

Event-Centric Natural Language Understanding

Muhao Chen¹, Hongming Zhang², Qiang Ning³, Manling Li⁴, Heng Ji^{4,5} and Dan Roth¹

¹Department of Computer and Information Science, University of Pennsylvania

²Department of Computer Science and Engineering, Hong Kong University of Science and Technology

³Allen Institute for AI

⁴Department of Computer Science, University of Illinois at Urbana-Champaign

⁵Amazon Alexa AI

1 Goal of Tutorial

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 of events, (ii) various methods for automated extraction, conceptualization and prediction of events and their relations, (iii) induction of event processes and properties, and (iv) a wide range of NLU and commonsense understanding tasks that benefit from aforementioned techniques. We will conclude the tutorial by outlining emerging research problems in this area.

Keywords: Event-centric NLU, information extraction, commonsense understanding, structured learning and inference.

2 Tutorial Description

Human languages always involve the description of real-world events. Therefore, understanding events plays a critical role in natural language understanding (NLU). For example, narrative prediction benefits from learning the causal relations of events to predict what happens next in a story (Chaturvedi, Peng, and Roth 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, Chen, and Bui 2020). 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, Lu, and Roth 2012) and text summarization (Daumé III and Marcu 2006). 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 NLU 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 understand 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 NLU applications.

3 Outline of Tutorial Content

This tutorial presents a systematic overview of recent advances in event-centric NLU 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 NLU applications of various domains, and will outline emerging research challenges that may catalyze further investigation on this topic. The detailed contents are outlined below.

Motivation [20min]

We will define the main research problem and motivate the topic by presenting several real-world applications based on

event-centric NLU.

Background of Events and Their Representations [40min]

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, Fillmore, and Lowe 1998), event knowledge graphs (Zhang et al. 2020c), event processes (Chambers and Jurafsky 2008), event language models (Peng, Chaturvedi, and Roth 2017), and more recent work on event meaning representation via question-answer pairs (He, Lewis, and Zettlemoyer 2015; Michael et al. 2018), event network embeddings and event time expression embeddings (Goyal and Durrett 2019).

Event-Centric Information Extraction [45min]

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), and joint entity, relation and event extraction (Lin et al. 2020). Besides, we will also discuss methods that identify temporal and causal relations of primitive events (Ning, Wu, and Roth 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.

Event-centric Prediction [45min]

We will then present recent works on machine comprehension and prediction of events. Specifically, people are trying to understand events from different angles. For example, many efforts have been devoted into modeling event narratives (Peng, Chaturvedi, and Roth 2017; Chaturvedi, Peng, and Roth 2017b; Lee and Goldwasser 2019) such that they can successfully predict missing events in an event sequence. Besides, another important event understanding angle is conceptualization (Zhang et al. 2020a), which aims at understanding the super-sub relations between events (Glavaš et al. 2014). 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).

Event-centric Commonsense Knowledge Acquisition [40min]

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. 2020b) of events.

Emerging Research Problems [20min]

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

4 History, Expected Attendance and Prerequisite Knowledge

This tutorial has not been presented elsewhere. Based on the level of interest in this topic, we expect around 100 participants. No specific background knowledge is assumed of the audience.

5 Tutorial Instructors

Muhao Chen is a postdoctoral fellow in Department of Computer and Information Science, UPenn. He received his Ph.D. in Computer Science from UCLA in 2019, where he was supported by the UCLA Dissertation Fellowship. Muhao's research focuses on data-driven machine learning approaches for processing structured and unstructured data, and extending their applications to natural language understanding, knowledge base construction, computational biology and medicine. Particularly, he is interested in developing knowledge-aware learning systems with generalizability and requiring minimal supervision. He has published over 40 papers in leading venues of these areas, and has received a SIGBio Best Student Paper Award at ACM BCB 2020. 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 a research scientist on the AllenNLP team at the Allen Institute for AI (AI2). 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 <http://qing2.web.engr.illinois.edu/>.

Manling Li is a third-year Ph.D. student at the Computer Science Department of the University of Illinois at Urbana-Champaign (UIUC). She obtained her master's degree in Computer Science from the Institute of Computing Technology, Chinese Academy of Sciences in 2018. Manling has won the Best Demo Paper Award at ACL'20 about extracting knowledge from multimedia data, and has more than 20 publications on knowledge extraction and reasoning. 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). Ad-

ditional information is available at <https://blender.cs.illinois.edu/hengji.html>.

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, AAAI, 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 AAAI 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 AAAI'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:

- Muhao Chen:
 - AAAI'20: Recent Advances of Transferable Representation Learning.
- Heng Ji:
 - 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.
- Dan Roth:
 - Data Science Summer Institute (DSSI) 2007, 2008, 2010, 2011, 2012. A tutorial on Machine Learning in Natural Language Processing.

- 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 research 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

Aldawsari, M., and Finlayson, M. 2019. Detecting subevents using discourse and narrative features. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 4780–4790. Florence, Italy: Association for Computational Linguistics.

Bach, E. 1986. The algebra of events. *Linguistics and philosophy* 9(1):5–16.

Baker, C. F.; Fillmore, C. J.; and Lowe, J. B. 1998. The berkeley framenet project. In *COLING-ACL*, 86–90.

Berant, J.; Srikumar, V.; Chen, P.-C.; Vander Linden, A.; Harding, B.; Huang, B.; Clark, P.; and Manning, C. D. 2014. Modeling biological processes for reading comprehension. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 1499–1510. Doha, Qatar: Association for Computational Linguistics.

Chambers, N., and Jurafsky, D. 2008. Unsupervised learning of narrative event chains. In *Proceedings of ACL-08: HLT*, 789–797. Columbus, Ohio: Association for Computational Linguistics.

Chaturvedi, S.; Peng, H.; and Roth, D. 2017a. Story comprehension for predicting what happens next. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, 1603–1614. Copenhagen, Denmark: Association for Computational Linguistics.

Chaturvedi, S.; Peng, H.; and Roth, D. 2017b. Story Comprehension for Predicting What Happens Next. In *In proceedings of the Conference on Empirical Methods in Natural Language Processing*.

Chen, M.; Zhang, H.; Wang, H.; and Roth, D. 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.

Cybulska, A., and Vossen, P. 2014. Guidelines for ecb+ annotation of events and their coreference. In *Technical Report*. Technical Report NWR-2014-1, VU University Amsterdam.

Daumé III, H., and Marcu, D. 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*, 305–312. Sydney, Australia: Association for Computational Linguistics.

Ding, X.; Zhang, Y.; Liu, T.; and Duan, J. 2015. Deep learning for event-driven stock prediction. In *Twenty-fourth international joint conference on artificial intelligence*.

Do, Q.; Lu, W.; and Roth, D. 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*, 677–687. Jeju Island, Korea: Association for Computational Linguistics.

Glavaš, G.; Šnajder, J.; Moens, M.-F.; and Kordjamshidi, P. 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)*, 3678–3683. Reykjavik, Iceland: European Languages Resources Association (ELRA).

Goyal, T., and Durrett, G. 2019. Embedding time expressions for deep temporal ordering models. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics 2019 (ACL2019)*.

- He, L.; Lewis, M.; and Zettlemoyer, L. 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*, 643–653.
- Huang, L.; Ji, H.; Cho, K.; Dagan, I.; Riedel, S.; and Voss, C. 2018. Zero-shot transfer learning for event extraction. In *Proc. The 56th Annual Meeting of the Association for Computational Linguistics (ACL2018)*.
- Lee, I.-T., and Goldwasser, D. 2019. Multi-relational script learning for discourse relations. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 4214–4226. Florence, Italy: Association for Computational Linguistics.
- Li, T.; Gupta, V.; Mehta, M.; and Srikumar, V. 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)*, 3924–3935. Hong Kong, China: Association for Computational Linguistics.
- Li, M.; Zareian, A.; Zeng, Q.; Whitehead, S.; Lu, D.; Ji, H.; and Chang, S.-F. 2020a. Cross-media structured common space for multimedia event extraction. In *Proc. The 58th Annual Meeting of the Association for Computational Linguistics (ACL2020)*.
- Li, M.; Zeng, Q.; Lin, Y.; Cho, K.; Ji, H.; May, J.; Chambers, N.; and Voss, C. 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)*.
- Lin, Y.; Ji, H.; Huang, F.; and Wu, L. 2020. A joint neural model for information extraction with global features. In *Proc. The 58th Annual Meeting of the Association for Computational Linguistics (ACL2020)*.
- Liu, H., and Singh, P. 2004. Conceptnet: a practical commonsense reasoning tool-kit. *BT technology journal* 22(4):211–226.
- Liu, J.; Cheng, F.; Wang, Y.; Shindo, H.; and Matsumoto, Y. 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.
- Michael, J.; Stanovsky, G.; He, L.; Dagan, I.; and Zettlemoyer, L. 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*, 560–568.
- Ning, Q.; Wu, H.; and Roth, D. 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)*, 1318–1328. Melbourne, Australia: Association for Computational Linguistics.
- P. D. Mourelatos, A. 1978. Events, processes, and states. *Linguistics and Philosophy* 2:415–434.
- Peng, H.; Chaturvedi, S.; and Roth, D. 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*, 173–183.
- Rashkin, H.; Sap, M.; Allaway, E.; Smith, N. A.; and Choi, Y. 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)*, 463–473. Melbourne, Australia: Association for Computational Linguistics.
- Sap, M.; Bras, R. L.; Allaway, E.; Bhagavatula, C.; Lourie, N.; Rashkin, H.; Roof, B.; Smith, N. A.; and Choi, Y. 2019. ATOMIC: an atlas of machine commonsense for if-then reasoning. In *The Thirty-Third AAAI Conference on Artificial Intelligence*, 3027–3035.
- Subburathinam, A.; Lu, D.; Ji, H.; May, J.; Chang, S.-F.; Sil, A.; and Voss, C. 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)*.
- Wang, H.; Chen, M.; Zhang, H.; and Roth, D. 2020. Temporal common sense acquisition with minimal supervision. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Yang, H.; Chua, T.-S.; Wang, S.; and Koh, C.-K. 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*, 33–40.
- Zhang, H.; Chen, M.; Wang, H.; Song, Y.; and Roth, D. 2020a. Open-domain process structure induction. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Zhang, H.; Khashabi, D.; Song, Y.; and Roth, D. 2020b. Transoms: From linguistic graphs to commonsense knowledge. In *Proceedings of IJCAI 2020*, 4004–4010.
- Zhang, H.; Liu, X.; Pan, H.; Song, Y.; and Leung, C. W. 2020c. ASER: A large-scale eventuality knowledge graph. In *WWW '20: The Web Conference 2020, Taipei, Taiwan, April 20-24, 2020*, 201–211.
- Zhang, T.; Chen, M.; and Bui, A. 2020. Diagnostic prediction with sequence-of-sets representation learning for clinical event. In *Proceedings of the 18th International Conference on Artificial Intelligence in Medicine (AIME)*.
- Zhou, B.; Ning, Q.; Khashabi, D.; and Roth, D. 2020. Temporal common sense acquisition with minimal supervision. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 7579–7589. Online: Association for Computational Linguistics.