Temporal Information Retrieval: Challenges and Opportunities

Omar Alonso
Microsoft Corp.
Mountain View, CA
omar.alonso@microsoft.com

Jannik Strötgen
Institute of Computer Science
University of Heidelberg, Germany
stroetgen@uni-hd.de

Ricardo Baeza-Yates
Yahoo! Research
Barcelona, Spain
rbaeza@acm.org

Michael Gertz
Institute of Computer Science
University of Heidelberg, Germany
gertz@uni-hd.de

ABSTRACT

Time is an important dimension of any information space. It can be very useful for a wide range of information retrieval tasks such as document exploration, similarity search, summarization, and clustering. Traditionally, information retrieval applications do not take full advantage of all the temporal information embedded in documents to provide alternative search features and user experience. However, in the last few years there has been exciting work on analyzing and exploiting temporal information for the presentation, organization, and in particular the exploration of search results.

In this paper, we review the current research trends and present a number of interesting applications along with open problems. The goal is to discuss interesting areas and future work for this exciting field of information management.

Categories and Subject Descriptors

I.2.7 [**Artificial Intelligence**]: Natural Language Processing—*Language models, Text analysis*

Keywords

temporal information, information retrieval

1. INTRODUCTION

Time clearly plays a central role in any information space, and it has been studied in several areas like information extraction, topic-detection, question-answering, query log analysis, and summarization. Time and temporal measurements can help recreating a particular historical period or describing the chronological context of a document or a collection of documents. As an extension to existing ranking techniques, which are primarily based on popularity or reputation, time can be in particular valuable for exploring search results along well-defined timelines and at multiple time granularities due to the key characteristics of temporal information:

- Temporal information is well-defined: Assuming two points in time or two intervals, the relationship between them can be identified, e.g., the relationship can be of the types before, overlap, or after [3].
- Temporal information can be normalized: Regardless of the used terms or the used language, every temporal expression referring to the same semantics can be normalized to the same value in some standard format. This property makes temporal information term- and language-independent.
- Temporal information can be organized hierarchically: Temporal expressions can be of different granularities, e.g., of type day ("May 20, 2011") or of type year ("2011"). Due to the fact that years consist of months, and months and weeks consist of days, temporal expressions can be mapped to coarser granularities based on the hierarchy of temporal expressions.

Using these key characteristics, temporal information about documents can be used to develop time-specific information retrieval and exploration applications. The most obvious type of temporal information associated with a document is its creation time or the date of its last modification. This kind of information, which is directly accessible through the metadata of a document, can already be used for several tasks such as time-aware search or temporal clustering. However, the document creation time is only valuable in a specific context such as the news domain. In other areas, and even in the news domain itself, a lot of temporal information is neglected if the document creation time is used as the only time information associated with a document. This is because there is a lot of temporal information latently available in a document's text. Assume a news document reports on an event that is already dated. Then, if only the

TWAW 2011, March 28, 2011, Hyderabad, India.

document creation time is taken into account, the information when the event occurred is ignored. But to make use of such information, so-called temporal taggers are applied to extract and normalize temporal expressions contained in documents.

The remainder of the paper is organized as follows. After a discussion of how time appears in documents and how it is possible to extract such temporal data, in Section 3, we survey research on temporal tagging. In Section 4, we present the current research trends on temporal information retrieval. We then describe application areas and challenges. Finally, we present our concluding remarks.

2. TIME IN DOCUMENTS

As indicated in the introduction, there is a lot of temporal information in any collection of documents, be it ranked documents in a hit list or a corpus of topic specific documents. To take advantage of such time related information in particular for document exploration purposes, in a document processing step, it is important to extract this information, anchor it in time, compute some (aggregated) measures, and make all this information explicit to subsequent exploration tasks.

In this section, we give a description of the different types of temporal information mentioned in documents (Sec. 2.1), explain how temporal expressions can be realized in natural language (Sec. 2.2), and demonstrate how they can be extracted and normalized (Sec. 2.3).

2.1 Types of Temporal Information

Temporal expressions mentioned in text documents can be grouped into four types according to TimeML [27], the standard markup language for temporal information: date, time, duration, and set. While duration expressions are used to provide information about the length of an interval (e.g., "three years" in *they have been traveling through the U.S. for three years*), set expressions inform about the periodical aspect of an event (e.g., "twice a week" in *she goes to the gym twice a week*). In contrast, time and date expressions (e.g., "3 p.m." or "January 25, 2010") both refer to a specific point in time – though in a different granularity.

An interesting key feature of temporal information is that it can be normalized to some standard format. Assuming a Gregorian calendar as representation of time, expressions of time and date can be directly placed on a timeline. A date is then typically represented as YYYY-MM-DD, e.g., the expression "January 25, 2010" is normalized to "2010-01-25". However, the normalization is not always as simple as in this example, but depends on the way temporal information is expressed in a document, which will be discussed in the next paragraph.

2.2 Occurrences of Temporal Expressions

There are many different ways how to express temporal information of the types date and time in documents. Similar to the work by Schilder and Habel [31], we distinguish between explicit, implicit, and relative temporal expressions.

Explicit temporal expressions refer to a specific point in time. Note that this point in time can be of different granularities. For example, the expression "January 25, 2010" refers to a specific day while the expression "November 2005" refers to a specific month. The key characteristic of explicit temporal expressions is that they can be normalized without any further knowledge. That is, the expression itself contains all the information needed for normalization and is thus fully specified.

In contrast, *relative temporal expressions* cannot be normalized without taking into account some context information. For example, the expression "today" cannot be normalized without knowing the corresponding reference time. This reference time can either be the document creation time or another temporal expression in the document. Typically, in news articles, the document creation time is important and often used as reference time. Note that this kind of information is directly accessible in form of a timestamp through the metadata of a document. The expression "yesterday" in *Thousands of prisoners in Egypt managed to escape from prison yesterday* can be normalized to "2011-01-29" if we know the document creation time to be "2011-01-30". In other types of documents, the reference time is usually represented by another temporal expression in the document. In general, there are many occurrences of relative temporal expressions. While sometimes the reference time is sufficient for normalization, further information is needed if the relation to the reference time has to be identified as well. For example, "on Monday" can either refer to the previous or to the next Monday with respect to the reference time. An indicator for determining the relationship can be the tense of the sentence with future tense and present tense indicating an after-relationship to the reference time and past tense a before-relationship. Figure 1 shows some parts of a news article containing explicit and relative temporal expressions and illustrate what kind of context information is needed for normalizing the relative expressions.

The third type of temporal expressions are *implicit expressions* such as names of holidays or events. These expressions can be anchored on a timeline if a mapping of the expression to its normalized value is available. For example, "New Year's Day 2002" can be normalized to "2002-01-01" since "New Year's Day" always refers to January 1. In addition, there are expressions for which a temporal function has to be applied. "Labor Day", for example, refers to the first Monday in September so that "Labor Day 2009" can be normalized to "2009-09-07" if we know this day to be the first Monday in September 2009.

Although there are many different ways to refer to a specific point in time, all expressions referring to the same point in time shall be normalized to the same value in the standard format. This normalization process is one of the tasks of so-called temporal taggers, as described in the next paragraph.

2.3 Temporal Tagging

Temporal tagging is a specific task in named entity recognition and normalization. The goals of so-called temporal taggers are (i) the extraction of temporal expressions and (ii) the normalization of these expressions to some standard format. As this standard format, TIMEX2 and TIMEX3 are often used. While TIMEX2 tags include pre- and post-modifiers of the temporal expression itself (e.g., dependent clauses) and allow for nested temporal expressions [11], such modifiers and nested tags are not included by TIMEX3 tags. Instead, TIMEX3 is part of the TimeML markup language in which further annotation types are available for capturing more complex temporal semantics. Nevertheless, although there are significant differences between TIMEX2

Document Creation Time: 1998-04-18

Hungarian astronaut Bertalan Farkas is leaving for the United States to start a new career, he said today. . . . On May 22, 1995, Farkas was made a brigadier general, and the following year he was appointed military attache . . . However, cited by District of Columbia traffic police in December for driving under the influence of . . .

Figure 1: Examples of temporal expressions in a news article with explicit (transparent boxes) and relative (solid boxes) expressions. Arrows indicate what kind of context information is needed to normalize the temporal expression.

and TIMEX3, they are very similar in many ways and a detailed analysis goes beyond the scope of this paper.[1] According to the TimeML annotation guidelines, a TIMEX3 tag contains, among others, the following information about a temporal expression:

- offset: the start and end position of the expression in the document
- type: whether the expression is of type date, time, duration, or set
- value: the normalized value of the expression

To identify this information, i.e., to extract and normalize temporal expressions, temporal taggers are applied after the text is preprocessed. Usually, sentence and token boundaries are detected and a part-of-speech tag is associated with every token. This information can then be used by the temporal tagger to identify temporal expressions. The first goal, i.e., the identification of the boundaries of temporal expressions, can be seen as typical classification task. Therefore, there has been some work on addressing this problem by applying machine learning techniques (e.g., [15, 38]). The classification problem can be described in the following way: For every token t, decide whether t is outside (O) of temporal expressions, inside (I) a temporal expression, or the beginning (B) of a temporal expression. The well-known IOB-format can be used for annotating tokens according to their property.

In addition to machine-learning approaches, there are several rule-based approaches to extract temporal expressions (e.g., [24, 34]). These are usually based on regular expressions although they may use other information about the text as well, such as part-of-speech information.

The more difficult task is the normalization of the temporal expressions. While explicit expressions can be normalized without further knowledge, the normalization of relative expressions is challenging. As described above, context information has to be identified to determine the correct reference time and the temporal relation between a temporal expression and its reference time. While there are rule-based and machine learning based approaches for the extraction of temporal expressions, the normalization is usually done in a rule-based way by all temporal taggers.

Due to their importance for temporal information retrieval, we give an overview of existing temporal taggers and their quality in the next section. In addition, we present resources for evaluating temporal taggers and survey temporal evaluation challenges organized so far.

[1]For further information on temporal annotation according to TimeML and differences between TIMEX2 and TIMEX3, see `http://www.timeml.org`.

3. RESEARCH ON TEMPORAL TAGGING

Temporal processing of text documents in terms of the extraction and normalization of temporal expressions as well as the extraction of temporal relations between events is very important for several NLP tasks requiring a deep understanding of language such as question answering or document summarization. Due to this fact, there has been significant research in temporal annotation of text documents. The markup language TimeML has become an ISO standard for temporal annotation [27], and the TimeBank corpus was developed [28]. The latest version of the TimeBank corpus contains 183 news articles and can be regarded as the gold standard for temporal annotation. However, there has been important research activity before, and several evaluation challenges have been held to bring forward research in the area of temporal information extraction as described in the following section.

3.1 Evaluation Challenges

The earliest competitions for the extraction of temporal expressions have been the named entity recognition tasks of the Message Understanding Conferences MUC 1995 and 1997 [8, 12]. A combination of the extraction and the normalization was introduced in the ACE (Automatic Content Extraction) time expression recognition and normalization (TERN) challenges in 2004, 2005, and 2007[2]. Several temporal taggers have been developed by the participants of these challenges (see Section 3.2). Often, the TERN 2004 and 2005 corpora[3] are used to compare the quality of temporal taggers. The TERN corpora are annotated with respect to the TIMEX2 annotation guidelines [11].

A further indication of the importance of temporal annotation and the activity in the research domain are the TempEval challenges. Motivated by the importance of temporal annotation for many NLP tasks, TempEval was organized the first time as one task of the SemEval workshop in 2007 [39]. In this competition, the organizers provided annotated text documents based on the TimeBank corpus. While the annotations of events and temporal expressions were given, the task for the tools to be developed was to identify temporal relations between events and the document creation time, between events and temporal expressions, and between two events.

In 2010, the full task of identifying all temporal related expressions and relations was faced in the follow-up challenge. That is, for TempEval-2, two further tasks were added [41]: the extraction and normalization of temporal expressions and of events. In addition, the discovery of relations between two events was split into two tasks, namely the identification of relations between two main events in consecu-

[2]`http://www.itl.nist.gov/iad/mig/tests/ace/`.

[3]The TERN development corpora are available through the Linguistic Data Consortium. See: `http://fofoca.mitre.org/tern.html`.

tive sentences and relations between two events where one event syntactically dominates the other event. The TempEval corpora are based on the TimeBank corpus and annotated according the TimeML annotation guidelines, i.e., using TIMEX3 tags for temporal expressions. A further novelty in the second TempEval challenge was that the tasks were offered not only in English but in six languages. However, only two languages where addressed by the participants, namely English and Spanish. Nevertheless, thanks to the creation of an annotation standard, a gold standard corpus, and competitions such as the TempEval challenges, there has been significant improvements in temporal relation identification and temporal tagging. Some existing temporal taggers and their quality is presented in the next paragraph by comparing their results in the TempEval-2 challenge.

3.2 Temporal Taggers

Having applied a temporal tagger on a document collection, the previously hidden temporal information is made available for tasks such as temporal relation extraction or temporal clustering. One often applied temporal tagger is GuTime, which is part of the Tarsqi tool kit [40]. It is based on the TempEx tagger, which was the first temporal tagger for the extraction of temporal expressions and their normalizations [22]. GuTime was developed as automatic evaluation tool for TimeML and extends the capabilities of the TempEx tagger. It was evaluated on the TERN 2004 training corpus and achieves F-measures of 85%, 78%, and 82% for lenient and strict detection and for normalization, respectively.

In the TempEval-2 challenge, eight teams participated in the task for temporal expression extraction and normalization for English documents. The best-performing system was HeidelTime with an F-Score of 86% for the extraction and an accuracy of 85% for the normalization [34]. For Spanish documents, the best result for the extraction was an F-Score of 91% [16], while another system achieved the highest accuracy for normalization (83%) [42]. While both machine-learning and rule-based approaches were applied for the extraction, the normalization was done in a rule-based way by all systems. As the best performing system, HeidelTime uses rules consisting of an extraction and a normalization part. Thus, all temporal expressions that are identified are normalized as well. Due to the strict separation of the code and the rules, HeidelTime is applicable for multi-lingual temporal tagging[4].

Although there has been significant advances in temporal tagging, there is still room for improvement, especially when switching the processing language or the domain of the document collection. For example, Mazur and Dale recently presented a new corpus for research on temporal expressions containing long, narrative-style documents, namely Wikipedia articles describing the historical course of wars [25]. Using their temporal tagger, they show that the normalization of temporal expressions in such documents is very challenging due to the rich discourse structure and the huge number of often underspecified temporal expressions in these documents compared to the usually used short news documents.

[4]For details, see `http://dbs.ifi.uni-heidelberg.de/stixx/`.

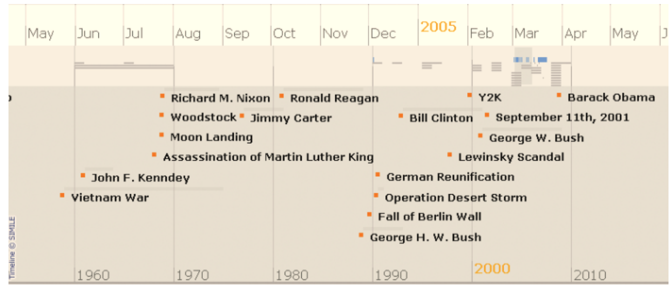

Figure 2: Annotation of a timeline by workers using crowdsourcing.

4. RESEARCH TRENDS

Research work on fully utilizing the temporal information embedded in the text of documents for exploration and search purposes is very recent. The work by Alonso et al. presents an approach for extracting temporal information and how it can be used for clustering search results [5]. Berberich et al. describe a model for temporal information needs [7]. Figure 2 shows the annotated timeline for the NYTimes data set for the latter reference using the Timeline widget[5]. These last two projects rely on crowdsourcing, mainly using Amazon Mechanical Turk, for evaluating parts of their work.

News sources have been the primary focus of a number of projects on exploiting time information in documents. For example, the Time Frames project realizes an approach to augment news articles by extracting time information [14]. Google's news timeline[6] is an experimental feature that allows a user to explore news by time.

Extensions to document operations such as comparing the temporal similarity of two documents in the context of news articles is presented by Makkonen and Ahonen-Myka [17]. An interesting approach that combines topic detection and tracking with timelines as a browsing interface is presented by Swan and Allan [37]. Time information is also used in temporal mining of blogs to extract useful information [29]. New research has also emerged for *future retrieval* where temporal information is used for searching the future [6].

There is exciting research on adding a time dimension to certain applications like news summaries [2], temporal patterns [33], and temporal Web search [26]. The special issue on temporal information processing gives a clear map of current directions [20]. Harvesting temporal information and how it can be used for entities and relationships is also a very recent rich area [43].

Closely related to information extraction is the recent research on *temporal annotations*, which is covered in depth in the book by Mani et al. [21]. Identification of time related information depends heavily on the language and the corpora, so traditional information extraction systems tend to fall short in terms of temporal extraction. Based on the latest advances, new research is emerging for automatic assignment of document event-time periods and automatic tagging of news messages using entity extraction [31].

Now, we outline a number of applications that can benefit from leveraging more temporal information either by temporal expressions or timestamps. For each application, we

[5]`http://simile.mit.edu/`
[6]`http://www.newstimeline.googlelabs.com`

describe why it is important and present a number of challenges.

4.1 Exploratory Search

Research in exploratory search systems has gained a lot of attention lately as they add a significant user interface component to help users search, navigate, and discover new facts and relationships. As the amount of information on the Web keeps growing, exploratory search interfaces are starting to surface. That said, it is not clear how to leverage temporal information. A few problems are:

- How to expose temporal information in exploratory search systems?
- What's the best way of presenting temporal information as a retrieval *cue*?
- For which data sources, besides news, does exploratory search make sense?
- Is e-discovery a vertical application that can benefit from temporal information?

4.2 Micro-blogging and Real-time Search

Micro-blogging sites like Twitter have gained a lot of attention lately as the ultimate mechanism to broadcast what's going on. Due to its nature, a typical message is very short and its lifespan is basically the *crowd* interest about that particular event be a football final game or an earthquake.

In the case of Twitter, it is very difficult to beat the timely broadcasting of an important event if one compares this to a news article. Each tweet has a timestamp but the organization of such information is still not clear. In the news context, the reporter has to write an article that contains a few paragraphs and submit the final version through some content management version that would push it to an external website so a search engine can hopefully crawl and index it in time. In parallel, if a tweet is so important by the time the reporter is finishing with the article, the main idea would be trending in Twitter, therefore highlighting its importance at a world scale. This is very similar to the traditional notion of topic detection and tracking [1, 18], with one key difference: speed to detect that the topic is important and therefore a candidate for trending. Some problems are:

- What is the best way to provide a timeline of events in micro-blogging?
- What is the lifespan of the main event?
- How fast and precise can one detect trending events?
- What is the fraction of new content on the topic stream?

4.3 Temporal Summaries

There has been seminal work on temporal summaries of news topics by Allan [2] that shows how important temporal information is. One extension is to generate time sensitive summaries that can be used as temporal snippets [4].

By design, the main goal of a snippet (or caption) is to present a document surrogate that the user can quickly scan in the search results page without the need to click and read the full content of a document. There is a limit to the number of lines of text that the snippet should present. Interesting questions include:

- When to show a timestamp or temporal expressions?
- Should the snippet present the matching lines in a timeline?
- Is a temporal summary a good surrogate for a document?
- For which kind of queries is a temporal summary appropriate?
- Should temporal summaries be query independent?

4.4 Temporal Clustering

The notion of clustering search results by temporal attributes has been presented in [5]. Preliminary results indicate that users are interested in dissecting a document collection by time. At the same time it is not clear for which kind of scenarios besides "research-like" questions this approach would work. Key issues are:

- Can we identify documents that are contemporary and therefore related?
- Which chronons can be more useful for clustering?
- How can we cluster micro-blogging data by time?
- Is a timeline the best way to cluster search results?

4.5 Temporal Querying

The temporal information extracted from documents can directly be used to allow the user of a search engine to constrain his/her query in a temporal manner. That is, in addition to a textual part, a query contains a temporal part. For example, in addition to "world war" a temporal constraint like "1944-1945" could be specified. The user would obviously expect documents about World War II as results for his query. The objective when using a combination of a text and a temporal query can thus be formulated in the following way: The more both parts of the query are satisfied, i.e, the more the textual and the temporal parts fit to a document, the higher should be the rank of this document. The main problems for such a combination of constraints is the following:

- How can a combined score for the textual part and the temporal part of a query be calculated in a reasonable way?
- Should a document in which the "textual match" and the "temporal match" are far away from each other be penalized?
- What about documents satisfying one of the constraints but "slightly" fail to satisfy the other constraint?

4.6 Temporal Question Answering

To be able to answer time-related questions, a question answering system has to know when specific events took place. For this, temporal information can be associated with extracted facts from text documents [26]. While this may be applicable for famous facts and events, question answering systems are often faced with imperfect temporal information. For this, identifying relationships between events described in documents is important as it is for many further NLP tasks (see Section 3.1). Especially historic events tend to have a gradual beginning and ending so that knowing the temporal relationship between two events may allow to answer a temporal query although no explicit temporal information is associated with the events [30]. Research issues are:

- How can inconsistent temporal information be dealt with?
- How can temporal reasoning be executed if temporal relationships are missing?

4.7 Temporal Similarity

A related research question to temporal querying is temporal document similarity. Instead of comparing a temporal query with the temporal information of a document, two documents can be compared with respect to their temporal similarity. The main problem arising here is what makes two documents temporally similar? This leads to the following questions:

- Should two documents be considered similar if they cover the same temporal interval?
- Should the temporal focus of the documents be important for their temporal similarity?
- Can two documents be regarded as temporally similar if one contains a small temporal interval of the other document in a detailed way?

4.8 Timelines and User Interfaces

One important use of time entities of a document is to create a sorted list of events, a timeline. A timeline can be shown as a list of vertical textual items or visualized in many different ways. For example, as in Yahoo!'s Correlator[7]. More sophisticated visualizations allow to focus on specific named entities with respect to time like in Yahoo!'s News Explorer [10, 19]. Here, interesting questions are:

- What is the appropriate way to present a timeline?
- Is a linear timeline the only way to present and anchor documents in time?
- How can one leverage document temporal measures to present a good display?
- Are there specific visualizations or user interfaces that can benefit from temporal information?

4.9 Searching in Time

Time entities can also be used to search in documents or log files that can be used to search the past for different purposes such as digital forensics, historical analysis or linguistic analysis. We can even search the future [6, 13], for example, in news for events that are scheduled or may happen in the future. This idea is supported in the Yahoo!s News Explorer tool already mentioned [19]. Microsoft Academic Search[8] is an example of presenting publications and citations in a timeline. Some problems are:

- Besides news, what other sources would one like to use to search in the past and/or the future?
- How far does one need to go back in time?
- Can we improve bibliographic search instead of just sorting by publication date?
- How can we evaluate the quality of such systems?

[7]`http://correlator.sandbox.yahoo.com/`.
[8]`http://academic.research.microsoft.com/`.

4.10 Web Archiving

The goal of Web archiving is to collect and store digital content so that it is accessible for future tasks. Besides the detection of spam, which can be dealt with analyzing the evolution of the link structure of web pages [9], a main challenge in Web archiving is to take care of the temporal coherence of Web pages since it is not possible to collect all pages at the same time. Thus, the content of parts of the collection may change during the crawling process. In [32], Spaniol et al. introduce a coherence framework to overcome the temporal diffusion of the Web crawls, i.e., to minimize the risk of incoherences. Nevertheless, open problems remain:

- How can temporal information be used to predict which pages are likely to change over time?
- How can temporal coherence be achieved for any point in time or time interval?

4.11 Spatio-temporal Information Exploration

Recently, there has been some research on combining spatial and temporal information extracted from documents for exploration tasks [23, 36]. In the same way as temporal information can be normalized using a timeline, spatial information can be normalized according its latitude/longitude information. To extract geographic expressions from documents, so-called geo taggers can be applied. Combining the information extracted from a temporal tagger with the information extracted from a geo tagger allows the exploration of documents according to the events mentioned in the text since events usually happen at some specific time and place. A system for the exploration of such spatio-temporal information from documents is TimeTrails [35]. Some questions are:

- What's the best way to represent maps and time?
- Which kinds of scenarios can benefit from spatio-temporal exploration?

5. CONCLUDING REMARKS

Temporal information embedded in documents in the form of temporal expressions offer an interesting means to further enhance the functionality of current information retrieval applications.

We have presented a number of examples and scenarios where temporal information can be very useful. We have identified research trends in this new area and a number of interesting practical applications as well as problems.

The problems we outline are difficult because they include several areas of computer science, mainly information retrieval, natural language processing, and user interfaces. Moreover, several of them are multidisciplinary because they touch issues related to psychology or design, to mention just two, making them even more challenging.

6. REFERENCES

[1] J. Allan, editor. *Topic Detection and Tracking: Event-based Information Organization*. Kluwer Academic Publishers, Norwell, MA, USA, 2002.

[2] J. Allan, R. Gupta, and V. Khandelwal. Temporal Summaries of New Topics. In *Proceedings of the 24th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '01)*, pages 10–18, 2001.

[3] J. F. Allen. Maintaining Knowledge about Temporal Intervals. In *Communications of the ACM*, 26(11):832–843, 1983.

[4] O. Alonso, R. Baeza-Yates, and M. Gertz. Effectiveness of Temporal Snippets. In *Proceedings of the Workshop on Web Search Result Summarization and Presentation (WSSP 09)*, pages 1–4, 2009.

[5] O. Alonso, M. Gertz, and R. Baeza-Yates. Clustering and Exploring Search Results Using Timeline Constructions. In *Proceedings of the 18th ACM International Conference on Information and Knowledge Management (CIKM '09)*, pages 97–106, 2009.

[6] R. Baeza-Yates. Searching the Future. In *Proceedings of the ACM SIGIR 2005 Workshop on Mathematical/Formal Methods in Information Retrieval (MF/IR 05)*, pages 1–6, 2005.

[7] K. Berberich, S. Bedathur, O. Alonso, and G. Weikum. A Language Modeling Approach for Temporal Information Needs. In *Proceedings of the 32nd European Conference on Information Retrieval Research (ECIR '10)*, pages 13–25, 2010.

[8] N. Chinchor. Overview of MUC-7/MET-2. In *Proceedings of the 7th Conference on Message Understanding (MUC '97)*, pages 1–11, 1997.

[9] Y. Chung, M. Toyoda, and M. Kitsuregawa. A Study of Link Farm Distribution and Evolution Using a Time Series of Web Snapshots. In *Proceedings of the 5th International Workshop on Adversarial Information Retrieval on the Web (AIRWeb '09)*, pages 9–16, 2009.

[10] G. Demartini, M. Missen, R. Blanco, and H. Zaragoza. TAER: Time-aware Entity Retrieval. Exploiting the Past to Find Relevant Entities in News Articles. In *Proceedings of the 19th ACM International Conference on Information and Knowledge Management (CIKM '10)*, pages 1517–1520, 2010.

[11] L. Ferro, L. Gerber, I. Mani, B. Sundheim, and G. Wilson. TIDES - 2005 Standard for the Annotation of Temporal Expressions. 2005. `http://fofoca.mitre.org/annotation_guidelines/2005_timex2_standard_v1.1.pdf`

[12] R. Grishman and B. Sundheim. Design of the MUC-6 Evaluation. In *Proceedings of the 6th Conference on Message Understanding (MUC '95)*, pages 1–11, 1995.

[13] A. Jatowt, K. Kanazawa, S. Oyama, and K. Tanaka. Supporting Analysis of Future-related Information in News Archives and the Web. In *Proceedings of the 9th Joint Conference on Digital Libraries (JCDL '09)*, pages 115–124, 2009.

[14] D. Koen and W. Bender. Time Frames: Temporal Augmentation of the News. In *IBM Systems Journal*, 39(4): 597–616, 2000.

[15] O. Kolomiyets and M.-F. Moens. Meeting TempEval-2: Shallow Approach for Temporal Tagger. In *Proceedings of the Workshop on Semantic Evaluations: Recent Achievements and Future Directions (SEW '09)*, pages 52–57, 2009.

[16] H. Llorens, E. Saquete, and B. Navarro. TIPSEM (English and Spanish): Evaluating CRFs and Semantic Roles in TempEval-2. In *Proceedings of the 5th International Workshop on Semantic Evaluation (SemEval '10)*, pages 284–291, 2010.

[17] J. Makkonen and H. Ahonen-Myka. Utilizing Temporal Information in Topic Detection and Tracking. In *Proceedings of 7th European Conference on Research and Advanced Technology for Digital Libraries (ECDL '03)*, pages 393–404, 2003.

[18] J. Makkonen, H. Ahonen-Myka, and M. Salmenkivi. Topic Detection and Tracking with Spatio-temporal Evidence. In *Proceedings of the 25th European Conference on Information Retrieval Research (ECIR '03)*, pages 251–265, 2003.

[19] M. Matthews, P. Tolchinsky, R. Blanco, J. Atserias, P. Mika, and H. Zaragoza. Searching Through Time in the New York Times. In *Proceedings of the Fourth Workshop on Human-Computer Interaction and Information Retrieval (HCIR '10)*, pages 41–44, 2010.

[20] I. Mani, J. Pustejovsky, and B. Sundheim. Introduction to the Special Issue on Temporal Information Processing. In *ACM Transactions on Asian Language Information Processing* 3(1): 1–10, 2004.

[21] I. Mani, J. Pustejovsky, and R. Gaizauskas, editors. *The Language of Time*. Oxford University Press, New York, NY, USA, 2005.

[22] I. Mani and G. Wilson. Robust Temporal Processing of News. In *Proceedings of the 38th Annual Meeting on Association for Computational Linguistics (ACL '00)*, pages 69–76, 2000.

[23] B. Martins, H. Manguinhas, and J. Borbinha. Extracting and Exploring the Geo-temporal Semantics of Textual Resources. In *Proceedings of the IEEE International Conference on Semantic Computing (ICSC '08)*, pages 1–9, 2008.

[24] P. Mazur and R. Dale. The DANTE Temporal Expression Tagger. In *Proceedings of the 3rd Language and Technology Conference (LTC '09)*, pages 245–257, 2009.

[25] P. Mazur and R. Dale. WikiWars: A New Corpus for Research on Temporal Expressions. In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing (EMNLP '10)*, pages 913–922, 2010.

[26] M. Pasca. Towards Temporal Web Search. In *Proceedings of the 2008 ACM Symposium on Applied Computing (SAC '08)*, pages 1117–1121, 2008.

[27] J. Pustejovsky, J. M. Castaño, R. Ingria, R. Sauri, R. J. Gaizauskas, A. Setzer, G. Katz, and D. R. Radev. TimeML: Robust Specification of Event and Temporal Expressions in Text. In *Proceedings of the AAAI Spring Symposium on New Directions in Question Answering*, pages 28–34, 2003.

[28] J. Pustejovsky, P. Hanks, R. Sauri, A. See, R. Gaizauskas, A. Setzer, D. Radev, B. Sundheim, D. Day, L. Ferro, and M. Lazo. The TIMEBANK

Corpus. In *Proceedings of Corpus Linguistics Conference*, pages 647–656, 2003.

[29] A. Qamra, B. Tseng, and E. Chang. Mining Blog Stories Using Community-based and Temporal Clustering. In *Proceedings of the 15th ACM International Conference on Information and Knowledge Management (CIKM '06)*, pages 58–67, 2006.

[30] S. Schockaert1, D. Ahn, M. De Cock, and E. Kerre. Question Answering with Imperfect Temporal Information. In *Proceedings of the 7th Conference on Flexible Query Answering Systems (FQAS 06)*, pages 647-658, 2006.

[31] F. Schilder and C. Habel. From Temporal Expressions to Temporal Information: Semantic Tagging of News Messages. In *Proceedings of the Workshop on Temporal and Spatial Information Processing (TASIP '01)*, pages 65–72, 2001.

[32] M. Spaniol, D. Denev, A. Mazeika, G. Weikum, and P. Senellart. Data Quality in Web Archiving. In *Proceedings of the 3rd Workshop on Information Credibility on the Web (WICOW '09)*, pages 19–26, 2009.

[33] B. Shaparenko, R. Caruana, J. Gehrke, and T. Joachims. Identifying Temporal Patterns and Key Players in Document Collections. In *Proceedings of the IEEE ICDM Workshop on Temporal Data Mining: Algorithms, Theory and Applications (TDM '05)*, pages 165–174, 2005.

[34] J. Strötgen and M. Gertz. HeidelTime: High Quality Rule-based Extraction and Normalization of Temporal Expressions. In *Proceedings of the 5th International Workshop on Semantic Evaluation (SemEval '10)*, pages 321-324, 2010.

[35] J. Strötgen and M. Gertz. TimeTrails: A System for Exploring Spatio-Temporal Information in Documents. In *Proceedings of the 36th International Conference on Very Large Data Bases (VLDB '10)*, pages 1569–1572, 2010.

[36] J. Strötgen, M. Gertz, and P. Popov. Extraction and Exploration of Spatio-temporal Information in Documents. In *Proceedings of the 6th Workshop on Geographic Information Retrieval (GIR '10)*, pages 1–8, 2010.

[37] R. Swan and J. Allan. TimeMine: Visualizing Automatically Constructed Timelines. In *Proceedings of the 23rd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '00)*, page 393, 2000.

[38] N. UzZaman and J. Allen. TRIPS and TRIOS System for TempEval-2: Extracting Temporal Information from Text. In *Proceedings of the 5th International Workshop on Semantic Evaluation (SemEval '10)*, pages 276–283, 2010.

[39] M. Verhagen, R. Gaizauskas, F. Schilder, M. Hepple, G. Katz, and J. Pustejovsky. SemEval-2007 Task 15: TempEval Temporal Relation Identification. In *Proceedings of the 4th International Workshop on Semantic Evaluation (SemEval '07)*, pages 75–80, 2007.

[40] M. Verhagen and J. Pustejovsky. Temporal Processing with the TARSQI Toolkit. In *Proceedings of the 22nd International Conference on Computational Linguistics (Coling '08)*, pages 189–192, 2008.

[41] M. Verhagen, R. Sauri, T. Caselli, and J. Pustejovsky. SemEval-2010 Task 13: TempEval-2. In *Proceedings of the 5th International Workshop on Semantic Evaluation (SemEval '10)*, pages 57–62, 2010.

[42] M. T. Vicente-Díez, J. Moreno-Schneider, and P. Martínez. UC3M System: Determining the Extent, Type and Value of Time Expressions in TempEval-2. In *Proceedings of the 5th International Workshop on Semantic Evaluation (SemEval '10)*, pages 329–332, 2010.

[43] G. Weikum and M. Theobald From Information to Knowledge: Harvesting Entities and Relationships from Web Sources. In *Proceedings of the 29th ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems of Data (PODS' 10)*, pages 65–76, 2010.