EventPlus

A Temporal Event Understanding Pipeline

Mingyu Derek Ma*, Jiao Sun*, Mu Yang, Kung-Hsiang Huang,
Nuan Wen, Shikhar Singh, Rujun Han, Nanyun Peng

    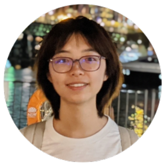   

Try our demo at:
kairos-event.isi.edu

Code available at:
github.com/pluslabnlp/eventplus

Event Understanding

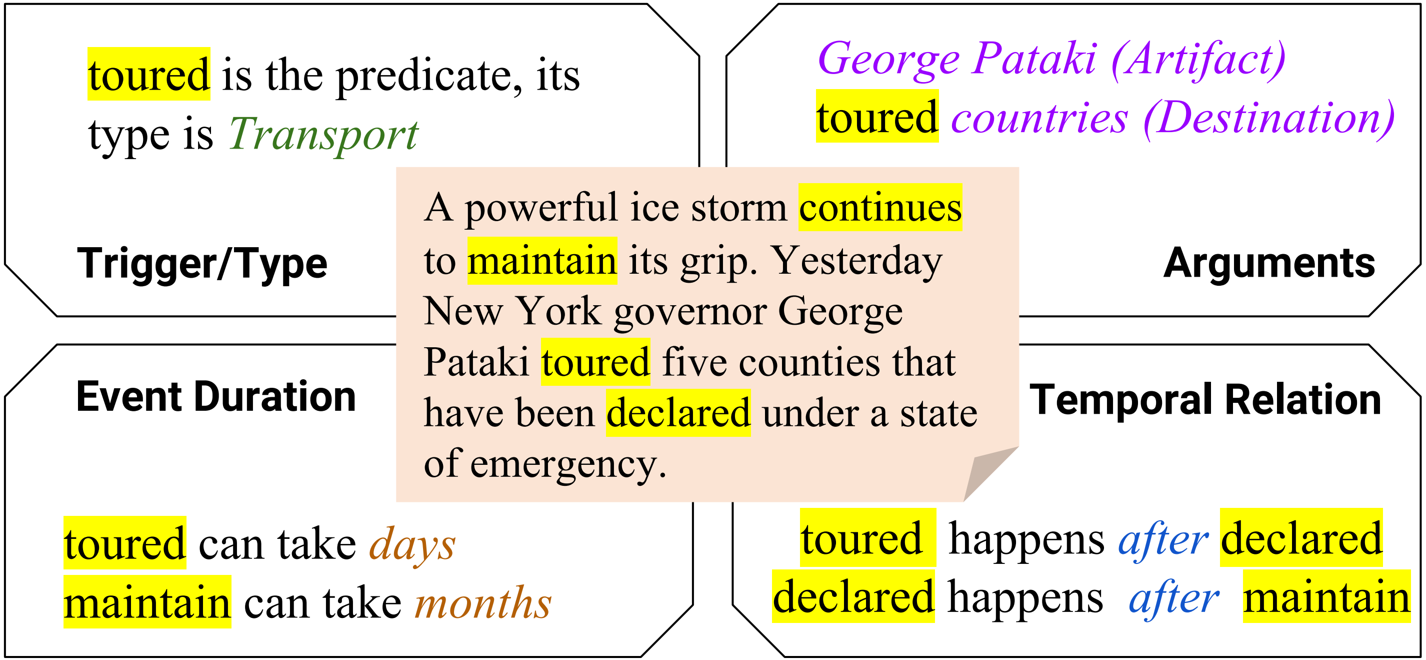

Token-level Tasks

- Tokenization
- Lemmatization
- POS tagging

Sentence-level Tasks

- Syntactic parsing
- Semantic role labeling

Semantic-level Tasks

- For example, events and temporal relations
- Separate works for event extraction, temporal relation and duration extraction, not comprehensive and coherent

Tao et al., 2013; Ning, 2019; Manning et al., 2014; Khashabi et al., 2018; Peng et al., 2015; Noji and Miyao, 2016

EventPlus

- The first event pipeline system providing comprehensive event understanding capabilities
- Each component in EventPlus has comparable performance to the state-of-the-art

Components

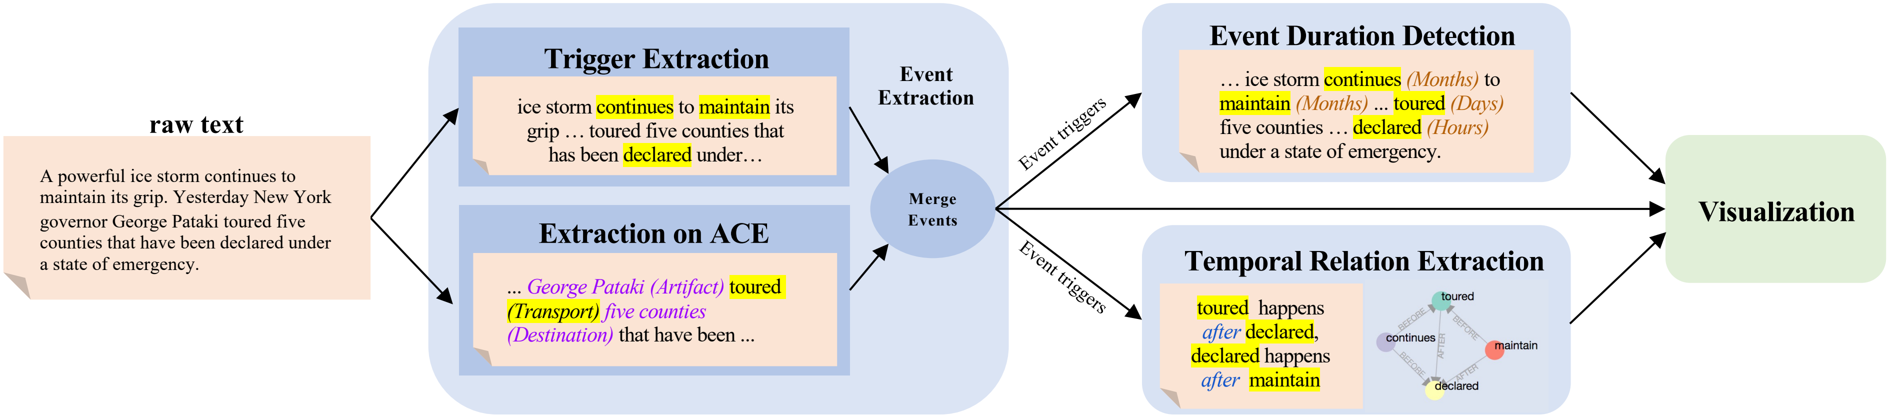

Event definition

- Trigger word
  - Broader coverage
- Complex structure
  - Richer information

Components

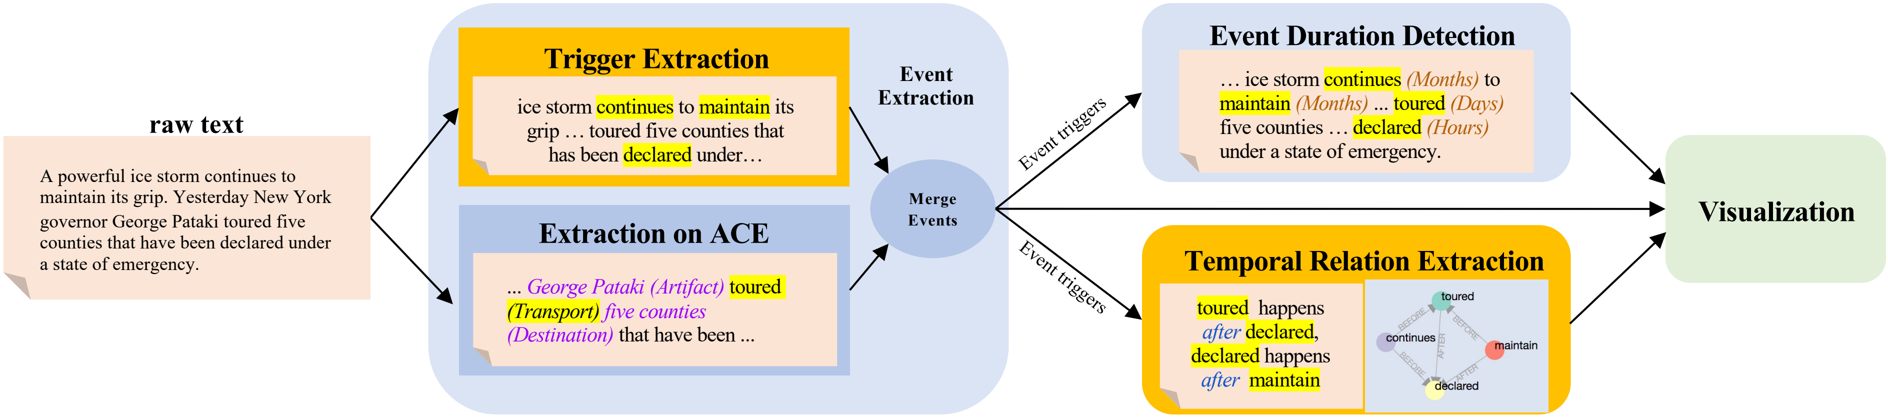

Component 1: Multi-task Learning of Event Trigger and Temporal Relation Extraction

- Event trigger extraction: raw text -> list of words
- Temporal relation extraction: list of words -> relations among them
- Intuition: event relation signals can help to distinguish event/non-event tokens
- BERT embedding + BiLSTM layer

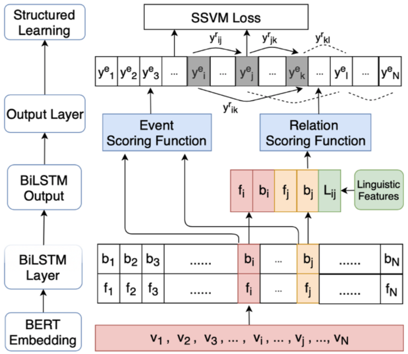

Han et al., EMNLP'19

Components

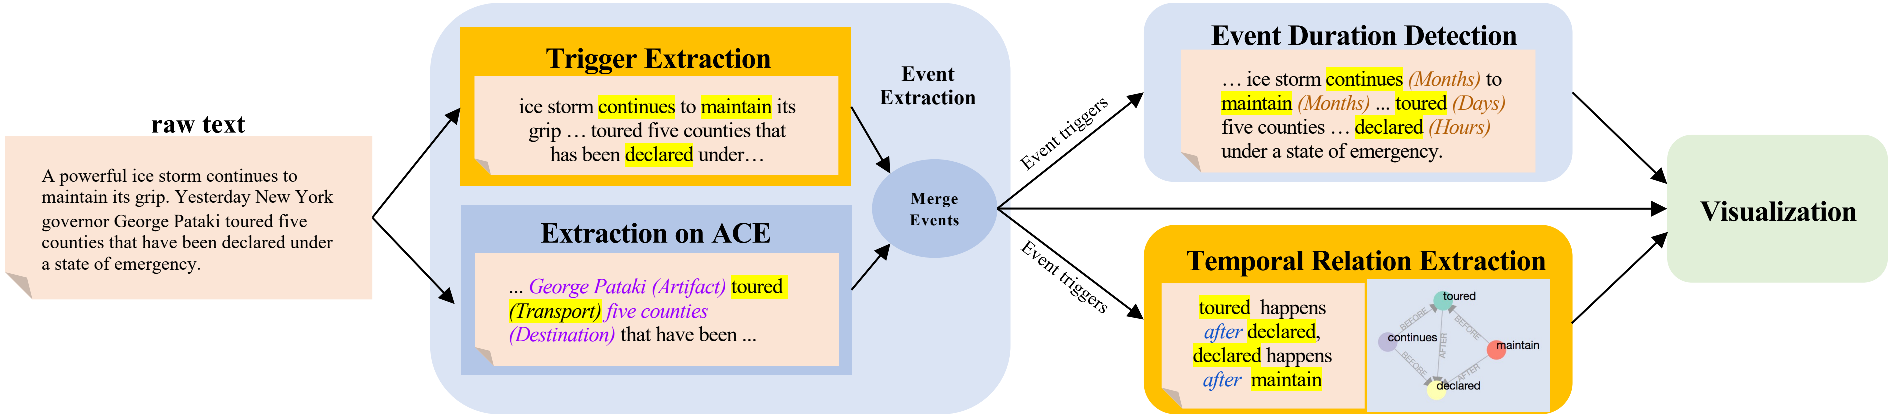

Component 1: Multi-task Learning of Event Trigger and Temporal Relation Extraction

Corpus	Model	F1
TB-Dense	Chambers et al. (2014)	87.4
	Han et al. (2020a)	90.3
	Ours	**90.8**
MATRES	Ning et al. (2018b)	85.2
	Ours	**87.8**

Performance of event trigger extraction

Corpus	Model	F1
TB-Dense	Vashishtha et al. (2019)	56.6
	Meng and Rumshisky (2018)	57.0
	Ours	**64.5**
MATRES	Ning et al. (2018b)	65.9
	Ning et al. (2018a)	69.0
	Ours	**75.5**

Performance of temporal relation extraction

Components

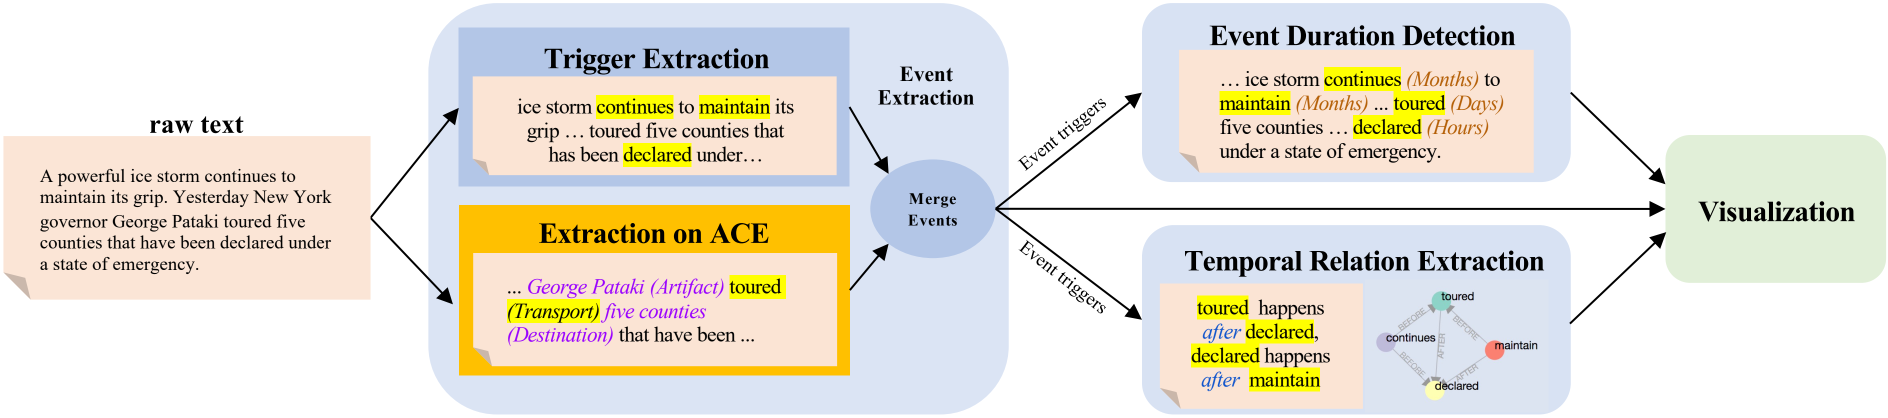

Component 2: Event Extraction on ACE Ontology

- Raw text -> semantic rich information of events
- Trained with ACE2005 corpus
- Multi-task learning of trigger detection, argument role detection and entity detection with shared BERT encoder
- Inference constraints
  - Entity-Argument constraint
  - Entity-Trigger constraint
  - Valid Trigger-Argument constraint

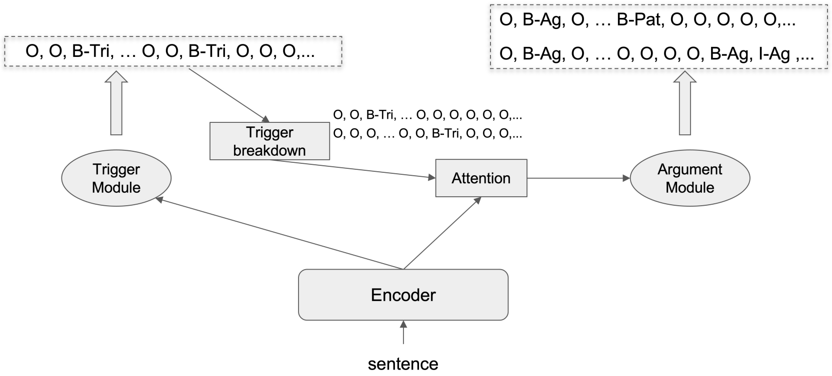

Components

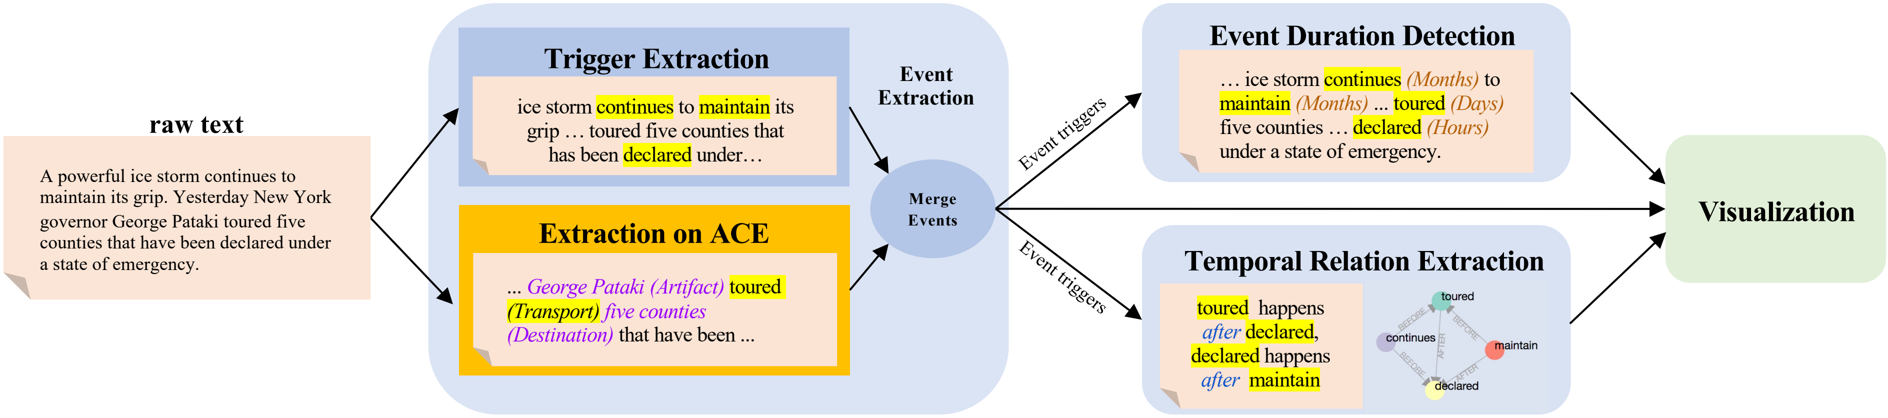

Component 2: Event Extraction on ACE Ontology

	NER	Trigger ID	Trigger CLS	Arg ID	Arg CLS
OneIE (Lin et al., 2020) w/ predicted entities	90.2	**78.2**	**74.7**	**59.2**	**56.8**
Ours w/ predicted entities	**91.3**	75.77	72.45	57.74	55.65

Components

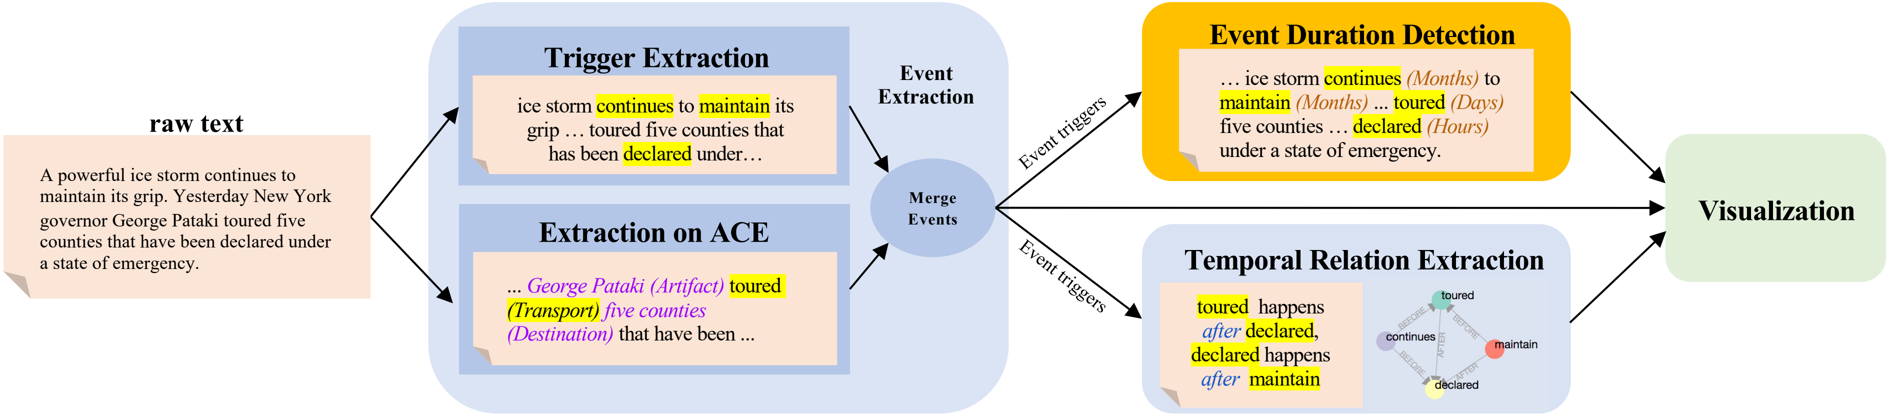

Component 3: Event Duration Detection

- Classifies event triggers into duration categories
- Model 1: Fine-tuned BERT
- Model 2: ELMo embeddings followed by attention layers to get attended representation of the event predicate (Vashishtha et al., 2019)

- Dataset 1: UDS-T (Vashishtha et al., 2019) over 11 duration categories (`instant, seconds, minutes, hours, days, weeks, months, years, decades, centuries`)
- Dataset 2: Typical-Duration (Pan et al., 2006) over 7 duration categories (`seconds, minutes, hours, days, weeks, months, years`)
- Dataset 3: ACE-Duration over 7 duration categories (`seconds, minutes, hours, days, weeks, months, years`)

Components

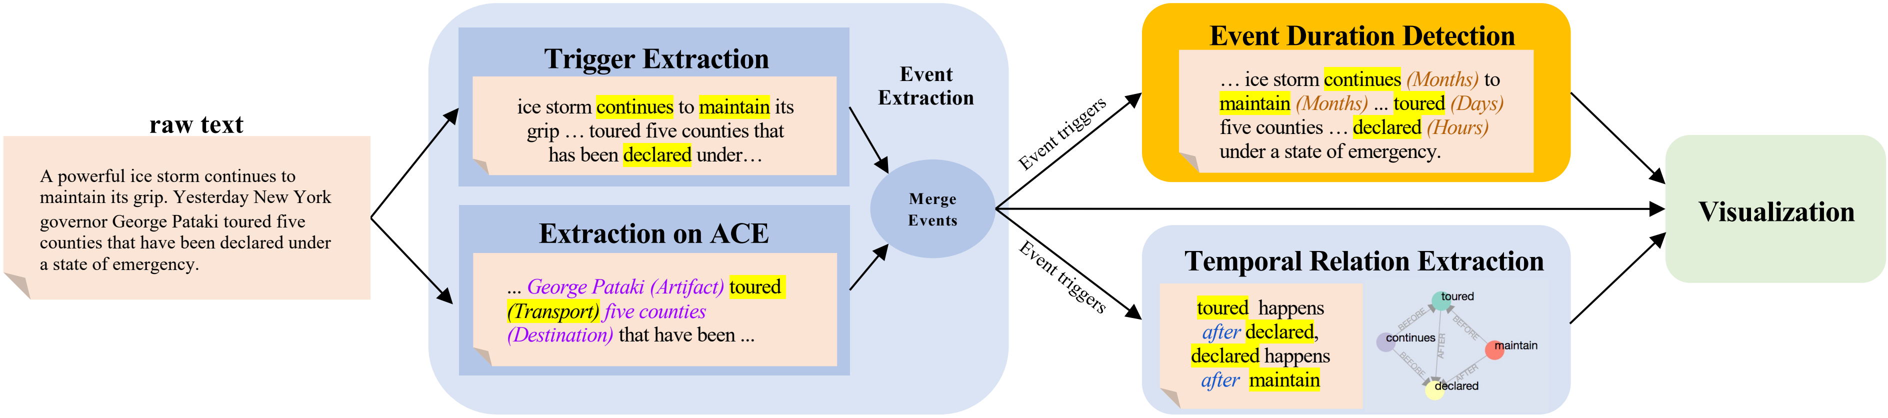

Component 3: Event Duration Detection

Model	Typical-Duration Acc	Typical-Duration Acc-c	Typical-Duration Corr	ACE-Duration Acc	ACE-Duration Acc-c	ACE-Duration Corr
UDS-T (U)	0.20	0.54	0.59	0.38	0.68	0.62
UDS-T (T)	0.52	0.79	0.71	0.47	0.67	0.50
UDS-T (T+U)	0.50	0.76	0.68	**0.49**	0.74	0.66
BERT (T)	**0.59**	**0.81**	**0.75**	0.31	0.67	0.64
BERT (T+U)	0.56	**0.81**	0.73	0.45	**0.79**	**0.70**

- Dataset 1: UDS-T (Vashishtha et al., 2019) over 11 duration categories (`instant, seconds, minutes, hours, days, weeks, months, years, decades, centuries`)
- Dataset 2: Typical-Duration (Pan et al., 2006) over 7 duration categories (`seconds, minutes, hours, days, weeks, months, years`)
- Dataset 3: ACE-Duration over 7 duration categories (`seconds, minutes, hours, days, weeks, months, years`)

Components

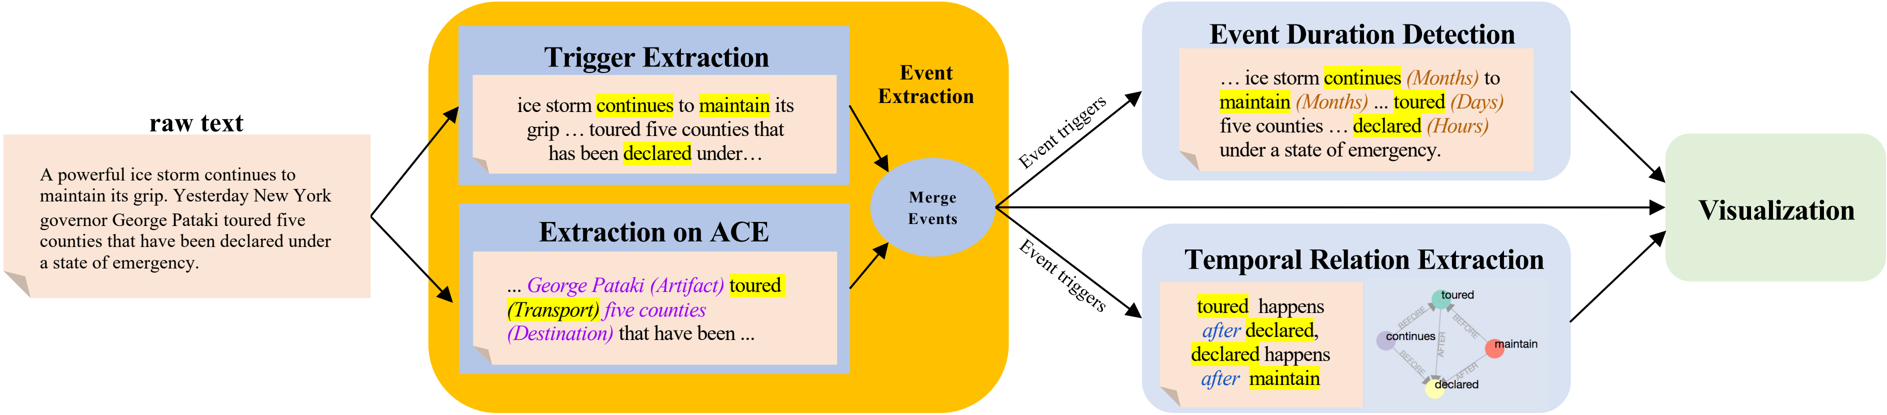

Component 4: Negation and Speculation Cue Detection and Scope Resolution

*The United States is **not** considering sending troops to Mozambique*
*The United States **might** send troops to Mozambique*

- Identifies speculation and negation events
- Step 1: cue detection
- Step 2: scope resolution
- BERT-based cue detection model and BERT-based scope resolution model (Khandelwal and Sawant, 2020)
- Train and test on SFU Review dataset with negation and speculation annotations
- 0.92 F1 score for cue detection, 0,88 F1 score for token-level scope resolution

Pipeline Design

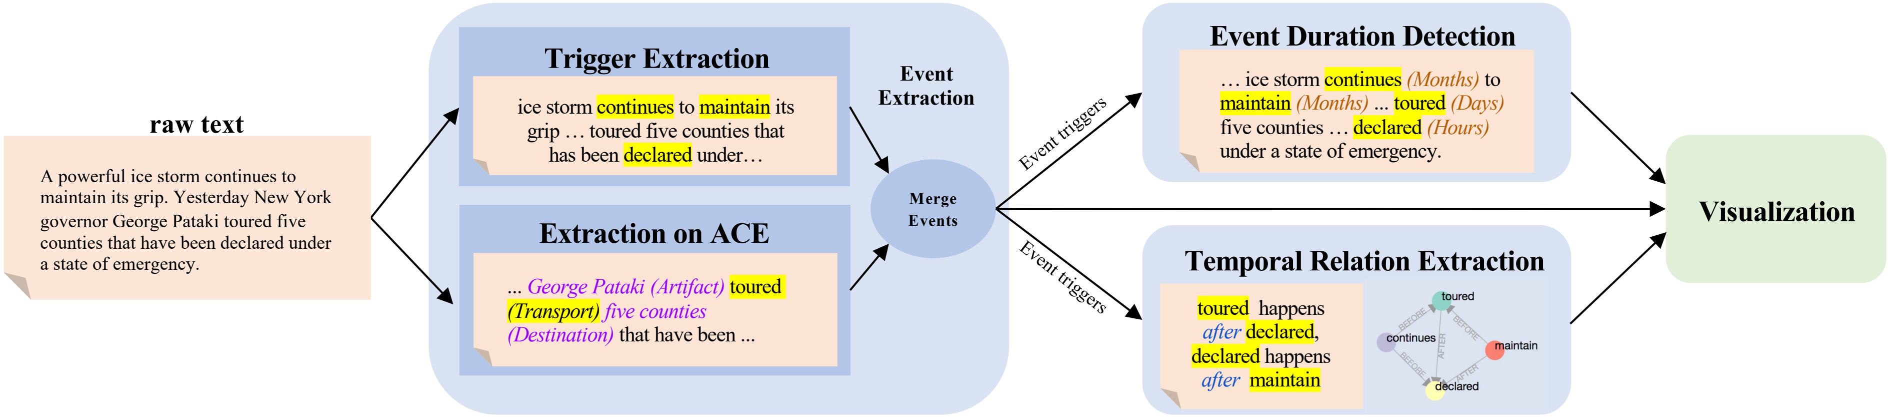

Interface Design

https://kairos-event.isi.edu

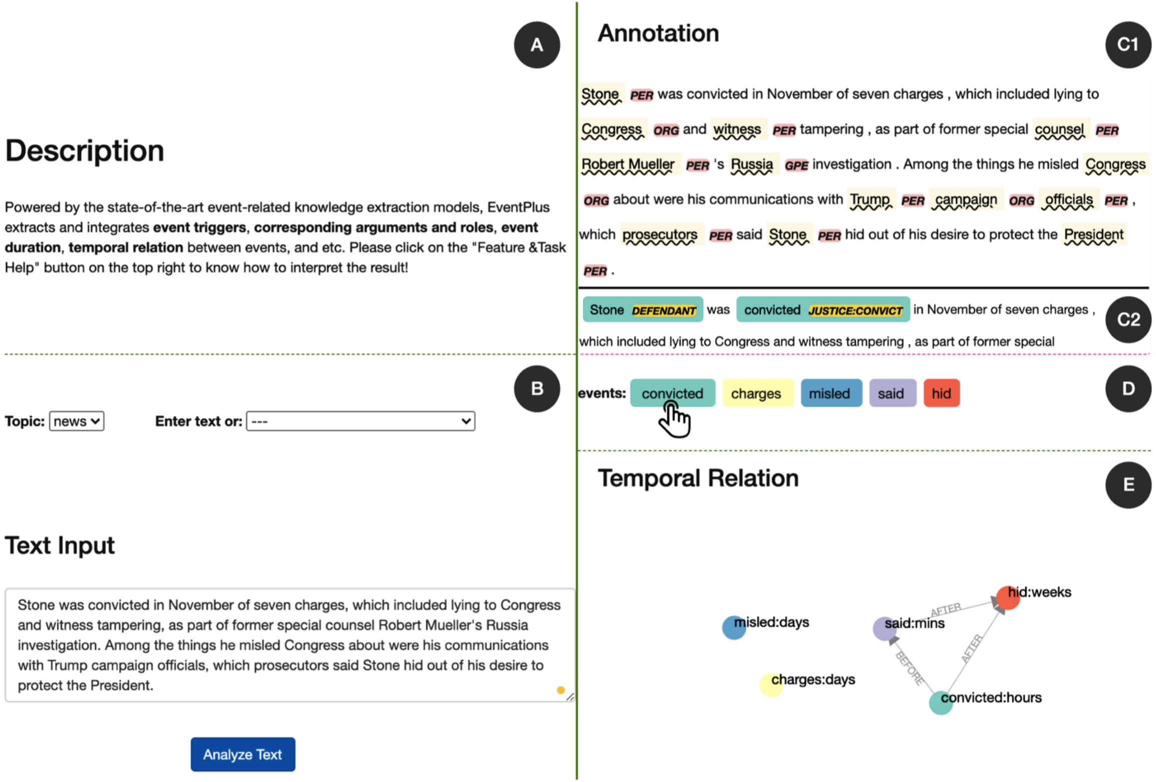

Extension to Biomedical Domain

Each component of EventPlus can be easily extended to other domains

Two approaches to extend event extraction to biomedical domain

- Multi-domain training (MDT) with GENIA
- Replace current component with an in-domain event extraction component: SciBERT-FT (Huang et al., 2020)

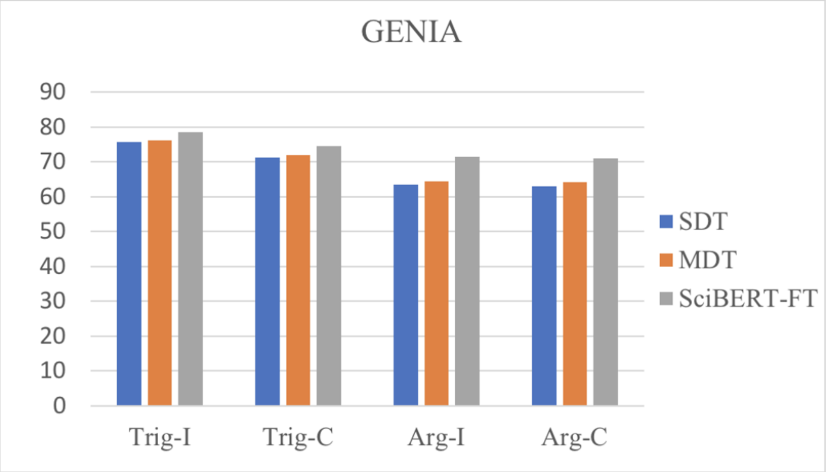

System Demo

Thanks!

EventPlus

A Temporal Event Understanding Pipeline

Mingyu Derek Ma*, Jiao Sun*, Mu Yang, Kung-Hsiang Huang,
Nuan Wen, Shikhar Singh, Rujun Han, Nanyun Peng