

A Bilingual Corpus of Inter-linked Events

Tommaso Caselli, Nancy Ide, Roberto Bartolini

ILC-CNR, Dep. Computer Science, Vassar College, ILC-CNR
Via Moruzzi, 1 56123 Pisa ITALY - Poughkeepsie, New York USA
tommaso.caselli@ilc.cnr.it, ide@cs.vassar.edu, roberto.bartolini@ilc.cnr.it

Abstract

This paper describes the creation of a bilingual corpus of inter-linked events for Italian and English. Linkage is accomplished through the Inter-Lingual Index (ILI) that links ItalWordNet with WordNet. The availability of this resource, on the one hand, enables contrastive analysis of the linguistic phenomena surrounding events in both languages, and on the other hand, can be used to perform multilingual temporal analysis of texts. In addition to describing the methodology for construction of the inter-linked corpus and the analysis of the data collected, we demonstrate that the ILI could potentially be used to bootstrap the creation of comparable corpora by exporting layers of annotation for words that have the same sense.

1. Introduction

Identification of temporal relations in texts is required to improve the performance of natural language processing (NLP) applications such as Information Extraction or Open-Domain Question-Answering. To perform this task, one must first identify and classify events in the text. In this paper we describe the creation of a bi-lingual corpus of Italian and English that is inter-linked by event types. The resource was created from the monolingual English TimeBank corpus (Pustejovsky et al., 2003) and an Italian counterpart that is comparable to TimeBank in terms of content and annotations. Linkage is accomplished through the Inter-Lingual Index (ILI) that in turn links ItalWordNet with WordNet.

In this paper, we describe our methodology for constructing the inter-linked corpus and provide a contrastive analysis of the data. We also demonstrate how the ILI could potentially be used to bootstrap the creation of comparable corpora by exporting layers of annotation for words that have the same sense.

2. TimeML

TimeML is a mark-up language specifically designed for the annotation of events, temporal expressions and their ordering. With respect to other existing annotation schemes, like (Katz and Arosio, 2001), (Filatova and Hovy, 2001), (Schilder and Habel, 2001), TimeML presents some important advantages, in particular:

- it identifies signals, i.e. textual elements which make explicit the relation holding between two entities (e.g. temporal prepositions and connectives, like *for*, *during*, *before*...);
- it identifies a wide range of linguistic expressions realizing events (verbs, nouns, including *nomina actionis* and other types of nominalizations, stative adjectives and some occurrences of prepositional phrases) and assigns them to one of the seven classes or *types*, i.e. ASPECTUAL, REPORTING, LSTATE, LACTION, PERCEPTION, STATE, and OCCUR-

RENCE, according to semantic and syntactic criteria¹. For instance, an event is assigned to LACTION if it describes an action and introduces another event as its argument, which is explicitly stated in the text;

- it creates dependencies between events, and between events and times.

The term *event* is very broadly defined in TimeML as any “situation that happens or occurs” (Sauri et al., 2006), including states, that is, circumstances in which something obtains or holds true if it is *temporally located* in the text. The constraint that an event is temporally located in the text restricts the set of states to temporary states, such as “being hungry”, and to states linked (explicitly or implicitly) to a temporal expression in the text. Absolute or permanent states, like “being tall”, are excluded. However, although TimeML’s definition of *event* is useful because it aims at capturing every event instance in a text, we will show that over-interpretation of the definition can lead to biases in the annotation process.

An advantage of the event classes proposed in TimeML is that they are language independent and indicate the relationships the event participates such as *factual*, *reported*, *intensional*.

3. TimeBank 1.2 and the Italian TimeBank

The TimeBank 1.2. corpus consists of 183 news articles (61K words) drawn from various sources and including broadcast news, newswire, and the Penn TreeBank2 *Wall Street Journal* texts. All of the texts are annotated with temporal information, including events, temporal expressions, and temporal links between events and temporal expressions, using the TimeML specifications. A total of 7,935 events are annotated in TimeBank, with an agreement of 0.81 on partial match (0.78 for total match), and 0.67 for event classification.

The Italian TimeBank (ITB) consists of 171 newspaper articles (62K words) collected from the Italian Treebank (Montemagni et al., 2003) and the PAROLE corpus (Marinelli et

¹See (Sauri et al., 2006) for a full description of the TimeML.

al., 2003)². The articles were carefully selected to be comparable in content to the TimeBank in order to enable contrastive analysis of temporal phenomena and events. The articles were then manually annotated for events using the TimeML annotation scheme. Annotation of the entire corpus is on-going; to date, 1,755 events (307 temporal expressions and 303 signals) have been annotated in ITB.

To apply TimeML to Italian, it was necessary to adapt TimeML to account for language-specific phenomena, including the addition of more detailed information in attributes on EVENT tags and the handling of modal verbs. In particular, the values of the TENSE attribute were expanded to include *imperfect*, and two new attributes were introduced: VERB_FORM, for non-finite verb forms, and MOOD, to indicate the presence of inflectional modality on the verb. In addition, to reflect the fact that in Italian modal verbs are considered real verbs and not auxiliaries, they are annotated with EVENT tags.

To reduce disagreement among annotators, we developed a mapping between the TimeML event classes and the PAROLE/SIMPLE/CLIPS³ semantic types for events (Caselli et al., 2007), which provided us with a set of heuristics to assist annotators in assigning the appropriate TimeML class to each event in the Italian corpus. Evaluation of inter-annotator agreement shows an accuracy rate of 81.17% on event match and 0.84 on event classification.

4. Linking events: methodology

Linkage of events in TimeBank and ITB is accomplished through the Inter-Lingual Index (ILI) developed in the EuroWordNet Project. The ILI is effectively an unstructured version of WordNet that is used as a “hub” through which WordNet1.5 synsets are associated with synsets in wordnets in other languages, thereby linking the lexical realizations of what may be regarded as language-independent concepts in multiple languages. In ItalWordNet, each ILI has been augmented with several semantic relations, such as *eq_synonym*, *eq_hyperonym*, *eq_cause*..., thus providing not only a mapping, but also more specific information about the relation(s) between synsets in English and Italian.

As an initial experiment, a subset of events identified in TimeBank was chosen, representing the event types given in Table 1.

OCCURRENCE	24%
I.ACTION	29%
I.STATE	33%
STATE	11%
ASPECTUAL	31%
REPORTING	12%

Table 1: Event types in TimeBank.

The resulting set consisted of 1,835 event instances, including 1,177 verbs and 658 nominalizations, represent-

²The Italian TimeBank is currently under development, to date, 13,352 words of text have been annotated.

³Referred to as SIMPLE in the remainder of this paper.

ing approximately 20% of the annotated events in TimeBank. Two annotators then manually assigned WordNet2.0 senses to each of the TimeBank event instances in the set. Inter-annotator agreement was over 91%, which is significantly higher than inter-annotator agreement for general sense assignment. This is likely due to the limited domain of the texts in the corpus and the fact that the occurrences were further constrained to include only events. For this study, we retained only those instances where both annotators agreed on the WordNet sense assignment, yielding a total of 1,686 sense-annotated events.

Assignment of senses to the Italian data was done semi-automatically, since a large part of the Italian Treebank has already been manually annotated with ItalWordNet senses. Manual annotation was done only for events to which no sense had been previously assigned. A total of 1,253 event tokens (71.73% of the total) were assigned a sense, including 778 verbs and 463 nominalizations and nouns. Events realized by prepositional phrases (e.g. *a bordo* [on board]) and numbers were not assigned a sense.

Each sense tag in TimeBank and ITB was then associated with the ILI link provided in ItalWordNet. The ILI link was determined automatically and restricted to the *eq_synonym* and *eq_near_synonym* relations, thereby linking only those events that have (exactly or approximately) the same meaning. Because the ILI links synsets rather than individual words, a relationship between an English word sense and an Italian word sense enabled linkage between all words in the synset for each language (e.g., on the basis of a link between *comfort1* and *rincuorare1*, we were able to link *comfort1*, *console1*, *solace1*, *soothe1* and *rincuorare1*, *consolare1*, *confortare1*, *rianimare2*). As a result, the procedure assigned a hypothetical reliable ILI link to 1,103 events in TimeBank, for a total of 115 different event synsets, and 1,250 events in ITB for a total of 653 event synsets.

5. Experiments and analysis

We conducted two sets of experiments using the event links. We first analyzed events with same ILI between which a *eq_synonym* or *eq_near_synonym* relation holds in TimeBank and ITB, in order to determine the extent to which the introduction of wordnet senses is useful for event identification. We also experimented with the use of the ILI as a bootstrapping device for the creation of comparable language resources by exporting layers of annotation for words that have the same sense.

5.1. Similar events in the two corpora

We first examined events realized by verbs and nouns independently, and later evaluated the usefulness of wordnet senses for event identification for both verbs and nouns. We automatically extracted all verbs and nouns with a common ILI in the two corpora, which resulted in 56 common event synsets. *Partial matches*, that is, cases where a single synset in English corresponds more than one synset in Italian and at least one pair of English-Italian synsets is mapped to the same ILI, and vice versa, were counted as common synsets, reflecting the fact that we considered the relations of *eq_synonym* and *eq_near_synonym* as equivalent. The relatively low number of common events is due

to the sparseness of data in the Italian corpus.

5.1.1. Verbs

We identified 35 common event types realized by verbs. Our goal was to verify the degree to which events that have the same or nearly the same meaning are semantically homogeneous, i.e. are assigned to the same TimeML class or classes. The initial analysis was restricted to 25 event synsets with a significant number of occurrences (5 or more) in both the English and the Italian corpora.

For each event token, we analyzed its *semantic pattern* or basic “argument structure” coupled with the semantic class (e.g., human) of each argument, as well as the thematic roles of the arguments (e.g., Agent, a Patient, a Theme, Experiencer. . .) together with subvalency features such as the fact that an argument is realized by an entire clause, a determined or undetermined NP, etc.⁴. Table 2 provides a sample of the results. We identified a total of 30 patterns among the common events. The results support the semantic homogeneity hypothesis in that in almost all cases (93.22%) there is a positive correlation between event meanings, semantic patterns, and TimeML class(es).

It is interesting to note that in five cases, we identified more than one semantic pattern occurring in both languages. Consider, for instance, ILI = 68138 in Table 2. The same event gives rise to two different semantic patterns: [*person/object: Underspec.*] E [*event: Theme*] and [*organization/object: Underspec.*] E [*entity: Manner*]. Cases such as these represent instances of event/verb subcategorization that offer insights into the effects that the semantic types of the arguments have both on event classification and event semantics. Moreover, this pattern of variation occurs in both languages, suggesting a parallelism between similarity in meaning and similarity of concepts.

Fewer than 10% of events with the same ILI and same semantic patterns differ in the assignment of the TimeML class. However, closer examination of these cases reveals that many of these discrepancies can be attributed to factors other than a real difference between event realization in the two languages. For example, consider the following:

1. ILI = 1490118, corresponding to WordNet *result1* and *derivare2*, *nascere9* in ItalWordNet. These events exhibit the same semantic pattern in all of the occurrences in the two corpora, i.e. [*event*] E [*event*], but differ in the assignment of the TimeML class: in particular, they are always annotated as OCCURRENCE in Italian, while in English they are assigned to either OCCURRENCE or ASPECTUAL;
2. ILI = 1432563, corresponding to WordNet *seek3* and ItalWordNet *cercare2*. Both have the same semantic pattern, i.e. [*person/organization*] E [*event*]; in English they are assigned to the LACTION class, while in Italian they are assigned to L.STATE;
3. ILI = 1381843, corresponding to WordNet *control1* and ItalWordNet *controllare3*. Both have the same semantic pattern, i.e. [*organization/person: Agent*] E

⁴Our analysis is, in fact, comparable to that carried out in Corpus Pattern Analysis (CPA).

[*organization: Patient*], but in English are assigned to STATE, while in Italian to OCCURRENCE.

4. ILI = 599498/519664, which correspond to WordNet *comment1* and ItalWordNet *osservare2*, *commentare2* and *rilevare3*. Both exhibit the same semantic pattern, i.e. [*person/organization*] E [*CLAUSE/event*], but in English are assigned to LACTION while in Italian to L.STATE.

The first case represents both an instance of inconsistency in annotation due to an incorrect sense interpretation by the TimeML annotators, and an incorrect semantic mapping between the English and the Italian senses. In particular, the verb *result1* means “terminate, end, ensue”, and the correct class assignment (according to the semantic pattern) should be ASPECTUAL. The assignment of the OCCURRENCE class results from an incorrect sense interpretation of *result1*, as “come about or follow as a consequence”, a sense which is absent in WordNet 2.0, but which has been included in WordNet 3.0. In addition, the ILI mapping between *result1* and *derivare2* is incorrect, since *derivare2* means “being caused by, follow as a consequence”. The correct mapping of *result1* in Italian should be *risolversi2*, which belongs to a different synset with respect to *derivare2*. Consequently, all the English events corresponding to *result1* should be assigned to the ASPECTUAL class, thus avoiding inconsistencies with respect to the semantic patterns. On the other hand, *derivare2*, which is a causal relation, is correctly assigned to OCCURRENCE in accordance to TimeML specifications. The *eq_synonym* link between *derivare2* and *result1* in ItalwordNet should therefore be corrected to reflect this. At first, cases like *seek3* - *cercare2* (example 2), *control1* - *controllare3* (example 3), and *comment1* - *osservare2*, *commentare2*, *rilevare3* (example 4) would seem to support the notion that similar events with the same semantic patterns can belong to different classes, thus contradicting the semantic homogeneity hypothesis. However, a closer analysis of these three cases shows that other issues are involved, namely, the exploitation of the mapping between TimeML classes and SIMPLE semantic types in the event annotation for Italian, and incorrect sense interpretation. The TimeML - SIMPLE mapping and heuristics work in this way: once the event is assigned to its semantic type, either it is classified directly or its classification is subordinated to the realization of a certain pattern.

In 2, the discrepancy between the TimeML class in English and in Italian is due to the fact that *cercare2* is wrongly assigned to the semantic type of *Modal Event* in SIMPLE instead of *Purpose Act*. In this case, it is a mistake in SIMPLE which leads to inconsistency of the data. On the other hand cases, like that in 4 suggest that the heuristics need to be refined to take into consideration not only the semantic type of the event, but also its lexical aspect or *aktionsaart*, i.e. whether it is a process, state, or transition. In fact, *osservare2*, *commentare2*, and *rilevare3* are ambiguous between state and process when taken out of context. According to the SIMPLE mapping heuristics they should be assigned to the L.STATE class. However, the L.STATE class is explicitly reserved for events that

ILI	Semantic Pattern	ITB Occur.	TB Occur.	TimeML Class
1517254-438385	[person/organization: Agent] E [event: Theme]	6	8	ASPECTUAL
1517254-438385	[object: Agent] E [event: Theme]	2	5	ASPECTUAL
68138	[person/object: Underspec.] E [event: Theme]	3	2	ASPECTUAL
68138	[organization/object: Underspec.] E [entity: Manner]	3	2	STATE
1546054	[person/organization: Experiencer] E [location]	5	8	STATE
97597-91455	[object: Patient] E [entity: Amount]	7	6	OCCURRENCE
580469	[person/organization: Agent] E [CLAUSE: Theme]	3	7	REPORTING

Table 2: Event synsets, semantic patterns and TimeML class.

denote a stative situations. So, the occurrences in our corpus to be classified as LSTATE must be states. On the contrary, when computing the *aktionsaart* values, they are all processes making it impossible to assign the class LSTATE⁵, in favor of LACTION, as in English.

Incorrect sense assignment to an event may also lead to an incorrect TimeML class assignment, as is the case for *controll* and *controllare*³. Both events have the same meaning, i.e. “exercise authoritative control or power over”, for which the set of participants, as the data show in 3, includes an *Agent* and a *Patient*. If a verb has an *Agent* as its thematic roles, cannot be considered a state. The event sense and the presence of this kind of “actors” claim for a wrong assignment in English. The event is not to be classified as STATE but as an OCCURRENCE, resolving the contradiction.

5.1.2. Nouns

We extracted only 11 common event types for nouns, all of them being nominalizations of a corresponding event verb in the subset. Our goal was, again, to test the semantic homogeneity hypothesis, but we were also interested in looking for cases of complex or dotted types (Pustejovsky, 1995). We are particularly interested in nominalizations, which normally represent reifications of the corresponding predicates (Gaeta, 2002).

A first interesting though trivial observation emerging from the data is the fact that nominal events in both Italian and English tend to belong to two classes: OCCURRENCE and STATE (80% of all instances). The classes LACTION and LSTATE are the only others assigned to nominal events in our data, although this may be at least partly attributable to data sparseness. LACTION and LSTATE are triggered when the event nominal either “governs” another event in cases of complex NP (i.e. NP + PP, where the NP of the PP is an event) or a purpose clause, and they are strictly linked to the semantic properties (*aktionsaart*) of the corresponding predicate, for example, “[...] i suoi *tentativi*[LACTION] di *distorcere*[OCCURRENCE] la realtà” “[...] his *attempts* to *distort* the truth.”]

We found that the assignment of wordnet senses was useful for identifying incorrect or inconsistent annotations of event nominals. In the Italian data, incorrect sense annotation of event nominals revealed missing semantic types in

SIMPLE. For instance, the nominal “aumento [increase]” exists in SIMPLE only in its eventive reading, which is just one of the possible senses in ItalWordNet. This led annotators to mark it as an event even when it has the meaning of quantity or amount. Similar problems occur in English. However, in the English data, we noticed an over-extension of the notion “nominalization = event”. Consider, for instance:

- *increase_n*: 3 of 11 occurrences in the subcorpus are marked as events, but the meaning is “the amount by which something increases”;
- *payment_n*: 8 of 10 occurrences are incorrectly marked as events, but the meaning is “a sum of money paid”.

In some cases, the identification of the wordnet sense is not enough to determine whether a nominal has an eventive reading or not. For example, in both WordNet and ItalWordNet there is not a clear-cut eventive sense of “*agreement_n*” - “*accordo_n, intesa_n*”, such as “the act of agreeing”⁶. In TimeBank, 31 of 32 occurrences of *agreement_n* are marked as events. All of these occurrences were assigned sense “agreement1”, i.e. “the statement - oral or written - of an exchange of promises”. In ITB, only seven of the sixteen occurrences *accordo_n* and *intesa_n* are marked as events, and all seven are assigned the same sense, i.e. “*accordo3, intesa3*” which is the *eq_synonym* of *agreement1*. It is interesting to note the differences in the number of occurrences marked as events (almost 100% in English; 43% in Italian), and that the occurrences of “*accordo_n, intesa_n*” that were not marked as events in ITB have the same sense (“*accordo3, intesa3*”) as the occurrences annotated as events. This shows that in Italian, “*accordo3, intesa3*” cannot be systematically interpreted as events, but instead represent instances of a dotted type, [EVENT ⊗ DOCUMENT], whose senses are not differentiated in WordNet and ItalWordNet, but may be activated by the co-text. For example, most of the unmarked occurrences in Italian co-occur with verbs like *firmare* [sign], *stipulare* [draw up], *siglare* [initial], which activate the DOCUMENT reading. “Incremental nominals” such as *decline, increase...*, where the eventive sense is blocked by the specification in the text of the amount or quantity, exhibits the same behavior as *agreement1*.

⁵Further data are required to completely validate this explanation.

⁶This sense is present in the ILI, but is restricted to oral agreement.

We believe that the over-extension of the eventive reading to almost all occurrences of nominalization in the English data is due to biases on the part of the annotators. In English, nominals are usually obtained via conversion from the verb, without any change in the aspect of the word itself. Also, cases such as *agreement1*, are clearly ambiguous, and a double reading (eventive *vs.* non-eventive) is always possible. This suggests that annotation schemes for events may need to be augmented to provide explicit means to mark ambiguous cases.

5.2. Automatically retrieved events: the ILI as a bootstrapping device

The second set of experiments focused on the use of the ILI and wordnet senses as a bootstrapping strategy for the creation of comparable corpora. The key idea is that words in two (or more) languages that have the same meaning share the same semantic properties, thus enabling the import of different layers of annotation from one corpus to another. To verify the validity of this claim, we used the entire corpus of the Italian TreeBank. A total of 62,522 words were manually tagged with a sense from ItalWordNet, including 9,832 verbs and 44,957 nouns. To carry out this experiment we developed a system which takes as input the events augmented with wordnet senses from TimeBank, and gives as output an additional layer of annotation, i.e. it creates the EVENT tag in the Italian TreeBank. The system workflow is the following: for each event in TimeBank, the system searches its database for its wordnet sense. If found, the system checks for the corresponding ILI, *eq_synonym* and *eq_near_synonym* relations in two other different databases according to the POS of the TimeBank event (i.e. noun or verb). The system searches the Italian TreeBank for words labeled with the corresponding sense, keeping track of the POS, and adding the event tag in the TreeBank if a match is found.

5.2.1. Evaluation

In this experiment we explored the use of the ILI as a bootstrapping device. The procedure described in the previous section identified 3,700 events (6.7.% of all sense annotated words) in Italian TreeBank, 1,183 of which are considered as “probable” events.⁷ In addition, 58 new event types not previously identified in the ITB were retrieved.

The results of this experiment confirmed the over-extension hypothesis by revealing several examples of incorrectly marked nominalizations, including *movement4*, *account1*, *earnings1*. In addition, we saw that sense assignment does not necessarily disambiguate the eventive and non-eventive readings of nominals. The identification of such cases can, in turn, provide material for studying the behavior of dotted types and pragmatic coercive effects. For example, neither the English “*indication1*” nor the Italian “*segnale1*”, whose synsets are linked through the ILI, denote an event. However, as the automatically retrieved occurrences in the Treebank of *segnale1* show, it can acquire an eventive reading,

⁷We say “probable” because, as illustrated in the previous sections, it is not always the case that event words with same sense also denote the same entity.

in particular when *segnale1* is followed by a clause or an event noun, such as “processo” [process] or “ripresa” [renewal].

Finally, a remark on the event extent is needed. Due to the way senses are annotated, it can happen that we have only a partial match between the event extent in TimeBank and the corresponding event-word in the Italian corpus. We see such cases when stative events are realized by prepositional phrases, for example, “*in declino*” [in decline], where the imported layer of annotation partially match the real event extent, assigning the tag event only to the noun head; e.g.: “*declino*” [decline].

The use of the ILI to automatically retrieve events in other languages shows promise as a means to lessen the cost of creating multilingual resources annotated for events. Although the process requires human intervention to validate the results, manual effort is significantly reduced because the set of “probable events” requiring validation is restricted to those words whose event reading is not explicitly stated in WordNet.

6. Conclusion

In this paper we have illustrated a methodology for linking comparable corpora in different languages by means of wordnet senses and the ILI. The resulting resource can be used for contrastive analysis of events as well as multilingual temporal analysis of texts. The data, in particular event verbs, provide support to the homogeneity hypothesis, i.e. that there exists a semantic homogeneity between similar events in different languages, in terms of their meaning, semantic patterns - including semantic preferences for thematic roles - and TimeML classes. In addition, the data suggest that there is a high degree of consistency among annotators in both Italian and English. We have shown that cases of non-correspondence between semantic pattern and TimeML class(es) are not due to differences in event realizations, but rather to incorrect sense interpretation and missing or wrong information in other resources.

Sense assignment to events has been shown to improve the accuracy of event annotation in the source and target corpora, by revealing inconsistencies and problems in the mapping between the Italian and the English wordnet sense, and facilitating event identification. In addition, it has been noticed that the presence of wordnet senses in the texts enables annotators to more easily and accurately identify those instances that satisfy the criteria for an event as proposed in TimeML, since they have the WordNet glosses for reference. The analysis of the data also suggests modifications to TimeML, in particular, the need for a tag to mark instances of ambiguous cases, so that the event tag is restricted to clear instances of events. Finally, we have shown how the ILI may be used as a semi-automatic bootstrapping device for creating comparable resources by importing layers of annotation for words that have the same sense.

7. Acknowledgements

Our thanks to Drew Nobile and James Schlessinger for their sense annotation of the English corpus.

8. References

- T. Caselli, I. Prodanof, N. Ruimy, and N. Calzolari. 2007. Mapping simple and timeml: improving event identification and classification using a semantic lexicon. In P. Bouillon, L. Danlos, and K. Kanzaki, editors, *Proceedings of the 4th International Conference on Generative Approaches to the Lexicon (G.L.2007)*.
- E. Filatova and E. Hovy. 2001. Assigning time-stamps to event-clauses. In *Proceedings of the ACL-EACL 2001 Workshop for Temporal and Spatial Information Processing*.
- L. Gaeta. 2002. *Quando i nomi compaiono come verbi. Un saggio di morfologia naturale*. FrancoAngeli.
- G. Katz and F. Arosio. 2001. The annotation of temporal information in natural language sentences. In *Proceedings of the ACL-EACL 2001 Workshop for Temporal and Spatial Information Processing*.
- R. Marinelli, L. Biagini, R. Bindi, S. Goggi, M. Monachini, P. Orsolini, E. Picchi, S. Rossi, N. Calzolari, and Zampolli A. 2003. Italian parole corpus. In A. Zampolli, N. Calzolari, and L. Cignoni, editors, *Linguistica Computazionale, Computational Linguistics in Pisa, special Issue*, volume XVIII-XIX, pages 401–421. IEPI, Pisa-Roma.
- S. Montemagni, F. Barsotti, M. Battista, N. Calzolari, O. Corazzar, A. Lenci, V. Pirelli, A. Zampolli, F. Fanciulli, M. Massetani, R. Raffaelli, R. Basili, M. T. Paziienza, D. Saracino, F. Zanzotto, N. Mana, F. Pianesi, and R. Delmonte. 2003. The syntactic-semantic treebank of italian. an overview. In A. Zampolli, N. Calzolari, and L. Cignoni, editors, *Linguistica Computazionale, Computational Linguistics in Pisa, special Issue*, volume XVIII-XIX, pages 461–493. IEPI, Pisa-Roma.
- J. Pustejovsky, J. Castao, R. Ingria, R. Sauri, R. Gaizauskas, A. Setzer, and G. Katz. 2003. Timeml: Robust specification of event and temporal expressions in text. In *Fifth International Workshop on Computational Semantics (IWCS-5)*.
- J. Pustejovsky. 1995. *The Generative Lexicon*. MIT Press.
- N. Ruimy, M. Monachini, E. Gola, N. Calzolari, M.C. Del Fiorentino, M. Ulivieri, and S. Rossi. 2003. A computational semantic lexicon of italian: *SIMPLE*. In A. Zampolli, N. Calzolari, and L. Cignoni, editors, *Linguistica Computazionale, Computational Linguistics in Pisa, special Issue*, volume XVIII-XIX, pages 821–864. IEPI, Pisa-Roma.
- R. Sauri, J. Littman, B. Knippen, R. Gaizauskas, A. Setzer, and J. Pustejovsky, 2006. *TimeML Annotation Guidelines version 1.2.1*. <http://www.timeml.org/site/publications/specs.html>.
- F. Schilder and C. Habel. 2001. From temporal expressions to temporal information: Semantic tagging of news messages. In *Proceedings of the ACL-EACL 2001 Workshop for Temporal and Spatial Information Processing*.