

Event-Centric Natural Language Processing

Muhao Chen¹, Hongming Zhang^{2,3}, Qiang Ning⁴,
Manling Li⁵, Heng Ji^{4,5}, Kathleen McKeown⁶, Dan Roth²

¹Department of Computer Science, USC

²Department of Computer and Information Science, UPenn

³Department of Computer Science and Engineering, HKUST

⁴Amazon Alexa AI

⁵Department of Computer Science, UIUC

⁶Department of Computer Science, Columbia University

Abstract

This tutorial targets researchers and practitioners who are interested in AI technologies that help machines understand natural language text, particularly real-world events described in the text. These include methods to extract the internal structures of an event regarding its protagonist(s), participant(s) and properties, as well as external structures concerning memberships, temporal and causal relations of multiple events. This tutorial will provide audience with a systematic introduction of (i) knowledge representations and acquisition of events, (ii) various methods for automated extraction, conceptualization, coreference resolution and prediction of events and their relations, (iii) induction of event processes and properties, and (iv) a wide range of NLP and commonsense understanding tasks that benefit from aforementioned techniques. We will conclude the tutorial by outlining emerging research problems in this area.

1 Introduction

Human languages always involve the description of real-world events. Therefore, understanding events plays a critical role in NLP. For example, narrative prediction benefits from learning the causal relations of events to predict what happens next in a story (Chaturvedi et al., 2017a); machine comprehension of documents may involve understanding of events that affect the stock market (Ding et al., 2015), describe natural phenomena (Berant et al., 2014) or identify disease phenotypes (Zhang et al., 2020d). In fact, event understanding also widely finds its important use cases in tasks such as open-domain question answering (Yang et al., 2003), intent prediction (Rashkin et al., 2018), timeline construction (Do et al., 2012), text summarization (Daumé III and Marcu, 2006) and misinformation detection (Fung et al., 2021). Since events are not just simple, standalone predicates, frontier research

on event understanding generally faces two key challenges. One challenge is to precisely induce the relations of events, which describe memberships, co-reference, temporal orders and causality of events. The other is to comprehend the inherent structure and properties of an event, concerning its participants, granularity, location and time.

In this tutorial, we will comprehensively review existing paradigms for event-centric knowledge representation in literature, and focus on their contributions to NLP tasks. Beyond introducing partial-label and unsupervised learning approaches for event extraction, we will discuss recent constrained learning and structured inference approaches for multi-faceted event-event relation extraction from text. We will also review recent data-driven methods for event prediction tasks, including event process induction and conceptualization, and how an event-centric language model benefits narrative prediction. In addition, we will illustrate how distantly supervised approaches help resolve temporal and causal commonsense understanding of events, and how they can be applied to construct a large-scale eventuality knowledge base. Participants will learn about recent trends and emerging challenges in this topic, representative tools and learning resources to obtain ready-to-use models, and how related models and techniques benefit end-use NLP applications.

2 Outline of Tutorial Content

This **half-day** tutorial presents a systematic overview of recent advances in event-centric NLP technologies. We will begin with motivating this topic with several real-world applications, and introduce the main research problems. Then, we will introduce methods for automated extraction of events as well as their participants, properties and relations from natural language text. Based on the

extracted eventuality knowledge, we will explain how various prediction tasks, including the completion of an event complex, conceptualization and consolidation of event processes, can be resolved. We will also discuss commonsense understanding of events, with a focus on the temporal and cognitive aspects. Moreover, we will exemplify the use of aforementioned technologies in NLP applications of various domains, and will outline emerging research challenges that may catalyze further investigation on this topic. The detailed contents are outlined below.

2.1 Motivation [20min]

We will define the main research problem and motivate the topic by presenting several real-world applications based on event-centric NLP. This seeks to provide 20 minutes of presented content to motivate the main topic of this tutorial.

2.2 Background of Events and Their Representations [30min]

We will start the tutorial by introducing the essential background knowledge about events and their relations, including the definitions, categorizations, and applications (P. D. Mourelatos, 1978; Bach, 1986). In the last part of the introduction, we will talk about widely used event representation methods, including event schemas (Baker et al., 1998; Li et al., 2020b, 2021a), event knowledge graphs (Zhang et al., 2020c), event processes (Chambers and Jurafsky, 2008), event language models (Peng et al., 2017), and more recent work on event meaning representation via question-answer pairs (He et al., 2015; Michael et al., 2018), event network embeddings (Zeng et al., 2021) and event time expression embeddings (Goyal and Durrett, 2019). This part is estimated to take 30 minutes.

2.3 Event-centric Information Extraction [40min]

We will introduce unsupervised and zero-shot techniques for parsing the internal structures of verb and nominal events from natural language text, which also involves methods for automatic salient event detection (Liu et al., 2018), joint entity, relation and event extraction (Lin et al., 2020), and graph neural networks based encoding and decoding for information extraction (Zhang and Ji, 2021). Then we will discuss the recent research trend to extend information extraction from sentence-level

to document-level (Du and Cardie, 2020; Li et al., 2021b). Besides, we will also discuss methods that identify temporal and causal relations of primitive events (Ning et al., 2018), and membership relations of multi-granular events (Aldawsari and Finlayson, 2019). Specifically, for data-driven extraction methods, we will present how constrained learning (Li et al., 2019) and structured prediction are incorporated to improve the tasks by enforcing logic consistency among different categories of event-event relations (Wang et al., 2020). We will also cover various cross-domain (Huang et al., 2018), cross-lingual (Subburathinam et al., 2019) and cross-media (Li et al., 2020a) structure transfer approaches for event extraction. This part is estimated to be 40 minutes.

2.4 Understanding Event Processes [35min]

We will then present recent works on machine comprehension and prediction on event processes/sequences. Specifically, people are trying to understand the progress of events from different angles. For example, many efforts have been devoted into modeling event narratives (Peng et al., 2017; Chaturvedi et al., 2017b; Lee and Goldwasser, 2019) such that they can successfully predict missing events in an event process. Besides, another important event understanding angle is conceptualization (Zhang et al., 2020a), which aims at understanding the super-sub relations between a coarse-grained event and a fine-grained event process (Glavaš et al., 2014). In this context, the machine could also be expected to generate the event process given a goal (Zhang et al., 2020a), infer the goal given the process (Chen et al., 2020), and capture the recurrence of events in a process (Zhu et al., 2021). Last but not least, event coreference, which links references to the same event together, also plays a critical role in understanding events (Cybulska and Vossen, 2014). This part should last for 35 minutes.

2.5 Event-centric Commonsense Knowledge Acquisition [35min]

Commonsense reasoning is a challenging yet important research problem in the AI community and one key challenge we are facing is the lack of satisfactory commonsense knowledge resources about events. Previous resources (Liu and Singh, 2004) typically require laborious and expensive human annotations, which are not feasible on a large scale. In this tutorial, we introduce recent

progress on modeling commonsense knowledge with high-order selectional preference over event knowledge and demonstrates that how to convert relatively cheap event knowledge, which can be easily acquired from raw documents with linguistic patterns, to precious commonsense knowledge defined in ConceptNet (Zhang et al., 2020b). Beyond that, we will also introduce how to automatically acquire other event-centric commonsense knowledge including but not limited to temporal properties (Zhou et al., 2020), intentions (Chen et al., 2020), effects (Sap et al., 2019) and graph schemas (Li et al., 2020c) of events. This part is estimated to be 35 minutes.

2.6 Event Summarization [30min]

In addition to specific, individual events, we are also interested in large-scale events that unfold over time. Over the past year, we saw many examples of such events, including COVID-19, the vaccine roll-out, the Black Lives Matter movement and the US presidential election. In this tutorial, we will present methods for tracking such events over time and generating summaries that provide updates as an event unfolds. The task of identifying and tracking events was first introduced in the Topic Detection and Tracking challenge (Allan et al., 1998). Recent work has explored new methods for tracking and visualizing such events over time (e.g., (Laban and Hearst, 2017; Miranda et al., 2018; Staykovski et al., 2019; Saravanakumar et al., 2021)), in some cases generating summaries that contain information on what is new (e.g., (Kedzie et al., 2015, 2018)) and in other cases, exploring timeline summarization, ordering events and generating summaries that are placed along a timeline (e.g., (Wang et al., 2015; Binh Tran et al., 2013; Chen et al., 2019; Nguyen et al., 2014)) We will also consider how these are related to summarization of an event that takes place within a single day, a problem that falls within the category of multi-document summarization (e.g., (Liu and Lapata, 2019; Fabbri et al., 2019)), as typically there may be many articles covering the same event. By using multiple articles as input, a summarizer can present different perspectives on the same event as well as identify salient information that is highlighted many in different ways across the set of input articles. This part is scheduled for 30 minutes

2.7 Emerging Research Problems [20min]

Event-centric NLP impacts on a wide spectrum of knowledge-driven AI tasks, and is particularly knotted with commonsense understanding. We will conclude the tutorial using 20 minutes by presenting some challenges and potential research topics in applying eventuality knowledge in downstream tasks (e.g., reading comprehension, dialogue generation, and event timeline generation), and grounding eventuality knowledge to visual modalities, and challenges for cross-document event consolidation with human-defined schemas.

3 Specification of the Tutorial

The proposed tutorial is considered a **cutting-edge** tutorial that introduces the recent advances in an emerging area of NLP. The presented topic has not been covered by previous ACL/EMNLP/NAACL/EACL/COLING tutorials in the past 4 years. This tutorial has not been presented elsewhere, while a more AI-flavored version with a subset of the contents has been planned in parallel at AAI 2021, to be presented by a subset of the instructors. We estimate that at least 60% of the works covered in this tutorial are from researchers other than the instructors.

Audience and Prerequisites While no specific background knowledge is assumed of the audience, it would be best for the attendees to know about basic deep learning technologies, pre-trained word embeddings (e.g. Word2Vec) and language models (e.g. BERT). The following is a reading list that could help provide background knowledge to the audience before attending this tutorial:

- Emmon Bach. The algebra of events. *Linguistics and philosophy*. 9(1):5-16, 1986.
- Nathanael Chambers. Event Schema Induction with a Probabilistic Entity-Driven Model. *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2013
- Tao Li, Vivek Gupta, Maitrey Mehta, and Vivek Srikumar. A Logic-Driven Framework for Consistency of Neural Models. *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 2019.

- Ying Lin, Heng Ji, Fei Huang and Lingfei Wu. A Joint Neural Model for Information Extraction with Global Features. Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL), 2020.
- Haoyu Wang, Muhao Chen, Hongming Zhang, and Dan Roth. Joint Constrained Learning for Event-Event Relation Extraction. Proceedings of the 2020 Empirical Methods in Natural Language Processing (EMNLP), 2020.
- Nathanael Chambers and Dan Jurafsky. Un-supervised learning of narrative event chains. In Proceedings of the 46th Annual Meeting of the Association for Computational Linguistics (ACL), 2008.
- Hongming Zhang, Muhao Chen, Haoyu Wang, Yangqiu Song, and Dan Roth. Open-domain Process Structure Induction. Proceedings of the 2020 Empirical Methods in Natural Language Processing (EMNLP), 2020.;
- Maarten Sap, Ronan Le Bras, Emily Allaway, Chandra Bhagavatula, Nicholas Lourie, Hannah Rashkin, Brendan Roof, Noah A. Smith, and Yejin Choi. Atomic: An atlas of machine commonsense for if-then reasoning. In Proceedings of the AAAI Conference on Artificial Intelligence (AAAI). 2019.
- Ben Zhou, Qiang Ning, Daniel Khashabi, and Dan Roth. Temporal Common Sense Acquisition with Minimal Supervision. Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL), 2020.

Open Access All the teaching materials are openly available at <https://cogcomp.seas.upenn.edu/page/tutorial.202108/>.

4 Tutorial Instructors

Muhao Chen is an Assistant Research Professor at the Department of Computer Science, USC. Prior to USC, he was a postdoctoral fellow at UPenn. He received his Ph.D. from the Department of Computer Science at UCLA in 2019, and B.S. in Computer Science from Fudan University in 2014. His research focuses on data-driven machine learning approaches for processing structured data, and knowledge acquisition from unstructured data. Particularly, he is interested in developing

knowledge-aware learning systems with generalizability and requiring minimal supervision, and with concrete applications to natural language understanding, knowledge base construction, computational biology and medicine. Muhao has published over 40 papers in leading AI, NLP and Comp Bio venues. His work has received a best student paper award at ACM BCB, and a best paper award nomination at CoNLL. Additional information is available at <http://muhaochen.github.io>.

Hongming Zhang is currently a third-year Ph.D. student at HKUST and a visiting scholar at UPenn. Hongming has received Hong Kong Ph.D. Fellowship and Microsoft Research Asia Fellowship to support his research on commonsense reasoning and open domain event understanding. He has published more than ten papers on related topics in top-tier conferences. Additional information is available at <http://www.cse.ust.hk/~hzhangal/>.

Qiang Ning is currently an applied scientist at Alexa AI. Qiang was a research scientist on the AllenNLP team at the Allen Institute for AI from 2019-2020. Qiang received his Ph.D. in Dec. 2019 from the Department of Electrical and Computer Engineering at the University of Illinois at Urbana-Champaign (UIUC). He obtained his master’s degree in biomedical imaging from the same department in May 2016. Before coming to the United States, Qiang obtained two bachelor’s degrees from Tsinghua University in 2013, in Electronic Engineering and in Economics, respectively. He was an “Excellent Teacher Ranked by Their Students” across the university in 2017 (UIUC), a recipient of the YEE Fellowship in 2015 (College of Engineering at UIUC), a finalist for the best paper in IEEE ISBI’ 15, and also won the National Scholarship at Tsinghua University in 2012. Additional information is available at <https://qiangning.info/>.

Manling Li is a third-year Ph.D. student at the Computer Science Department of the University of Illinois at Urbana-Champaign (UIUC). Manling has won the Best Demo Paper Award at ACL’20, the Best Demo Paper Award at NAACL’21, C.L. Dave and Jane W.S. Liu Award, and has been selected as Mavis Future Faculty Fellow. She has more than 20 publications on knowledge extraction and reasoning from multimedia data. Additional information is available at <https://limanling.github.io>.

Heng Ji is a Professor at Computer Science Department, and an affiliated faculty member at Electrical

and Computer Engineering Department of University of Illinois at Urbana-Champaign. She is also an Amazon Scholar. She received her B.A. and M. A. in Computational Linguistics from Tsinghua University, and her M.S. and Ph.D. in Computer Science from New York University. Her research interests focus on Natural Language Processing, especially on Multimedia Multilingual Information Extraction, Knowledge Base Population and Knowledge-driven Generation. She was selected as “Young Scientist” and a member of the Global Future Council on the Future of Computing by the World Economic Forum in 2016 and 2017. The awards she received include “AI’s 10 to Watch” Award by IEEE Intelligent Systems in 2013, NSF CAREER award in 2009, Google Research Award in 2009 and 2014, IBM Watson Faculty Award in 2012 and 2014 and Bosch Research Award in 2014-2018. She was invited by the Secretary of the U.S. Air Force and AFRL to join Air Force Data Analytics Expert Panel to inform the Air Force Strategy 2030. She is the lead of many multi-institution projects and tasks, including the U.S. ARL projects on information fusion and knowledge networks construction, DARPA DEFT Tinker Bell team and DARPA KAIROS RESIN team. She has coordinated the NIST TAC Knowledge Base Population task since 2010. She has served as the Program Committee Co-Chair of many conferences including NAACL-HLT2018. She is elected as the North American Chapter of the Association for Computational Linguistics (NAACL) secretary 2020-2021. Her research has been widely supported by the U.S. government agencies (DARPA, ARL, IARPA, NSF, AFRL, DHS) and industry (Amazon, Google, Bosch, IBM, Disney). Additional information is available at <https://blender.cs.illinois.edu/hengji.html>.

Kathleen R. McKeown is the Henry and Gertrude Rothschild Professor of Computer Science at Columbia University and the Founding Director of the Data Science Institute, serving as Director from 2012 to 2017. She is also an Amazon Scholar. In earlier years, she served as Department Chair (1998-2003) and as Vice Dean for Research for the School of Engineering and Applied Science (2010-2012). A leading scholar and researcher in the field of natural language processing, McKeown focuses her research on the use of data for societal problems; her interests include text summarization, question answering, natural language generation,

social media analysis and multilingual applications. She has received numerous honors and awards, including American Academy of Arts and Science elected member, American Association of Artificial Intelligence Fellow, a Founding Fellow of the Association for Computational Linguistics and an Association for Computing Machinery Fellow. Early on she received the National Science Foundation Presidential Young Investigator Award, and a National Science Foundation Faculty Award for Women. In 2010, she won both the Columbia Great Teacher Award—an honor bestowed by the students—and the Anita Borg Woman of Vision Award for Innovation. Additional information is available at <http://www.cs.columbia.edu/~kathy/>.

Dan Roth is the Eduardo D. Glandt Distinguished Professor at the Department of Computer and Information Science, University of Pennsylvania, and a Fellow of the AAAS, ACM, AAI, and the ACL. In 2017 Roth was awarded the John McCarthy Award, the highest award the AI community gives to mid-career AI researchers. Roth was recognized “for major conceptual and theoretical advances in the modeling of natural language understanding, machine learning, and reasoning.” Roth has published broadly in machine learning, natural language processing, knowledge representation and reasoning, and learning theory, and has developed advanced machine learning based tools for natural language applications that are being used widely. Roth has given tutorials on these and other topics in all ACL and AAI major conferences. Until February 2017 Roth was the Editor-in-Chief of the Journal of Artificial Intelligence Research (JAIR). He was the program chair of AAI’11, ACL’03 and CoNLL’02, and serves regularly as an area chair and senior program committee member in the major conferences in his research areas. Prof. Roth received his B.A Summa cum laude in Mathematics from the Technion, Israel, and his Ph.D. in Computer Science from Harvard University in 1995. Additional information is available at <http://www.cis.upenn.edu/~danroth/>.

The presenters of this tutorial have given the following tutorials at leading international conferences and venues in the past:

- Muhao Chen:
 - AAI’21: Event-Centric Natural Language Understanding.
 - AAI’20: Recent Advances of Transferable

Representation Learning.

- Heng Ji:
 - AACL’21: Event-Centric Natural Language Understanding.
 - Multi-lingual Entity Discovery and Linking. Tutorial at the 17th China National Conference on Computational Linguistics (CCL2018) and The 6th International Symposium on Natural Language Processing based on Naturally Annotated Big Data (NLP-NABD2018).
 - ACL’18: Multi-lingual Entity Discovery and Linking.
 - Information Extraction and Knowledge Base Population, Invited course for the 10th Russian Summer School in Information Retrieval, 2016.
 - SIGMOD’16: Automatic Entity Recognition and Typing in Massive Text Data.
 - ACL’15: Successful Data Mining Methods for NLP.
 - ACL’14: Wikification and Beyond: The Challenges of Entity and Concept Grounding.
 - Wikification and Beyond: The Challenges of Entity and Concept Grounding, Advanced Disciplines Lecture at NLPCC’14.
 - COLING’12: Temporal Information Extraction and Shallow Temporal Reasoning.
- Kathleen R. McKeown:
 - COLING’86: Natural Language Generation and User Modelling.
 - ACL’86: Natural Language Generation.
- Dan Roth:
 - Data Science Summer Institute (DSSI) 2007, 2008, 2010, 2011, 2012. A tutorial on Machine Learning in Natural Language Processing.
 - AACL’21: Event-Centric Natural Language Understanding.
 - ACL’20: Commonsense Reasoning for Natural Language Processing.
 - AACL’20: Recent Advances of Transferable Representation Learning.
 - ACL’18: The Conference of the Association on Computational Linguistics. A tutorial on Multi-lingual Entity Discovery and Linking.
 - EACL’17: The European Conference of the Association of Computational Linguistics; A tutorial on Integer Linear Programming Formulations in Natural Language Processing.
- AACL’16: The Conference of the Association for the Advancement of Artificial Intelligence; A tutorial on Structured Prediction.
- ACL’14: The International Conference of the Association on Computational Linguistics. A tutorial on Wikification and Entity Linking.
- AACL’13: The AACL Conference on Artificial Intelligence. Information Trustworthiness.
- COLING’12: The International Conference on Computational Linguistics. A Tutorial on Temporal Information Extraction and Shallow Temporal Reasoning.
- NAACL’12: The North American Conference of the Association on Computational Linguistics. A Tutorial on Constrained Conditional Models: Structured Predictions in NLP.
- NAACL’10: The North American Conference of the Association on Computational Linguistics. A Tutorial on Integer Linear Programming Methods in NLP.
- EACL’09: The European Conference of the Association on Computational Linguistics. A Tutorial on Constrained Conditional Models.
- ACL’07: The International Conference of the Association on Computational Linguistics. A Tutorial on Textual Entailment.

Acknowledgement

This material is based upon work supported in part by U.S. DARPA KAIROS Program Nos. FA8750-19-2-1004, U.S. DARPA AIDA Program No. FA8750-18-2-0014, Air Force Nos. FA8650-17-C-7715 and FA8750-20-2-10002. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of DARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein.

References

- Mohammed Aldawsari and Mark Finlayson. 2019. [Detecting subevents using discourse and narrative features](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4780–4790, Florence, Italy. Association for Computational Linguistics.
- James Allan, Jaime G Carbonell, George Doddington, Jonathan Yamron, and Yiming Yang. 1998. Topic detection and tracking pilot study final report.

- Emmon Bach. 1986. The algebra of events. *Linguistics and philosophy*, 9(1):5–16.
- Collin F. Baker, Charles J. Fillmore, and John B. Lowe. 1998. The berkeley framenet project. In *COLING-ACL*, pages 86–90.
- Jonathan Berant, Vivek Srikumar, Pei-Chun Chen, Abby Vander Linden, Brittany Harding, Brad Huang, Peter Clark, and Christopher D. Manning. 2014. [Modeling biological processes for reading comprehension](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1499–1510, Doha, Qatar. Association for Computational Linguistics.
- Giang Binh Tran, Mohammad Alrifai, and Dat Quoc Nguyen. 2013. [Predicting relevant news events for timeline summaries](#). In *Proceedings of the 22nd International Conference on World Wide Web*, pages 91–92.
- Nathanael Chambers and Dan Jurafsky. 2008. [Unsupervised learning of narrative event chains](#). In *Proceedings of ACL-08: HLT*, pages 789–797, Columbus, Ohio. Association for Computational Linguistics.
- Snigdha Chaturvedi, Haoruo Peng, and Dan Roth. 2017a. [Story comprehension for predicting what happens next](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1603–1614, Copenhagen, Denmark. Association for Computational Linguistics.
- Snigdha Chaturvedi, Haoruo Peng, and Dan Roth. 2017b. [Story Comprehension for Predicting What Happens Next](#). In *In proceedings of the Conference on Empirical Methods in Natural Language Processing*.
- Muhao Chen, Hongming Zhang, Haoyu Wang, and Dan Roth. 2020. “what are you trying to do?” semantic typing of event processes. In *Proceedings of the 24th Conference on Computational Natural Language Learning (CoNLL 2020)*. Association for Computational Linguistics.
- Xiuying Chen, Zhangming Chan, Shen Gao, Meng-Hsuan Yu, Dongyan Zhao, and Rui Yan. 2019. [Learning towards abstractive timeline summarization](#). In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019*, pages 4939–4945. ijcai.org.
- Agata Cybulska and Piek Vossen. 2014. Guidelines for ecb+ annotation of events and their coreference. In *Technical Report*. Technical Report NWR-2014-1, VU University Amsterdam.
- Hal Daumé III and Daniel Marcu. 2006. [Bayesian query-focused summarization](#). In *Proceedings of the 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics*, pages 305–312, Sydney, Australia. Association for Computational Linguistics.
- Xiao Ding, Yue Zhang, Ting Liu, and Junwen Duan. 2015. [Deep learning for event-driven stock prediction](#). In *Twenty-fourth international joint conference on artificial intelligence*.
- Quang Do, Wei Lu, and Dan Roth. 2012. [Joint inference for event timeline construction](#). In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 677–687, Jeju Island, Korea. Association for Computational Linguistics.
- Xinya Du and Claire Cardie. 2020. Event extraction by answering (almost) natural questions. In *Proc. The 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP2020)*.
- Alexander Fabbri, Irene Li, Tianwei She, Suyi Li, and Dragomir Radev. 2019. [Multi-news: A large-scale multi-document summarization dataset and abstractive hierarchical model](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1074–1084, Florence, Italy. Association for Computational Linguistics.
- Yi Fung, Christopher Thomas, Revanth Gangi Reddy, Sandeep Polisetty, Heng Ji, Shih-Fu Chang, Kathleen McKeown, Mohit Bansal, and Avi Sil. 2021. [Infosurgeon: Cross-media fine-grained information consistency checking for fake news detection](#). In *Proc. The Joint Conference of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (ACL-IJCNLP 2021)*.
- Goran Glavaš, Jan Šnajder, Marie-Francine Moens, and Parisa Kordjamshidi. 2014. [HiEve: A corpus for extracting event hierarchies from news stories](#). In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC-2014)*, pages 3678–3683, Reykjavik, Iceland. European Languages Resources Association (ELRA).
- Tanya Goyal and Greg Durrett. 2019. Embedding time expressions for deep temporal ordering models. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics 2019 (ACL2019)*.
- Luheng He, Mike Lewis, and Luke Zettlemoyer. 2015. Question-answer driven semantic role labeling: Using natural language to annotate natural language. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 643–653.
- Lifu Huang, Heng Ji, Kyunghyun Cho, Ido Dagan, Sebastian Riedel, and Clare Voss. 2018. Zero-shot transfer learning for event extraction. In *Proc. The 56th Annual Meeting of the Association for Computational Linguistics (ACL2018)*.
- Chris Kedzie, Kathleen McKeown, and Hal Daumé III. 2018. [Content selection in deep learning models of](#)

- summarization. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1818–1828, Brussels, Belgium. Association for Computational Linguistics.
- Chris Kedzie, Kathleen McKeown, and Fernando Diaz. 2015. [Predicting salient updates for disaster summarization](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1608–1617, Beijing, China. Association for Computational Linguistics.
- Philippe Laban and Marti Hearst. 2017. [newsLens: building and visualizing long-ranging news stories](#). In *Proceedings of the Events and Stories in the News Workshop*, pages 1–9, Vancouver, Canada. Association for Computational Linguistics.
- I-Ta Lee and Dan Goldwasser. 2019. [Multi-relational script learning for discourse relations](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4214–4226, Florence, Italy. Association for Computational Linguistics.
- Manling Li, Sha Li, Zhenhailong Wang, Lifu Huang, Kyunghyun Cho, Heng Ji, and Jiawei Han. 2021a. Future is not one-dimensional: Complex event schema induction via graph modeling. In *arxiv2104.06344*.
- Manling Li, Alireza Zareian, Qi Zeng, Spencer Whitehead, Di Lu, Heng Ji, and Shih-Fu Chang. 2020a. Cross-media structured common space for multimedia event extraction. In *Proc. The 58th Annual Meeting of the Association for Computational Linguistics (ACL2020)*.
- Manling Li, Qi Zeng, Ying Lin, Kyunghyun Cho, Heng Ji, Jonathan May, Nathanael Chambers, and Clare Voss. 2020b. Connecting the dots: Event graph schema induction with path language modeling. In *Proc. The 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP2020)*.
- Manling Li, Qi Zeng, Ying Lin, Kyunghyun Cho, Heng Ji, Jonathan May, Nathanael Chambers, and Clare Voss. 2020c. Connecting the dots: Event graph schema induction with path language modeling. In *Proc. The 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP2020)*.
- Sha Li, Heng Ji, and Jiawei Han. 2021b. Document-level event argument extraction by conditional generation. In *Proc. The 2021 Conference of the North American Chapter of the Association for Computational Linguistics - Human Language Technologies (NAACL-HLT2021)*.
- Tao Li, Vivek Gupta, Maitrey Mehta, and Vivek Srikumar. 2019. [A logic-driven framework for consistency of neural models](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3924–3935, Hong Kong, China. Association for Computational Linguistics.
- Ying Lin, Heng Ji, Fei Huang, and Lingfei Wu. 2020. A joint neural model for information extraction with global features. In *Proc. The 58th Annual Meeting of the Association for Computational Linguistics (ACL2020)*.
- Hugo Liu and Push Singh. 2004. Conceptnet: a practical commonsense reasoning tool-kit. *BT technology journal*, 22(4):211–226.
- Jun Liu, Fei Cheng, Yiran Wang, Hiroyuki Shindo, and Yuji Matsumoto. 2018. [Automatic error correction on Japanese functional expressions using character-based neural machine translation](#). In *Proceedings of the 32nd Pacific Asia Conference on Language, Information and Computation*, Hong Kong. Association for Computational Linguistics.
- Yang Liu and Mirella Lapata. 2019. [Hierarchical transformers for multi-document summarization](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5070–5081, Florence, Italy. Association for Computational Linguistics.
- Julian Michael, Gabriel Stanovsky, Luheng He, Ido Dagan, and Luke Zettlemoyer. 2018. Crowdsourcing question-answer meaning representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 560–568.
- Sebastião Miranda, Artūrs Znotiņš, Shay B. Cohen, and Guntis Barzdins. 2018. [Multilingual clustering of streaming news](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4535–4544, Brussels, Belgium. Association for Computational Linguistics.
- Kiem-Hieu Nguyen, Xavier Tannier, and Veronique Moriceau. 2014. [Ranking multidocument event descriptions for building thematic timelines](#). In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, pages 1208–1217, Dublin, Ireland. Dublin City University and Association for Computational Linguistics.
- Qiang Ning, Hao Wu, and Dan Roth. 2018. [A multi-axis annotation scheme for event temporal relations](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1318–1328, Melbourne, Australia. Association for Computational Linguistics.
- Alexander P. D. Mourelatos. 1978. Events, processes, and states. *Linguistics and Philosophy*, 2:415–434.
- Haoruo Peng, Snigdha Chaturvedi, and Dan Roth. 2017. A joint model for semantic sequences: Frames, entities, sentiments. In *Proceedings of*

- the 21st Conference on Computational Natural Language Learning (CoNLL 2017), Vancouver, Canada, August 3-4, 2017, pages 173–183.
- Hannah Rashkin, Maarten Sap, Emily Allaway, Noah A. Smith, and Yejin Choi. 2018. [Event2Mind: Commonsense inference on events, intents, and reactions](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 463–473, Melbourne, Australia. Association for Computational Linguistics.
- Maarten Sap, Ronan Le Bras, Emily Allaway, Chandra Bhagavatula, Nicholas Lourie, Hannah Rashkin, Brendan Roof, Noah A. Smith, and Yejin Choi. 2019. ATOMIC: an atlas of machine commonsense for if-then reasoning. In *The Thirty-Third AAAI Conference on Artificial Intelligence*, pages 3027–3035.
- Kailash Karthik Saravanakumar, Miguel Ballesteros, Muthu Kumar Chandrasekaran, and Kathleen McKeown. 2021. [Event-driven news stream clustering using entity-aware contextual embeddings](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 2330–2340, Online. Association for Computational Linguistics.
- Todor Staykovski, Alberto Barrón-Cedeno, Giovanni Da San Martino, and Preslav Nakov. 2019. Dense vs. sparse representations for news stream clustering. In *Text2Story@ ECIR*, pages 47–52.
- Ananya Subburathinam, Di Lu, Heng Ji, Jonathan May, Shih-Fu Chang, Avirup Sil, and Clare Voss. 2019. Cross-lingual structure transfer for relation and event extraction. In *Proc. 2019 Conference on Empirical Methods in Natural Language Processing and 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP2019)*.
- Haoyu Wang, Muhao Chen, Hongming Zhang, and Dan Roth. 2020. Joint constrained learning for event-event relation extraction. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Lu Wang, Claire Cardie, and Galen Marchetti. 2015. [Socially-informed timeline generation for complex events](#). In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1055–1065, Denver, Colorado. Association for Computational Linguistics.
- Hui Yang, Tat-Seng Chua, Shuguang Wang, and Chun-Keat Koh. 2003. [Structured use of external knowledge for event-based open domain question answering](#). In *Proceedings of the 26th annual international ACM SIGIR conference on Research and development in informaion retrieval*, pages 33–40.
- Qi Zeng, Manling Li, Tuan Lai, Heng Ji, Mohit Bansal, and Hanghang Tong. 2021. Gene: Global event network embedding. In *Proc. NAACL-HLT2021 Workshop on Graph-Based Natural Language Processing*.
- Hongming Zhang, Muhao Chen, Haoyu Wang, Yangqiu Song, and Dan Roth. 2020a. Open-domain process structure induction. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Hongming Zhang, Daniel Khashabi, Yangqiu Song, and Dan Roth. 2020b. Transomcs: From linguistic graphs to commonsense knowledge. In *Proceedings of IJCAI 2020*, pages 4004–4010.
- Hongming Zhang, Xin Liu, Haojie Pan, Yangqiu Song, and Cane Wing-Ki Leung. 2020c. ASER: A large-scale eventuality knowledge graph. In *WWW '20: The Web Conference 2020, Taipei, Taiwan, April 20-24, 2020*, pages 201–211.
- Tianran Zhang, Muhao Chen, and Alex Bui. 2020d. Diagnostic prediction with sequence-of-sets representation learning for clinical event. In *Proceedings of the 18th International Conference on Artificial Intelligence in Medicine (AIME)*.
- Zixuan Zhang and Heng Ji. 2021. Abstract meaning representation guided graph encoding and decoding for joint information extraction. In *Proc. The 2021 Conference of the North American Chapter of the Association for Computational Linguistics - Human Language Technologies (NAACL-HLT2021)*.
- Ben Zhou, Qiang Ning, Daniel Khashabi, and Dan Roth. 2020. [Temporal common sense acquisition with minimal supervision](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7579–7589, Online. Association for Computational Linguistics.
- Cunchao Zhu, Muhao Chen, Changjun Fan, Guangquan Cheng, and Yan Zhan. 2021. Learning from history: Modeling temporal knowledge graphs with sequential copy-generation networks. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*.