

Document Level Time-anchoring for TimeLine Extraction

Egoitz Laparra, Itziar Aldabe, German Rigau

IXA NLP group, University of the Basque Country (UPV/EHU)

{egoitz.laparra, itziar.algabe, german.rigau}@ehu.eus

Abstract

This paper investigates the contribution of document level processing of time-anchors for TimeLine event extraction. We developed and tested two different systems. The first one is a baseline system that captures explicit time-anchors. The second one extends the baseline system by also capturing implicit time relations. We have evaluated both approaches in the SemEval 2015 task 4 *TimeLine: Cross-Document Event Ordering*. We empirically demonstrate that the document-based approach obtains a much more complete time anchoring. Moreover, this approach almost doubles the performance of the systems that participated in the task.

1 Introduction

Temporal relation extraction has been the topic of different SemEval tasks (Verhagen et al., 2007; Verhagen et al., 2010; UzZaman et al., 2013; Llorens et al., 2015) and other challenges as the 6th i2b2 NLP Challenge (Sun et al., 2013). These tasks focused mainly on the temporal relations of the events with respect to other events or time expressions, and their goals are to discover which of them occur before, after or simultaneously to others. Recently, SemEval 2015 included a novel task regarding temporal information extraction (Minard et al., 2015). The aim of SemEval 2015 task 4 is to order in a TimeLine the events in which a target entity is involved and presents some significant differences with respect to previous exercises. First, the temporal information must be recovered from different sources in a cross-document way. Second, the TimeLines are focused on the events involving just a given entity. Finally, unlike previous challenges, SemEval 2015 task 4 requires a quite complete time anchoring. This work focuses

mainly on this latter point. We show that the temporal relations that explicitly connect events and time expressions are not enough to obtain a full time-anchor annotation and, consequently, produce incomplete TimeLines. We propose that for a complete time-anchoring the temporal analysis must be performed at a document level in order to discover implicit temporal relations. We present a preliminary approach that obtains, by far, the best results on the main track of SemEval 2015 task 4.

2 Related work

The present work is closely related to previous approaches involved in TempEval campaigns (Verhagen et al., 2007; Verhagen et al., 2010; UzZaman et al., 2013; Llorens et al., 2015). In these works, the problem can be seen as a classification task for deciding the type of the temporal link that connects two different events or an event and a temporal expression. For that reason, the task has been mainly addressed using supervised techniques. For example, (Mani et al., 2006; Mani et al., 2007) trained a MaxEnt classifier using training data which were bootstrapped by applying temporal closure. (Chambers et al., 2007) focused on event-event relations using previously learned event attributes. More recently, (D'Souza and Ng, 2013) combined hand-coded rules with some semantic and discourse features. (Laokulrat et al., 2013) obtained the best results on TempEval 2013 annotating sentences with predicate-role structures, while (Mirza and Tonelli, 2014) affirm that using a simple feature set results in better performances.

However, recent works like (Chambers et al., 2014) have pointed out that these tasks cover just a part of all the temporal relations that can be inferred from the documents. Furthermore, time-anchoring is just a part of the works presented above. Our approach aims to extend these strategies and it is based on other research lines

involving the extraction of implicit information (Palmer et al., 1986; Whittemore et al., 1991; Tetreault, 2002). Particularly, we are inspired by recent works on Implicit Semantic Role Labelling (ISRL) (Gerber and Chai, 2012) and very specially on the work by (Blanco and Moldovan, 2014) who adapted the ideas about ISRL to focus on modifiers, including arguments of time, instead of core arguments or roles. As the SemEval 2015 task 4 does not include any training data we decided to develop a deterministic algorithm of the type of (Laparra and Rigau, 2013) for ISRL.

3 TimeLine: Cross-Document Event Ordering

In the SemEval task 4 *TimeLine: Cross-Document Event Ordering* (Minard et al., 2015), given a set of documents and a target entity, the aim is to build a TimeLine by detecting the events in which the entity is involved and anchoring these events to normalized times. Thus, a TimeLine is a collection of ordered events in time relevant for a particular entity. TimeLines contain relevant events in which the target entity participates as ARG0 (i.e. agent) or ARG1 (i.e. patient) as defined in Prop-Bank (Palmer et al., 2005).¹ The target entities can be *people, organization, product* or *financial* entities and the annotation of time anchors is based on TimeML.

For example, given the entity *Steve Jobs*, a TimeLine contains the events with the associated ordering in the TimeLine and the time anchor:

1	2004	18135-7-fighting
2	2005-06-05	1664-2-keynote
...		
4	2011-08-24	18315-2-step_down

The dataset used for the task is composed of articles from Wikinews. The trial data consists of 30 documents about “Apple Inc.” and gold standard TimeLines for six target entities. The test corpus consists of 3 sets of 30 documents around three topics and 38 target entities. The topics are “Airbus and Boeing”, “General Motors, Chrysler and Ford” and “Stock Market”.

The evaluation used in the task is based on the metric previously introduced in TempEval-3 (UzZaman et al., 2013). The metric captures the tem-

poral awareness of an annotation (UzZaman and Allen, 2011) based on temporal closure graphs. In order to calculate the precision, recall and F1 score, the TimeLines are first transformed into a graph representation. For that, the time anchors are represented as TIMEX3 and the events are related to the corresponding TIMEX3 by means of the SIMULTANEOUS relation type. In addition, BEFORE relation types are created to represent that one event happens before another one and SIMULTANEOUS relation types to refer to events happening at the same time. The official scores are based on the micro-average of F1 scores.

The main track of the task (Track A) consists of building TimeLines providing only the raw text sources. Two systems participated in the task. The organisers also defined a Track B where gold event mentions were given. In this case, two different systems sent results. For both tracks, a sub-track in which the events are not associated to a time anchor was also presented.

In this work, we focus on the main track of the task. We believe the main track is the most challenging one as no annotated data is provided. Indeed, WHUNLP 1 was the best run and achieved an F1 of 7.28%.

Three runs were submitted. The WHUNLP team used the Stanford CoreNLP and they applied a rule-based approach to extract the entities and their predicates. They also performed temporal reasoning.² The remaining two runs were submitted using the SPINOZA_VU system (Caselli et al., 2015). They performed entity resolution, event detection, event-participant linking, coreference resolution, factuality profiling and temporal processing at document and cross-document level. Then, the TimeLine extractor built a global timeline between all events and temporal expressions regardless of the target entities and then it extracted the target entities for the TimeLines. The participants also presented an out of the competition system which anchors events to temporal expressions appearing not only in the same sentence but also in the previous and following sentences.

4 Baseline TimeLine extraction

In this section we present a system that builds TimeLines which contain events with explicit time-anchors. We have defined a three step pro-

¹For more information consult <http://tinyurl.com/owuyuybb>

²Unfortunately, the task participants did not submit a paper with the description of the system.

cess to build TimeLines. Given a set of documents and a target entity, the system first obtains the events in which the entity is involved. Second, it obtains the time-anchors for each of these events. Finally, it sorts the events according to their time-anchors. For steps 1 and 2 we apply a pipeline of tools (cf. section 4.1) that provides annotations at different levels.

4.1 NLP processing

Detecting mentions of events, entities and time expressions in text requires the combination of various Natural Language Processing (NLP) modules. We apply a generic pipeline of linguistic tools that includes Named-Entity Recognition (NER) and Disambiguation (NED), Co-reference Resolution (CR), Semantic Role Labelling (SRL), Time Expressions Identification (TEI) and Normalization (TEN), and Temporal Relation Extraction (TRE). The NLP processing is based on the NewsReader pipeline (Agerri et al., 2014a), version 2.1. Next, we present the different tools in our pipeline.

Named-Entity Recognition (NER) and Disambiguation (NED): We perform NER using the *ixa-pipe-nerc* that is part of *IXA pipes* (Agerri et al., 2014b). The module provides very fast models with high performances, obtaining 84.53 in F1 on CoNLL tasks. Our NED module is based on *DBpedia Spotlight* (Daiber et al., 2013). We have created a NED client to query the *DBpedia Spotlight* server for the Named entities detected by the *ixa-pipe-nerc* module. Using the best parameter combination, the best results obtained by this module on the TAC 2011 dataset were 79.77 in precision and 60.67 in recall. The best performance on the AIDA dataset is 79.67 in precision and 76.94 in recall.

Coreference Resolution (CR): In this case, we use a coreference module that is loosely based on the *Stanford Multi Sieve Pass* system (Lee et al., 2011). The system consists of a number of rule-based sieves that are applied in a deterministic manner. The system scores 56.4 F1 on CoNLL 2011 task, around 3 points worse than the system by (Lee et al., 2011).

Semantic Role Labelling (SRL): SRL is performed using the system included in the *MATE-tools* (Björkelund et al., 2009). This system reported on the CoNLL 2009 Shared Task a labelled semantic F1 of 85.63 for English.

Time Expression Identification (TEI) and

Normalization (TEN): We use the time module from *TextPro* suite (Pianta et al., 2008) to capture the tokens corresponding to temporal expressions and to normalize them following *TIDES* specification. This module is trained on *TempEval3* data. The average results for English is: 83.81% precision, 75.94% recall and 79.61% F1 values.

Time Relation Extraction (TRE): We apply the temporal relation extractor module from *TextPro* to extract and classify temporal relations between an event and a time expression. This module is trained using *yamcha* tool on the *TempEval3* data. The result for relation classification on the corpus of *TempEval3* is: 58.8% precision, 58.2% recall and 58.5% F1.

4.2 TimeLine extraction

Our TimeLine extraction system uses the linguistic information provided by the pipeline. The process to extract the target entities, the events and time-anchors can be described as follows:

(1) **Target entity identification:** The target entities are identified by the NED module. As they can be expressed in several forms, we use the redirect links contained in *DBpedia* to extend the search of the events involving those target entities. For example, if the target entity is *Toyota* the system would also include events involving the entities *Toyota Motor Company* or *Toyota Motor Corp.* In addition, as the NED does not always provide a link to *DBpedia*, we also consider the matching of the wordform of the head of the argument with the head of the target entity.

(2) **Event selection:** We use the output of the SRL module to extract the events that occur in a document. Given a target entity, we combine the output of the NER, NED, CR and SRL to obtain those events that have the target entity as filler of their ARG0 or ARG1. We also set some constraints to select certain events according to the specification of the *SemEval* task. That is, we only return those events that are not negated and are not accompanied by modal verbs except *will*.

(3) **Time-anchoring:** We extract the time-anchors from the output of the TRE and SRL. From the TRE, we extract as time-anchors those relations between events and time-expressions identified as *SIMULTANEOUS*. From the SRL, we extract as time-anchors those ARG-TMP related to time expressions. In both cases we use the time-expression returned by the TEI module. The

tests performed on the trial data show that the best choice for time-anchoring is combining both options. For each time anchor we normalize the time expression using the output of the TEN module.

The TimeLine extraction process described following this approach builds TimeLines for events with explicit time-anchors. We call this system **BTE** and it can be seen as a baseline since we believe that the temporal analysis should be carried out at document level. Section 5 presents our strategy for improving the time-anchoring carried out by our baseline system.

5 Document level time-anchoring

The explicit time anchors provided by the NLP tools presented in Section 4.1 do not cover the full set of events involving a particular entity. That is, most of the events do not have an explicit time anchor and therefore are not captured as part of the TimeLine of that entity. Thus, we need to recover the time-anchors that appear implicitly in the text. In this preliminary work, we propose a simple strategy that tries to capture implicit time-anchors while maintaining the coherence of the temporal information in the document. As said in Section 2, this strategy follows previous works on Implicit Semantic Role Labelling.

The rationale behind the algorithm 1 is that by default the events of an entity that appear in a document tend to occur at the same time as previous events involving the same entity, except stated explicitly. For example, in Figure 1 all the events involving *Steve Jobs*, like *gave* and *announced*, are anchored to the same time-expression *Monday* although this only happens explicitly for the first event *gave*. The example also shows how for other events that occur in different times the time-anchor is also mentioned explicitly, like for those events that involve the entities *Tiger* and *Mac OS X Leopard*.

Algorithm 1 starts from the annotation obtained by the tools described in Section 4.1. For a particular entity a list of events (*eventList*) is created sorted by its occurrence in the text. Then, for each event in this list the system checks if that event has already a time-anchor (*eAnchor*). If this is the case, the time-anchor is included in the list of default time-anchors (*defaultAnchor*) for the following events of the entity with the same verb tense (*eTense*). If the event does not have an explicit time-anchor but the system has found

a time-anchor for a previous event belonging to the same tense (*defaultAnchor[eTense]*), this time-anchor is also assigned to the current event (*eAnchor*). If none of the previous conditions satisfy, the algorithm anchors the event to the **Document Creation Time** (DCT) and sets this time-expression as the default time-anchor for the following events with the same tense.

Algorithm 1 Implicit Time-anchoring

```

1: eventList = sorted list of events of an entity
2: for event in eventList do
3:   eAnchor = time anchor of event
4:   eTense = verb tense of event
5:   if eAnchor not NULL then
6:     defaultAnchor[eTense] = eAnchor
7:   else if defaultAnchor[eTense] not NULL then
8:     eAnchor = defaultAnchor[eTense]
9:   else
10:    eAnchor = DCT
11:    defaultAnchor[eTense] = DCT
12:   end if
13: end for

```

Note that the algorithm 1 strongly depends on the tense of the events. As this information can be only recovered from verbal predicates, this strategy cannot be applied to events described by nominal predicates. For these cases just explicit time-anchors are taken into account.

The TimeLine is built ordering the events according to the time-anchors obtained both explicitly and implicitly. We call this system **DLT**.

6 Experiments

We have evaluated our two TimeLine extractors on the main track of the SemEval 2015 task 4. Two systems participated in this track, **WHUNLP** and **SPINOZAVU**, with three runs in total. Their performances in terms of Precision (P), Recall (R) and F1-score (F1) are presented in Table 6. We also present in italics additional results of both systems. On the one hand, the results of a corrected run of the WHUNLP system provided by the SemEval organizers. On the other hand, the results of an out of the competition version of the SPINOZAVU team explained in (Caselli et al., 2015). The best run is obtained by the corrected version of **WHUNLP_1** with an F1 of 7.85%. The low figures obtained show the intrinsic difficulty of the task, specially in terms of Recall.

Apple Computer CEO and co-founder **Steve Jobs** gave his annual opening keynote to the World Wide Developers Conference (WWDC) at Moscone Center in San Francisco, California on **Monday**...

Moving on, **Jobs** announced that there have been 2 million copies of **Tiger** sold in the **6 weeks** that it has been available....

Steve announced that **Mac OS X Leopard** would be released in **2007**

Figure 1: Example of document-level time-anchoring.

Table 6 also contains the results obtained by our systems. We present two different runs. On the one hand, we present the results obtained using just the explicit time-anchors provided by **BTE**. As it can be seen, the results obtained by this run are similar to those obtained by **WHUNLP.1**. On the other hand, the results of the implicit time-anchoring approach (**DLT**) outperforms by far our baseline and all previous systems applied to the task. To check that these results are not biased by the time-relation extractor we use in our pipeline (TimePro), we reproduce the performances of BTE and DLT using another system to obtain the time-relations. For this purpose we have used CAEVO by (Chambers et al., 2014). The results obtained in this case show that the improvement obtained by our proposal is quite similar, regardless of the time-relation extractor chosen.

System	P	R	F1
SPINOZAVU-RUN-1	7.95	1.96	3.15
SPINOZAVU-RUN-2	8.16	0.56	1.05
WHUNLP.1	14.10	4.90	7.28
<hr/>			
<i>OC.SPINOZA.VU</i>	-	-	7.12
<i>WHUNLP.1</i>	14.59	5.37	7.85
<hr/>			
BTE	26.42	4.44	7.60
DLT	20.67	10.95	14.31
<hr/>			
BTE.caevo	17.56	4.86	7.61
DLT.caevo	17.02	12.09	14.13

Table 1: Results on the SemEval-2015 task

The figures in Table 6 seem to prove our hypothesis. In order to obtain a full time-anchoring annotation, the temporal analysis must be carried out at a document level. The TimeLine extractor almost doubles the performance by just including a straightforward strategy as the one described in Section 5. As expected, Table 6 shows that this

improvement is much more significant in terms of Recall.

7 Conclusion and future-work

In this work we have shown that explicit temporal relations are not enough to obtain a full time-anchor annotation of events. We have proved the need of a temporal analysis at document level. For that, we have proposed a simple strategy that acquires implicit relations and it obtains a more complete time-anchoring.³ The approach has been evaluated on the TimeLine extraction task and the results show that the performance can be doubled when using implicit relations. As future work, we plan to explore in more detail this research line by applying more sophisticated approaches in the temporal analysis at document level.

However, this is not the only research line that we want to go in depth. The errors that the tools of the pipeline are producing have a direct impact on the final result of our TimeLine extractors. In a preliminary analysis, we have noticed that this is specially critical when detecting the events given a target entity. Our pipeline does not detect all mentions of the target entities. That is why we are planning an in-depth error analysis of the pipeline in order to find the best strategy to improve on the linguist analyses and the TimeLine extraction.

8 Acknowledgment

We are grateful to the anonymous reviewers for their insightful comments. This work has been partially funded by SKaTer (TIN2012-38584-C06-02) and NewsReader (FP7-ICT-2011-8-316404), as well as the READERS project with the financial support of MINECO (CHIST-ERA READERS project - PCIN-2013-002-C02-01).

³Publicly available at <http://adimen.si.ehu.es/web/DLT>

References

- Rodrigo Agerri, Itziar Aldabe, Zuhaitz Beloki, Egoitz Laparra, Maddalen Lopez de Lacalle, German Rigau, Aitor Soroa, Antske Fokkens, Ruben Izquierdo, Marieke van Erp, Piek Vossen, Christian Girardi, and Anne-Lyse Minard. 2014a. Event detection, version 2. Newsreader Deliverable 4.2.2. <http://www.newsreader-project.eu/files/2012/12/NWR-D4-2-2.pdf>.
- Rodrigo Agerri, Josu Bermudez, and German Rigau. 2014b. IXA pipeline: Efficient and Ready to Use Multilingual NLP tools. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC-2014)*. 00013.
- Anders Björkelund, Love Hafdel, and Pierre Nugues. 2009. Multilingual semantic role labeling. In *Proceedings of the Thirteenth Conference on Computational Natural Language Learning: Shared Task, CoNLL '09*, pages 43–48, Boulder, Colorado, USA.
- Eduardo Blanco and Dan Moldovan. 2014. Leveraging verb-argument structures to infer semantic relations. In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics*, pages 145–154, Gothenburg, Sweden.
- Tommaso Caselli, Antske Fokkens, Roser Morante, and Piek Vossen. 2015. SPINOZA_VU: An nlp pipeline for cross document timelines. In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*, pages 786–790, Denver, Colorado, June 4–5.
- Nathanael Chambers, Shan Wang, and Dan Jurafsky. 2007. Classifying temporal relations between events. In *Proceedings of the 45th Annual Meeting of the ACL on Interactive Poster and Demonstration Sessions, ACL'07*, pages 173–176, Prague, Czech Republic.
- Nathanael Chambers, Taylor Cassidy, Bill McDowell, and Steven Bethard. 2014. Dense event ordering with a multi-pass architecture. *Transactions of the Association for Computational Linguistics*, 2:273–284.
- Joachim Daiber, Max Jakob, Chris Hokamp, and Pablo N. Mendes. 2013. Improving efficiency and accuracy in multilingual entity extraction. In *Proceedings of the 9th International Conference on Semantic Systems (I-Semantics)*.
- Jennifer D'Souza and Vincent Ng. 2013. Classifying temporal relations with rich linguistic knowledge. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL'13*, pages 918–927, Atlanta, Georgia.
- Matthew Gerber and Joyce Chai. 2012. Semantic role labeling of implicit arguments for nominal predicates. *Computational Linguistics*, 38(4):755–798, December.
- Natsuda Laokulrat, Makoto Miwa, Yoshimasa Tsuruoka, and Takashi Chikayama. 2013. Uttime: Temporal relation classification using deep syntactic features. In *Second Joint Conference on Lexical and Computational Semantics (*SEM), Volume 2: Proceedings of the Seventh International Workshop on Semantic Evaluation (SemEval 2013)*, pages 88–92, Atlanta, Georgia, USA.
- Egoitz Laparra and German Rigau. 2013. Impar: A deterministic algorithm for implicit semantic role labelling. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (ACL 2013)*, pages 33–41.
- Heeyoung Lee, Yves Peirsman, Angel Chang, Nathanael Chambers, Mihai Surdeanu, and Dan Jurafsky. 2011. Stanford's multi-pass sieve coreference resolution system at the conll-2011 shared task. In *Proceedings of the Fifteenth Conference on Computational Natural Language Learning: Shared Task, CONLL Shared Task '11*, Portland, Oregon.
- Hector Llorens, Nathanael Chambers, Naushad UzZaman, Nasrin Mostafazadeh, James Allen, and James Pustejovsky. 2015. Semeval-2015 task 5: Qa tempeval - evaluating temporal information understanding with question answering. In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*, pages 792–800, Denver, Colorado, June.
- Inderjeet Mani, Marc Verhagen, Ben Wellner, Chong Min Lee, and James Pustejovsky. 2006. Machine learning of temporal relations. In *Proceedings of the 21st International Conference on Computational Linguistics and the 44th Annual Meeting of the Association for Computational Linguistics, ACL'06*, pages 753–760, Sydney, Australia.
- Inderjeet Mani, Ben Wellner, Marc Verhagen, and James Pustejovsky. 2007. Three approaches to learning tlinks in timeml. Technical report.
- Anne-Lyse Minard, Manuela Speranza, Eneko Agirre, Itziar Aldabe, Marieke van Erp, Bernardo Magnini, German Rigau, and Ruben Urizar. 2015. Semeval-2015 task 4: Timeline: Cross-document event ordering. In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*, pages 778–786, Denver, Colorado, June 4–5.
- Paramita Mirza and Sara Tonelli. 2014. Classifying temporal relations with simple features. In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics*, pages 308–317, Gothenburg, Sweden, April. Association for Computational Linguistics.
- Martha S. Palmer, Deborah A. Dahl, Rebecca J. Schiffman, Lynette Hirschman, Marcia Linebarger, and John Dowding. 1986. Recovering implicit information. In *Proceedings of the 24th annual meeting on Association for Computational Linguistics, ACL '86*, pages 10–19, New York, New York, USA.

- Martha Palmer, Daniel Gildea, and Paul Kingsbury. 2005. The proposition bank: An annotated corpus of semantic roles. *Computational Linguistics*, 31(1):71–106, March.
- Emanuele Pianta, Christian Girardi, and Roberto Zanoli. 2008. The textpro tool suite. In *Proceedings of the Sixth International Conference on Language Resources and Evaluation (LREC'08)*, Marrakech, Morocco, may.
- Weiyi Sun, Anna Rumshisky, and Ozlem Uzuner. 2013. Evaluating temporal relations in clinical text: 2012 i2b2 Challenge. *Journal of the American Medical Informatics Association*, 20(5):806–813, September.
- Joel R. Tetreault. 2002. Implicit role reference. In *International Symposium on Reference Resolution for Natural Language Processing*, pages 109–115, Alicante, Spain.
- Naushad UzZaman and James Allen. 2011. Temporal evaluation. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 351–356, Portland, Oregon, USA.
- Naushad UzZaman, Hector Llorens, Leon Derczynski, James Allen, Marc Verhagen, and James Pustejovsky. 2013. Semeval-2013 task 1: Tempeval-3: Evaluating time expressions, events, and temporal relations. In *Second Joint Conference on Lexical and Computational Semantics (*SEM), Volume 2: Proceedings of the Seventh International Workshop on Semantic Evaluation (SemEval 2013)*, SemEval '13, pages 1–9, Atlanta, Georgia, USA.
- Marc Verhagen, Robert Gaizauskas, Frank Schilder, Mark Hepple, Graham Katz, and James Pustejovsky. 2007. Semeval-2007 task 15: Tempeval temporal relation identification. In *Proceedings of the 4th International Workshop on Semantic Evaluations, SemEval '07*, pages 75–80, Prague, Czech Republic.
- Marc Verhagen, Roser Saurí, Tommaso Caselli, and James Pustejovsky. 2010. Semeval-2010 task 13: Tempeval-2. In *Proceedings of the 5th International Workshop on Semantic Evaluation, SemEval '10*, pages 57–62, Los Angeles, California.
- Greg Whittemore, Melissa Macpherson, and Greg Carlson. 1991. Event-building through role-filling and anaphora resolution. In *Proceedings of the 29th annual meeting on Association for Computational Linguistics, ACL '91*, pages 17–24, Berkeley, California, USA.