

# Event Search and Analytics

## Detecting Events in Semantically Annotated Corpora for Search & Analytics

Dhruv Gupta

Max Planck Institute for Informatics, Saarbrücken, Germany  
 Saarbrücken Graduate School of Computer Science, Saarbrücken, Germany  
 dhgupta@mpi-inf.mpg.de

### ABSTRACT

In this article, I present the questions that I seek to answer in my PhD research. I posit to analyze natural language text with the help of semantic annotations and mine important events for navigating large text corpora. Semantic annotations such as named entities, geographic locations, and temporal expressions can help us mine events from the given corpora. These events thus provide us with useful means to discover the locked knowledge in them. I pose three problems that can help unlock this knowledge vault in semantically annotated text corpora: *i.* identifying important events; *ii.* semantic search; and *iii.* event analytics.

### Keywords

Information Retrieval; Text Analytics; Text Summarization; Semantic Annotations; Diversity & Novelty; Semantic Search

## 1. INTRODUCTION

Information retrieval systems have largely relied on word statistics in text corpora to satisfy information needs of users by retrieving documents with high relevance for a given keyword query. In my PhD research I hypothesize that information needs of users can be satisfied to a greater extent by using *events* as a means of navigating text corpora. Events in our context would be an act performed by certain actor(s) at a specific location during a specific time interval. An example would be: *Usain Bolt won the gold medal at the 2008 summer Olympics in Beijing.* With the availability of annotators that can provide us with accurate semantic annotations in form of named entities, geographic locations, and temporal expressions; we can leverage the growing number of knowledge resources such as Wikipedia [37] and ontologies such as Freebase [12] to understand natural language text and mine important events. Formally the central hypothesis can be stated as follows:

**Central Hypothesis** Given text corpora with semantic annotations; traditional information retrieval models can be improved by utilizing knowledge about *events* and using *events* as proxies for information need.

As a toy example consider the following text snippet<sup>1</sup> with demonstrative semantic annotations in Figure 1 :

... Beijing(*Wiki : Beijing*)(*Geo : (39.55, 116.23)*) where Bolt (*Wiki : Usain Bolt*) announced himself to the world with two Olympic golds and two world records in 2008(*Time : [01 - 01 - 2008, 31 - 12 - 2008]*) ...

Figure 1: Semantically annotated text<sup>1</sup> snippet

In the text snippet (Figure 1), we obtain the named entity Usain Bolt whose mention has been identified and disambiguated to point to an external knowledge source. Also identified is a geographical location - Beijing, which is disambiguated and resolved to its geographical coordinates. Likewise the temporal expression 2008 has also been resolved to time range. Having these semantic annotations we can now devise algorithms that can deduce that the event is that of Usain Bolt winning Olympic competition in Beijing, China.

The goal of the proposed research is to leverage the semantic annotations for mining important events and use them to navigate text corpora. The research will find application in many domains of research such as *digital humanities*, in which social scientists are interested in computational history in large digital-born text collections. Anthropologists are interested in cultural and linguistic shifts that occur in such collections. Collectively we can allow *computational culturomics* [24] on corpora to study cultural trends. Events can also be used to link information in multiple and diverse text collections. In short, important events provide a way to create a gist from semantically annotated corpora, which otherwise is not possible through manual human effort.

**Outline.** The article consists of:

- a literature survey (Section 2);
- an overview of the research problems (Section 3);
- available corpora, test sources and evaluation measures for research (Section 5);
- discussion of few open technical problems (Section 6).

<sup>1</sup><http://www.bbc.com/sport/0/athletics/34032366>

## 2. RELATED WORK

In this section I discuss the progress already made in the area of analyzing different semantic annotations in isolation as well as in conjunction for some of the problems proposed.

**Temporal Information Retrieval and Extraction.** Researchers have considered only temporal annotations in text corpora to improve retrieval effectiveness by analyzing the time sensitivity of keyword queries and incorporating the time dimension in retrieval models. Some methods of analysis of time-sensitive queries rely on publication dates of documents [21, 22], while others also look at the temporal expressions in document contents [15]. Several works also take into account the time dimension for re-ranking documents [10] and diversifying them along time [11, 25]. One of the seminal works in extracting temporal events was by Ling and Weld [23]. They outline a probabilistic model to solve the problem of extracting relations from text with temporal constraints.

**Important Events in Annotated Corpora.** One of the most important seminal works in identifying existing and emerging events were the various tasks in *Topic Detection and Tracking* (TDT) [7]. The TDT program aimed to “search, organize and structure” broadcast news media from multiple sources. The five tasks laid within the ambit of TDT where topic tracking, link detection, topic detection, first story detection, and story segmentation. The task of topic tracking required to build a system to detect *on topic stories* from an evaluation corpus after being trained on a set *on topic* stories. The link detection task involved answering a boolean query to whether two given *stories* are related by a common topic. The topic detection task comprised of declaring new *topics* from incoming *stories* which had not been presented to the system. First story detection was another boolean decision task of determining whether the given *story* is a seed story (first-story) to create a new *topic* cluster. Story segmentation task required segmentation of an incoming stream of text into *stories*.

Focusing specifically on extracting and summarizing events in the future, Jatowt and Yeung [20] present a model-based clustering algorithm. The clustering considers both textual and temporal similarities. For computing temporal similarity, the authors model time as a probability distribution by utilizing different family of distributions based on whether its is singular time point, starting date or and ending date. The similarity is then computed using KL-divergence.

Radinsky et al. [26] present an algorithm *Pundit*, which based on the past events in text is able to predict a future event given a query to the system. The events are represented as multidimensional attributes such as time, geographic location and participating entities. The algorithm derives these events from external text collection and builds an *abstraction tree*, which is the result from hierarchical agglomerative clustering. In order to predict the future *Pundit* is trained to select the most similar cluster from the *abstraction tree* and produce an event representation.

The work by Yeung and Jatowt [38] tackles the problem of analysis of historical events in multiple large document collections. They utilize *latent Dirichlet allocation* to identify *topic* distributions along time. Thereafter they perform analytics to answer questions such as *i.* significant years and topics, *ii.* triggers that caused remembrance of the past and *iii.* historical similarity of countries.

Most recently, Abujabal and Berberich [5] present a system which identifies important events in text collections by counting frequent itemsets of sentences containing named entities and temporal expressions. For evaluation they resort to *Wikipedia’s* event directory as a ground truth.

**Semantic Search.** Summarizing text collections in a timeline visualization is a natural choice. Swan and Allan [31] present an approach for producing a timeline that depicts most important topics and events closely modeled on the *Topic Detection and Tracking* task. The algorithms analyzes features based on named entities and noun phrases. The analysis involves construction of  $2 \times 2$  contingency table on presence or absence of features, and subsequent measurement of  $\chi^2$  statistic for measuring significance of co-occurrence of a pair of features.

The seminal work by Baeza-Yates [8] proposed a *future retrieval* (FR) system. The FR system considers both text and temporal expressions to identify future events that might be relevant to an input query. Baeza-Yates outlined the components of a FR system to be composed of an *information extraction* (IE) module, *information retrieval* (IR) module, and a *text mining* (TM) module. The IE module would act as a temporal annotator; identify temporal expressions and normalize them. The IR system is designed to incorporate the time dimension in an index; thus retrieving documents with text and time similarity. The TM module would identify the most relevant topics given a time period. He presented a retrieval model, in which each document consists of a multiple temporal events. A temporal event consists of a time segment and its associated likelihood of occurring. The score of the document is thus obtained by its textual similarity and the maximum likelihood of all the temporal events in that document.

Bast and Buchhold [9] outline a joint index structure over ontologies and text. Which allows for fast semantic search and provide context sensitive auto-complete suggestions.

Events as a means of search document collections has also been explored by Strötgen and Gertz [28]. Events were modeled by the geographic location and time of their occurrence. For temporal queries expressed in simple natural language they outline an extended Backus-Naur form (EBNF) language that incorporates time intervals with standard boolean operations. Geographical queries are also modeled as EBNF language, however the input for them is a minimum bounding rectangle (MBR). Using this multidimensional querying model the user is able to visualize search results in form of events; which are additionally represented on a map.

Giving special attention to geographical information retrieval, Samet et al. [27] present a system *NewsStand*, that is able to resolve and pinpoint a news article based on the geographic information present in its content. They discuss various methods for toponym resolution, which is in essence disambiguating the geographic location based on its surface form in the news content. The system involves a streaming clustering algorithm that can keep track of emerging news in new locations and present them in a map-based interface.

**Event Analytics.** By disambiguating and linking named entities to ontologies, Hoffart et al. [17, 18] provide a framework for semantic search and performing analytics on them. They provide features for giving auto-complete suggestions in the form of similar entities for the input named-entity. In [17] they provide analytics that leverage accurate entity

Query: summer olympics			
Event	$c_1$	$c_2$	$c_3$
Words	micheal, phelps, bejing, china, tibet	london, usain, bolt, england, badminton	rio, brazil, copacabana, deodoro, maracanã
Time	[08-08-2008, 24-08-2008]	[27-07-2012, 12-08-2012]	[05-08-2016, 21-08-2016]
Location	$\langle \text{Beijing, China} \rangle$	$\langle \text{London, England} \rangle$	$\langle \text{Riode Janeiro, Brazil} \rangle$
Entities	$\langle \text{China} \rangle, \langle \text{Micheal_Phelps} \rangle$	$\langle \text{England} \rangle, \langle \text{Badminton} \rangle$	$\langle \text{Brazil} \rangle, \langle \text{Copacabana} \rangle$

Figure 2: Example important events for an example query *summer olympics*

counts and provide entity co-occurrence statistics which is helpful in analyzing semantically similar named-entities.

### 3. RESEARCH OBJECTIVES

Given the text corpora with semantic annotations, I describe three important research problems in this section: *i.* identifying important events; *ii.* using identified events for improving retrieval effectiveness; and *iii.* using identified events for analytics.

#### 3.1 Notation

Let us consider multiple corpora for the purpose of analysis. This allows us to capture frequently occurring events as well as link similar events across corpus. Given corpora

$$D = \bigcup_{k=1}^N D_k,$$

where each document  $d \in D$  consists of word sequence  $x$  at appropriate granularity (e.g. paragraph or sentence):

$$d = \bigcup_{i=1}^n x_i.$$

Further each  $x \in d$  contains semantic annotations in form of *i.* named entities ( $e$ ), *ii.* geographical location ( $g$ ), and *iii.* temporal expressions ( $t$ ). Additionally  $x$  also consists of the a bag of words  $\mathcal{W}$  drawn from a vocabulary  $\mathcal{V}$ . Formally represented as:

$$x = \langle \mathcal{E}, g, t, \mathcal{W} \rangle$$

#### 3.2 Problem Definition

The objective is to design a family of algorithms:

$$\text{EVENT}^*(X, Q, \alpha)$$

where  $X = \bigcup x$ ,  $Q$  represents an input query and  $\alpha \in \mathbb{R}^m$ , where  $\alpha$  is set of parameters.

The input query  $Q$  can be a combination of following input components: *i.* keyword query  $q$ , *ii.* time  $q_{time}$ , *iii.* geographical location  $q_{geo}$ , and *iv.* named entity  $q_{entity}$ .

Given the input, we need to design the algorithms  $\text{EVENT}^*$  according to the different problems. We discuss the design objectives for the three different purposes in this section.

**Identifying Important Events.** *Events* are the proposed building blocks for further text analysis. An *event* in our context is defined to be an activity or an act involving named entities that happens in a specific geographical location anchored to a specific time interval. Mathematically, given a multidimensional query :

$$Q = \langle q, q_{time}, q_{geo}, q_{entity} \rangle,$$

and a subset of highly relevant documents  $R \subseteq D$ , the algorithm for this purpose  $\text{EVENTDETECT}$  should produce a set of ordered events :

$$\mathcal{C} = \{c_1, c_2, \dots, c_k\},$$

where,  $c = \langle \mathcal{E}, g, t, \mathcal{W} \rangle$ . The event  $c$  is hence described by the participating named entities  $\mathcal{E}$ , its location  $g$ , its time of occurrence  $t$ , and frequently occurring contextual terms around these semantic annotations  $\mathcal{W}$ . This requires proposing a probability mass function,  $P(\mathcal{C}, R)$ , using which we can impose a total order on  $\mathcal{C}$ .

As an example consider the keyword-only query *summer olympics* to the processed corpora of news articles. The designed algorithm shall then identify the important events as in Figure 2.

**Diversifying and Summarizing Search Results** are retrieval tasks that try to address the information need underlying an ambiguous query at different levels of textual granularity. Each task tries to maximize the coverage of different information needs underlying the given ambiguous query. As information intents, we propose to use the mined set of *events*. Accomplishing these tasks would allow for automatic creation of *event timelines* or *entity biographies*. We briefly discuss an intuition of achieving the same.

When diversifying search results we would like to present users with *documents* such that the user finds *at least* one document that satisfies her information intent. For this we need to devise an algorithm  $\text{EVENTDIVERSE}$  which considers as an input  $Q$  and  $R \subseteq D$ . As an output it returns a set of documents  $S \subset R$  which cover all events in  $\mathcal{C}$ .

Summarizing search results would require us to construct an algorithm  $\text{EVENTSUMMARY}$  to piece together, *sentences*  $\hat{S} = \bigcup x$ , such that the text summary covers all events in  $\mathcal{C}$ .

**Semantic Search and Analytics.** The mined set of *events* can further be utilized for search and analytics. For this purpose we can utilize inherent hierarchy in the semantic annotations. For example a given year 2015 can be broken down to different months and subsequently days in those months. Similarly, we can utilize the *type hierarchies* in named entities. Such as Usain Bolt and Justin Gatlin are subtypes of *Athletes*. This can jointly be modeled by using the concept of a *data cube* [16] as shown in Figure 3.

Formally, given a query  $Q$ , the objective would to first model the mined set of events as a *data cube* and subsequently provide *data cube operations* [16]:

- roll ( $\circ$ ),
- slice ( $\ominus$ ),
- dice ( $\oplus$ ),
- drill up ( $\Delta$ ),
- drill down ( $\nabla$ ).

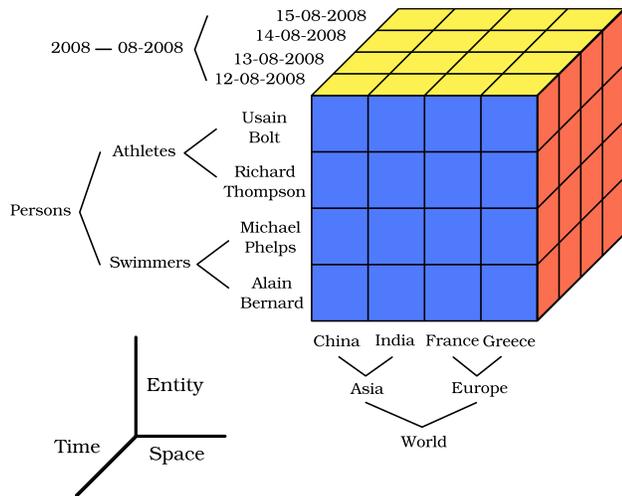


Figure 3: Example data cube based on set of events  $C$

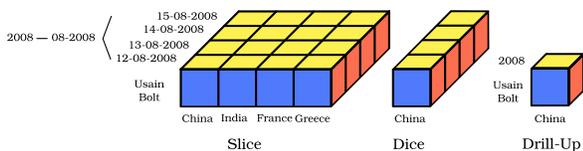


Figure 4: Example data cube operations for the query all races won during 2008 by usain bolt in china

As a concrete example consider the query **all races won during 2008 by usain bolt in china**. To produce an appropriate result the sequence of operations would be: first a slice on the entity **Usain Bolt**; second dice on **China**; and finally drill up to year 2008 (see Figure 4).

## 4. DATA

**Corpora.** There are several readily available massive data sets. They are available from news corporation such as the *New York Times* [4], *English Gigaword* [3]. These corpora have the benefit of being available with reliable publication dates and grammatically well-formed text. On larger scale are Web collections such as *ClueWeb'09* [1]/'12 [2], which are not always accompanied by reliable creation dates and many are ill-formed documents.

**Semantic Annotations.** The text corpora next need to be annotated for text mining. I explain how to obtain the different semantic annotations in the following paragraphs.

**Named Entities.** For disambiguating and linking named entities in text to an external knowledge source such as *Wikipedia* [37] or an ontology such as YAGO [30] or Freebase [12]; I use the AIDA system [19]. The AIDA system does named entity disambiguation and linking by leveraging contexts extracted from ontologies such as YAGO. For Web collections such *ClueWeb'09/'12* the entity disambiguation and linking has been released as *fact1: Freebase annotation of ClueWeb Corpora* [14].

**Geographical Locations** can be obtained by utilizing *geographic* named entities such as those known to be cities, countries, or continents. Geographical relations stored in an ontology can be used to resolve these locations to its geographical coordinates. Having obtained a set of coordinates, we can subsequently construct a geographical representation such as a *minimum bounding rectangle* over the coordinates.

**Temporal Expressions**, both implicit and explicit, can be extracted and normalized from text by using *temporal taggers* such as HeidelbergTime [29] or SUTime [13].

## 5. EVALUATION

To test our approach we need to construct query sets that contain an event description associated with the query; along with participating named entities, geographical locations where the event took place and relevant time interval associated with it. I describe a tentative approach to achieve this here.

**Test Data.** To evaluate the correctness of the various algorithms, I plan to use reliable encyclopedic resources on the Web such as *Wikipedia* [37] or other curated knowledge sources. For an objective evaluation, I propose the following different sources depending on the algorithm under evaluation.

- Identify important events
  - Events in a particular year/decade etc. pages available on *Wikipedia* [33].
  - Testing of past events can be done by extracting important topics from *Category* pages on various historical topics on *Wikipedia* [35].
  - Events in the future can be evaluated by using important infrastructure projects, engineering projects etc. These can be extracted from *Wikipedia* and other sources on the Internet.
  - Current events extracted from *Wikipedia* [36].
  - Alternatively, we can manually construct a list of prominent events and extract relevant information such as named entities, geographical location, and time from ontologies such as: YAGO [30], Freebase [12], etc.
- Diversifying and summarizing search events
  - Biographies of eminent personalities, for example United States presidents [32].
  - Historical timelines of various countries, for example for India [34].

**Structure.** Each event in our test bed is then composed of a fact with an accompanying query. Formally, a *fact* in our testbed is a 4-tuple extracted from one of the aforementioned sources:

$$\langle q, \mathcal{E}, g, t, \mathcal{W} \rangle$$

where  $q$  consists of keyword query describing the event,  $\mathcal{E}$  is a bag of participating entities,  $g$  is the geographic location,  $t$  is the time of its occurrence, and  $\mathcal{W}$  are important terms describing the event.

**Metrics.** Based on the structure of the testbed of events, metrics such as *precision*, *recall* and  $F_1$  can be utilized to measure the effectiveness of the algorithms for detecting important events in semantically annotated corpora. How effectively the algorithm diversifies documents along multiple dimensions can be evaluated by metrics such as  $\alpha$ -nDCG [6]. Quality of summaries can be measured by an automatic evaluation metric called *Rouge* [39].

## 6. DISCUSSION

I briefly present some open technical challenges that I will address along with the research objectives in my PhD dissertation.

**Mathematical Models.** One key aspect that occurs in the design of the algorithms is that of computational models for named entities, geographical locations and temporal expressions. What would be the most descriptive mathematical models for each of these semantic annotations?

**Similarity Functions.** Given a pair of named entities, geographical locations or temporal expressions; how can we efficiently compute the similarity between the same type of annotations?

**Efficiency & Scalability.** Identifying data structures for indexing corpora along with their semantic annotations, such that their asymptotic run times scale linearly with the size of the corpora.

**Evaluation.** Since evaluation of the solutions outlined are very subjective in nature; what are other reliable sources of objective ground truth? What other metrics can be employed to test the effectiveness of our methods?

## 7. CONCLUSION

In this article I laid out an outline of the research work that I envisage to carry out for my PhD dissertation. The research would in its culmination provide us methods to computationally extract world history as sequence of temporally ordered events and portray future events to take place from semantically annotated corpora. The research would also provide ways to perform semantic search and large scale event analytics on these annotated corpora. I further described already available resources that can be utilized for carrying out the research; test cases that can be built from encyclopedic resources on the Internet; and the metrics that can be utilized for evaluation.

## 8. REFERENCES

- [1] Clueweb'09. [Online; accessed 26-August-2015] <http://www.lemurproject.org/clueweb09.php/>.
- [2] Clueweb'12. [Online; accessed 26-August-2015] <http://www.lemurproject.org/clueweb12.php/>.
- [3] English gigaword. [Online; accessed 26-August-2015] <https://catalog.ldc.upenn.edu/LDC2003T05>.
- [4] New york times anotated corpus. [Online; accessed 26-August-2015] <https://catalog.ldc.upenn.edu/LDC2008T19>.
- [5] Abujabal A. and Berberich K. Important events in the past, present, and future. WWW'15 Companion Volume.
- [6] Agrawal R. et al. Diversifying search results. WSDM'09.
- [7] Allan J., editor. *Topic Detection and Tracking: Event-based Information Organization*. Kluwer Academic Publishers, Norwell, MA, USA, 2002.
- [8] Baeza-Yates R. Searching the future. SIGIR'05 Workshop MF/IR.
- [9] Bast H. and Buchhold B. An index for efficient semantic full-text search. CIKM'13.
- [10] Berberich K. et al. A language modeling approach for temporal information needs. ECIR'10.
- [11] Berberich K. and Bedathur S. Temporal diversification of search results. TAIA'13.
- [12] Bollacker K. et al. Freebase: A collaboratively created graph database for structuring human knowledge. SIGMOD'08.
- [13] Chang A. X. and Manning C. D. SUTIME: A library for recognizing and normalizing time expressions. LREC'12.
- [14] Ringgaard M. et al. Faccl: Freebase annotation of clueweb corpora, version 1 (release date 2013-06-26, format version 1, correction level 0), June 2013. <http://lemurproject.org/clueweb12/>.
- [15] Gupta D. and Berberich K. Temporal query classification at different granularities. SPIRE'15.
- [16] Han J. et al. *Data Mining: Concepts and Techniques*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 3rd edition, 2011.
- [17] Hoffart J. et al. AESTHETICS: analytics with strings, things, and cats. CIKM'14.
- [18] Hoffart J. et al. STICS: searching with strings, things, and cats. SIGIR'14.
- [19] Hoffart J. et al. Robust disambiguation of named entities in text. EMNLP'11.
- [20] Jatowt A. and Yeung C. A. Extracting collective expectations about the future from large text collections. CIKM'11.
- [21] Jones R. and Diaz F. Temporal profiles of queries. *ACM Trans. Inf. Syst.*, 25(3), July 2007.
- [22] Kanhabua N. and Nørvåg K. Determining time of queries for re-ranking search results. ECDL'10.
- [23] Ling X. and Weld D. S. Temporal information extraction. AAAI'10.
- [24] Michel J. B. et al. Quantitative analysis of culture using millions of digitized books. Science'10.
- [25] Nguyen T. N. and Kanhabua N. Leveraging dynamic query subtopics for time-aware search result diversification. ECIR'14.
- [26] Radinsky K. et al. Learning to predict from textual data. *J. Artif. Intell. Res. (JAIR)*, 45:641–684, 2012.
- [27] Samet H. et al. Reading news with maps by exploiting spatial synonyms. *Commun. ACM*, 57(10):64–77, 2014.
- [28] Strötgen J. and Gertz M. Event-centric search and exploration in document collections. JCDL'12.
- [29] Strötgen J. and Gertz M. Multilingual and cross-domain temporal tagging. *Language Resources and Evaluation*, 47(2):269–298, 2013.
- [30] Suchanek F. M. et al. Yago: A large ontology from wikipedia and wordnet. *Web Semant.*, 6(3):203–217, September 2008.
- [31] Swan R. C. and Allan J. Automatic generation of overview timelines. SIGIR'00.
- [32] Wikipedia. List of presidents of the united states — Wikipedia, the free encyclopedia, 2015. [Online; accessed 26-August-2015] [https://en.wikipedia.org/wiki/List\\_of\\_Presidents\\_of\\_the\\_United\\_States](https://en.wikipedia.org/wiki/List_of_Presidents_of_the_United_States).
- [33] Wikipedia. List of years — Wikipedia, the free encyclopedia, 2015. [Online; accessed 26-August-2015] [https://en.wikipedia.org/wiki/List\\_of\\_years](https://en.wikipedia.org/wiki/List_of_years).
- [34] Wikipedia. List of years in india — Wikipedia, the free encyclopedia, 2015. [Online; accessed 26-August-2015] [https://en.wikipedia.org/wiki/List\\_of\\_years\\_in\\_India](https://en.wikipedia.org/wiki/List_of_years_in_India).
- [35] Wikipedia. Portal:contents/history and events — Wikipedia, the free encyclopedia, 2015. [Online; accessed 26-August-2015] [https://en.wikipedia.org/wiki/Portal:Contents/History\\_and\\_events](https://en.wikipedia.org/wiki/Portal:Contents/History_and_events).
- [36] Wikipedia. Portal:current\_events — Wikipedia, the free encyclopedia, 2015. [Online; accessed 26-August-2015] [https://en.wikipedia.org/wiki/Portal:Contents/History\\_and\\_events](https://en.wikipedia.org/wiki/Portal:Contents/History_and_events).
- [37] Wikipedia. Wikipedia, the free encyclopedia, 2015. [Online; accessed 26-August-2015] <http://en.wikipedia.org/>.
- [38] Yeung C. A. and Jatowt A. Studying how the past is remembered: towards computational history through large scale text mining. CIKM'11.
- [39] Lin C. Y. Rouge: a package for automatic evaluation of summaries. ACL-04 workshop.