

ExcavatorCovid: Extracting Events and Relations from Text Corpora for Temporal and Causal Analysis for COVID-19

Bonan Min¹, Ben Rozonoyer¹, Haoling Qiu¹, Alex Zamanian¹,
Nianwen Xue², and Jessica MacBride¹

¹Raytheon BBN Technologies

²Brandeis University

bonan.min@raytheon.com

Abstract

Timely responses from policy makers to mitigate the impact of the COVID-19 pandemic rely on a comprehensive grasp of events, their causes, and their impacts. These events are reported at such a speed and scale as to be overwhelming. In this paper, we present ExcavatorCovid, a machine reading system that ingests open-source text documents (e.g., news and scientific publications), extracts COVID-19 related events and relations between them, and builds a Temporal and Causal Analysis Graph (TCAG). Excavator will help government agencies alleviate the information overload, understand likely downstream effects of political and economic decisions and events related to the pandemic, and respond in a timely manner to mitigate the impact of COVID-19. We expect the utility of Excavator to outlive the COVID-19 pandemic: analysts and decision makers will be empowered by Excavator to better understand and solve complex problems in the future. A demonstration video is available at <https://vimeo.com/528619007>.

1 Introduction

Timely responses from policy makers to mitigate the impact of the COVID-19 pandemic rely on a comprehensive grasp of events, their causes, and their impacts. Since the beginning of the COVID-19 pandemic, an enormous amount of articles are being published every day, that report many events¹ and studies related to COVID. It is very difficult, if not impossible, to keep track of these developing events or to get a comprehensive overview of the temporal and causal dynamics underlying these events.

To aid the policy makers in overcoming the information overload, we developed ExcavatorCovid (or Excavator for short), a system that will ingest

¹We define an event as any occurrence, action, process or state of affairs, following (O’Gorman et al., 2016).

open-source text sources (e.g., news articles and scientific publications), extract COVID-19 related events and relations between them, and build a Temporal and Causal Analysis Graph (TCAG). Excavator combines the following NLP techniques:

- Extracting events (§3) for types in our comprehensive COVID-19 event taxonomy (§2). Each event will have time and location if available in text, allowing analyses targeted at specific times or geographic regions of interest.
- Extracting three types of temporal and causal relations (§4) between pairs of events.
- Constructing a TCAG (§5) by assembling all events and relations, to provide a comprehensive overview of the events related to COVID-19 as well as their causes and impacts.
- Supporting trend and correlation analysis of events, via visualizing event popularity time series (§ 6) in the TCAG visualization.

Excavator produces a TCAG that is in a machine-readable JSON format and is also human-understandable (visualized via a web-based interactive User Interface), to support varied analytical and decision making needs. We hope that Excavator will aid government agencies in efforts to understand likely downstream effects of political and economic decisions and events related to the pandemic, and respond in a timely manner to mitigate the impact of COVID-19. The benefit of Excavator is realized through a comprehensive visualization of events and how they affect each other. We expect the utility of Excavator to outlive the COVID-19 pandemic: analysts and decision makers will be empowered by Excavator to better understand and solve complex problems in the future.

We first present our COVID-19 event taxonomy, and then we present details about event extraction, causal and temporal relation extraction, measuring




Figure 1: A partial illustration of the COVID-19 event taxonomy.

event popularity using news text as “quantitative data”, and the approach for constructing a TCAG. We then describe the system demonstration, present a quantitative analysis of the extractions, and conclude with recommended use cases.

2 Building a COVID-19 Event Taxonomy

COVID-19 affects many aspects of our political, economic, and personal lives. A comprehensive analysis requires an event taxonomy that categorizes the events related to COVID-19 in many sectors and domains. We developed a COVID-19 event taxonomy using a hybrid approach of manual curation with automated support: first, we run Stanza (Qi et al., 2020) on a large sample (10%) of the Aylieen coronavirus news dataset (§ 7) to tag verb and noun phrases that are likely to trigger events. Second, we represent each phrase as the average of the BERT (Devlin et al., 2019) contextualized embedding vectors of the subwords within each phrase, and then run committee-based clustering (Pantel and Lin, 2002) over the vector representations of the phrases to discover salient clusters. Finally, we review the frequently appearing clusters and define event types related to COVID-19.

The event taxonomy includes 76 event types. Each type comes with a name and a short description. Figure 1 illustrates several branches of the event taxonomy². The events come from a wide range of domains. We also manually added hyponymy relations via *is_a* links (e.g., COVID-19 *is_a* Virus) between pairs of event types.

3 Extracting Events

We developed a neural network model for extracting events defined in the COVID-19 event taxonomy (the *event classification* stage) and extracting

²The complete taxonomy is available at https://github.com/BBN-E/LearnIt/blob/master/inputs/domains/CORD_19/ontology/covid_event_ontology.yaml.

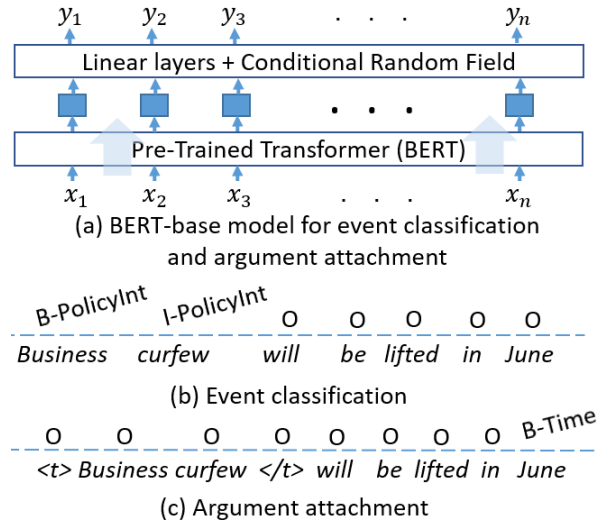


Figure 2: The BERT-based sequence tagging model for event classification and argument extraction. Figure (a) shows the architecture of the model, which takes a sequence of words x_1, x_2, \dots, x_n as input and outputs a sequence of tags y_1, y_2, \dots, y_n . Figure (b) and (c) shows an example for each of the two stages. “PolicyInt” is short for “PolicyIntervention”.

the location and time arguments (the *event argument extraction* stage), if they are mentioned in text, for each event mention. The structured representation (events with location and/or time) enables analyses of events targeting a specific time or location. Both stages use a BERT-based sequence tagging model. Figure 2(a) shows the model architecture. Given a sequence of tokens as input, the model extracts a sequence of tags, one per each token. We use the commonly used Begin-Inside-Outside (BIO) tags for both event types and event argument role types for the event classification and argument attachment tasks respectively.

Event classification: a sequence tagging model is trained to predict BIO tags of event types such that it identifies the event trigger span as well as the event type. Figure 2(b) shows an example.

Event argument extraction: similarly, another sequence tagging model is trained to predict BIO

Type	Counts	Type	Counts
COVID-19	2114	SocialDistancingMeasures	412
Virus	1028	TravelRestrictions	403
Pandemic	596	Disease	378
Unemployment	506	Death	355
Shortage	502	Lockdown	321

Table 1: Top-10 frequent events in the training dataset.

tags of argument role types, such that it identifies token spans of event arguments as well as their argument role types, with respect to a trigger that has already been identified in the *event classification* stage and marked in the input sentence in “< *t* > ... < /*t* >”. Figure 2(c) shows an example.

We run these two models in a pipeline: the event classification model is applied first to find event triggers and classify their types, then the event argument extraction model is applied to find location and time arguments for each event mention.

Training data curation. We use LearnIt rapid customization for event extraction (Chan et al., 2019) to curate a dataset for training the event classification model. Our developer spent about 13 minutes per event type to find, expand, and filter potential event triggers in a held-out 10% of the Aylieen coronavirus news corpus. Statistics of the curated data set are shown in Table 1 (we only show the top-10 most frequent event types for brevity). In total, there are 11814 mentions in 7159 sentences.

To train the argument extraction model, we use the related event-argument annotation from the ACE 2005 dataset (Doddington et al., 2004). We focus on location and time arguments³ and ignore other roles. At decoding time, after extracting the argument mentions for events, we apply the AWAKE (Boschee et al., 2014) entity linking system to resolve each location argument to a canonical geolocation, and use SERIF (Boschee et al., 2005) to resolve each time argument to a canonical time and then convert it to the month level. This allows us to perform analyses of events targeting a specific geolocation or month of interest.

4 Extracting Temporal and Causal Relations

We develop two approaches for extracting temporal and causal relations: a pattern-based approach and a neural network model. We take the union of the

³For example, *Place* and *Time* event argument roles in ACE can be used to train an argument-role model to extract location and time arguments, respectively.

Type	Subtype	Definition
Causes	Cause	Y happens because of X.
	Catalyst	If X, intensity of Y increases.
	Precondition	X must have occurred for Y to happen.
Mitigates	Mitigation	If X, intensity of Y decreases.
	Preventative	If X happens, Y can't happen.
Before	Before/after	X happens before/after Y.

Table 2: Causal and temporal relations between event X and Y.

outputs from both approaches to maximize recall. The list of causal and temporal relations extracted by the systems is shown in Table 2. Our extractors extract relations at the subtype level. However, we decided to merge the subtypes into types because (a) a user survey shows that users prefer to have a simplified definition of causality that only includes “event X causes (positively impacts) event Y” and “X mitigates (reduces/prevents) Y”, because finer-grained distinctions at sub-type level are difficult and less useful, and (b) merging the subtypes into types improves accuracy to near or above 0.8 as shown in Table 4, comparing to slightly below 0.7 at the sub-type level due to extraction approaches struggling to differentiate between the sub-types.

Pattern-based relation extraction. We applied the temporal and causal relation extraction patterns from LearnIt (Min et al., 2020). A pattern is either a lexical pattern, which is a sequence of words between a pair of events, e.g., “X leads to Y”⁴, or a proposition pattern, which is the (nested) predicate-argument structure that connects the pair of events. For example, “verb:cause[subject=X] [object=Y]” is the proposition counterpart of the lexical pattern “X causes Y”.

Neural relation extraction. We developed a mention pooling (Baldini Soares et al., 2019) neural model for causal and temporal relation extraction. Figure 3 shows the model architecture. Taking a sentence in which a pair of event mention spans are marked as input, the model first encodes the sentence with BERT (Devlin et al., 2019)⁵. For each of the left and right event mentions, it then uses average pooling over the BERT contextualized vectors of the words in the span to obtain fixed-dimension vectors V_1 and V_2 as the span representations. It then concatenates the input embeddings V_1 and V_2 with the element-wise difference $|V_1 - V_2|$ to generate the pair representation

⁴X and Y refer to the left and right arguments of a relation.

⁵The BERT-Base model is used.

$V = (V_1, V_2, |V_1 - V_2|)$. V is passed into a linear layer followed by a softmax layer to make the relation prediction. The model is trained with a blended dataset consisting of the Entities, Events, Simple and Complex Cause Assertion Annotation datasets⁶ released by LDC⁷, and 1.5K temporal relation instances generated by applying the LearnIt temporal relation extraction patterns to 10,000 sampled Gigaword (Parker et al., 2011) articles.

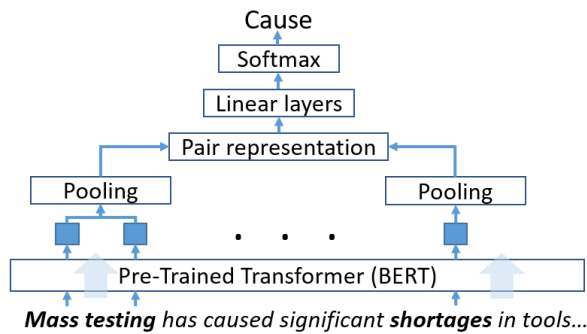


Figure 3: The neural model for causal and temporal relation extraction.

5 Constructing a TCAG

We aggregate all extracted events and causal and temporal relations across the corpus to construct a TCAG. The TCAG is visualized in the interactive visualization, in which each node is an event type and each edge is a causal or temporal relation⁸.

We use a simple approach to aggregate events: by default, all event mentions sharing the same type are grouped into a single node named by the type; we resort to the UI to allow the user to selectively focus on a specific location and/or time, such that the UI will only show a TCAG involving event mentions and causal relations between pairs of events for the location and/or time of interest.

6 Measuring Event Popularity through Time

The TCAG only provides a qualitative analysis of the temporal and causal relations between the COVID-related events. It will be more informative if we can measure the popularity of events through time to enable **trend analysis** (e.g., *does lockdown go up or down between January and May, 2020?*)

⁶The catalog IDs of the LDC datasets are LDC2019E48, LDC2019E61, LDC2019E70, LDC2019E82, LDC2019E83.

⁷www ldc.upenn.edu

⁸*is_a* relations are also added as dashed edges in the TCAG.

and **correlation analysis** (e.g., *will a stricter lockdown improve or deteriorate the economy?*).

In order to support these analyses, we produce a timeseries of a popularity score for each event type over time (a.k.a., event timeline). Extending our prior work (Min and Zhao, 2019), we define the popularity score for event type e at time t as:

$$Popularity(e)_t = \frac{1}{T} \sum_{t' \in [t-\frac{T}{2}, t+\frac{T}{2}]} \frac{N_{e,t'}}{cM_{t'}}$$

in which $N_{e,t}$ is the frequency of event e at month t . We calculate the moving average centered at each t with a sliding window of $T = 3$ months to reduce noise. M_t is the total number of articles published in month t . $c = 1/500$ is a normalizing constant. The raw event frequency counts can be inflated due to the increasing level of media activity. Therefore, we divide the raw counts by cM_t to normalize the counts so that they are comparable across different months.

7 System Demonstration

Datasets. We run Excavator on the following two corpora to produce a TCAG for COVID-19: the first corpus is 1.2 million articles⁹ from the Aylieen Coronavirus News Dataset¹⁰, which contains 1.6 million COVID-related articles published between November 2019 and July 2020 that are from ~440 news sources. We only kept the articles that are published between January and May 2020, since the corpus contains fewer articles in other months. The second corpus is the COVID-19 Open Research Dataset (Wang et al., 2020). It contains coronavirus-related research from PubMed’s PMC corpus, a corpus maintained by the WHO, and bioRxiv and medRxiv pre-prints. As of 11/08/2020, it contains over 300,000 scholarly articles.

We combine these two corpora because news and research articles are complementary: news are rich in real-world events and are up to date, while analytical articles contain more causal relationships. Therefore, combining them is likely to lead to a more comprehensive analysis and new insights.

Overall statistics of extractions. Excavator extracted 6.2 million event mentions of 59 types. Table 3 shows the event types that appear more than 50,000 times. We randomly sampled 100 event

⁹These articles do not overlap with the held-out set for training data curation.

¹⁰<https://aylien.com/blog/free-coronavirus-news-dataset>

Type	Counts	Type	Counts
COVID-19	2772.3	Travel	111.8
Death	730.0	FearOrPanic	94.6
Pandemic	689.2	Closures	92.3
Lockdown	417.2	TravelRestrictions	76.9
Isolation*	195.4	Shortage	68.3
DiseaseSpread	145.4	Conflict	55.5
Testing	130.7	Virus	54.8
Treatment	112.8	Symptom	54.0

Table 3: Frequent events extracted from the corpora (ranked by frequency reversely; numbers are in thousands). *Isolation refers to IsolationOrConfinement.

Type	Count	Precision
Causes	193,694	0.78
Mitigates	30,452	0.87
Before	2,030	0.81

Table 4: Causal and temporal relations extracted.

mentions, manually reviewed them, and found that the extracted events are 83% accurate. Excavator extracted 226,176 causal and temporal relations from the two corpora. A summary of the extracted relations and their precision¹¹ are shown in Table 4.

TCAG Visualization We developed an interactive visualization of the TCAG. Figure 4 shows a small part of the TCAG centered on the event Lockdown. Each node represents an event type in our COVID event taxonomy for which Excavator is able to extract events and track their popularity scores (§ 6) through time. The three types of relational edges (Causes, Mitigates and Before) are shown in different colors. The size of the nodes and the thickness of the edges indicate the relative frequency of the event types or relations in the log scale, respectively. For example, Figure 4 shows that Death is mentioned more frequently than Lockdown, and the causal relation {Lockdown, Causes, EconomicCrisis} appears more frequently than {Lockdown, Mitigates (“reduces”), AccessToHealthcare}. To support analysis focusing on a single event, we color the focused event in blue, events that cause or precede the focused event in orange, and events that the focused event causes or precedes in green.

Event popularity timeseries visualization For each node (event) in the TCAG visualization, we show its event popularity timeseries visualization on the side. Figure 5 shows 3 screenshots of the

¹¹Estimated by manually reviewing 40 instances per type

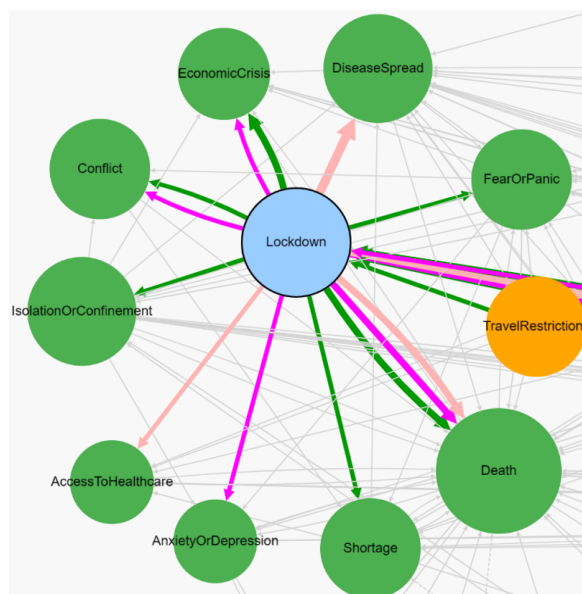


Figure 4: A screenshot of a partial TCAG centered on Lockdown. Green, pink, and purple edges shows Cause, Mitigate and Before relations, respectively. Blue, orange and green nodes show the focused node and nodes with incoming and outgoing edges (with respect to the focused node), respectively.

event popularity timeseries (§ 6) visualization between January and May 2020 for Lockdown, EconomicCrisis and COVID-19 respectively.

8 Recommended Use Cases

We describe 3 recommended use cases below. More details are in our demonstration video.

Use case 1: causal and temporal analysis. We can get a panoramic view of the underlying causal and temporal dynamics between events related to COVID from the overall TCAG. We can start by analyzing the causal or temporal relations centered at an event of interest. For example, Figure 4 shows a diverse range of effects and consequences of Lockdown, such as EconomicCrisis (economic), Shortage (supply-chain), FearOrPanic (mental), etc. Interestingly, the graph also reveals surprises such as {Lockdown, Causes, Death}: the UI shows supporting evidence such as “lockdown exacerbates deaths and chronic health problems associated with poverty, ...”. Furthermore, the TCAG shows that Lockdown mitigates DiseaseSpread but it also has a negative impact on the Economy, which will inform the decision makers that they will need to understand the economic trade-offs when implementing the Lockdown policy.

We can also analyze longer-distance causal path-




Figure 5: Event popularity timeseries (timeline) between 01-05/2020 for Lockdown, EconomicCrisis and COVID-19. X-axis shows months between January and May 2020. Y-axis shows event popularity scores.

ways consisting of two or more causal/temporal edges. For example, our demo video shows that COVID-19 causes or precedes (Before) Lockdown, and that Lockdown causes or precedes EconomicCrisis. This helps us understand details about how COVID causes EconomicCrisis.

Use case 2: trend and correlation analysis. We can inspect the event timeline for a node or an edge to perform a trend analysis and a correlation analysis, respectively. Figure 5 shows screenshots of the event popularity timeseries between January and May 2020 for Lockdown, EconomicCrisis and COVID-19. First, the user can click on a single event to perform a trend analysis: the popularity of Lockdown goes up continuously, indicating an upward trend in implementing lockdown policies in more geographic regions. The user can also click on an edge to perform a correlation analysis for a pair of events: when the user clicks on the edge {Lockdown, Causes, EconomicCrisis}, the UI shows a strong correlation between the two upward curves. For another edge “Lockdown mitigates COVID-19”, the UI shows a negative correlation near the end: as Lockdown rises, COVID-19 slightly falls towards the end.

Use case 3: analyses targeted at geolocations. The event timeline visualization also allows the user to see the timeline for geolocations such as each U.S. state individually, instead of the aggregate for the entire U.S.. Figure 6 is a screenshot showing the 10 timelines for Lockdown for the top-10 most frequently mentioned U.S. states. The screenshot shows that the curves for California and New York go much higher than other states. This roughly matches the stricter lockdown policies implemented in the two states during this time period, compared with other states. Such targeted analysis

is made possible because our events have *location* and *time* arguments. We can also make the TCAG only show events and relations for a specific state, if a user selects a state of interest in the UI.




Figure 6: Event popularity timeseries in 01-05/2020 for Lockdown, for top-10 frequently mentioned US states.

9 Related Work

Extracting events. Event extraction has been studied using feature-based approaches (Huang and Riloff, 2012; Ji and Grishman, 2008), or neural networks (Chen et al., 2015; Nguyen et al., 2016a; Wadden et al., 2019; Liu et al., 2020). GDELT (Lee-taru and Schrodt, 2013) creates an event database for the conflict and mediation domain. It has very few event types related to COVID-19. To adapt event extraction to new domains, Chen et al. (2019) developed a user-in-the-loop rapid event customization system. Nguyen et al. (2016b) proposed a neural model for event type extension given seed examples. Peng et al. (2016) developed a minimally supervised approach using triggers gathered from ACE annotation guideline.

Extracting causal and temporal relations. There are a lot of work in temporal (D’Souza

and Ng, 2013; Chambers et al., 2014; Ning et al., 2018b; Meng and Rumshisky, 2018; Han et al., 2019; Vashishtha et al., 2020; Wright-Bettner et al., 2020) and causal (Bethard and Martin, 2008; Do et al., 2011; Riaz and Girju, 2013; Roemmele and Gordon, 2018; Hashimoto, 2019) relation extraction. Mirza and Tonelli (2016) and Ning et al. (2018a) extract both in a single framework.

Constructing Causal Graphs from Text. Eidos (Sharp et al., 2019) uses a rule-based approach to extract causal relations to build a causal analysis graph, that has limited coverage on events related to COVID-19. LearnIt (Min et al., 2020) enables rapid customization of causal relation extractors. LearnIt does not focus on causal relations involving COVID-related events. This work also differs from these two in that we extract event arguments and temporal relations, and track event popularity.

10 Conclusion and Future Work

We present Excavator, a machine reading system that automatically constructs a Temporal and Causal Analysis Graph for COVID-19 by reading open-source text documents such as news and scientific publications. Our next steps are to integrate Modal Dependency Parsing (Yao et al., 2021) for event factuality assessment, and cross-lingual transfer learning (Nguyen et al., 2021) to make Excavator applicable to more languages.

Acknowledgements

This research is based upon work supported in part by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via Contract No.: 2021-20102700002. 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 ODNI, IARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes not withstanding any copyright annotation therein.

References

Livio Baldini Soares, Nicholas FitzGerald, Jeffrey Ling, and Tom Kwiatkowski. 2019. [Matching the blanks: Distributional similarity for relation learning](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*,

pages 2895–2905, Florence, Italy. Association for Computational Linguistics.

Steven Bethard and James H. Martin. 2008. [Learning semantic links from a corpus of parallel temporal and causal relations](#). In *Proceedings of ACL-08: HLT, Short Papers*, pages 177–180, Columbus, Ohio. Association for Computational Linguistics.

Elizabeth Boschee, Marjorie Freedman, Saurabh Khanwalkar, Anoop Kumar, Amit Srivastava, and Ralph Weischedel. 2014. [Researching persons & organizations: Awake: From text to an entity-centric knowledge base](#). In *2014 IEEE International Conference on Big Data (Big Data)*, pages 1030–1039. IEEE.

Elizabeth Boschee, Ralph Weischedel, and Alex Zamanian. 2005. [Automatic information extraction](#). In *Proceedings of the International Conference on Intelligence Analysis*, volume 71. Citeseer.

Nathanael Chambers, Taylor Cassidy, Bill McDowell, and Steven Bethard. 2014. [Dense event ordering with a multi-pass architecture](#). *Transactions of the Association for Computational Linguistics*, 2:273–284.

Yee Seng Chan, Joshua Fasching, Haoling Qiu, and Bonan Min. 2019. [Rapid customization for event extraction](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 31–36, Florence, Italy. Association for Computational Linguistics.

Yubo Chen, Liheng Xu, Kang Liu, Daojian Zeng, and Jun Zhao. 2015. [Event extraction via dynamic multi-pooling convolutional neural networks](#). In *ACL-IJCNLP2-2015*, pages 167–176, Beijing, China. Association for Computational Linguistics.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.

Quang Do, Yee Seng Chan, and Dan Roth. 2011. [Minimally supervised event causality identification](#). In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 294–303.

George R Doddington, Alexis Mitchell, Mark A Przybicki, Lance A Ramshaw, Stephanie M Strassel, and Ralph M Weischedel. 2004. [The automatic content extraction \(ace\) program-tasks, data, and evaluation](#). In *Lrec*, volume 2, pages 837–840. Lisbon.

Jennifer D’Souza and Vincent Ng. 2013. [Classifying temporal relations with rich linguistic knowledge](#). In *Proceedings of the 2013 Conference of the North*

- American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 918–927.
- Rujun Han, I-Hung Hsu, Mu Yang, Aram Galstyan, Ralph Weischedel, and Nanyun Peng. 2019. [Deep structured neural network for event temporal relation extraction](#). In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 666–106, Hong Kong, China. Association for Computational Linguistics.
- Chikara Hashimoto. 2019. [Weakly supervised multilingual causality extraction from Wikipedia](#). 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 2988–2999, Hong Kong, China. Association for Computational Linguistics.
- Ruihong Huang and Ellen Riloff. 2012. [Modeling textual cohesion for event extraction](#). In *AAAI-CAI, AAAI’12*, pages 1664–1670. AAAI Press.
- Heng Ji and Ralph Grishman. 2008. [Refining event extraction through cross-document inference](#). In *ACL-HLT-2008*, pages 254–262, Columbus, Ohio. Association for Computational Linguistics.
- Kalev Leetaru and Philip A Schrod. 2013. [Gdelt: Global data on events, location, and tone, 1979–2012](#). In *ISA annual convention*, volume 2, pages 1–49. Citeseer.
- Jian Liu, Yubo Chen, Kang Liu, Wei Bi, and Xiaojiang Liu. 2020. [Event extraction as machine reading comprehension](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1641–1651, Online. Association for Computational Linguistics.
- Yuanliang Meng and Anna Rumshisky. 2018. [Context-aware neural model for temporal information extraction](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 527–536, Melbourne, Australia. Association for Computational Linguistics.
- Bonan Min, Manaj Srivastava, Haoling Qiu, Prasanakumar Muthukumar, and Joshua Fasching. 2020. [Learnit: On-demand rapid customization for event-event relation extraction](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 13630–13631.
- Bonan Min and Xiaoxi Zhao. 2019. [Measure country-level socio-economic indicators with streaming news: An empirical study](#). 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 1249–1254, Hong Kong, China. Association for Computational Linguistics.
- Paramita Mirza and Sara Tonelli. 2016. [Catena: Causal and temporal relation extraction from natural language texts](#). In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 64–75.
- Minh Van Nguyen, Tuan Ngo Nguyen, Bonan Min, and Thien Huu Nguyen. 2021. [Crosslingual transfer learning for relation and event extraction via word category and class alignments](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Thien Huu Nguyen, Kyunghyun Cho, and Ralph Grishman. 2016a. [Joint event extraction via recurrent neural networks](#). In *NAACL-HLT-2016*, pages 300–309, San Diego, California. Association for Computational Linguistics.
- Thien Huu Nguyen, Lisheng Fu, Kyunghyun Cho, and Ralph Grishman. 2016b. [A two-stage approach for extending event detection to new types via neural networks](#). In *WRepLANLP*, pages 158–165, Berlin, Germany. Association for Computational Linguistics.
- Qiang Ning, Zhili Feng, Hao Wu, and Dan Roth. 2018a. [Joint reasoning for temporal and causal relations](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2278–2288, Melbourne, Australia. Association for Computational Linguistics.
- Qiang Ning, Hao Wu, Haoruo Peng, and Dan Roth. 2018b. [Improving temporal relation extraction with a globally acquired statistical resource](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 841–851, New Orleans, Louisiana. Association for Computational Linguistics.
- Tim O’Gorman, Kristin Wright-Bettner, and Martha Palmer. 2016. [Richer event description: Integrating event coreference with temporal, causal and bridging annotation](#). In *Proceedings of the 2nd Workshop on Computing News Storylines (CNS 2016)*, pages 47–56, Austin, Texas. Association for Computational Linguistics.
- Patrick Pantel and Dekang Lin. 2002. [Document clustering with committees](#). In *Proceedings of the 25th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 199–206.
- Robert Parker, David Graff, Junbo Kong, Ke Chen, and Kazuaki Maeda. 2011. [English gigaword fifth edition, 2011](#). *Linguistic Data Consortium, Philadelphia, PA, USA*.
- Haoruo Peng, Yangiu Song, and Dan Roth. 2016. [Event detection and co-reference with minimal supervision](#). In *EMNLP-2016*.

- Peng Qi, Yuhao Zhang, Yuhui Zhang, Jason Bolton, and Christopher D. Manning. 2020. [Stanza: A Python natural language processing toolkit for many human languages](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*.
- Mehwish Riaz and Roxana Girju. 2013. [Toward a better understanding of causality between verbal events: Extraction and analysis of the causal power of verb-verb associations](#). In *Proceedings of the SIGDIAL 2013 Conference*, pages 21–30, Metz, France. Association for Computational Linguistics.
- Melissa Roemmele and Andrew Gordon. 2018. [An encoder-decoder approach to predicting causal relations in stories](#). In *Proceedings of the First Workshop on Storytelling*, pages 50–59, New Orleans, Louisiana. Association for Computational Linguistics.
- Rebecca Sharp, Adarsh Pyarelal, Benjamin Gyori, Keith Alcock, Egoitz Laparra, Marco A. Valenzuela-Escárcega, Ajay Nagesh, Vikas Yadav, John Bachman, Zheng Tang, Heather Lent, Fan Luo, Mithun Paul, Steven Bethard, Kobus Barnard, Clayton Morrison, and Mihai Surdeanu. 2019. [Eidos, INDRA, & delphi: From free text to executable causal models](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 42–47, Minneapolis, Minnesota. Association for Computational Linguistics.
- Siddharth Vashishtha, Adam Poliak, Yash Kumar Lal, Benjamin Van Durme, and Aaron Steven White. 2020. [Temporal reasoning in natural language inference](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4070–4078, Online. Association for Computational Linguistics.
- David Wadden, Ulme Wennberg, Yi Luan, and Hananeh Hajishirzi. 2019. [Entity, relation, and event extraction with contextualized span representations](#). 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 5784–5789, Hong Kong, China. Association for Computational Linguistics.
- Lucy Lu Wang, Kyle Lo, Yoganand Chandrasekhar, Russell Reas, Jiangjiang Yang, Doug Burdick, Darrin Eide, Kathryn Funk, Yannis Katsis, Rodney Michael Kinney, Yunyao Li, Ziyang Liu, William Merrill, Paul Mooney, Dewey A. Murdick, Devvret Rishi, Jerry Sheehan, Zhihong Shen, Brandon Stilson, Alex D. Wade, Kuansan Wang, Nancy Xin Ru Wang, Christopher Wilhelm, Boya Xie, Douglas M. Raymond, Daniel S. Weld, Oren Etzioni, and Sebastian Kohlmeier. 2020. [CORD-19: The COVID-19 open research dataset](#). In *Proceedings of the 1st Workshop on NLP for COVID-19 at ACL 2020*, Online. Association for Computational Linguistics.
- Kristin Wright-Bettner, Chen Lin, Timothy Miller, Steven Bethard, Dmitriy Dligach, Martha Palmer, James H. Martin, and Guergana Savova. 2020. [Defining and learning refined temporal relations in the clinical narrative](#). In *Proceedings of the 11th International Workshop on Health Text Mining and Information Analysis*, pages 104–114, Online. Association for Computational Linguistics.
- Jiarui Yao, Haoling Qiu, Jin Zhao, Bonan Min, and Nianwen Xue. 2021. [Factuality assessment as modal dependency parsing](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1540–1550, Online. Association for Computational Linguistics.