

# Event Annotation Schemes and Event Recognition in Spanish Texts

Dina Wonsever<sup>1</sup>, Aiala Rosá<sup>1</sup>, Marisa Malcuori<sup>2</sup>, Guillermo Moncecchi<sup>1</sup>,  
and Alan Descoins<sup>1</sup>

<sup>1</sup> Instituto de Computación

Facultad de Ingeniería

Universidad de la República, Uruguay

{wonsever, aialar, gmonce}@fing.edu.uy, dekked@gmail.com

<sup>2</sup> Instituto de Lingüística, Uruguay

Facultad de Humanidades y Ciencias de la Educación

Universidad de la República

marisamalcuori@gmail.com

**Abstract.** This paper presents an annotation scheme for events in Spanish texts, based on TimeML for English. This scheme is contrasted with different proposals, all of them based on TimeML, for various Romance languages: Italian, French and Spanish. Two manually annotated corpora for Spanish, under the proposed scheme, are now available. While manual annotation is far from trivial, we obtained a very good event identification agreement (93% of events were identically identified by both annotators). Part of the annotated text was used as a training corpus for the automatic recognition of events. In the experiments conducted so far (SVM and CRF) our best results are in the state of the art for this task (80.3% of F-measure).

## 1 Introduction

The fact of processing texts, no matter the purpose of such task, involves dealing with certain properties of the discourse that need to be grasped. We have chosen to adopt a modular structure to account for these properties, expressing them by means of the analysis of different independent axes, nevertheless able to interact with each other. Even though this structure does not provide, in principle, a holistic view of the discourse, it does allow to work independently in each axis, while it enables others to be added, as they develop.

The proposed analysis axes are: Enunciation, Events-Factivity, Temporality, Rhetorical Structure. Two more axes of structural nature are added to these four: Syntax axis and Textual Structure (paragraph, section, title, etc.) axis. The analysis for each one of the first four modules or axes is expressed in an annotation scheme for corpus annotation. Machine Learning techniques are applied upon these annotated corpora in order to generate a discourse analyzer. In this work we present the results of a set of tasks performed within the *Events-Factivity* module. We propose an event

annotation scheme based on TimeML (called SIBILA), we contrast this scheme with other proposals for Romance languages and we report the results obtained in the automatic recognition of events.

## 2 Event Annotation on Texts

### 2.1 Definition of Event

A core aspect in the computational understanding of a text is the detection of event references, as they constitute the minimal units with propositional content. Events can be actions (carried out voluntarily by an agent), processes (events spontaneously set off or caused by a force external to the process, which can, in both cases, be punctual or have duration), or states (situations maintained along a period or that are permanent). Generic predications will also be considered as events for they refer to states of things, states about which it is asserted that they take place.

Even though the events are in general indicated by verb forms, there also exist nouns that designate events. These event nouns do not designate objects (whether physical or abstract) but occurrences or incidents as in the case of *accidente* [accident], *batalla* [battle], *cena* [dinner], *eclipse* [eclipse], *desfile* [parade], *muerte* [death], *nacimiento* [birth], *tempestad* [storm], among many others.

While the verb category, whether in a finite form or not, is a powerful indication for detecting events, clear morphosyntactic indicators are missing for nominal events. Also, under the same form it is possible to interpret a noun as denoting an event or an object: *El concierto empieza a las ocho. / El concierto en si menor para violonchelo* [The concert starts at eight. / Cello concerto in B minor]; *Durante la construcción se presentaron varios problemas. / La construcción data del siglo XIX* [Several problems arose during the construction. / The construction dates from the 19<sup>th</sup> century]. This ambiguity constitutes a difficulty for automatic recognition. Nevertheless, there exists a series of syntactic indications that help to recognize this kind of nouns: co-occurrence with verbs such as *tener lugar* [to take place] o *presenciar* [to witness]; with verbs or expressions indicating duration or aspectual phase such as *empezar* [to start], *comenzar* [to begin], *concluir* [to finish], *terminar* [to end], *durar* [to last], as it is shown in (1):

- (1) *Esto sucedía después de que se mirara con buenos ojos el fin del corte en Guleguaychú llevado a cabo sobre las 14 horas de la tarde de ayer. [This was happening after the end of the roadblock in Guleguaychú carried out around 14.00 hours yesterday afternoon was well regarded.]*

Besides, events can be expressed by means of other categories such as adjective, prepositional phrase, given that states can be designated by means of them, and also by the pronoun category when the referent is an event.

## 2.2 Annotation Scheme

The annotation scheme SIBILA, which dates from 2008, is an adaptation to Spanish of the TimeML scheme [12, 17]. Beyond the fact that adaptation is not a trivial task, the SIBILA scheme incorporates some innovative elements, the most important of which is the *factivity* attribute and its values. Starting from the SIBILA scheme a detailed annotation guide with lots of examples was made in order to guide annotators [24] and, likewise, reasons for the study of factivity and its relevant values [25] were established.

There currently exist other adaptations of TimeML for Romance languages such as Italian [7] and French [3], and there is also a Spanish version proposed by the TimeML team [20].

The adaptation for Spanish by means of the SIBILA scheme shares some attributes incorporated in the schemes above mentioned and it also includes, besides the *factivity* attribute, other changes about which we are going to briefly speak about next. Anyway, the SIBILA scheme is consistent with the proposal of TimeML.

Even though the scheme establishes, in addition to events, the annotation of other elements such as different kinds of indexes, aspectual and subordination links between events, temporal expressions and temporal links, in this occasion we will only refer to the events.

### Events and Their Attributes

A complete description of the *event* element is presented next, followed by the analysis of differences and similarities with regard to the rest of the schemes based on TimeML. Table 1 shows the event attributes and their values.

**Table 1.** Event attributes

Attribute	Value
<i>id</i>	unique identifier
<i>class</i>	OCCURRENCE   PERCEPTION   REPORT   ASPECT   STATE   INTENSIONAL_CAUSAL_ACTION   INTENSIONAL_STATE   EXISTENCE
<i>category</i>	VERB   NOUN   ADJECTIVE   PRONOMINAL   OTHER
<i>verb_form</i>	INFINITIVE   GERUND   PARTICIPLE   FINITE_FORM
<i>mood</i>	INDICATIVE   SUBJUNCTIVE   CONDITIONAL   IMPERATIVE

**Table 1.** (continued)

<i>time</i>	PAST   PRESENT   FUTURE
<i>determination</i>	DEFINITE   INDEFINITE   BARE
<i>modality</i>	Lexical item of a modality operator (free text)
<i>polarity</i>	NEG   POS
<i>factivity</i>	YES   NO   PROGRAMMED_FUTURE   NEGATED_FUTURE   POSSIBLE   INDEFINITE
<i>indexes</i>	references to indexes (ids)
<i>lex_item</i>	free text (CDATA)
<i>comments</i>	free text (CDATA)

As in the schemes proposed for French and Italian and in the Spanish version of TimeML, in SIBILA, *mode* and *verb form* attributes are incorporated in order to account for the flexive complexity of Romance languages. However, a significant difference shown by SIBILA relates to the value of the *time* attribute. Beyond the tense value assigned to finite forms by the tagger, the *time* attribute will take the value of PAST, PRESENT or FUTURE accordingly with the meaning that the verb form may have in the text in which it appears. So, it represents the semantic temporal value and not the syntactic tense value. For instance, a verb form like *descubre* [discovers] in *Colón descubre América en 1492* [Colon discovers America in 1492] will have for the time attribute the PAST value, even if it is a present verb form.

On the other hand, SIBILA incorporates the EXISTENCE value for the *class* attribute. In this way, it treats the copulative, existential and presentative verbs as events that operate predicating others' event existence. That is to say, when an event referred by a noun, an adjective or a prepositional phrase is part of a predicate with copulative verb or when an existential or presentative verb takes an argument that refers to an event, the copulative, existential, presentative or other verb elements that may act as such in the text will take the EXISTENCE value. In (2) and (3) we show in bold the events with existence value and underlined the subordinated events.

(2) *La estatal brasileña también **está interesada** en estaciones de servicio y otros activos de Esso en el resto del Cono Sur Americano, dijo durante un encuentro con periodistas en Río. [The Brazilian state-owned company is also interested in gas stations and other assets of Esso in the rest of South America, he/she said during a press conference in Rio.]*

(3) *Tal fue el caso de este lunes, en que **se registraron** durante 20 minutos fuertes nevadas en Colonia, según informó Canal 10. [That was the case of this Monday, when strong snowfalls during 20 minutes were recorded in Colonia, as reported by Channel 10.]*

It can also occur, as it is shown in (4), that nominalizations of the OCCURRENCE class may behave in a way similar to the predicates mentioned and introduce an event under the form of complement. In this case, they will also take the EXISTENCE value.

(4) *Se descartó la **ocurrencia** de nevadas en Montevideo. Sí pueden producirse precipitaciones de "agua nieve". [The occurrence of snowfalls in Montevideo was ruled out. "Sleet" falls may certainly happen.]*

A partially similar change is proposed for French [3] with the introduction of the new class EVENT\_CONTAINER for events. Predicates that take a nominal event as subject (*De nombreuses manifestations se sont produites dans la tournée du dimanche.*) belong to this class.

The scheme for French also introduces for the *class* attribute the CAUSE value to account for verbs that indicate a causal relationship between two events (*causer, provoquer, engendrer*, etc.). A similar change had already been proposed in the SIBILA scheme: the extension of the INTENSIONAL ACTION class under the name of CAUSAL INTENSIONAL ACTION, in order to give place, precisely, to this kind of verbs.

In the TimeML annotation guide[17] the description of the events is presented by means of two differentiated elements: *event* and *makeinstance*<sup>1</sup>, the second of which is an empty element. This information was unified in SIBILA in order to simplify the annotation task, which implied the creation of 2 elements by each registered event. An alternative solution was then proposed, the *lexical item* attribute for the case of elided events that TimeML resolved by means of the creation of another instance with the same reference. The *lexical item* attribute is optional and is used to register an event in the cases of ellipsis, that is to say, to register the instance of an event the mention of which is omitted, because the predicate that names it may be recovered by resorting to another mention in the text. The remaining attributes of the elided event (empty event) collect additional information associated to it, as it is shown in (5) and (6).

(5) *En el norte del país llovió abundantemente el sábado y <event lex\_item = "llovió"/> el domingo. [It rained heavily on Saturday and Monday in the north of the country.]*

(6) *El corte de ruta comenzó el día 14 y <event lex\_item = "corte"/> terminó una semana después. [The roadblock began on the 14<sup>th</sup> and ended a week later.]*

---

<sup>1</sup> Reference to the *makeinstance* element has disappeared in the last versions of TimeML [21].

### The *Factivity* Attribute

The *factivity* attribute represents the degree of certainty of the utterer with regard to the occurrence of the referred event. It follows then that any affirmation about the occurrence or not of an event remains circumscribed to an enunciation context.

(7) *Esto **dificulta** aún más el **diálogo** con el gobierno uruguayo quien **confirmó** ayer a través de la cancillería que no se **negociará** mientras permanezca algún **corte**. [This makes the dialogue with the Uruguayan government even more difficult; the Uruguayan government confirmed through the Ministry of Foreign Affairs that they will not negotiate while a roadblock remains in place.]*

In (7) the events<sup>2</sup> are in bold and an aspectual operator (*permanezca* [*remains*]) is underlined. Note that while some events are presented as occurred (*confirmó* [*confirmed*], *dificulta* [*makes difficult*], *corte* [*roadblock*]) others are uncertain (*diálogo* [*dialogue*]) and the eventual negotiation (*negociará* [*will negotiate*]) is presented as future and with negative polarity. This means that the occurrence of some event referring word is not enough to infer that such event has occurred or is occurring. It is also necessary to interpret these terms in their contexts of occurrence, where they can be affected by elements of negative polarity, or by modal operators, or by predicates that affect their veracity value, and by combinations of all of them. The property of an event of having occurred or not or of being occurring is not then an evident piece of information. In fact, it is necessary to make some kind of textual inference in order to determine it.

Annotators must, precisely, make those inferences and annotate the event by attributing to it one of the following values:

YES – performed event

NO – non performed event

PROGRAMMED\_FUTURE – event with high probability of taking place

NEGATED\_FUTURE – highly improbable event

POSSIBLE – event that might take place

INDEFINITE<sup>3</sup> – event about which it is not known whether it has taken place or not

An example for each of these values is offered next:

(8) a. *La ministra Daisy Tourné **anunció** que algunos reclusos del Compen **serán trasladados** al interior del país, para **aliviar** la superpoblación de ese centro carcelario. No se **conocen** más novedades. [Minister Daisy Tourné announced*

<sup>2</sup> The expression of the event usually contains more than a word. A term considered to be the nucleus of the event is annotated (and is shown highlighted).

<sup>3</sup> Note that the difference between the values POSSIBLE and UNDEFINITE is that the first one indicates that an event may take place in the future, while the second says that there are no elements to determine if an event took or did not take place: *Están valorando iniciar un paro* [POSSIBLE]. [They are evaluating to begin a strike] / *En ese momento valoraron iniciar un paro*[UNDEFINITE] [They evaluated at that time to initiate a strike].

*that some Compen prisoners will be transferred to the provinces, in order to relieve that prison's overpopulation. No further news are known.]*

*anunció* = YES

*serán trasladados* = PROGRAMMED\_FUTURE

*aliviar* = POSSIBLE

*se conocen* = NO

b. *La idea de la exposición "Shoá. Memoria y legado del Holocausto" surgió de tres jóvenes judíos que querían **transmitir** el legado recibido de los supervivientes del exterminio. [The idea of the "Shoá. Memoria y legado del Holocausto" exhibition came from three Jewish young people who wanted to transmit the legacy received from the Holocaust survivors.]*

*transmitir* = INDEFINITE

c. *El gobierno uruguayo confirmó ayer a través de la cancillería que no **se negociará** mientras permanezca algún corte. [The Uruguayan government confirmed yesterday through the Ministry of Foreign Affairs that they will not negotiate while a roadblock remains in place.]*

*se negociará* = NEGATED\_FUTURE

Although factivity is closely related to tense, modality and polarity, the association is not automatic. Thus, events with the same values for these attributes may exhibit different factivity values: *Celebro que **lleguen** [PROGRAMMED\_FUTURE] mañana [I am glad they are coming tomorrow] / Dudo que **lleguen** [POSSIBLE] mañana [I doubt they are coming tomorrow]; *Logró **cerrar** [YES] la puerta [He managed to close the door] / Olvidó **cerrar** [NO] la puerta [He forgot to close the door].**

In [19] there is a proposal to associate factivity to events, with a definition partially similar to ours. Besides, in this work it is developed a determinist algorithm for the calculus of factivity values, based on the fact that some relevant elements such as markers of polarity or modality, source introducing predicates and events selecting predicates, have been recognized and classified in a previous stage. But, to our knowledge, an attribute for factivity has not been previously included in an annotation scheme. We claim that this attribute will be useful for an effective recognition of this complex phenomenon.

### 3 The Annotated Corpora

#### 3.1 Description of the Corpora

The annotated corpus are constituted by journalistic and historical texts. Journalistic texts come from a corpus in Spanish created for the TempEval2<sup>4</sup> task, annotated on the basis of the TimeML scheme for Spanish. It was decided to annotate these texts in order to obtain a comparative parameter for Spanish.

---

<sup>4</sup> <http://www.timeml.org/tempeval2>

The corpus is formed by 11,986 tokens and 408 sentences. 1,677 events were annotated, most of them being verbs, nouns in second place, and lastly, a few of them being adjectives.

### 3.2 Agreement between Annotators

In order to evaluate the agreement between annotators we used the  $agr^5$  measure proposed in [18], defined as follows:

Let  $A$  and  $B$  be the portions of text marked as events by two annotators  $a$  and  $b$  respectively. The  $agr$  measure tells us which proportion of  $A$  was also marked by  $b$ . To be precise, agreement between  $b$  and  $a$  is computed as:

$$agr(a|b) = \frac{|A \text{ agreeing with } B|}{|A|}$$

The  $agr(a|b)$  measure corresponds to the *recall* if  $a$  is taken as *gold standard* and  $b$  as the labeling system, and to *precision* if  $b$  is the *gold standard* and  $a$  the system.

Agreement values between the annotators obtained are shown in table 2.

**Table 2.** Agreement between annotators

	Precision	Recall	F-measure
Global	91.6 %	93.0 %	92.3 %
Verbal Events	94.2 %	97.1 %	95.6 %
Nominal Events	85.8 %	88.7 %	87.3 %

We can see that the values are significantly lower for nouns than for verbs, as it was to be expected. Agreement values for other categories were not calculated for they are much less frequent in the corpus and, therefore, results would not be representative.

## 4 Machine Learning on the Corpus

### 4.1 Models for Learning

As a first experience of exploitation of the annotated corpus, we have developed a system that uses machine learning techniques for event recognition. The system only determines the text segments corresponding to events, a task that, for the particular case of nouns, is far from trivial. Recognition of segments referring to events was focused as a problem of sequential classification, using the usual system of labels B,I,O.

---

<sup>5</sup> The widely used Kappa measure was discarded as it suffers from various problems for sequential relatively scattered data [12].



We have used two learning methods radically different to generate classifiers: *Conditional Random Fields* (CRF) and an adaptation of *Support Vector Machine* (SVM) for problems of sequential classification.

CRF [9] is a discriminative model of sequential classification which, given a sequence  $x$  of observations, tries to obtain the sequence  $y$  of output labels that maximizes probability  $P(y|x)$ . This model has certain advantages over other models (of generative type, such as the *Hidden Markov Models*, HMM), for they do not need to calculate probability  $P(x)$  of the input sequence [9].

The SVM [22] model is not in principle a sequential classification method, although it can be adapted for that task. In the non-sequential case, SVM considers instances to be classified as points in a space with a certain dimension (possibly finite) and builds a lineal separator that partitions the space and divides the instances according to their class. In this way, the new instances will obtain their class according to the side of the hyperplane in which they are. Two modifications are necessary in order to apply this model to the sequential classification task. The first one is to be able to classify in more than two classes (SVM is a binary classification method), for which classifiers for each pair of classes are built, making then a pondered voting to determine the class to be assigned. The second one is to incorporate the rest of the elements of the sequence, in addition to the one that is being classified, to the classification. This is made by means of a technique called *forward parsing*, that uses labels assigned so far as attributes for subsequent classifications (proceeding from left to right in the sequence). For more details, consult [8].

70% of the total annotated corpus was used as training corpus in order to train classifiers. The remaining 30% was divided as follows: 15% as development corpus and 15% as testing corpus.

We used the CRFSuite<sup>6</sup> tool in order to train the classifier based on CRF and Yamcha<sup>7</sup> (a sequential classification tool) for the SVM classifier.

In both cases we used morphosyntactic attributes, some of which coming from the *Freeling* [1] tagger (token, lemma, POS-tag, number, mood and tense), and others associated with word structure (capital letters, last four letters). We considered a [-2,2] window centered on the token we wanted to classify.

## 4.2 Results

Results can be observed in table 3. The base line shown there was obtained by marking as an event every contiguous sequence of verbs and the nouns with the most frequent endings (4 final letters) among the nominal events of the training corpus. Results of agreement between annotators were used as top line.

---

<sup>6</sup> <http://www.chokkan.org/software/crfsuite>

<sup>7</sup> <http://chasen.org/~taku/software/yamcha>

**Table 3.** Classifiers' results (%) on the testing corpus

	<i>Precision</i>				<i>Recall</i>				F-measure			
	Base	Top	CRF	SVM	Base	Top	CRF	SVM	Base	Top	CRF	SVM
Global	67.1	91.6	81.7	<b>84.7</b>	57.3	93	72.4	<b>76.4</b>	61.8	92.3	76.7	<b>80.3</b>
Verbal Events	65.2	94.2	83.2	<b>84.2</b>	79.3	97.1	91.9	<b>98.5</b>	71.6	95.6	87.3	<b>90.8</b>
Nominal Events	63.3	85.8	71.8	<b>78.9</b>	27.9	82.7	41.2	<b>44.1</b>	38.8	87.3	52.3	<b>56.6</b>

As it can be seen in the table, the base line of 61,8% of F-Measure is broadly surpassed by both methods. Contrary to what might be expected, given the fact that CRF is the state of the art in several problems of sequential classification, the SVM model gives higher values than the CRF model in all cases. On the other hand, both classifiers are far from reaching the top line, for which the F-Measure is 92,3%.

The most frequent mistakes made by both classifiers are related to nominal events. In order to improve this result, strategies similar to those used in [15] will be tried for the detection of non-deverbal event nouns. With regard to the precision value of verb events, we think that it is affected by the inclusion in this class of participle forms that many times do not constitute events.

### 4.3 Comparison with Other Works

With regard to automatic recognition, the obtained results are very encouraging, being of the same order that the results produced by similar works applied to English (see table 4). As it is shown by the table, only one system reaches a F-Measure higher than the ours. This work [10] includes among the input attributes information about thematic roles. For the time being, it is not possible to have this kind of information for Spanish for there does not exist an automatic tagger for thematic roles.

An important difference between the works mentioned and ours is the size of the corpus used for learning. In our case, the training corpus contains about 8,500 tokens, while the rest of the systems, all of them based on TimeBank, have a corpus 7 times larger. Even though it is generally accepted that it is necessary to have a larger corpus, differences between sizes of corpora used, on the one hand, and similarity of the results obtained, on the other hand, suggest that it is not the size of the corpus the element that has more bearing on the results.

**Table 4.** Comparison with other systems

System	F-Measure
Our system	76.7% (CRF) / 80.3% (SVM)
Evita [14]	80.1%
Sim-Evita [2]	73.0%
Boguraev, Ando [4]	80.3%
Step [2]	75.9%
March, Baldwin [9]	76.4%
Llorens et al [8]	81.4%

## 5 Conclusions

The SIBILA annotation scheme was defined; it constitutes an adaptation of the TimeML event annotation scheme to Spanish with the addition of elements for event factivity annotation. The basic part of the scheme is maintained, but some changes that we think make SIBILA more suitable for this language are introduced. From a comparative study with works for other Romance languages it comes out that similar modifications were proposed independently. Modifications proposed by SIBILA do not imply a mismatch with TimeML, a SIBILA conversion to TimeML is completely feasible, with some loss of information. This is important because TimeML is becoming a standard in works in this field.

The SIBILA scheme was validated by the effective annotation of a first set of texts with more than 1,500 events. Event manual annotation is not an easy task, there exist several difficult cases for which it is still necessary to clarify the criteria to be followed by annotators. Anyway, the agreement measure between annotators is very good (92.3% of global F-measure), even for nouns, that constitute the most complex case (87.3% of F measure in event nouns).

As a first experience of exploitation of the annotated corpus, a system that uses machine learning techniques for event recognition was developed. The system only determines the text segments corresponding to events, a task that, for the particular case of nouns, is far from trivial. Two learning methods radically different were used to generate classifiers: *Conditional Random Fields* (CRF) and an adaptation of *Support Vector Machine* (SVM) for problems of sequential classification. Results obtained are encouraging, having obtained in the best case 80% of F-measure with SVM. This number improves a lot (90%) if we only consider the verb events; the best F-measure that we have obtained for nominal events is 56.6%.

A larger volume of text is being annotated; it will be used for conducting new experiments, as well as for carrying independent factivity learning experiments. Another future work will be the integration with the enunciation axis, based on t[16].

## References

1. Atserias, J., Casas, B., Comelles, E., González, M., Padró, L., Padró, M.: FreeLing 1.3: Syntactic and semantic services in an open-source NLP library. In: Proceedings of the Fifth International Conference on Language Resources and Evaluation (LREC) ELRA (2006)
2. Bethard, S., Martin, J.H.: Identification of event mentions and their semantic class. In: Proceedings of the 2006 Conference on Empirical Methods in Natural Language Processing, EMNLP 2006, pp. 146–154. Association for Computational Linguistics, Stroudsburg (2006)
3. Bittar, A., Amsili, P., Denis, P., Danlos, L.: French TimeBank: An ISO-TimeML Annotated Reference Corpus. In: Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: shortpapers, Portland, Oregon, pp. 130–134 (2011)
4. Boguraev, B., Kubota Ando, R.: TimeML-compliant text analysis for temporal reasoning. In: Proceedings of the 19th International Joint Conference on Artificial Intelligence, pp. 997–1003. Morgan Kaufmann Publishers Inc., San Francisco (2005)
5. Boguraev, B., Kubota Ando, R.: Effective Use of TimeBank for TimeML Analysis. In: Schilder, F., Katz, G., Pustejovsky, J. (eds.) *Annotating, Extracting and Reasoning about Time and Events*. LNCS (LNAI), vol. 4795, pp. 41–58. Springer, Heidelberg (2007)
6. Carletta, J.: Assessing agreement on classification tasks: The Kappa statistic. *Computational Linguistics* 22, 249–254 (1996)
7. Caselli, T., Bartalesi, V., Sprugnoli, R., Pianta, E., Prodanof, I.: Annotating Events, Temporal Expressions and Relations in Italian: the It-TimeML Experience for the Ita-TimeBank. In: Proceedings of the Fifth Law Workshop (LAW V), Portland, Oregon, pp. 143–151 (2011)
8. Kudo, T., Matsumoto, Y.: Chunking with Support Vector Machines. In: NAACL (2001)
9. Lafferty, J., McCallum, A., Pereira, F.: Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In: Proceedings of the Eighteenth International Conference on Machine Learning, ICML 2001, pp. 282–289. Morgan Kaufmann Publishers Inc., San Francisco (2001)
10. Llorens, H., Saquete, E., Navarro-Colorado, B.: TimeML events recognition and classification: learning CRF models with semantic roles. In: Proceedings of the 23rd International Conference on Computational Linguistics, COLING 2010, pp. 725–733. Association for Computational Linguistics, Stroudsburg (2010)
11. March, O., Baldwin, T.: Automatic event reference identification. In: Proceedings of the Australasian Language Technology Association Workshop 2008, páginas, Hobart, Australi, pp. 79–87 (2008)
12. Ben, M.: *Investigating Classification for Natural Language Processing Tasks*, Ph.D. Thesis, Cambridge University (2007)
13. Pustejovsky, J., Castaño, J., Ingria, R., Saurí, R., Gaizauskas, R., Setzer, A., Katz, G.: TimeML: Robust specification of event and temporal expressions in text. In: Fifth International Workshop on Computational Semantics, IWCS-5 (2003)

14. Pustejovsky, J., Hanks, P., Saurí, R., See, A., Gaizauskas, R., Setzer, A., Radev, D., Sundheim, B., Day, D., Ferro, L., Lazo, M.: The TIMEBANK Corpus. In: *Proceedings of Corpus Linguistics*, pp. 647–656 (2003)
15. Resnik, G., Bel, N.: Automatic Detection of Non-deverbal Event Nouns in Spanish. In: *Proceedings of the 5th International Conference on Generative Approaches to the Lexicon*, Istituto di Linguistica Computazionale, Pisa (2009)
16. Rosá, A., Wonsever, D., Minel, J.-L.: Comparación de dos métodos para la extracción de opiniones en textos en español. In: *Proceedings of IBERAMIA 2010, Workshop on Natural Language Processing and Web-based Technologies*, Bahía Blanca (2010)
17. Saurí, R., Knippen, R., Verhagen, M., Pustejovsky, J.: Evita: a robust event recog-nizer for QA systems. In: *Proceedings of the conference on Human Language Technology and Empirical Methods in Natural Language Processing, HLT 2005*, pp. 700–707. Association for Computational Linguistics, Stroudsburg (2005)
18. Saurí, R., Littman, J., Knippen, B., Gaizauskas, R., Setzer, A., Pustejovsky, J.: TimeML Annotation Guidelines Version 1.2.1 (2006)
19. Saurí, R.: A Factuality Profiler for Eventualities in Text. PhD Dissertation. Brandeis University (2008)
20. Saurí, R., Batiukova, O., Pustejovsky, J.: Annotating Events in Spanish TimeML Annotation Guidelines. Version TempEval-2010 (2009)
21. Saurí, R., Goldberg, L., Verhagen, M., Pustejovsky, J.: Annotating Events in English TimeML Annotation Guidelines. Version TempEval-2010 (2009)
22. Vapnik, V.: *The Nature of Statistical Learning Theory*. Springer, New York (1995)
23. Wiebe, J., Wilson, T., Cardie, C.: Annotating expressions of opinions and emotions in language. In: *Language Resources and Evaluation* (2005)
24. Wonsever, D., Malcuori, M., Rosá, A.: Sibila: esquema de anotación de eventos. Technical Report 08–11, Biblioteca InCo PEDECIBA (2008) ISSN: 0797–6410
25. Wonsever, D., Malcuori, M., Rosá, A.: Factividad de los eventos referidos en textos. Technical Report 09–12, Biblioteca InCo PEDECIBA (2009) ISSN: 0797–6410