

Learning to Select the Next Reasonable Mention for Entity Linking

Jian Sun,^{1,2} Yu Zhou,^{1,2,3} Chengqing Zong,^{1,2}

¹ School of Artificial Intelligence, University of Chinese Academy of Sciences,
Beijing 100049, P.R.China

² National Laboratory of Pattern Recognition, Institute of Automation, CAS

³ Beijing Fanyu Technology Co., Ltd, Beijing, China

jian.sun@nlpr.ia.ac.cn, yzhou@nlpr.ia.ac.cn, cqzong@nlpr.ia.ac.cn

Abstract

Entity linking aims to establish a link between entity mentions in a document and the corresponding entities in knowledge graphs (KGs). Previous work has shown the effectiveness of global coherence for entity linking. However, most of the existing global linking methods based on sequential decisions focus on how to utilize previously linked entities to enhance the later decisions. In those methods, the order of mention is fixed, making the model unable to adjust the subsequent linking targets according to the previously linked results, which will cause the previous information to be unreasonably utilized. To address the problem, we propose a novel model, called DyMen, to dynamically adjust the subsequent linking target based on the previously linked entities via reinforcement learning, enabling the model to select a link target that can fully use previously linked information. We sample mention by sliding window to reduce the action sampling space of reinforcement learning and maintain the semantic coherence of mention. Experiments conducted on several benchmark datasets have shown the effectiveness of the proposed model.

Introduction

As a critical step in the task of natural language understanding, entity linking (EL) has been widely used in many fields, such as question answering (Bordes et al. 2015), knowledge fusion (Yan et al. 2020), and information extraction (Hoffmann et al. 2011). This task is challenging because mentions are usually ambiguous, and it is difficult to link with the entities in the knowledge base only from the surface of mention.

To alleviate this problem, various EL models have been proposed (Deng et al. 2020; Yang et al. 2019; Zhang et al. 2020; Ganea and Hofmann 2017; Le and Titov 2018; Cao et al. 2018), which can be divided into local model and global model. The local model mainly focuses on contextual words around the mention, and the global model concerns the topic coherence. Some global models (Ganea and Hofmann 2017; Le and Titov 2018) calculate the pairwise coherence between all candidate entities, which suffers from high computational complexity and introduces some noises (Yang et al. 2019). Therefore, some studies convert the

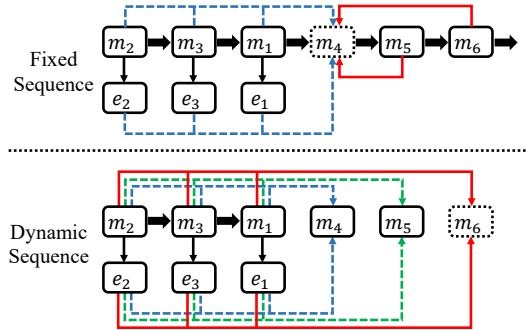


Figure 1: A comparative example of fixed sequence and dynamic sequence.

global linking into a sequential decisions problem (Yang et al. 2019; Fang et al. 2019; Gu et al. 2021), which use previous information to disambiguate subsequent mentions. These studies use the natural order in the original document, the similarity between mentions, or the difficulty of entity linking to sort mentions and get a fixed sequence, on which a reinforcement learning algorithm or supervised learning method is used to link the entity. However, the above fixed sequence is problematic. For example, in Figure 1, m_i is a mention in the document, and e_i is a candidate entity that has already been linked to m_i . m_4 , m_5 , and m_6 are the mentions to be linked. In the fixed sequence, m_4 is the next target to be linked, but sometimes m_1 , m_2 , and m_3 do not provide enough critical information to disambiguate m_4 . On the contrary, the critical information used to disambiguate m_4 may be m_5 or m_6 . Meanwhile, although m_1 , m_2 , and m_3 cannot be used to disambiguate m_4 , they may be enough to disambiguate m_6 . Obviously, the methods based on fixed sequence can not deal with the above situation.

In this paper, to solve the above problems, we propose a global entity linking model where mention can be selected dynamically based on the results of previously linked entities. Instead of deciding which key mention can be used to disambiguate the current mention, we use the previously linked mention and entity information to select a target that can make the most of the previous information. For example, in the dynamic sequence in Figure 1, m_2 , m_3 , and m_1 are the mentions that have been linked. Instead of choosing

which mentions can be used to disambiguate m_4 , we choose those that can make full use of m_2 , m_3 , and m_1 for disambiguation, such as m_6 . We utilize a policy network to choose a linking target that can fully use previously linked information. Through a sliding window, the model can sample a fixed number of mentions, which is beneficial to reducing the action search space of reinforcement learning and maintaining a certain semantic coherence between the mentions. To better explore the importance of the order of mention, we optimize our model from three different perspectives. We conduct experiments on different datasets, and the results show that our model outperforms state-of-the-art systems.

Related Work

Entity linking has gained increasing attention in recent years, which can be divided into two major models: local and global models. The local model utilizes the information around mention to disambiguate mentions. However, sometimes the context information is sparse, and it is difficult to make use of this sparse information to complete entity disambiguation (Yang et al. 2019). Because of this, many global disambiguation models are proposed from the perspective of topic coherence. Some studies (Ganea and Hofmann 2017; Le and Titov 2018; Chen et al. 2020) assume that the target entities of all mentions in a document shall be related and utilize fully-connected pairwise Conditional Random Field model and exploit loopy belief propagation to estimate the max-marginal probability. This assumption is unreasonable. Sometimes, some entities in a document are unrelated to each other. With the popularity of graph neural networks, some studies (Wu et al. 2020; Cao et al. 2018) apply graph neural networks (GNN) to encode mention-entity graphs. These approaches need to build a subgraph for all candidate entities of mentions and apply graph neural networks to integrate both local contextual features and global coherence information for entity linking. A high-quality subgraph is a key to the linking method based on GNN. Some studies (Yang et al. 2019; Fang et al. 2019; Gu et al. 2021) convert the global linking into a sequential decisions problem. The core idea is to use previously linked entities to enhance the future mention disambiguation. The limitation of these methods is that the order of mention sequence is fixed while selecting the entity, and the reasonable link objects can not be selected dynamically according to the previous information, which will cause the previous information can not be used reasonably. Moreover, it can easily lead to error accumulation. In an extreme case, most of the previously linked results are incorrect, and these incorrect entities may negatively affect future link steps. Intuitively, the order of mention is related to the performance of subsequent entity linking and is the basis of entity linking.

Methodology

Task Definition

Formally, given a set of entity mentions $M = \{m_1, m_2, \dots, m_t\}$ in a document D , each mention $m_t \in D$ has a set of candidate entities $C_t = \{e_t^1, e_t^2, \dots, e_t^n\}$. Entity

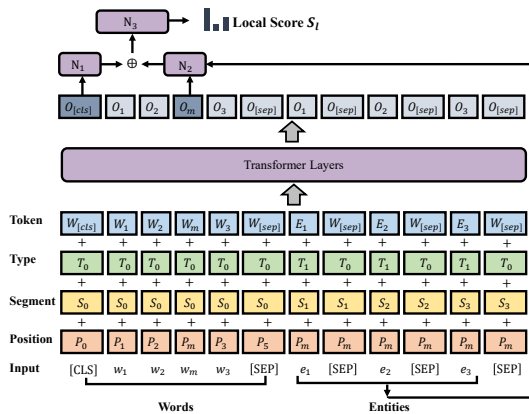


Figure 2: The architecture of local entity linking model. w_m is the mention and w_i is the word in context surrounding mention w_m .

linking maps each mention m_t to the correct entity. Our entity linking model consists of two models: local model and global model. The local model only utilizes the context of mentions to disambiguate mentions. In contrast, the global model uses previously linked mentions to disambiguate subsequent mentions according to the topic consistency between mentions.

Local Model

In this section, we apply two different local models to the entity linking model. The first is the local model in (Ganea and Hofmann 2017), which measures the relevance of entity candidates of each mention independently. The local score function is as follows:

$$\Psi(e_t^i, \hat{w}_t) = \mathbf{e}_t^{i\top} \mathbf{B}_1 f(\hat{w}_t) \quad (1)$$

where \mathbf{e}_t^i is the embedding of entity e_t^i , \mathbf{B}_1 is a diagonal matrix and $f(\hat{w}_t)$ is the feature representation of local context surrounding m_t , \hat{w}_t is the set of local context words.

We also design a new local entity linking model based on Transformer (Vaswani et al. 2017) encoder, as shown in Figure 2, which utilizes the mechanism of attention. Inspired by the popular pre-train language model such as BERT (Devlin et al. 2018), we introduce two special symbols [CLS] and [SEP] to construct the input data. The special symbol [CLS] is utilized to encode the information of the entire sequence by self-attention, and the symbol [SEP] is used to distinguish between words and entities. Similar to the methods in BERT, we concatenate local context w_t and candidate entities $\{e_t^1, e_t^2, \dots, e_t^n\}$ with [CLS] and [SEP] tokens as the input sequence in the following format:

$$[\text{CLS}] w_1, w_2, \dots, w_t [\text{SEP}] e_1 [\text{SEP}], \dots, e_n [\text{SEP}] \quad (2)$$

where $w_t \in \hat{w}_t$ is the word in the local context, and $e_i \in \{e_t^1, e_t^2, \dots, e_t^n\}$ is the candidate entity. The input representation is constructed by summing up the following four parts:

(1) **Token embedding** represents the embedding of the word or entity. We represent the embedding of word and entity as W and E , respectively. Here, $E_i = (w_{e_i} + \mathbf{B}_2 \mathbf{e}_i)/2$,

w_{e_i} is the surface form of the entity and the \mathbf{B}_2 is to migrate entities from entity space to word space. (2) **Type embedding** T_k is used to distinguish the type of token. There are two types, words type T_0 and entity type T_1 . (3) **Segment embedding** S_l is used to distinguish different candidate entities and contexts, where $l \in \{0, 1, \dots, n\}$. (4) **Position embedding** P represents the position of the token in context. It is worth noting that there is no clear sequence relationship between different candidate entities. So, we use the position embedding of mention in context to represent the position embedding of all candidate entities. In this way, candidate entities can capture the same position information as mention. We feed the above four parts into Transformer layers and get the output of words. We use \mathbf{N}_3 to get the result.

$$\begin{aligned} \mathbf{N}_1 &= \sigma(f_2(\text{Drop}(\text{ReLU}(f_1(\mathbf{o}_{cls})))))) \\ \mathbf{N}_2 &= [\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_n] \mathbf{o}_m^\top \\ \mathbf{N}_3 &= \sigma(f_4(\text{Drop}(\text{ReLU}(f_3([\mathbf{N}_1; \mathbf{N}_2]))))) \end{aligned} \quad (3)$$

where \mathbf{o}_{cls} is the output representation of [CLS]. f_i is the fully connected layer. σ is softmax function. $\text{Drop}(\cdot)$ is the dropout operation (Srivastava et al. 2014). $\text{ReLU}(\cdot)$ is the rectifier nonlinearity (Nair and Hinton 2010). \mathbf{N}_1 , \mathbf{N}_2 , and $\mathbf{N}_3 \in \mathbb{R}^{1 \times n}$. Finally, we get the local score $\Psi_m = \mathbf{N}_3$.

Global Model

Our global entity linking model consists of the dynamic mention selection model and the candidate entity selection model. The dynamic mention selection model utilizes the policy networks to select the corresponding mention, and the candidate entity selection model maps the mention to the correct target entity. The architecture of the global model is shown in Figure 3.

Dynamic Mention Selection The basic idea of dynamic mention selection is to use information (state) about the previously selected mention and the linked entities to select a reasonable mention from unresolved mentions. We consider the dynamic mention selection model as an agent. The agent utilizes the policy network to learn a stochastic policy and select an action (mention) according to the policy at each state. The stochastic policy can ensure the agent can explore the action space. Now, we will introduce the detail of how to select the target mention.

States. We treat the previously selected mentions and entities as state information. Specifically, each state $S_t = \{[m_0; e_0], [m_1; e_1], \dots, [m_t; e_t]\}$ is a set of mentions and entities, where $[m_1; e_1]$ represent the concatenation of the resolved mention m_1 and linked entity e_1 . Note m_1 indicates the first mention selected according to the dynamic selection model, not the first mention appearing in the text. m_0 and e_0 are initial states.

Actions. The set A_t of possible actions a_t at step t consists of those unresolved mentions. Initially, this set consists of all the mentions, making it difficult for agents to select an appropriate action. Meanwhile, choosing the next action from the above set of actions will greatly disarrange the original order of mentions in the document. We believe the original order of mentions is beneficial to understanding mention and

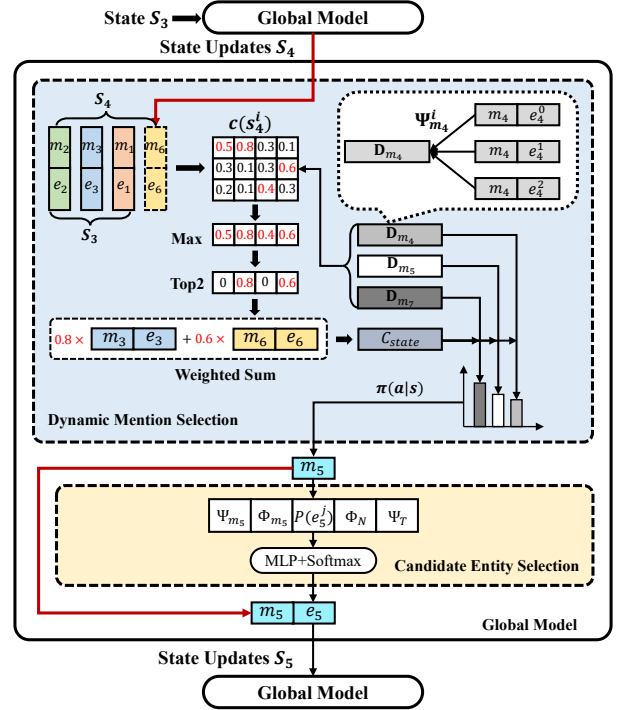


Figure 3: The architecture of global model.

subsequent entity linking. Therefore, we design a local fine-tuned action space selection method, which considers maintaining contextual coherence and reducing action space. The method is shown in Figure 4, the subscript k of a_k refers to the position of mention in the text.

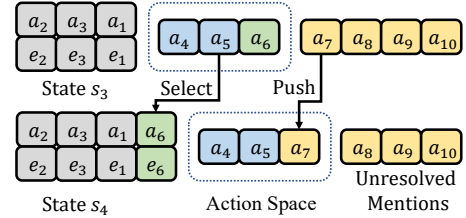


Figure 4: An illustration of local fine-tuned action space selection method. The main idea is to use a fixed size sliding window to constantly update the action space.

Policy Network. After defining the state and action, we can select an action from the action space through a policy network $\pi(a|s)$. We apply a neural attention mechanism to the policy network. Formally, we calculate the relevance score of each element in the state to the action space.

$$c(s_t^i) = \max_{a_t \in A_t} \mathbf{D}_{a_t}^\top \mathbf{B}_3 s_t^i \quad (4)$$

$$\mathbf{D}_{a_t} = \sum_{j=1}^n \Psi_{a_t}^j [a_t; \mathbf{e}_t^j] \quad (5)$$

where s_t^i is the $i + 1$ -th element in the state S_t . \mathbf{D}_{a_t} is the concatenation of the mention and the candidate entities. \mathbf{e}_t^j

refers to j -th candidate entity that corresponds to mention a_t . $\Psi_{a_t}^j$ is the local score of j -th candidate entity. \mathbf{B}_3 is a parameterized diagonal matrix. We choose Top K elements in S_t as evidence to select the subsequent target mention.

$$c(\hat{s}_t) = \text{TopK}\{c(s_t^0), c(s_t^1), \dots, c(s_t^t)\} \quad (6)$$

Then, the attention weights are as follows

$$w(\hat{s}_t^j) = \frac{\exp[c(\hat{s}_t^j)]}{\sum_{\hat{s}_t^k \in \hat{s}_t} \exp[c(\hat{s}_t^k)]} \quad (7)$$

Finally, we get the probability distribution of each action.

$$\pi_\theta(a_t|S_t) = \text{Softmax}\left(\sum_{\hat{s}_t^j \in \hat{s}_t} w(\hat{s}_t^j) \mathbf{D}_{a_t}^\top \mathbf{B}_4 \hat{s}_t^j\right) \quad (8)$$

where \mathbf{B}_4 is a learnable diagonal matrix. Next, we need to design the reward function to guide the learning of policy for mention prediction.

Rewards. In our framework, since one mention selection has a long-term impact on subsequent entity linking, we use a delayed reward to improve the overall performance. In this paper, we will design the reward functions from different perspectives to explore the influence of the order of mention on the subsequent entity selection. When designing the reward function, we follow the principle that if the current mention is linked to the correct entity, we think that the previous information is enough to disambiguate the current mention; if not, we think that the previous information is not enough to disambiguate the current motion, and we need to choose a new mention to disambiguate. In the end, we will get an **ideal sequence**: all the mentions in the front of the sequence are correctly linked, and there are a few incorrect linked mentions in the back of the sequence, which need other information to disambiguate.

$$\begin{aligned} s_1: & 1, \mathbf{0}, 1, \mathbf{0}, 1, 1, 1 \\ s_2: & 1, 1, 1, 1, \mathbf{0}, \mathbf{0}, \mathbf{0} \\ s_3: & 1, 1, \mathbf{0}, \mathbf{0}, 1, 1, 1 \end{aligned}$$

Figure 5: An example of the reward functions. s_1 , s_2 , and s_3 are different predictions for the same sequence, where ‘1’ means the link is correct and ‘0’ means the link is wrong.

(1) The first reward function R1 is as follows,

$$R1(t) = \frac{\gamma^{L-t}}{L} (-L + I_{\text{first_error}}) \quad (9)$$

where γ is the discount factor. L is the number of mentions, and $I_{\text{first_error}}$ is the index of the first mention with the incorrect linking in the mention sequence. For example, in Figure 5, we can calculate $R1(T) = (-L + I_{\text{first_error}})$, and we get $R1(T)_{s_1} = -6$, $R1(T)_{s_2} = -3$, and $R1(T)_{s_3} = -5$. Finally, the rewards for different moments are $R1(t)_{s_2} > R1(t)_{s_3} > R1(t)_{s_1}$. The reward function can guide the policy network to give priority to learning those easy to be linked mentions.

(2) The second reward function R2 is designed by adjusting the transition reward. We can get the transition reward

$\lambda_{\tau-1, \tau}$ between mentions based on the result of entity linking. For instance, we have a mention sequence $\{m_0, m_1\}$ where the mention has been linked. We select a new mention m_2 through the policy network and link it to an entity. If the linking is correct, there will be a large transition reward between m_1 and m_2 . If the linking is wrong, we will get a small transition reward.

$$R2(t) = \frac{\gamma^{L-t}}{L} \sum_{\tau=1}^L \lambda_{\tau-1, \tau} \quad (10)$$

where $\lambda_{\tau-1, \tau}$ is the transition reward between action $a_{\tau-1}$ and a_τ , where $\lambda_{\tau-1, \tau} \in \{\lambda_{T, T}, \lambda_{T, F}, \lambda_{F, F}, \lambda_{F, T}\}$. $\lambda_{T, F}$ refers to the transition reward of situations where the previous entity prediction is correct, and the next prediction is wrong. ‘T’ and ‘F’ are True and False, respectively. For example, for s_2 in Figure 5, we get $R2(t)_{s_2} = \frac{\gamma^{L-t}}{L} (3\lambda_{T, T} + 1\lambda_{T, F} + 2\lambda_{F, F})$. This reward function reflects the dynamic selection effect as a whole. If there are more wrong links, the reward function will be smaller, indicating that some mention may not make use of reasonable information.

(3) We also consider both the number of predictions wrong and their position in the sequence to design the reward function R3.

$$R3(t) = \frac{\gamma^{L-t}}{L} \sum_{d \in \text{error}} -1 + (I_d - L)/L \quad (11)$$

where I_d is the index of mention that has the incorrect prediction in the sequence. For the sequences in Figure 5, we get $R3(T)_{s_1} = -24/7$, $R3(T)_{s_2} = -27/7$, and $R3(T)_{s_3} = -23/7$. Finally, $R3(t)_{s_3} > R3(t)_{s_1} > R3(t)_{s_2}$. This reward function takes link accuracy as the primary optimization objective ($R3(t)_{s_3} > R3(t)_{s_2}$ and $R3(t)_{s_1} > R3(t)_{s_2}$). In the case of the same accuracy, the model is optimized in the direction of constructing the **ideal sequence** ($R3(t)_{s_3} > R3(t)_{s_1}$).

Candidate Entity Selection Our dynamic mention selection methods can be transplanted to all entity linking models based on sequential decisions as the basis for subsequent candidate entity selection modules. Here, we adopt the candidate entity selection method in Yang et al. (2019), which is also a neural attention method.

$$\Phi(e_i, \hat{C}_i) = \mathbf{e}_i^\top \mathbf{B}_5 f(\hat{C}_i) \quad (12)$$

where \hat{C}_i denote the list of previously linked entities $\{e_1, e_2, \dots, e_t\}$, $f(\cdot)$ is a neural attention function. We first calculate the similarity between each candidate entity and a linked entity e_i in \hat{C}_i , and take a maximum value as the score of e_i . Then we sort the entities in \hat{C}_i according to their scores and select the Top K entities. Finally, the embedding of these words is weighted and summed according to their scores to get $f(\hat{C}_i)$. We also take mention-entity prior $P(e_t^j | m_t)$ (Ganea and Hofmann 2017; Le and Titov 2018; Yang et al. 2019), the mention-entity type characteristics $\Psi_T(m_t, e_t^j)$ (Yang et al. 2019), and the characteristics $\Phi(e_{t+1}, \mathcal{N}(\hat{C}_t))$ (Yang et al. 2019) of the neighborhood entities mentioned as the basis for our entity linking, where

$P(e_t^j|m_t)$ is the probability of selection e_t^j conditioned on m_i , $\mathcal{N}(\hat{C}_t)$ is the neighborhood set of all entities in \hat{C}_t in the knowledge graph. We concatenate the above four features and local feature Ψ_{m_t} as the final feature and input the final feature into a feed-forward network to get the final probability of candidate entities $\hat{P}_\theta(e_t^i|m_t)$.

Model Learning

In the dynamic mention selection module, we get the $R(t)$, which is the expected reward at time step t . Our goal is to maximize the expected reward over the mention sequence:

$$J_{ms}(\theta) = \sum_t \sum_{a_t} \pi_\theta(a_t|S_t)R(t) \quad (13)$$

We use the REINFORCE algorithm (Sutton and Barto 2011) to optimize the objective function $J(\theta)$.

$$\theta \leftarrow \theta + \alpha \sum_t R(t) \nabla_\theta \log \pi_\theta(a_t|S_t) \quad (14)$$

In the entity selection stage, we adopt the supervised learning models for the linking model:

$$L_{es}(\theta) = \sum_{D \in \mathcal{D}} \sum_{m_t \in D} \sum_{e \in C_t} h(m_t, e) \quad (15)$$

$$h(m_t, e) = \max(0, \beta - \hat{P}_\theta(e_t^*|m_t) + \hat{P}_\theta(e_t^i|m_t))$$

where θ is the model parameter, e_t^* is the ground truth entity. Finally, we get the final loss by summing the individual loss:

$$L(\theta) = L_{es}(\theta) + \gamma_1 |J_{ms}(\theta)| \quad (16)$$

where $\gamma_1 > 0$. In the training stage, the entities in state S_t and \hat{C}_t are all ground truth entities, which can ensure that the historical information is correct.

Experiment

We conduct experiments on the following datasets: (1) For in-domain scenario, we use AIDA-CoNLL dataset (Hoffart et al. 2011): this dataset contains AIDA-train for training, AIDA-A for validation, and AIDA-B (AI) for testing; (2) For out-domain scenario, we trained on AIDA-train and evaluated on five popular test sets: MSNBC (MS), AQUAINT (AQ), ACE2004 (AC), which were cleaned and updated by Guo and Barbosa (2018), CWEB (CW) and WIKI (WI) which were automatically extracted from ClueWeb and Wikipedia respectively. The statistics of these datasets are shown in Table 1. ‘Men’ is the number of mentions in all the documents. ‘Doc’ is the number of documents. ‘Pre-doc’ is the average number of mentions in each document. ‘Gold recall’ is the percentage of mentions for which the candidate entities contain the ground truth entity.

Experimental Settings

We set the dimensions of word embedding to 300 and used GloVe embedding (Pennington, Socher, and Manning 2014). The entity embedding is publicly released by Ganea and Hofmann (2017). We use the following parameter values: $K = 7$, the size of the hidden layer is

Dataset	Men	Doc	Pre-doc	Gold recall
AIDA-train	18448	946	19.5	-
AIDA-A	4791	216	22.1	97.2
AIDA-B	4485	231	19.4	98.5
MSNBC	656	20	32.8	98.5
AQUAINT	727	50	14.5	93.9
ACE2004	257	36	7.1	91.4
CWEB	11154	320	34.8	91.3
WIKI	6821	320	21.3	92.6

Table 1: Statistics of datasets used in experiments.

set to 100, the discount factor $\gamma = 0.9$. The method of generating candidate entities is equal to previous work (Le and Titov 2018). The value range of weight γ_1 of multi-task learning is [0.001, 0.00075, 0.0005, 0.00025, 0.0001]. In R2, we set $\{\lambda_{T,T}, \lambda_{T,F}, \lambda_{F,F}, \lambda_{F,T}\}$ into the following different parameter sets $\{0, -2, -1, 0\}$ and $\{0, -L\hat{P}_{\tau-1, \tau}, -L\hat{P}_{\mu-1, \mu}, 0\}$, where $-\hat{P}_{\tau-1, \tau}$ refers to the probability of the incorrect linked entity. In the Transformer-based model, we use 4 layers of the encoder with 6 attention heads, the vector size is 50 in self-attention, and the hidden size and feed-forward layer size are 300 and [300, 600]. The range of the sliding window is set to [2, 3, 4, 5, 6, 7, L]. We use the Adam (Kingma and Ba 2014) to train our model for 300 epochs with a learning rate of $2e-4$ until validation accuracy exceeds 90%, afterward setting it to $1e-4$.

Experimental Results

Methods	AIDA-B
Deep-ED (Ganea and Hofmann 2017)	92.22
FGS2EE (Hou et al. 2020)	92.63
DGCN (Wu et al. 2020)	93.13
BERT-Ent-Sim (Chen et al. 2020)	93.54
DRL (Fang et al. 2019)	94.30
DCA-SL (Yang et al. 2019)	94.64
Attn-global (Le and Titov 2018)	93.07
Tran-global (300)	93.42±0.20
Tran-global (600)	93.75±0.17
Tran-DyMen (R1)	94.26±0.10
Tran-DyMen (R2-1)	94.15±0.14
Tran-DyMen (R2-2)	93.18±0.20
Tran-DyMen (R3)	93.30±0.10
Attn-DyMen (R1)	94.90±0.10
Attn-DyMen (R2-1)	94.82±0.05
Attn-DyMen (R2-2)	94.80±0.04
Attn-DyMen (R3)	94.90±0.10

Table 2: F1 scores (%) on AIDA-B.

Following previous work, we first report the micro F1 scores on AIDA-B and other test sets. Table 2 shows the results on the AIDA-B. ‘Attn’ is the first local model and ‘Tran’ is the Transformer-based local model. ‘global’ is the global model in Le and Titov (2018). ‘(300)’ means

System	MSNBC	AQUAINT	ACE2004	CWEB	WIKI
Deep-ED (Ganea and Hofmann 2017)	93.70	88.50	88.50	77.90	77.50
FGS2EE (Hou et al. 2020)	94.26	88.47	90.70	77.41	77.66
DGCN (Wu et al. 2020)	92.50	89.40	90.60	81.20	77.60
BERT-Ent-Sim (Chen et al. 2020)	93.40	89.80	88.90	77.90	80.10
DRL (Fang et al. 2019)	92.80	87.50	91.20	78.50	82.80
DCA-SL (Yang et al. 2019)	94.57	87.38	89.44	73.47	78.16
Attn-global (Le and Titov 2018)	93.90	88.30	89.90	77.50	78.00
Tran-global	94.10±0.16	89.45±0.50	90.95±0.40	77.03±0.35	76.14±0.30
Tran-DyMen (R1)	94.41±0.10	88.13±0.20	90.14±0.40	74.07±0.10	75.95±0.12
Tran-DyMen (R2-1)	94.87±0.15	88.19±0.22	90.14±0.40	74.15±0.20	75.46±0.23
Tran-DyMen (R2-2)	94.41±0.15	87.85±0.20	90.94±0.40	73.23±0.23	75.57±0.21
Tran-DyMen (R3)	94.41±0.15	88.29±0.10	90.54±0.40	73.27±0.20	75.46±0.32
Attn-DyMen (R1)	95.07±0.20	87.44±0.50	91.35±0.00	75.07±0.20	78.24±0.20
Attn-DyMen (R2-1)	94.92±0.29	87.83±0.23	90.54±0.40	74.80±0.27	78.02±0.32
Attn-DyMen (R2-2)	95.03±0.15	87.53±0.34	91.35±0.40	74.57±0.17	78.06±0.31
Attn-DyMen (R3)	94.72±0.15	87.72±0.31	90.54±0.40	75.22±0.30	78.27±0.10

Table 3: The performance comparison on cross-domain datasets using F1 score.

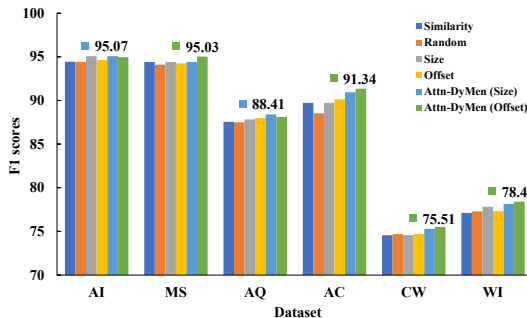


Figure 6: The performance comparison of different ranking methods of mention.

that the feed-forward layer size are 300. ‘R2-1’ means we adopt the parameter sets $\{0, -2, -1, 0\}$ in the reward function R2, and ‘R2-2’ means we adopt the parameter sets $\{0, -L\hat{P}_{\tau-1, \tau}, -L\hat{P}_{\mu-1, \mu}, 0\}$ in the reward function R2. We can see that our dynamic mention selection method can achieve the highest performance in-KB accuracy.

Table 3 shows the results on five cross-domain datasets. It is worth noting that FGS2EE, DGCN, BERT-Ent-Sim, and DRL use the additional text from Wikipedia data for training. In FGS2EE and BERT-Ent-Sim, the global entity linking model is the same as Le and Titov (2018). Comparing Attn-global and Tran-global, our Transformer-based local model has improved performance on MSNBC, AQUAINT, and ACE2004. Our model performs well even when compared to FGS2EE and BERT-Ent-Sim with extra text, especially on ACE2004. Comparing Attn-DyMen and DCA-SL, where the only difference between them is that the order of mention in DCA-SL is fixed, and we can see that the performance of the dynamic mention selection model is improved on all datasets. At the same time, Tran-DyMen also only performs better on AQUAINT.

γ_1		AI	MS	AQ	AC	CW	WI
1e-3	R1	94.7	94.6	87.6	90.1	75.5	77.5
	R2	94.9	94.4	87.1	90.1	75.3	78.0
	R3	95.0	94.7	87.6	90.5	74.3	78.0
7.5e-4	R1	95.0	94.1	87.1	90.5	75.0	77.7
	R2	94.9	94.6	86.9	90.1	75.6	77.8
	R3	94.8	94.9	87.0	90.5	74.7	77.7
5e-4	R1	94.8	94.4	87.8	90.5	74.4	77.8
	R2	94.8	94.7	87.0	90.9	74.5	78.1
	R3	94.8	94.6	88.1	89.7	75.0	77.9
2.5e-4	R1	94.9	95.2	86.3	91.4	74.3	78.1
	R2	95.0	95.2	87.1	90.9	74.1	78.0
	R3	94.9	94.9	87.7	90.1	74.4	77.8
1e-4	R1	94.8	95.2	86.9	90.9	74	77.8
	R2	94.8	95.0	87.1	90.9	74.5	78.0
	R3	95.0	94.7	87.71	90.1	74.6	78.3

Table 4: The performance comparison of different γ_1 .

Experimental Analysis

(1) **Order of mention.** Figure 6 shows the performance comparison of different ranking methods of mention. ‘Similarity’ refers to sorting by the similarity between mentions. ‘Random’ means to break up the order of the mention in the original document. ‘Size’ first links mentions with smaller candidate sizes, as they tend to be easier to link. ‘Offset’ is the natural order in the original document. ‘Attn-DyMen (Size)’ means we execute our Attn-DyMen model based on the ‘Size’ mention sequence. We can find the results of ‘Attn-DyMen (Offset)’ in Table 3. We can see that the dynamic mention selection method is superior to the fixed mention sequence method.

(2) **Parameter γ_1 .** Table 4 shows the performance comparison of different γ_1 . We set the size of the window to 4, and ‘R2’ is R2-2. It can be seen that the parameters have dif-

ferent effects on different reward functions and datasets. For example, in terms of data, MSNBC achieves better performance on smaller γ_1 , while CWEB requires larger parameters γ_1 . In terms of the reward function, we take the average on all datasets, and R1 achieves better performance on larger γ_1 , while R2-2 requires smaller parameters γ_1 .

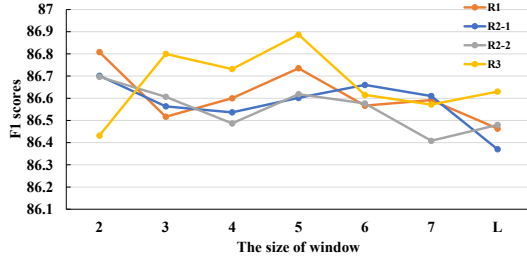


Figure 7: Results averaged over all datasets.

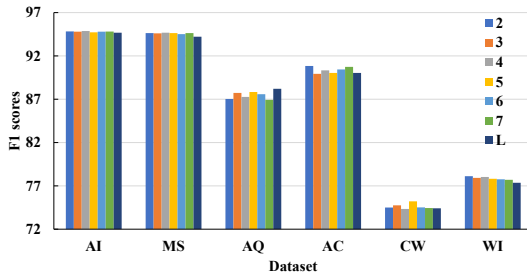


Figure 8: Results averaged over all rewards.

(3) **Size of the window.** Figure 7 and 8 shows the performance comparison of different window sizes, here $\gamma_1 = 0.0001$. We can see that different reward functions and datasets have different sensitivity to the size of the window. In general, some datasets require dynamic selection across the entire mention sequence, such as AQUAINT. Other datasets need to select mention within a limited window, such as ACE2004, CWEB, and WIKI. In terms of the reward function, all methods achieve the best performance in the limited window, which shows the effectiveness of the sliding window.

Ablation Studies

(1) **Transformer-based local model.** We design ablation experiments on the Transformer-based local model. It can be observed from Table 5 that the N_1 module are essential for the Transformer-based local model. Compared with Trans-global, position encoding, type encoding, and segment encoding can improve the model significantly.

(2) **Feature fusion.** To fairly compare the sequential decisions and the pairwise coherence global entity linking method (Ganea and Hofmann 2017), we design ablation studies on feature fusion. We remove the mention-entity type feature $\Phi_T(m_t, e_t^j)$ and the neighborhood entities feature $\Psi(e_{t+1}, \mathcal{N}(C_t))$ mentioned in . Table 6 shows that the pairwise coherence method is superior to the sequential decisions with non-dynamic mention selection (Yang et al.

Methods	AI	MS	AQ	AC	CW	WI	Avg.
Trans-global	93.3	94.1	90.0	90.9	76.9	75.5	86.8
w/o Posi	93.1	94.3	89.6	89.7	77.7	75.4	86.6
w/o Type	93.6	94.1	89.2	90.5	77.0	75.2	86.6
w/o Seg	93.5	94.1	89.5	90.1	76.8	75.8	86.6
w/o N_1	74.2	89.8	84.8	86.9	70.4	64.3	78.4
w/o N_2	93.7	94.0	89.0	90.9	76.9	75.4	86.7

Table 5: Ablation studies on Trans-global.

2019) in all datasets. Compared with the pairwise coherence method, the dynamic mention selection method slightly improves some data, but it also has an obvious decline in AIDA-B and WIKI. We think that the reason may be the shortcomings of the sequential decisions itself. When we disambiguate a mention, the sequential method can only be based on the previously linked entities, which may lead to the model can not make full use of all the information in a real sense, especially in the initial stage of the linking, there is little information to utilize. Meanwhile, the existing sequence-based one-way decision methods cannot correct potentially incorrect linked entities. Our approach can alleviate the problem of information inadequacy by adjusting the order of mention and selecting the next target based on historical information. However, it cannot solve the problem caused by the one-way decision.

Methods	AI	MS	AQ	AC	CW	WI
Attn-global	93.1	93.9	88.3	89.9	77.5	78.0
DCA(w/o T-K)	91.5	93.8	87.7	89.7	76.7	77.3
R1 (w/o T-K)	92.0	94.0	88.5	90.1	77.1	77.2
R2-1(w/o T-K)	92.0	94.0	89.2	90.1	76.6	77.0
R2-2(w/o T-K)	91.5	94.1	89.4	90.1	77.3	76.5
R3(w/o T-K)	92.1	93.8	88.6	89.7	77.3	77.3

Table 6: Ablation studies on feature fusion.

Conclusion and Future Work

This paper introduces a dynamic mention selection global entity linking model. It can dynamically select the next target according to the previously linked entities. Specifically, we design a policy network that can take the linked mentions and entities as historical information and select an appropriate mention from the unresolved mentions to perform entity linking. In selecting mention, we use a sliding window sampling mention to reduce the action sampling space of reinforcement learning and maintain the semantic coherence of mention. Finally, we optimize our model from three different perspectives. Our dynamic mention selection method can be transplanted to all entity linking models based on sequential decisions. Extensive experiments on several public benchmarks demonstrate the effectiveness of our model. In future work, we will design a model to take full advantage of global information between entities and correct for potentially incorrect linked results.

References

- Bordes, A.; Usunier, N.; Chopra, S.; and Weston, J. 2015. Large-scale simple question answering with memory networks. *arXiv preprint arXiv:1506.02075*.
- Cao, Y.; Hou, L.; Li, J.; and Liu, Z. 2018. Neural Collective Entity Linking. In *Proceedings of the 27th International Conference on Computational Linguistics*, 675–686.
- Chen, S.; Wang, J.; Jiang, F.; and Lin, C.-Y. 2020. Improving entity linking by modeling latent entity type information. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 7529–7537.
- Deng, Z.; Li, Z.; Yang, Q.; Liu, Q.; and Chen, Z. 2020. Improving Entity Linking with Graph Networks. In Huang, Z.; Beek, W.; Wang, H.; Zhou, R.; and Zhang, Y., eds., *Web Information Systems Engineering – WISE 2020*, volume 12342, 343–354. Springer International Publishing. ISBN 978-3-030-62004-2 978-3-030-62005-9. Series Title: Lecture Notes in Computer Science.
- Devlin, J.; Chang, M.; Lee, K.; and Toutanova, K. 2018. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *CoRR*, abs/1810.04805.
- Fang, Z.; Cao, Y.; Li, Q.; Zhang, D.; Zhang, Z.; and Liu, Y. 2019. Joint Entity Linking with Deep Reinforcement Learning. In *The World Wide Web Conference on - WWW '19*, 438–447. ACM Press. ISBN 978-1-4503-6674-8.
- Ganea, O.-E.; and Hofmann, T. 2017. Deep Joint Entity Disambiguation with Local Neural Attention. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, 2619–2629. Copenhagen, Denmark: Association for Computational Linguistics.
- Gu, Y.; Qu, X.; Wang, Z.; Huai, B.; Yuan, N. J.; and Gui, X. 2021. Read, Retrospect, Select: An MRC Framework to Short Text Entity Linking. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(14): 12920–12928.
- Guo, Z.; and Barbosa, D. 2018. Robust named entity disambiguation with random walks. *Semantic Web*, 9(4): 459–479.
- Hoffart, J.; Yosef, M. A.; Bordino, I.; Fürstena, H.; Pinkal, M.; Spaniol, M.; Taneva, B.; Thater, S.; and Weikum, G. 2011. Robust disambiguation of named entities in text. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, 782–792.
- Hoffmann, R.; Zhang, C.; Ling, X.; Zettlemoyer, L.; and Weld, D. S. 2011. Knowledge-based weak supervision for information extraction of overlapping relations. In *Proceedings of the 49th annual meeting of the association for computational linguistics: human language technologies*, 541–550.
- Hou, F.; Wang, R.; He, J.; and Zhou, Y. 2020. Improving Entity Linking through Semantic Reinforced Entity Embeddings. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 6843–6848. Association for Computational Linguistics.
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Le, P.; and Titov, I. 2018. Improving Entity Linking by Modeling Latent Relations between Mentions. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 1595–1604. Melbourne, Australia: Association for Computational Linguistics.
- Nair, V.; and Hinton, G. E. 2010. Rectified linear units improve restricted boltzmann machines. In *ICML*.
- Pennington, J.; Socher, R.; and Manning, C. D. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, 1532–1543.
- Srivastava, N.; Hinton, G.; Krizhevsky, A.; Sutskever, I.; and Salakhutdinov, R. 2014. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. 15(56): 1929–1958.
- Sutton, R. S.; and Barto, A. G. 2011. Reinforcement learning: An introduction.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention is all you need. *arXiv preprint arXiv:1706.03762*.
- Wu, J.; Zhang, R.; Mao, Y.; Guo, H.; Soflaei, M.; and Huai, J. 2020. Dynamic graph convolutional networks for entity linking. In *Proceedings of The Web Conference 2020*, 1149–1159.
- Yan, J.; Wang, Y.; Xiang, L.; Zhou, Y.; and Zong, C. 2020. A Knowledge-driven Generative Model for Multi-implication Chinese Medical Procedure Entity Normalization. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*.
- Yang, X.; Gu, X.; Lin, S.; Tang, S.; Zhuang, Y.; Wu, F.; Chen, Z.; Hu, G.; and Ren, X. 2019. Learning Dynamic Context Augmentation for Global Entity Linking. In *Proceedings of EMNLP-IJCNLP*.
- Zhang, Z.; Sind, X.; Liu, T.; Fang, Z.; and Li, Q. 2020. Joint Entity Linking and Relation Extraction with Neural Networks for Knowledge Base Population. In *2020 International Joint Conference on Neural Networks (IJCNN)*, 1–8. ISSN: 2161-4407.