Recent Advances in Distantly Supervised Relation Extraction

William Wang
Department of Computer Science
University of California, Santa Barbara

Joint work with Jiawei Wu, Lei Li, Pengda Qin, Weiran Xu.

CIPS Summer School 2018

Beijing, China

Agenda

- Motivation
- Challenges in Semi-Supervised Learning
- Reinforced Co-Training (Wu et al., NAACL 18)
- Reinforced Distant Supervision Relation Extraction (Qin et al., ACL 18a)
- DSGAN (Qin et al., ACL 18b)
- Conclusions

Motivation

- Most of the existing successful stories of deep learning are still based on supervised learning.
- For example, object recognition, machine translation, text classification.
- However, in many applications, it is not realistic to obtain large amount of labeled data.
- We need to leverage unlabeled data.

A Classic Example of Semi-Supervised Learning

- Co-Training (Blum and Mitchell, 1998)

Given:

- a set L of labeled training examples
- a set U of unlabeled examples

Create a pool U' of examples by choosing u examples at random from U

Loop for k iterations:

Use L to train a classifier h_1 that considers only the x_1 portion of x

Use L to train a classifier h_2 that considers only the x_2 portion of x

Allow h_1 to label p positive and n negative examples from U'

Allow h_2 to label p positive and n negative examples from U'

Add these self-labeled examples to L

Randomly choose $2p + 2n$ examples from U to replenish U'

Challenges

- The two classifiers in co-training have to be independent.
- Choosing highly-confident self-labeled examples could be suboptimal.
- Sampling bias shift is common.

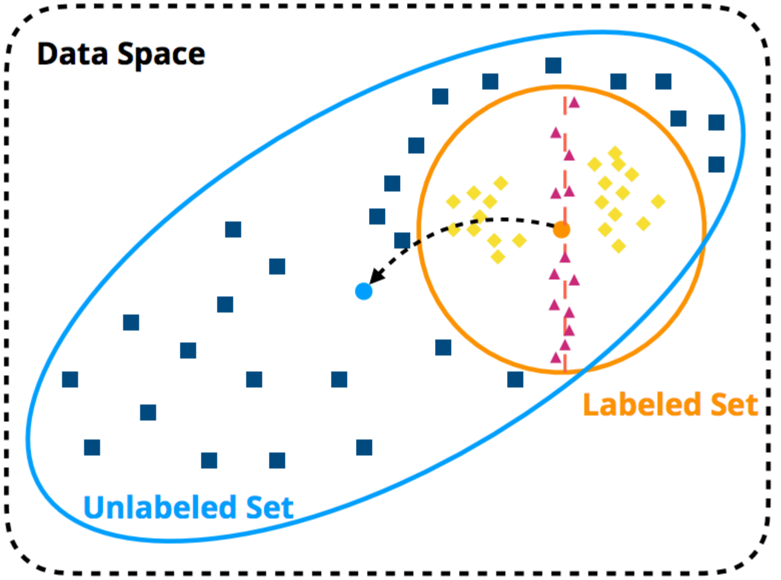

Our Approach: Reinforced SSL

- Assumption: not all the unlabeled data are useful.
- Idea: performance-driven semi-supervised learning that learns an unlabeled data selection policy with RL, instead of using random sampling.

- 1. Partition the unlabeled data space
- 2. Train a RL agent to select useful unlabeled data
- 3. Reward: change in accuracy on the validation set

Reinforcement Learning

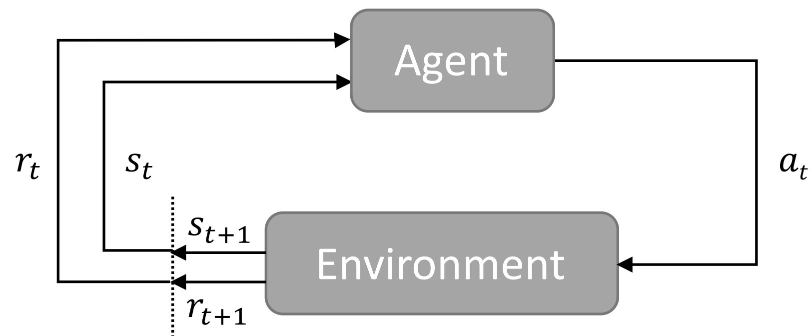

Deep Q-Network

Unlabeled Data

Reinforced Co-Training (Wu et al., NAACL 2018)

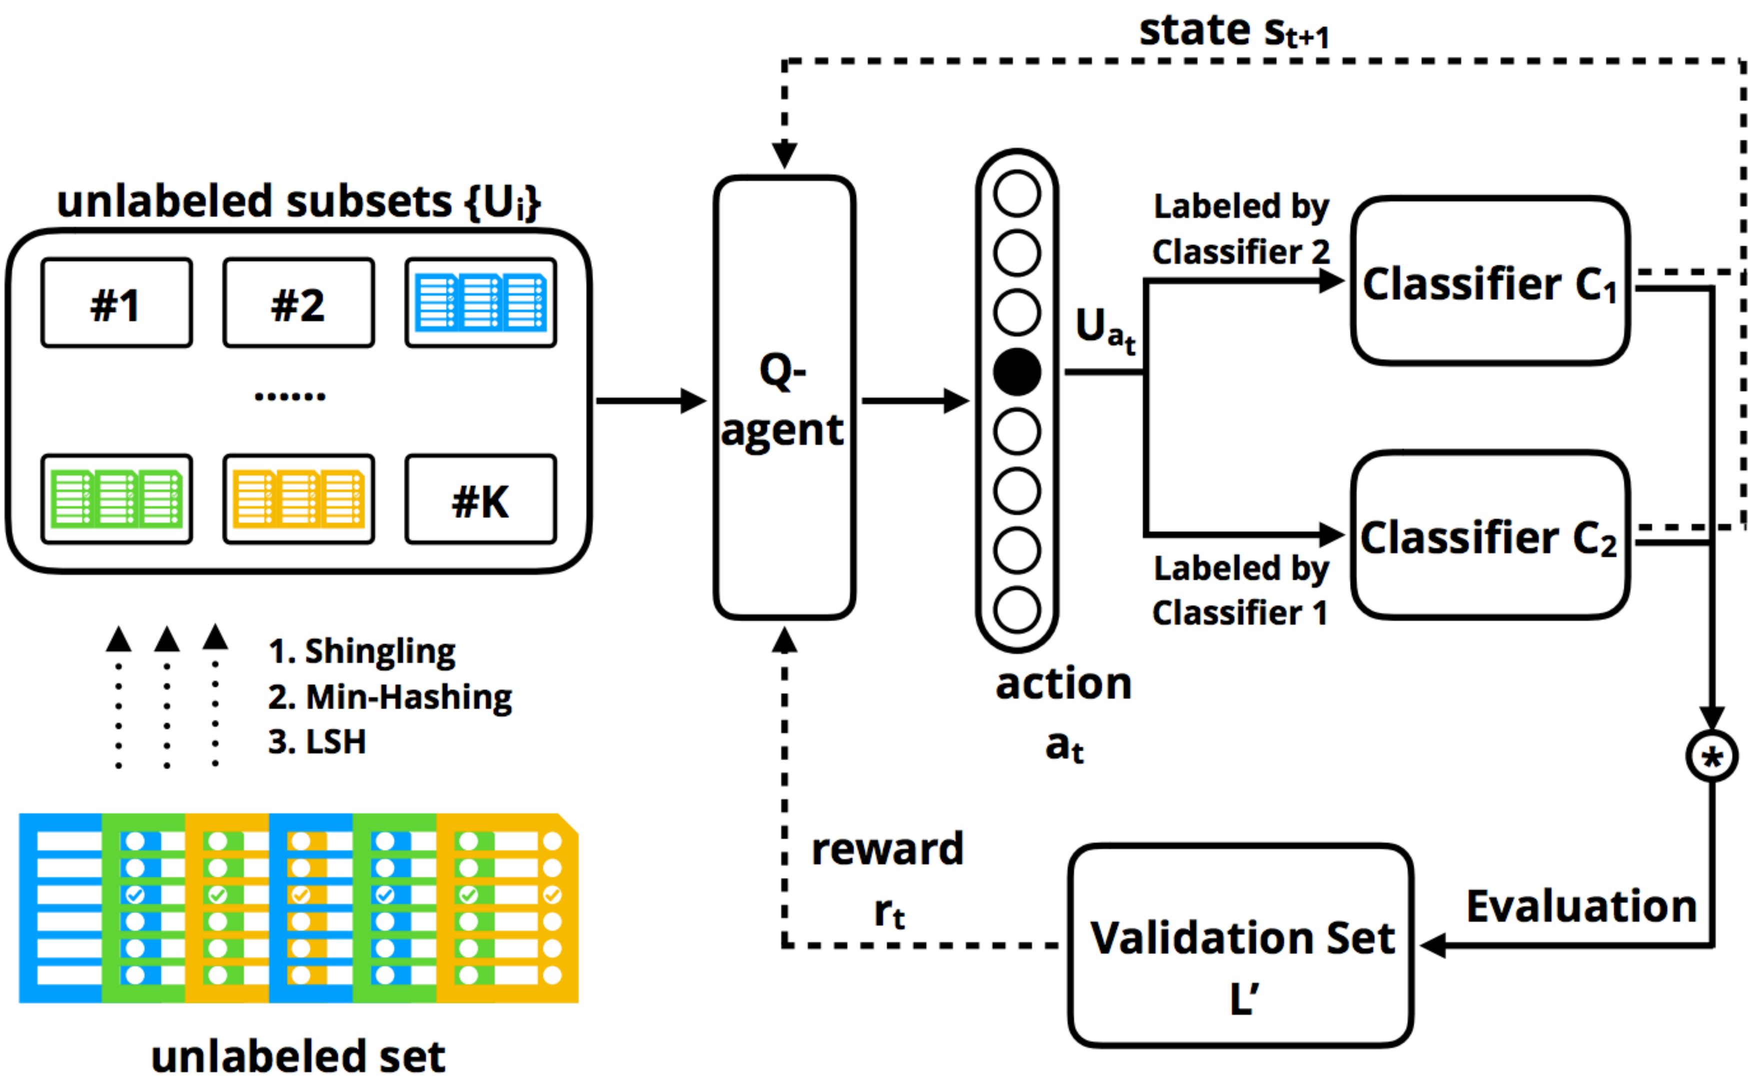

Deep Q-Learning

The Q-network parameters θ are learned by optimizing:

$$L_i(\theta_i) = \mathbb{E}_{s,a}[(V(\theta_{i-1}) - Q(s,a;\theta_i))^2], \tag{8}$$

where i is an iteration of optimization and

$$V(\theta_{i-1}) = \mathbb{E}_{s'}[r + \gamma \max_{a'} Q(s',a';\theta_{i-1})|s,a]. \tag{9}$$

Experiment 1: Clickbait Detection

Dataset	#Tweets	#Clickbait	#Non-Clickbait
Training	2,495	762	1,697
Validation	9,768	2,380	7,388
Test	9,770	2,381	7,389
Unlabeled	80,012	N/A	N/A

Table 1: Statistics of Clickbait Dataset.

Experiment 1: Clickbait Detection

Methods	Prec.	Recall	F1 Score
Self-attentive biGRU	0.683	0.649	0.665
CNN (Document)	0.537	0.474	0.503
Standard Co-Training	0.418	0.433	0.425
Performance Co-Training	0.581	0.629	0.604
CoTrade Co-Training	0.609	0.637	0.623
Sequence-SSL	0.595	0.589	0.592
Region-SSL	0.674	0.652	0.663
Adversarial-SSL	0.698	**0.691**	0.694
Reinforced Co-Training	**0.709**	0.684	**0.696**

Table 2: The experimental results on clickbait dataset. Prec.: precision.

Experiment 2: Generic Text Classification

Dataset	AG's News	DBpedia
#Classes	4	14
#Training	12,000	56,000
#Validation	12,000	56,000
#Test	7,600	70,000
#Unlabeled	96,000	448,000

Table 4: Statistics of the Text Classification Datasets.

Experiment 2: Generic Text Classification

Methods	**AG's News**	**DBpedia**
CNN (Training+Validation)	28.32%	9.53%
CNN (All)	8.69%	0.91%
Standard Co-Training	26.52%	7.66%
Performance Co-Training	21.73%	5.84%
CoTrade Co-Training	19.06%	5.12%
Sequence-SSL	19.54%	4.64%
Region-SSL	18.27%	3.76%
Adversarial-SSL	8.45%*	0.89%*
Reinforced Co-Training	**16.64%**	**2.45%**

Table 5: The experimental results on generic text classification datasets. * Adversarial-SSL is trained on full labeled data after pre-training.

Conclusion

- We proposed a novel RL framework for semi-supervised learning
- Strong results in SSL text classification
- Also showed effectiveness in relation extraction

Deep Reinforcement Learning for Distantly Supervised Relation Extraction

Outline

- Motivation
- Algorithm
- Experiments
- Conclusion

Outline

- **Motivation**
- Algorithm
- Experiments
- Conclusion

Relation Extraction

Relation Type with
Labeled Dataset

Plain Text Corpus
(Unstructured Info) → Classifier → Entity-relation Triple
(Structured Info)

Relation Type without
Labeled Dataset

Distant Supervision

"If two entities participate in a relation, any sentence that contains those two entities might express that relation." (Mintz, 2009)

Distant Supervision

Freebase

Data(x): <Belgium, Nijlen>
Label(y): /location/contains

Target Corpus (Unlabeled)

DS Label(y): /location/contains

DS Data(x): Nijlen is a municipality located in the Belgian province of Antwerp.

Wrong Labeling

- **Within-Sentence-Bag Level**
  - **Hoffmann et al., ACL 2011.**
  - **Surdean et al., ACL 2012.**
  - **Zeng et al., ACL 2015.**
  - **Li et al., ACL 2016.**

- **Entity-Pair Level**
  - **None**

Wrong Labeling

- Place_of_Death

  i. Some New York city mayors – William O'Dwyer, Vincent R. Impellitteri and Abraham Beame – were born abroad.

  ii. Plenty of local officials have, too, including two New York city mayors, James J. Walker, in 1932, and William O'Dwyer, in 1950.

❖ Entity-Pair Level

Wrong Labeling

- ❖ Most of entity pairs only have several sentences

- ❖ Lots of entity pairs have repetitive sentences

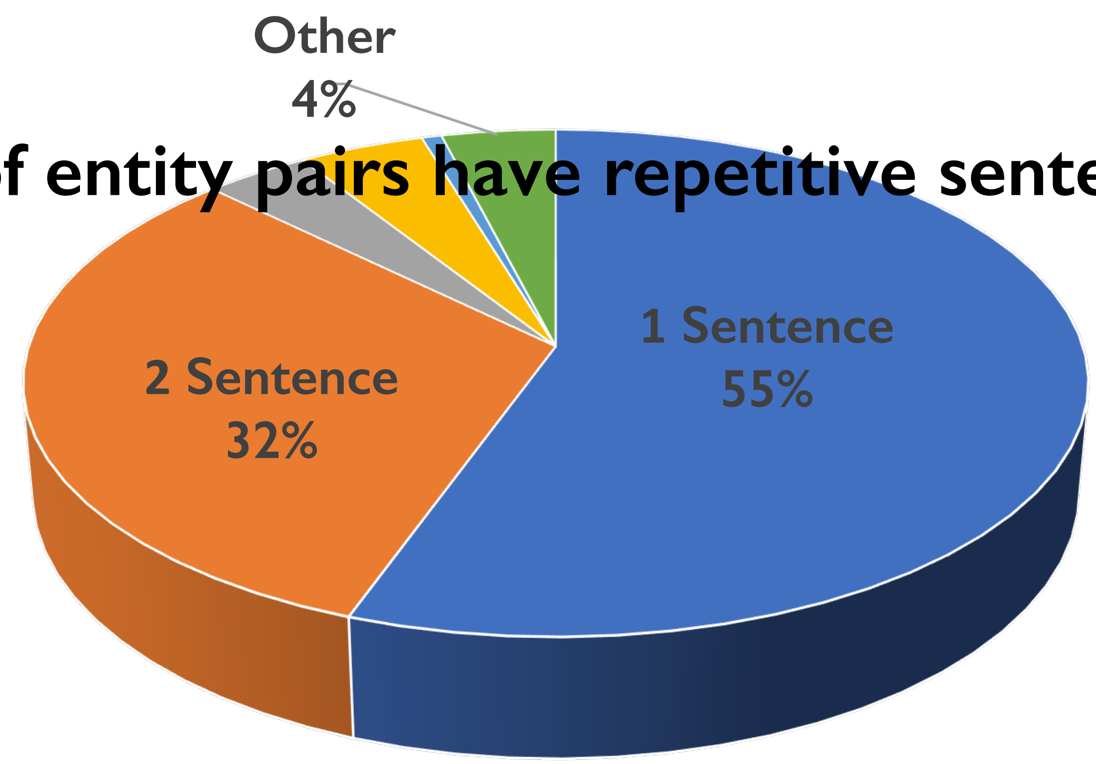

Outline

- Motivation
- Algorithm
- Experiments
- Conclusion

Requirements

Entity-Pair Level Wrong Labeling Problem

⇩

Sentence-Level Indicator

Without Supervised Information

General Purpose and Offline Process

⇩

Learn a Policy to Denoise the Training Data

Overview

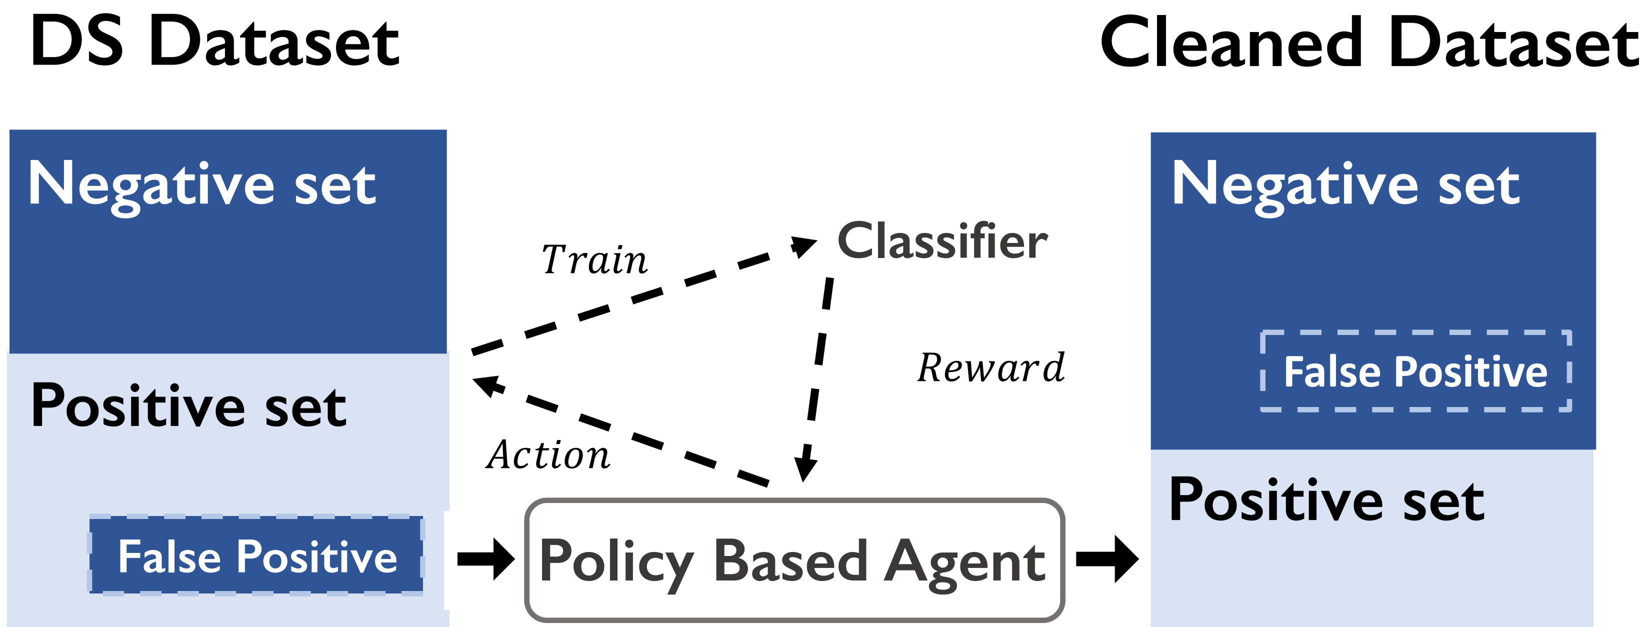
DS Dataset
Cleaned Dataset
Negative set
Positive set
False Positive
Train
Classifier
Reward
Action
Policy Based Agent
Negative set
False Positive
Positive set

Deep Reinforcement Learning

- **State**
  - Sentence vector
  - The average vector of previous removed sentences
- **Action**
  - Remove & retain
- **Reward**
  - ???

Deep Reinforcement Learning

- One relation type has an agent

- Sentence-level
  - Positive: Distantly-supervised positive sentences
  - Negative: Randomly sampled

- Split into training set and validation set

Deep Reinforcement Learning

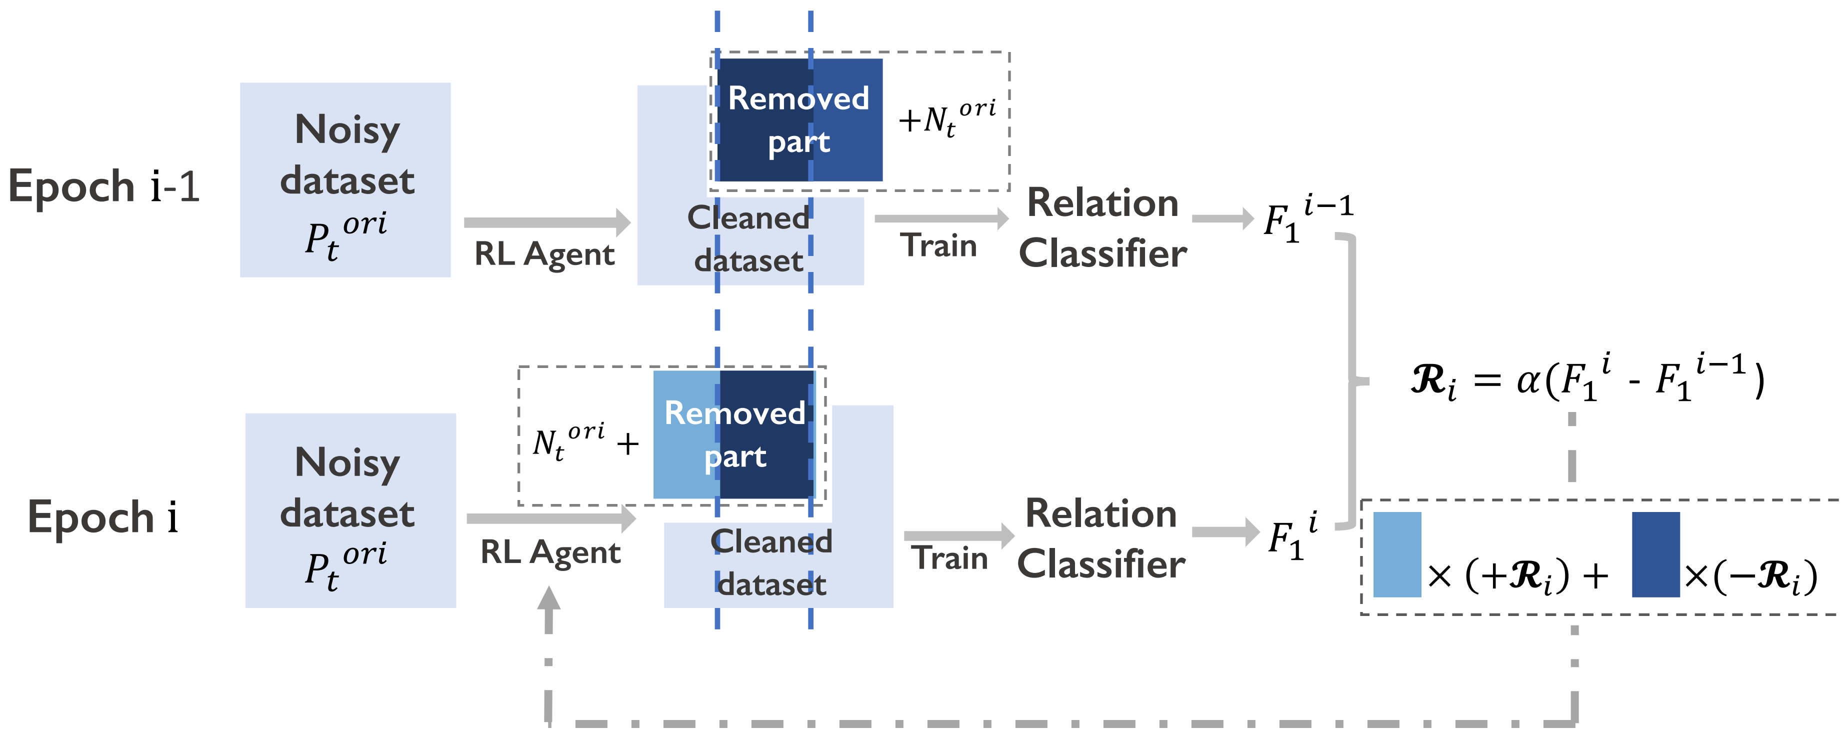

Reward

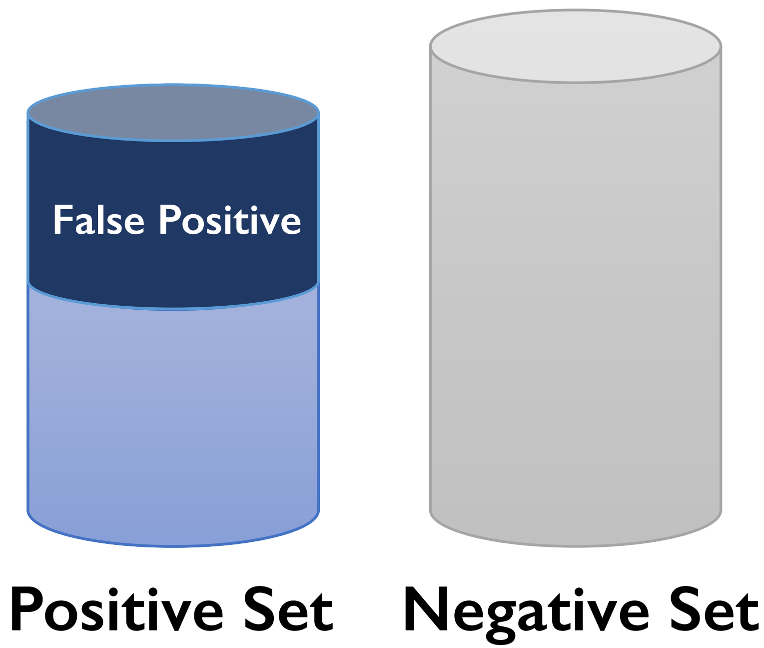

- Accurate
- Steady
- Fast
- Obvious

Reward

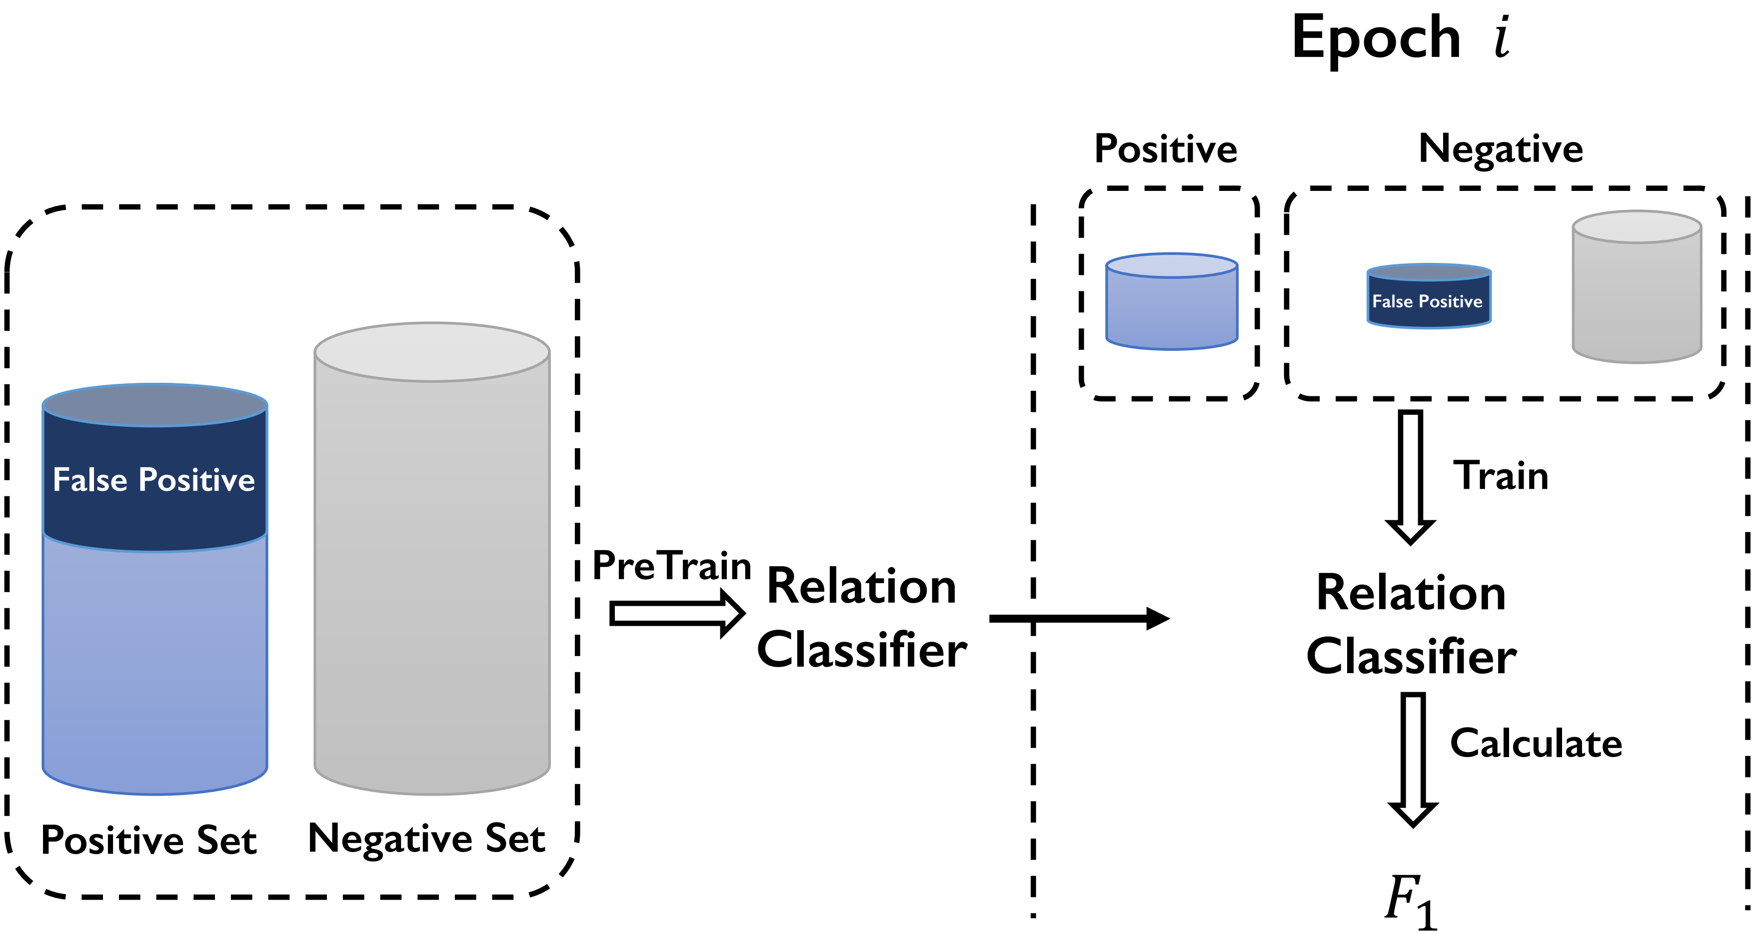
Epoch i
Positive
Negative
False Positive
Train
Relation Classifier
Calculate
F_1
False Positive
Positive Set
Negative Set
PreTrain
Relation Classifier

Outline

- Motivation
- Algorithm
- **Experiments**
- Conclusion

Evaluation on a Synthetic Noise Dataset

- ❖ Dataset: SemEval-2010 Task 8
- ❖ True Positive: Cause-Effect
- ❖ False Positive: Other
- ❖ True Positive + False Positive: 1331 samples

Evaluation on a Synthetic Noise Dataset

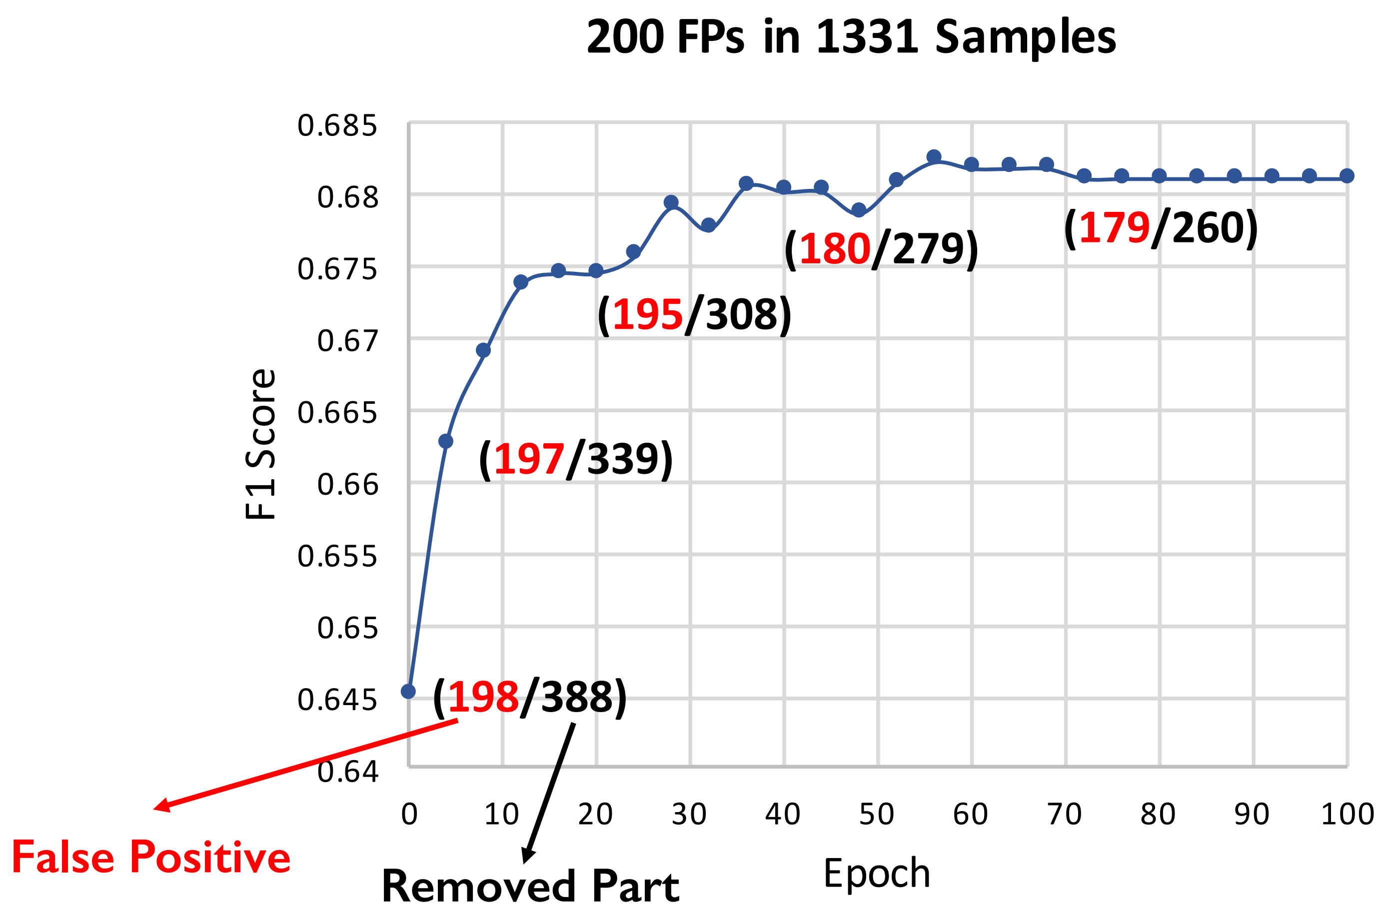

Evaluation on a Synthetic Noise Dataset

0 FPs in 1331 samples

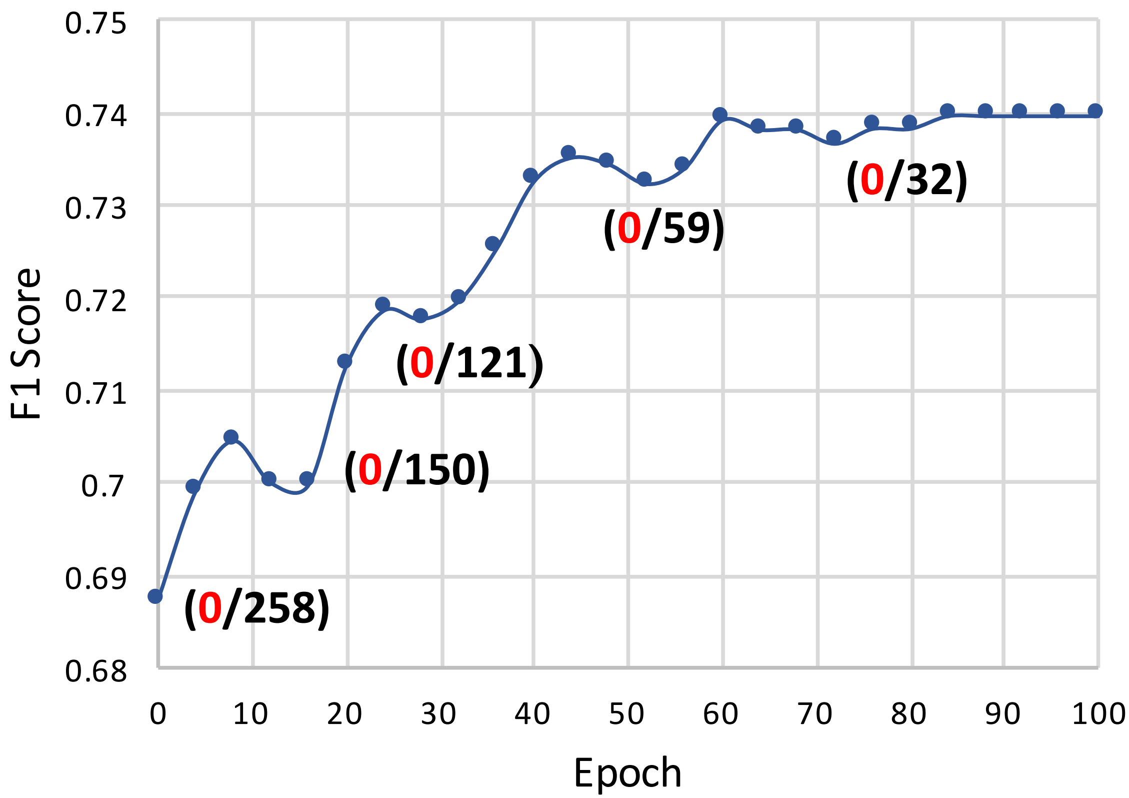

Distant Supervision on NYT Freebase Dataset

- CNN+ONE, PCNN+ONE
  - Distant supervision for relation extraction via piecewise convolutional neural networks. (Zeng et al., 2016)

- CNN+ATT, PCNN+ATT
  - Neural relation extraction with selective attention over instances. (Lin et al., 2016)

Distant Supervision

CNN-based

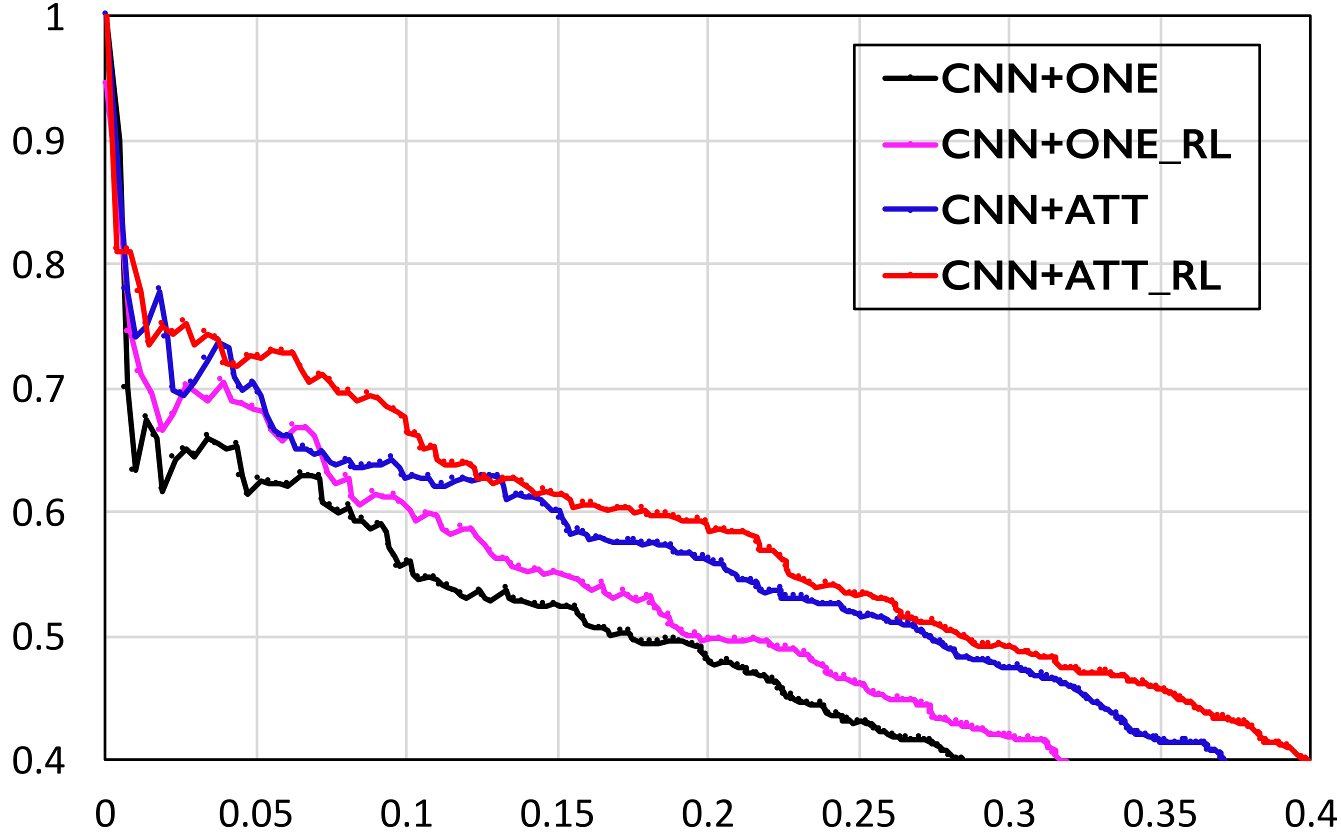

Distant Supervision

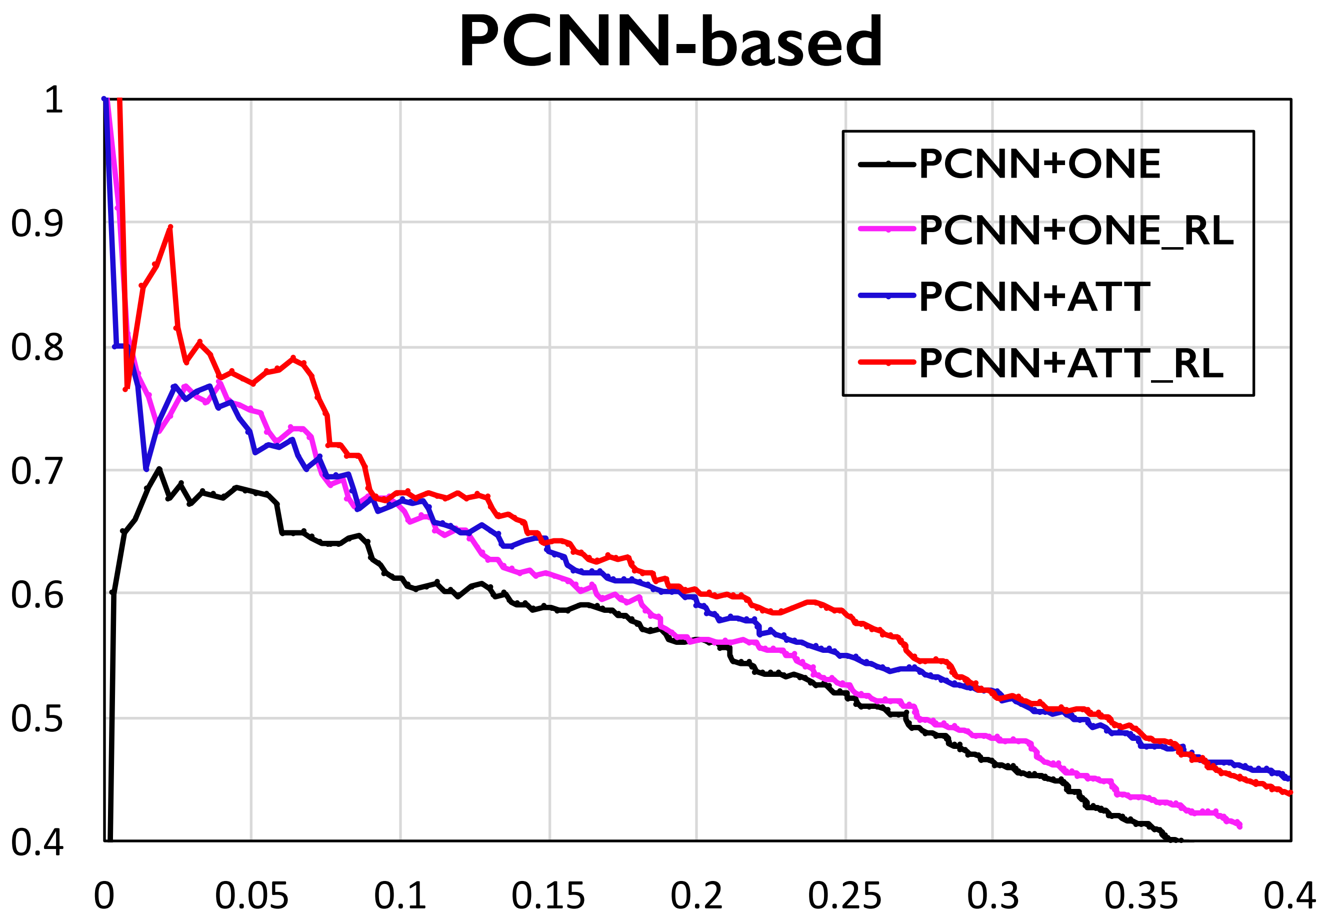

Outline

- Motivation
- Algorithm
- Experiments
- Conclusion

Conclusion

- ❖ We propose a deep reinforcement learning method for robust distant supervision relation Extraction.

- ❖ Our method is model-agnostic.

- ❖ Our method boost the performance of recently proposed neural relation extractors.

DSGAN: Adversarial Learning for Denoising Distantly Supervised Relation Extraction (Qin et al., ACL 2018b)

Distant Supervision Data Distribution

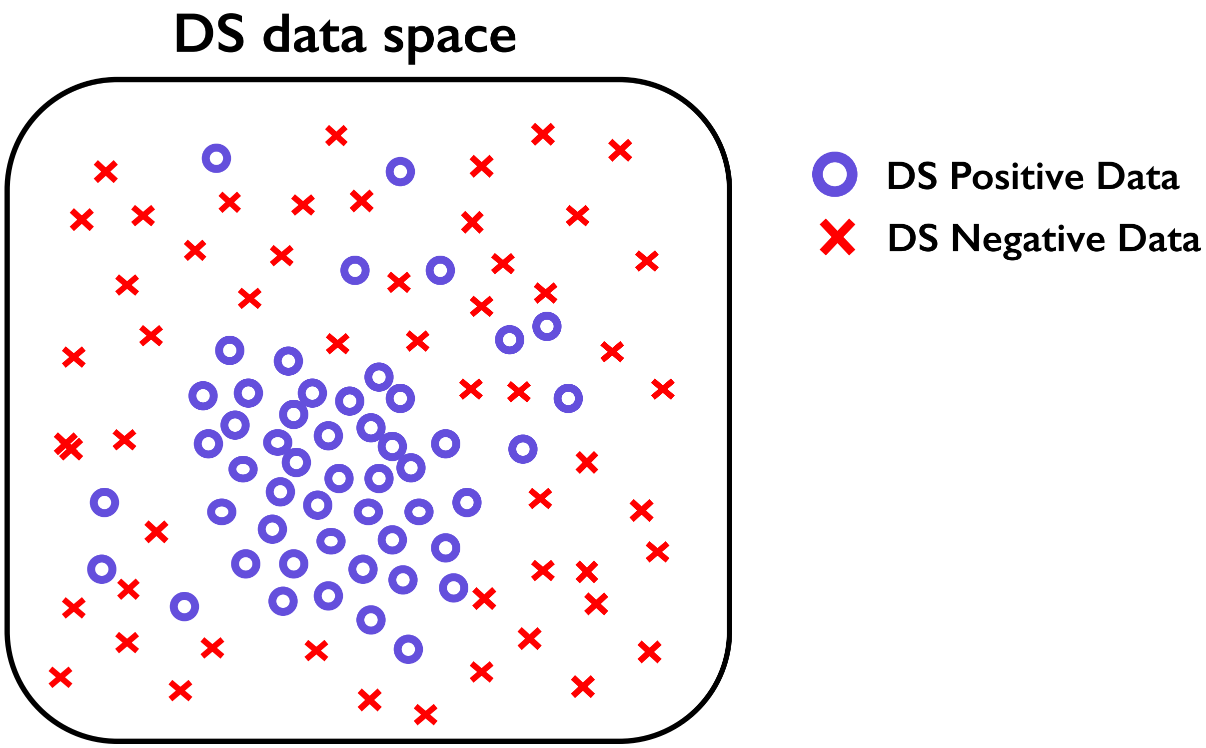

Data Distribution

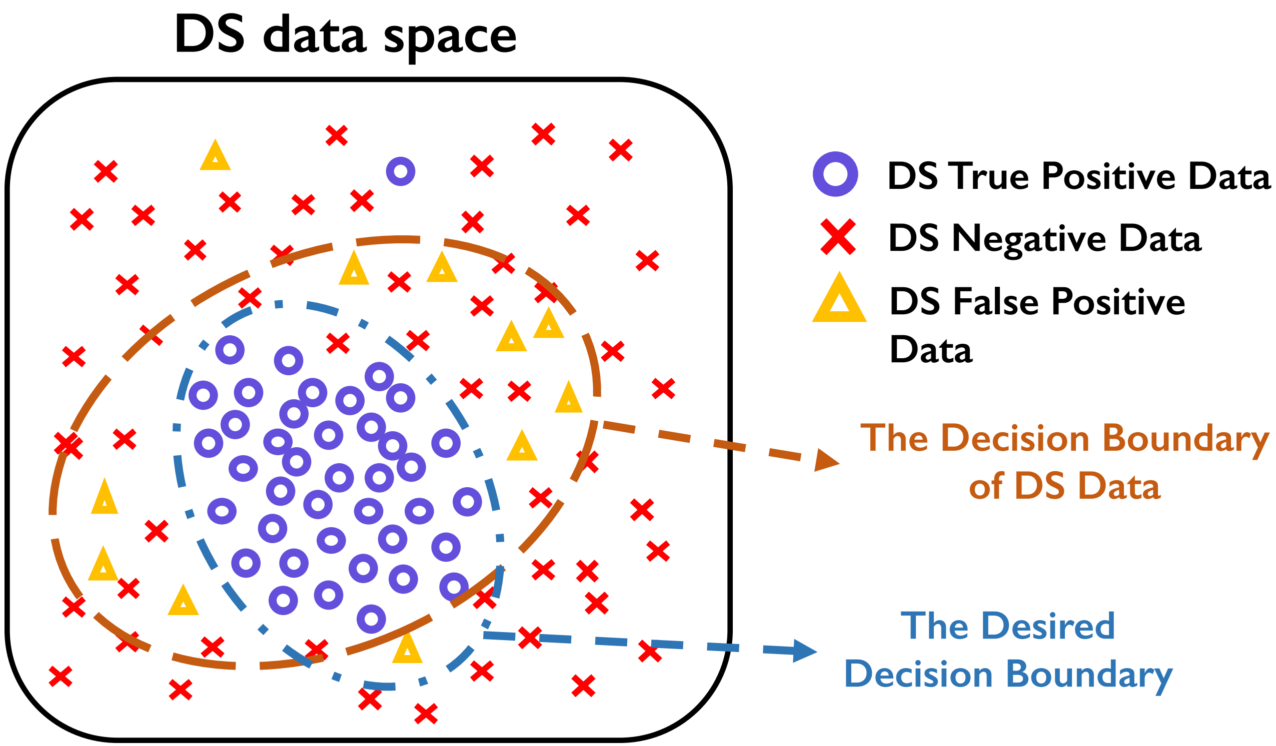

Adversarial Training

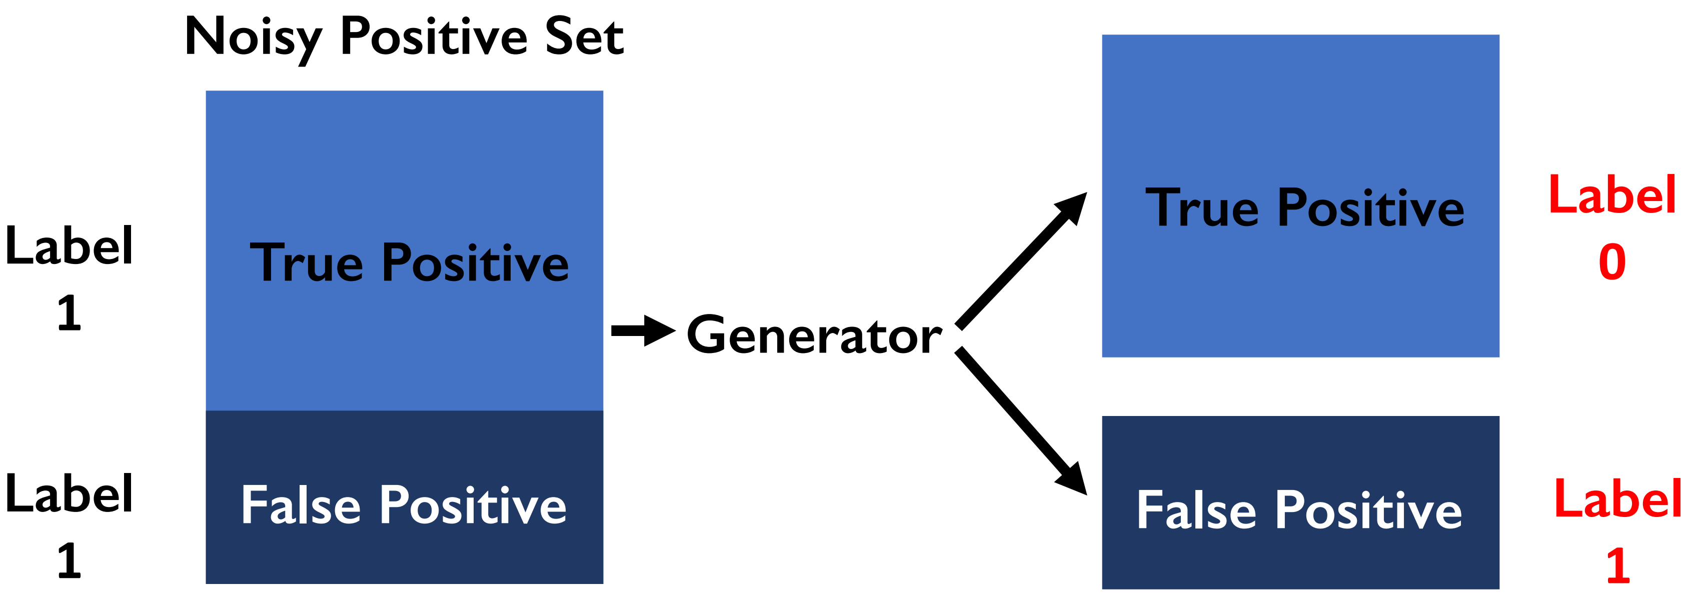

DSGAN (Qin et al., ACL 2018b)

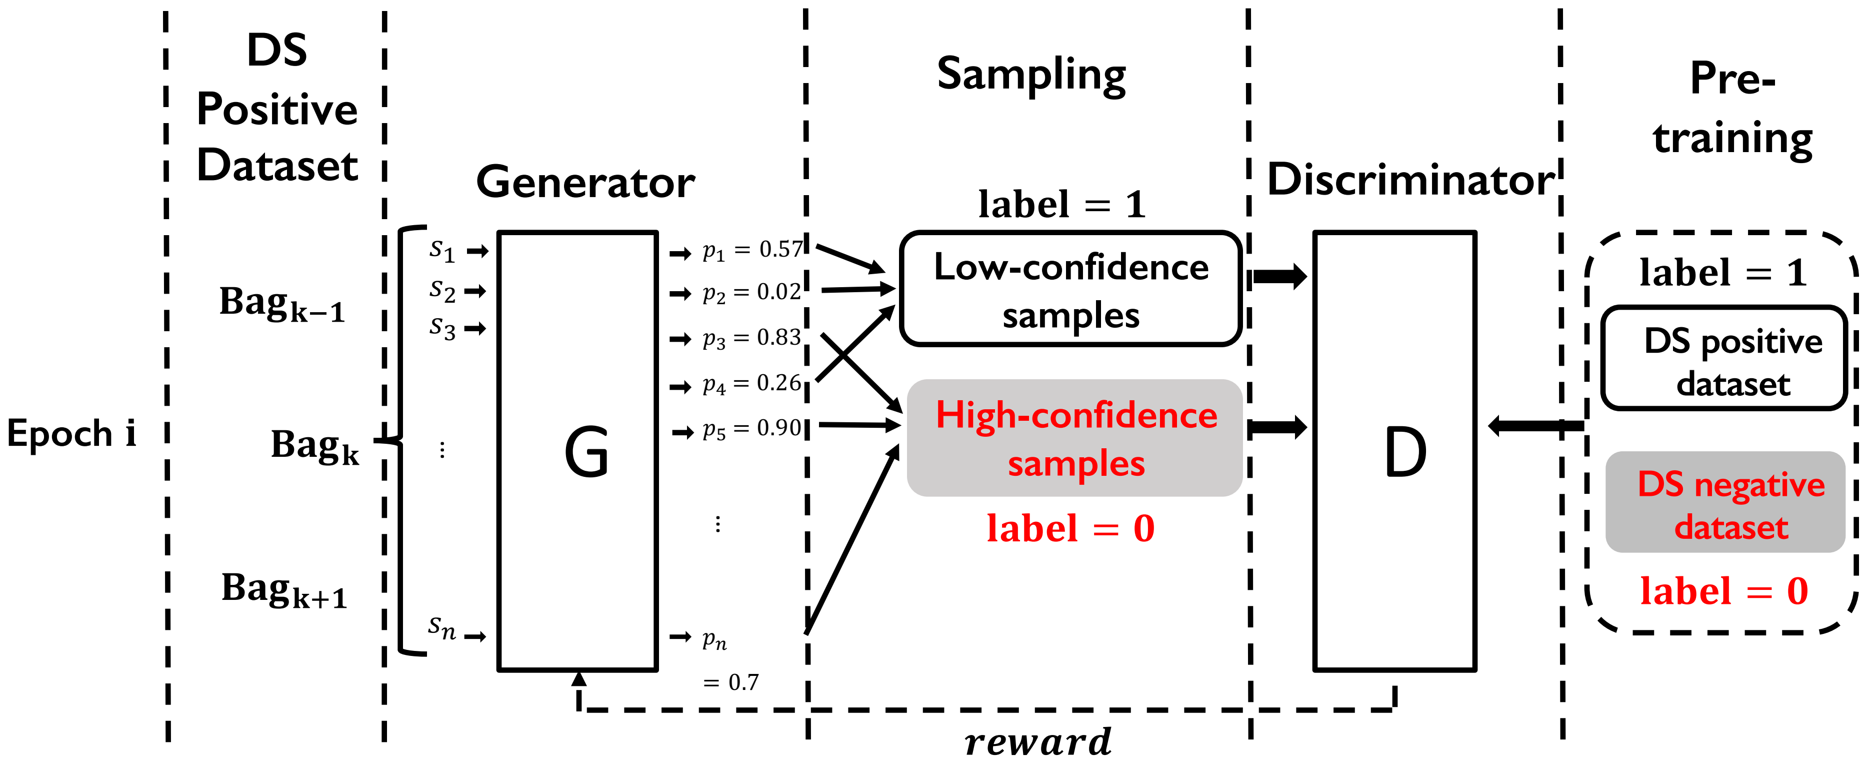

Characteristics

❖ **Sentence-Level Noise Reduction**

❖ **Training Without Supervised Information**

❖ **Model-Agnostic**

Distant Supervision Relation Extraction

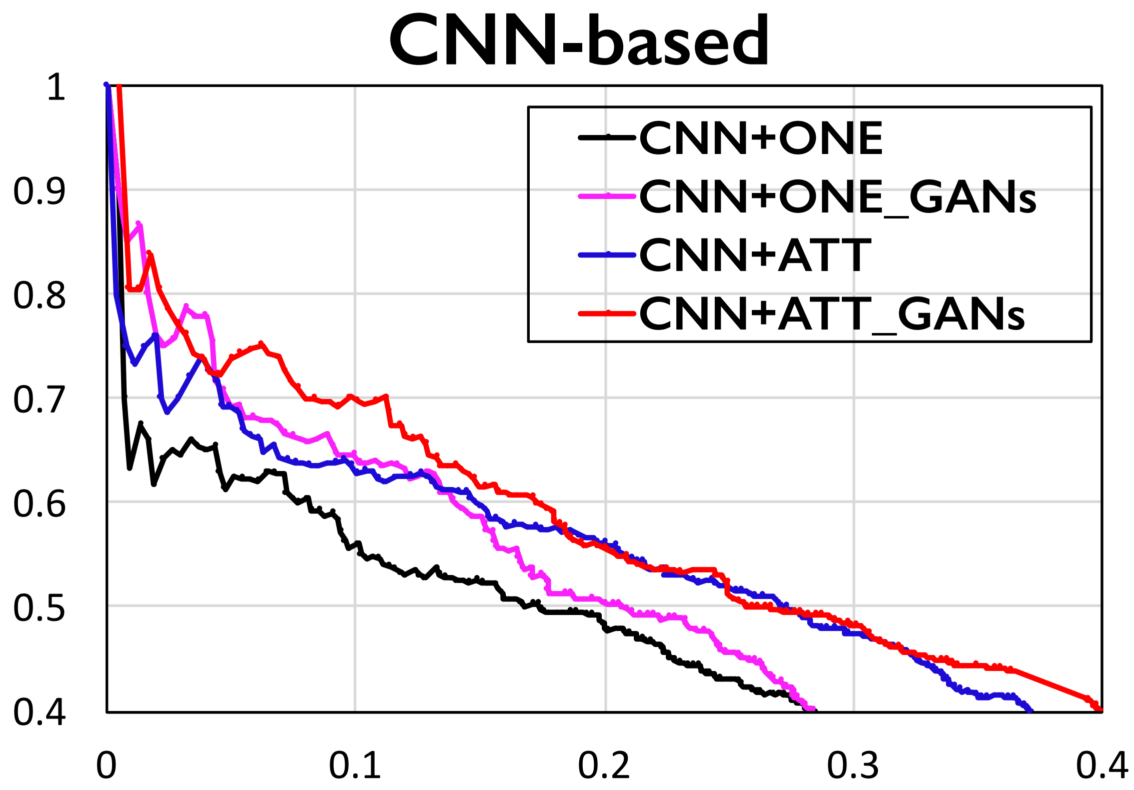

Distant Supervision Relation Extraction

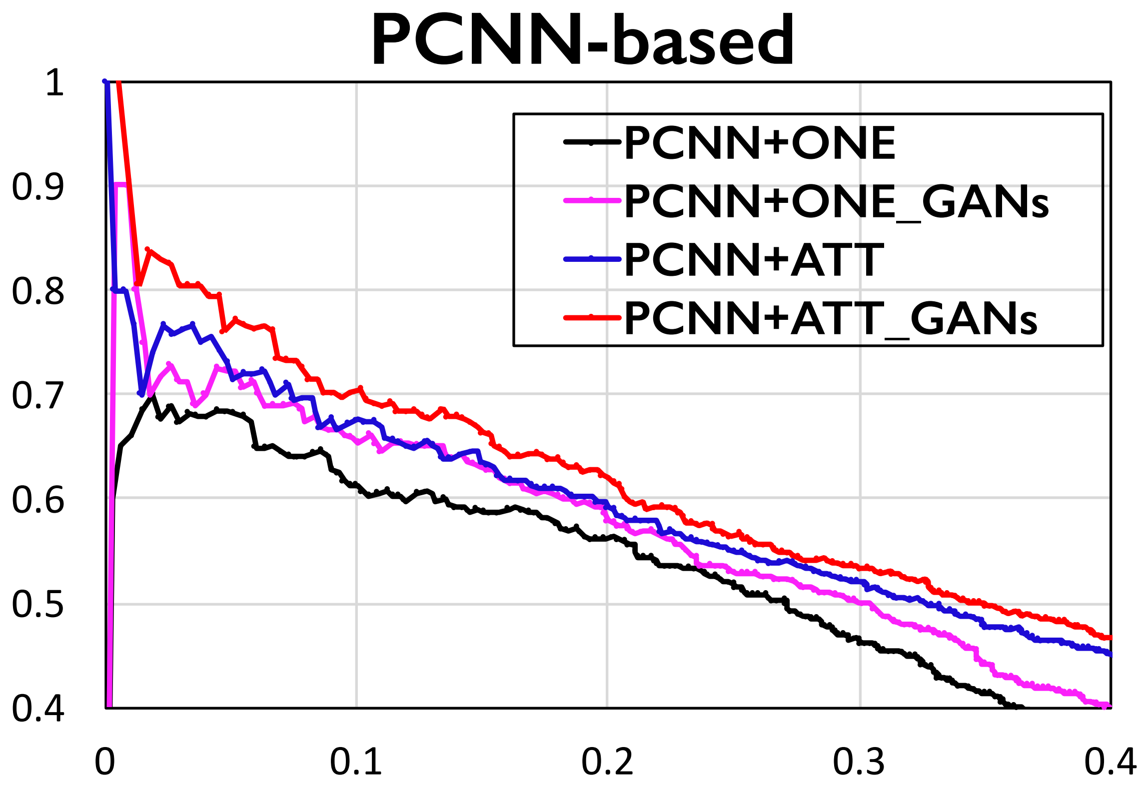

Conclusion

- We introduce Reinforced Co-Training, a new approach that combines reinforcement learning and semi-supervised learning.
- We show that in weakly-supervised relation extraction, reinforcement learning can be utilized to de-noise the training signals.
- Adversarial learning serves as a joint learning framework, and it can also be applied to de-noising distantly supervised IE data.

Thanks!

http://nlp.cs.ucsb.edu