

Abstract: Understanding events and communicating about them are fundamental human activities. However, it's much more difficult to remember event-related information compared to entity-related information. For example, most people in the United States will be able to answer the question "Which city is Columbia University located in?", but very few people can give a complete answer to "Who died from COVID-19". Human-written history books are often incomplete and highly biased because "History is written by the victors". For example, the history textbooks used in Japanese schools barely mention details about the Nanjing Massacre, and even so contain a lot of false information. In this talk I will describe an ambitious ongoing project on automatic history book generation. We propose a new research direction on event-centric knowledge base construction from multimedia multilingual sources. Our minds represent events at various levels of granularity and abstraction, which allows us to quickly access and reason about old and new scenarios. Progress in natural language understanding and computer vision has helped automate some parts of event understanding but the current, first-generation, automated event understanding is overly simplistic since it is local, sequential and flat. Real events are hierarchical and probabilistic. Understanding them requires knowledge in the form of a repository of abstracted event schemas (complex event templates), understanding the progress of time, using background knowledge, and performing global inference. Our approach to second-generation event understanding builds on an incidental supervision approach to inducing an event schema repository that is probabilistic, hierarchically organized and semantically coherent. Low level primitive components of event schemas are abundant, and can be part of multiple, sparsely occurring, higher-level schemas. Consequently, we combine bottom-up data driven approaches across multiple modalities with top-down consolidation of information extracted from a smaller number of encyclopedic resources. This facilitates inducing higher-level event representations analysts can interact with, and allow them to guide further reasoning and extract events by constructing a novel structured cross-media common semantic space. When complex events unfold in an emergent and dynamic manner, the multimedia multilingual digital data from traditional news media and social media often convey conflicting information. To understand the many facets of such complex, dynamic situations, we have developed various novel methods to induce hierarchical narrative graph

schemas and apply them to enhance end-to-end joint neural Information Extraction, event coreference resolution and event time prediction.

Bio: 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. Before that she was an associate professor at Queens College and Graduate Center of City University of New York, and then Edward P. Hamilton Development Chair Professor at Rensselaer Polytechnic Institute. 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 is the associate editor for IEEE/ACM Transaction on Audio, Speech, and Language Processing, and 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, Facebook, Bosch, IBM, Disney).