Enriching TimeBank: Towards a More Precise Annotation of Temporal Relations in a Text

Volker Gast*, **Lennart Bierkandt***, **Stephan Druskat**†, **Christoph Rzymski***
*Friedrich Schiller University Jena, English Department, Ernst-Abbe-Platz 8, D-07743 Jena
†Humboldt University of Berlin, Department of German Studies and Linguistics, Dorotheenstr. 24, D-10117 Berlin
volker.gast@uni-jena.de, post@lennartbierkandt.de, stephan.druskat@uni-jena.de, christoph.rzymski@uni-jena.de

Abstract

We propose a way of enriching the TimeML annotations of TimeBank by adding information about the Topic Time in terms of Klein (1994). The annotations are partly automatic, partly inferential and partly manual. The corpus was converted into the native format of the annotation software GraphAnno and POS-tagged using the Stanford bidirectional dependency network tagger. On top of each finite verb, a FIN-node with tense information was created, and on top of any FIN-node, a TOPICTIME-node, in accordance with Klein's (1994) treatment of finiteness as the linguistic correlate of the Topic Time. Each TOPICTIME-node is linked to a MAKEINSTANCE-node representing an (instantiated) event in TimeML (Pustejovsky et al., 2005), the markup language used for the annotation of TimeBank. For such links we introduce a new category, ELINK. ELINKs capture the relationship between the Topic Time (TT) and the Time of Situation (TSit) and have an aspectual interpretation in Klein's (1994) theory. In addition to these automatic and inferential annotations, some TLINKs were added manually. Using an example from the corpus, we show that the inclusion of the Topic Time in the annotations allows for a richer representation of the temporal structure than does TimeML. A way of representing this structure in a diagrammatic form similar to the T-Box format (Verhagen, 2007) is proposed.

Keywords: inferential annotation, Topic Time, GraphAnno

1. Introduction

TimeML[1] is a markup language which allows us, among other things, to annotate and analyse the temporal structure of a text and to represent it diagrammatically, e.g. in the T-Box format (Verhagen, 2007). It has been used to annotate TimeBank (Pustejovsky et al., 2003),[2] a collection of news articles from various sources, including the Wall Street Journal. While TimeML is an unrivaled standard in corpus markup, the annotations are somewhat sparse from a theoretical point of view, as they only provide information about the temporal order of events, but not about the 'Reference Time' in terms of Reichenbach (1947), or the 'Topic Time' in terms of Klein (1994) (Derczynski and Gaizauskas, 2013; Gast et al., 2015b). In this paper we propose a way of enriching TimeBank in such a way that it allows us to represent and extract temporal relations in a more precise way.

We start in Section 2. with some remarks on tense and aspect. In Section 3., we describe the annotation procedure, using a sentence from TimeBank for illustration. Section 4. contains the conclusions.

2. Theories of tense and aspect

Traditionally, tense is regarded as a relation between some event and the moment of utterance (Leech, 1971). Accordingly, the past tense is assumed to be used to describe events that are located before the moment of utterance. However, since Reichenbach (1947) at the latest, it has been known that tense and aspect cannot be adequately analysed without taking into account a third component, labeled 'Reference Time' by Reichenbach (1947), and explicated by Klein (1994) as 'Topic Time'. The Topic Time is "the time span to which a speaker's claim is confined" (Klein, 1994, 6). According to this view, tense is not a (binary) relation between an event and a point in time (the time of utterance); it is a ternary relation involving, in addition to the time of utterance/TU and the time of situation/TSit, a Topic Time/TT.

In English, the Topic Time has proven an indispensable component of temporal analysis, even at a descriptive level,[3] as it allows us to differentiate between the Simple Past ('I worked a lot') and the Present Perfect ('I have worked too much'). Both tenses refer to events located (partly or entirely) before the moment of utterance (TU); they differ in that the Simple Past views the event from a past perspective – the Topic Time is located before the moment of utterance (TT < TU, cf. 1a) – while the Present Perfect takes the perspective of the present moment, which means that the Topic Time includes the time of utterance (TT $\supseteq$ TU, cf. 2a). The Topic Time is also essential for the interpretation of aspect. In the case of the Simple aspect, TT and TSit have a non-empty intersection and TT is not fully included in TSit. This relation is labeled 'AT' by Klein (1994), and is here represented as '@'[4] (cf. 1b). The Perfect aspect indicates that TSit is anterior to TT (or TT posterior to TSit, cf. 2b).

(1) Yesterday$_{TT}$ I worked$_{TSit}$ a lot.
  a. TT < TU (→ past tense)
  b. TT @ TSit (→ simple/non-progressive aspect)

(2) Now$_{TT}$ I have worked$_{TSit}$ too much.
  a. TT $\supseteq$ TU (→ present tense)
  b. TT > TSit (→ perfect aspect)

The Progressive aspect is used when the Topic Time is fully included in the time of the situation (cf. 3b).

(3) [Yesterday between 4 and 5]$_{TT}$ I was working$_{TSit}$.
  a. TT < TU (→ past tense)
  b. TT $\subset$ TSit (→ progressive aspect)

[1] `http://www.timeml.org`

[2] `https://catalog.ldc.upenn.edu/LDC2006T08`

[3] For instance, it is used by Huddleston and Pullum (2002), where it is called 'time referred to (T$_r$)'.

[4] A @ B iff $A \cap B \neq \emptyset \wedge \neg (A \subset B)$.

```
The White House <EVENT eid=”e1” class=”REPORTING”>said</EVENT> President Bush has <EVENT eid=”e2” class=”I_ACTION”>approved</EVENT>
duty-free <EVENT eid=”e25” class=”STATE”>treatment</EVENT> for imports of certain types of watches that are n’t <EVENT eid=”e3”
class=”OCCURRENCE”>produced </EVENT> in ”significant quantities” in the U.S., the Virgin Islands and other U.S. possessions.

The <EVENT eid=”e26” class=”OCCURRENCE”>action</EVENT> <EVENT eid=”e5” class=”OCCURRENCE”>came</EVENT> in response to a petition
<EVENT eid=”e6” class=”OCCURRENCE”>filed</EVENT> by Timex Inc. for changes in the U.S. Generalized System of Preferences for imports from developing
nations. <SIGNAL sid=”s28”>Previously</SIGNAL> , watch imports were <EVENT eid=”e10” class=”I_ACTION”>denied</EVENT> such duty-free <EVENT
eid=”e29” class=”STATE”>treatment</EVENT>.

Timex had <EVENT eid=”e11” class=”OCCURRENCE”>requested</EVENT> duty-free treatment for many types of watches, covered by 58 different U.S. tariff
classifications.

<MAKEINSTANCE eventID=”e1” eiid=”ei1989” tense=”PAST” aspect=”NONE” polarity=”POS” pos=”VERB”/>
<MAKEINSTANCE eventID=”e2” eiid=”ei1990” tense=”PRESENT” aspect=”PERFECTIVE” polarity=”POS” pos=”VERB”/>
<MAKEINSTANCE eventID=”e25” eiid=”ei1991” tense=”NONE” aspect=”NONE” polarity=”POS” pos=”NOUN”/>
<MAKEINSTANCE eventID=”e3” eiid=”ei1992” tense=”PRESENT” aspect=”NONE” polarity=”NEG” pos=”VERB”/>
<MAKEINSTANCE eventID=”e26” eiid=”ei1993” tense=”NONE” aspect=”NONE” polarity=”POS” pos=”NOUN”/>
<MAKEINSTANCE eventID=”e5” eiid=”ei1994” tense=”PAST” aspect=”NONE” polarity=”POS” pos=”VERB”/>
<MAKEINSTANCE eventID=”e6” eiid=”ei1995” tense=”PASTPART” aspect=”NONE” polarity=”POS” pos=”VERB”/>
<MAKEINSTANCE eventID=”e10” eiid=”ei1996” tense=”PAST” aspect=”NONE” polarity=”POS” pos=”VERB”/>
<MAKEINSTANCE eventID=”e29” eiid=”ei1997” tense=”NONE” aspect=”NONE” polarity=”POS” pos=”NOUN”/>
<MAKEINSTANCE eventID=”e11” eiid=”ei1998” tense=”PAST” aspect=”PERFECTIVE” polarity=”POS” pos=”VERB”/>

<TLINK lid=”l7” relType=”IDENTITY” eventInstanceID=”ei1993” relatedToEventInstance=”ei1990”/>
<TLINK lid=”l8” relType=”BEGUN_BY” eventInstanceID=”ei1993” relatedToEventInstance=”ei1994”/>
<TLINK lid=”l10” relType=”AFTER” eventInstanceID=”ei1994” relatedToEventInstance=”ei1995”/>
<TLINK lid=”l12” relType=”BEFORE” eventInstanceID=”ei1996” relatedToEventInstance=”ei1993” signalID=”s28”/>
<TLINK lid=”l14” relType=”BEFORE” eventInstanceID=”ei1998” relatedToEventInstance=”ei1990”/>
```

Figure 1: The first four sentences from document wsj_0026.tml with MAKEINSTANCE and TLINK-elements

3. Annotating temporal relations

3.1. An example from TimeBank

As both tense and aspect are defined as relationships between the Topic Time and another interval (either TU or TSit), TT is obviously the primary ‘anchor’ for temporal interpretation. We will illustrate its role with an example from the TimeBank corpus. Figure 1 above shows the first four sentences from document wsj0026.tml, with the MAKEINSTANCE-nodes and the TLINKs. (4) below contains the part that is relevant to the following discussion in plain text, with tokens referenced by TLINKs underlined.

(4) . . . President Bush has approved duty-free treatment for imports of certain types of watches . . . . The action came in response to a petition filed by Timex Inc. . . . . Previously, watch imports were denied such duty-free treatment. Timex had requested duty-free treatment for many types of watches, . . .

The temporal structure as retrievable from the TimeML annotations in Figure 1 is displayed in Figure 2 in the T-Box format (Verhagen, 2007). The TLINKs of TimeML indicate that the events encoded by *requested* and *denied* precede those encoded by *approved* and *action*, and that *came* denotes the onset of the *action* (of approval). They do not show, however, that *approved* has “current relevance” (Leech, 1971), indicated by the present tense of the auxiliary *has*. This kind of information is part of the ‘enhanced’ annotation proposed in the present paper.

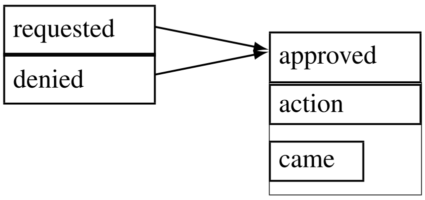

Figure 2: Fragment (4) in T-Box format

3.2. Preprocessing and automatic annotation

We imported the 132 Wall Street Journal documents of TimeBank 1.2 with lower case ‘wsj’ in the file name into GraphAnno (Gast et al., 2015b).[5] This part of the corpus comprises 37,908 tokens. In GraphAnno, the data is represented as a graph both program-internally and in terms of visualization. Figure 3 shows the beginning of the second sentence in (4) above.

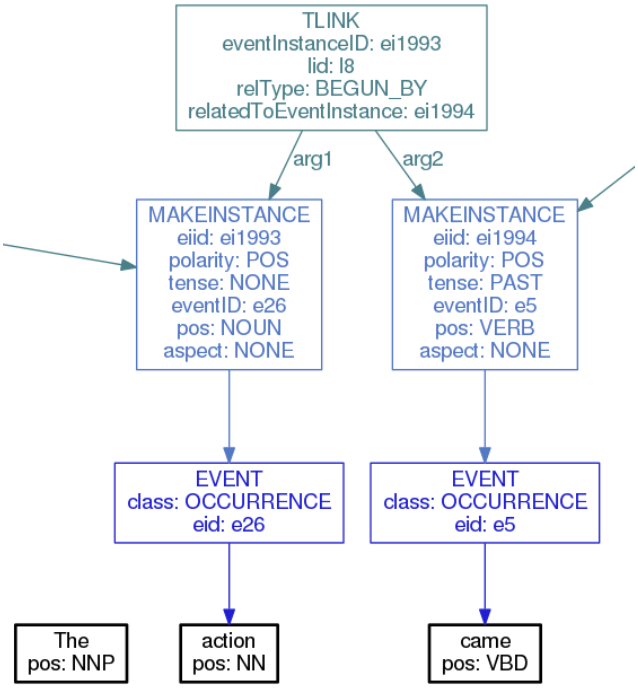

Figure 3: Beginning of the second sentence of (4)

[5] GraphAnno was originally designed as a prototype of another tool, Atomic (Druskat et al., 2014), but as it has proven very useful, we have continued to use it for annotation projects (Gast et al., 2015a; Gast et al., 2015b). It is available on GitHub (https://github.com/LBierkandt/graph-anno).

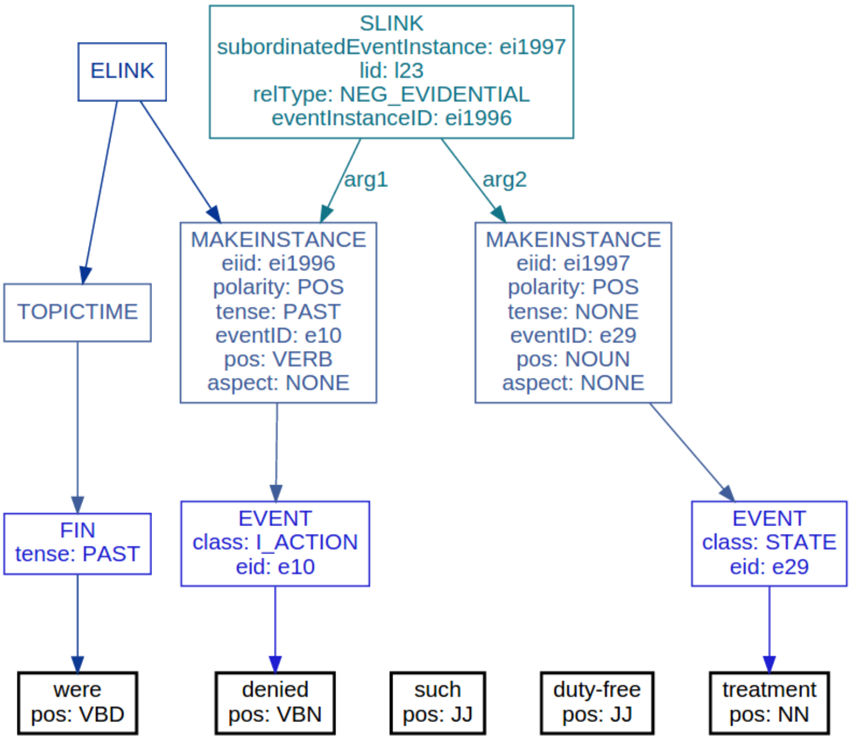

Figure 4: A part of (4), represented as a graph

We tagged the corpus with the Stanford bidirectional dependency network tagger. On the basis of the POS-tags, finite verb forms could be identified automatically and a FIN-node was created on top of each finite verb (cf. Figure 4). In Klein's (1994) theory, finiteness is regarded as the morphosyntactic carrier of information on the Topic Time. Note that FIN-nodes were also created on verbs already linked to a MAKEINSTANCE-node in TimeBank.

3.3. Inferential annotation

While we used a Python-interface to GraphAnno for preprocessing and automatic annotation, the following processes of annotation were carried out in GraphAnno. We call such annotations 'inferential', as they are based on information implied by other annotations in the corpus and consist in a 'transfer' of information between corpus elements. Inferential annotations can be carried out globally in GraphAnno, as will be detailed in Section 3.4.

Analogously to MAKEINSTANCE-nodes dominating EVENT-nodes in TimeML, we created a TOPICTIME-node on top of any FIN-node generated by the automatic annotation step. As has been mentioned, Klein (1994) treats aspect as a relationship between the Topic Time and the time of situation, TSit. Translated into an annotation scheme, this means that aspect is represented by links between a TOPICTIME-node and a MAKEINSTANCE-node. We have chosen the label 'ELINK' for this type of link (with 'E' standing for 'event').[6] ELINKs were added between TOPICTIME-nodes and EVENT-nodes dominating the same token or, in the case of periphrastic forms (e.g. Present Perfect, Progressive aspect, etc.), between TOPICTIME-nodes and the next MAKEINSTANCE-node to the right, within a window of three tokens.[7]

This created 2,453 ELINKs. The VP *were denied such duty-free treatment* from (4), thus annotated, is shown in Figure 4.

As MAKEINSTANCE-nodes contain temporal annotations (key-value pairs for tense and aspect), we could compare the temporal categories assigned on the basis of the POS-tagging with the manual annotations of TimeBank. The annotations diverged in 236 of 2,453 cases, mostly because the temporal interpretation of modals such as *could* was not retrievable from the POS-information alone. We therefore relied on the manual annotations of TimeBank for further processing.

In order to integrate the TOPICTIME-nodes into the temporal graph we had to link them to the network of MAKEINSTANCE-nodes. This was achieved via the aspectual values retrievable from those nodes. TimeML distinguishes four values for the key 'aspect': NONE, PROGRESSIVE, PERFECTIVE and PERFECTIVE_PROGRESSIVE. With the exception of PERFECTIVE, each of these values implies that there is overlap between the Topic Time and the time of the situation, TSit. Accordingly, we assigned all of these categories to the more general class of 'OVERLAP' (which differs from Klein's 'AT'-relation in that it also covers cases of full inclusion). For the PERFECTIVE cases, we assigned the relation 'AFTER', as the Topic Time is located after TSit in such cases (Klein, 1994). Given that we intended to reduce the number of relations to two, we reversed AFTER-edges, changing them into BEFORE-edges.

[6] 'ALINK' may have been more appropriate, but this name is used for information relating to *aktionsart* in TimeML already.

[7] This restriction was necessary because some non-finite verb forms are not annotated in TimeBank, and an unrestricted association between finite forms and non-finite verbs to their right would have led to continued misalignments.

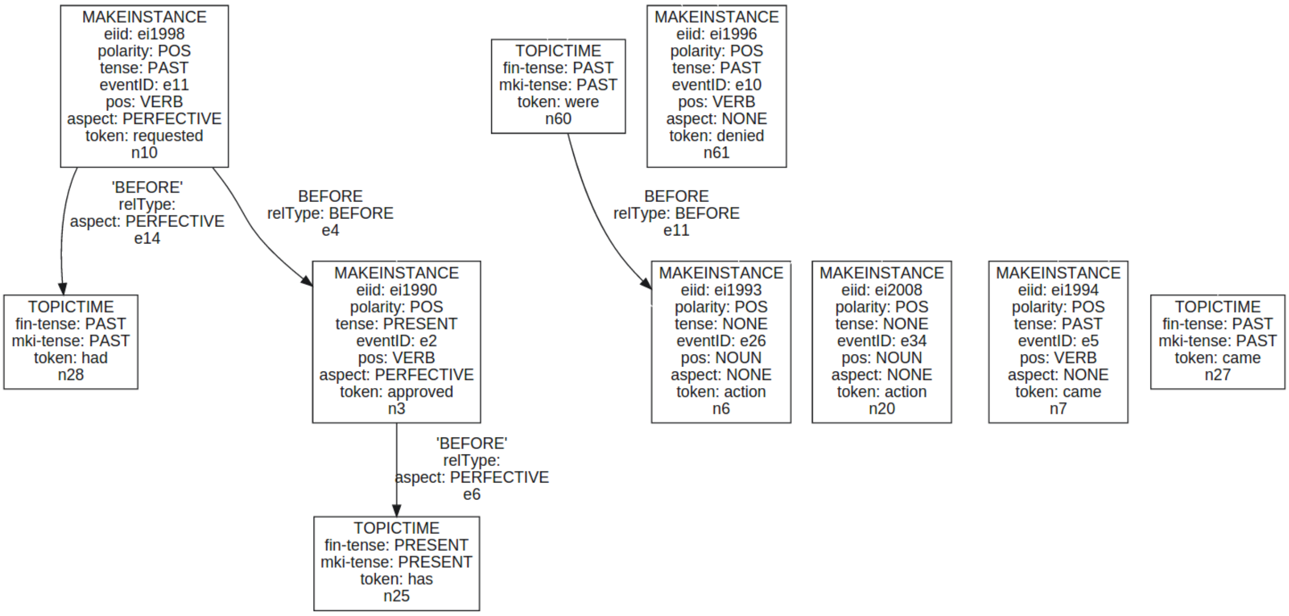

Figure 5: Partial graph corresponding to (4)

As we are primarily interested in the temporal arrangement of Topic Times, not of events (or times of situations/TSits/MAKEINSTANCE-nodes), we inferred temporal order relations from the existing annotations on edges linking MAKEINSTANCE-nodes to each other. We first subsumed the specific relation types (relType-attributes) holding between MAKEINSTANCE-nodes in TimeML under the general categories OVERLAP and BEFORE/AFTER, just like the four aspectual categories of TimeBank. The category of OVERLAP contains all relations except BEFORE and AFTER as well as their 'immediate' counterparts, IBEFORE and IAFTER (e.g. INCLUDES, IS_INCLUDED, DURING). This allowed us to infer TLINKs between TOPICTIME-nodes and MAKEINSTANCE-nodes on the basis of rules like the following (cf. Section 3.4. for a technical description of this process):

(5) IF (TT OVLP MKI$_1$ AND MKI$_1$ BEFORE MKI$_2$)
{TT BEFORE MKI$_2$}

(6) IF (TT OVLP MKI$_1$ AND MKI$_2$ BEFORE MKI$_1$)
{MKI$_2$ BEFORE TT}

Moreover, as in the case of ELINKS, we transformed all order relations into BEFORE-relations (i.e., AFTER-links were reversed).

The resulting graph contains some information about the temporal ordering of Topic Times already. It can be visualized in GraphAnno by assigning TOPICTIME-nodes, MAKEINSTANCE-nodes as well as TLINKs to a specific layer, and by applying a filter hiding all the rest from the graph. OVERLAP-relations are represented as invisible nodes dominating their arguments, which leads to the arrangement of overlapping Topic Times or TSits on a given horizontal layer. BEFORE-relations are realized as edges, leading to a vertical arrangement of temporally ordered nodes. For sentence (4), this gives us the representation shown in Figure 5 (the tokens dominated by TOPICTIME and MAKEINSTANCE-nodes are indicated in these nodes as values of the key 'token').

The temporal structure of (4) can be represented more clearly as shown in Figure 6, in a format similar, but not identical, to T-Box. Left-to-right arrangement indicates temporal ordering, and the vertical arrangement corresponds to the occurrence of the event descriptions in the text. TOPICTIME-nodes are represented in dashed boxes, MAKEINSTANCE-nodes in boxes with solid lines (the finite verb *came* is therefore enclosed by both a solid and a dashed box, as it is dominated by both a FIN- and a MAKEINSTANCE-node, cf. Section 3.2.). ELINKs are represented by solid lines, TLINKs are not represented explicitly, but are reflected in horizontal ordering.

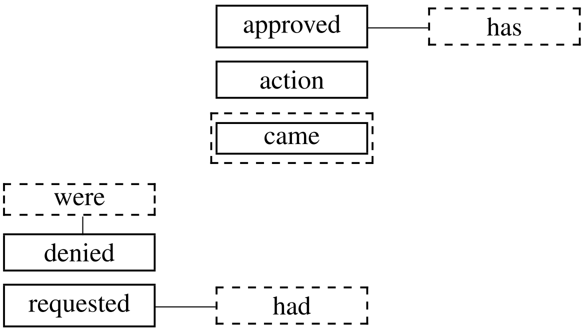

Figure 6: Temporal structure of the fragment in (4)

While the representation in Figure 6 disregards some of the information retrievable from the T-Box format (e.g. relationships of inclusion), it is richer in that it provides information about the Topic Time for each finite predication. For example, it shows that the event denoted by the verb *approved* provides information about the Topic Time associated with the auxiliary *has*, i.e., the moment of utterance. As is well known, such "current relevance" (Leech, 1971) is a crucial ingredient of the Present Perfect in English and distinguishes *I broke my leg* from *I've broken my leg*.

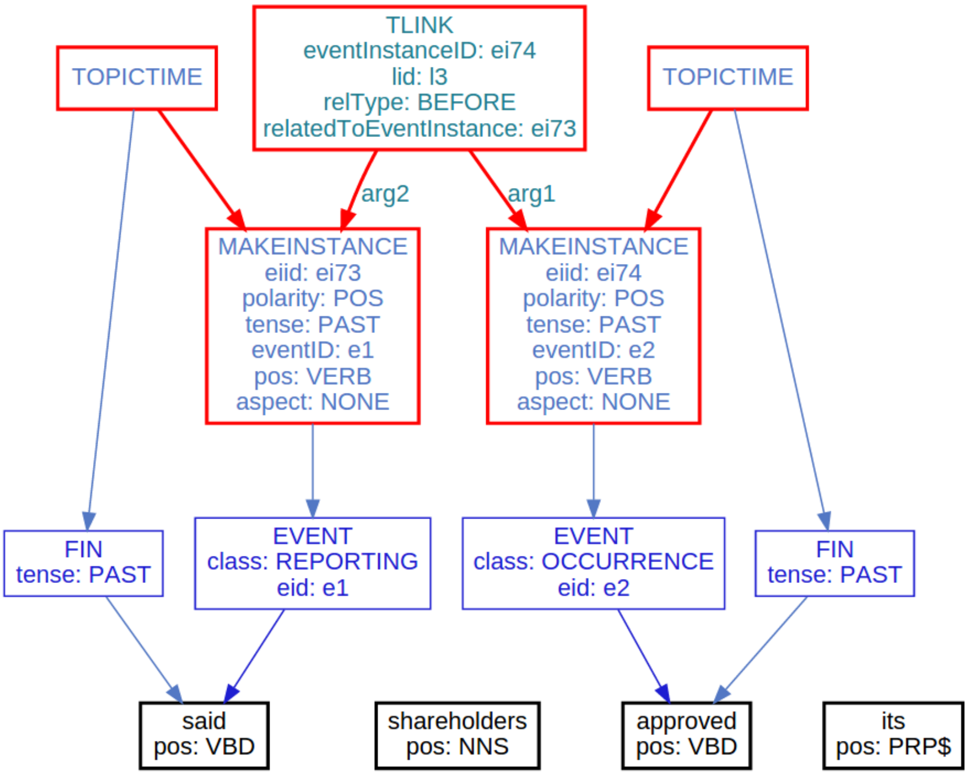

Figure 7: A graph fragment as defined in (10)

The representation in Figure 6 is still quite similar to the one in Figure 2. As we will show in Section 3.5., however, the presence of TOPICTIME-nodes in the annotation graph prepares the ground for further, manual annotations. Before turning to a further semantic enrichment of the corpus we will describe the process of inferential annotation from a technical point of view.

3.4. Inferential annotation in GraphAnno

It is one of the main assets of GraphAnno that it combines the visualization and annotation of corpus data with search as well as export functionalities (Gast et al., 2015b; Gast et al., 2015a). Its search function can moreover be used for what we call 'inferential annotation'. This type of annotation consists in identifying specific patterns in the graph and (globally) applying annotations to these patterns.

By pressing F7 a search window is opened where graph fragments can be defined. We can simply search for nodes or edges with specific attributes, as in (7) and (8).

(7)
```
node cat:TIMEX3 & value:1989-11-02
```
(finds TIMEX3-nodes referring to 02/11/1989)

(8)
```
edge cat:TLINK & relType:IAFTER
```
(finds TLINKs with an 'immediately after' relation)

More complex graph fragments can be defined by assigning names to nodes or edges for cross-referencing. Names are arbitrary but carry a @-prefix. For example, in order to find a TLINK governing a TIMEX3-node and a MAKEINSTANCE-node we define the two nodes and assign names to them, and then specify an edge linking the two nodes (the edge can also be assigned a name, here '@edg'):

(9)
```
node @tmx cat:TIMEX3
node @mki cat:MAKEINSTANCE
edge @edg@tmx@mki
```

The hits of the search process are graph fragments meeting the relevant conditions. Their elements can now be used as arguments of GraphAnno's annotation commands, for example `a` (annotate), `p` (create parent node), `e` (create edge), etc. (Gast et al., 2015b; Gast et al., 2015a). We will illustrate this procedure with the inferential annotation process that we used to transfer temporal relations holding between MAKEINSTANCE-nodes to the relevant TOPICTIME-nodes (TLINK-nodes had been transformed into edges at this stage). Here is the code:

(10)
```
node @tlink cat:TLINK
node @arg1 cat:MAKEINSTANCE
     & !aspect:PERFECTIVE
edge @ea@tlink@arg1 cat:arg1

node @tt1 cat:TOPICTIME
edge @ex@tt1@arg1

node @tt2 cat:TOPICTIME
edge @ey@tt2@arg2

node @arg2 cat:MAKEINSTANCE
     & !aspect:PERFECTIVE
edge @eb@tlink@arg2 cat:arg2

e @tt1 @tt2
  cat:"#{@tlink[’relType’]}" f
```

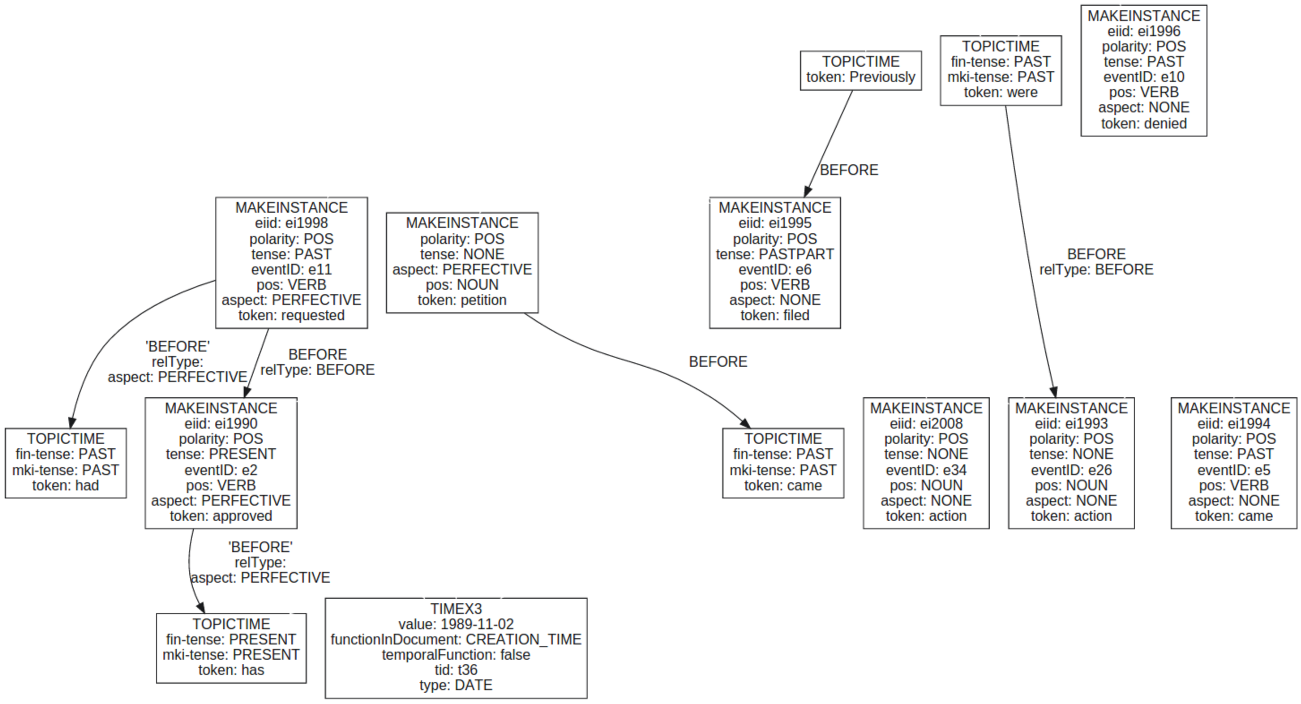

Figure 8: The manually enriched graph fragment corresponding to (4)

The first block finds graph fragments in which a TLINK (assigned the name `@tlink`) dominates a MAKEINSTANCE-node (`@arg1`) with an aspectual value other than PERFECTIVE. The edge linking the two nodes is assigned the name `@ea`, and it is required to be of category 'arg1' (the attributes 'arg1' and 'arg2' are not contained in TimeML but were added during the import into GraphAnno, to keep the directedness of the edges as they were transformed into nodes govering their arguments).

The second block defines a TOPICTIME-node (`@tt`) governed by the TLINK-node `@tlink` and assigns the name `@ex` to the linking edge. The third and forth block repeat the same procedure for another pair of MAKEINSTANCE- and TOPICTIME-nodes.

When searching for such nodes in GraphAnno, the hits are highlighted, as shown in Figure 7.

The last line of the code in (10) has the same syntax as a 'common' annotation command in GraphAnno (Gast et al., 2015b; Gast et al., 2015a). There are two differences, however. The nodes are referenced by the names assigned to them in the query, and the inferential annotation procedure may use features of other graph elements as arguments. The command `e` produces an edge from `@tt1` to `@tt2` (the two TOPICTIME-nodes). It is, moreover, assigned a category. This category is the value for the key `relType` of the `@tlink`-node (as the program is written in Ruby, the query language partially uses Ruby syntax). In this way, the TLINK-node dominating two MAKEINSTANCE-nodes is transformed into an edge linking the two TOPICTIME-nodes to each other.

3.5. Manual annotations

Returning to matters of content, we will illustrate how the corpus can be annotated further – manually – after adding automatic and inferential annotations.

GraphAnno runs in a browser window and is operated with simple, often one-character commands as illustrated in Section 3.4.. The type of annotation that we will focus on now concerns the creation of additional temporal links between elements in the annotation graph. TLINKs between two nodes `n1` and `n2` can be created with the command in (11):

```
(11)  p n1 n2 cat:TLINK
```

To make things easier, we can simply define a macro for the key-value pair `cat:TLINK` – say, `x` – and thus create TLINKs with `p n1 n2 x`.

Among the most obvious candidates to be linked to Topic Times are the TIMEX3-nodes referring to the CREATION_TIME or the PUBLICATION_TIME.[8] Moreover, we could add further TLINKs between adverbs introducing Topic Times, such as *previously* in (4), and non-finite predicates that cannot be inferentially linked to any Topic Time, e.g. the post-nominal attribute *filed*, modifying *petition*. Just adding these three TLINKs makes the structure considerably richer. A view of the relevant graph fragment is shown in Figure 8 above. Using our diagrammatic format of representation, the fragment can be represented as shown in Figure 9 below.

4. Conclusions

We have aimed to show how the TimeBank corpus can be enriched with temporal annotations in the spirit of Klein (1994), an analysis of tense and aspect that is widely considered as a standard among theoretical linguists and that allows us to represent temporal structure in a richer way than a traditional, binary analysis. A crucial ingredient of our 'enhanced' annotation scheme is the annotation of a Topic Time, "the time span to which a speaker's claim is confined"

[8] Note that these annotations can of course be carried out automatically, but not within GraphAnno, which presently supports inferential annotation only for coherent graph fragments.

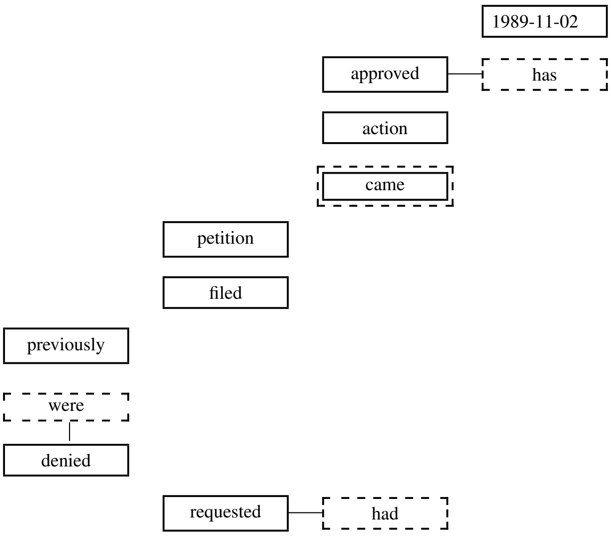

Figure 9: More elaborate temporal structure of (4)

(Klein, 1994). The Topic Time not only allows us to capture semantic differences between, for instance, the Present Perfect and the Simple Past, it also provides an important anchor and reference point for further temporal annotations. Our annotations have largely been automatic and inferential, and we have tried to show that with a few additional manual annotations a considerable gain in granularity can be achieved. Inferential annotation is a process that is greatly faciliated by GraphAnno (Gast et al., 2015b; Gast et al., 2015a). As GraphAnno is easy to use and does not require any specific computational skills, we hope to inspire theoretical linguists to carry out similar studies, thus contributing to the slow but steady process of convergence between ‘small data’ theoretical studies and more applied studies based on ‘big data’.

5. Acknowledgements

We are indebted to three anonymous reviewers for valuable comments. Financial support from the German Science Foundation (DFG, grant GA-1288/5-1) is gratefully acknowledged.

6. Bibliographical References

Derczynski, L. and Gaizauskas, R. (2013). Empirical validation of Reichenbach’s tense framework. In *Proceedings of the 10th International Conference on Computational Semantics*.

Druskat, S., Bierkandt, L., Gast, V., Rzymski, C., and Zipser, F. (2014). Atomic: An open-source software platform for multi-level corpus annotation. In Josef Ruppert et al., editors, *Proceedings of the 12th Konferenz zur Verarbeitung natürlicher Sprache (KONVENS 2014)*, pages 228–234.

Gast, V., Bierkandt, L., and Rzymski, C. (2015a). Annotating modals with GraphAnno, a configurable lightweight tool for multi-level annotation. In *Proceedings of the Workshop on Models for Modality Annotation, held in conjunction with IWCS 11, 2015*, pages 19–28, Stroudsburg, PA.

Gast, V., Bierkandt, L., and Rzymski, C. (2015b). Creating and retrieving tense and aspect annotations with graphanno, a lightweight tool for multi-level annotation. In H. Bunt, editor, *Proceedings of the 11th Joint ACL-ISO Workshop on Interoperable Annotation*, pages 23–28, Tilburg. Tilburg Center for Cognition and Communication.

Huddleston, R. and Pullum, G. (2002). *The Cambridge Grammar of the English Language*. Cambridge University Press, Cambridge.

Klein, W. (1994). *Time in Language*. Routledge, London.

Leech, G. (1971). *Meaning and the English Verb*. Taylor & Francis, London.

Pustejovsky, J., Hanks, P., Saurí, R., See, A., Gaizauskas, R., Setzer, A., Radev, D., Sundheim, B., Day, D., Ferro, L., and Lazo., M. (2003). The timebank corpus. *Corpus Linguistics*, pages 647–656.

Pustejovsky, J., Ingria, R., Saurí, R., Castaño, J., Littman, J., Gaizauskas, R., and Setzer, A. (2005). The specification language TimeML. In Inderjeet Mani, et al., editors, *The Language of Time: A Reader*. Oxford University Press, Oxford.

Reichenbach, H. (1947). *Elements of Symbolic Logic*. Macmillan & Co, New York.

Verhagen, M. (2007). Drawing timeml relations with tbox. In F. et al. Schilder, editor, *Reasoning about Time and Events*, pages 7–28. Springer, Heidelberg.