Robust Distant Supervision Relation Extraction via Deep Reinforcement Learning

BUPT

UC SANTA BARBARA

Pengda Qin, Weiran Xu and William Wang

Outline

- Motivation
- Algorithm
- Experiments
- Conclusion

Outline

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

Relation Extraction

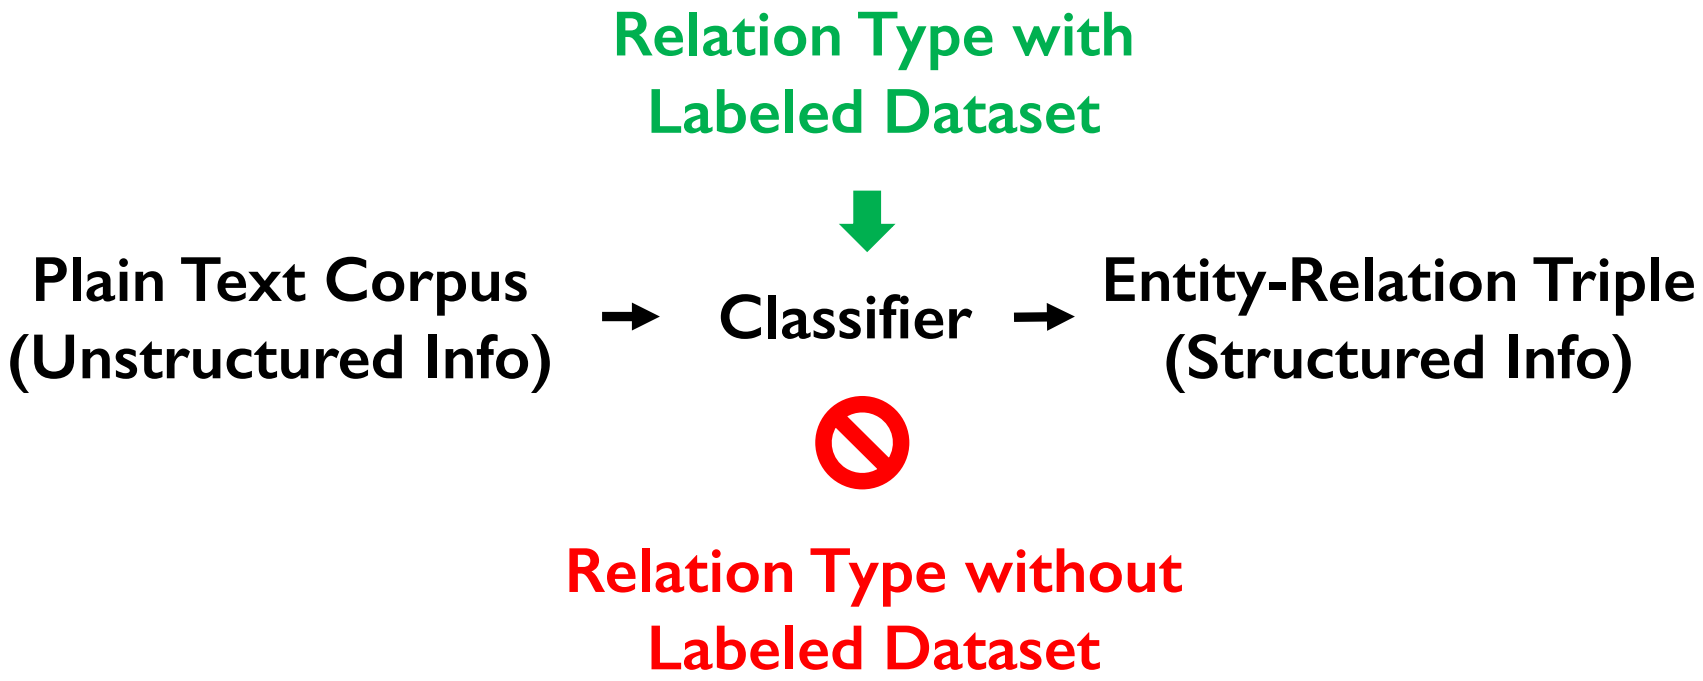

Distant Supervision

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

Distant Supervision

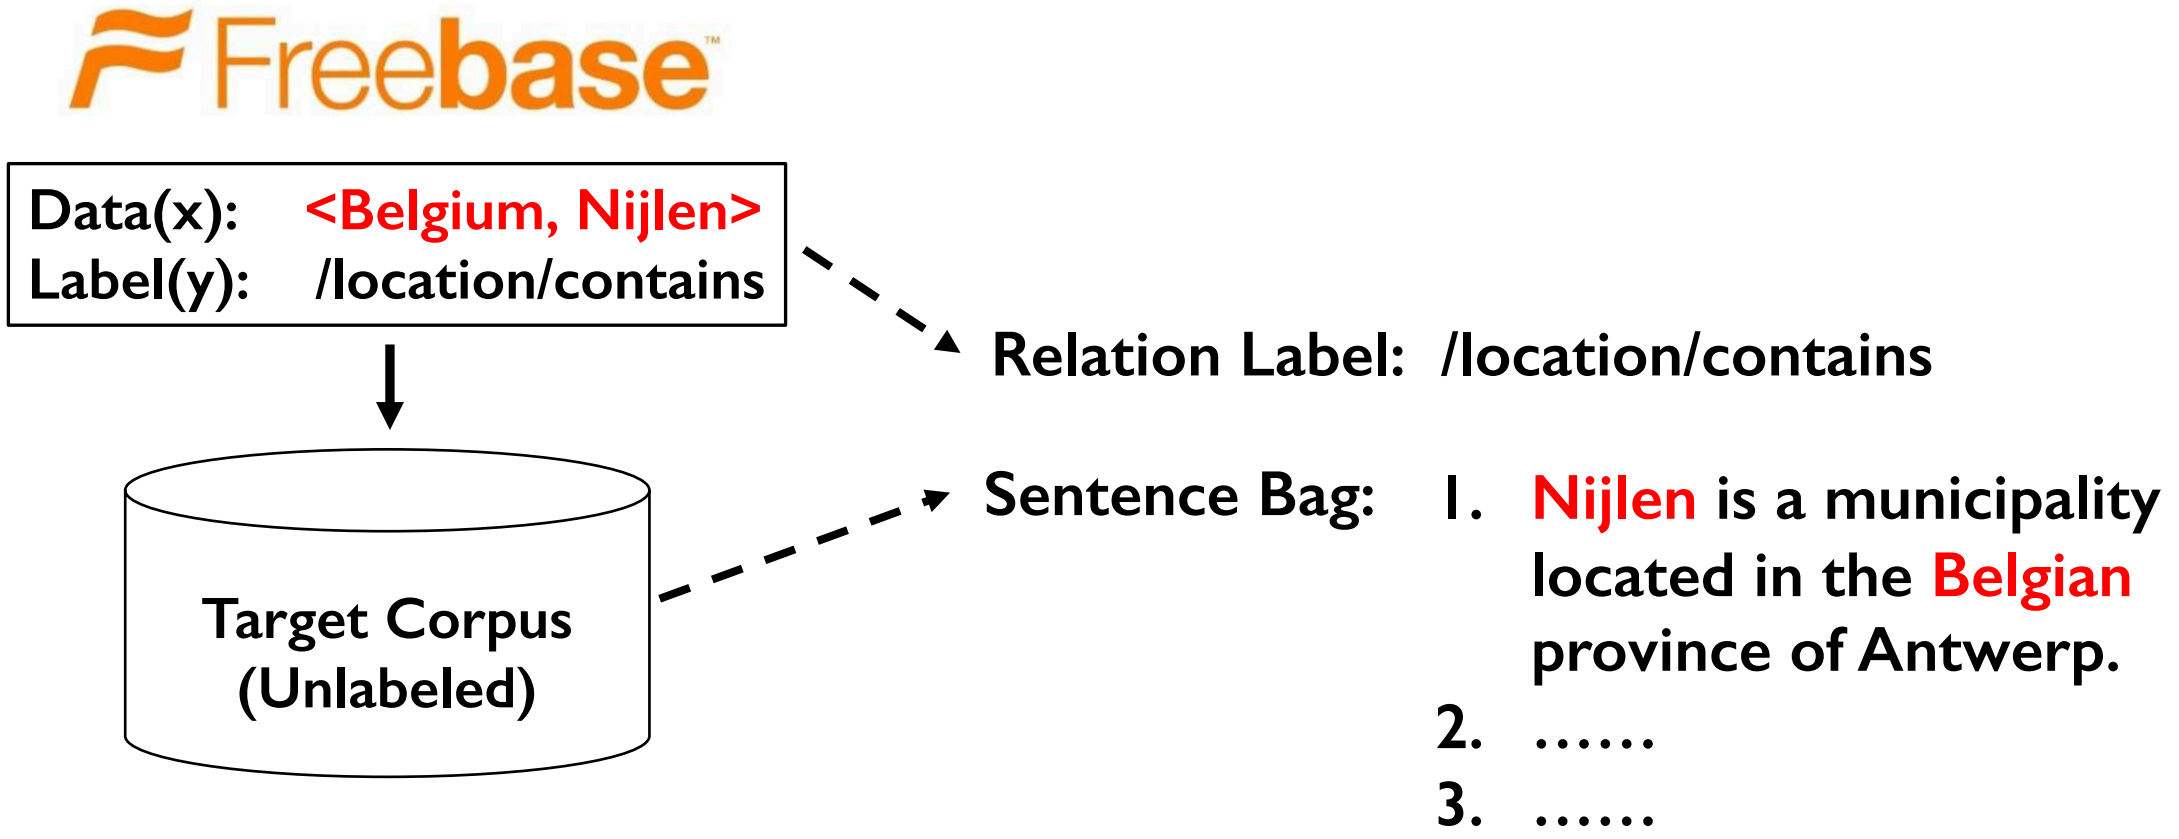

Wrong Labeling

- 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 (William O'Dwyer, New York city)

  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

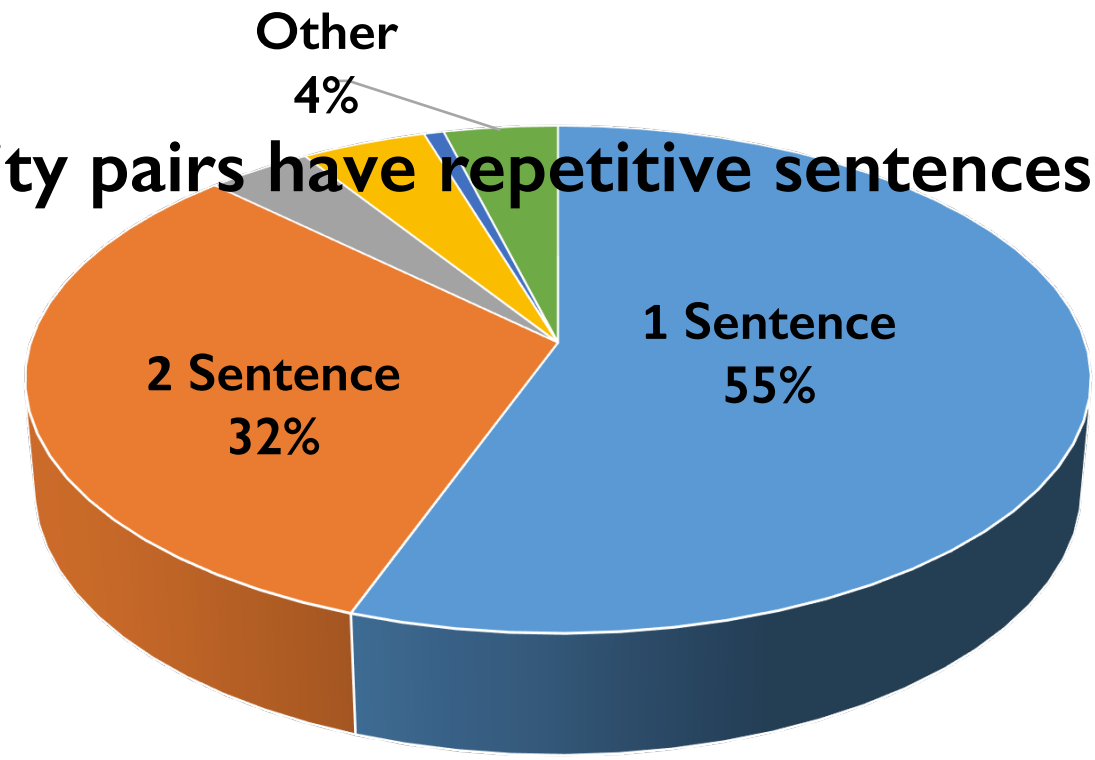

Outline

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

Overview

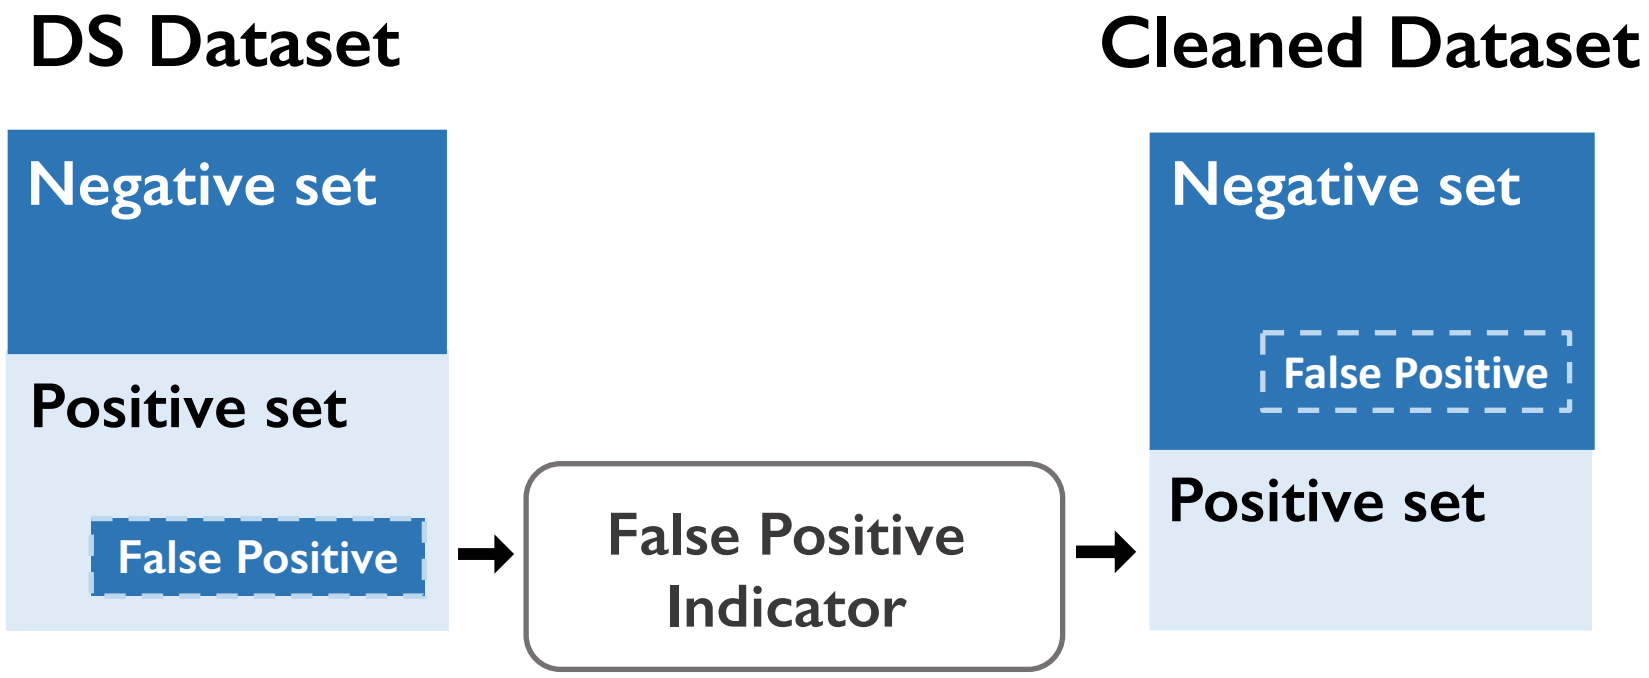
DS Dataset
Negative set
Positive set
False Positive
False Positive Indicator
Cleaned Dataset
Negative set
False Positive
Positive set

Requirements

False-Positive Indicator

⇩

Sentence-Level Indicator

Without Supervised Information

General Purpose and Offline Process

⇩

Learn a Policy to Denoise the Training Data

Overview

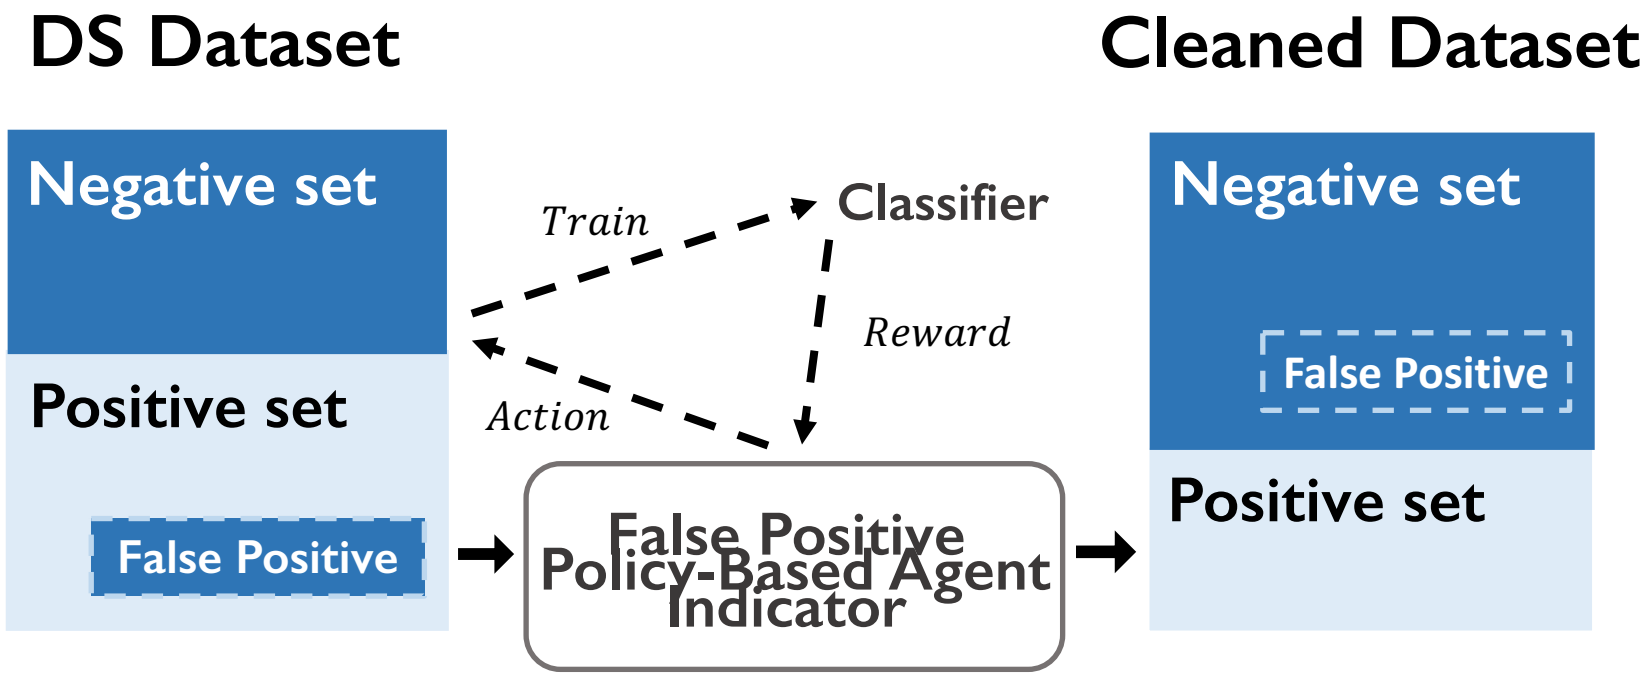
DS Dataset
Cleaned Dataset
Negative set
Positive set
False Positive
Train
Classifier
Reward
Action
False Positive
Policy-Based Agent
Indicator
Negative set
False Positive
Positive set

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: Sampled from other relations
- Split into training set and validation set

Deep Reinforcement Learning

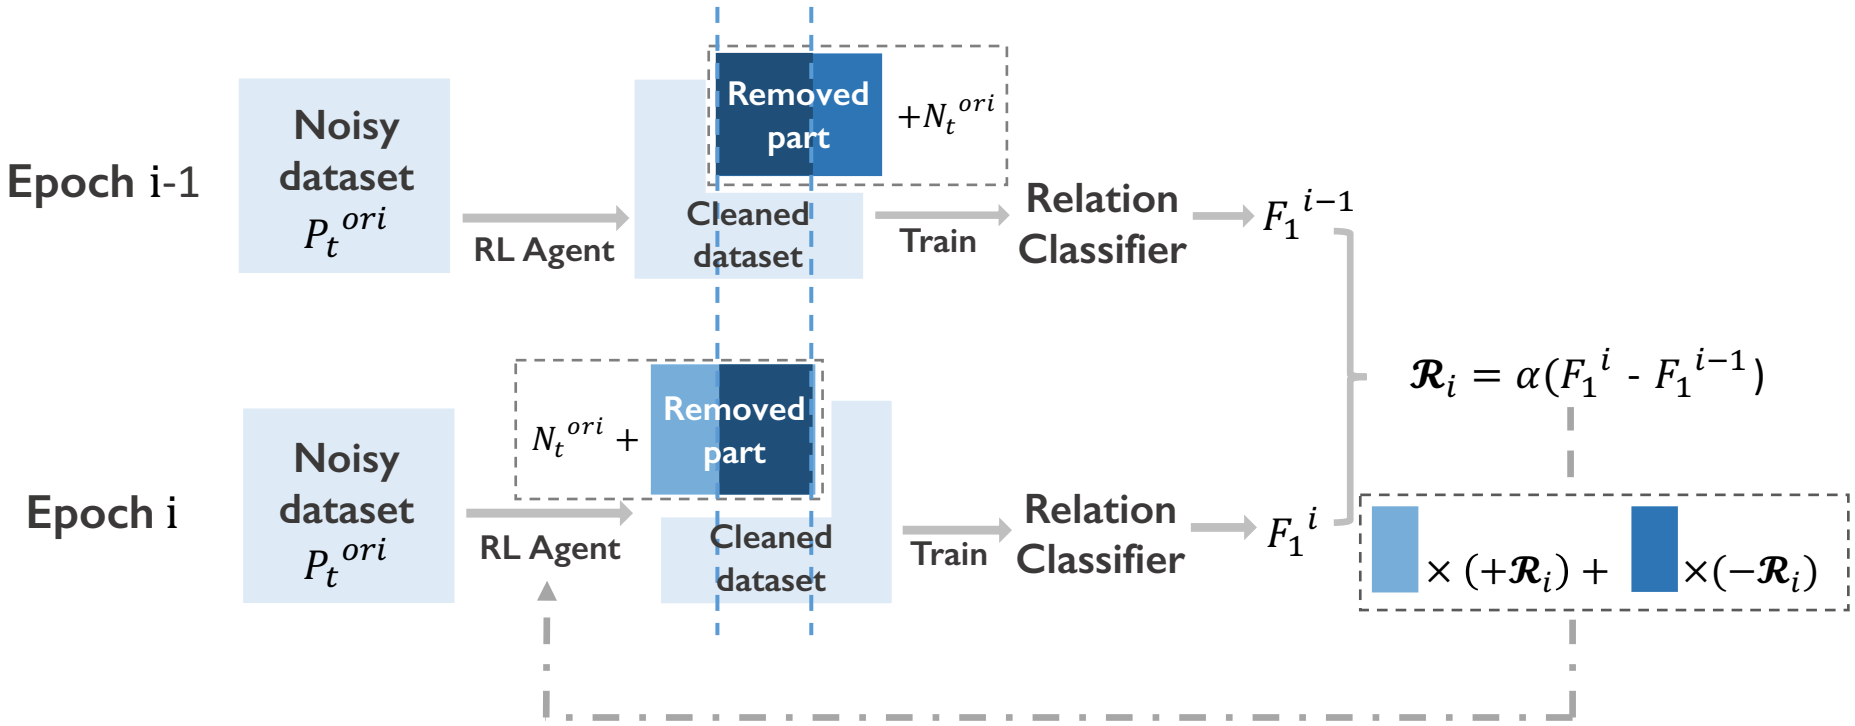

Reward

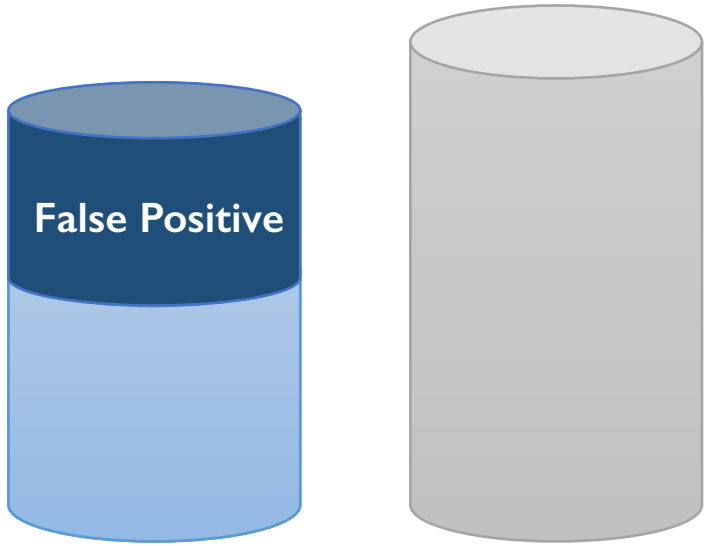

Positive Set Negative Set

- Accurate
- Steady
- Fast
- Obvious

Reward

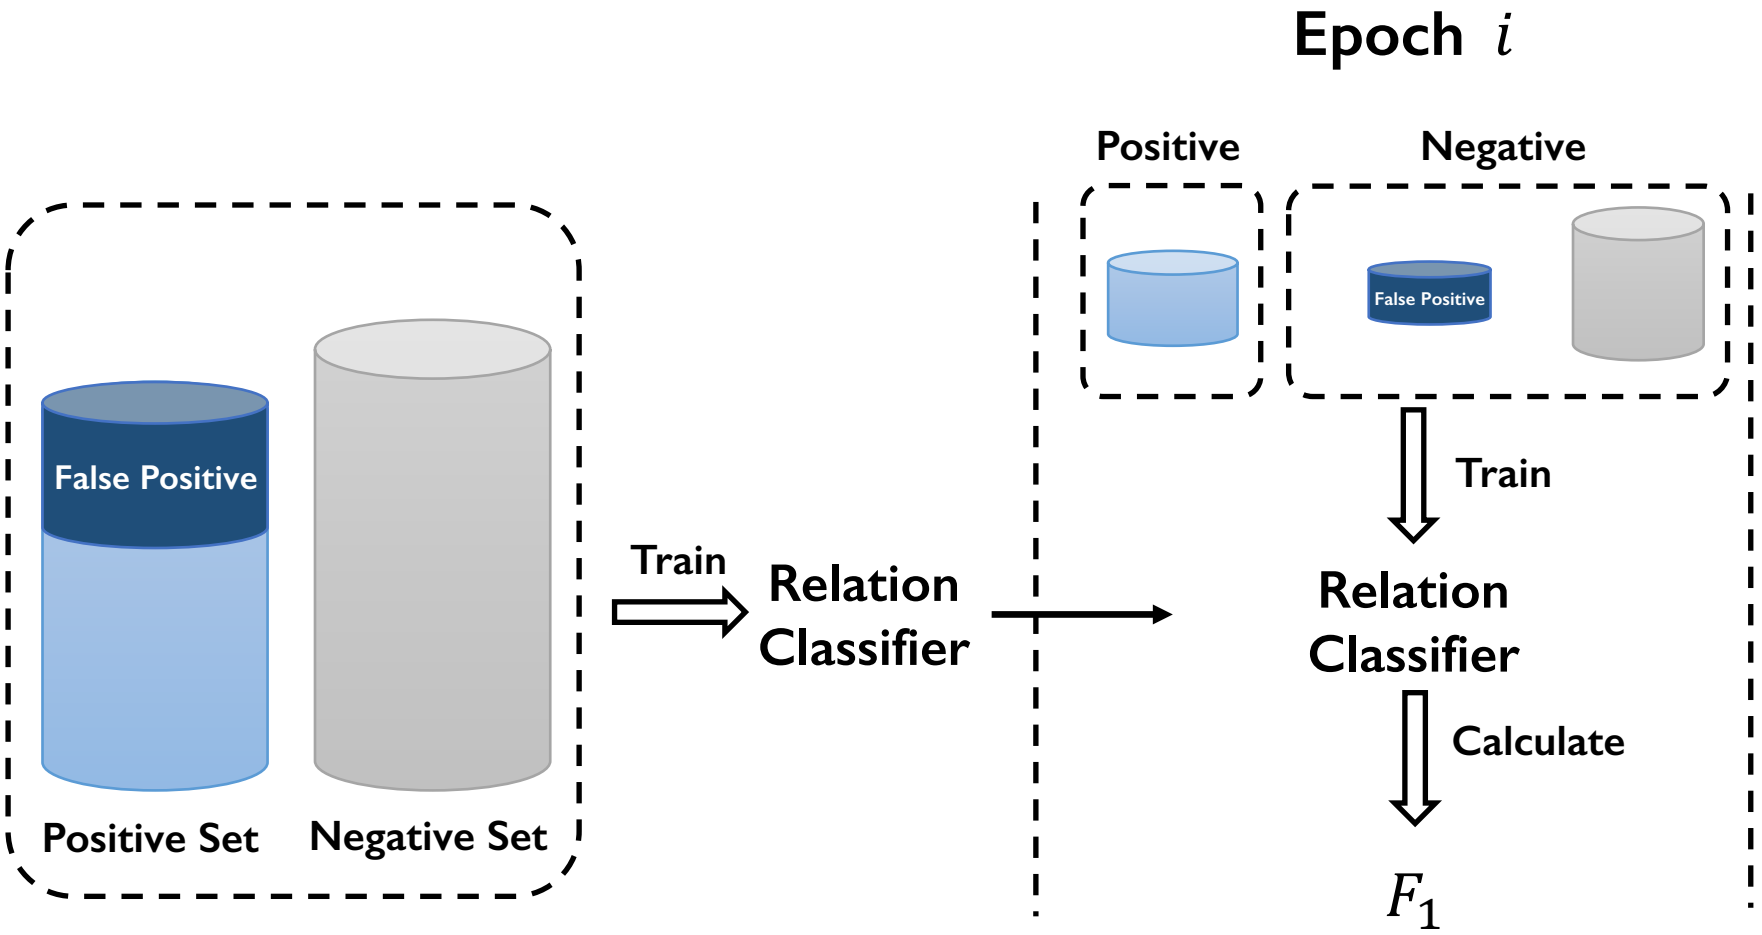
Epoch i
Positive
Negative
False Positive
Positive Set
Negative Set
Train
Relation Classifier
Train
Relation Classifier
Calculate
F_1

Outline

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

Evaluation on a Synthetic Noise Dataset

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

Evaluation on a Synthetic Noise Dataset

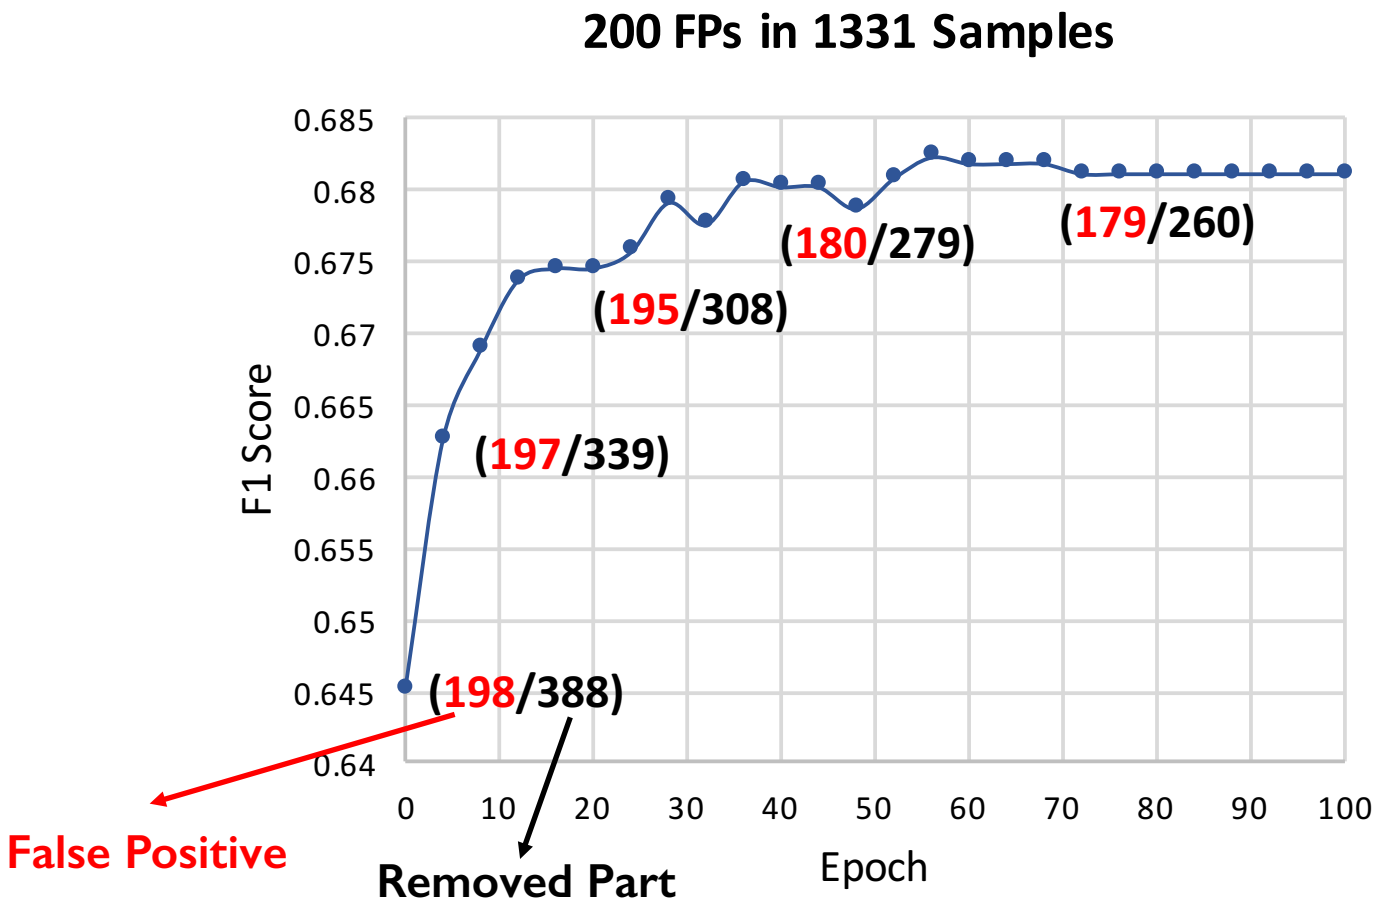

Evaluation on a Synthetic Noise Dataset

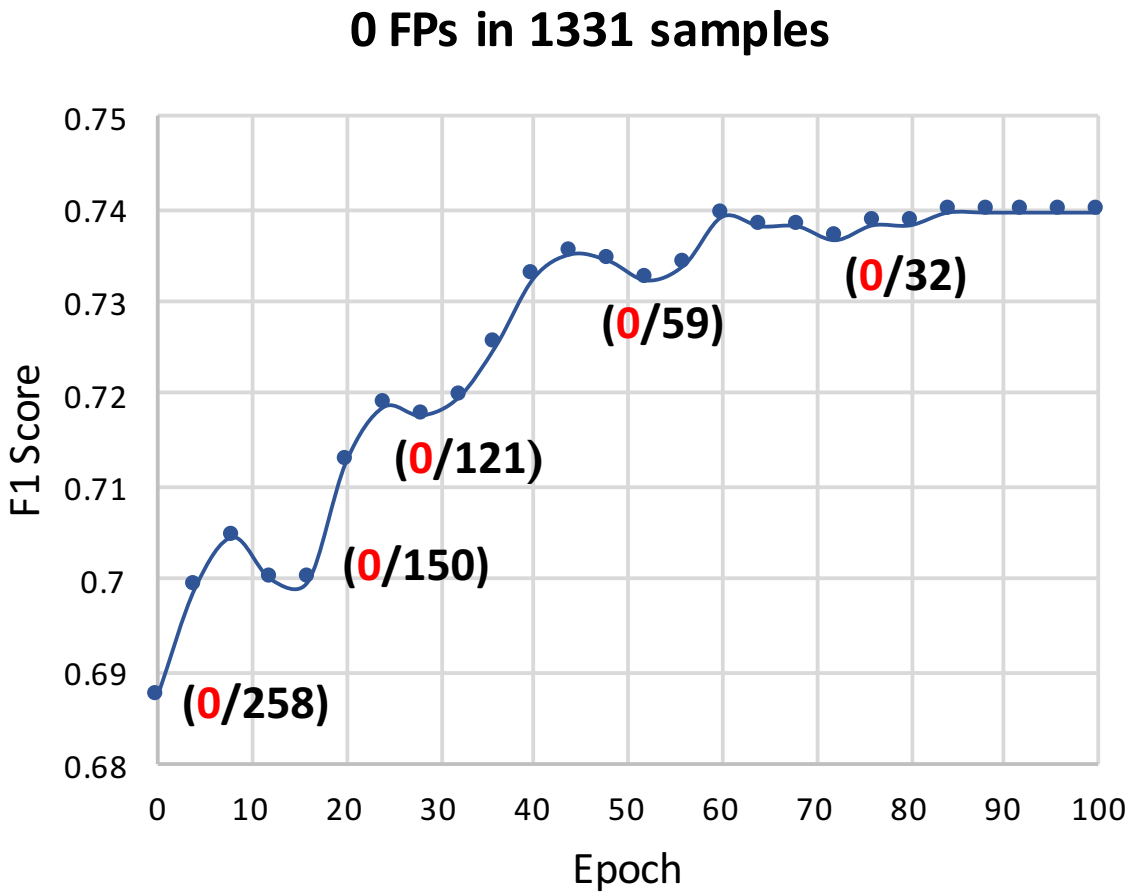

Distant Supervision

- **Dataset: Riedel et al., 2010**
  - http://iesl.cs.umass.edu/riedel/ecml/
- **CNN+ONE, PCNN+ONE**
  - **Distant supervision for relation extraction via piecewise convolutional neural networks. (Zeng et al., 2015)**
- **CNN+ATT, PCNN+ATT**
  - **Neural relation extraction with selective attention over instances. (Lin et al., 2016)**

Distant Supervision

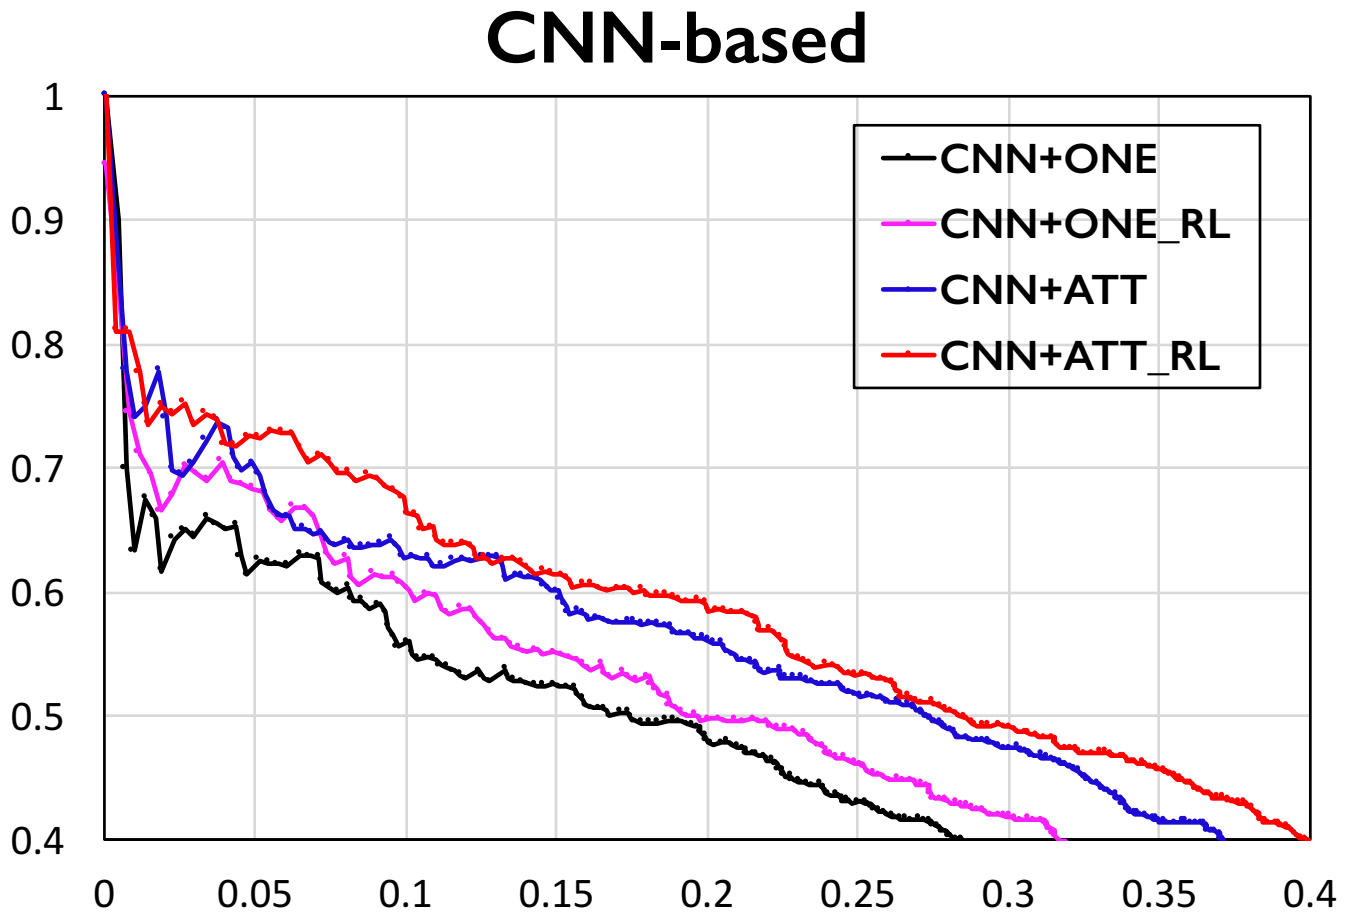

Distant Supervision

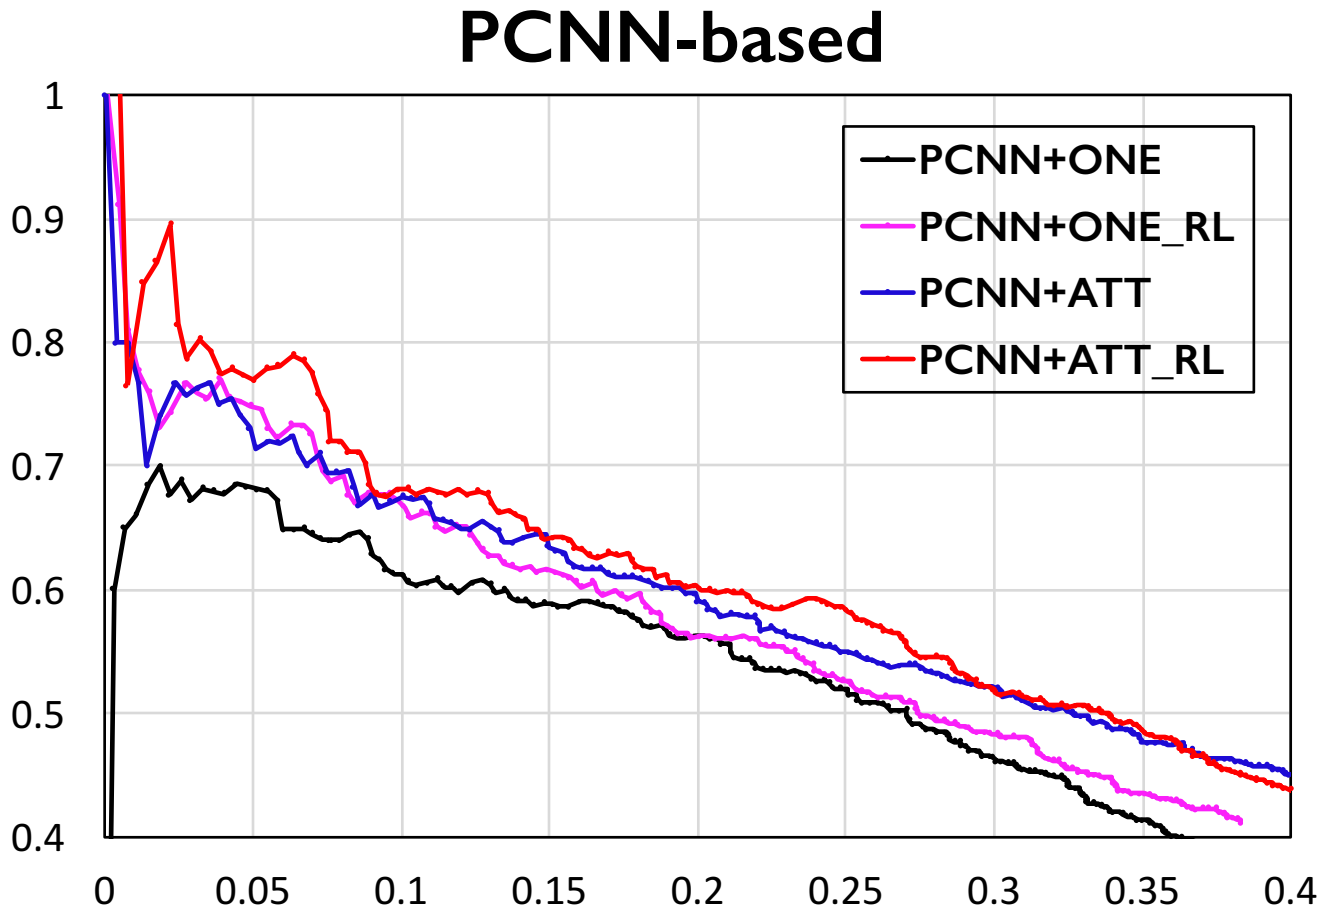

Outline

- 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.

Thank you!

Q&A