

# RESIN-EDITOR: A Schema-guided Hierarchical Event Graph Visualizer and Editor

Khanh Duy Nguyen<sup>1\*</sup>, Zixuan Zhang<sup>1\*</sup>, Reece Suchocki<sup>2</sup>, Sha Li<sup>1</sup>,  
Martha Palmer<sup>2</sup>, Susan Brown<sup>2</sup>, Jiawei Han<sup>1</sup>, Heng Ji<sup>1</sup>

<sup>1</sup>University of Illinois Urbana-Champaign

<sup>2</sup>University of Colorado Boulder

{knguye71, zixuan11, sha12, hanj, hengji}@illinois.edu  
{reece.suchocki, susan.brown, martha.palmer}@colorado.edu

## Abstract

In this paper, we present RESIN-EDITOR, an interactive event graph visualizer and editor designed for analyzing complex events. Our RESIN-EDITOR system allows users to render and freely edit hierarchical event graphs extracted from multimedia and multi-document news clusters with guidance from human-curated event schemas. RESIN-EDITOR's unique features include hierarchical graph visualization, comprehensive source tracing, and interactive user editing, which is more powerful and versatile than existing Information Extraction (IE) visualization tools. In our evaluation of RESIN-EDITOR, we demonstrate ways in which our tool is effective in understanding complex events and enhancing system performance. The source code<sup>1</sup>, a video demonstration<sup>2</sup>, and a live website<sup>3</sup> for RESIN-EDITOR have been made publicly available.

## 1 Introduction

Complex events, such as an international negotiation or a disease outbreak, can take place over a prolonged period of time spanning from weeks to even months. Typically, such complex events can be further deconstructed into many atomic sub-events, where each occurs at a specific time and place. Fully modeling such complex events requires an understanding of the temporal, logical, and hierarchical connections among many sub-events, making the task of event modeling difficult for existing Information Extraction (IE) visualization and analysis tools. Recent research efforts (Wang et al., 2020; Du et al., 2022; Wang et al., 2022) extend the event understanding paradigm to model event-event relations and propose to use event schemas to

guide the organization of complex event structures and predict possible future events. However, convenient visualization and editing tools specifically designed to handle such types of complex events are still largely undeveloped.

A visualization and editing tool is essential for developing and improving an IE system. By utilizing such a tool, system developers can conduct a qualitative error analysis, which involves human analysts carefully investigating any missing or spurious errors, tracing these inaccuracies back to the original documents, and record any proposed modifications or suggestions. Existing tools designed for atomic events (Sheng et al., 2018; Li et al., 2019) are limited when directly applied to complex events, which are much more intricate with complicated event-event relations and multi-document multi-media source inputs. Furthermore, prior analysis tools lack support for interactive user editing, which poses a challenge for model developers to comprehend any real feedback from actual users.

To address the previously discussed challenges and accelerate research of complex events, we present the RESIN-EDITOR, an analysis interface that allows users to visualize, analyze, and freely edit the multi-granularity graphical event narratives extracted from multi-document multimedia news clusters. Compared with existing analysis interfaces, our tool has the following three significant advantages: 1) **Hierarchical Graph Visualization**. Our tool is able to generate a comprehensive visualization of the extracted hierarchical event graph (as shown in Figures 1 and 2). This graph not only includes atomic events and entities, but also includes all event-event hierarchical, temporal, and logical relations of the entire complex event. Such an informative visualization can offer human analysts a rapid and comprehensive understanding of any complex event, significantly improving the efficiency of our error analysis process. Moreover, our graphical visualization is initialized with

\*Equal Contribution.

<sup>1</sup>Data and source code: <https://github.com/blender-nlp/RESIN-Editor>

<sup>2</sup>Video demonstration: <https://www.youtube.com/watch?v=fmW-GwPMrw0>

<sup>3</sup>Live website: <https://blender-nlp.github.io/RESIN-Editor/>

carefully-curated schemas, and the extracted events are grounded to the schema whenever possible, which improves the reliability of our error analysis. 2) **Comprehensive Source Tracing**. Our tool is able to show multimedia provenances for all extracted events, entities, and relations, including source documents, text spans, images, and bounding boxes. These trace-back provenances are presented in a clear-structured and tabular format, which greatly simplifies the users' task of understanding and evaluating the performance of the extraction system. 3) **Interactive User Editing**. Most importantly, with our analysis tool, users are allowed to edit and manipulate almost every extracted element depicted in the visualized event graph and save their edited event graph for downstream usage. This includes not only the node attributes of events and entities, such as names, types, and description sentences but also the graphical structures, like the entity-entity relations and event-event temporal links. Through open and direct interactions with intended users, our tool equips model developers with a comprehensive and precise understanding of actual user requirements.

Our main contributions can be summarized as:

- A novel development to a multimedia analysis tool for complex event understanding, which enables multi-granularity visualization, comprehensive source document tracing, and interactive user editing.
- Extensive empirical studies show that our tool is significantly more effective than existing tools in understanding and analyzing complex events.

## 2 Backend System Overview

The RESIN-EDITOR system provides an interface for event graph visualization with the added capability of graph editing. The visualized graph is an instantiated schema graph, which results from matching an extracted event graph with an event schema. An event schema is a prototypical description of the usual progression of a particular type of complex event, which can be induced from historical news documents.

In this section, we present an overview of how this event graph is generated, with the main components being multimedia event extraction, cross-

document event organization, and event prediction in hierarchical schemas.

**Multimedia Event Extraction** The input to our system is a multimedia news cluster. This cluster includes multiple news documents with related images from a variety of news sources, all describing a single news event. From this cluster, we first perform event extraction for atomic events. Following the convention of ACE-2005<sup>4</sup>, each event is composed of an event trigger as well as several event arguments. Since we would like our system to work for all news scenarios, we employ GLEN (Zhan et al., 2023), the event trigger extraction system that works for 3k+ event types, making it the most comprehensive event detection system to date. We then extract event arguments with a template-based generation approach (Li et al., 2021) distilled from a large language model that is prompted to extract arguments as code structures (Wang et al., 2023). To handle visual input, we use (Li et al., 2020) to extract events from images and merge them with events extracted from text. We represent the atomic events from the multimedia event extraction system as star-shaped graphs such as the *Diagnosis* event shown in Figure 1.

**Cross-Document Event Organization** After extracting atomic events and entities, all event and entity mentions that correspond to the same real-life event and entity are merged through coreference resolution (Lai et al., 2021), and linked to Wikidata through entity linking (Lai et al., 2022). We then conduct event-event temporal relation extraction (Jin et al., 2023) to organize the events into a temporal event graph. We refer to the resulting event graph as the instance event graph, which serves as a semantic representation of the input news cluster in aggregate. The event graph exemplifies unfolding narratives within the given news cluster, as well as connects entities across articles.

**Matching and Prediction with Hierarchical Schemas** After extracting the instance event graph, we perform event instantiation and prediction with the help of event schemas. The schema is typically depicted as a graph in which the nodes are generalized event types linked by inter-event relations, and our task is to instantiate the schema with our instance event graph and make predictions for those unmatched events. As compared

<sup>4</sup><https://www ldc.upenn.edu/collaborations/past-projects/ace>

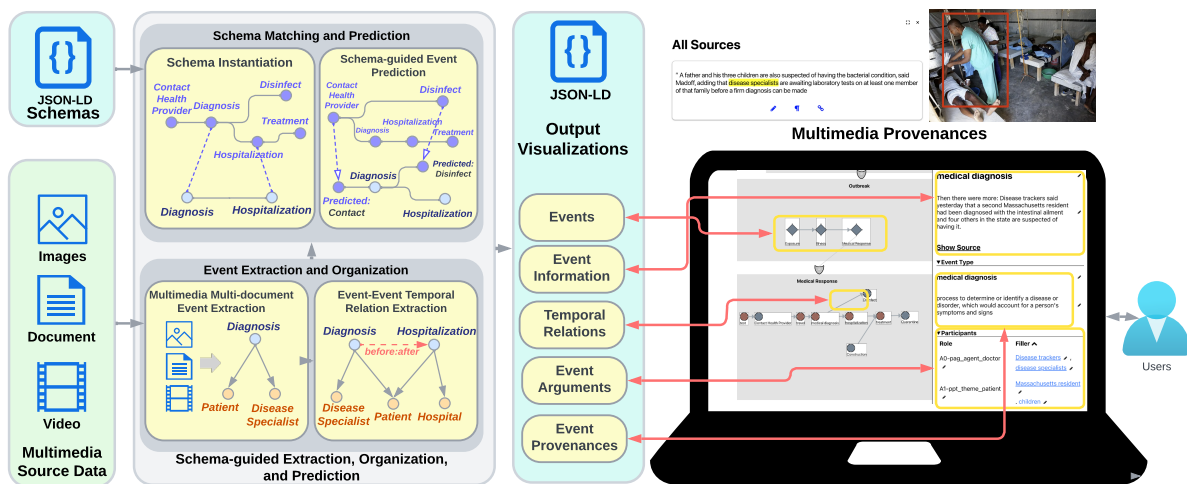


Figure 1: Overview of our RESIN-EDITOR system. Given a multimedia multi-document news cluster and the schema graph, our system extracts the atomic events, organizes them into an instance event graph, performs schema-guided matching and prediction, and finally interactively visualizes the results for user editing.

to previous systems (Du et al., 2022; Wang et al., 2022), we now facilitate event matching and prediction on hierarchical schema structures, where the events in the schema are presented with parent-child relations to represent the different granularity of events. All of the extracted event nodes in our system are grounded in the structured event hierarchy, either by matching to a schema node, or being attached as a child event to a schema node. Event and entity nodes in our extracted event graph are matched top-down, enabling event prediction based on our graphical neural network event prediction system (Wang et al., 2022).

### 3 Graphical Interface

Our graphical interface visualizes the instantiated schema graph after matching and prediction along with the source documents that provided provenance for the extracted events. The interface (Figure 2) is organized in the panel hierarchy illustrated in Figure 3. The graph panel, which is the canvas for showing the event graph, is the primary access point. The overlay panels, such as the entity table, provenance, and mini-map, enhance the exploration experience with contextually relevant information or provide navigation support.

#### 3.1 Interactive Graph Panel

We adopt a graph-based representation for the schema-matching results to provide a streamlined and efficient error analysis interface for event analysts. For each event, we use color coding to illustrate whether the event is extracted from the

source document (source-only), from the schema (predicted), or found in both (matched).

To assist analysts or system developers in making informed decisions, we find it critical to present the relations between events clearly and disentangle different types of relations. We support 3 types of relations among events in our visualization tool:

- **Hierarchical event-event relations.** Parent-child relations between events are visualized by arranging them on the vertical axis: parent events are shown above their children events. Primitive events, which are leaf nodes in the graph, are displayed as circular nodes, whereas parent events are displayed as diamond-shaped nodes. Initially, we only show the top-level events as an overview of the scenario. Users can click on the parent event to expand the graph and dive deeper into the details of the event.
- **Temporal relations.** We order events with temporal relations along the horizontal axis from left to right. In the case of misaligned temporal relations, the user can change the event order by editing the temporal edge.
- **Logical relations.** We support the usage of logical relations defined on a group of events to enforce stronger logical constraints. The *OR gate* signifies that one or more events in the group can occur, the *AND gate* necessitates the occurrence of all events, and the *XOR gate* permits exclusively one event to happen. Given the common use of digital logic gates

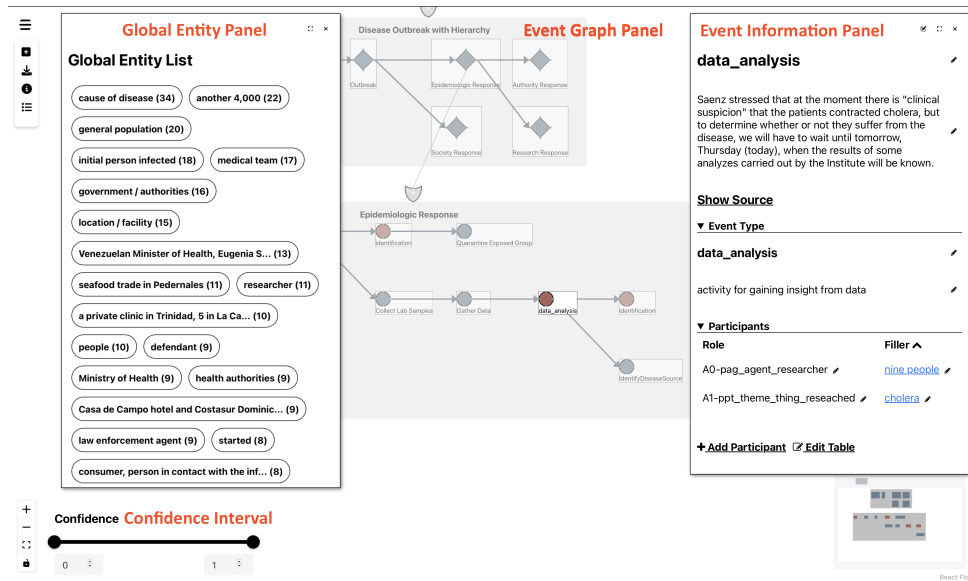


Figure 2: The RESIN-EDITOR interface with *Cholera Decease Outbreak* event graph.




Figure 3: The View Panel Hierarchy illustrates the call sequence of the panels. The most significant panel is located at the top, and subsequent panels are activated from their immediate parent panels.

in computer science and the conceptual analogy they provide, we employ them to depict logical gates in our event graph (Figure 4).

To improve navigation efficiency in the Graph Panel, we include direct navigation through drag panels and event dragging, a zoom function for detailed exploration, and a minimap for a macroscopic overview of the graph. More detailed information and context can be accessed by clicking on an event node, which triggers an information panel.

### 3.2 Information Panel

The individual observer without assistance has significant limitations regarding how much information they can perceive, process, and retain (Miller, 1956). To avoid interface clutter, we partition the event's information into three distinct sections:

**Event Description** Include the event name and the description of the event. Furthermore, descriptions of events can be viewed swiftly by hovering over the events on the Graph Panel.




Figure 4: After an *Illness* event, possible subsequent events include *Symptoms*, *Illness Outcomes*, or *Confirmation*. Following the *Illness Outcomes* event, only a single *Death Outcomes* event can occur. The *Death Outcomes* event then invariably leads to both *Death* and *Funeral* events.

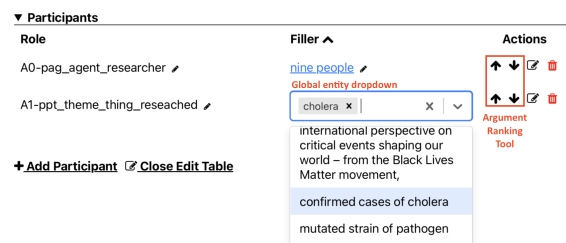


Figure 5: Updating the argument lists of the *Data Analysis* event, incorporating the *Confirmed Cases of Cholera* entity into the theme research.






Figure 6: Multimedia provenances support *Disease Specialist* entity.

**Event Type** The event type information is either from the event schema or the event detection system. We use the XPO ontology (Spaulding et al., 2023) for the scope of event types covered by our system. This information might prove helpful in cases where the event description is insufficiently descriptive.

**Argument Table** Each event is associated with multiple argument roles as defined by the ontology, and each argument role can have multiple fillers. We present this information in the argument table which consists of rows of arguments (e.g., location, places, participants) and their associated entities, as shown in Figure 5. We include only the entities backed by sources, leaving out generic ones. In editing mode, the user can reorder the arguments per their needs, ensuring that the most relevant arguments are prioritized at the top. Furthermore, the edit mode provides a comprehensive list of entities from which analysts can add any relevant missing entities as arguments. Each information display includes an editing function, allowing analysts to adjust the presented event information while analyzing schema-matching results.

### 3.3 Provenance Panel

Identifying the source document that validates schema matching is essential for analysts during error analysis because it helps to verify any inaccuracies either from the IE system or in matching process. Hence, in the provenance panel (Figure 6), we display the text and image provenances for source-only events, matched events, and entities. We utilize text highlighting within the paragraph to indicate text provenances, allowing the user to find the relevant citation in the source document easily. Additionally, image provenances employ a bounding box to indicate the referenced entity in the image. Under each piece of provenance, we

provide three functions:

**Edit the provenance mentions.** Analysts can substitute inaccurate text spans in the source document with the correct trigger span for text provenance. For image provenance, the bounding box can be modified to suit the entity in the image better. For instance, in Figure 6, analysts can realign and resize the bounding box to correspond with the nurse image. This adjustment allows for adequately presenting the *Disease Specialist* entity inside the provenance of the *Diagnosis* event.

**Expand the provenance context.** Analysts can expand and view the entire paragraph containing the extracted provenance, but only for text provenance.

**Go to source document.** The feature enables tracing back to the original documents containing the sentence, providing analysts with additional contextual information about the provenance. This tool becomes essential when analysts need to authenticate the provenance’s integrity.

### 3.4 Event Filtering Tools

To ease navigation through numerous graph events, we develop these two event filters:

**Global Entity List** This feature aids analysts in isolating entity-related events by reducing the opacity of unrelated events. For instance, selecting *Cholera* entity in Figure 2 emphasizes the connected *Data Analysis* event, while other events become dim. Moreover, the global entities list ranks entities based on their occurrence in graph events, highlighting those most frequently mentioned across all events.

**Confidence Interval** The schema-matching process between the instance graph and the schema assigns a confidence value to each event from its schema matching and prediction component (Wang et al., 2022). This tool is beneficial when analysts aim to discover matching events with less certainty. By selecting a lower confidence interval, analysts can monitor all events within that range and investigate those with low confidence values.

## 4 Experiments

We present our empirical experiments with multiple document clusters to evaluate the effectiveness of our editor.

**Evaluation Setup** We use two high-impact news scenarios *Disease Outbreak* and *Terrorist Attacks* with four newsworthy complex events, where the detailed statistics are shown in Table 1.

**Results** Based on the results shown in Table 2, RESIN-EDITOR significantly outperforms the traditional method of direct source document analysis, showing considerable improvements across all main IE tasks in our backend system. Such an improvement is particularly notable in the prediction task, probably because the original documents offer limited supporting information for predicting new events. However, our visualizer can provide a well-organized graphical event timeline, which is greatly beneficial for comprehending the events and making accurate predictions.

News Cluster	# Docs	# Images
Cholera outbreak, Dominica	13	114
E. coli outbreak, USA	16	118
Legionnaires outbreak, USA	14	66
Mogadishu bombings, Somalia	13	57

Table 1: Statistics of the news document clusters.

Tasks	Prec	Rec	F1
Event Triggers (documents)	75.5	69.0	72.1
Event Triggers (visualizer)	84.9	76.3	<b>80.4</b>
Event Arguments (documents)	68.0	58.6	63.0
Event Arguments (visualizer)	84.0	72.4	<b>77.8</b>
Schema Matching (documents)	67.9	62.3	65.0
Schema Matching (visualizer)	71.1	64.5	<b>67.6</b>
Prediction (documents)	62.9	50.0	55.7
Prediction (visualizer)	93.1	66.7	<b>77.8</b>

Table 2: The improved IE task performances using our designed visualizer compared with directly looking into the source documents. We focus on the four main IE tasks in our backend system, and the performances are characterized in precision, recall, and F1 scores.

## 5 Related Work

In the IE visualization domain, (Jenkins et al., 2023; Vacareanu et al., 2022) offer event-focused indexed documents, while others (Ma et al., 2021; Ning et al., 2018) use temporal-relation entity graphs. However, both methods lack hierarchical event organization which leads to clutter in the graph display. Additionally, despite presenting entity provenance triggers, they are not able to assist in amending extraction errors.

In terms of schema visualization, Mishra et al. (2021) introduce a graph-based interface tailored

to depict schemas, emphasizing event order and entity relations. While the interface allows users to edit schemas, the user must come up with editing suggestions based on his/her own knowledge without any help from the interface. More recently, Zhang et al. (2023) introduce a human-in-the-loop interface that employs pre-trained LLMs to help create schema graphs detailing event sequences in scenarios. Whereas the use of LLMs can alleviate the manual effort of the user, the LLM-generated schema is not grounded in real news documents and may be prone to hallucination.

Existing tools mainly focus on either IE extraction or schema visualization. Inspired by the RESIN-11 system (Du et al., 2022), which generates schema-guided event graphs from multi-modal news documents, we introduce a unified tool that integrates both the schema and the extracted event graph, providing a comprehensive visualization for complex event analysis. By matching the extracted event graph with the human-curated schema, we are able to effectively organize the complex event and guide users in providing informed feedback on both the IE results and the schema.

## 6 Conclusion and Future Work

To the best of our knowledge, RESIN-EDITOR is the first graphical user interface for editing and visualizing schema-guided event graphs. Our tool assists users by highlighting the matching results, allowing them to fix discrepancies between event attributes and their argument entities and resolve disparities between extracted entities and their multimedia provenance. Our results show that our solution outperforms the traditional method of error straight from their sources in primary IE tasks. While we recognize the importance of comparing with other tools, our current experimentation is limited due to time constraints. We plan to conduct extensive experiments in the future to validate our results further.

In future work, we aim to develop a ranking system that considers the relevance of multimedia provenance so that users can quickly view the most related sentence/image. We also plan to integrate the event grounding tool, so that users can easily add new events by providing natural language descriptions, and the system will provide event type suggestions and fill in the argument template.

## Acknowledgements

This research is based on work supported by U.S. DARPA KAIROS Program No. FA8750-19-2-1004. 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

- Xinya Du, Zixuan Zhang, Sha Li, Pengfei Yu, Hongwei Wang, Tuan Lai, Xudong Lin, Ziqi Wang, Iris Liu, Ben Zhou, Haoyang Wen, Manling Li, Darryl Hannan, Jie Lei, Hyounghun Kim, Rotem Dror, Haoyu Wang, Michael Regan, Qi Zeng, Qing Lyu, Charles Yu, Carl Edwards, Xiaomeng Jin, Yizhu Jiao, Ghazaleh Kazeminejad, Zhenhailong Wang, Chris Callison-Burch, Mohit Bansal, Carl Vondrick, Jiawei Han, Dan Roth, Shih-Fu Chang, Martha Palmer, and Heng Ji. 2022. [RESIN-11: Schema-guided event prediction for 11 newsworthy scenarios](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: System Demonstrations*, pages 54–63, Hybrid: Seattle, Washington + Online. Association for Computational Linguistics.
- Chris Jenkins, Shantanu Agarwal, Joel Barry, Steven Fincke, and Elizabeth Boschee. 2023. [Massively multi-lingual event understanding: Extraction, visualization, and search](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations)*, pages 247–256, Toronto, Canada. Association for Computational Linguistics.
- Xiaomeng Jin, Haoyang Wen, Xinya Du, and Heng Ji. 2023. [Toward consistent and informative event-event temporal relation extraction](#). In *Proceedings of the First Workshop on Matching From Unstructured and Structured Data (MATCHING 2023)*, pages 23–32, Toronto, ON, Canada. Association for Computational Linguistics.
- Tuan Lai, Heng Ji, Trung Bui, Quan Hung Tran, Franck Dernoncourt, and Walter Chang. 2021. [A context-dependent gated module for incorporating symbolic semantics into event coreference resolution](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3491–3499, Online. Association for Computational Linguistics.
- Tuan Lai, Heng Ji, and ChengXiang Zhai. 2022. [Improving candidate retrieval with entity profile generation for Wikidata entity linking](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 3696–3711, Dublin, Ireland. Association for Computational Linguistics.
- Manling Li, Ying Lin, Joseph Hoover, Spencer Whitehead, Clare Voss, Morteza Dehghani, and Heng Ji. 2019. [Multilingual entity, relation, event and human value extraction](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 110–115, Minneapolis, Minnesota. Association for Computational Linguistics.
- Manling Li, Alireza Zareian, Qi Zeng, Spencer Whitehead, Di Lu, Heng Ji, and Shih-Fu Chang. 2020. [Cross-media structured common space for multimedia event extraction](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2557–2568, Online. Association for Computational Linguistics.
- Sha Li, Heng Ji, and Jiawei Han. 2021. [Document-level event argument extraction by conditional generation](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 894–908, Online. Association for Computational Linguistics.
- Mingyu Derek Ma, Jiao Sun, Mu Yang, Kung-Hsiang Huang, Nuan Wen, Shikhar Singh, Rujun Han, and Nanyun Peng. 2021. [EventPlus: A temporal event understanding pipeline](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: Demonstrations*, pages 56–65, Online. Association for Computational Linguistics.
- George A. Miller. 1956. [The magical number seven, plus or minus two: Some limits on our capacity for processing information](#). *Psychological Review*, 63(2):81–97.
- Piyush Mishra, Akanksha Malhotra, Susan Windisch Brown, Martha Palmer, and Ghazaleh Kazeminejad. 2021. [A graphical interface for curating schemas](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing: System Demonstrations*, pages 159–166, Online. Association for Computational Linguistics.
- Qiang Ning, Ben Zhou, Zhili Feng, Haoruo Peng, and Dan Roth. 2018. [CogCompTime: A tool for understanding time in natural language](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 72–77, Brussels, Belgium. Association for Computational Linguistics.
- Yongpan Sheng, Zenglin Xu, Yafang Wang, Xiangyu Zhang, Jia Jia, Zhonghui You, and Gerard de Melo. 2018. [Visualizing multi-document semantics via open domain information extraction](#). In *Machine*

*Learning and Knowledge Discovery in Databases - European Conference, ECML PKDD 2018, Dublin, Ireland, September 10-14, 2018, Proceedings, Part III*, volume 11053 of *Lecture Notes in Computer Science*, pages 695–699. Springer.

Elizabeth Spaulding, Kathryn Conger, Anatole Gershman, Rosario Uceda-Sosa, Susan Windisch Brown, James Pustejovsky, Peter Anick, and Martha Palmer. 2023. [The darpa wikidata overlay: Wikidata as an ontology for natural language processing](#). In *Workshop on Interoperable Semantic Annotation (ISA-19)*, page 12.

Robert Vacareanu, George C.G. Barbosa, Enrique Noriega-Atala, Gus Hahn-Powell, Rebecca Sharp, Marco A. Valenzuela-Escárcega, and Mihai Surdeanu. 2022. [A human-machine interface for few-shot rule synthesis for information extraction](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: System Demonstrations*, pages 64–70, Hybrid: Seattle, Washington + Online. Association for Computational Linguistics.

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 (EMNLP)*, pages 696–706, Online. Association for Computational Linguistics.

Hongwei Wang, Zixuan Zhang, Sha Li, Jiawei Han, Yizhou Sun, Hanghang Tong, Joseph P. Olive, and Heng Ji. 2022. [Schema-guided event graph completion](#). In *The 4th Conference on Automated Knowledge Base Construction (AKBC 2022)*.

Xingyao Wang, Sha Li, and Heng Ji. 2023. [Code4Struct: Code generation for few-shot event structure prediction](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3640–3663, Toronto, Canada. Association for Computational Linguistics.

Qiusi Zhan, Sha Li, Kathryn Conger, Martha Palmer, Heng Ji, and Jiawei Han. 2023. [Glen: General-purpose event detection for thousands of types](#).

Tianyi Zhang, Isaac Tham, Zhaoyi Hou, Jiaxuan Ren, Leon Zhou, Hainiu Xu, Li Zhang, Lara Martin, Rotem Dror, Sha Li, Heng Ji, Martha Palmer, Susan Windisch Brown, Reece Suchocki, and Chris Callison-Burch. 2023. [Human-in-the-loop schema induction](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations)*, pages 1–10, Toronto, Canada. Association for Computational Linguistics.