Temporal information extraction for temporal question answering

Frank Schilder & Christopher Habel

Department for Informatics (AB WSV), University of Hamburg
Vogt-Kölln-Str. 30, D-22527 Hamburg
{schilder,habel}@informatik.uni-hamburg.de

Abstract

This paper outlines work-in-progress regarding temporal information extraction, and highlights issues to be discussed by future research in temporal question answering. Based on Schilder & Habel (2001), we first present a semantic tagging system for temporal expressions and discuss how the temporal information conveyed by these expressions can be extracted. We then discuss extensions to the temporal tagger currently being implemented.

1. Introduction

Temporal question answering requires the extraction of temporal information encoded in natural language text. The information to be extracted includes first of all explicit temporal information such as dates and time expressions. Additionally and more importantly expressions that denote events also have to be extracted and correctly annotated.
Work aiming at the development of temporal taggers that extract the time denoted by expressions such as *01.01.2003* or *last Thursday* has been carried out over the last couple of years. The next major challenge is the extraction of event information and the combination of this information with the temporal information supplied by a temporal tagger. A temporal question answering system needs to have access to the event information and the temporal relations between them. Such a system not only has to deal with When-questions, but also with relative temporal questions such as "*what happened before the attack*? "
This paper starts with the description of a semantic tagging system that extracts temporal information from news messages (Schilder & Habel, 2001). This temporal tagging system marks temporal expressions that are defined as chunks of text that express some sort of direct or inferred temporal information. We assume the following prerequisites for a successful temporal question answering task:

- derivation of the meaning of the temporal expressions
- identification of the event
- anchoring of the event (with respect to a time line and/or a different event)

In this position paper we focus on the second and the third prerequisite. The remainder of the paper is structured as follows: First, we briefly describe the system presented in Schilder & Habel (2001). Then, we raise several questions regarding further extensions. Finally, we describe current research and future directions.

2. Temporal and event information extraction

The set of the temporal expressions tagged by our tagger, described in (Schilder and Habel, 2001), includes dates (e.g. 08.04.2001), prepositional phrases (PPs) containing some time expression (e.g. *on Friday*), and verbs referring to a situation (e.g. *opened*). Related work by Mani and Wilson (2000) focuses only on the core temporal expressions and neglects the temporal information conveyed by prepositions (e.g. *Friday* vs. *by Friday*).
The main part of the system is a temporal expression tagger that employs finite state transducers based on hand-written rules. The tagger was trained on economic news articles obtained from two German newspapers and an online news agency (*Financial Times Deutschland*, *die tageszeitung* and the online news service *comdirect*).
Since we focus on a particular text domain (i.e. news articles), the classification of temporal expressions can be kept to a manageable set of classes

2.1 Classification of temporal expressions

The main distinction we make is between time-denoting and event-denoting expressions. The first group comprises chunks expressing temporal information that can be stated with reference to a calendar or clock system. Syntactically speaking, these expressions are mainly expressed by prepositional, adverbial or noun phrases (e.g. *on Friday* or *today,* or *the fourth quarter*). The second group, event-denoting expressions, refers to events. These expressions have an implicit temporal dimension, since all situations possess a temporal component. For these expressions, however, there is no direct or indirect link to the calendar or clock system. These expressions are verb or noun phrases (e.g. *increased* or *the election*).

Time-denoting expressions.

Temporal reference can be expressed in three different ways:

Explicit reference. Date expressions such as *08.04.2001* refer explicitly to entries of a calendar system. Also, time expressions such as *3 p.m.* or *Midnight* denote a precise moment in our temporal representation system.

Indexical reference. All temporal expressions that can only be evaluated via a given index time are called indexical. Expressions such as *today*, *by last week* or *next Saturday* need to be evaluated with respect to the article's time stamp.

Vague reference. Some temporal expressions express only vague temporal information and it is rather difficult to precisely place the information expressed on a time line. Expressions such as *in several weeks*, *in the evening* or *by Saturday the latest* cannot be represented by points or exact intervals in time.

For the given domain of news articles, the extraction of a time stamp for the given article is very important. This time stamp represents the production time of the news information and is used by the other temporal expressions as an index time to compute the correct temporal meaning of the expression. Note that an explicit date expression such as *24.12.* can only be evaluated with respect to the year that the article was written. This means that even an explicit temporal expression can contain some degree of indexicality.

Event-denoting expressions

Two types of event-denoting expressions have to be distinguished, on the one hand, sentences, and, on the other, specific noun phrases. In the former case, the verb is the lexical bearer of information about the event in question, in the latter case, specific nouns, especially those created by nominalization, refer to an event.

Since temporal information is the topic of the system described in this paper, only a subset of event-denoting nouns have to be considered. These expressions - as *election* in the phrase *after the election* - that serve as temporal reference pointers in building the temporal structure of a news article, can be marked by a specific attribute in their lexical entry. Furthermore, in the text classes we have investigated, there is a small number of *event nouns*, which are used as domain dependent pointers to elements of temporal structures. For the domain of business and stock market news, phrases such as *opening of the stock exchange*, *opening bell*, or *the close* are examples of domain specific event expressions.

2.2 Representation of temporal information: the time domain

The primary purpose of the present paper is to anchor the temporal information obtained from natural language expressions in news messages in *absolute time*, i.e. in a linearly ordered set of abstract time-entities, which we call *time-set* in the following. One of the major tasks in this anchoring process is to augment the temporal information in case of indexical and vague temporal descriptions (see section 4.3 for more details). Since these expressions do not specify an individual time-entity of the time-set, it is necessary to add temporal information until the temporal entity build up from natural language is fully specified, i.e. can be anchored in the time-set.

Definition of temporal relations

Temporal relations are explicitly marked by temporal prepositions (e.g. *before*, *on* or *by*). We use the following seven temporal relations: *before, after, incl, at, starts, finishes, excl*. The preposition *on* as in *on Friday*, for instance, denotes the inclusion relation *incl*, whereas the preposition *by* as in *by Friday* is represented as *finishes*.

Note that the seven temporal relations employed by the current version can be characterized by sets of Allen's topological interval relations (Allen, 1983) as described in table 1.[1]

before	*{b, m}*
after	*{bi, mi}*
incl	*{d, s, f, eq}*
at	*{di, si, fi, eq}*
starts	*{s}*
finishes	*{f}*
excl	*{b, bi, m, mi}*

Table 1: The temporal relations used

Using Hobbs' 'Ontology of time', which is suitable as an ontological basis for Allen-like representations of the time domain (cf. Hobbs, 2002) we distinguish two types of temporal entities, *instants* and *intervals*. Common sense, and thus natural language, does not consistently differentiate between *instants* and *intervals*: A day, e.g. 01.01.2002, sometimes called Euro-day or E-day, can from a global long-term economic perspective be seen as an instant, namely the beginning of the Euro-phase of the EU economy. On the other hand, from the perspective of planning the actions needed to get ATMs ready for the public to use, the same day is seen as an interval. In other words, being an instant or

[1] Allen (1983) proposes a temporal reasoning system that contains all 13 conceivable relations between intervals: *b (efore), m (eets), o (verlaps), s (tarts), d (uring), f (inishes),* the 6 inverse relations *bi, mi, oi, si, di* and *fi* and *eq (ual).*

being an interval is a matter of levels of granularity.[2]

2.3 Representation of temporal information: the event domain

The domain of "eventualities" possesses a complex ontological structure, with the subtypes event, process, and state (cf. Bach, 1986; Hobbs 2002). In the present paper we focus on events and processes, those types of eventualities that are most relevant in analyzing economic news articles, i.e. members of that text class which constitutes the application domain of our investigation. (To be consistent with the terminology used in the scientific community of 'information systems' and 'question answering', we use "eventuality" only as a technical term, and prefer "event domain" to "eventuality domain".)

The core temporal property of events is that they are located in time. In other words, to every event *e* exists a temporal entity *t*, in formal notation "*t = event_time(e)*" which can intuitively be described as "the time when *e* happens". In, Hobbs' (2002) ontology of time and events, this relation between events and temporal entities is established by the predicate 'time-span-of'.[3]

Since people as well as machines only seldom know precisely which temporal entity is the *eventuality time* of an event mentioned in a text, an interpretation strategy, which aims to anchor events in the time line, would mostly be without success. Instead, another strategy is chosen, which focuses on characterizing the temporal location of one eventuality relative to another. The projection of events to the time domain given by *event_time* induces for each topological relation between temporal entities a corresponding temporal relation between events. Therefore, we use the same inventory of relations listed in table 1, to express temporal relations between eventualities.

Akin to Setzer and Gaizauskas' (2002) "time-event graph", we propose a two layer representation for time and events, consisting of an event-layer that only holds information about relations between events, and a time-layer that includes temporal entities, especially those induced by *event_time,* temporal relations between those entities, and—if possible—anchoring in the time line, which can be seen as a specific kind of temporal relationship.

A major advantage for in separating the two layers is that they are subject to different types of processes. The time-layer is the representational section, which is dealt with in temporal reasoning and temporal reference processing, e.g. in temporal anaphora resolution. Furthermore, the time-layer mediates in temporal reasoning about events.

On the other hand, the event-layer enables some types of time-induced causal reasoning and can lead to the establishment of causal links: E.g., the information given by the sentence "*Treasury prices rose sharply after the Greenspan speech"* [CNN Money Market, January 11, 2002: 5:29 p.m. ET] should be accessible to when-questions as well as why-questions.[4] Furthermore, the event-layer contains information about mereological relations between events, i.e. about *part_of* relationships. This type of information is also essential for the adequate treatment of subtypes of eventualities, e.g. *events* vs. *processes*, which are fundamental for treating aspectual information (see section 4.2).

3. The temporal tagger

The temporal tagger presented in Schilder & Habel (2001) employed a cascade of Finite State Transducers (FST). The following sections provide a brief introduction to this technique before the overall system architecture of the tagger and some modifications to the original system are explained in more detail.[5]

3.1 Preliminaries

The temporal expression chunks are extracted via an FST. FSTs are basically automata that have transitions labelled with a translation instruction. A label of the form *a:b* indicates such a translation from *a* to *b*. Take, for example, the simple FST in figure 1. If the input contains the sequence of the three subsequent characters *F*, *S*, and *T*, the output produced is these same three characters put into brackets. The input stream "*FSTs are basically automata*" is, for instance, translated, into "*[FST]s are basically automata*".

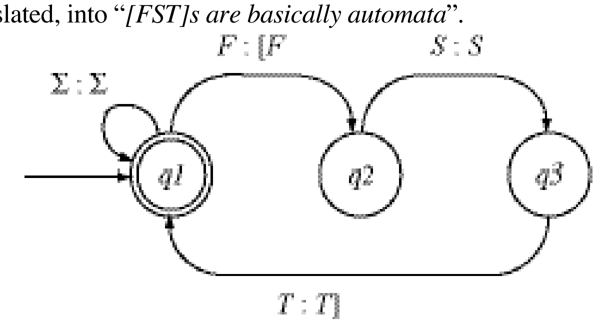

Figure 1: A simple FST

[2] In the present paper we do not go in the details of granular approaches of time and events. Cf. Hobbs (1985) on granularity in general, Bettini, Jajodia & Wang (2000) on time granularity, and Schilder & Habel (2001) on the treatment of granularity in temporal tagging.

[3] Hobbs (2002) uses a predicate instead of a function to avoid problems in treating noncontiguous eventualities. We skip this topic in the present paper.

[4] Note, that we use here "causal" in a common sense way, not based on a systematic theory of causality. More detailed work on causal reasoning and causal reference for information extraction is in progress.

[5] The most recent version of the temporal tagging system is written in SWI-PROLOG 5.0.10.

on Monday *(time-denoting expression)*	<CHUNK id = t43 type = time sem = 20:01:04:02TXX-XX-XX:incl > on Monday </CHUNK>
ftd.de, Fr, 16.3.2001, 11:00 *(document time stamp)*	<CHUNK id = t1 type = time ag = 'FTD' sem = 20:01:03:16T11-00-XX > ftd.de, Fr, 16.3.2001, 11:00 </CHUNK>
said *(event-denoting expression)*	<CHUNK id = e23 type = verbal sem = say> said </CHUNK>

Table 2: Examples of tagged temporal expressions

3.2 The annotation of temporal information

The FSTs defined are fed by the output of a Part of Speech (POS) tagger.[6] The POS tagger specifies the syntactic categories and a lemma for every word of the input text. The syntactic information is then stored in an XML file.[7] Given the derived syntactic categories and the lemma information for every word of the text, several FSTs specialised into different classes of temporal expressions are run.

Temporal Expressions. One FST consisting of 15 states and 61 arcs tags all occurrences of time-denoting temporal expressions. The POS information stored in an XML file as well as a predefined class of temporal lemmas are used by this FST. The class of temporal lemmas used include days of the week (e.g. *Friday*), months (e.g. *April*) as well as general temporal descriptions such as *midday*, *week* or *year*. Since German is a very productive language regarding compound nouns, a simple morphological analyzing tool was integrated into this FST as well. This tool captures expressions such as *Rekordjahr* ('record year') or *Osterferien* ('Easter holiday').

The extracted temporal expression chunks are marked by the `CHUNK` tag and an attribute `type = time`. See the first row of table 2 for an example. Note that the attribute `sem` carries the semantic value of the temporal expressions . The possible values for sem are explained in section 3.3.

Document time stamp. The document time stamp for a given article is crucial for the computation of almost all temporal expressions (e.g. *now*). In particular, this index time is indispensable for the computation of all temporal expressions that express an indexical reference (see the second row of table 2).

Verbal descriptions. Another FST that contains 4 states and 27 arcs marks all verbs as previously tagged by the POS tagger. As already pointed out, these temporal expressions denote an event. The tag for such expressions is `<CHUNK type = verbal> </CHUNK>` (see table 2; third row).

Nominal descriptions. All noun chunks are marked as well. This information will become more important when the aspectual class of the event is determined (cf. Section 4.2). The noun chunks may have an attribute event, if the noun denotes an event, such as *the election*. These nouns are also used to denote events mentioned in the text when combined with time-denoting expressions, as in *after the election in May*.

3.3 System output

While reading the output stream from the FSTs temporal inferences are drawn by the system. In particular, expressions bearing indexical references are resolved and the event descriptions are matched with the time denoting temporal expressions. The values for the attribute `sem` are temporal expressions according to the standard ISO format (i.e. `YYYY-MM-DDTHH:MM:SS`). The granularity level for a given temporal expressions is reflected by the number of Xs the ISO expression contains. The year *2003,* for instance, has the value `2003-XX-XXTXX:XX:XX` which indicates the granularity level year.

[6] A decision-tree-based POS tagger developed by Schmid (1994) was integrated into the system.
[7] Some of the XML and HTML handling predicates the system uses stem from the PiLLoW package developed by Manuel Hermenegildo and Daniel Cabeza (URL www.clip.dia.fi.upm.es/miscdocs/pillow/pillow.html).

```
<chunk id=n53 type=nominal>
        Die neuen Länder
<chunk>
<chunk id=e34 type=verbal sem = werden>
        würden
<chunk>
        neben der von
<chunk id=n54 type=nominal>
        den Schlichtern
<chunk>
<chunk id=e35 type=nominal sem=Tariferhöhung>
  <chunk id=e36 type=nominal sem=vorgeschlagen>
          vorgeschlagenen
  <chunk>
        Tariferhöhung
<chunk>
        zusätzlich von
<chunk id=e36 type=nominal sem=Angleichung>
        der Angleichung
<chunk>
        der Ost- an
<chunk id=n55 type=nominal>
        die Westtarife
<chunk>
<chunk id=t26 type=time sem=2007-XX-XXTXX:XX:XXPY:end:finishes>
        bis Ende 2007
<chunk>
<chunk id=e34 type=verbal sem=belasten>
        belasten
<chunk>

<tempRel rel=finishes id1=et36 id2=t26/>
<eventTime event=e36 time=et23/>
```

```
[Die neuen Länder ] [würden]
neben der von [den Schlichtern ]
[[vorgeschlagenen]
[Tariferhöhung]] zusätzlich von
[der Angleichung] der Ost- an
[die Westtarife] {bis Ende 2007}
[belastet].
```

Figure 2: Temporal and event information annotated

In addition to the ISO expression, periods are explicitly marked. *Two months*, for example, is annotated by `P2M`. Vague expressions such as *afternoon* are added to the ISO expression as well. The semantic value for *Friday afternoon* may be `2003-01-10TXX:XX:XX:afternoon` given a particular index time.[8]

Finally, the temporal relations holding between the events and times expressed by the text are stored as well.

After all expressions have been tagged, an HTML file is produced highlighting the respective expressions. See the annotation of an example sentence in figure.[9]

(1) Die neuen Länder würden neben der von den
The new States would aside from the from the
Schlichtern vorgeschlagenen Tariferhöhung zusätzlich
mediator suggested pay scale rise additionally
von der Angleichung] der Ost- an die Westtarife
from the alignment the East at the West-pay scale
bis Ende 2007 belastet.
until end 2007 burden.

'The pay rise suggested by the mediator would be an extra burden for the new states because of the alignment of pay scales between the eastern to the western part of Germany

[8] Note that we give *afternoon* a vague value, whereas Mani and Wilson (2000) suggest a definite time point (i.e. 16:30).

[9] Time-denoting expressions are indicated by a dark (or magenta) background, while event-denoting expressions are indicated by a very light (or yellow) background for verbal descriptions, while noun event chunks are medium light (or turquoise). Noun chunks that do not denote an event possess a darker (or green) background that is not as dark as the background for time-denoting expressions.

until the end of 2007.'

4. The event tagger

The event tagger extracts event-denoting expressions. The current version of the event tagger only annotates verbs or nouns as event-denoting expressions. However, an event description is actually based on the entire clause. Consider the following minimal pair:

(2) a. The company went bankrupt.

b. The companies went bankrupt.

In (8b), the plural noun affects the aspectual class of the entire event description (see section 4.2). Hence, it seems desirable to annotate the entire clause instead of only the verb. The verb and nouns should also be annotated, if an automatic derivation of aspectual classes should be added for later extensions of the event tagger.

Although it is probably not debatable to annotate verbs as carriers of event information, nouns may also denote events. NPs are normally viewed as references to entities, such as in *Peter* or *chair*. However, there are NPs that refer to an event, such as *election*. A test that distinguishes between event-denoting and entity-denoting nouns involves the combination of temporal modifiers (e.g. *the 2000 election* vs. **the 2000 Peter*).

In the following we describe work on an event tagger we are currently developing. Future extensions, including the automatic classification of events according to their aspectual properties , are described in section 4.2.

4.1 Event-denoting nouns

Within the current work, we focus on the extraction of nouns that can carry temporal information. The starting point is the temporal tagger presented in Schilder and Habel (2001). Based on a domain analysis (financial world), an event ontology is being created. In addition, we are exploring a bootstrapping mechanism in order to extract previously unknown event nouns. We assume that noun chunks adjacent to temporal expressions are most likely modified by these expressions (e.g. *the election 2002*).

4.2. Extracting aspectual information

Verbs can be divided into different aspectual classes (event vs. state). The event ontology being created could also contain further information regarding the aspectual class for every event. Specific domains do have a class of reoccurring events that could be compiled beforehand after doing a more detailed corpus investigation. However, this would be generated by hand and consequently very work intensive.

Instead it is desirable to automatically derive aspectual classes in a robust way. A first attempt at the automatic derivation of aspectual classes can be found in Siegel (1999). He used a set of indicators in order to improve the classification of aspectual classes. Using this as a starting point, we plan to incorporate more theoretical approaches (e.g. Moens and Steedman, 1988) in order to achieve a more theoretical and sound output from the system.

Employing the knowledge of aspectual information of events can be very useful when it comes to the combination of temporal and event information. Several linguistic tests determine the aspectual class of an event (Dowty, 1979). An event, for example, can be used with a time frame interval, such as *within 2 hours*, whereas a state can only be combined with a time duration interval, such as *for 2 hours*.

(3) a. Peter lost $200,000 within/*for 2 hours.

b. Peter was rich *within/for 2 hours.

(3a) is only felicitous with a time frame adverbial, whereas (3b) goes together with a time duration adverbial. Note that (3b) may work with the time frame adverbial if *was* is reinterpreted as *became*.

5. Combining temporal and event information

5.1 Semantics for temporal expressions

With respect to processing temporal information, the crucial distinction between time-denoting and event-denoting expressions is that event-denoting expressions lack the direct link to temporal entities. An event-denoting expression (e.g. a verb) refers to an event of a certain type. The verb *to meet*, for instance, can be formalized as $meet(e_1)$. In order to add the temporal information to the event, a function *temp* is defined that gives back the time when the event occurred (i.e. eventuality time). A time-denoting expression such as *on Monday* that is combined with the event description carries some temporal information that can further specify the run time *temp(e1)* of the event `e1`.

Semantics for temporal prepositions

PPs are the carriers of temporal relations. The semantics for a preposition is, therefore, as follows: *rel(t, e)*. For each preposition a temporal relation *rel* was defined. The preposition *by* expresses, for instance, the `finishes` relation, as in *by Friday*. Temporal expressions that do not contain a preposition are assumed to express an inclusion relation, as in *Die Pflegeversicherung war* ***1995*** *[...] in Kraft getreten* (`the statutory health insurance coverage of nursing care for the infirm took effect in 1995').

Derivation of meaning

The temporal information expressed by a sentence as in example sequence (1) is derived via unification of the semantic attributes derived for the temporal expression chunks.

(4) [Die US-Technologiebörse Nasdaq]$_n$
The US-technology stock market Nasdaq
[hatte]$_v$ {amFreitag} mit [einem Minus]$_n$
had on Friday with a minus
von [3,11 Prozent]$_n$ bei [1782 Punkten]$_v$ [geschlossen]$_v$.
of 3.11 percent at 1782 points closed.
'The Nasdaq closed with a minus of 3.11 percent at 1782 points on Friday'.

Several temporal expressions are marked by the tagger: nominal expressions (e.g. *einem Minus* ('a minus')), verbal descriptions (e.g. *geschlossen* ('closed')) and one temporal relation expressed by *am Montag* ('on Monday') .The first two expressions are entity or event-denoting expressions. The latter expression is a time-denoting expression that consists of a preposition and a time-denoting expression that is stored by the FST. The derivation of the semantics for this expression is done during the tagging process for the temporal expressions (i.e. `2003-01-10TXX:XX:XX:incl`).

First, the preposition *am* ('on'), denoting an inclusion relation between an event and a time, is processed. The expressed temporal relation is represented by a PROLOG list (i.e. `[incl,[E,T]]`). After having processed the subsequent noun referring to a time (i.e. *Friday*), the following semantic representation is obtained via unification: `sem = [incl,[E,t1]]`, where `t1` refers to the time *Friday* refers to.

In the next step, the verbal expression tagger combines the temporal information `t1` with the event representation for *geschlossen*. The following semantic representation is assign to the verb *geschlossen* during the tagging of the verbal expressions: `sem = close etime = [_,[et23,_]]`. This means that event `e23` is of type closing and the eventuality-time `et23` of this event stands in some to-be-specified relation with another expression. Next, the temporal information extracted by the FST specialised in time-denoting expressions is unified with the value of the `temp`-attribute. The result is `[incl,[et23,t1]]`.

So far, only one temporal relation has been determined: that the event of closing happened within a time frame of one day. Since *Freitag* contains an indexical reference, this reference has to be resolved. The document time stamp is needed here. All references regarding this index time are resolved during the generation of the HTML output file. Accordingly, the following time stamp is generated for *am Freitag*: sem = 2003-01-10TXX:XX:XX The `time` information is underspecified because the current granularity level is `GL-day`.

This intermediate PROLOG representation also has to be put into proper XML format. The temporal link between the event `e23` and the time `t1` is encoded as follows:

```
<tempRel rel = incl id1=et23 id2=t1/>
<eventTime event= e23 time=et23/>
```

The anchoring of temporal information is done via the mapping of the temporal expressions to the verb chunk(s) or to an event-denoting noun chunk. The mapping process ensures that temporal relations link the temporal expression with the verbal description in a sentence or clause unless a nominal expression is closer to the temporal expression:

(5) …[eine Lohnerhöhung]$_{event}$ von [2,4 Prozent]$_n$
a pay rise of 2.4 percent
{ab Januar}.
from January onwards.

A relatively high number of noun chunks denoting an event was found in our test corpus. 17.5% of noun chunks denoted an event. And even more importantly, 26.09% of temporal expressions were linked to a nominal[10] event expression and not to a verbal one. Consequently, a current research goal is the extraction of nominal event descriptions.

A first set of indicators for this type of expressions includes the morphological analysis of nouns. The suffix *–ung,* for instance, quite often signals, in German, an event description, as in *Lohnerhöhung* ('pay rise').

Since not every sentence contains a temporal expression, we do not force every event to be put into an order along a time line. Instead, we leave the temporal ordering underspecified unless explicit information can be derived. We believe that this approach avoids the problems with inter-annotator accuracy that one encounters when every event has to be ordered on a time line.

The ultimate goal of our research is to automatically extract event-event and time-event relationships rather than a time line annotated by the events described in the text (cf. Setzer and Gaizauskas, 2002).

5.2 Evaluation

In Schilder & Habel (2001), we evaluated the temporal expression tagger with respect to a small corpus consisting of 10 news articles taken from Financial Times Deutschland. Precision and recall rates regarding the recognition of simple temporal expressions and complex temporal expression phrases were around 90. The accuracy of the derived semantic meaning of the correctly tagged temporal expressions was around 85.

Recently, we tested a new version of the tagger regarding the linking of temporal and event information as described in the previous section.

Tagging results

The improved version of the temporal tagger and the event

[10] Some adjective constructions were also linked to a temporal expression (e.g. die von dem Schlichter *vorgeschlagene* Lohnerhöhung 'the pay rise *suggested* by the mediator').

tagger were evaluated regarding the following four tasks:

1. nominal chunks
2. verbal chunks
3. nominal chunks that denote an event
4. linking of temporal and event information

Table 3 contains the results of the evaluation with respect to the two first tasks. There were a total of 589 nominal and 240 verbal expressions previously annotated.

	Nominal chunks	Verbal chunks
Precision	98.10	95.11
Recall	96.60	93.75

Table 3: Extraction of nominal and verbal chunks

The results for these tasks are comparable to noun and verb chunkers that employ machine learning or statistical methods (cf. Zhang, Damerau & Johnson, 2002; Skut & T. Brants, 1998).

The extraction of nominal event descriptions, however, is not such a simple task. We annotated 101 event-denoting nouns in our corpus. Recall was 44.55 and Precision was 66.18. Since we only used a few morphological and context cues for determining an event description, these results can be taken as a base line for future investigations.

Finally, we checked the linking of temporal and event information. As a base line algorithm to this task we took the approach taken in Schilder & Habel (2001), assigning the temporal meaning to the verbal description in the same clause. Precision for this base line was 63.77 (recall= 98.55).

Using the event tagger, an improvement was observed. The linking task was carried out with a precision of 78.26 (recall=100).

	Event tagger and linking results		
	Event nouns	Event-time linking (base line)	Event-time linking (with event tagger)
Precision	66.18	63.77	78.26
Recall	44.55	98.55	100

Table 4: Performance of the event tagger and the linking task

6. Current and future approaches to detection of temporal and event information

Based on the experience gained with the temporal and the event tagger we would like to address the following questions, all of which seem important for future directions in temporal tagging and temporal question answering.

6.1 Annotating temporal and event information: TimeML

The recently proposed specification language TimeML (Pustejovsky et al., 2002) offers a guideline for the annotation of temporal and event information. Similar to our approach, TimeML keeps the representation of temporal expressions and events separate: temporal expressions are tagged by `TIMEX3`, whereas events are annotated by `EVENT`. In addition, two further tags are used: `SIGNAL` and LINK.

Events are annotated by the <EVENT> tag. Events are described as situations that happen or occur and include states which are situations in which something holds true. Additional information, annotated as attributes, encompasses the event class (e.g. OCCURRENCE or STATE), the tense and the aspect (e.g. PROGRESSIVE or PERFECTIVE).

Temporal information such as *12:00, 24. December 2002* or *3 days* is tagged by <TIMEX3>. The definition of <TIMEX3> is based on the TIMEX specification given by Ferro et al. (2001) and Setzer (2001).

Function words that signal the relation between temporal objects are tagged by <SIGNAL>. In particular temporal prepositions and connectives indicate such a relation (e.g. before, during or when). This tag was first introduced by Setzer (2001). With TimeML, polarity indicators (e.g. no, not) and temporal quantification words (e.g. twice, three times) are also annotated as signals.

The <LINK> tag specifies the exact relation that holds between the temporal elements in a text. The <LINK> tag can either indicate a temporal (i.e. TLINK), a subordination (i.e. SLINK) or an aspectual (i.e. ALINK) relation.

How a sentence is annotated with TimeML is shown in the following example.

(6) John left 2 days before the attack.

The annotation process extracts two events (i.e. left and party), one temporal expression (i.e. 2 days) and one signal word (i.e. before). In addition, two event instances ei1 and ei2 are annotated. This is necessary because verbs and nouns can also have a generic reading as in *Leaving is always sad*. The generic reading does not invoke an event instance.

The links established between these two events is signalled by the temporal preposition *before*. Hence, a temporal relation is stated in the TLINK tag.

The attributes for the tags are not of any importance for the remainder of this paper. The usage of these attributes is explained in more detail in Pustejovsky et al. (2002).

```
John
<EVENT eid="e1" class="OCCURRENCE"
tense="PAST" aspect="PERFECTIVE">
left
</EVENT>
<MAKEINSTANCE eiid="ei1" eventID="e1"/>
<TIMEX3 tid="t1" type="DURATION"
```

```
value="P2D" temporalFunction="false">
2 days
</TIMEX3>
<SIGNAL sid="s1">
before
</SIGNAL>
the
<EVENT eid="e2" class="OCCURRENCE"
tense="NONE" aspect="NONE">
attack
</EVENT>
<MAKEINSTANCE eiid="ei2" eventID="e2"/>
<TLINK eventInstanceID="ei1"
signalID="s1" relatedToEvent="ei2"
relType="BEFORE" magnitude="t1"/>
```

The development of a TimeBank corpus that can serve as a gold standard is very encouraging. Having large corpora annotated with temporal and event data will help to make progress in the field of temporal information extraction. However, will the proposed annotation also be useful for the temporal question answering task? Since only the head of the VP, NP etc[11] is marked as `<EVENT>`, it is questionable how a requested event can be extracted in its entirety. Instead, the clause containing the noun and verb chunks should be tagged.
Another problem we see is the unclear treatment of `<SIGNAL>` tags. According to the annotation guidelines, the important trigger words such as *in*, *before* etc. are not always tagged. In particular, when the trigger words occur in a more complex expression (e.g. *two days before the end* of *May*) the guidelines state that the trigger word should not be tagged.

6.2 Future challenges and possible approaches

Having presented our current work in the extraction of events from news texts and having reviewed recent work in annotation of temporal and event information, we would like to come back to the connection of temporal information extraction and temporal question answering. We showed that temporal and event information can be explicitly encoded in temporal expressions such as *December 31*, *Monday* etc. Recent approaches to temporal tagging focus on these expressions in order to annotate or to automatically extract and to anchor the respective expression in time (Setzer and Gaizauskas, 2001, Schilder and Habel, 2001, Filatova and Hovy, 2001). However, many linguistic expressions other than explicitly temporal expressions have a temporal dimension:

- nouns (e.g. *the election 2002*, *the changeover, the culmination*)
- verbs (e.g. *is about to*, *finished*, *completed*)
- adverbs (e.g. *twice*, *quickly*)
- prepositions (e.g. *on Monday*, *until last year, on board*)

Our current work is focussed on the extraction of nouns that denote an event. An annotation of a small test corpus has highlighted the need for a deeper analysis of the data, because the usage of a finite state transducer that relied on a couple of morphological and context rules was not able to achieve satisfactory precision and recall rates. What other methods could be implemented in to tag these more complex temporal expressions? We suggest that only a combination of a deeper semantic analysis and other robust techniques will help improve the accuracy of the event tagger.
After the successful detection of a temporal expression, the meaning of the expression has to be computed. Assuming that a time stamp for the particular text can be extracted, the values for expressions such as *Monday* can be derived and expressed by an ISO expression. However, the temporal information carried by nouns, for instance, is far more difficult to determine. It would be useful to derive the default event time for every type of eventuality (e.g. the default event time for *election* is 1 day).
Since robust methods that only draw from morphological or syntactic clues seem to be limited with respect to the task at hand, additional knowledge sources must be exploited. Online dictionaries such as WordNet or event ontologies developed for special domains could be useful here (e.g. the eXtensible Business Reporting Language (XBRL)).
Another main challenge for temporal tagging is the ordering of the described events with respect to the extracted temporal information. Current approaches (Setzer and Gaisauskas, 2001, Filatova and Hovy, 2001) aim for the ordering of each event along a time line. Corpora studies and automatic extraction studies, however, have shown that this task is very difficult to achieve.
Our proposal to assume an underspecified temporal relation unless explicit temporal information is given may circumvent the problems observed by these corpora studies. However, it is unclear how much useful temporal information can be extracted so that a temporal question answering system can find the correct answer among the annotated data.

[11] The following syntactical structures can denote an event according to the TimeML annotation guidelines: tensed and untensed verbs, nouns, adjectives, predicative clauses and prepositions.

References

J.F. Allen (1983). Maintaining Knowledge about Temporal Intervals. *Communications of the ACM, 26(1)*. 832-843.

E. Bach (1986). The algebra of events. *Linguistics and Philosophy, 9*. 5–16.

C. Bettini, Jajodia, S. & Wang, S. X. (2000). *Time granularities in databases, data mining, and temporal reasoning*. Berlin: Springer-Verlag.

D.R. Dowty (1979). *Word Meaning and Montague Grammar.* D.Reidel: Dordrecht.

L. Ferro, Mani, I., Sundheim, B. & Wilson, G. (2001). Tides temporal annotation guidelines. Technical Report MTR 01W0000041, MITRE (2001) Version 1.0.2.

E. Filatova and E. Hovy (2001). Assigning Time-Stamps To Event-Clauses. In *Proceedings of the ACL-2001 Workshop on Temporal and Spatial Information Processing*, ACL 2001. Toulouse, France, 6-11 July. (pp. 88-95).

J.R. Hobbs (1985). Granularity. In *Proceedings IJCAI-85, Los Angeles, CA*. (pp. 432–435).

J.R. Hobbs (2002). Toward an ontology of time for the semantic web. In *Proc. of the Workshop "Annotation Standards for Temporal Information in Natural Language", LREC 2002*, 27th May 2002, Las Palmas, Canary Islands - Spain. 28–35.

I. Mani and G. Wilson (2000). Robust temporal processing of news. In *Proceedings of the 38th Annual Meeting of the ACL*, Hong Kong.

M. Moens & M. Steedman (1988). A computational ontology and temporal reference. *Computational Linguistics, 14(2)*.

J. Pustejovsky, Sauri, R., Setzer, A., Ingria, B. (2002). TimeML annotation guidelines. Technical report, TERQAS report (2002) available at http://time2002.org.

A. Setzer (2001). *Temporal Information in Newswire Articles: an Annotation Scheme and Corpus Study*. PhD thesis, University of Shefield.

A. Setzer and R. Gaizauskas (2002). On the Importance of Annotating Temporal Event-Event Relations in Text. In *Proceedings of LREC 2002, Workshop on Annotation Standards for Temporal Information in Natural Language*.

A. Setzer and R. Gaizauskas (2001). A Pilot Study on Annotating Temporal Relations in Text. In *Proceedings of the ACL-2001 Workshop on Temporal and Spatial Information Processing*, ACL 2001, Toulouse, France, 6-11 July.

F. Schilder and Chr. Habel (2001). From Temporal Expressions To Temporal Information: Semantic Tagging Of News Messages. In *Proceedings of the ACL-2001 Workshop on Temporal and Spatial Information Processing*, ACL 2001, Toulouse, France, 6-11 July. pp. 65-72.

H. Schmid (1994). Probabilistic part-of-speech tagging using decision trees. In *Proceedings of International Conference on New Methods in Language Processing*.

E. V. Siegel (1999). Corpus-Based Linguistic Indicators for Aspectual Classification. In *Proceedings of the 37th Annual Meeting of the Association for Computational Linguistics*, University of Maryland, College Park, MD, USA.

W. Skut & T. Brants (1998). A Maximum Entropy partial parser for unrestricted text. In *Proceedings of the 6th ACL Workshop on Very Large Corpora (WVLC),* p. 143-151, Montréal, Canada.

WordNet (2003, 10. January) a lexical database for the English language [WWW document], URL http://www.cogsci.princeton.edu/~wn/

XBRL (2003, 10. January) XBRL the XML based business reporting standard [WWW document], URL http://www.xbrl.org

T. Zhang, F. Damerau & D. Johnson (2002), Text Chunking based on a Generalization of Winnow. In *Journal of Machine Learning* Research, volume 2 (March), 615-637.