Adversarial training for multi-context joint entity and relation extraction

Giannis Bekoulis Johannes Deleu Thomas Demeester Chris Develder

Ghent University - imec, IDLab
Department of Information Technology

Problem
Joint entity recognition and relation extraction

Example

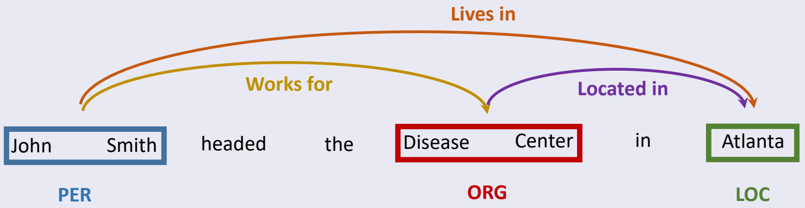

Tasks

1. Entity recognition
2. Relation extraction

Our idea

(1) + (2) with adversarial training

Outline

Idea

Regularization method to improve the robustness of neural network methods by adding small perturbations in the training data.

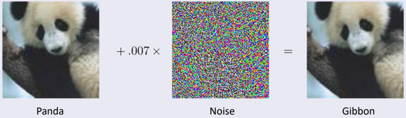

Figure: Goodfellow et al. (2015).

In NLP

- Text classification (Miyato et al., 2017)
- Relation extraction (Wu et al., 2017)
- POS tagging (Yasunaga et al., 2018)

Outline

Baseline model

Joint architecture

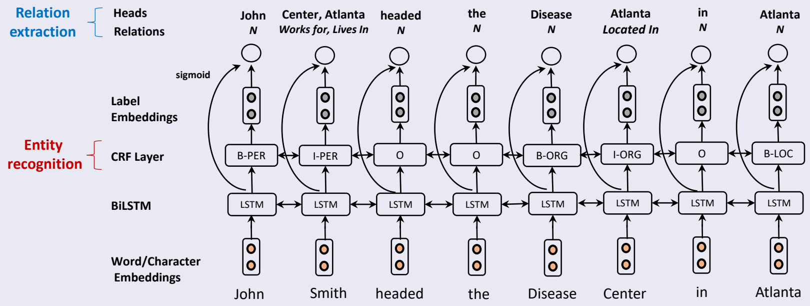

Relation extraction

Multi-head selection

Outline

Training

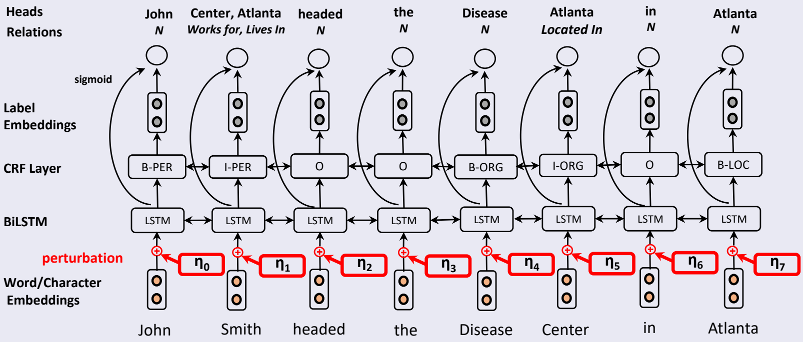

- Adding worst case noise from the perspective of the loss
- $\eta_{adv} = \text{argmax}_{\|\eta\| \leq \epsilon} \mathcal{L}_{\text{Joint}}(w + \eta; \hat{\theta}) \rightarrow \eta_{adv} = \epsilon g / \|g\|$ with $g = \nabla_w \mathcal{L}_{\text{Joint}}(w; \hat{\theta})$ (Goodfellow et al., 2015)

Outline

Experimental results
Datasets

ACE04

EMP-ORG

The United Nations' Secretary-General, Kofi Annan, opened the exhibition with a speech.

ORG PER

CoNLL04

Works for

Robert Bernero, chief of waste disposal for the NRC.

PER ORG

Experimental results

Datasets

DREC

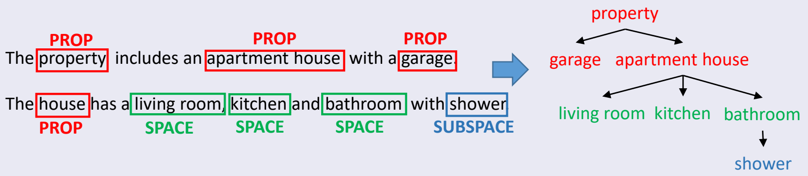

ADE

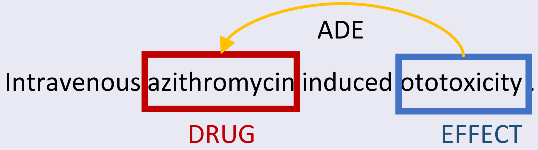

Experimental results

Performance close or better compared to feature based models

	Settings	Features	Entity	Relation	Overall F_1	
ACE 04	Miwa and Bansal (2016)	✓	81.80	48.40	65.10	
	Katiyar and Cardie (2017)	✗	79.60	45.70	62.65	$\simeq$
	baseline	✗	81.16	47.14	64.15	
	baseline + AT	✗	**81.64**	**47.45**	**64.54**	
CoNLL 04	Gupta et al. (2016)	✓	92.40	69.90	81.15	
	Gupta et al. (2016)	✗	88.80	58.30	73.60	
	Adel and Schütze (2017)	✗	82.10	62.50	72.30	$\simeq$
	baseline EC	✗	**93.26**	67.01	80.14	
	baseline EC + AT	✗	93.04	**67.99**	**80.51**	
	Miwa and Sasaki (2014)	✓	80.70	61.00	70.85	
	baseline	✗	83.04	61.04	72.04	$>$
	baseline + AT	✗	**83.61**	**61.95**	**72.78**	
DREC	Bekoulis et al. (2018)	✗	79.11	49.70	64.41	
	baseline	✗	82.30	52.81	67.56	
	baseline + AT	✗	**82.96**	**53.87**	**68.42**	
	baseline	✗	81.39	52.26	66.83	
	baseline + AT	✗	**82.04**	**53.12**	**67.58**	
ADE	Li et al. (2016)	✓	79.50	63.40	71.45	
	Li et al. (2017)	✓	84.60	71.40	78.00	
	baseline	✗	86.40	74.58	80.49	$\gg$
	baseline + AT	✗	**86.73**	**75.52**	**81.13**	

	Settings	Features	Entity ⇓	Relation ⇓	Overall F_1
ACE 04	Miwa and Bansal (2016)	✓	81.80	48.40	65.10
	Katiyar and Cardie (2017)	✗	79.60	45.70	62.65
	baseline	✗	81.16	47.14	64.15
	baseline + AT	✗	**81.64**	**47.45**	**64.54**
CoNLL 04	Gupta et al. (2016)	✓	92.40	69.90	81.15
	Gupta et al. (2016)	✗	88.80	58.30	73.60
	Adel and Schütze (2017)	✗	82.10	62.50	72.30
	baseline EC	✗	**93.26**	67.01	80.14
	baseline EC + AT	✗	93.04	**67.99**	**80.51**
	Miwa and Sasaki (2014)	✓	80.70	61.00	70.85
	baseline	✗	83.04	61.04	72.04
	baseline + AT	✗	**83.61**	**61.95**	**72.78**
DREC	Bekoulis et al. (2018)	✗	79.11	49.70	64.41
	baseline	✗	82.30	52.81	67.56
	baseline + AT	✗	**82.96**	**53.87**	**68.42**
	baseline	✗	81.39	52.26	66.83
	baseline + AT	✗	**82.04**	**53.12**	**67.58**
ADE	Li et al. (2016)	✓	79.50	63.40	71.45
	Li et al. (2017)	✓	84.60	71.40	78.00
	baseline	✗	86.40	74.58	80.49
	baseline + AT	✗	**86.73**	**75.52**	**81.13**

Experimental results

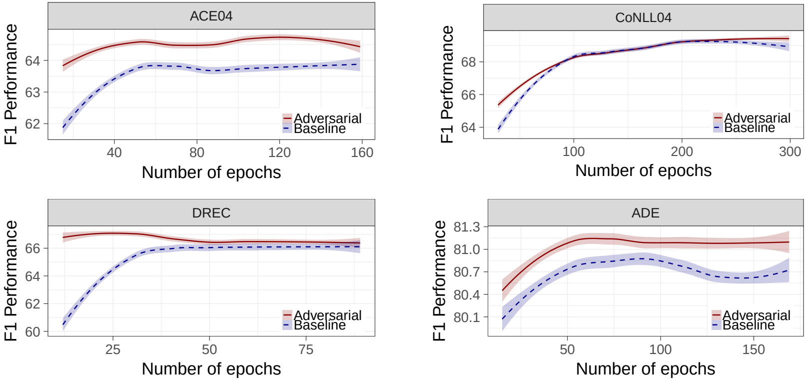

- AT outperforms the neural baseline model consistently across multiple and diverse datasets
- Improvement of AT depends on the dataset

Outline

Conclusions

- Investigate the consistent effectiveness of AT as a regularization method over a multi-context baseline joint model
- Large scale experimental evaluation
- Improvement for each task separately, as well as the overall performance of the baseline joint model

Code: `https://github.com/bekou/multihead_joint_entity_relation_extraction`

References I

Adel, H. and Schütze, H. (2017). Global normalization of convolutional neural networks for joint entity and relation classification. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, Copenhagen, Denmark. Association for Computational Linguistics.

Bekoulis, G., Deleu, J., Demeester, T., and Develder, C. (2018). An attentive neural architecture for joint segmentation and parsing and its application to real estate ads. *Expert Systems with Applications*, 102:100–112.

Goodfellow, I., Shlens, J., and Szegedy, C. (2015). Explaining and harnessing adversarial examples. In *Proceedings of the International Conference on Learning Representations*, San Diego, USA.

Gupta, P., Schütze, H., and Andrassy, B. (2016). Table filling multi-task recurrent neural network for joint entity and relation extraction. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 2537–2547.

Katiyar, A. and 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 (Volume 1: Long Papers)*, Vancouver, Canada.

Li, F., Zhang, M., Fu, G., and Ji, D. (2017). A neural joint model for entity and relation extraction from biomedical text. *BMC Bioinformatics*, 18(1):1–11.

Li, F., Zhang, Y., Zhang, M., and Ji, D. (2016). Joint models for extracting adverse drug events from biomedical text. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*, pages 2838–2844, New York, USA. IJCAI/AAAI Press.

Li, Q. and Ji, H. (2014). Incremental joint extraction of entity mentions and relations. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 402–412, Baltimore, USA.

Miwa, M. and 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 (Volume 1: Long Papers)*, pages 1105–1116, Berlin, Germany.

References II

Miwa, M. and 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*, pages 1858–1869, Doha, Qatar. Association for Computational Linguistics.

Miyato, T., Dai, A. M., and Goodfellow, I. (2017). Adversarial training methods for semi-supervised text classification. In *Proceedings of the International Conference on Learning Representations*, Toulon, France.

Nadeau, D. and Sekine, S. (2007). A survey of named entity recognition and classification. *Lingvisticae Investigationes*, 30(1):3–26.

Wu, Y., Bamman, D., and Russell, S. (2017). Adversarial training for relation extraction. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1778–1783, Copenhagen, Denmark. Association for Computational Linguistics.

Yasunaga, M., Kasai, J., and Radev, D. (2018). Robust multilingual part-of-speech tagging via adversarial training. In *Proceedings of the 16th Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, New Orleans, USA.

Zelenko, D., Aone, C., and Richardella, A. (2003). Kernel methods for relation extraction. *Journal of Machine Learning Research*, 3:1083–1106.