Temporal Relation Classification using a Model of Tense and Aspect

Leon Derczynski
University of Sheffield
leon@dcs.shef.ac.uk

Robert Gaizauskas
University of Sheffield
robertg@dcs.shef.ac.uk

Abstract

Determining the temporal order of events in a text is difficult. However, it is crucial to the extraction of narratives, plans, and context. We suggest that a simple, established framework of tense and aspect provides a viable model for ordering a subset of events and times in a given text. Using this framework, we investigate extracting features that represent temporal information and integrate these in a machine learning approach. These features improve event-event ordering.

1 Introduction

It is important to understand time in language. The ability to express and comprehend expressions of time enables us to plan, to tell stories, and to discuss change in the world around us.

When we automatically extract temporal information, we are often concerned with events and times – referred to collectively as temporal **intervals**. We might ask, for example, *"Who is the current President of the USA?."* In order to extract a single contemporary answer to this question, we need to identify events related to persons becoming president and the times of those events. Crucially, however, we also need to identify the ordering between these events and times, by assigning a temporal relation type (from e.g. Allen (1983)). This last task, **temporal relation typing**, is challenging (UzZaman et al., 2013; Bethard et al., 2015), and is the focus of this paper.

When events are expressed as verbs, tense and aspect are used to convey temporal features of these events. Thus, it is intuitive that tense and aspect will be of value in determining the type of temporal relation that holds between two verb events, and evidence in human-annotated corpora supports this intuition.

Event-event relations are often the hardest to label (Derczynski, 2015). Around 45% of links in TempEval-3 (UzZaman et al., 2013) event-event ordering tasks cannot reliably be labelled automatically.

Temporal relations involving at least one argument with tense or aspect information are prevalent. Verb-verb links make up around a third of TimeBank's temporal relations,[1] and tensed verb-verb links the largest share of that set, so of all verb-verb relations, the majority are between two tensed verbs.

Data-driven approaches to the relation typing task are hampered in two ways. Firstly, there is a shortage of ground truth training data. This leads to low volumes of instances for many combinations of tense and aspect values for pairs of events, hampering automatic hypothesis learning (Lapata and Lascarides, 2006). Secondly, the range of tense and aspect expression in TimeML is relatively limited, describing three "tenses"[2] (past and past participle, present and present participle, and future) and three "aspects" (none, perfective and progressive). This markup language may be insufficiently descriptive to capture relations implied by variations in linguistic use of tense and aspect.

Reichenbach (1947) offers a theoretical framework for analysis of tense and aspect that can be used to predict constraints on temporal orderings between verb events based on their tense and aspect, and also between times and tensed verbs. Applying Reichenbach's framework requires tense and aspect information, which may yet be usefully available in existing corpora.

In this paper, we describe an approach to using Reichenbachs model to generate features for

[1] TimeBank is a corpus semantically annotated for temporal information in TimeML (Pustejovsky et al., 2003; Pustejovsky et al., 2004)

[2] In TimeML v1.2, the tense attribute of events has values that are conflated with verb form. This conflation is deprecated in newer versions of TimeML, post-TimeBank.

a machine learning approach to temporal relation typing and report an experiment showing it brings modest improvement.

2 Reichenbachian Tenses

Reichenbach details nine tenses (see Table 1). The tenses detailed by Reichenbach are past, present or future, and may take a simple, anterior or posterior form. In English, these apply to single finite verbs and to verbal groups consisting of head verb and auxiliaries. The tense system describes abstract time points for each tensed verb – event time E, speech/utterance time S, and reference time R – and how they may interact, both for a single verb and with other events.

In Reichenbach's view, different tenses specify different relations between E, R and S. Table 1 shows the six tenses conventionally distinguished in English. As there are more than six possible ordering arrangements of S, E and R, some English tenses might suggest more than one arrangement. Reichenbach's tenses also suffer from this ambiguity when converted to $S/E/R$ structures, albeit to a lesser degree. When following Reichenbach's tense names, it is the case that for past tenses, R always occurs before S; in the future, R is always after S; and in the present, S and R are simultaneous. Further, "anterior" suggests E before R, "simple" that R and E are simultaneous, and "posterior" that E is after R. The flexibility of this framework is sufficient to allow it to account for a very wide set of tenses, including all those described by Song and Cohen (1988), and this is sufficient to account for the observed tenses in many languages. Past, present and future tenses imply $R < S$, $R = S$ and $S < R$ respectively. Anterior, simple and posterior tenses imply $E < R$, $E = R$ and $R < E$ respectively.

2.1 Verb Interactions

While each tensed verb involves a speech, event and reference time, multiple verbs may share one or more of these points. For example, all narrative in a news article usually has the same speech time (that of document creation). Further, two events linked by a temporal conjunction (e.g. *after*) are very likely to share the same reference time. Basic methods of linking between verb events or linking verbs to fixed points on a time scale are described below.

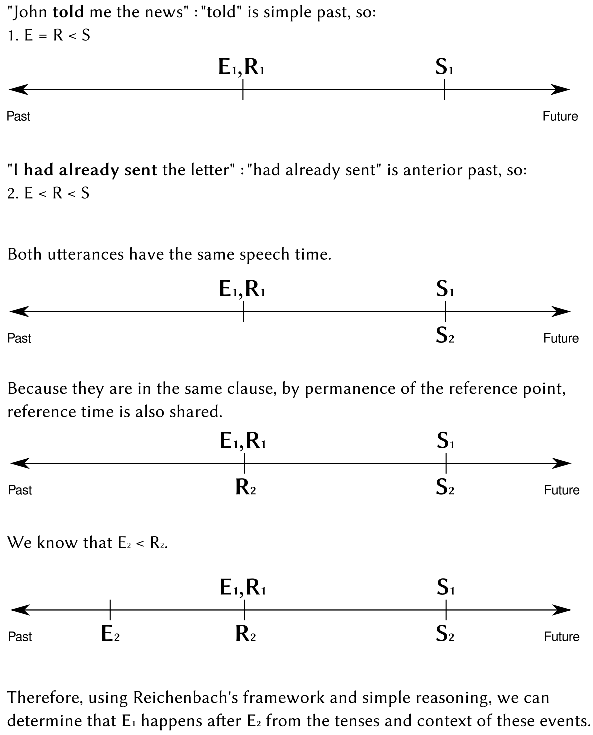

Figure 1: An example of permanence of the reference point.

2.2 Special Properties of the Reference Point

The reference point R has two special uses. These relate to verbs in the same *temporal context* and to the effect of time expressions on verbs. Reichenbach relies on a notion of "same temporal context" without ever defining it precisely. It could be similar to the concept put forward by Dowty (1986) with **temporal discourse interpretation principle** (TDIP). Below we operationalise the concept in several ways to mean either "same sentence" or "adjacent sentence pairs", though other interpretations are also possible.

Permanence Firstly, when sentences are combined to form a compound sentence, tensed main verbs interact, and implicit grammatical rules require tenses to be adjusted. These rules operate such that R is the same in all cases in the sequence. Reichenbach names this principle *permanence* of the reference point. Figure 1 contains an example of this principle.

Positional Secondly, when temporal expressions (such as a TimeML TIMEX3 of type DATE, but not DURATION) occur in the same clause as a verbal event, the temporal expression does not (as one might expect) specify event time E, but instead is used to position reference time R. This is named *positional* use of the reference point.

Relation	*Reichenbach's Tense Name*	*English Tense Name*	*Example*
E<R<S	Anterior past	Past perfect	*I had slept*
E=R<S	Simple past	Simple past	*I slept*
R<E<S			
R<S=E	Posterior past		*I expected that I would sleep*
R<S<E			
E<S=R	Anterior present	Present perfect	*I have slept*
S=R=E	Simple present	Simple present	*I sleep*
S=R<E	Posterior present	Simple future	*I will sleep (Je vais dormir)*
S<E<R			
S=E<R	Anterior future	Future perfect	*I will have slept*
E<S<R			
S<R=E	Simple future	Simple future	*I will sleep (Je dormirai)*
S<R<E	Posterior future		*I shall be going to sleep*

Table 1: Reichenbach's tenses; from Mani et al. (2005)

e1 ↓; e2 →	**Sim Past**	**Pos Past**	**Ant Pres**	**Sim Pres**	**Ant Fut**	**Sim Fut**
Sim Past	vague	after	vague	after	after	after
Pos Past	before	vague	vague	vague	after	after
Ant Pres	vague	vague	vague	after	vague	after
Sim Pres	before	vague	vague	overlap	vague	after
Ant Fut	before	before	vague	vague	vague	after
Sim Fut	before	before	before	before	before	vague

Table 2: Verb-verb event orderings based on the Reichenbachian tenses that map directly to those in TimeML. Cell values describe the `e1` [*rel*] `e2` relationship.

In Example 1, the reference point is determined positionally with an explicit time (*10 o'clock*).

(1) *It was 10 o'clock, and Sarah had brushed her teeth.*

The verb group *had brushed* is anterior past tense; that is, $E < R < S$. The event is complete before the reference time – that is, at any point until *10 o'clock* – and so the relation between the event and timex can be determined (*brushed* BEFORE *10 o'clock*).

2.3 Feature Extraction

Two interpretations of the model are used in feature extraction. Firstly, a simple view is taken assuming permanence of the reference point. This provides a constraint dependent on the pairing of Reichenbachian tenses used, and is detailed in Table 2. Secondly, an advanced interpretation is used, following Derczynski and Gaizauskas (2013). This approach fully populates all Reichenbachian tense combinations using Freksa's temporal semi-interval algebra (Freksa, 1992) to derive a (large) temporal constraint table, which for space reasons is omitted here.

In all cases, the gold standard tense and aspect features annotated on the events in TimeBank are used as the basis for Reichenbachian representations.

3 The Framework in TLINK Typing

TimeML provides some of the information that Reichenbach's framework alone does not cater for and vice versa. A combination of the two may lead to better labelling performance, but relying on Reichenbach's framework alone for rule-based temporal relation label constraint is insufficient. However, the framework has shown to inform prior systems effectively (Chambers et al., 2014). The situations we examine are those where two verb events occur in the same temporal context, where a timex directly influences a verb event, and also verb events that report other verb events.

Reichenbach's framework is used as a linguistic model that generates temporal ordering features, which are added to a base feature set. The base features are those as in Mani et al. (2007), i.e.:

For each event: text; TimeML tense and aspect; modality; cardinality; polarity; event class; part-of-speech tag.

For each event pair: booleans for: are events in the same sentence; are events in adjacent sentences; do events have the same TimeML aspect, and again for tense; does event 1 textually precede event 2.

	Base features		Extended features	
Classifier	**Acc**	**Err. red.**	**Acc**	**Err. red.**
MCC	48.04%	-	48.04%	-
Maxent	57.47%	22.86%	**57.65%**	23.19%
ID3	56.52%	21.14%	**57.47%**	22.86%
N.Bayes	58.31%	24.37%	**58.72%**	25.12%

Table 3: Using Reichenbach-suggested event ordering features representing permanence of the reference point, considering only same-sentence TLINKs, using the advanced interpretation. 562 instances.

	Base features		Extended features	
Classifier	**Acc**	**Err. red.**	**Acc**	**Err. red.**
MCC	44.87%	-	44.87%	-
Maxent	62.28%	31.58%	**62.55%**	32.07%
ID3	**59.21%**	26.01%	58.74%	25.16%
N.Bayes	56.96%	21.92%	**57.58%**	23.05%

Table 4: Reichenbach-suggested event ordering feature representing permanence of the reference point, same-sentence and adjacent-sentence TLINKs. 858 instances.

3.1 Same Context Event-Event Links

Reichenbach's framework provides information for ordering events in the same temporal context (same context event-event relations, SCEE). This applies to any two verb events that have a shared reference point.

Verb events are those in TimeML that have a POS attribute of VERB. We exclude those with a TENSE of NONE or INFINITIVE. Shared reference points are assumed for event-event links having both arguments in the same or adjacent sentences.

4 Experimental Results

We conducted an experiment to test the utility of the Reichenbach-motivated temporal ordering features in a supervised learning approach to the temporal relation typing task. The goal is to find a way to incorporate Reichenbach's framework into a machine learning model. The experiment was conducted with 10-fold cross validation, using from TimeBank v1.2. Links in each document were never shared across a split (i.e., splits were made at document level). Experiments were conducted with relation folding, where the set of temporal relation types is reduced; e.g. AFTER and BEFORE can be switched between by flipping their argument order – A BEFORE B and B AFTER A are equivalent. The impact of Reichenbach's framework is measured by comparing classifier performance on SCEE links using the basic feature set and using the basic feature set plus the new feature. Features representing the text (i.e. lexical form) of events were removed as they consistently harmed performance, likely due to the sparsity of their values. Results are shown in Table 3. In this instance, the extended features provide a performance boost regardless of classifier choice. This shows that the framework can be integrated into a machine learning model for temporal relation typing. However, the improvements are modest. This can be attributed to a variety of factors salient to the relation typing task.

Firstly, the sizes of datasets, while not tiny, are still small. More temporally-annotated data will help here, though larger corpora using the same annotation standard are hard to come by. Next, Reichenbach can be applied with full accuracy to a tiny number of cases (where it makes an unambiguous suggestion) (Chambers et al., 2014), but this is only the first attempt to use it for constraining (rather than specifying) the target temporal relation type. Last, temporal context is not defined precisely but rather approximated. This is likely to affect results, and so we investigate further.

In the next case, the scope of temporal context is broadened to include cases where events are in adjacent sentences. Results are shown in Table 4. Here, classifiers with inductive biases toward the independence assumption do better with the extended feature set.

In both cases, there was a consistent performance increase from almost all classifiers with the introduction of the feature derived from the advanced interpretation of Reichenbach's framework. The performance increase was consistent when assuming that event-event relations in the same sentence are also in the same temporal context. The increase is smaller when context is stretched to adjacent sentences. We attribute this to weaknesses in modelling context, a task that others have also tackles (Miller et al., 2013) that remains an open and interesting research problem.

5 Conclusion

Reichenbach's framework for tense and aspect is intuitive, and of utility in typing temporal relations. Automatic identification of where the framework applies remains difficult. One question is how to formally define and annotate temporal context.We investigate two approximations for temporal context, both of which are useful. The other question is how to map Reichenbach's framework to features based on a common semantic annotation standard. We proposed two ways of using Reichenbach's framework to generate features for machine learning of temporal relations, which improved relation typing performance in this difficult task. The framework suggests helpful constraint of relation types in cases where verbs are in the same context, helping in the difficult task of automatic temporal relation typing.

References

J. Allen. 1983. Maintaining knowledge about temporal intervals. *Communications of the ACM*, 26(11):832–843.

S. Bethard, L. Derczynski, J. Pustejovsky, and M. Verhagen. 2015. Semeval-2015 task 6: Clinical TempEval. In *Proc. SemEval*.

N. Chambers, T. Cassidy, B. McDowell, and S. Bethard. 2014. Dense event ordering with a multi-pass architecture. *Transactions of the Association for Computational Linguistics*, 2:273–284.

L. Derczynski and R. Gaizauskas. 2013. Empirical Validation of Reichenbach's Tense Framework. In *Proceedings of the 10th Conference on Computational Semantics*, pages 71–82. ACL.

L. Derczynski. 2015. *Automatically Ordering Events and Times in Text*. Studies in Computational Intelligence. Springer.

D. Dowty. 1986. The effects of aspectual class on the temporal structure of discourse: semantics or pragmatics? *Linguistics and philosophy*, 9(1):37–61.

C. Freksa. 1992. Temporal reasoning based on semi-intervals. *Artificial intelligence*, 54(1):199–227.

M. Lapata and A. Lascarides. 2006. Learning sentence-internal temporal relations. *Journal of Artificial Intelligence Research*, 27(1):85–117.

I. Mani, J. Pustejovsky, and R. Gaizauskas. 2005. *The Language of Time: A Reader*. Oxford University Press.

I. Mani, B. Wellner, M. Verhagen, and J. Pustejovsky. 2007. Three approaches to learning TLINKS in TimeML. Technical Report CS-07-268, Brandeis University, Waltham, MA, USA.

T. A. Miller, S. Bethard, D. Dligach, S. Pradhan, C. Lin, and G. K. Savova. 2013. Discovering narrative containers in clinical text. *Proc. ACL*, page 18.

J. Pustejovsky, P. Hanks, R. Sauri, A. See, R. Gaizauskas, A. Setzer, D. Radev, B. Sundheim, D. Day, L. Ferro, et al. 2003. The Timebank Corpus. In *Corpus Linguistics*, pages 647–656.

J. Pustejovsky, B. Ingria, R. Sauri, J. Castano, J. Littman, and R. Gaizauskas. 2004. The Specification Language TimeML. In *The Language of Time: A Reader*, pages 545–557. Oxford University Press.

H. Reichenbach. 1947. The tenses of verbs. In *Elements of Symbolic Logic*. Dover Publications.

F. Song and R. Cohen. 1988. The interpretation of temporal relations in narrative. In *Proceedings of the 7th National Conference of AAAI*.

N. UzZaman, H. Llorens, L. Derczynski, J. Allen, M. Verhagen, and J. Pustejovsky. 2013. SemEval-2013 Task 1: TempEval-3: Evaluating Time Expressions, Events, and Temporal Relations. In *Proc. SemEval*, pages 1–9. Association for Computational Linguistics.