



---

## Event Summarization

Kathleen McKeown

Department of Computer Science

Columbia University

**Aug 2021**

**ACL Tutorials**

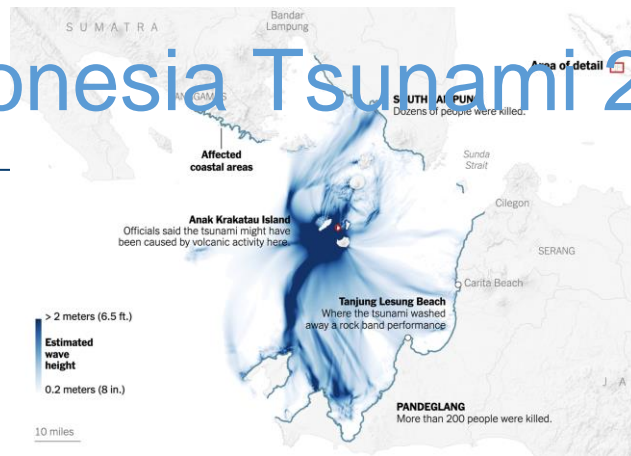
**Event-Centric Natural Language Processing**

# A Different Kind of Event



- Large-scale events
  - Composed of many fine-grained events
  - Also called *complex events*
- Elections, disasters, disease outbreaks, economic impacts

# Indonesia Tsunami 2018





## FEATURE: Five years after tsunami, Indonesia better prepared for catastrophes



UN Photo/E. Schneider | Aerial view of the vast destruction of the Indonesian coast caused by the Indian Ocean tsunami in 2004.

# What Might be Useful?

---



- Track an event
- Provide updates
- Timeline summarization
- Event summarization

# What Might be Useful?

---



- Track an event
- Provide updates
- Timeline summarization
- Event summarization

# Tracking an Event in the News

---



- Did a new major event start?
- Do news articles continue to refer to the event?
- Has the event forked into two sub-events?
- Typically done at the article level using some form of clustering

## Dataset - Miranda et al (2018)

### Multilingual news stream clustering dataset

- English portion – 21k articles from 815 event topics
- Articles annotated with language, timestamp and event topic cluster
- Contains training set for tuning clustering parameters
- Training and test sets are temporally disjoint
- SOTA results reported for comparative evaluation

Statistic	Dataset Split	
	Training	Evaluation
# Articles	12,233	8,726
# Tokens / Article	434 ( $\pm$ 364)	521 ( $\pm$ 495)
# Events	593	222
# Articles / Event	21 $\pm$ 32	39 $\pm$ 88
Earliest Article Date	Dec 18, 2013	Feb 2, 2014
Latest Article Date	Nov 2, 2014	Aug 25, 2015



- Batch-clustering

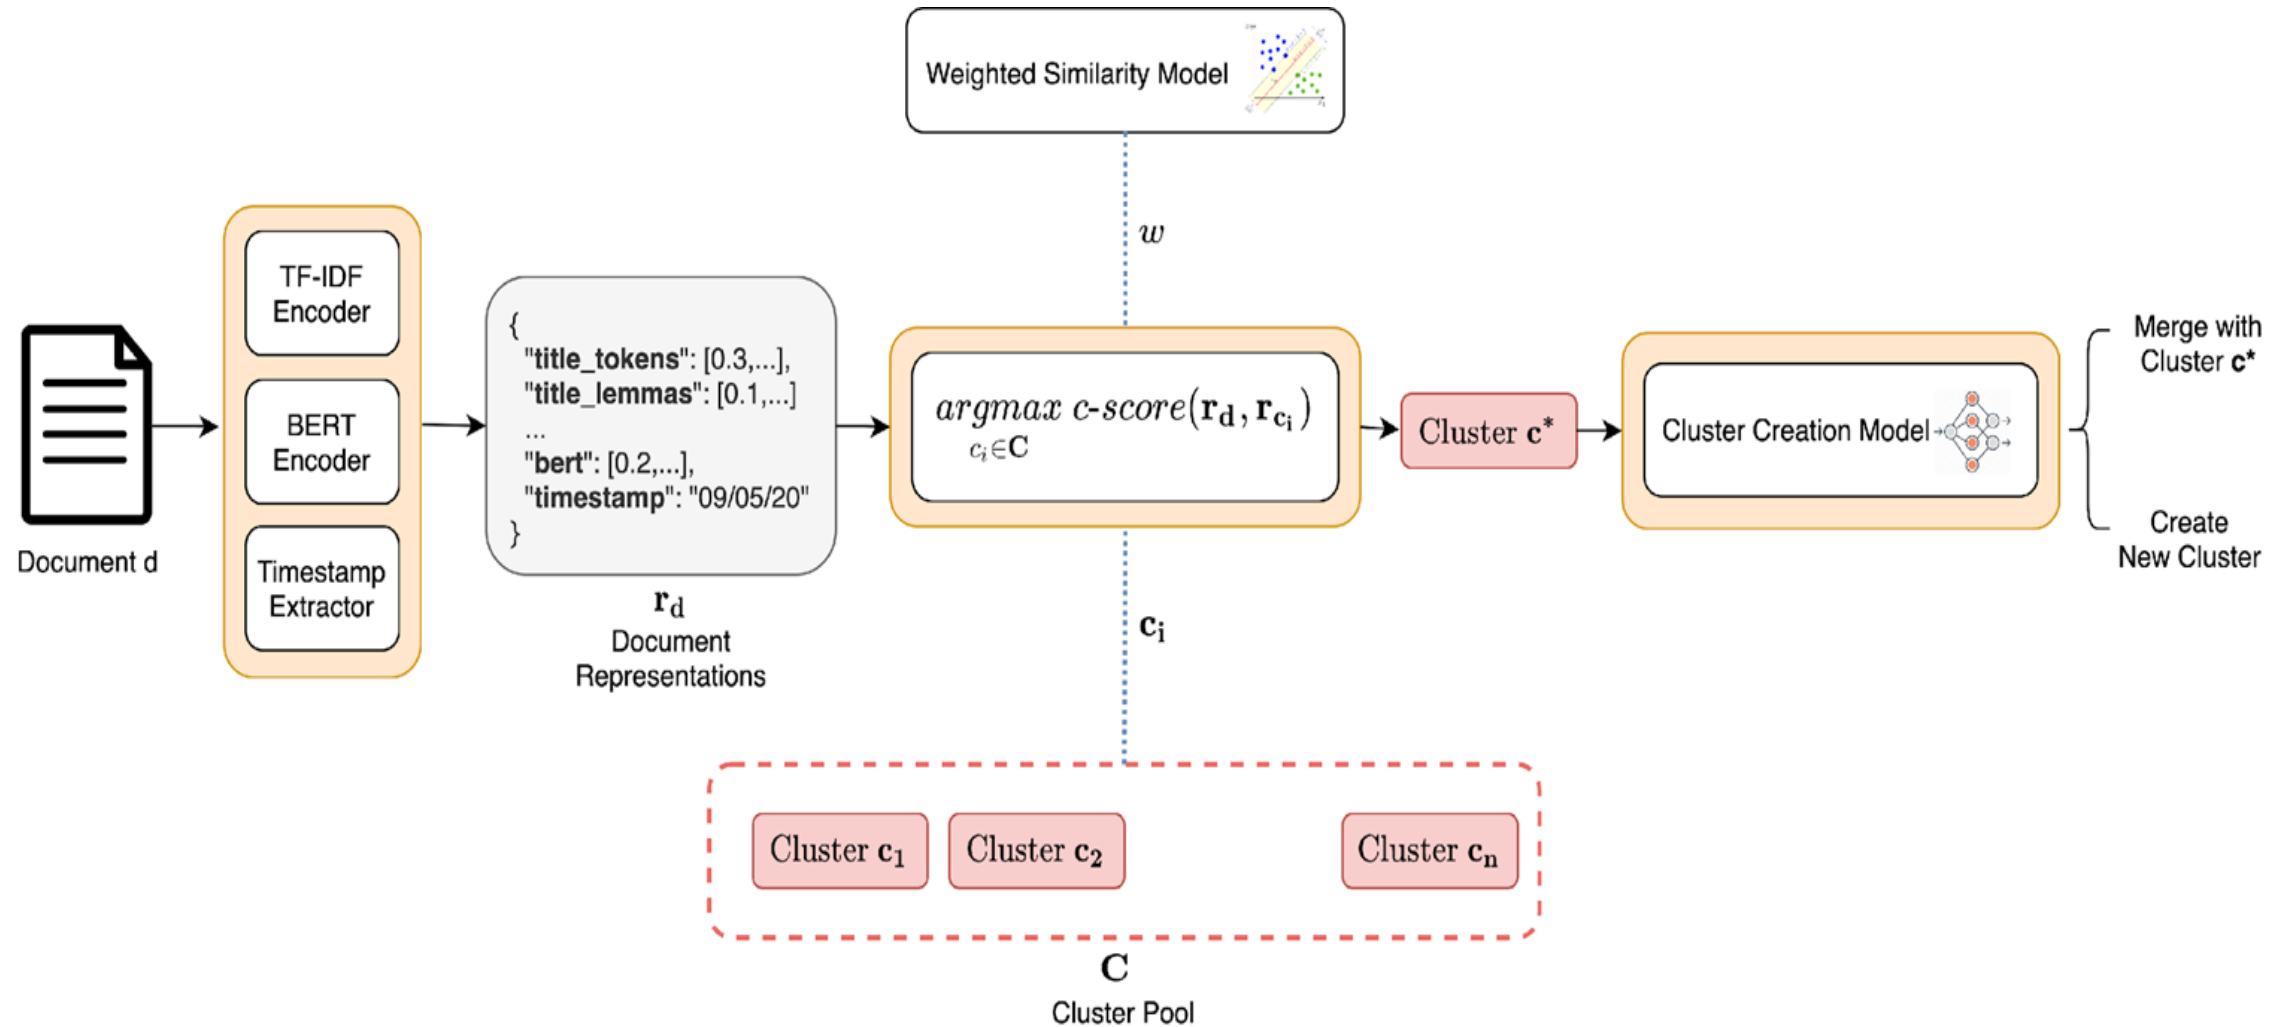
- Document stream is batched temporally (with partial overlap between batches)
- Each batch is clustered and clusters across batches are linked using similarity
- TF-IDF representation of document text and Louvian algorithm for clustering
- [Laban and Hearst \(2017\)](#), [Staykovski et al \(2019\)](#), [Linger and Hajaiej \(2020\)](#)

- Stream clustering

- Online clustering paradigm where new documents are clustered incrementally
- Robust to temporal density variations in document stream
- Doesn't require the tuning of the batch size and overlap hyperparameters
- [Miranda et al \(2018\)](#), [Saravanakumar et al \(2021\)](#)

- Vectors represent title, body, title+body
- TF\*IDF Term Frequency \* Inverse Document Frequency ([Miranda et al 2018](#), [Staykovski et al 2019](#))
- Dense representations ([Staykovski et al 2019](#))
  - Staykovski could not get improvement over TF\*IDF alone
- TF\*IDF plus fine-tuned dense representations ([Saravanakumar et al 2021](#))
  - Entity aware, event aware
- Time
  - Document time-stamp, cluster time max, average, min

# Model Architecture (Saravanakumar et al 2024)



# Results (Saravanakumar et al 2021)



Model	Clusters Count (True Count - 222)	B-Cubed Metrics			
		Precision	Recall	$F_1$ Score	
Laban and Hearst (2017)	873	94.37	85.58	89.76	
Miranda et al. (2018)	326	94.27	90.25	92.36	
Staykovski et al. (2019)	484	95.16	93.66	94.41	
Linger and Hajaiej (2020)	298	94.19	93.55	93.86	
Ours - TF-IDF	530	93.50	80.23	86.36	
Ours - TF-IDF (out-of-order)	413	90.57	87.51	89.01	1
Ours - TF-IDF + Time	222	87.57	96.27	91.72	
Ours - E-S-BERT	452	79.76	60.77	68.98	
Ours - E-S-BERT + Time	471	92.70	74.69	82.73	2
Ours - TF-IDF + P-BERT + Time	196	83.12	<b>97.26</b>	89.63	
Ours - TF-IDF + C-BERT + Time	321	83.10	91.33	87.03	3
Ours - TF-IDF + S-BERT + Time	247	88.30	96.10	92.04	
Ours - TF-IDF + E-S-BERT	433	89.40	86.99	88.18	
Ours - TF-IDF + E-S-BERT (out-of-order)	384	91.15	88.60	89.86	4
Ours - TF-IDF + E-S-BERT + Time	276	<b>94.28</b>	95.25	<b>94.76</b>	

Table 1: Results of clustering performance for different document representation strategies as compared against contemporary models. P-BERT refers to pre-trained BERT; C-BERT refers to BERT fine-tuned on event classification S-BERT refers to BERT fine-tuned using triplet loss on event similarity; E-S-BERT refers to entity aware BERT fine-tuned on event similarity.

# What Might be Useful?


---



- Track an event
  - Visualize
- Provide updates
- Timeline summarization
- Event summarization



# Visualization (Laban & Hearst 2017)



- Using the batch clustering approach to tracking, stories across time are identified (TF\*IDF representation)
- Noun phrase keywords selected from the cluster as labels

# Detail View (Laban & Hearst 2017)

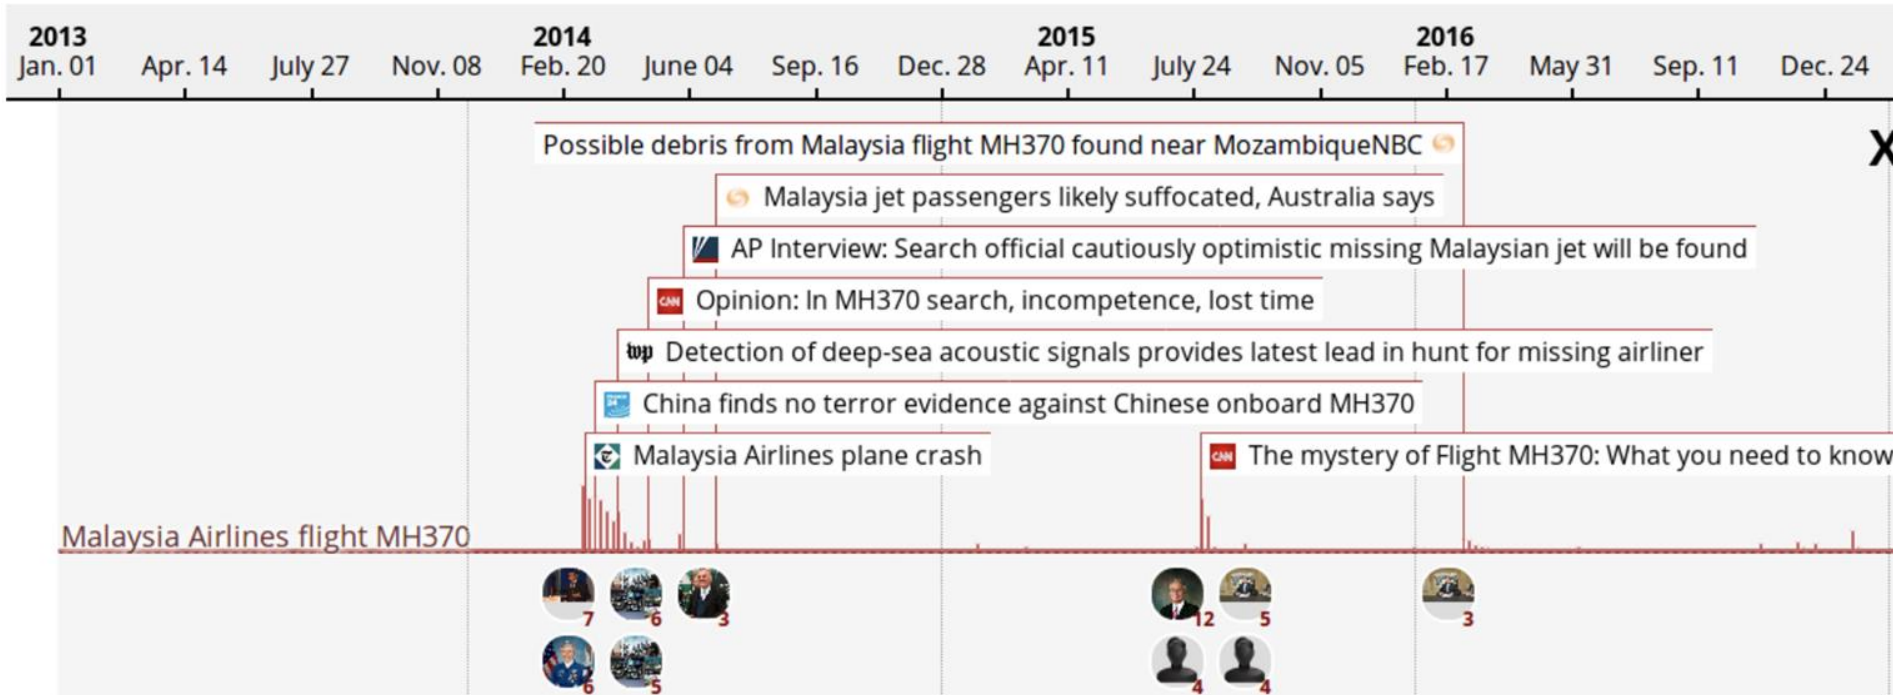


Figure 4: Timeline of the “Malaysia Airline flight MH370”, it has large time gaps with no articles.

Headlines of articles at timepoints with a lot of information are selected to provide a detailed view

# What Might be Useful?

---



- Track an event
- Provide updates
- Timeline summarization
- Event summarization



# Update Summarization


## Data from NIST: 2011 – 2013 Web Crawl, 11 categories



- Trec Temporal Summarization Track
  - ☐ TREC KBA Stream Corpus
  - ☐ hourly web crawl
  - ☐ October 2011 - February 2013
  - ☐ 16.1TB!
- Training Data drawn from Wikipedia
- The update summarizer must decide when to give the update (and thus, how many articles to include in the summary)



:15



nbcdfw.com

local • news • classifieds

headlines: Rothko at the Modern ... city insists brown water safe to drink


The U.S. Pacific Tsunami Warning Center said there was a possibility of a local U.S. Pacific Tsunami Warning Center said there was a possibility of a local tsunami, within 100 or 200 miles of the epicenter, but they were not issuing an immediate warning for the broader region.

The magnitude-7.5 quake, about 20 miles deep, was centered off the town of Champerico.

People fled buildings in Guatemala City, in Mexico City and in the capital of the Mexican state of Chiapas, across the border from Guatemala.

Would you like to contribute to this story? [Start a discussion.](#)

:30



**nbcdfw.com**  
local • news • classifieds

headlines: Rothko at ...

The U.S. said the Pacific there within but the warni

The miles Cham

People in Me

Mexican state of Chiapas, across the border from Guatemala.

Would you like to contribute to this story? [Start a discussion.](#)

**ny1.com**  
local • news • classifieds

headlines: G train stuck forever ... weather on the 1's ... rats eat tourists

A 7.4 magnitude earthquake struck off the coast of Guatemala Wednesday, the U.S. Geological Survey reported.


The epicenter was 124 miles west southwest of Guatemala City.

Reuters reported that the quake could be felt as far away as Mexico City. There were no immediate reports of injury or damage .



Chris Kedzie (Columbia U.)

# 1:00



nbcdf

headlines: Kothko at  
The U.S. said the Pacific there within but the warni  
headlines: G tra  
A 7.4 ma the coast U.S. Geolo  
The epic southwes  
The miles Cham  
People in Me  
Mexican state of border from Guaten  
Would you like to story? [Start a disc](#)

headlines: fixy bike festiv  
GUATEMALA  
Survey says that hit off the Pacific the capital and as Mexico City  
The U.S. Pacific said there was a tsunami, within epicenter, but the an immediate v region.  
The magnitude- deep, was cente Champerico.  
People fled buil Mexico City and


ktla.com  
local • news • classifieds  
headlines: fire in south land ... earthquake in Guatemala ... accident on the 5  
TODAY'S BRIEF

- Greeks protesting austerity measures are clashing with riot police in Athens.
- The U.S. Geological Survey says that a strong earthquake has hit off the Pacific coast of Guatemala, rocking the capital and shaking buildings as far away as Mexico City and El Salvador.
- The election behind them, U.S. investors dumped stocks Wednesday and turned their focus to a world of problems - tax increases and spending cuts that could stall the nation's economic recovery and a deepening recession in Europe.





1:00



## hour 1 updates

- The U.S. Geological Survey says that a strong earthquake has hit off the Pacific coast of Guatemala, rocking the capital and shaking buildings as far away as Mexico City and El Salvador.
- The magnitude-7.5 quake, about 20 miles deep, was centered off the town of Champerico.

Would you like to  
story? [Start a discussion](#)

Champerico.  
People fled buildings  
Mexico City and

focus to a world of problems - tax increases  
and spending cuts that could stall the nation's  
economic recovery and a deepening recession  
in Europe.


# Temporal Summarization Approach



At time  $t$ :

- Predict **salience** for input sentences
  - A domain specific language model to score sentences by event typicality
  - Location to indicate current event
  - Semantics: event type synonyms, hyernyms, hyponyms
- Remove **redundant** sentences
- Cluster and select **exemplar sentences** incorporating salience prediction as a prior

# Saliency Helps



- Saliency predictions lead to high precision quickly
- Saliency predictions allow us to more quickly recover more information

# What Might be Useful?

---



- Track an event
- Provide updates
- Timeline summarization
- Event summarization

# Event Timeline Summarization



## March 16<sup>th</sup>, 2014

- **Crimeans** vote in a referendum to rejoin Russia or return to its status under the 1992 constitution.

## March 17<sup>th</sup>, 2014

- **The Crimean parliament** officially declared independence and requested full accession to the Russian Federation.
- **Obama** declared sanctions on Russian officials considered responsible for the crisis.
- **The leader of the pro-Russian organization** “Youth Unity” was arrested.

## March 18<sup>th</sup>, 2014

- **President Obama** warned Vladimir Putin that further provocations by Russia could isolate and diminish its influence.
- **One pro-Russian soldier** was killed in the Simferopol incident.

... summaries for other dates ...

\* Comment A: The “Crimean Parliament”, headed by an ethnic Russian separatist who was elected leader of parliament AFTER pro-Russian armed forces occupied the parliamentary chambers, has voted for Crimea to be annexed into Russia...

\* Comment B: *Does the West and US have a policy at all? The Obama administration has warned of “increasingly harsh sanctions”, but it is unlikely that Europe will comply...*

\* Comment C: *Sanctions are effective and if done in unison with the EU...*



# Event Timeline Summarization



March 16<sup>th</sup>, 2014

- **Crimeans** vote in a referendum to rejoin Russia or return to its status under the 1992 constitution.

March 17<sup>th</sup>, 2014

- **The Crimean parliament** officially declared independence and requested full accession to the Russian Federation.
- **Obama** declared sanctions on Russian officials considered responsible for the crisis.
- **The leader of the pro-Russian organization** “Youth Unity” was arrested.

March 18<sup>th</sup>, 2014

- **President Obama** warned Vladimir Putin that further provocations by Russia could isolate and diminish its influence.
- **One pro-Russian soldier** was killed in the Simferopol incident.

... summaries for other dates ...

\* Comment A: The “Crimean Parliament”, headed by an ethnic Russian separatist who was elected leader of parliament AFTER pro-Russian armed forces occupied the parliamentary chambers, has voted for Crimea to be annexed into Russia...

\* Comment B: *Does the West and US have a policy at all? The Obama administration has warned of “increasingly harsh sanctions”, but it is unlikely that Europe will comply...*

\* Comment C: *Sanctions are effective and if done in unison with the EU...*

# Data: collected by Wang et al



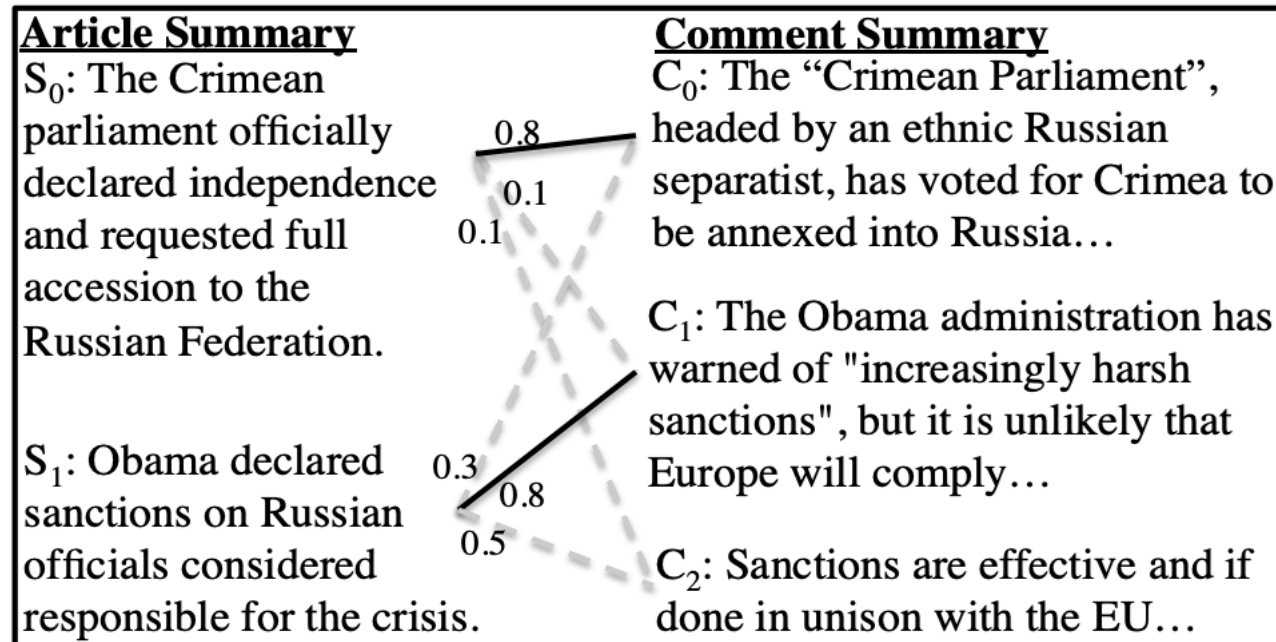
	Time span	# articles	# comments
MH370	03/08-6/30	955	406,646
Ukraine	03/08-6/30	3,779	646,961
Israel-Gaza	07/20-9/30	909	322,244
NSA	3/23-6/30	145	60,481

- Comments in addition to articles
- Summaries are for an article on the event on day x (i.e., not an update)

# Approach



- Joint learning for sentence importance
  - Article summary
  - Comment summary
- Use of event threading to connect an article summary to previous events in the timeline
- Graph –based algorithm



(Wang et al 2015)

- How well could students answer questions about an event (with and without event threads)?

Answer Type	Avg $\pm$ STD	Rated 5 (%)	Rated 4 (%)
No Thread	2.58 $\pm$ 1.20	7%	23%
With Threads	3.29 $\pm$ 1.28	17%	26%

Table 6: Human evaluation on the informativeness of answers written after reading timelines *with threads* vs. with *no thread*. Answers written with access to threads are rated higher (3.29) than the ones with no thread (2.58).

(Wang et al 2015)

- 
- The diagram illustrates the proposed neural network architecture for event-based sentence classification. The input consists of a sequence of words (MJ, was, born, in, ...) and a sequence of events (Event1, Event2, Event3). The words are processed by a Word-attention mechanism, which generates a sequence of hidden states ( $h'_t$ ). The events are processed by an Event-attention mechanism, which generates a sequence of event representations ( $e_t$ ). The hidden states and event representations are combined to produce the final output ( $\hat{y}_t$ ).
- The architecture includes the following components:
- Key-Value Structure:** A table with columns for Key, Global Value, and Local Value. The Key column lists Time1, Time2, and Time3. The Global Value and Local Value columns contain corresponding values.
  - Initial State:** A red dashed circle labeled  $h_0$ .
  - Event-attention:** A yellow box that takes the event sequence as input and outputs  $e_t$ .
  - Word-attention:** A grey box that takes the word sequence as input and outputs  $h'_t$ .
  - Hidden States:** A sequence of red boxes labeled  $h'_t$  and  $h'_{t-1}$ .
  - Output:** A sequence of orange boxes labeled  $\hat{y}_t$  and  $\hat{y}_{t-1}$ .
- The flow of information is as follows:
- The word sequence is processed by the Word-attention mechanism to generate the hidden state sequence  $h'_t$ .
  - The event sequence is processed by the Event-attention mechanism to generate the event representation sequence  $e_t$ .
  - The hidden state  $h'_{t-1}$  and the event representation  $e_t$  are combined to produce the final output  $\hat{y}_t$ .

(Chen et al 2019)

- Query focused event ranking timeline summarization (Nguyen et al 2014)

# What Might be Useful?

---



- Track an event
- Provide updates
- Timeline summarization
- Event summarization



- A summary of a single event on a single day drawn from multiple documents
- Datasets:
  - Multi-News ([Fabbri et al 2019](#))
    - 56,216 articles/summary pairs drawn from newser.com
    - Summaries professionally written
    - > 1500 news sites
  - WikiSum ([Liu et al 2018](#))
    - Input: Wikipedia title (topic)
    - Articles: non-Wikipedia articles referenced from Wikipedia
    - Summary: Lead of the Wikipedia article

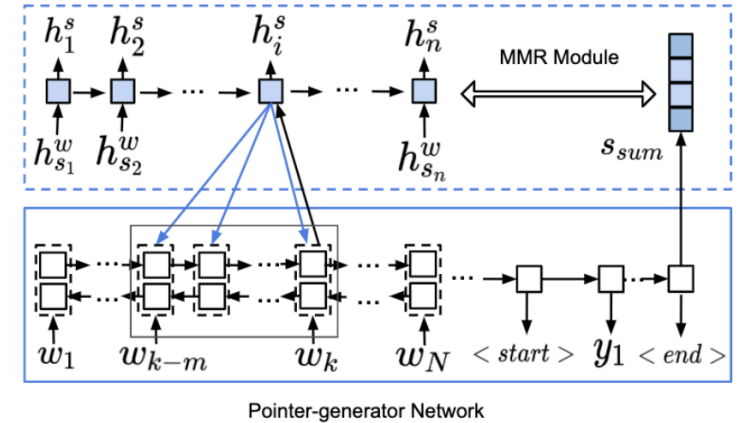
**newser**  
READ LESS **KNOW MORE**



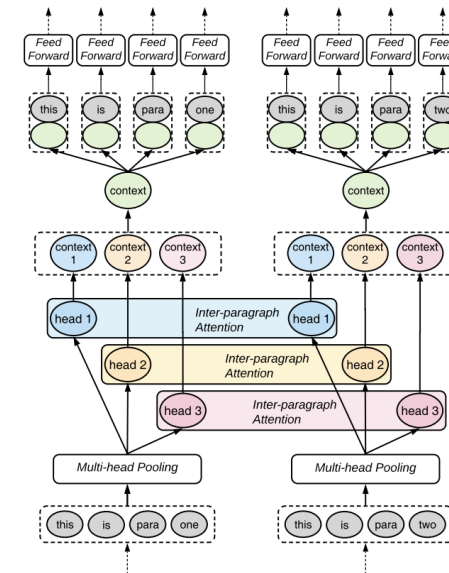
# Multi-document Summarization Models



- Both hierarchical
- Architectures
  - Integrate Maximal Marginal Relevance and standard pointer-generator network to learn parameters for similarity and redundancy (Fabbri et al 2019)
  - Extractive/abstractive models
    - Rank paragraphs in comparison to title using semantic similarity (Lui et al 2018)
    - Rank paragraphs in comparison to summary augmented with graph learning representing semantic similarity between paragraphs and discourse relations (Liu et al 2019)



Fabbri et al 2019



Liu et al 2019

# What Have We Learned?



- Research on complex events focuses on many tasks
- An event can be represented by
  - A cluster of articles ([tracking](#))
  - A single article ([timeline summarization](#))
  - Multiple articles since last update ([update summarization](#))
  - Multiple articles on one day ([multidocument summarization](#))
- What about multiple large-scale events in one article? Is granularity at the right level?
- Could we use temporal relations between events?
- How can we use multiple media, multiple languages, different viewpoints

# Bibliography



- Alexander Fabbri, Irene Li, Tianwei She, Suyi Li, and Dragomir Radev. 2019. Multi-news: A large-scale multi-document summarization dataset and abstractive hierarchical model. In *Proceedings ACL 2019*, Florence, Italy.
- Chris Kedzie, Kathleen McKeown, and Fernando Diaz. 2015. Predicting salient updates for disaster summarization. In *Proceedings ACL 2015*. Beijing, China.
- Chris Kedzie, Kathleen McKeown, and Hal Daume' III. 2018. Content selection in deep learning models of summarization. In *Proceedings EMNLP 2018*, Brussels, Belgium.
- Philippe Laban and Marti Hearst. 2017. newsLens: building and visualizing long-ranging news stories. In *Proceedings of the Events and Stories in the News Workshop*, Vancouver, Canada.
- Mathis Linger and Mhamed Hajaiej. 2020. Batch clustering for multilingual news streaming. arXiv preprint arXiv:2004.08123.
- Peter J. Liu, Mohammad Saleh, Etienne Pot, Ben Goodrich, Ryan Sepassi, Łukasz Kaiser, and Noam Shazeer. Generating wikipedia by summarizing long sequences, in *Proceedings of ICLR 2018*, Vancouver, BC.
- Yang Liu and Mirella Lapata. 2019. Hierarchical transformers for multi-document summarization. In *Proceedings ACL 2019*, Florence, Italy.
- Sebastião Miranda, Arturs Znotins, Shay B. Cohen, and Guntis Barzdins. 2018. Multilingual clustering of streaming news. In *Proceedings EMNLP 2018*, Brussels, Belgium.
- Kiem-Hieu Nguyen, Xavier Tannier, and Veronique Moriceau. 2014. Ranking multidocument event descriptions for building thematic timelines. In *Proceedings COLING 2014*, Dublin, Ireland.
- Kailash Karthik Saravanakumar, Miguel Ballesteros, Muthu Kumar Chandrasekaran, and Kathleen McKeown. 2021. Event-driven news stream clustering using entity-aware contextual embeddings. In *Proceedings EACL 2021*, online.
- Todor Staykovski, Alberto Barro'n-Cedeno, Giovanni Da San Martino, and Preslav Nakov. 2019. Dense vs. sparse representations for news stream clustering. In *Text2Story@ ECIR*.
- Lu Wang, Claire Cardie, and Galen Marchetti. 2015. Socially-informed timeline generation for complex events. In *Proceedings NAACL/HLT 2015*, pages 1055–1065, Denver, Colorado.

Thank you