WVALI: Temporal Relation Identification by Syntactico-Semantic Analysis

Georgiana Puşcaşu[†§]
[†]Research Group in Computational Linguistics
University of Wolverhampton, UK
[§]Department of Software and Computing Systems
University of Alicante, Spain
georgie@wlv.ac.uk

Abstract

This paper reports on the participation of University of Wolverhampton and University of Alicante at the SemEval-2007 TempEval evaluation exercise. TempEval consisted of three tasks involving the identification of event-time and event-event temporal relations. We participated in all three tasks with TICTAC (Syntactico-Semantic Temporal Annotation Cluster), a system comprising both knowledge based and statistical techniques. Our system achieved the highest strict and relaxed scores for tasks A and B, and the highest relaxed score for task C.

1 Introduction

TempEval comprises novel tasks concerned with the identification of temporal relations between events and temporal expressions (TEs). The evaluation exercise includes three tasks testing the capability of participating systems to relate an event and a TE located in the same sentence (task A), an event and the TE representing the Document Creation Time (DCT) (task B), and two events located in neighbouring sentences (task C). We tackle all tasks with a mix of knowledge based and statistical techniques incorporated in our system TICTAC.

Our approach for discovering intrasentential temporal relations relies on sentence-level syntactic trees and on a bottom-up propagation of the temporal relations between syntactic constituents, by employing syntactical and lexical properties of the constituents and the relations between them. A temporal reasoning mechanism is afterwards employed to relate the two targeted temporal entities to their closest ancestor and then to each other. Conflict resolution heuristics are also applied.

In establishing a temporal relation between an event and the Document Creation Time (DCT), the temporal expressions directly or indirectly linked to that event are first analysed and, if no relation is detected, the temporal relation with the DCT is propagated top-down in the syntactic tree.

Inter-sentence temporal relations are discovered by first applying several heuristics that involve the temporal expressions and the tensed verbs of the two clauses containing the main events to be temporally related, and then by using statistical data extracted from the training corpus that revealed the most frequent temporal relation between two tensed verbs characterised by the tense information.

This paper presents the techniques employed for the three TempEval tasks (Sections 2, 3 and 4 correspond to the tasks A, B and C). The evaluation results are presented and discussed in Section 5. Conclusions are drawn in the last section.

2 Task A

Task A at TempEval involved the automatic identification of the temporal relations holding between events and all temporal expressions appearing in the same sentence. The events and TEs were annotated in the source in accordance with the TimeML standard (Pustejovsky et al., 2003a).

For all tasks, the set of temporal relations to be predicted includes: OVERLAP, BEFORE, AFTER,

Proceedings of the 4th International Workshop on Semantic Evaluations (SemEval-2007), pages 484–487,
Prague, June 2007. ©2007 Association for Computational Linguistics

BEFORE-OR-OVERLAP, OVERLAP-OR-AFTER and VAGUE.

Figure 1 depicts the processing stages involved in the identification of the temporal relation given the event, the TE and the sentence they are in. The sentence is first annotated with morpho-syntactic and functional dependency information by employing Conexor's FDG Parser (Tapanainen and Jaervinen, 1997).

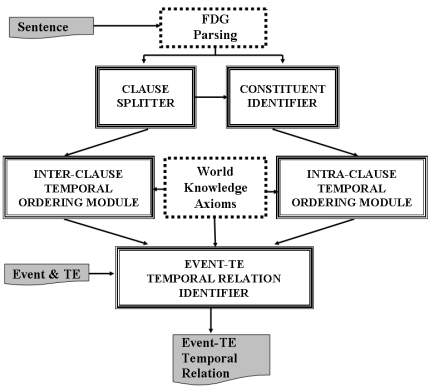

Figure 1: Task A processing stages

A clause splitter previously developed by the author is then used to detect clause boundaries and to establish the dependencies between the resulting clauses by relying on formal indicators of coordination and subordination and, in their absence, on the functional dependency relation predicted by the FDG parser. On the basis of the morpho-syntactic information we identify in each clause a set of temporally-relevant constituents (verb phrase VP, noun phrases NPs, prepositional phrases PPs, non-finite verbs and adverbial TEs).

The identified constituents and the syntactic tree of the corresponding clause are afterwards employed in a recursive bottom-up process of finding the temporal order between directly linked constituents. Each constituent is linked only with the constituent it syntactically depends on with one of the predefined temporal relations. The temporal relation is decided on the basis of heuristics that involve parameters such as: semantic properties of the two constituents' heads (whether their root forms denote reporting or aspectual start/end events), the type of the two constituents, the syntactic relation holding between them, presence of temporal signals (e.g. prepositions like *before*, *after*, *until*), the tense of the clause VP and the temporal relation between any clause TE and the DCT. This process will result in a path of temporal relations connecting every clause constituent with the clause VP.

Each pair of clauses involved in a dependency relation are then temporally related by means of their VPs' tenses, of the dependency relation between them and of their property of being reporting events or not. The underlying hypothesis is that the clause binding elements and the tenses of the two VPs provide a natural way to establish temporal relations between two syntactically related clauses. For example, in the case of an *if*-clause, its temporal relation with the superordinate clause is BEFORE. In this way, each syntactic tree branch connecting a non-root node with its father gets tagged with a temporal relation (Figure 2).

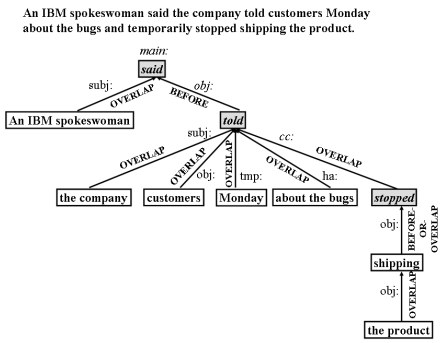

Figure 2: Temporally tagged parse tree

The final stage involves the detection of the temporal relation between a certain event and a certain TE, both situated in the sentence processed as above. The two entities are first tested to determine if they comply with world knowledge axioms that would predict their temporal relation. For example if the TE refers to a date that is previous to the DCT, and the event is a Future tensed verb, then the event-TE temporal relation is obviously AFTER. If no axiom applies to the two entities, a temporal reasoning mechanism is employed to relate the two targeted temporal entities to their closest ancestor and then to each other. If conflicts occur in relating one entity to the ancestor, priority is given to the relation linked to the entity, but if the conflict is between the event-ancestor and the TE-ancestor temporal relations, the TE-ancestor relation wins.

3 Task B

Task B consisted of the identification of temporal relations between events and the DCT. The processing stages for solving task B follow the course of the ones involved in task A, with the only difference that the inter-clause and intra-clause temporal ordering modules no longer order clauses/constituents with respect to each other and in a bottom-up manner, but with respect to the DCT going top-down through the syntactic tree and employing the knowledge gained at task A about the relative ordering between same clause constituents.

Whenever establishing a temporal relation between a constituent and the DCT, the TEs directly linked to it or situated in the same clause with it are first analysed and, if no relation can be detected, the temporal relation with the DCT is propagated top-down in the syntactic tree using the father node's temporal relation with the DCT and the temporal relation between the two constituents. For clause VPs, the relation with the DCT is found on the basis of the VP tense, the superordinate clause's VP tense, the syntactic relation connecting the clause with its superordinate and the relation between the superordinate clause's VP and the DCT.

4 Task C

For task C each pair of events signalled by the main verbs of two consecutive sentences needs to be temporally linked. This time, besides the events and TEs, the main verb in the matrix clause (matrix verb) of each sentence is also annotated in the documents.

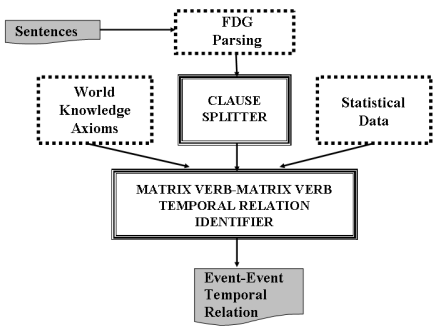

Figure 3: Task C processing stages

Figure 3 illustrates the task C processing flow. The two sentences are first parsed using Conexor's FDG Parser and then clause boundaries are identified. Due to the fact that we have noticed cases when the annotated matrix verb was not the central verb of the main clause, we have considered as matrix verb the tensed verb of the clause including the annotated matrix verb.

All TEs situated in the same clause with each matrix verb are investigated and if through these TEs and the relations between them and the matrix verbs we are able to predict a temporal relation then this relation represents the system output.

At the next stage the semantic properties of the two matrix verbs are checked to detect whether they denote reporting events or not.

If both matrix verbs are reporting events then their tense information is used to predict a relation.

If only one matrix verb is a reporting event, then we look at the TEs linked to the other matrix verb to see if we can predict the relation to the DCT. The assumption is that a reporting event is located temporally simultaneous with the DCT and, if a relation between the other event and the DCT can be established by means of surrounding TEs, then this is the relation providing us the output. If the non-reporting event can not be positioned in time with respect to the DCT by analysing surrounding TEs, then its relation with the DCT will be the one established by solving task B.

The most complicated case is the one in which both matrix verbs are non-reporting events. This case is solved by extracting statistics from the training documents, statistics involving the number of occurrences of a certain pair of verb tenses with a certain predicted temporal relation. The extracted statistics are then reconciled for tense pairs with more possible temporal relations, in the sense that if the first two most frequent possibilities have very similar frequencies, then the reconciliation is performed according to Table 1. In this manner a temporal relation is associated to each tense pair and, consequently, the temporal relation between the two matrix verbs is identified.

5 Results

The test corpus consists of 20 articles from TimeBank (Pustejovsky et al., 2003b). The performance is assessed with three evaluation

Temporal Relation	Temporal Relation	Reconciled Relation
OVERLAP	BEFORE-OR-OVERLAP	BEFORE-OR-OVERLAP
OVERLAP	BEFORE	BEFORE-OR-OVERLAP
OVERLAP	OVERLAP-OR-AFTER	OVERLAP-OR-AFTER
OVERLAP	AFTER	OVERLAP-OR-AFTER
BEFORE	BEFORE-OR-OVERLAP	BEFORE-OR-OVERLAP
AFTER	OVERLAP-OR-AFTER	OVERLAP-OR-AFTER
VAGUE	any relation	any relation

Table 1: Reconciliation between temporal relations

metrics (precision, recall, f-measure) and two scoring schemes (strict, relaxed). The strict scoring scheme counts only exact matches, while the relaxed one gives credit to partial semantic matches too.

The following three tables illustrate for each task the results our team obtained at TempEval, the baseline, the minimum and maximum values achieved by participating systems.

For each of the three tasks, the baseline is established by the most frequent temporal relation encountered in that task's training data. In the case of task A the most frequent temporal relation present in the training data is OVERLAP, in the case of task B BEFORE and for task C OVERLAP.

TASK A	STRICT SCORE			RELAXED SCORE		
	P	R	F	P	R	F
WVALI	*0.62*	*0.62*	***0.62***	*0.64*	*0.64*	***0.64***
BASELINE	0.49	0.49	**0.49**	0.51	0.51	**0.51**
MIN	0.53	0.25	**0.34**	0.60	0.30	**0.41**
MAX	0.62	0.62	**0.62**	0.64	0.64	**0.64**

Table 2: Results for task A

Our system achieved the highest strict and relaxed f-measure scores in tasks A and B, with the task B results substantially above the baseline (18%).

TASK B	STRICT SCORE			RELAXED SCORE		
	P	R	F	P	R	F
WVALI	*0.80*	*0.80*	***0.80***	*0.81*	*0.81*	***0.81***
BASELINE	0.62	0.62	**0.62**	0.62	0.62	**0.62**
MIN	0.73	0.57	**0.66**	0.74	0.62	**0.71**
MAX	0.80	0.80	**0.80**	0.84	0.81	**0.81**

Table 3: Results for task B

TASK C	STRICT SCORE			RELAXED SCORE		
	P	R	F	P	R	F
WVALI	*0.54*	*0.54*	***0.54***	*0.66*	*0.66*	***0.66***
BASELINE	0.42	0.42	**0.42**	0.46	0.46	**0.46**
MIN	0.42	0.42	**0.42**	0.56	0.56	**0.56**
MAX	0.55	0.55	**0.55**	0.66	0.66	**0.66**

Table 4: Results for task C

Despite the challenges posed by task C, our system achieved the best relaxed score among all participants, as well as a high strict score.

6 Conclusion

This paper presented our approach and participation in the TempEval evaluation exercise involving the identification of event-time and event-event temporal relations. We propose an approach mainly based on syntactical properties, combining knowledge-based and statistical techniques, all included in our automatic temporal annotation system TICTAC. Our system participated in all three TempEval tasks.

When compared to the other systems participating in this competition, we have obtained the highest results both in the strict and relaxed scoring schemes in the case of tasks A and B, as well as in the relaxed scoring scheme for task C. Therefore, we conclude that the proposed approach is appropriate for the TempEval tasks and we plan to find ways of improving the system's performance.

Several future work directions emerge naturally from a first look and shallow analysis of the results. Firstly, we would like to carry out an in-depth study of other possible correlations between syntax and temporality. Secondly, we aim at exploiting apart from the syntax of the analysed text, more of its semantics.

Acknowledgments

The author would like to thank Patricio Martinez Barco, Andres Montoyo and Constantin Orasan for their advice and support.

References

J. Pustejovsky, J. Castano, R. Ingria, R. Sauri, R. Gaizauskas, A. Setzer, and G. Katz. 2003a. TimeML: Robust Specification of Event and Temporal Expressions in Text. In *Proceedings of the Fifth International Workshop on Computational Semantics (IWCS-5)*.

J. Pustejovsky, P. Hanks, R. Sauri, A. See, R. Gaizuaskas, A. Setzer, D. Radev, B. Sundheim, D. Day, L. Ferro, and M. Lazo. 2003b. The TIMEBANK Corpus. In *Proceedings of Corpus Linguistics 2003*.

P. Tapanainen and T. Jaervinen. 1997. A non–projective dependency parser. In *Proceedings of the 5th Conference of Applied Natural Language Processing, ACL*.