

# A Fine-grained and Noise-aware Method for Neural Relation Extraction

ADVISOR: JIA-LING, KOH

SOURCE: CIKM 2019

SPEAKER: SHAO-WEI, HUANG

DATE: 2020/05/04

# OUTLINE

- Introduction
- Method
- Experiment
- Conclusion

# INTRODUCTION

- **Relation Extraction (RE)** : Extracting semantic relations between two entities from the text corpus.

(Ex): Donald Trump is the 45th President of United States.



President of

# INTRODUCTION

## ➤ **Supervised RE :**

- Heavily relies on human-annotated data to achieve outstanding performance.
- Limited in size and domain specific, preventing large-scale supervised relation extraction.

# INTRODUCTION

## ➤ Distant supervision RE :

- Automatically generates large-scale training data through knowledge base and plain texts.

Relations In KB	<i>President_of</i> ( <b>Donald Trump</b> , <b>United States</b> )
Sentences in Plain texts	<p>S1: <b>Donald Trump</b> is the 45th President of the <b>United States</b>.</p> <p>S2: <b>Donald Trump</b> was born in the <b>United States</b>.</p> <p>S3: <b>Donald Trump</b> believes the <b>United States</b> has incredible potential.</p>

} **President of**  
(A training bag)

# INTRODUCTION

## ➤ Challenge in Distant supervision(1/2) :

- Multi-instance multi-label problem (MIML) .

Relations In KB	<i>Place_of_birth</i> (Donald Trump, United States) <i>President_of</i> (Donald Trump, United States)
Sentences in Plain texts	<b>S1</b> : Donald Trump is the 45th President of the United States. ( <i>President_of</i> ) <b>S2</b> : Donald Trump was born in the United States. ( <i>Place_of_birth</i> )

Multi-label problem

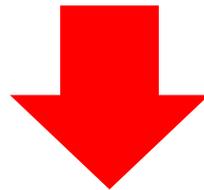
Relations In KB	<i>Place_of_birth</i> (Donald Trump, United States) <i>President_of</i> (Donald Trump, United States)
Sentences in Plain texts	S3: Donald Trump believes the United States has incredible potential. (-)

Multi-instance problem

# INTRODUCTION

## ➤ **Challenge in Distant supervision(2/2) :**

- The assigned **relation labels** are annotated at **bag-level** (a set of sentences) **instead of sentence-level**.



Reinforcement Learning Model

# INTRODUCTION

➤ **Goal :**

The most informative  
sentence for each relation



⋮



# OUTLINE

- ▶ Introduction
- ▶ Method
- ▶ Experiment
- ▶ Conclusion

# REINFORCEMENT LEARNING

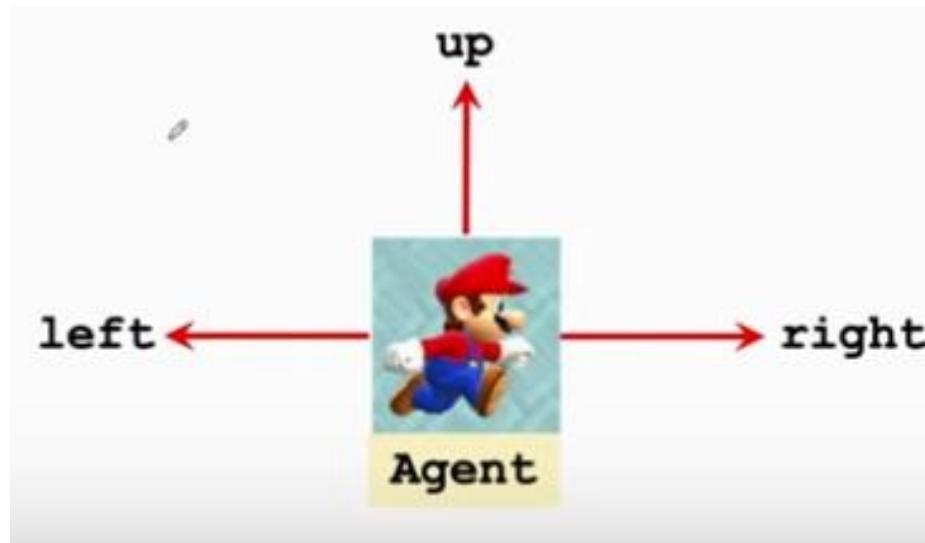
10

<https://www.youtube.com/watch?v=vmkRMvhCW5c&t=1557s>

State

Action; Agent

Reward



Collect a coin:  $R = +1$

Win the game:  $R = +10000$

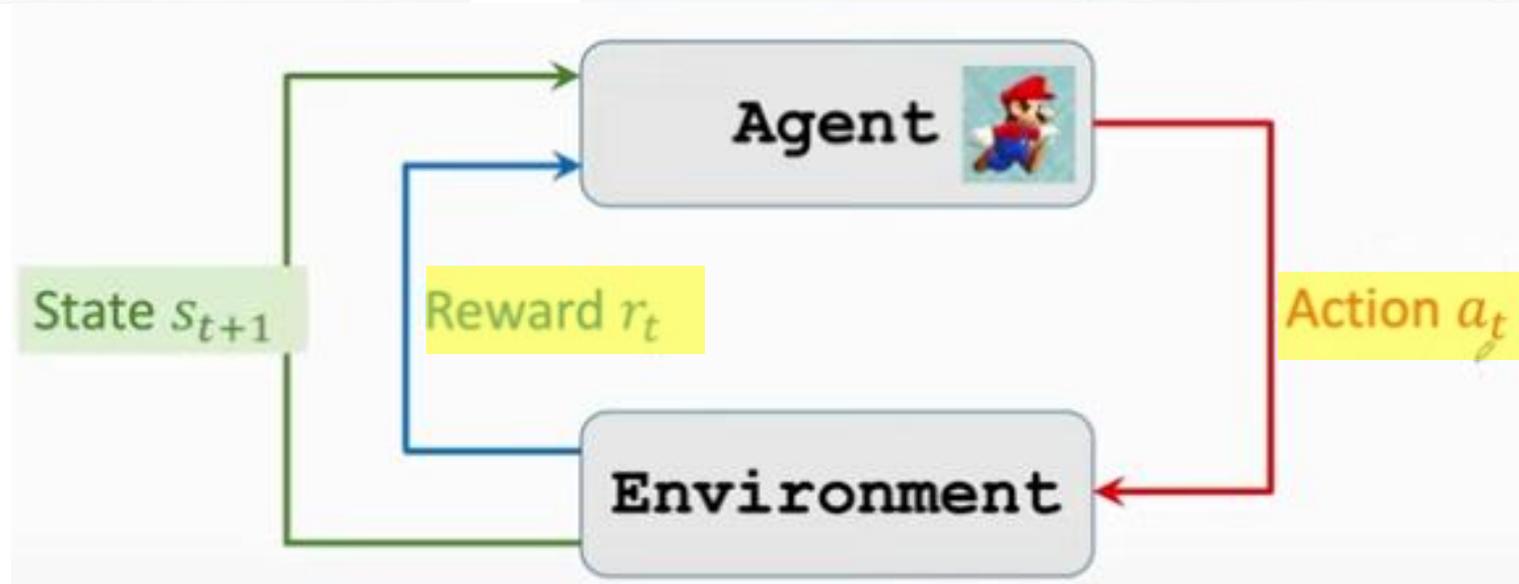
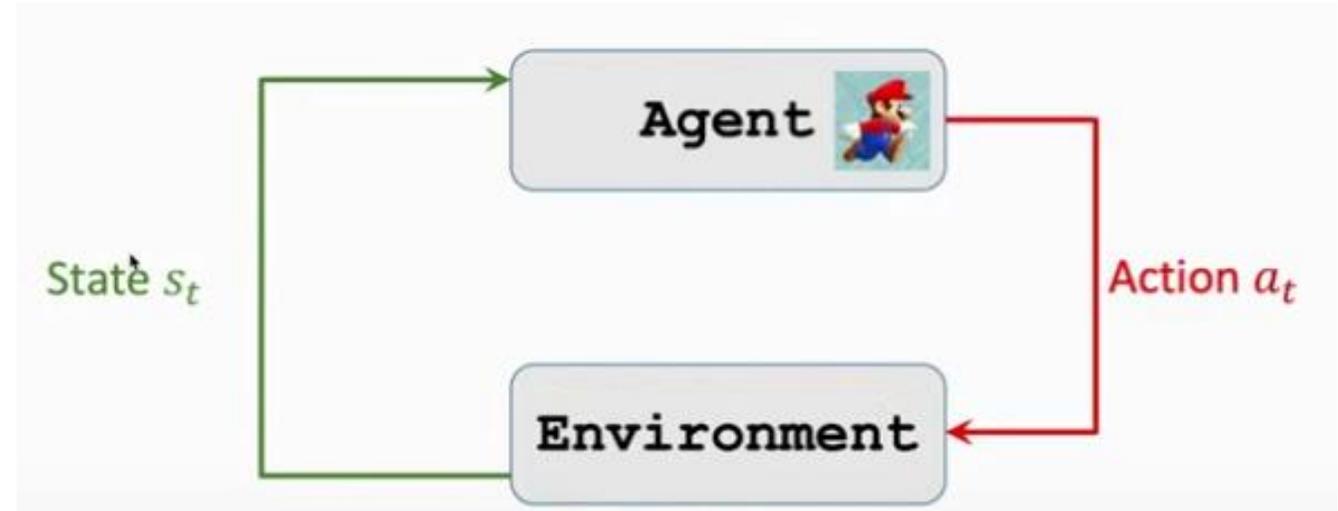
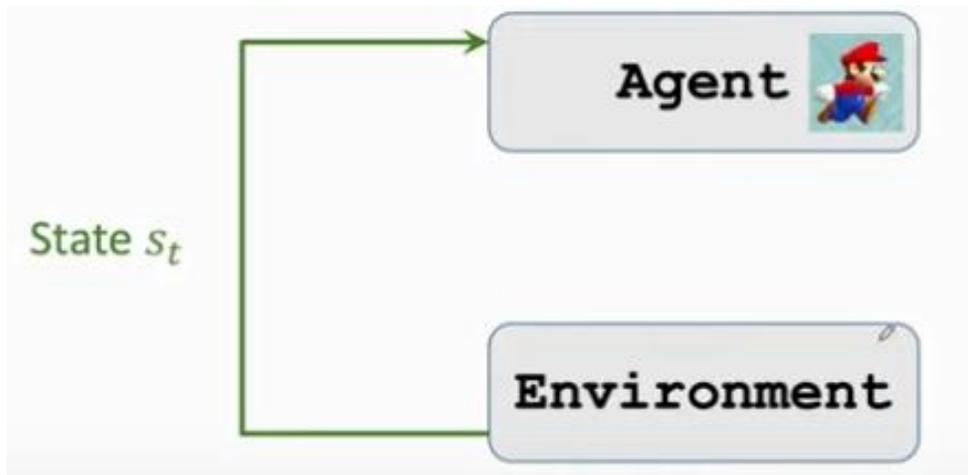
Touch a Goomba:  $R = -10000$   
(game over).

Nothing happens:  $R = 0$

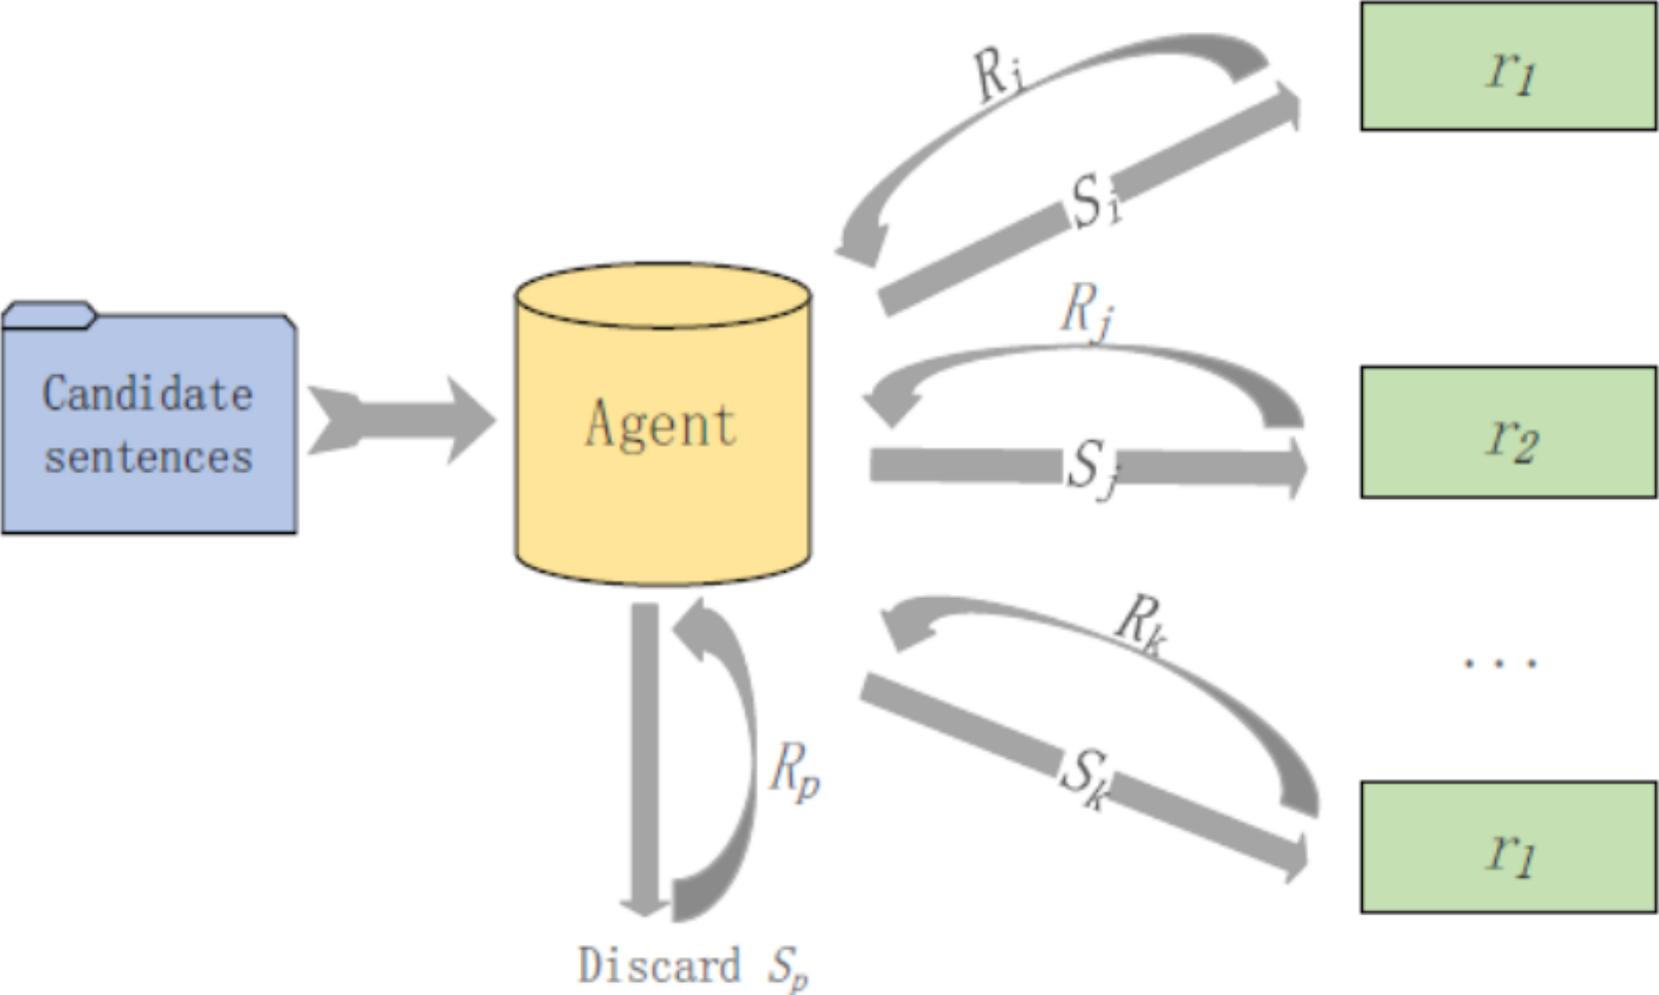
# REINFORCEMENT LEARNING

11

<https://www.youtube.com/watch?v=vmkRMvhCW5c&t=1557s>



# FRAMEWORK



# METHOD

## Notation and Problem definition

➤ Let  $B = \{ \langle e_1, e_2 \rangle, (r_1, \dots, r_l), \{S_1, \dots, S_n\} \}$  be a training bag.

An **entity pair**

The **relations** that link  
the entity pair **in KB**

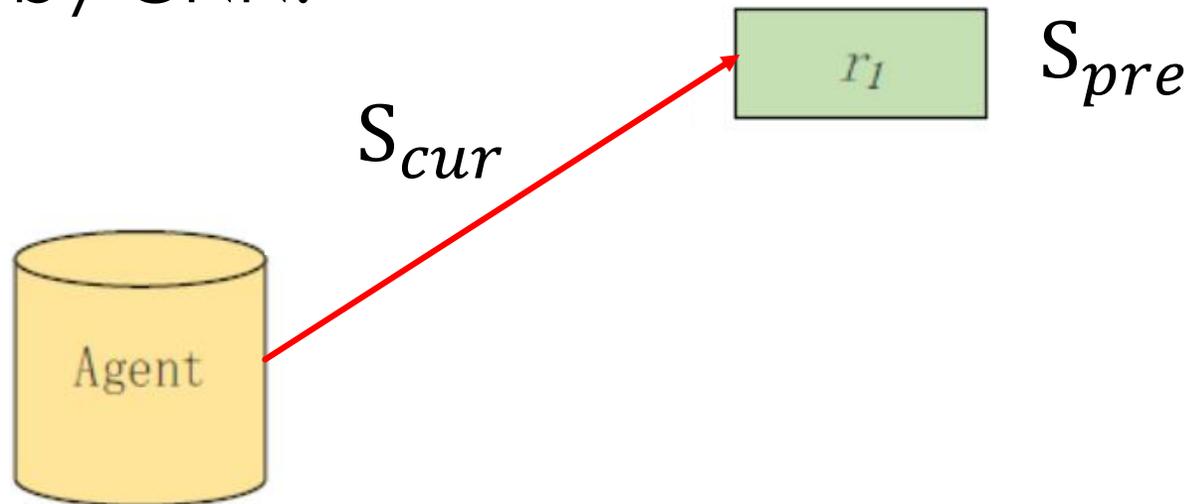
The **sentences** from corpus  
which **mention this entity pair**

➤ **Problem definition** : Figuring out **the most expressive sentence for each relation in B.**

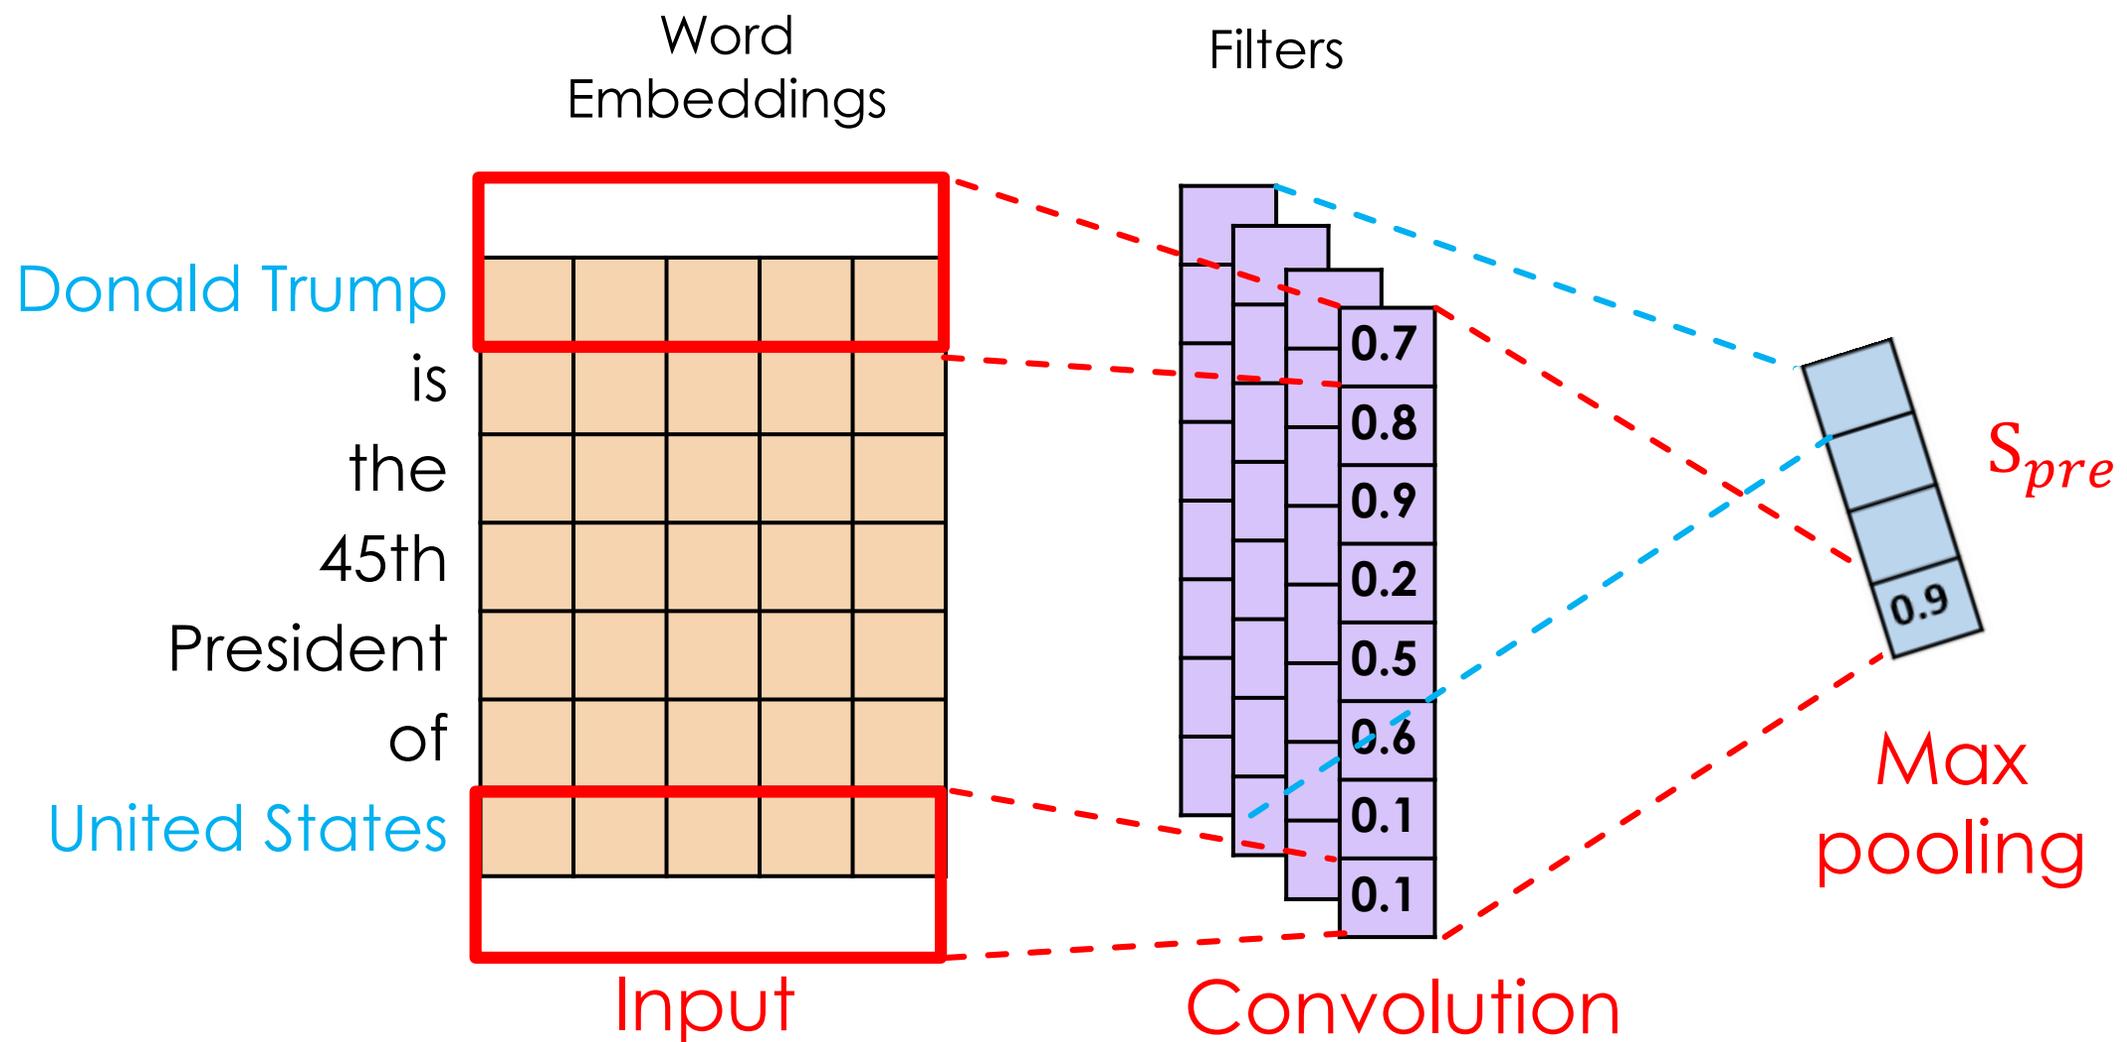
(in order to promote the extractor's performance)

# METHOD State

- **Embedding of target entity pair** :  $[e_1, ; e_2]$
- **Encoding for previously chosen sentence** :  $S_{pre}$ ,  
encoded by CNN.

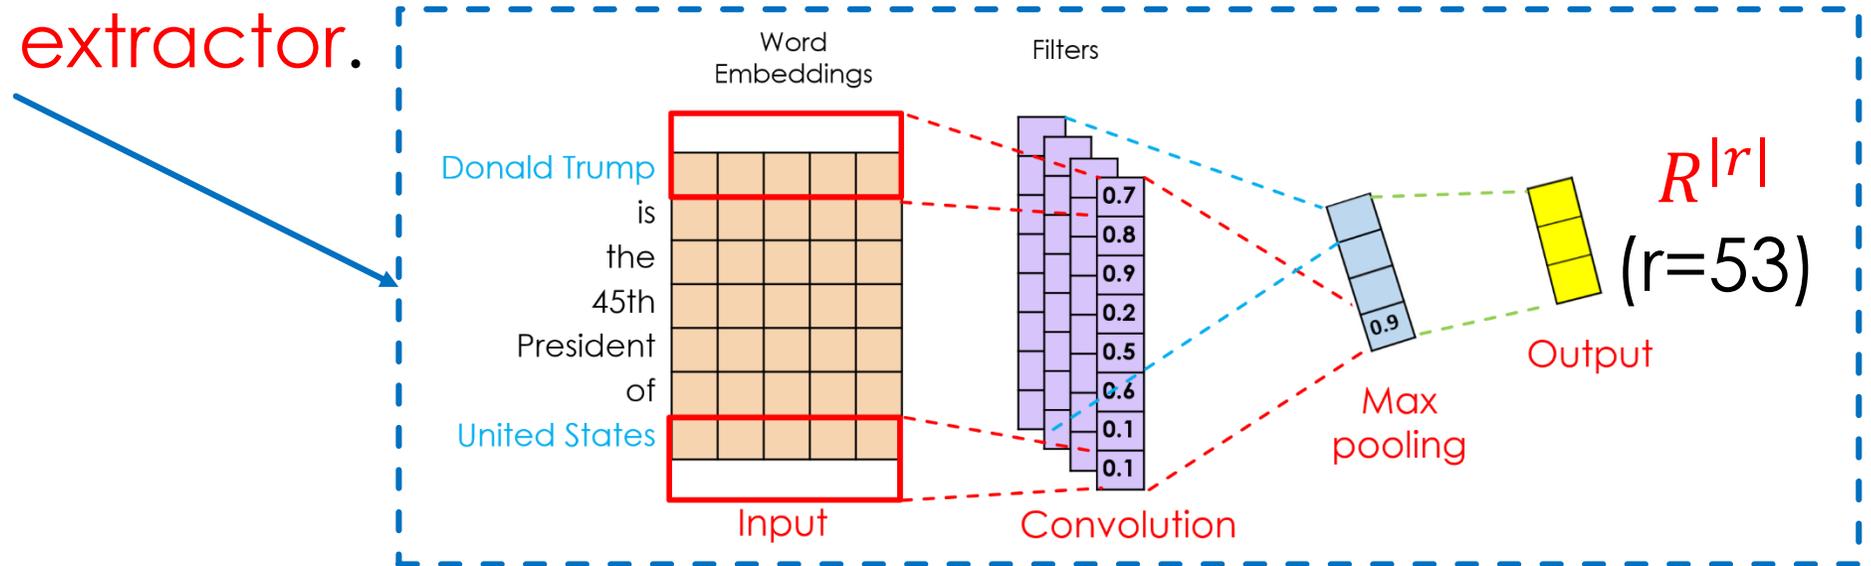


# METHOD **State**



➤ **Confidence score for  $S_{pre}$**  :  $p(r | S_{pre}; \theta)$ , calculated by

relation extractor.



➤ **Encoding for current sentence** :  $S_{cur}$

➤ **Confidence score for  $S_{cur}$**  :  $p(r | S_{cur}; \theta)$

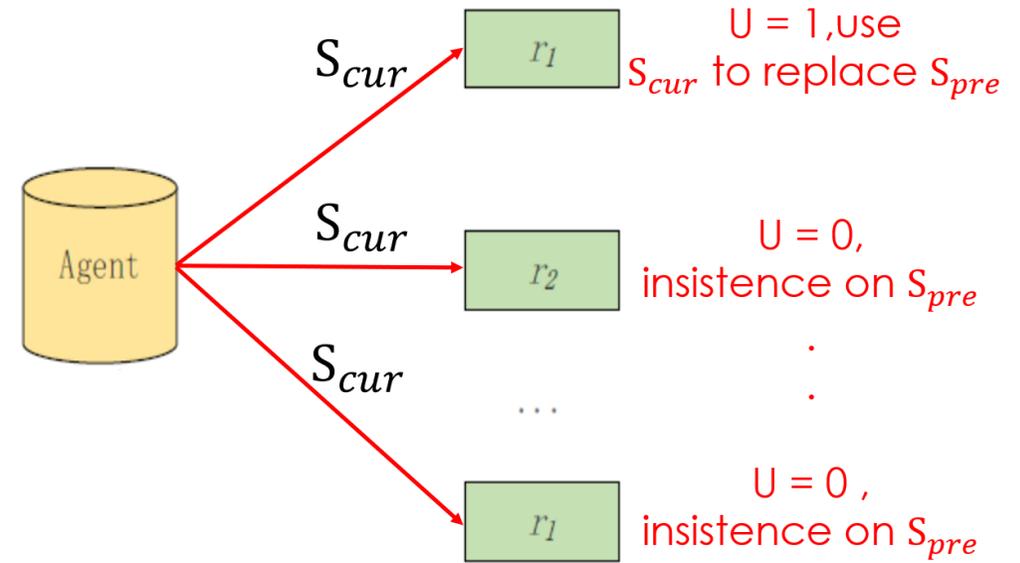
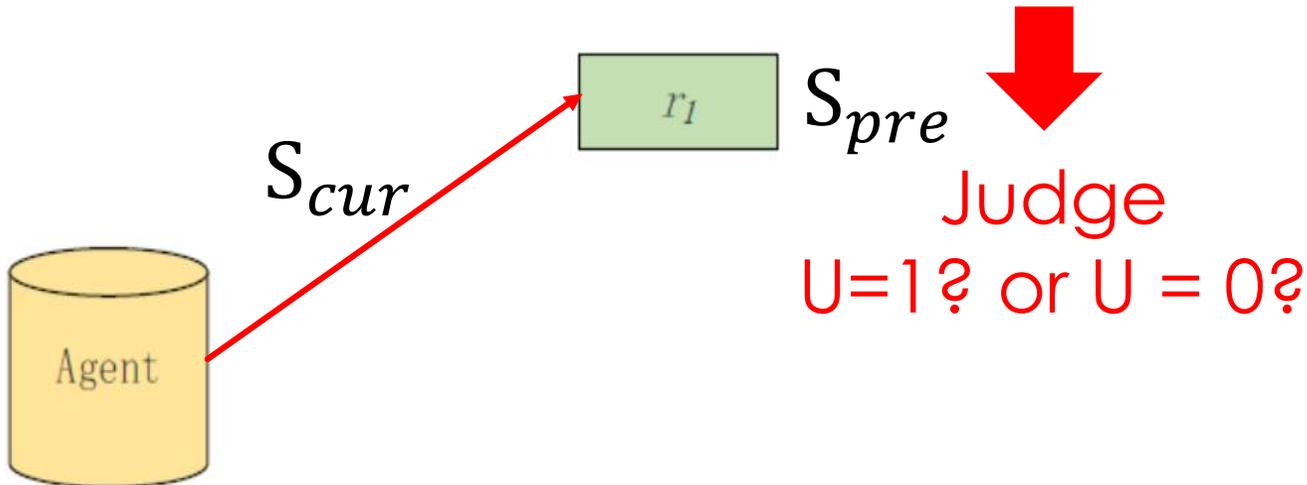
➤ **St** :  $[e_1 ; e_2 ; S_{pre} ; p(r | S_{pre}; \theta) ; S_{cur} ; p(r | S_{cur}; \theta)]$

# METHOD **Action**

➤ **First part of action** : U, decide whether to adopt the current sentence to replace the previously chosen.

- $U = 1 \rightarrow$  Yes.
- $U = 0 \rightarrow$  No.

$$Q(St, U) = f(W_u(f(W \cdot St)) + b_u)$$



## ➤ Reasonable assumptions for each bag :

- Expressed-at-least-once
- Express-at-most-one

## ➤ Compete mechanism :

- When **more relations simultaneously** intend to update the  $S_{pre}$  with the  $S_{cur}$ .
- **Only update** the relation which **has highest  $Q(S_t, U = 1)$ .**

# METHOD **Action**

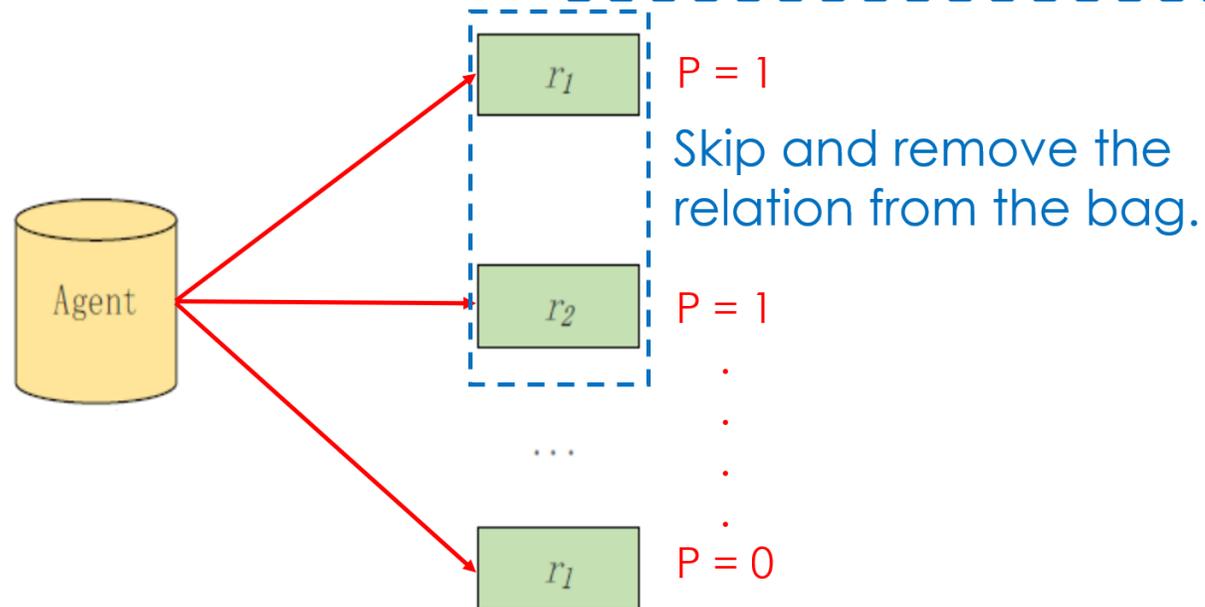
19

➤ **Second part of action** : P, decide the relation **whether to stop the search action.**

- $P = 1 \rightarrow$  Yes ( **believe has picked out the expressive sentence** for this relation ) .

$$Q(St, P) = f(W_p(f(W \cdot St)) + b_p)$$

- $P = 0 \rightarrow$  No.



# METHOD **Action**

➤ Takes the action  $A^*$  that equals  $\operatorname{argmax}_A Q(S_t, A)$  :

- $$Q(S_t, A; \eta) = \begin{cases} Q(S_t, U) = f(W_u(f(W \cdot S_t)) + b_u) \\ Q(S_t, P) = f(W_p(f(W \cdot S_t)) + b_p) \end{cases}$$

# METHOD Reward

21

## ➤ Reward :

- When execute a action → Get a reward.
- The objective of reinforcement learning : By maximize total reward to learn Q function.

$$Re(S_t, A) = \begin{cases} p(r_t|S_{cur}; \theta) - p(r_t|S_{pre}; \theta) & (\neg T) \wedge (U = 1) & (+)/(-) \\ p(r_t|S_{pre}; \theta) - p(r_t|S_{cur}; \theta) & (\neg T) \wedge (U = 0) \wedge (\arg \max_r p(r|S_{cur}; \theta) = r_t) & (+)/(-) \\ 0 & (\neg T) \wedge (U = 0) \wedge (\arg \max_r p(r|S_{cur}; \theta) \neq r_t) & \\ p(r_t|S_{cur}; \theta) & (T) \wedge (U = 1) \wedge (\arg \max_r p(r|S_{cur}; \theta) = r_t) & (+) \\ p(r_t|S_{cur}; \theta) - \max(p(r|S_{cur}; \theta)) & (T) \wedge (U = 1) \wedge (\arg \max_r p(r|S_{cur}; \theta) \neq r_t) & (-) \\ p(r_t|S_{pre}; \theta) & (T) \wedge (U = 0) \wedge (\arg \max_r p(r|S_{pre}; \theta) = r_t) & (+) \\ p(r_t|S_{pre}; \theta) - \max(p(r|S_{pre}; \theta)) & (T) \wedge (U = 0) \wedge (\arg \max_r p(r|S_{pre}; \theta) \neq r_t) & (-) \end{cases}$$

# METHOD

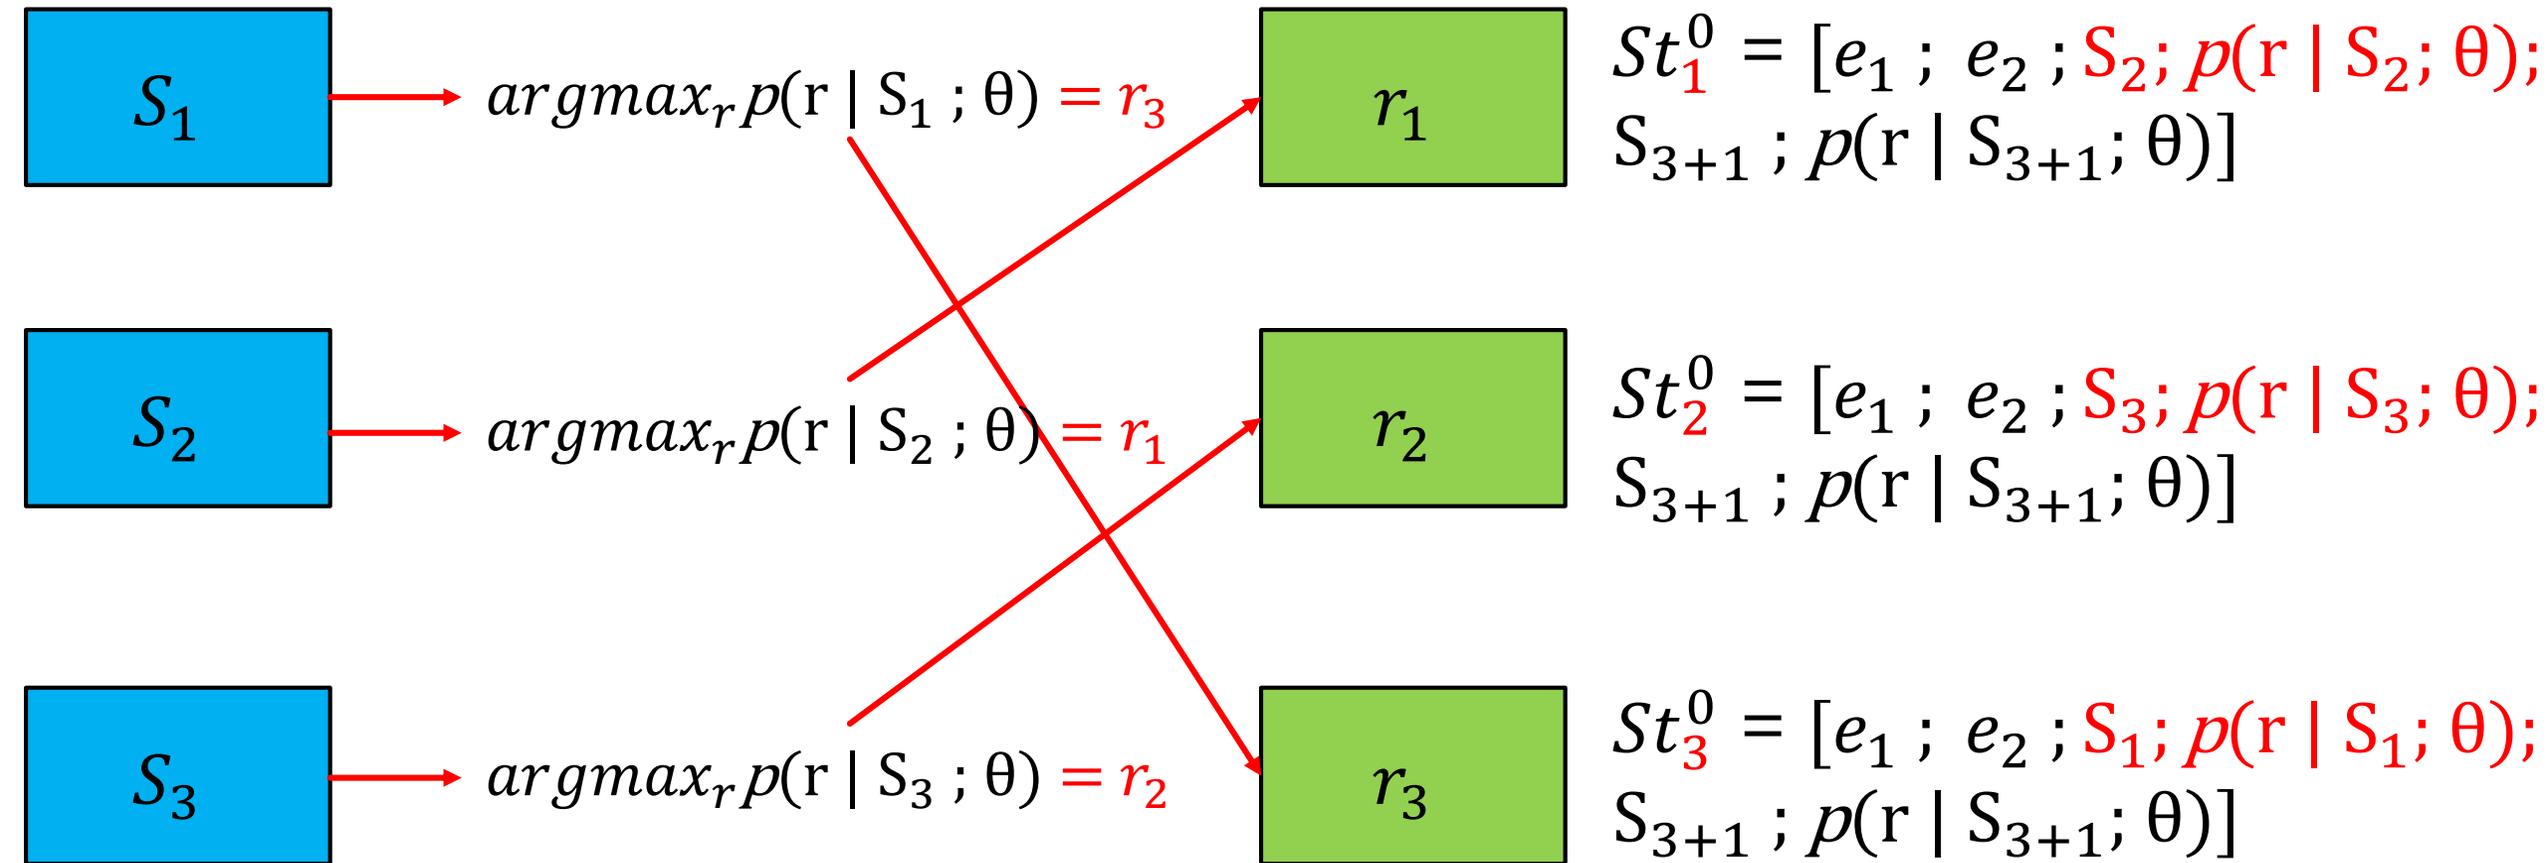
Initial states  $St_i^0$  for all episodes

- Initial states  $St_i^0$  for all relations of a bag are the same :
  - $B = \{ \langle e_1, e_2 \rangle, (r_1, \dots, r_l), \{S_1, \dots, S_n\} \}, i \in \{1, 2, \dots, l\}$
  - $St_i^0 = [e_1 ; e_2 ; 0 ; 0 ; S_1 ; p(r | S_1; \theta)],$  for all  $i$ .
  - The output values of Q-function are the same too.

# METHOD Initial states $St_i^0$ for all episodes

23

## ➤ heuristic initialization :



# METHOD

- **Reinforcement learning algorithm for MIML :**
  - In paper Algorithm 1.

➤ Optimization of the extractor :

● **Maximize**  $J(\theta) = \sum_{i=1}^{|\text{TS}|} \log p(y_i | x_i; \theta)$

$x_i$  is a sentence in TS with relation  $y_i$

➤ Optimization of the reinforcement learning :

● **minimize**  $Loss = \sum_{i=1}^{|\text{Z}|} (z_i - Q(St_i, A_i; \eta))^2$

Regarded as accurate value of  $Q(St_i, A_i)$

● 
$$z_i = \begin{cases} \text{Reward} & St^{i+1} \in \text{terminal} \\ \text{Reward} + \alpha \max_{A'} Q(St^{i+1}, A') & St^{i+1} \notin \text{terminal} \end{cases}$$

# METHOD

- **Joint training for extractor and reinforcement learning :**
  - In paper Algorithm 2.

# OUTLINE

- ▶ Introduction
- ▶ Method
- ▶ Experiment
- ▶ Conclusion

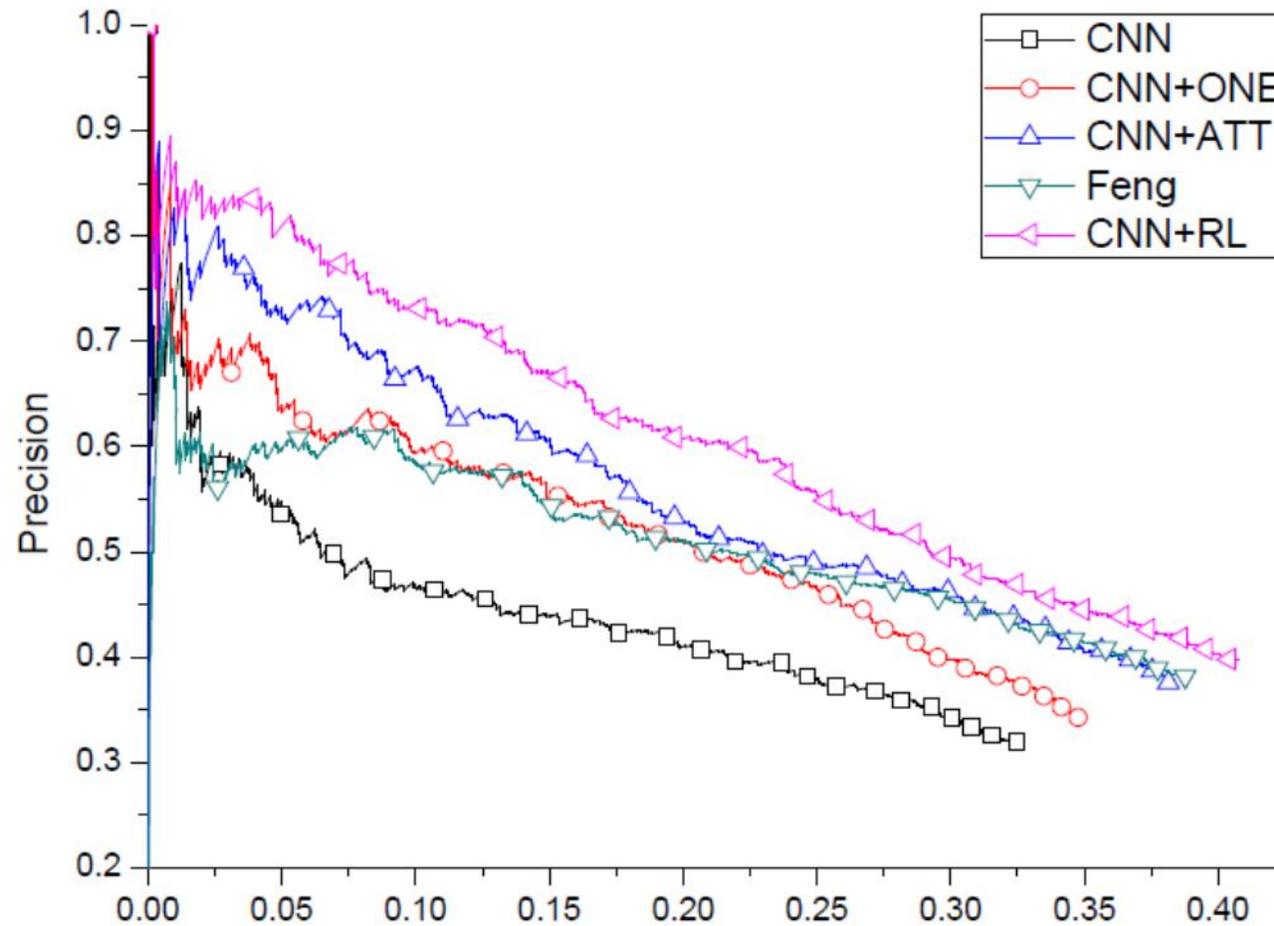
# EXPERIMENT

## Dataset

- **NYT+Freebase** : Aligning **entities and relations in Freebase** with the **corpus of New York Times**.
  - NYT in 2005-2006 → Training data
  - NYT in 2007 → Testing data

# EXPERIMENT

## Performance with NN methods



# EXPERIMENT

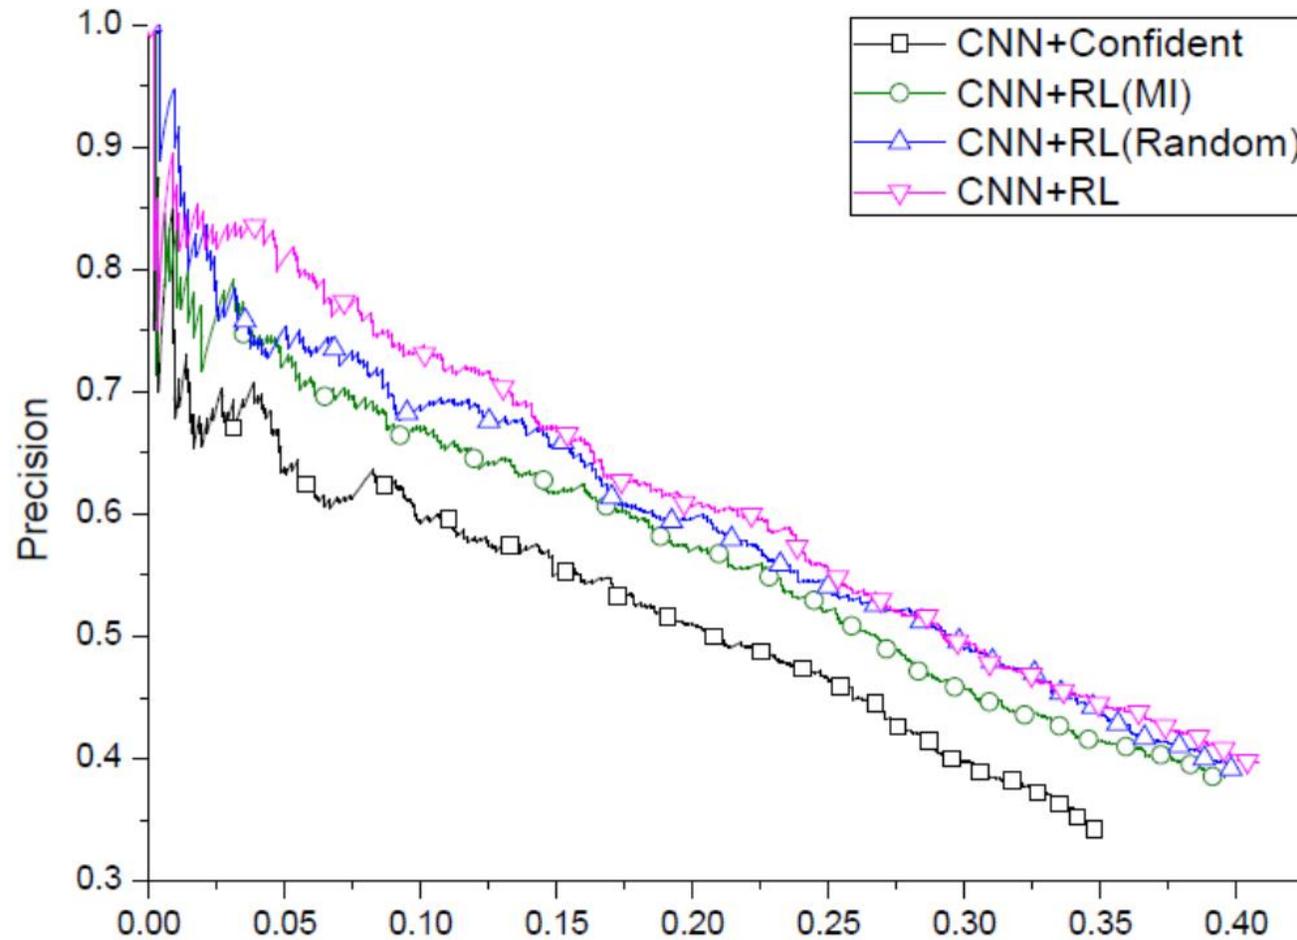
30

## Performance with NN methods

P@N(%)	100	200	300	500	mean
<i>CNN</i>	59.00	53.00	47.67	46.20	51.47
<i>CNN+ONE</i>	69.00	61.00	61.33	56.80	62.03
<i>CNN+ATT</i>	76.00	70.50	66.67	59.40	68.14
<i>Feng</i>	57.00	59.50	60.67	56.00	58.29
<i>CNN+RL</i>	<b>82.00</b>	<b>76.50</b>	<b>72.00</b>	<b>64.40</b>	<b>73.73</b>

# EXPERIMENT

## Reasonability of model design



# EXPERIMENT

32

## Case study

Method	Sentences	/location/contains	/country/capital
<i>Feng</i>	<i>S1: according to local legend , recounted by the Africa scholar Stephen Ellis in his book “ the mask of anarchy , ” a baby born</i>	choose	choose
<i>CNN+RL</i>	<i><u>in Monrovia , Liberia ’s capital , miraculously spoke english straight from the womb</u></i>	not choose	choose
<i>Feng</i>	<i>S2: my sister was living in <u>the Monrovia suburb of Paynesville , Liberia</u> , with her family and a handful of</i>	choose	choose
<i>CNN+RL</i>	<i>orphans and other refugees from the liberian civil war.</i>	choose	not choose

# EXPERIMENT

## Case study

		/person/place_lived	/deceased_person /place_of_death
<i>Feng</i>	S3: <u>Jane Jacobs</u> , the activist who took him on , <u>now lives</u> <u>in Toronto</u> .	choose	choose
<i>CNN+RL</i>		choose	<b>not choose</b>
<i>Feng</i>	S4: <u>Jane Jacobs</u> , the writer and thinker who brought penetrating eyes and ingenious insight to the sidewalk ballet of her own greenwich village street and came up with a book that challenged and changed the way people view cities , <u>died</u> <u>yesterday in Toronto</u> , where she moved in 1968 .	not choose	not choose
<i>CNN+RL</i>		not choose	<b>choose</b>

# OUTLINE

- ▶ Introduction
- ▶ Method
- ▶ Experiment
- ▶ Conclusion

# CONCLUSION

- Craft reinforcement learning to solve MIML problem, and generate the sentence-level annotated signal in distant supervised relation extraction.
- Then these chosen expressive sentences serve as training instances to feed the extractor.
- We conduct extensive experiments and the experimental results demonstrate that our model can effectively alleviate MIML problem and achieve the new state-of-the-art performance.