

Task-Oriented Extraction of Temporal Information: The Case of Clinical Narratives

Rob Gaizauskas Henk Harkema* Mark Hepple Andrea Setzer
University of Sheffield
Sheffield, UK
{robertg, henk, hepple, andrea}@dcs.shef.ac.uk

Abstract

Most recent work on temporal relation extraction from text has addressed text drawn from the newswire domain and has attempted to extract all temporal relational information, as specified by proposed temporal annotation schemes such as TimeML. In this paper we explore the task of extracting restricted amounts of temporal information in support of an information extraction application in the medical domain, specifically that of extracting information about times of clinical investigations (X-rays, ultrasounds, etc.) from clinic letters. We describe the task, the corpus and evaluation data we have assembled, a baseline algorithm for extracting temporal relations between temporal expressions and clinical investigation events, and present evaluation results for the algorithm. Overall scores of precision 73.83% and recall 58.70% are promising for a simple baseline approach and suggest that extracting only a restricted subset of the temporal information available in a text may be a sensible way to proceed in the context of specific applications.

1 Introduction

While the task of extracting *all* the temporal information from a text is a worthwhile long term objective for natural language analysis, there may be many cases where extracting *limited* information is of specific utility. For example, for information extraction applications in specific domains it may be the case that temporal information is required only for certain classes of events, and that effort may be concentrated on these cases. In this paper, we consider the case of extracting information for the class of *investigation* events, such as scans, x-rays, and so forth, in clinical narratives. We first discuss the application setting for our work and the

framework for temporal analysis we adopt, as well as previous work on automatically extracting temporal information that has been carried out within this framework. We then describe the specific task we address, which involves relating investigation events mentioned in clinical narratives to dates (including both document dates and dates given by time expressions), before describing the data we have developed to aid system development and evaluation, which consists of real clinical narratives that have been annotated to a level that accords with the intended task. We next describe a simple ‘baseline’ algorithm for extracting this restricted level of temporal relational information, and present results evaluating the performance of this approach. This algorithm is seen as a simple base upon which more elaborate and effective algorithms can be built, which exploit linguistic cues and lexical/domain knowledge in assigning temporal relations, rather than simplistic default heuristics. Whilst the reported evaluation is too limited to draw firm conclusions, the results are sufficiently promising to lend support for the viability of the programme we pursue. We go on to describe some potential enhancements of the baseline algorithm that we intend to implement and evaluate in our on-going work.

2 Information Extraction for Biomedical Research and Clinical Care

2.1 CLEF Services

Our work on temporal information extraction is carried out in the context of the Clinical e-Science Framework Services (CLEF-S) project [12].¹ This project aims to establish methodologies and a technical infrastructure for assembling and managing repositories of clinical patient data, for the purposes of biomedical research and clinical care.

An important aspect of building patient data repositories is information capture. Much of the key information

*Now at Cognia: henk.harkema@cognia.com

¹See www.clinical-esience.org

regarding patient treatment resides in textual, narrative patient notes dictated by doctors, such as case notes, lab reports, and discharge summaries. Although the final clinical diagnosis of a patient may be represented within the structured (i.e. non-narrative) information of an electronic patient record (EPR) – as well as repeatedly in the text of letters written between members of the health care team – much valuable clinical information remains locked in the narratives, such as earlier provisional diagnoses, when a relapse of a disease occurred, what symptoms the patient experienced, when treatment was changed and why, and why investigations were ordered. The unstructured format of the narratives and their volume make it difficult to survey even a single patient’s complete record; it is practically infeasible to aggregate over the records of groups of patients of the size required to carry out clinical research.

To address this ‘information capture’ barrier, CLEF employs Information Extraction (IE) techniques to identify medically relevant classes of entities and events mentioned in documents, as well as significant relations that hold between these entities and events. In particular, the IE component contains rules and resources to identify entities and events such as *drugs*, *problems* (i.e. symptoms, diseases, diagnoses), *loci* (i.e. anatomical locations, physiological functions, mental processes), *investigations* (e.g. X-ray) and *interventions* (e.g. surgery). The relationships extracted include *location* (i.e. that a problem is located in or affects a particular locus), *finding* (i.e. that an investigation reveals the existence of a particular problem), and *target* (i.e. that an investigation or intervention targets a particular locus). The extracted entities, events and relationships are structured in a pre-defined format and imported into the patient data repository, where the information is accessible for further computational processing to support, for example, the generation of patient summaries or the formulation of answers to clinical queries.

2.2 Patient Chronicle

Within CLEF it is proposed that the information available in a patient’s clinical record as stored in the repository, both from the structured data and narrative texts, should be integrated into a *patient chronicle*. The chronicle is a coherent overview of the significant events in the patient’s medical history, i.e. covering their condition, diagnosis and treatment over the period of care. Such chronicles have the potential to be helpful in regard to both clinical care and research. For the former, for example, a patient’s chronicle might be used to generate a textual summary of the key aspects of the patient’s history to be read by a clinician who is newly involved in the patient’s care. For the latter, consider that many of the questions for which a clinical researcher might seek to find answers from a large database

of patient records are ones that require not just aggregation over multiple patients, but which are fundamentally stated in terms that relate to the time-course of patients’ conditions, treatments and outcomes. For example, a clinical researcher who has a hypothesis in mind about some significant medical effect, might look for initial indications of the correctness of the hypothesis by asking questions such as “How many patients with Stage II adenocarcinoma who were treated with tamoxifen had tumour recurrence within five years?” or “For all patients with cancer of the pancreas, compare the percentage alive at five years for those who had a course of gemcitabine with those who didn’t.”

The structured data component of a patient’s clinical record will cover all or most of the noteworthy medical events occurring during a patient’s clinical history, such as major diagnoses, the initiation and discontinuation of drug treatments, and investigations such as X-rays, together with associated information, e.g. the body region that was X-rayed. These events will all be clearly time-stamped in the structured data, allowing them to be readily mapped onto the time-line of the patient chronicle, in effect providing a solid “backbone” for the chronicle.

As discussed in section 2.1, there is additional valuable information to be found in narrative documents that will not be found in the structured data. This information must be extracted and linked into the chronicle ‘backbone’ that is provided by the structured data. Since many important medical events in the course of a patient’s treatment are mentioned in multiple documents and most documents will only include partial descriptions of these events, information extraction over the set of these documents will produce a collection of potentially fragmented and duplicated descriptions of medical entities and events. To integrate the additional information extracted from narratives into the patient chronicle, these various bits and pieces have to be merged and assigned to intervals or time points on the time-line of the chronicle. This process involves extracting temporal information about events from the narratives, and using this and other information to map the events extracted from the narratives onto their corresponding, time-stamped, events in the structured data wherever possible.

3 A Framework for Temporal Analysis

Recent years have seen an increased interest in temporally-aware Natural Language Processing. We will here discuss two important developments in the area, which are, firstly, the specification of the TimeML temporal annotation scheme, and secondly, the release of the TIMEBANK corpus, consisting of text annotated in accordance with the TimeML scheme.

TimeML is an XML-based mark-up language for annotating temporal and event expressions and their relation-

ships in natural language [11].² Commonly in language, events are introduced by verbs. In our particular application area, however, most of the events are introduced by nominals, as we will see later in the paper. In TimeML, event expressions are marked up using EVENT tags. TimeML allows for a subdivision of events into classes, although we will make no use of such distinctions in this paper. The TIMEX3 tag is used to mark up explicit temporal expressions of different types such as times, dates, durations, etc. Besides a *type* attribute, TIMEX3 expressions also have a *value* attribute, which records the value of the temporal expression according to ISO 8601, the International Standard for the representation of dates and times. For example, the value of the temporal expression *July 16, 1997* is “1997-07-16.” TimeML provides links of various kinds that serve to anchor events to temporal expressions and to order events in text relative to one another. We will make use of TLINKS to represent temporal relationships holding between events and times. TLINKS have a *relType* attribute, making explicit the type of relationship holding between the linked entities, e.g. whether an event occurs before or after a given time. TimeML also permits the annotation of textual elements signalling temporal relationships, e.g. temporal prepositions such as *before* and *after*. We will not use these signals in this paper.

TIMEBANK is a human-annotated corpus marked up according to the TimeML annotation scheme [10].³ The corpus contains a wide variety of texts from the news domain with careful, detailed annotations of terms denoting events, temporal expressions and temporal signals, and links between them denoting temporal relations. It is a valuable empirical resource for research into the way that temporal information is organised and conveyed within text. It also provides a *gold standard* for measuring the performance of systems that attempt to identify temporal information automatically in text from the newswire domain. The current version of the TIMEBANK corpus includes 186 newswire articles, containing around 8300 annotated events, 1400 temporal expressions, 6000 links, and 2100 annotated signals.

4 Related Work

A number of approaches have been taken to the automatic identification of temporal relations, usually in systems that simultaneously recognise events and time expressions. The automatic TLINK annotation, the focus of this paper, is handled differently in each case. For example, the TARSQI system [13] uses GUTenLINK (developed at Georgetown University) for TLINK annotation. GUTenLINK itself uses hand-developed syntactic and lexical rules [13]. In contrast, the approach of Boguraev and Ando [1]

combines the use of finite state grammars with machine learning techniques. Both these approaches aim at annotating the full extent of TimeML mark up, as exemplified in TIMEBANK, i.e. they seek to annotate and evaluate all event and time expressions in a text and to establish all temporal relations between them.

Another example of temporal information extraction is provided by the Carsim system [8], which automatically converts narratives about traffic accidents into animated 3D scenes. Animation requires a temporal ordering to be established over the events recognised in the narratives. Temporal relations between events are identified using a hybrid system which comprises a statistical component, based on decision trees, and a small set of hand-written heuristics.

The temporal dimension of medical information is crucially important to its correct use and interpretation. Consequently, work on representing, querying, and reasoning about time-oriented data is an active area of research within medical informatics. Much of the work in this area focuses on the specification and validation of formal frameworks for modeling temporal aspects of medical information (e.g. [7, 2, 3]), rather than on methods for the automatic extraction of temporal information from medical narratives, and so does not bear directly on the specific task described in this paper. Earlier pioneering work on the use of information extraction within the medical domain, such as [5], addresses the extraction of factual medical information, having no particular concern with temporal information.

5 Task and Data

5.1 Task

As noted above, this paper focusses on the automatic extraction of temporal information from narrative texts. The task of identifying *all* events and time expressions, and *all* the salient temporal relations that obtain amongst them — which is essentially the task undertaken in work that uses TimeBank for its gold standard — is very difficult. For our purposes, within the CLEF-S project, we want to identify a more restricted annotation which is sufficient to enable the process of chronicle construction, but which can more effectively be achieved automatically. Furthermore, it is hoped that this restriction of the task will make it possible to use methods that might otherwise be infeasible, e.g. the use of manually authored rules that are geared to the specific linguistic properties of lexical items that play an important role within the domain.

Our restriction of the task has a number of aspects. Firstly, we limit the set of events considered to be only those significant to CLEF’s clinical domain, e.g. events such as investigations and therapeutic interventions. In the preliminary work reported in this paper, we consider only investi-

²See also www.timeml.org

³See www.timeml.org/site/timebank/timebank.html

gations. We allow the smallest temporal unit that is handled within the system to be the (calendar) day. Since we further assume that any investigation takes place within a single day, it follows that such events can be maximally temporally resolved by assigning them to a single date. We restrict our attention to assigning TLINKs only between (relevant) events and dates, including both dates given by time expressions appearing in the text (within the same sentence as the event), and also, as a special case, the date of the document. This means that there is no explicit identification of temporal relations *between* events. Furthermore, we use a more restricted set of temporal relations (or TLINK types) than is used in TimeML, including only the relations of *before*, *after* and *is_included*. Knowing whether an investigation took place before or after a certain date, or is included within it, in particular the date of the letter that mentions the investigation, is very helpful in mapping an investigation extracted from a narrative onto the corresponding investigation in the structured data (cf. section 2.2). For example, if a letter mentions an X-ray investigation that occurred before the date of the letter, then we can immediately eliminate all X-rays in the structured data that are dated after the letter as potential targets in resolving the textual X-ray mention.

5.2 Corpus and Evaluation Data

Our data set consists of more than 332,000 textual clinical documents for almost 37,000 cancer patients. These documents include general case notes, clinic letters, lab reports, and discharge summaries. Apart from a simple pseudonymization step in which the names of patients and medical personnel have been overwritten, these are truthful manual transcriptions of the notes as dictated by doctors. For our current task we are using clinic letters. These letters are accounts of consultations between doctors and patients. They describe the general medical situation of the patient, and include discussions of the results of recent investigations such as X-rays and ultrasounds, often in comparison to results of earlier investigations, and the scheduling of future investigations.

To facilitate our work on temporal information extraction, we have created a small corpus of patient letters which have been annotated to a level consistent with the task requirements as described in the preceding subsection. In particular, the patient letters (narrative data) for a number of patients were annotated, concentrating on CLEF investigations, temporal expressions and the temporal relations holding between them. This work was done by a combination of manual and automatic methods. CLEF investigations were annotated automatically, using a simple gazeteer list. Temporal expressions were annotated and evaluated by the GUTime tagger, developed at Georgetown University, which annotates in accordance with the TimeML TIMEX3

standard⁴ and also covers a variety of temporal modifiers and European date formats.⁵ After these automatic steps, we manually annotated the temporal relations holding between CLEF investigations and temporal expressions appearing in the same sentence, and between investigations and the date of the letter.

	All	Document date	Local TE
all TLINKs	201	149	52
before	73	67	6
after	68	64	4
is_included	60	18	42

Table 1. The distribution of TLINKs in the evaluation corpus

The final evaluation corpus contains the patient letters of 5 patients, numbering 252 patient letters in total, of which 98 contain CLEF investigations. These 98 documents contained 159 CLEF investigations and 605 temporal expressions, amongst which 201 TLINKs were asserted during manual annotation. Table 1 shows a breakdown of these TLINKs according to the temporal relation they represent (i.e. *before*, *after*, or *is_included*). The figures are shown firstly for all TLINKs together and then subdivided into the cases where the TLINKs are between a CLEF investigation and the date of the document and those where the link is between a CLEF investigation and a date given by a temporal expression that appears within the same sentence. Observe that for the case of local temporal expressions the temporal relation is largely *is_included*. Further, for the case of links to document date *before* and *after* relations occur in almost equal proportion, while the *is_included* relations occur significantly less frequently.

6 Evaluating a Baseline Approach to Extracting Restricted Temporal Information

Sentences in patient documents provide various kinds of temporal information, ranging from explicit temporal expressions denoting a specific date to tense and aspect features of verbs. Sentences mentioning an event in association with an absolute expression of time, e.g. *She had a mastectomy on 23/5/89*, typically indicate a link of some appropriate type between the event and the time denoted by the temporal expression. Some sentences contain temporal expressions which must be interpreted relative to the date

⁴www.timeml.org

⁵GUTime extends the capabilities of the TempEx tagger [9], developed at MITRE, which annotates according to the ACE TIMEX2 standard (see timex2.mitre.org).

of the letter, as in, for example, *She underwent an X-ray last month* and *Her recent X-ray shows no further complications*. For sentences which lack explicit temporal expressions, a temporal relation between an event and the date of the letter may be inferred based on the tense and aspect of the verb to which the event is grammatically related. Thus, for example, for the sentence *The tumour was removed successfully*, there will be a link of the type *before* between the “removal” event and the date of the letter in which this sentence occurs.

In the light of these observations, we propose a relatively simple, heuristically-based algorithm for assigning temporal links amongst the temporal entities found within narratives. The algorithm uses only *intra-sentential* evidence in assigning these temporal links — including links that are posited between events occurring within sentences and the global “document date” of the letter. Whilst such an approach is unrealistically simple, it permits a base system to be implemented, to which various possible enhancements may be applied and empirically evaluated. A number of such enhancements are outlined later in the paper.

In what follows, we present the baseline algorithm, and sketch an implementation that realises this approach, before presenting some evaluation results produced using this implementation. Note that in this evaluation, our focus is on the determination of temporal relations, i.e. the assignment of TLINKs, and how well this task can be achieved. So that we can address this question separately from the issue of how well event and time expressions can be recognised, we have chosen to provide the component that assigns TLINKs with *idealised input*, which is derived from the gold standard data, in which the event and TIMEX3 annotations are retained, i.e. so that these are, in effect, ‘perfectly’ recognised. We believe that this approach will better allow us to investigate how different mechanisms may contribute to effective TLINK assignment, without confusing effects from imperfect recognition of events/times.

6.1 A Baseline Algorithm

The algorithm is shown in pseudo-code as Algorithm 1 and proceeds as follows. We assume that all time and event expressions (event nominals and verbs) that are deemed relevant for the TLINK task are part of the input (see section 5.2), and also that time expressions have been evaluated, and normalised to their ISO values. For each relevant event E, the algorithm seeks time expressions within the same sentence, and (if found) asserts a temporal link between E and the nearest of them, with a default type *is_included* (this default is motivated by the distributional properties of the data observable in Table 1). If the time has an ISO value, this allows a temporal link to be inferred and asserted between E and the date of the document (‘DoD’). If no link

```

DoD ← Date of document D;
foreach sentence S in D do
  foreach event E in S do
    if S contains a time expression T then
      Add a TLINK L between E and the closest time
      expression T' to E in S with TLINK_TYPE(L)
      = "is_included";
    if T has an ISO value then
      Add a TLINK L between E and DoD;
      if T < DoD then
        | TLINK_TYPE(L) = "before";
      else
        if T > DoD then
          | TLINK_TYPE(L) = "after";
        else
          | TLINK_TYPE(L) =
          "is_included";
    if a TLINK L does not already exist between E and
    DoD then
      if E is an event nominal then
        Tense ← Tense of nearest tensed verb to E
        in S;
        Aspect ← Aspect of nearest tensed verb to
        E in S;
      else
        Tense ← Tense of E;
        Aspect ← Aspect of E;
      Add a TLINK L between E and DoD;
      if Tense = "future" then
        | TLINK_TYPE(L) = "after";
      else
        if Tense = "past" or (Tense = "present"
        and Aspect = "perfective") then
          | TLINK_TYPE(L) = "before";
        else
          | TLINK_TYPE(L) = "unknown";

```

Algorithm 1: A simple heuristic algorithm for assigning temporal relations

	All		Doc. date		Local TE	
	P	R	P	R	P	R
all	72.83	58.70	74.07	53.69	70.37	73.07
before	70.37	52.05	70.37	56.71	0	0
after	84.29	42.64	84.29	45.31	0	0
is_included	68.00	85.00	65.00	72.22	70.37	90.47

Table 2. Precision and recall figures in %

between E and the DoD is generated this way (either because no time expressions were present within the sentence, or the linked time expression has no ISO value), then the algorithm attempts to infer the temporal relation of E and the DoD based on the tense and aspect information within the sentence — either that of the verb corresponding to the event, or for event nominals, that of the closest tensed verb.

6.2 The Implementation

The system is constructed as a pipeline of different modules, which are mostly assembled from pre-existing components used in previous IE research at Sheffield, and assembled together within the GATE [4] environment.

The preprocessing module comprises tokenisation, sentence splitting, POS tagging and morphological analysis. In a complete run-time system, this phase would also include components for recognising and handling event and time expressions, but for the experiments reported here, as noted above, we have ‘idealised’ recognition of these expressions in the current implementation, drawing annotations from the gold standard corpus.

The SUPPLE module is a robust parser, which can return *partial*, i.e. fragmentary, parses of sentences where necessary. Its output is a predicate-argument representation, or quasi-logical form (QLF), for each sentence or fragment [6]. Note that SUPPLE is included in the implementation in part for secondary technical reasons (it produces the right form of input for the next module), but mainly with a view to future work on *extensions* of the algorithm for assigning temporal relations, i.e. since the baseline algorithm does not rely on either the syntactic or semantic output of the parser.

The Discourse Interpreter is a component that integrates the QLF representations of multiple sentences into a discourse representation, and in which processes such as coreference resolution are performed. The internal formalism of this module allows for the specification of inference rules that apply over QLF representations. It is in terms of such rules that the baseline algorithm for assigning temporal relations has been implemented.

The Annotation Writer is a component for outputting the results of information extraction, commonly in the form of

IE templates. For the current implementation, it outputs a record of the TLINKs that have been assigned within the Discourse Interpreter for subsequent evaluation against the gold standard corpus.

6.3 Results and Discussion

Table 2 shows the results of applying the baseline approach to the evaluation corpus, given as precision and recall scores for TLINK identification. It is important to bear in mind that these results are based on idealised input, i.e. input in which the event and time expressions over which temporal relations are inferred are exactly as in the evaluation corpus. Even so, the overall results seem reasonably promising, given the simplistic nature of the baseline algorithm, and we may hope that these results can be significantly improved upon through appropriate enhancements of the basic method.

The results in this table are broken down in a similar fashion to Table 1, i.e. in terms of the particular temporal relation that TLINKs represent, and in terms of whether events are related to the date of the document or to dates given by temporal expressions within the narrative.

In the rightmost columns of Table 2 the precision and recall scores for the *before* and *after* relations are 0 because of the the default heuristic used in the baseline algorithm, by which the *is_included* relation is always asserted between CLEF investigation events and time expressions in the same sentence (recall that this heuristic was adopted because of the high proportion of cases where it would be valid). The recall is not 100% because in some cases there is more than one temporal expression in the same sentence and the algorithm always asserts, sometimes incorrectly, the *is_included* relation to the one *closest* to the CLEF investigation event.

In the central columns for TLINKS between CLEF investigations and the DoD, the figures reflect cases where (1) the TLINK is asserted indirectly as a consequence of a TLINK being established between a CLEF investigation and a temporal expression in the same sentence which is interpretable with respect to the DoD and (2) there is no (ISO-interpretable) temporal expression in the event-bearing sen-

tence and the TLINK is asserted based on tense and aspect information on the event verb, or the nearest tensed verb in case the event is nominalised. Unfortunately, without more detailed data for how the algorithm is performing in these different cases, which is not at present available, no further conclusions may be drawn.

The precision and recall figures in the leftmost columns for all TLINKS are simply the microaveraged precision and recall figures from the DoD and Local TE columns to their right.

7 Future Work

The work we have reported in this paper is at a preliminary stage. Further work will include more extensive evaluation, failure analysis, and the exploration of various refinements to the algorithm. The refinements we envisage can be divided into two classes: generic enhancements and domain-specific enhancements. The generic enhancements take advantage of the fact that our approach uses a syntactic parser and also a discourse interpreter that can perform inter-sentential event coreference resolution. The domain-specific enhancements involve modelling the characteristics of specific lexical items that play an important role within the domain. Some specific cases are as follows.

1. The current algorithm uses the closest tensed verb to relate a nominalised event to the date of the document (DoD). The enhanced algorithm could make use of the information provided by the parser to identify the verb to which the nominalised event is grammatically linked, i.e. of which the event is either the logical subject or object. Using the closest tensed verb could still be a fallback rule for cases where no such relation can be found.
2. Similarly, instead of relating an event to the closest time, the algorithm could link the event to a time to which it is syntactically related, perhaps with a preference for certain syntactic patterns over others.
3. Instead of assigning the `is_included` relationship by default when relating an event to a time, we could take advantage of temporal adverbials and prepositions to infer the appropriate temporal relation.
4. Preliminary analysis of our evaluation corpus suggests that the tense of a verb does not allow the relationship between the related event and the date of the document to be reliably inferred. For example *I have arranged an x-ray* does not place the x-ray before the DoD but after. The enhanced algorithm will include specialised rules for dealing with the cases for which this non-standard behaviour is observed.
5. Intersentential event coreference could be used to propagate temporal relation information between multiple mentions of the same event.

Also in future work, we will look at applying the above approach, in both its baseline and enhanced versions, to texts from the TIMEBANK corpus, to see how performance compares against that in the restricted CLEF domain.

8 Conclusion

We have proposed that a restricted version of temporal information extraction, which limits its attention to a subset of the temporal information that concerns a restricted set of domain-relevant events, may provide an adequate basis for fulfilling particular domain relevant tasks. We have illustrated this viewpoint in relation to the task of constructing patient chronicles from clinical records in the CLEF-Services project. We have proposed a baseline algorithm for assigning temporal relations in this restricted domain, for which evaluation results are provided, and described a number of enhancements to this baseline approach that will be explored in our on-going research. In the end, the correctness of our main proposal, for the utility of restricted task-oriented temporal information extraction, will only be determined empirically. It may ultimately prove necessary for *all* the temporal relational information within a text to be extracted for any subpart of it to be adequately extracted.

Acknowledgements

This research has been done as part of the CLEF-Services project sponsored by the UK Medical Research Council and has been carried out in collaboration with University of Manchester, University College London, the Open University, the University of Cambridge, the Royal Marsden Hospital and the London Institute of Genetic Medicine. The authors would like to thank Ian Roberts for programming help in generating the figures in the data tables.

References

- [1] B. Boguraev and R. Ando. TimeML-Compliant Text Analysis for Temporal Reasoning. In *Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence (IJCAI)*, Edinburgh, Scotland, August 2005.
- [2] C. Combi and G. Pozzi. HMAP – a temporal data model managing intervals with different granularities and indeterminacy from natural language sentences. *VLDB Journal*, 9, 2001.
- [3] C. Combi and Y. Shahar. Temporal reasoning and temporal data maintenance in medicine: issues and challenges. *Computers in Biology and Medicine*, 27(5), 1997.

- [4] H. Cunningham, D. Maynard, K. Bontcheva, and V. Tablan. GATE: A framework and graphical development environment for robust NLP tools and applications. In *Proceedings of the 40th Anniversary Meeting of the Association for Computational Linguistics*, 2002.
- [5] C. Friedman, P. Alderson, J. Austin, J. Cimino, and S. Johnson. A General Natural-Language Textprocessor for Clinical Radiology. *Journal of the American Medical Informatics Association*, 1(2):161–174, 1994.
- [6] R. Gaizauskas, M. Hepple, H. Saggion, M. A. Greenwood, and K. Humphreys. SUPPLE: A practical parser for natural language engineering applications. In *Proceedings of the 9th International Workshop on Parsing Technologies IWPT2005*, Vancouver, 2005.
- [7] G. Hripcsak, L. Zhou, S. Parsons, A. K. Das, and S. B. Johnson. Modeling electronic discharge summaries as a simple temporal constraint satisfaction problem. *Journal of the American Medical Informatics Association*, 12(1), 2005.
- [8] R. Johansson, A. Berglund, M. Danielsson, and P. Nugues. Automatic Text-to-Scene Conversion in the Traffic Accident Domain. In *Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence (IJCAI)*, Edinburgh, Scotland, August 2005.
- [9] I. Mani and G. Wilson. Robust Temporal Processing of News. In *Proceedings of the 38th Annual Meeting of the Association for Computational Linguistics (ACL2000)*, Hong Kong, October 2000.
- [10] J. Pustejovsky, D. Day, L. Ferro, R. Gaizauskas, P. Hanks, M. Lazo, , R. Saurí, A. See, A. Setzer, and B. Sundheim. The TIMEBANK Corpus. In *Proceedings of Corpus Linguistics 2003*, Lancaster, March 2003.
- [11] J. Pustejovsky, B. Ingria, R. Saurí, J. Castaño, J. Littman, R. Gaizauskas, G. K. A. Setzer, and I. Mani. The specification language TimeML. In I. Mani, J. Pustejovsky, and R. Gaizauskas, editors, *The Language of Time: A Reader*. Oxford University Press, Oxford, 2005.
- [12] A. Rector, A. Taweel, J. Rogers, D. Ingram, D. Kalra, R. Gaizauskas, M. Hepple, J. Milan, R. Powers, D. Scott, and P. Singleton. Joining up Health and BioInformatics: e-Science meets e-Health. In *Proceedings of the 3rd UK e-Science All Hands Meeting*, Nottingham, UK, 2004.
- [13] M. Verhagen, I. Mani, R. Sauri, J. Littman, R. Knippen, S. Jang, A. Rumshisky, J. Phillips, and P. Pustejovsky. Automating Temporal Annotation with TARSQI. In *Proceedings of the ACL*, Ann Arbor, June 2005.