
Must-read papers on NRE

NRE: Neural Relation Extraction.

Contributed by Tianyu Gao and Xu Han.

We released OpenNRE, an open-source framework for neural relation extraction. This repository provides several relation extraction methods and an easy-to-use training and testing framework.

Reviews

1. Nguyen Bach, Sameer Badaskar. **A review of relation extraction.** [paper]
2. Shantanu Kumar. 2017. **A survey of deep learning methods for relation extraction.** [paper]
3. Sachin Pawar, Girish K. Palshikara, Pushpak Bhattacharyya. 2017. **Relation extraction: a survey.** [paper]

Datasets

You can download most of the following datasets in [json](#) format from OpenNRE.

Sentence-Level Relation Extraction

1. **ACE 2005 Dataset.** [link] [paper]
2. **SemEval-2010 Task 8 Dataset.** [link] [paper]
3. **TACREDD.** [link] [paper]

Distantly Supervised Relation Extraction Datasets

1. **NYT Dataset.** [link] [paper]

Few-shot Relation Extraction Datasets

1. **FewRel.** [link] [1.0 paper] [2.0 paper]

Document-Level Relation Extraction Datasets

1. **DocRED.** [link] [paper]

Papers

Pattern-Based Methods

1. Stephen Soderland, David Fisher, Jonathan Aseltine, and Wendy Lehnert. 1995. **Crystal inducing a conceptual dictionary**. In Proceedings of IJCAI. [paper]
2. Jun-Tae Kim and Dan I. Moldovan. 1995. **Acquisition of linguistic patterns for knowledge-based information extraction**. TKDE. [paper]
3. Scott B Huffman. 1995. **Learning information extraction patterns from examples**. In Proceedings of IJCAI. [paper]
4. Mary Elaine Califf and Raymond J. Mooney. 1997. **Relational learning of pattern-match rules for information extraction**. In Proceedings of CoNLL. [paper]
5. Andrew Carlson, Justin Betteridge, Bryan Kisiel, Burr Settles, Estevam R Hruschka, and Tom M Mitchell. 2010. **Toward an architecture for never-ending language learning**. In Proceedings of AAAI. [paper]
6. Ndapandula Nakashole, Gerhard Weikum, and Fabian Suchanek. 2012. **PATTY: A taxonomy of relational patterns with semantic types**. In Proceedings of EMNLP-CoNLL. [paper]
7. Shun Zheng, Xu Han, Yankai Lin, Peilin Yu, Lu Chen, Ling Huang, Zhiyuan Liu, and Wei Xu. 2019. **DIAG-NRE: A neural pattern diagnosis framework for distantly supervised neural relation extraction**. In Proceedings of ACL. [paper]

Statistical Methods

Feature-Based

1. Nanda Kambhatla. 2004. **Combining lexical, syntactic, and semantic features with maximum entropy models for extracting relations**. [paper]
2. Guodong Zhou, Jian Su, Jie Zhang, and Min Zhang. 2005. **Exploring various knowledge in relation extraction**. In Proceedings of ACL, pages 427–434. [paper]
3. Jing Jiang and ChengXiang Zhai. 2007. **A systematic exploration of the feature space for relation extraction**. In Proceedings of NAACL, pages 113–120. [paper]
4. Dat PT Nguyen, Yutaka Matsuo, and Mitsuru Ishizuka. 2007. **Relation extraction from wikipedia using subtree mining**. In Proceedings of AAAI, pages 1414–1420. [paper]

Kernel-Based

1. Aron Culotta and Jeffrey Sorensen. 2004. **Dependency tree kernels for relation extraction**. In Proceedings of ACL, page 423. [paper]

-
2. Razvan C Bunescu and Raymond J Mooney. 2005. **A shortest path dependency kernel for relation extraction.** In Proceedings of EMNLP, pages 724–731. [paper]
 3. Shubin Zhao and Ralph Grishman. 2005. **Extracting relations with integrated information using kernel methods.** In Proceedings of ACL, pages 419–426. [paper]
 4. Raymond J Mooney and Razvan C Bunescu. 2006. **Subsequence kernels for relation extraction.** In Proceedings of NIPS, pages 171–178. [paper]
 5. Min Zhang, Jie Zhang, Jian Su, and Guodong Zhou. 2006. **A composite kernel to extract relations between entities with both flat and structured features.** In Proceedings of ACL, pages 825–832. [paper]
 6. Mengqiu Wang. 2008. **A re-examination of dependency path kernels for relation extraction.** In Proceedings of IJCNLP, pages 841–846. [paper]

Graphical Models

1. Dan Roth and Wen-tau Yih. 2002. **Probabilistic reasoning for entity & relation recognition.** In Proceedings of COLING. [paper]
2. Sunita Sarawagi and William W Cohen. 2005. **Semimarkov conditional random fields for information extraction.** In Proceedings of NIPS, pages 1185–1192. [paper]
3. Xiaofeng Yu and Wai Lam. 2010. **Jointly identifying entities and extracting relations in encyclopedia text via a graphical model approach.** In Proceedings of ACL, pages 1399–1407. [paper]

Embedding Models

1. Jason Weston, Antoine Bordes, Oksana Yakhnenko, and Nicolas Usunier. 2013. **Connecting language and knowledge bases with embedding models for relation extraction.** In Proceedings of EMNLP, pages 1366–1371. [paper]
2. Sebastian Riedel, Limin Yao, Andrew McCallum, and Benjamin M Marlin. 2013. **Relation extraction with matrix factorization and universal schemas.** In Proceedings of NAACL, pages 74–84. [paper]
3. Matthew R Gormley, Mo Yu, and Mark Dredze. 2015. **Improved relation extraction with feature-rich compositional embedding models.** In Proceedings of EMNLP, pages 1774–1784. [paper]
4. Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. 2013. **Translating embeddings for modeling multirelational data.** In Proceedings of NIPS, pages 2787–2795. [paper]
5. Zhen Wang, Jianwen Zhang, Jianlin Feng, and Zheng Chen. 2014. **Knowledge graph embedding by translating on hyperplanes.** In Proceedings of AAAI. [paper]

-
6. Yankai Lin, Zhiyuan Liu, Maosong Sun, Yang Liu, and Xuan Zhu. 2015. **Learning entity and relation embeddings for knowledge graph completion**. In Proceedings of AAAI. [paper]

Neural Methods

Recursive Neural Networks

1. Richard Socher, Brody Huval, Christopher D Manning, and Andrew Y Ng. 2012. **Semantic compositionality through recursive matrix-vector spaces**. In Proceedings of EMNLP, pages 1201–1211. [paper]
2. Makoto Miwa and Mohit Bansal. 2016. **End-to-end relation extraction using lstms on sequences and tree structures**. In Proceedings of ACL, pages 1105–1116. [paper]

Convolutional Neural Networks

1. Chunyang Liu, Wenbo Sun, Wenhan Chao, and Wanxiang Che. 2013. **Convolution neural network for relation extraction**. In Proceedings of ICDM, pages 231–242. [paper]
2. Daojian Zeng, Kang Liu, Siwei Lai, Guangyou Zhou, and Jun Zhao. 2014. **Relation classification via convolutional deep neural network**. In Proceedings of COLING, pages 2335–2344. [paper]
3. Cicero Nogueira dos Santos, Bing Xiang, and Bowen Zhou. 2015. **Classifying relations by ranking with convolutional neural networks**. In Proceedings of ACL-IJCNLP, pages 626–634. [paper]
4. Thien Huu Nguyen and Ralph Grishman. 2015. **Relation extraction: Perspective from convolutional neural networks**. In Proceedings of the NAACL Workshop on Vector Space Modeling for NLP, pages 39–48. [paper]

Recurrent Neural Networks

1. Dongxu Zhang and Dong Wang. 2015. **Relation classification via recurrent neural network**. arXiv preprint arXiv:1508.01006. [paper]
2. Thien Huu Nguyen and Ralph Grishman. 2015. **Combining neural networks and log-linear models to improve relation extraction**. arXiv preprint arXiv:1511.05926. [paper]
3. Ngoc Thang Vu, Heike Adel, Pankaj Gupta, et al. 2016. **Combining recurrent and convolutional neural networks for relation classification**. In Proceedings of NAACL, pages 534–539. [paper]
4. Shu Zhang, Dequan Zheng, Xinchun Hu, and Ming Yang. 2015. **Bidirectional long short-term memory networks for relation classification**. In Proceedings of PACLIC, pages 73–78. [paper]

Graph Neural Networks

-
1. Yuhao Zhang, Peng Qi, and Christopher D. Manning. 2018. **Graph convolution over pruned dependency trees improves relation extraction**. In Proceedings of EMNLP, pages 2205–2215. [paper]
 2. Hao Zhu, Yankai Lin, Zhiyuan Liu, Jie Fu, Tat-Seng Chua, and Maosong Sun. 2019. **Graph neural networks with generated parameters for relation extraction**. In Proceedings of ACL, pages 1331–1339. [paper]
 3. Zhijing Jin, Yongyi Yang, Xipeng Qiu, Zheng Zhang. 2020. **Relation of the Relations: A New Paradigm of the Relation Extraction Problem**. In Proceedings of EMNLP. [paper]

Attention

1. Peng Zhou, Wei Shi, Jun Tian, Zhenyu Qi, Bingchen Li, Hongwei Hao, and Bo Xu. 2016. **Attention-based bidirectional long short-term memory networks for relation classification**. In Proceedings of ACL, pages 207–212. [paper]
2. Linlin Wang, Zhu Cao, Gerard De Melo, and Zhiyuan Liu. 2016. **Relation classification via multi-level attention cnns**. In Proceedings of ACL, pages 1298–1307. [paper]
3. Minguang Xiao and Cong Liu. 2016. **Semantic relation classification via hierarchical recurrent neural network with attention**. In Proceedings of COLING, pages 1254–1263. [paper]

Word & Position Embedding

1. Joseph Turian, Lev Ratinov, and Yoshua Bengio. 2010. **Word representations: a simple and general method for semi-supervised learning**. In Proceedings of ACL, pages 384–394. [paper]
2. Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. **Distributed representations of words and phrases and their compositionality**. In Proceedings of NIPS, pages 3111–3119. [paper]
3. Daojian Zeng, Kang Liu, Siwei Lai, Guangyou Zhou, and Jun Zhao. 2014. **Relation classification via convolutional deep neural network**. In Proceedings of COLING, pages 2335–2344. [paper]

Shortest Dependency Path

1. Yang Liu, Furu Wei, Sujian Li, Heng Ji, Ming Zhou, and WANG Houfeng. 2015. **A dependency-based neural network for relation classification**. In Proceedings of ACL-IJCNLP, pages 285–290. [paper]
2. Yan Xu, Lili Mou, Ge Li, Yunchuan Chen, Hao Peng, and Zhi Jin. 2015. **Classifying relations via long short term memory networks along shortest dependency paths**. In Proceedings of EMNLP, pages 1785–1794. [paper]

Universal Schema

1. Patrick Verga, David Belanger, Emma Strubell, Benjamin Roth, and Andrew McCallum. 2016. **Multilingual relation extraction using compositional universal schema**. In Proceedings of NAACL, pages 886–896. [paper]
2. Patrick Verga and Andrew McCallum. 2016. **Row-less universal schema**. In Proceedings of ACL, pages 63–68. [paper]
3. Sebastian Riedel, Limin Yao, Andrew McCallum, and Benjamin M Marlin. 2013. **Relation extraction with matrix factorization and universal schemas**. In Proceedings of NAACL, pages 74–84. [paper]

Transformer and BERT

1. Jinhua Du, Jingguang Han, Andy Way, and Dadong Wan. 2018. **Multi-level structured self-attentions for distantly supervised relation extraction**. In Proceedings of EMNLP, pages 2216–2225. [paper]
2. Patrick Verga, Emma Strubell, and Andrew McCallum. 2018. **Simultaneously self-attending to all mentions for full-abstract biological relation extraction**. In Proceedings of NAACL-HLT, pages 872–884. [paper]
3. Shanchan Wu and Yifan He. 2019. **Enriching pre-trained language model with entity information for relation classification**. arXiv preprint arXiv:1905.08284. [paper]
4. Livio Baldini Soares, Nicholas FitzGerald, Jeffrey Ling, and Tom Kwiatkowski. 2019. **Matching the blanks: Distributional similarity for relation learning**. In Proceedings of ACL, pages 2895–2905. [paper]

Distant Supervision

1. Mike Mintz, Steven Bills, Rion Snow, and Dan Jurafsky. 2009. **Distant supervision for relation extraction without labeled data**. In Proceedings of ACL/IJCNLP, pages 1003–1011. [paper]
2. Truc-Vien T Nguyen and Alessandro Moschitti. 2011. **End-to-end relation extraction using distant supervision from external semantic repositories**. In Proceedings of ACL, pages 277–282. [paper]
3. Bonan Min, Ralph Grishman, Li Wan, Chang Wang, and David Gondek. 2013. **Distant supervision for relation extraction with an incomplete knowledge base**. In Proceedings of NAACL, pages 777–782. [paper]

Selecting Informative Instances

-
1. Sebastian Riedel, Limin Yao, and Andrew McCallum. 2010. **Modeling relations and their mentions without labeled text**. In Proceedings of ECML-PKDD, pages 148–163. [paper]
 2. Raphael Hoffmann, Congle Zhang, Xiao Ling, Luke Zettlemoyer, and Daniel S Weld. 2011. **Knowledgebased weak supervision for information extraction of overlapping relations**. In Proceedings of ACL, pages 541–550. [paper]
 3. Mihai Surdeanu, Julie Tibshirani, Ramesh Nallapati, and Christopher D Manning. 2012. **Multi-instance multi-label learning for relation extraction**. In Proceedings of EMNLP, pages 455–465. [paper]
 4. Daojian Zeng, Kang Liu, Yubo Chen, and Jun Zhao. 2015. **Distant supervision for relation extraction via piecewise convolutional neural networks**. In Proceedings of EMNLP, pages 1753–1762. [paper]
 5. Yankai Lin, Shiqi Shen, Zhiyuan Liu, Huanbo Luan, and Maosong Sun. 2016. **Neural relation extraction with selective attention over instances**. In Proceedings of ACL, pages 2124–2133. [paper]
 6. Yuhao Zhang, Victor Zhong, Danqi Chen, Gabor Angeli, and Christopher D Manning. 2017. **Positionaware attention and supervised data improve slot filling**. In Proceedings of EMNLP, pages 35–45. [paper]
 7. Xu Han, Pengfei Yu, Zhiyuan Liu, Maosong Sun, and Peng Li. 2018c. **Hierarchical relation extraction with coarse-to-fine grained attention**. In Proceedings of EMNLP, pages 2236–2245. [paper]
 8. Yang Li, Guodong Long, Tao Shen, Tianyi Zhou, Lina Yao, Huan Huo, and Jing Jiang. 2019. **Self attention enhanced selective gate with entity-aware embedding for distantly supervised relation extraction**. arXiv preprint arXiv:1911.11899. [paper]
 9. Linmei Hu, Luhao Zhang, Chuan Shi, Liqiang Nie, Weili Guan, and Cheng Yang. 2019. **Improving distantly-supervised relation extraction with joint label embedding**. In Proceedings of EMNLP-IJCNLP, pages 3812–3820. [paper]

Incorporating Extra Context

1. Guoliang Ji, Kang Liu, Shizhu He, Jun Zhao, et al. 2017. **Distant supervision for relation extraction with sentence-level attention and entity descriptions**. In AAAI, pages 3060–3066. [paper]
2. Xu Han, Zhiyuan Liu, and Maosong Sun. 2018b. **Neural knowledge acquisition via mutual attention between knowledge graph and text**. In Proceedings of AAAI. [paper]
3. Ningyu Zhang, Shumin Deng, Zhanlin Sun, Guanying Wang, Xi Chen, Wei Zhang, and Huajun Chen. 2019. **Long-tail relation extraction via knowledge graph embeddings and graph convolution networks**. In Proceedings of NAACL-HLT, pages 3016–3025. [paper]
4. Jianfeng Qu, Wen Hua, Dantong Ouyang, Xiaofang Zhou, and Ximing Li. 2019. **A fine-grained**

-
- and noise-aware method for neural relation extraction.** In Proceedings of CIKM, pages 659–668. [paper]
 5. Patrick Verga, David Belanger, Emma Strubell, Benjamin Roth, and Andrew McCallum. 2016. **Multilingual relation extraction using compositional universal schema.** In Proceedings of NAACL, pages 886–896. [paper]
 6. Yankai Lin, Zhiyuan Liu, and Maosong Sun. 2017. **Neural relation extraction with multilingual attention.** In Proceedings of ACL, pages 34–43. [paper]
 7. Xiaozhi Wang, Xu Han, Yankai Lin, Zhiyuan Liu, and Maosong Sun. 2018. **Adversarial multilingual neural relation extraction.** In Proceedings of COLING, pages 1156–1166. [paper]

Sophisticated Mechanisms

1. Ngoc Thang Vu, Heike Adel, Pankaj Gupta, et al. 2016. **Combining recurrent and convolutional neural networks for relation classification.** In Proceedings of NAACL, pages 534–539. [paper]
2. Iz Beltagy, Kyle Lo, and Waleed Ammar. 2019. **Combining distant and direct supervision for neural relation extraction.** In Proceedings of NAACL-HLT, pages 1858–1867. [paper]
3. Tianyu Liu, Kexiang Wang, Baobao Chang, and Zhifang Sui. 2017. **A soft-label method for noisetolerant distantly supervised relation extraction.** In Proceedings of EMNLP, pages 1790–1795. [paper]
4. Jun Feng, Minlie Huang, Li Zhao, Yang Yang, and Xiaoyan Zhu. 2018. **Reinforcement learning for relation classification from noisy data.** In Proceedings of AAAI, pages 5779–5786. [paper]
5. Xiangrong Zeng, Shizhu He, Kang Liu, and Jun Zhao. 2018. **Large scaled relation extraction with reinforcement learning.** In Proceedings of AAAI, pages 5658–5665. [paper]
6. Yi Wu, David Bamman, and Stuart Russell. 2017. **Adversarial training for relation extraction.** In Proceeding of EMNLP, pages 1778–1783. [paper]
7. Xu Han, Zhiyuan Liu, and Maosong Sun. 2018. **Denoising distant supervision for relation extraction via instance-level adversarial training.** arXiv preprint arXiv:1805.10959. [paper]

Few-Shot Learning

1. Xu Han, Hao Zhu, Pengfei Yu, ZiyunWang, Yuan Yao, Zhiyuan Liu, and Maosong Sun. 2018d. **Fewrel: A large-scale supervised few-shot relation classification dataset with state-of-the-art evaluation.** In Proceedings of EMNLP, pages 4803–4809. [paper]
2. Tianyu Gao, Xu Han, Hao Zhu, Zhiyuan Liu, Peng Li, Maosong Sun, and Jie Zhou. 2019. **FewRel 2.0: Towards more challenging few-shot relation classification.** In Proceedings of EMNLP-IJCNLP, pages 6251–6256. [paper]
3. Tianyu Gao, Xu Han, Zhiyuan Liu, Maosong Sun. 2019. **Hybrid Attention-Based Prototypical Networks for Noisy Few-Shot Relation Classification.** In Proceedings of AAAI. [paper]

-
4. Livio Baldini Soares, Nicholas FitzGerald, Jeffrey Ling, and Tom Kwiatkowski. 2019. **Matching the blanks: Distributional similarity for relation learning**. In Proceedings of ACL, pages 2895–2905. [paper]
 5. Zhi-Xiu Ye and Zhen-Hua Ling. 2019. **Multi-level matching and aggregation network for few-shot relation classification**. In Proceedings of ACL, pages 2872–2881. [paper]

Document-Level Relation Extraction

1. Yuan Yao, Deming Ye, Peng Li, Xu Han, Yankai Lin, Zhenghao Liu, Zhiyuan Liu, Lixin Huang, Jie Zhou, and Maosong Sun. 2019. **DocRED: A large-scale document-level relation extraction dataset**. In Proceedings of ACL, pages 764–777. [paper]
2. Michael Wick, Aron Culotta, et al. 2006. **Learning field compatibilities to extract database records from unstructured text**. In Proceedings of EMNLP. [paper]
3. Matthew Gerber and Joyce Chai. 2010. **Beyond Nom-Bank: A study of implicit arguments for nominal predicates**. In Proceedings of ACL, pages 1583–1592. [paper]
4. Kumutha Swampillai and Mark Stevenson. 2011. **Extracting relations within and across sentences**. In Proceedings of RANLP. [paper]
5. Katsumasa Yoshikawa, Sebastian Riedel, et al. 2011. **Coreference based event-argument relation extraction on biomedical text**. J. Biomed. Semant. [paper]
6. Chris Quirk and Hoifung Poon. 2017. **Distant supervision for relation extraction beyond the sentence boundary**. In Proceedings of EACL, pages 1171–1182. [paper]
7. Wenyuan Zeng, Yankai Lin, Zhiyuan Liu, and Maosong Sun. 2017. **Incorporating relation paths in neural relation extraction**. In Proceedings of EMNLP, pages 1768–1777. [paper]
8. Fenia Christopoulou, Makoto Miwa, and Sophia Ananiadou. 2018. **A walk-based model on entity graphs for relation extraction**. In Proceedings of ACL, pages 81–88. [paper]
9. Nanyun Peng, Hoifung Poon, Chris Quirk, Kristina Toutanova, and Wen-tau Yih. 2017. **Cross-sentence n-ary relation extraction with graph LSTMs**. TACL, 5:101–115. [paper]
10. Linfeng Song, Yue Zhang, et al. 2018. **N-ary relation extraction using graph-state lstm**. In Proceedings of EMNLP. [paper]
11. Hao Zhu, Yankai Lin, Zhiyuan Liu, Jie Fu, Tat-Seng Chua, and Maosong Sun. 2019. **Graph neural networks with generated parameters for relation extraction**. In Proceedings of ACL, pages 1331–1339. [paper]

Open Information Extraction

1. Michele Banko, Michael J Cafarella, Stephen Soderland, Matthew Broadhead, and Oren Etzioni. 2007. **Open information extraction from the web**. In Proceedings of IJCAI, pages 2670–2676.

-
- [paper]
2. Anthony Fader, Stephen Soderland, and Oren Etzioni. 2011. **Identifying relations for open information extraction**. In Proceedings of EMNLP, pages 1535–1545. [paper]
 3. Mausam, Michael Schmitz, Stephen Soderland, Robert Bart, and Oren Etzioni. 2012. **Open language learning for information extraction**. In Proceedings of EMNLP-CoNLL, pages 523–534. [paper]
 4. Luciano Del Corro and Rainer Gemulla. 2013. **Clausie: clause-based open information extraction**. In Proceedings of WWW, pages 355–366. [paper]
 5. Gabor Angeli, Melvin Jose Johnson Premkumar, and Christopher D. Manning. 2015. **Leveraging linguistic structure for open domain information extraction**. In Proceedings of ACL-IJCNLP, pages 344–354. [paper]
 6. Gabriel Stanovsky and Ido Dagan. 2016. **Creating a large benchmark for open information extraction**. In Proceedings of EMNLP, pages 2300–2305. [paper]
 7. Mausam Mausam. 2016. **Open information extraction systems and downstream applications**. In Proceedings of IJCAI, pages 4074–4077. [paper]
 1. Lei Cui, Furu Wei, and Ming Zhou. 2018. **Neural open information extraction**. In Proceedings of ACL, pages 407–413. [paper]

Relation Discovery

1. Limin Yao, Aria Haghighi, Sebastian Riedel, and Andrew McCallum. 2011. **Structured relation discovery using generative models**. In Proceedings of EMNLP, pages 1456–1466. [paper]
2. Diego Marcheggiani and Ivan Titov. 2016. **Discretestate variational autoencoders for joint discovery and factorization of relations**. TACL, 4:231–244. [paper]
3. Yusuke Shinyama and Satoshi Sekine. 2006. **Preemptive information extraction using unrestricted relation discovery**. In Proceedings of NAACL, pages 304–311. [paper]
4. Hady Elsahar, Elena Demidova, Simon Gottschalk, Christophe Gravier, and Frederique Laforest. 2017. **Unsupervised open relation extraction**. In Proceedings of ESWC, pages 12–16. [paper]
5. Ruidong Wu, Yuan Yao, Xu Han, Ruobing Xie, Zhiyuan Liu, Fen Lin, Leyu Lin, and Maosong Sun. 2019. **Open relation extraction: Relational knowledge transfer from supervised data to unsupervised data**. In Proceedings of EMNLP-IJCNLP, pages 219–228. [paper]