Bi-LSTM module in Eq. (2), and $g(\cdot)$, $f^l(\cdot)$ are non-linear functions implemented by MLP. Low-level strategy uses a randomized strategy $\pi : S \rightarrow \mathcal{A}$ to stochastically sample the probabilities output from the softmax layer to obtain the action a_t of the current time step t .

$$a_t \sim \pi(a_t | \mathbf{s}_t^l; o_{t'}) = \text{softmax}(\mathbf{W}_\pi[o_{t'}] \mathbf{s}_t^l), \tag{8}$$

where \mathbf{W}_π denotes an array of relation matrices \mathcal{R} .

Reward: the reward r_t^l received by the present time step t is shown in Eq (8):

$$r_t^l = \lambda(y_t) \cdot \text{sgn}(a_t = y_t(o_{t'})), \tag{9}$$

where the immediate reward r_t^l is provided when the action a_t is sampled by simply the prediction error gold-standard annotation. The function $y_t(o_{t'})$ is the gold-standard entity tag conditioned on the predicted relationship type $o_{t'}$, $\lambda(y)$ is a weighting function for low-weight non-entity tag, denoted as follows:

$$\lambda(y) = \begin{cases} 1, & \text{if } y \neq N \\ \alpha, & \text{if } y = N. \end{cases} \quad (10)$$

The small α results to less reward on words that are not entities.

Quantum hierarchical strategy learning models

Similar to HRL-RE, QHRL-RE optimizes the strategy by maximizing the expected decay cumulative return:

$$J(\theta_{\mu,t}) = E_{s^h, o, r^h \sim \mu(o|s^h)} \left[\sum_{k=t}^T \gamma^{k-t} r_k^h \right], \quad (11)$$

where high-level strategy μ is parameterized by θ_{μ} , γ denotes the decay factor in RL.

Unlike HRL-RE, QHRL-RE calculates the expected discounted cumulative return for the low-level model using the following formula:

$$J(\theta_{\pi,t}; o'_t) = E_{s^l, a, r^l \sim \pi(a|s^l; o'_t)} \left(r_k^l + \left((1 - \epsilon) J(\theta_{\mu,t}) - \epsilon \max_{\mu \in \mathcal{O}} J(\theta_{\mu,t}) \right) \right), \quad (12)$$

where low-level strategy π is parameterized by θ_{π} , ϵ represents a hyperparameter, and $A(\theta_{\pi,t}; o'_t)$ represents the advantage function.

We decompose the expected decay cumulative rewards into a Bellman equation:

$$R^{\mu} \left(s_t^h, o_t \right) = E \left[\sum_{j=0}^{N-1} r_{t+j}^h \gamma^j R^{\mu} \left(s_{t+N}^h, o_{t+N} \right) | s_t^h, o_t \right], \quad (13)$$

$$R^{\pi} \left(s_t^l, a_t; o_t \right) = E \left[r_t^l + \gamma R^{\pi} \left(s_{t+1}^l, a_t; o_t \right) | s_t^l, o_t \right], \quad (14)$$

where N denotes the time steps of the entity identification component task started under the current high-level strategy o_t .

The gradient for the high-level policy is defined as follows:

$$\nabla_{\theta_{\mu}} J(\theta_{\mu,t}) = E_{s^h, o, r^h \sim \mu(o|s^h)} \left[R^{\mu} \left(s_t^h, o_t \right) \nabla_{\theta_{\mu}} \mu \left(o | s_t^h \right) \right]. \quad (15)$$

Unlike HRL-RE, QHRL-RE adopts the following Equation to update the gradient for the low-level policy:

$$\nabla_{\theta_{\pi}} J(\theta_{\pi,t}; o'_t) = E_{s^l, a, r^l \sim \pi(a|s^l; o'_t)} \left[R^{\pi} \left(s_t^l, a_t; o_t \right) \nabla_{\theta_{\pi}} \pi \left(a | s_t^l; o_t \right) A(\theta_{\pi,t}; o'_t) \right]. \quad (16)$$

Algorithm 1 Quantum Hierarchical Reinforcement Learning for Relation Extraction (QHRL-RE)

```

1: Compute  $h_t$  for each entity in the sentence by Quantum Bi-LSTM;
2: Initialize state  $s_0^h \leftarrow 0$  and time step  $t \leftarrow 0$ ;
3: for  $i \leftarrow 1$  to  $TextLength$  do
4:    $t \leftarrow t + 1$ ;
5:   Compute  $s_t^h$  through Equation (1);
6:   Sample  $o_t$  from  $s_t^h$  by Equation (4);
7:   Acquire reward  $r_t^h$  by Equation (5);
8:   if  $o_t \neq NR$  then
9:     for  $j \leftarrow 1$  to  $TextLength$  do
10:       $t \leftarrow t + 1$ ;
11:      Compute  $s_t^l$  through Equation (6);
12:      Sample  $a_t^l$  from  $s_t^l$  by Equation (7);
13:      Get reward  $r_t^l$  by Equation (8);
14:    end for
15:    Get low-level ultimate reward  $r_{fin}^l$ ;
16:  end if
17: end for
18: Get high-level reward  $r_{fin}^h$  by Equation (5);
19: Optimize the model through Equation (14) and Equation (15);

```

Experiments

Experimental setup

The dataset used in this article is the New York Times (NYT) corpus, sourced from distant supervision research, and it contains noisy relationship data [61, 62]. The corpus has two versions: (1) the traditional version generated by aligning the original data with Freebase relationships [61]. (2) A thin version of which the test set was manually annotated. We call the traditional version the NYT10 and the thin version the NTY11 [62].

Evaluation criterion: We evaluate the performance of this method using precision, recall, and micro F1 scores.

Baselines: We choose the following several entity extraction methods as the baselines.

FCM ([29]): A pipeline method that combines manually crafted lexicalized language context with word embeddings for entity and relationship extraction.

MultiR ([62]): A distant supervision approach that uses multiple weighted instances to handle noisy labels in training data.

CoType ([63]): A single approach that embeds entities, relationships, text features, and type labels into representations, treating the extraction task as a global embedding problem.

SPTree ([19]): A joint extraction approach that employs bidirectional sequential and bidirectional tree-structured LSTM-RNN in a single model to discover entities and relationships.

Tagging: A joint extraction approach that discovers entities and relationships using new labeling patterns.

CopyR ([64]): A Seq2Seq learning approach that utilizes multiple decoders to generate triplets for collectively extracting entities and relationships.

HRL-RE ([23]): A method based on HRL that breaks down the entire extraction task into high-level relations detection subtasks and low-level entity extraction subtasks.

Entities and relationships extraction

Table 1 presents the experimental results for relationship extraction. Since all methods were trained on noisy data, it is worth noting that there is a prominent difference in performance between the noisy dataset (NYT10) and the clean dataset (NYT11). It can be found that our algorithm QHRL-RE outperforms other entity relationship extraction approaches on both datasets. Crucially, the scores on the NYT10 dataset are much higher than those on the NYT11 dataset, indicating that the presented approach is more robust to noisy data.

Superimposed entities and relations extraction

We showcase the effectiveness of our method in discovering superimposed entities and relationships on two test sets: NYT11-plus and NYT10-sub. Here, we categorize superimposed entities and relationships into two types.

- Type 1: One entity participates in multiple relationships in the same sentence.
- Type 2: The identical entity pair in a sentence is associated with disparate relations.

Table 2 displays the manifestation of different entity and relationship extraction methods in extracting superimposed entities and relationships. The results of the experiment on the NYT10-sub dataset indicate that our method outperforms the HRL-RE approach. Furthermore, compared to our

Table 1 The experimental results for entity and relationship extraction

Model	NYT10			NYT11		
	Prec	Rec	F_1	Prec	Rec	F_1
FCM	–	–	–	0.431	0.292	0.348
MultiR	–	–	–	0.325	0.302	0.314
Cotype	–	–	–	0.483	0.382	0.427
SPTree	0.488	0.554	0.518	0.517	0.543	0.533
Tagging	0.588	0.376	0.459	0.471	0.483	0.472
CopyR	0.562	0.451	0.514	0.348	0.531	0.432
HRL-RE	0.716	0.581	0.641	0.528	0.536	0.525
QHRL-RE	0.736	0.624	0.673	0.586	0.579	0.602

Bold values highlight the maximum value under the same evaluation criteria

Table 2 Manifestation comparison for discovering superimposed entities and relationships

Model	NYT10-sub			NYT11-plus		
	Prec	Rec	F_1	Prec	Rec	F_1
FCM	–	–	–	0.233	0.198	0.221
MultiR	–	–	–	0.238	0.212	0.224
Cotype	–	–	–	0.288	0.251	0.269
SPTree	0.271	0.313	0.291	0.463	0.227	0.305
Tagging	0.253	0.234	0.242	0.291	0.217	0.244
CopyR	0.393	0.265	0.312	0.3297	0.225	0.263
HRL-RE	0.812	0.473	0.622	0.436	0.332	0.367
QHRL-RE	0.835	0.496	0.647	0.487	0.374	0.406

Bold values highlight the maximum value under the same evaluation criteria

QHRL-RE and HRL-RE methods, other relationship extraction methods could improve when dealing with noisy data in handling the 2nd class of superimposed entities and relationships. This suggests that traditional joint extraction methods could be more effective in solving the problem of superimposed entity relationships. Therefore, our method is better suited to address the 2nd class of superimposed entity relationship problems in noisy data. Additionally, experimental results on the NYT11-plus dataset show that our method outperforms other entity relationship extraction algorithms in extracting Type 1 superimposed entities and relationships in clean data. In a word, our algorithm can extract both superimposed entities and relations more effectively.

To verify the result of our method, a sample, "Arthur Lee, the leader of Love, was born in Memphis and lived there until 1952." stochastically chosen from the dataset is exhibited in Table 3. There are three categories of relationships in this example, and the corresponding triplets are \langle Arthur Lee, /person/location/place_birth, Memphis \rangle , \langle Arthur Lee, /person/location/place_lived, Memphis \rangle and \langle Arthur Lee, /person/leader_of/organization, Love \rangle . Our method detects three relations triumphantly, and the results are exhibited in Table 3. Take the first of these relations as an example. When the high-level relation detection subtask scans the sentence to "born in," it detects the word as a relationship indicator and identifies the relationship as "/person/location/place_birth." Then, the low-level subtask starts to scan the sentence. When it scans the word "Arthur Lee," it identifies this word as the source entity, and when it scans the word "Memphis," it identifies this word as the target entity.

Interaction between the two levels of component tasks

The results of the experiments in Table 4 demonstrate that our approach outperforms other relationship extrac-

Table 3 Example of relation detection

Relations type	Results of detection
/person/location /place_birth	[Arthur Lee] <i>source entity</i> , the leader of Love, was [born in] <i>relation indicator</i> [Memphis] <i>target entity</i> and lived there until 1952
/person/location /place_lived	[Arthur Lee] <i>source entity</i> , the leader of Love, was born in [Memphis] <i>target entity</i> and [lived there] <i>relation indicator</i> until 1952
/person/leader of /organization	[Arthur Lee] <i>source entity</i> , the [leader of] <i>relation indicator</i> [Love] <i>target entity</i> , was born in Memphis and lived there until 1952

Table 4 Comparison of experiment results for relationship prediction

Model	NTY11			NTY11-plus		
	Prec	Rec	F ₁	Prec	Rec	F ₁
FCM	0.501	0.475	0.487	0.441	0.326	0.376
MultiR	0.463	0.436	0.446	0.421	0.337	0.369
Cotype	0.556	0.554	0.553	0.485	0.412	0.446
SPTree	0.648	0.612	0.627	0.697	0.341	0.458
CopyR	0.476	0.712	0.573	0.625	0.423	0.503
HRL-RE-Env	0.673	0.673	0.673	0.574	0.322	0.411
HRL-RE	0.652	0.652	0.652	0.624	0.453	0.519
QHRL-RE-Env	0.687	0.636	0.691	0.576	0.312	0.408
QHRL-RE	0.655	0.647	0.661	0.632	0.472	0.553

Bold values highlight the maximum value under the same evaluation criteria

tion approaches in the relationship detection task on both datasets. Particularly, the improvement in the NYT11-plus dataset is more pronounced, indicating that our approach is better suited for discovering multiple relationships from sentences. Thus, embedding entities as relationship parameters in relationship detection can better leverage relationship information in the text.

The performance on the NYT11 dataset exhibits slight variations when the low-level policy is omitted separately from models HRL-RE-Ent and QHRL-RE-Ent. This is because nearly every sentence in this test set contains almost only one relationship. In such cases, the interaction between high-level and low-level component task policies has minimal impact on relationship detection results. In contrast, there is a significant difference in the NYT11-plus dataset, indicating that QHRL-RE and the hierarchical reinforcement learning-based QHRL-RE can capture dependencies between multiple extraction tasks. Furthermore, this interaction can increase the rewards for high-level component task policies. Thus, the entity and relationship extraction methods based on HRL intensify the interaction between relationship detection and entity identification.

Conclusion

This paper presents a new relations extraction method, quantum hierarchical reinforcement learning for relation extraction (QHRL-RE), which incorporates the quantum computing advantages of quantum entanglement and superposition into a hierarchical reinforcement learning relation extraction model. Specifically, drawing inspiration from the breakthroughs of quantum reinforcement learning in speech recognition and control domains, we employ quantum long short-term memory (QLSTM) network models for encoding and decoding representations in relation extraction tasks. These QLSTM models can better capture long-term dependencies in unstructured text data. Our proposed method utilizes a hybrid quantum-classical approach, which iteratively optimizes tasks applicable to relation extraction while harnessing the enhanced expressive power conferred by quantum superposition. In this way, our QHRL-RE approach is more effective for discovering superimposed entities and relations from unstructured text. Experiments on the commonly used datasets show that our method performs better than the selected baselines. As future work, our QHRL-RE method can be generalized to many other pairwise or triple-wise extraction tasks, such as aspect-opinion mining or ontology induction.

Acknowledgements This work is supported by the Research Foundation for Advanced Talents 31401529 of Henan University of Technology. This work is also supported by the Key Scientific Research Projects of Higher Education Institutions in Henan Province under Grant 24A520014, the Open Fund of Key Laboratory of Grain Information Processing and Control (Henan University of Technology), Ministry of Education (NO. KFJJ2023015). This work is partly supported by the National Natural Science Foundation of China under Grants (No. 62006071) and the Science and Technology Project of the Science and Technology Department of Henan Province (No. 212102210149). This work is also supported by the College Students' Innovative Entrepreneurial Training Plan Program PX-38233887, PX-38233886, and PX-38233884, respectively. The authors would like to thank Ruiyuan Zhang for helpful discussions on topics related to this work.

Data availability In our paper, data will be made available on request.

Declarations

Conflict of interest The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Mintz M, Bills S, Snow R, Jurafsky D (2009) Distant supervision for relation extraction without labeled data. In: *ACL 2009, Proceedings of the 47th Annual Meeting of the Association for Computational Linguistics and the 4th International Joint Conference on Natural Language Processing of the AFNLP*, 2-7 August, Singapore, The Association for Computer Linguistics, pp. 1003–1011
- Nadeau D, Sekine S (2007) A survey of named entity recognition and classification. *Linguisticae Investigationes* 30(1):3–26
- de Jesús Rubio J, Hernandez MA, Rosas FJ, Orozco E, Balcazar R, Pacheco J (2024) Genetic high-gain controller to improve the position perturbation attenuation and compact high-gain controller to improve the velocity perturbation attenuation in inverted pendulums. *Neural Netw* 170:32–45
- Chiang H-S, Chen M-Y, Huang Y-J (2019) Wavelet-based eeg processing for epilepsy detection using fuzzy entropy and associative petri net. *IEEE Access* 7:103255–103262
- López-González A, Campaña JM, Martínez EH, Contro PP (2020) Multi robot distance based formation using parallel genetic algorithm. *Appl Soft Comput* 86:105929
- Fader A, Zettlemoyer L, Etzioni O (2014) Open question answering over curated and extracted knowledge bases. In: *The 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '14*, New York, NY, USA - August 24 - 27, ACM, pp. 1156–1165
- Glavaš G, Šnajder J (2014) Event graphs for information retrieval and multi-document summarization. *Expert Syst Appl* 41(15):6904–6916
- de Jesús Rubio J (2023) Bat algorithm based control to decrease the control energy consumption and modified bat algorithm based control to increase the trajectory tracking accuracy in robots. *Neural Netw* 161:437–448
- de Jesús Rubio J, Garcia D, Sossa H, Garcia I, Zacarias A, Mujica-Vargas D (2023) Energy processes prediction by a convolutional radial basis function network. *Energy* 284:128470
- Luan Y, He L, Ostendorf M, Hajishirzi H (2018) Multi-task identification of entities, relations, and coreference for scientific knowledge graph construction. In: *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, Brussels, Belgium, October 31 - November 4, The Association for Computational Linguistics, pp. 3219–3232
- Mújica-Vargas D (2021) Superpixels extraction by an intuitionistic fuzzy clustering algorithm. *J Appl Res Technol* 19(2):140–152
- Huang C-C, Lu Z (2016) Community challenges in biomedical text mining over 10 years: success, failure and the future. *Brief Bioinformatics* 17(1):132–144
- Fei H, Ren Y, Zhang Y, Ji D, Liang X (2021) Enriching contextualized language model from knowledge graph for biomedical information extraction. *Briefings in Bioinformatics* 22(3)
- Zhao W, Zhao Y, Jiang X, He T, Liu F, Li N (2021) Efficient multiple biomedical events extraction via reinforcement learning. *Bioinformatics* 37(13):1891–1899
- Nadeau D, Sekine S (2007) A survey of named entity recognition and classification. *Linguisticae Investigationes* 30(1):3–26
- Li F, Zhang M, Fu G, Ji D (2017) A neural joint model for entity and relation extraction from biomedical text. *BMC Bioinformatics* 18(1):198:1-198:11
- Wang S, Zhang Y, Che W, Liu T (2018) Joint extraction of entities and relations based on a novel graph scheme. In: *Lang J (Ed.), Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018*, July 13-19, Stockholm, Sweden, Morgan Kaufmann, pp. 4461–4467
- Miwa M, Sasaki Y (2014) Modeling joint entity and relation extraction with table representation. In: *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014*, October 25-29, Doha, Qatar, A meeting of SIGDAT, a Special Interest Group of the ACL, The Association for Computational Linguistics, pp. 1858–1869
- Miwa M, Bansal M (2016) End-to-end relation extraction using lstms on sequences and tree structures. In: *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016*, August 7-12, Berlin, Germany, Volume 1: Long Papers, The Association for Computer Linguistics
- Zheng S, Wang F, Bao H, Hao Y, Zhou P, Xu B (2017) Joint extraction of entities and relations based on a novel tagging scheme. In: *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017*, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers, The Association for Computational Linguistics, pp. 1227–1236
- Katiyar A, Cardie C (2017) Going out on a limb: Joint extraction of entity mentions and relations without dependency trees. In: *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017*, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers, Association for Computational Linguistics, pp. 917–928
- Huang P, Zhao X, Takanobu R, Tan Z, Xiao W (2020) Joint event extraction with hierarchical policy network. In: *Proceedings of the 28th international conference on computational linguistics*, pp. 2653–2664
- Takanobu R, Zhang T, Liu J, Huang M (2019) A hierarchical framework for relation extraction with reinforcement learning. In: *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019*, Honolulu, Hawaii, USA, January 27 - February 1, AAAI Press, pp. 7072–7079
- Yang C-HH, Qi J, Chen SY-C, Chen P-Y, Siniscalchi SM, Ma X, Lee C-H (2021) Decentralizing feature extraction with quantum convolutional neural network for automatic speech recognition. In: *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE, pp. 6523–6527
- Di Sipio R, Huang J-H, Chen SY-C, Mangini S, Worring M (2022) The dawn of quantum natural language processing. In: *ICASSP*

- 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), IEEE, pp. 8612–8616
26. Metz F, Bukov M (2023) Self-correcting quantum many-body control using reinforcement learning with tensor networks. *Nat Mach Intell* 5(7):780–791
 27. Ma H, Dong D, Ding SX, Chen C (2022) Curriculum-based deep reinforcement learning for quantum control. *IEEE Transactions on Neural Networks and Learning Systems*
 28. Chen SY-C, Yoo S, Fang Y-LL (2022) Quantum long short-term memory. In: *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE, pp. 8622–8626
 29. Gormley MR, Yu M, Dredze M (2015) Improved relation extraction with feature-rich compositional embedding models. In: *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing, EMNLP 2015*, Lisbon, Portugal, September 17–21, The Association for Computational Linguistics, pp. 1774–1784
 30. Tang J, Qu M, Wang M, Zhang M, Yan J, Mei Q (2015) Line: Large-scale information network embedding. In: *Proceedings of the 24th International Conference on World Wide Web, WWW 2015*, Florence, Italy, May 18–22, ACM, pp. 1067–1077
 31. Zhou P, Shi W, Tian J, Qi Z, Li B, Hao H, Xu B (2016) Attention-based bidirectional long short-term memory networks for relation classification. In: *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016*, August 7–12, Berlin, Germany, Volume 2: Short Papers, The Association for Computer Linguistics
 32. Tuo M, Yang W (2023) Review of entity relation extraction. *Journal of Intelligent & Fuzzy Systems (Preprint)* 1–15
 33. Zhou Q, Zhang Y, Ji D (2023) Distantly supervised relation extraction with kb-enhanced reconstructed latent iterative graph networks. *Knowl Based Syst* 260:110108
 34. Javeed A (2023) A hybrid attention mechanism for multi-target entity relation extraction using graph neural networks. *Mach Learn Appl* 11:100444
 35. Li Q, Ji H (2014) Incremental joint extraction of entity mentions and relations. In: *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics, ACL 2014*, June 22–27, Baltimore, MD, USA, Volume 1: Long Papers, The Association for Computational Linguistics, pp. 402–412
 36. Björne J, Heimonen J, Ginter F, Airola A, Pahikkala T, Salakoski T (2011) Extracting contextualized complex biological events with rich graph-based feature sets. *Comput Intell* 27(4):541–557
 37. Narasimhan K, Yala A, Barzilay R (2016) Improving information extraction by acquiring external evidence with reinforcement learning. In: *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016*, Austin, Texas, USA, November 1–4, The Association for Computational Linguistics, pp. 2355–2365
 38. Katiyar A, Cardie C (2016) Investigating lstms for joint extraction of opinion entities and relations. In: *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016*, August 7–12, Berlin, Germany, Volume 1: Long Papers, The Association for Computational Linguistics, pp. 919–929
 39. Zhang M, Zhang Y, Fu G (2017) end-to-end neural relation extraction with global optimization. In: *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP 2017*, Copenhagen, Denmark, September 9–11, The Association for Computational Linguistics, pp. 1730–1740
 40. Qin P, Xu W, Wang WY (2018) Robust distant supervision relation extraction via deep reinforcement learning. In: *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018*, Melbourne, Australia, July 15–20, Volume 1: Long Papers, The Association for Computational Linguistics, pp. 2137–2147
 41. Feng Y, Zhang H, Hao W, Chen G (2017) Joint extraction of entities and relations using reinforcement learning and deep learning. *Comput Intell Neurosci* 7643065(1–7643065):11
 42. Feng J, Huang M, Zhao L, Yang Y, Zhu X (2018) Reinforcement learning for relation classification from noisy data. In: *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI-18)*, the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2–7, Vol. 32, AAAI Press
 43. Dong D, Chen C, Li H, Tam T-J (2008) Quantum reinforcement learning. *IEEE Trans Syst Man Cybern* 38(5):1207–1220
 44. Chen SY-C, Yang C-HH, Qi J, Chen P-Y, Ma X, Goan H-S (2020) Variational quantum circuits for deep reinforcement learning. *IEEE Access* 8:141007–141024
 45. Lockwood O, Si M (2020) Reinforcement learning with quantum variational circuit. In: *Proceedings of the AAAI conference on artificial intelligence and interactive digital entertainment*, Vol. 16, pp. 245–251
 46. Skolik A, Jerbi S, Dunjko V (2022) Quantum agents in the gym: a variational quantum algorithm for deep q-learning. *Quantum* 6:720
 47. Schenk M, Combarro EF, Grossi M, Kain V, Li KSB, Popa M-M, Vallecorsa S (2022) Hybrid actor-critic algorithm for quantum reinforcement learning at cern beam lines, arXiv preprint [arXiv:2209.11044](https://arxiv.org/abs/2209.11044)
 48. Lan Q (2021) Variational quantum soft actor-critic, arXiv preprint [arXiv:2112.11921](https://arxiv.org/abs/2112.11921)
 49. Qiu Y, Liu R, Lee RS (2024) The design and implementation of a deep reinforcement learning and quantum finance theory-inspired portfolio investment management system. *Expert Syst Appl* 238:122243
 50. Heimann D, Hohenfeld H, Wiebe F, Kirchner F (2022) Quantum deep reinforcement learning for robot navigation tasks, arXiv preprint [arXiv:2202.12180](https://arxiv.org/abs/2202.12180)
 51. Jerbi S, Gyurik C, Marshall S, Briegel H, Dunjko V (2021) Parametrized quantum policies for reinforcement learning. *Adv Neural Inform Process Syst* 34:28362–28375
 52. Hsiao J-Y, Du Y, Chiang W-Y, Hsieh M-H, Goan H-S (2022) Unentangled quantum reinforcement learning agents in the openai gym, arXiv preprint [arXiv:2203.14348](https://arxiv.org/abs/2203.14348)
 53. Sutton RS, Barto AG (2018) *Reinforcement learning: An introduction*. MIT press
 54. Vezhnevets AS, Osindero S, Schaul T, Heess N, Jaderberg M, Silver D, Kavukcuoglu K (2017) Feudal networks for hierarchical reinforcement learning. In: *Proceedings of the 34th International Conference on Machine Learning, ICML 2017*, Sydney, NSW, Australia, 6–11 August, Vol. 70, PMLR, pp. 3540–3549
 55. Mohri M, Rostamizadeh A, Talwalkar A (2018) *Foundations of machine learning*. MIT press
 56. Sutton RS, Precup D, Singh S (1999) Between mdps and semi-mdps: A framework for temporal abstraction in reinforcement learning. *Artif Intell* 112(1–2):181–211
 57. Li R, Cai Z, Huang T, Zhu W (2021) Anchor: The achieved goal to replace the subgoal for hierarchical reinforcement learning. *Knowl Based Syst* 225:107128
 58. Zhu X, Zhang R, Zhu W (2022) Mdm options discovery for accelerating exploration in sparse-reward domains. *Knowl Based Syst* 241:108151
 59. Ou W, Luo B, Wang B, Zhao Y (2024) Modular hierarchical reinforcement learning for multi-destination navigation in hybrid crowds. *Neural Netw* 171:474–484
 60. Luo J, Xu C, Geng X, Feng G, Fang K, Tan L, Schaal S, Levine S (2024) Multi-stage cable routing through hierarchical imitation learning. *IEEE Transactions on Robotics*

61. Riedel S, Yao L, McCallum A (2010) Modeling relations and their mentions without labeled text. In: Machine Learning and Knowledge Discovery in Databases, European Conference, ECML PKDD 2010, Barcelona, Spain, September 20-24, 2010, Proceedings, Part III, Vol. 6323, Springer, pp. 148–163
62. Hoffmann R, Zhang C, Ling X, Zettlemoyer LS, Weld DS (2011) Knowledge-based weak supervision for information extraction of overlapping relations. In: The 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies, Proceedings of the Conference, 19-24 June, Portland, Oregon, USA, The Association for Computer Linguistics, pp. 541–550
63. Ren X, Wu Z, He W, Qu M, Voss CR, Ji H, Abdelzaher TF, Han J (2017) Cotype: Joint extraction of typed entities and relations with knowledge bases. In: Proceedings of the 26th International Conference on World Wide Web, WWW 2017, Perth, Australia, April 3-7, ACM, pp. 1015–1024
64. Zeng X, Zeng D, He S, Liu K, Zhao J (2018) Extracting relational facts by an end-to-end neural model with copy mechanism. In: Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, Volume 1: Long Papers, Association for Computational Linguistics, pp. 506–514

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.