Towards the Application of TimeML in Clinical Guidelines

Reinhardt Wenzina, Katharina Kaiser

Abstract

The modeling of computer-interpretable clinical guidelines is a laborious task. In this project we propose a two-step approach to explore whether the TimeML specification language can build a solid foundation to enable an automatic modeling process by temporal reasoning. The comparison of the main linguistic differences between clinical guidelines and clinical narratives leads to the conclusion that existing research results of the latter are only partially applicable. Finally, the goals, methods, and expected results – which should ease the modeling task in the long run – are discussed.

1 Introduction

Clinical practice guidelines (CPGs) (Field and Lohr 1990) are documents containing recommendations that describe the appropriate care for the management of patients with a specific clinical condition. In order to deploy them in some kind of computerized tool (e.g., a clinical decision support system) they have to be represented in a specialized language (e.g., Asbru, PROforma, GLIF (Clercq et al. 2004)). Several editing/authoring tools exist with these languages, but still the authoring task remains complex and labor-intensive and requires comprehensive knowledge in medicine and computer science. For this reason, various approaches have been developed to deal with an automated modeling using natural language processing (NLP) and information extraction methods (Kaiser et al. 2007, 2011; Wenzina and Kaiser 2013). Thereby, rules based on linguistic and semantic information are developed to identify activities in the text, link them with conditions controlling them, and connect and order activities based on some keywords or their hierarchical position in the document. Major challenges in modeling – whether manually or (semi-)automatically – that have to be tackled, are that documents are often long and confusing, concepts are vaguely or incompletely described, and the text contains redundancies that have to be identified.

Other attempts using NLP in the medical domain focus on the processing of documents such as clinical narratives. Although these documents differ from CPGs in the nature of the text, the methods applied also have to deal with underspecified expressions. Several projects focusing on these aspects are using TimeML (Pustejovsky et al. 2004) for annotation of events and temporal expressions (Mele and Sorgente 2011; UzZaman and Allen 2010a). TimeML has become the de-facto standard for annotation of events and temporal expressions in natural language and can be used for reasoning on events. As CPGs describe activities (corresponding to events) and contain (temporal) expressions for ordering and relating them, we propose applying TimeML for annotation. Having a CPG annotated in TimeML could ease the further modeling in a computer-interpretable guideline's (CIG) language by providing a semi-structured document containing additional information that can then be transformed and authored by hand or even automatically. In this paper we want to describe and discuss our considerations about the challenges of modeling a CPG and the dealing with temporal relations.

In the following section we motivate the use of TimeML for annotation and temporal reasoning on documents such as CPGs. In section 3 we give a short overview of TimeML followed by a detailed problem description in section 4. Section 5 describes the approach we are intending to apply, whereas section 6 describe the goals we want to achieve. Section 7 concludes our ideas.

2 Motivation

As temporal aspects play an important part within clinical guidelines (Terenziani et al. 2008) (e.g., the description of care-paths) the use of temporal reasoning methods based on a temporal representation of the guideline may lead towards an automatic modeling approach. To start with, a temporal representation of a document consists of temporal expressions, concept primitives, and temporal relations (Sun et al. 2013b) in order to handle the various vague and/or complex temporal dimensions. The following sentence of a guideline gives a first insight into this temporal complexity: "If post-dinner or fasting glucose values achieve the threshold for insulin treatment as a consequence of increasing the amount of carbohydrates, starting insulin would be the best option".

In recent years, the specification language TimeML for annotating events and temporal expressions in natural language processing gained recognition. The language addresses the problems of (a) identifying an event (for example in the sentence above *"achieve"* or *"increase"*) and anchoring it in time, (b) chronologically arranging events (e.g., the temporal relation between *"achieve"* and *"increase"* is expressed by the phrase *"as a consequence"*) , (c) reasoning with contextually underspecified temporal expressions (temporal functions such as *"post-dinner"* or *"fasting"*), and (d) reasoning about the duration of events (Pustejovsky et al. 2004). Therefore, TimeML may also fulfill our requirements for annotating the various information aspects of clinical guidelines.

Consequently, we hypothesize that

Reasoning over temporal annotations – specified by TimeML and annotated by adaptable tools (e.g., TARSQI Toolkit (Verhagen and Pustejovsky 2010)) – can support the automatic transformation process of a guideline into its computer interpretable representation.

This support comprises the extension of authoring tools (to identify relevant sentences) as well as the automatic generation of parts of the CIG model in the long run.

3 Background and Related Work

The markup language TimeML is a specification formalism for annotating events and temporal relations in narrative texts and is widely spread in the natural language processing community. The fact that a revised and interoperable version of TimeML called ISO-TimeML (Pustejovsky et al. 2010) is expected to be published as an international standard for temporal annotations by ISO, demonstrates its importance. Generally, TimeML defines events as situations that *happen* or *occur* and categorizes the relations among them as temporal, subordinated, or aspectual (Pustejovsky et al. 2004). Originally, TimeML was applied to the TimeBank corpus containing 186 news articles and was in the following extended to other domains.

Domain independent additions to the TimeML language were proposed by UzZaman and Allen (2010a). They implemented the TRIOS system containing a semantic parser to extract events, their linguistic features, and relations. OntoTimeFL, a formalism for the reasoning of complex events – also built on TimeML – categorizes the events as narrative, intentional, and causal (Mele and Sorgente 2011). During the THYME[1] project a temporal relation annotation scheme and annotation guidelines for clinical free texts based on TimeML were developed.

The TARSQI Toolkit (TTK) was one of the first implementations which generated TimeML-compliant annotations in order to enable temporally based questions about events in news articles (Verhagen and Pustejovsky 2010). In MED-TTK – an extension of the toolkit – the TTK's time tagger was modified in order to improve the identification methods of temporal references in medical notes (Reeves et al. 2013). The latest version of the toolkit concentrates on, among other extensions, the application to the medical domain and the introduction of narrative containers - with the implementation still in progress (Verhagen and Pustejovsky 2012).

Gooch (2012) used external resources (e.g., the UMLS – the Unified Medical Language System (Lindberg et al. 1993)) in order to categorize selected medical concepts as events to be formalised in TimeML expressions. The solution was finally evaluated against a corpus of clinical discharge summaries.

Methods in natural language processing support the modeling process of clinical guidelines. A rule-based information extraction approach in order to identify medical activities in CPGs was presented by Kaiser et al. (2011). A set of semantic patterns representing activities based on semantic relations allowed the detection of large parts of control flow related aspects. The extraction of temporal relations between the activities was not part of the method. This shortcoming was partly resolved by Thorne et al. (2013). They focused on a supervised approach in order to recognize process fragments in guidelines and succeeded in the identification of simple before/after temporal relations between medical tasks. However, the modeling of the control flow still is an open issue.

[1] https://clear.colorado.edu/TemporalWiki

4 Problem Description

TimeML has already been successfully applied to medical texts (e.g., clinical narratives). However, clinical practice guidelines (CPGs) differ in many ways (major differences are discussed below) and therefore, existing research results can only partially be transferred.

A clinical narrative describes a specific clinical case/the disease and patient level from the therapist's view. It also acts as a tool for self-reflection of current practice and supports the discussion of a case among colleagues (UW Health 2014). CPGs, on the other hand, are defined as "systematically developed statements to assist practitioners and patient decisions about appropriate healthcare for specific circumstances" (Field and Lohr 1990). Hence, they have no relation to an existing patient. They serve as manuals for patient-groups and care personnel, such as physicians or nurses.

In natural language processing clinical narratives are difficult to handle because (1) a third of the text consists of sentence fragments which can hardly be identified by conventional language parsers (Savova et al. 2009), (2) medical terms are often abbreviated, and (3) the style of the language is fairly different to domain independent English texts (Sun et al. 2013b). On the contrary, CPGs are published as linguistic mature documents containing highly sophisticated and complex sentence constructions.

Clinical narratives describe the progression of illnesses and their related events chronologically. These events, such as laboratory tests, doctor's visits, administered procedures, and records from medication logs are explicitly dated (Sun et al. 2013b). Based on these absolute time points, the events can be temporally related to each other and temporal reasoning mechanisms can be applied. Clinical guidelines, however, mostly specify only relative time points (e.g., first trimester of pregnancy) which additionally are often expressed in a vague way.

One possibility to derive temporal relations between events in narrative texts is to analyze the tense used in the sentence. This allows for example, the determination if a medication is historically used, currently taken, or prescribed for future use in clinical narratives. (Raghavan et al. 2012) used the tense of the verb as a feature (amongst others) for classifying medical events into coarse time-bins by means of Conditional Random Fields. The following sentence from a guideline, "Administration of inhaled steroids at or above 400 mcg a day of BDP or equivalent may be associated with systemic side-effects", shows that the tense of the modal verb does not show any hint that the administration has to occur before the side-effects may appear. Hence, the tense of the verb used is of limited suitability for identifying temporal relations.

5 Approach

The general objective of this project is to explore whether temporal annotations, as specified in TimeML, facilitate the manual as well as the automatic translation process of CPGs into their computer interpretable models. Therefore, we follow a strategy based on sentence-level as well as document-level perspectives. This *'think big, start small'* strategy allows us to approach this challenge from two different views and presumably helps to prove our hypothesis.

In the following subsections the two perspectives are discussed in detail, whereas in section 6 we will show the corresponding methods and results.

5.1 Sentence-Level Perspective: Subordination Links

TimeML specifies different types of temporal relations (TLINK, SLINK, and ALINK) representing the various dependencies that exist between temporal concept primitives (Pustejovsky et al. 2004). The subordination link (SLINK) is used for contexts with a subordination relation between two *events* (Saurí et al. 2006b). The relation is classified according to their types *modal, factive, counterfactive, evidential, negative evidential, or conditional*. In CPGs such subordination relations represent a prominent aspect of the control flow and are often found in sentences expressing conditional activities (corresponds to SLINK type *conditional*) in contrast to clinical narratives, where mainly facts are reported. As no detailed research has been done on subordination links in the medical domain so far, we will set our focus onto this specific aspect.

A few examples of subordination links from clinical guidelines[2] are shown below. The introducing *event* is marked "in bold" and the consequence, which takes the place of the subordinated *event*, is "underlined".

(1) *If there is no **response** the drug should be <u>discontinued</u>.*
(2) ***Women with pain** should be <u>re-examined</u> after two hours.*
(3) *The partogram should be <u>used</u> once labour is **established**.*

Subordination links may be found all over the guideline document – not only in the descriptions of care paths but e.g., also for presenting background information (see sentence number three). Hence, they have to be filtered concerning their control flow relatedness. A closer categorization of the involved *events* may help to distinguish these.

5.2 Document-Level Perspective: The General Approach

If the TimeML specification is applied to a complete guideline document, every contained temporal concept (e.g., EVENT, TIMEX3, TLINK, etc.) will be annotated. The challenge now is to use, adapt, or extend these annotations to make it possible to generate parts of the formal representation of a CPG automatically. In order to reduce the complexity of this approach, the output of previous

[2] The events are identified compliant to the guideline presented in (Sun et al. 2013a).

projects, especially the application of TimeML on clinical narratives, will be reused. Nevertheless, several challenges, which are discussed below, have to be mastered before doing so.

Temporal reasoning is based on the anchoring of events to relative or absolute points of time. However, the identification of these temporal anchors in guidelines is a challenging task. An example: The document creation time (DCT) which is used as an anchor in the newswire domain has not really an informative value for clinical guidelines - exceptionally perhaps for the practitioner but surely not for the sequence of activities in clinical care paths. A possible solution to this problem could be the introduction of the concept of Narrative Time (Pustejovsky and Stubbs 2011). It describes the current temporal anchor for events in a guideline and changes during the reading process. Miller et al. (2013) already discussed this concept for clinical narratives.

In CPGs various information dimensions containing *events* can be identified (e.g., *control, data, time, evidence, resource, patient,* etc.) but not all of them are temporally related to each other. These *events* are not positioned in time or in relation to other temporally located *events* in the document (e.g., *'In clinical practice, women often require 1600 to 2200 kcal per day.'*). The handling of such *Generics* (UzZaman and Allen 2010b) must be defined.

The fact that clinical guidelines not only contain patient and treatment related information but also explanatory notes like: *'The diagnostic criteria for gestational diabetes have been a matter of debate since the publication of the IADPSG recommendations in 2010.'* it is worth investigating if the various extensions (e.g., Sun et al. 2013a) to the TimeML specification formalism are consistently applicable to CPGs.

In guidelines, clinical activities are often recommended to different patient groups (e.g., specified by the age) whereas in clinical narratives the described facts and procedures are always related to one and the same patient. The sentence: *'...the dose of inhaled steroids should be increased to 800 micrograms/day in adults or 400 micrograms day in children (5-12 years), if not already on these doses.'* shows a complex example concerning this matter. As you can see from the sentence above, the patient specifications in this case are even mixed within the same sentence. In order to represent this information in a formal model, a new concept like a patient-container should be introduced.

6 Goals

The long-term goal of this project is the design of an annotation scheme (for temporal relations based on TimeML for clinical practice guidelines) and built upon, the development of methods to automatically generate formal models of CPGs. Based on the two perspectives presented above, this general goal has to be split into two subgoals.

6.1 *Goal I: Automatic Identification of Subordination Links*

Method. In our experiments we set the focus on the discovering of subordination links (SLINK). Therefore, we (1) identify condition-action sentences concerning control-flow related aspects as candidates for subordination links in a guideline which we have already used in a former project, (2) manually annotate the selected sentences complying to the TimeML specification for SLINKs and find the limitations of the formalism in our application area, (3) extend/adapt the rules of the TimeML annotation guideline in this respect, (4) analyze existing application frameworks which automatically generate TimeML annotations and allow the extension of their rule-sets to those elaborated in step four, (5) manually annotate an existing corpus of condition-action sentences according to step three, apply the prototype of step four, and evaluate the results in order to draw a general conclusion about our algorithm.

Constraints. SLINK annotations depend on the proper annotation of events. Hence, we act on the assumption that temporal expressions (TIMEX3) and *events* (based on the annotation guideline developed in Sun et al. 2013a) have been annotated correctly.

Expected results. These are (1) a prototype with the implemented rules, (2) an evaluation of the rules after applying them on the selected corpus, and (3) a general conclusion to continue our work on the long-term goal.

Evaluation. A comparison of the identified subordination links with a gold standard (developed in step 5) facilitates the calculation of quantitative measures like recall, precision, and accuracy.

Current status. Step (1) of our method presented above has already been finished. The selection of a functioning framework to be able to focus on the implementation of the rule set and not on the coding process is a crucial task and therefore, we rescheduled our steps and continued with step four. The requirement specification for a tool (described in step four) comprises (1) the extraction and normalization of temporal expressions, (2) the identification of concept primitives, and (3) the classification of temporal relations in narrative texts (Sun et al. 2013b). We selected the TARSQI-Toolkit (TTK)[3] for our project as it fulfills these requirements and all TTK modules already use the TimeML annotation language. Additionally, its stable performance is comparable with other tools (e.g. TRIPS (UzZaman and Allen 2010b)).

Figure 1 shows the structure of the toolkit (Verhagen and Pustejovsky 2010). The Slinket module is responsible for the correct identification of subordination links and relies on a combination of lexical and syntactic knowledge (Saurí et al. 2006a). Hence, it will be extended to our rule set.

Currently we are working on step (2).

[3] We used the current version 1.0.

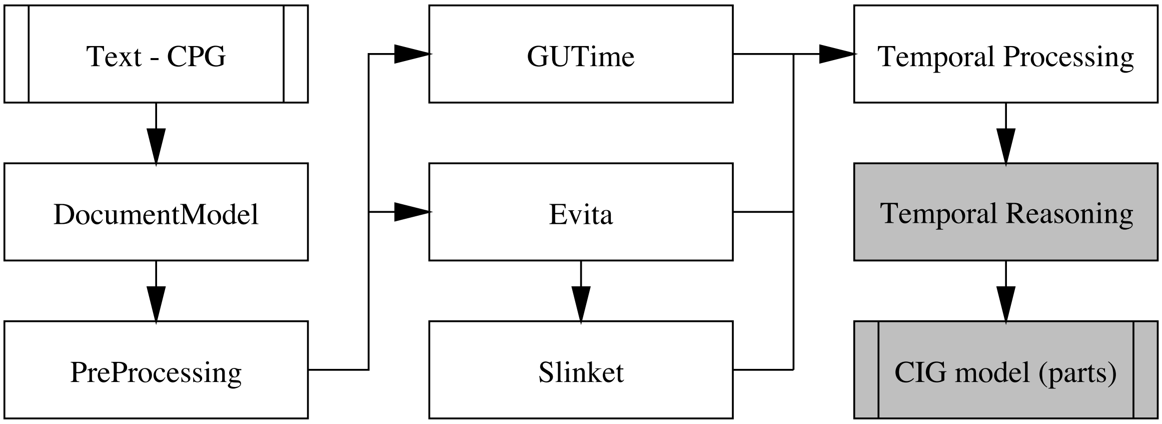

Fig. 1: Structure of the TTK (*GUTime*: temporal expression tagger; *Evita*: event recognition tool; *Slinket*: subordination link categorizer; *Temporal Processing*: responsible for temporal relations) and future extensions (colored in gray).

6.2 *Goal II: Extension of the TimeML Specification Language for the Automatic Generation of CIGs*

Method. We will (1) manually annotate a guideline based on TimeML and its extensions (Sun et al. 2013a), (2) apply a temporal reasoning mechanism to explore which temporal relations can be inferred for the formal guideline model, (3) identify the concept primitives modeled with TimeML which can be directly transferred to the computerized guideline model (e.g., ASBRU (Shahar et al. 1998)), and (4) find out which extensions for the TimeML specification are needed to successfully realize steps two and three.

Answers to the following research questions are essential for realizing step four:

1. Does the introduction of new types of temporal links help to identify the different temporal dimensions in a formal model of a guideline?
2. Does the concept of the narrative scope solve the problem of temporal anchoring?
3. How can the different views of the stakeholders (patient, physician) of a guideline be modeled in the specification language?
4. Which influence does the structure of the document (e.g., title, headlines) have to the temporal relations?
5. How can generic events be handled?
6. Which external resources (e.g., UMLS Semantic Network) can support the automatic annotation process of medical concepts?

Expected results. We want to (1) draw a general conclusion concerning the suitability of TimeML for our goal, and in case of a positive answer (2) extend the TimeML specification, (3) prepare an annotation guideline for upcoming projects, and (4) implement a prototype based on existing frameworks (see Figure 1) to show the feasibility of our results.

Evaluation. As an ASBRU model is represented as an XML document, we use an XML distance measure (e.g., Long et al. 2005) to show the similarity between the inferred model and the gold standard.

Current status. Has not been started yet.

7 Conclusion

The aim of this research-in-progress paper was to present a first step towards the application of TimeML to clinical practice guidelines in order to support their translation process into their corresponding formal models. After discussing the main differences between clinical narratives and CPGs, we found out that existing research results are only partially applicable. Subsequently, we gave an overview of our research tasks (goals, methods, expected results) and discussed the current state of our work. If it is possible to prove our hypothesis, the modeling of CPGs will be simplified substantially.

Acknowledgement. This research was carried out as part of project no. TRP71- N23 funded by the Austrian Science Fund (FWF) and of the MobiGuide project partially funded by the European Commission under the 7th Framework Program, grant #287811.

References

de Clercq P. A., Blom J. A., Korsten H. H. M., Hasman A. (May 2004) METHODOLOGICAL REVIEW: Approaches for Creating Computer-interpretable Guidelines That Facilitate Decision Support. In: Artif. Intell. Med. 31(1), pp. 1–27 ISSN: 0933-3657 DOI: 10.1016/j.artmed.2004.02.003 http://dx.doi.org/10.1016/j.artmed.2004.02.003

Field M. J., Lohr K. N. (1990) Clinical Practice Guidelines:Directions for a New Program. The National Academies Press ISBN: 9780309076661 http://www.nap.edu/openbook.php?record_id=1626

Gooch P. (2012) A lightweight, pattern-based approach to identification and formalisation of TimeML expressions in clinical narratives. In: http://openaccess.city.ac.uk/2412/

Kaiser K., Akkaya C., Miksch S. (2007) How Can Information Extraction Ease Formalizing Treatment Processes in Clinical Practice Guidelines? A Method and Its Evaluation. In: Artificial Intelligence in Medicine 39(2), pp. 151–163

Kaiser K., Seyfang A., Miksch S. (2011) Identifying Treatment Activities for Modelling Computer-Interpretable Clinical Practice Guidelines. In: Lecture Notes in Computer Science. Springer Berlin Heidelberg, pp. 114–125 ISBN: 978-3-642-18049-1 DOI: 10.1007/978-3-642-18050-7_9

Lindberg D., Humphreys B., McCray A. (1993) The Unified Medical Language System. In: Methods of Information in Medicine 32(4), pp. 281–291

Long J., Schwartz D. G., Stoecklin S. (2005) An XML Distance Measure. In: Proceedings of The 2005 International Conference on Data Mining (DMIN'05, pp. 119–125

Mele F., Sorgente A. (2011) A Formalism for Temporal Annotation and Reasoning of Complex Events in Natural Language.. In: Vol. 771. CEUR-WS.org http://dblp.uni-trier.de/db/conf/aiia/dart2011.html#MeleS11

Miller T., Bethard S., Dligach D., Pradhan S., Lin C., Savova G. (2013) Discovering Temporal Narrative Containers in Clinical Text. In: Sofia, Bulgaria http://www.aclweb.org/anthology/W13-1903

Pustejovsky J., Stubbs A. (2011) Increasing Informativeness in Temporal Annotation.. In: Linguistic Annotation Workshop. The Association for Computer Linguistics, pp. 152–160 ISBN: 978-1-932432-93-0 http://dblp.uni-trier.de/db/conf/acllaw/law2011.html#PustejovskyS11

Pustejovsky J., CastaÒo J. M., Ingria R., Saurí R., Gaizauskas R. J., Setzer A., Katz G., Radev D. R. (2nd Mar. 2004) TimeML: Robust Specification of Event and Temporal Expressions in Text.. In: AAAI Press, pp. 28–34 ISBN: 1-57735-184-3 http://dblp.uni-trier.de/db/conf/ndqa/ndqa2003.html#PustejovskyCISGSKR03

Pustejovsky J., Lee K., Bunt H., Romary L. (2010) ISO-TimeML: An International Standard for Semantic Annotation. In: European Language Resources Association (ELRA), Valletta, Malta ISBN: 2-9517408-6-7

Raghavan P., Fosler-Lussier E., Lai A. M. (2012) Temporal Classification of Medical Events. In: BioNLP '12. Montreal, Canada, pp. 29–37 http://dl.acm.org/citation.cfm?id=2391123.2391128

Reeves R. M., Ong F. R., Matheny M. E., Denny J. C., Aronsky D., Gobbel G. T., Montella D., Speroff T., Brown S. H. (2013) Detecting temporal expressions in medical narratives.. In: I. J. Medical Informatics 82(2), pp. 118–127 http://dblp.uni-trier.de/db/journals/ijmi/ijmi82.html#ReevesOMDAGMSB13

Saurí R., Verhagen M., Pustejovsky J. (2006a) SlinkET: A partial modal parser for events. In: In Language Resources and Evaluation Conference, LREC 2006

Saurí R., Littman J., Gaizauskas R., Setzer A., Pustejovsky J. (2006b) TimeML Annotation Guidelines, Version 1.2.1

Savova G., Bethard S., Styler W., Martin J., Palmer M., Masanz J., Ward W. (2009) Towards temporal relation discovery from the clinical narrative.. In: AMIA Annual Symposium proceedings 2009, pp. 568–572 ISSN: 1942-597X http://view.ncbi.nlm.nih.gov/pubmed/20351919

Shahar Y., Miksch S., Johnson P. (1998) The ASGAARD Project: A Task-Specific Framework for the Application and Critiquing of Time-Oriented Clinical Guidelines. In: Artificial Intelligence in Medicine, pp. 29–51

Sun W., Rumshisky A., Uzuner O. (2013a) Annotating temporal information in clinical narratives. In: Journal of Biomedical Informatics 46, Supplement(0) 2012 i2b2 {NLP} Challenge on Temporal Relations in Clinical Data, S5 –S12 ISSN: 1532-0464 DOI: http://dx.doi.org/10.1016/j.jbi.2013.07.004 http://www.sciencedirect.com/science/article/pii/S1532046413001032

Sun W., Rumshisky A., Uzuner O. (Sept. 2013b) Temporal reasoning over clinical text: the state of the art.. In: Journal of the American Medical Informatics Association : JAMIA 20(5), pp. 814–819 ISSN: 1527-974X DOI: 10.1136/amiajnl-2013-001760 http://dx.doi.org/10.1136/amiajnl-2013-001760

Terenziani P., German E., Shahar Y. (2008) The temporal aspects of clinical guidelines. In: Stud Health Technol Inform 139, pp. 81–100

Thorne C., Cardillo E., Montali M., Eccher C., Calvanese D. (2013) Process Fragment Recognition in Clinical Documents. In: Proceedings of the 13th Conference of the Italian Association for Artificial Intelligence (AI*IA 2013)

UW Health (2014) Clinical Narrative. http://www.uwhealth.org/advancement-recognition/clinical-narrative/20484

UzZaman N., Allen J. (2010a) TRIOS-TimeBank Corpus: Extended TimeBank Corpus with Help of Deep Understanding of Text. In: European Language Resources Association (ELRA), Valletta, Malta ISBN: 2-9517408-6-7

UzZaman N., Allen J. F. (2010b) Extracting Events and Temporal Expressions from Text.. In: ICSC. IEEE, pp. 1–8 ISBN: 978-1-4244-7912-2 http://dblp.uni-trier.de/db/conf/semco/icsc2010.html#UzZamanA10

Verhagen M., Pustejovsky J. (4th June 2010) Temporal Processing with the TARSQI Toolkit.. In: COLING (Demos), pp. 189–192 http://dblp.uni-trier.de/db/conf/coling/coling2008d.html#VerhagenP08

Verhagen M., Pustejovsky J. (2012) The TARSQI Toolkit.. In: LREC. European Language Resources Association (ELRA), pp. 2043–2048 ISBN: 978-2-9517408-7-7 http://dblp.uni-trier.de/db/conf/lrec/lrec2012.html#VerhagenP12

Wenzina R., Kaiser K. (2013) Identifying Condition-Action Sentences Using a Heuristic-based Information Extraction Method. In: Proceedings of the Joint International Workshop: KR4HC'13+ProHealth'13, pp. 17–29