

MEANTIME, the NewsReader Multilingual Event and Time Corpus

Anne-Lyse Minard*, Manuela Speranza*, Ruben Urizar[‡], Begoña Altuna[‡],
Marieke van Erp[◇], Anneleen Schoen[◇], Chantal van Son[◇]

*Fondazione Bruno Kessler, Trento, Italy

{minard,manspera}@fbk.eu

[‡] University of the Basque Country (UPV/EHU), Spain

{ruben.urizar,begona.altuna}@ehu.eus

[◇] Vrije Universiteit Amsterdam, the Netherlands

{marieke.van.erp,a.m.schoen,c.m.van.son}@vu.nl

Abstract

In this paper, we present the NewsReader MEANTIME corpus, a semantically annotated corpus of Wikinews articles. The corpus consists of 480 news articles, i.e. 120 English news articles and their translations in Spanish, Italian, and Dutch. MEANTIME contains annotations at different levels. The document-level annotation includes markables (e.g. entity mentions, event mentions, time expressions, and numerical expressions), relations between markables (modeling, for example, temporal information and semantic role labeling), and entity and event intra-document coreference. The corpus-level annotation includes entity and event cross-document coreference.

Semantic annotation on the English section was performed manually; for the annotation in Italian, Spanish, and (partially) Dutch, a procedure was devised to automatically project the annotations on the English texts onto the translated texts, based on the manual alignment of the annotated elements; this enabled us not only to speed up the annotation process but also provided cross-lingual coreference. The English section of the corpus was extended with timeline annotations for the SemEval 2015 *TimeLine* shared task. The *First CLIN Dutch Shared Task* at CLIN26 was based on the Dutch section, while the EVALITA 2016 *FactA* (Event Factuality Annotation) shared task, based on the Italian section, is currently being organized.

Keywords: benchmark for NLP technologies, parallel corpus, semantic annotation, cross-lingual coreference

1. Introduction

The NewsReader MEANTIME (Multilingual Event AND TIME) corpus is a semantically annotated corpus of 480 English, Italian, Spanish, and Dutch news articles. It was created within the NewsReader project,¹ whose goal is to build a multilingual system for reconstructing storylines across news articles to provide policy and decision makers with an overview of what happened, to whom, when, and where. Semantic annotations in the MEANTIME corpus span multiple levels, including entities, events, temporal information, semantic roles, and intra-document and cross-document event and entity coreference.

The English section of the corpus consists of articles taken from Wikinews,² and the Spanish, Italian, and Dutch sections are translations of the same texts. This ensures access to non-copyrighted articles for the evaluation of the NewsReader system and enables a fine-grained comparison of natural language processing tools across the languages.

Semantic annotation on the English texts was performed manually. For the annotation in Italian, Spanish, and (partially) Dutch, we devised a procedure to automatically project the annotations available in the English texts onto the translated texts, based on the manual alignment of the annotated elements, speeding up the annotation process considerably.

The remainder of this article is organized as follows. We review related work in Section 2. In Section 3 and 4 we present the different annotation levels and the annotation process. In Section 5 we provide a description of the cor-

pus. Finally, we conclude presenting some uses of the corpus.

2. Related work

A number of semantically annotated corpora have been made available to train and evaluate text processing systems, especially in relation to evaluation exercises on specific shared tasks, such as Named Entity Recognition and Entity Coreference (Linguistic Data Consortium,), Temporal Processing (UzZaman et al., 2013), and Semantic Role Labeling (Surdeanu et al., 2008). Additionally, several resources annotated at the corpus level exist, such as the ECB+ corpus for cross-document event coreference (Cybulska and Vossen, 2014) and the corpus used for cross-document entity coreference at the Web People Search task at SemEval 2007 (Artiles et al., 2007).

Among Italian corpora, a set of news articles from the local Italian newspaper “L’Adige” has been annotated with different levels of information: (i) entities and entity coreference in I-CAB³ (Magnini et al., 2006); (ii) events, time expressions and temporal relations in EVENTI⁴ (Caselli et al., 2014); (iii) event factuality in Fact-Ita Bank⁵ (Minard et al., 2014); and (iv) cross-document entity coreference for a specific subset of person entities in CRIPCO⁶ (Bentivogli et al., 2008). To the best of our knowledge there exist no Italian corpora with either semantic role labeling or event cross-document coreference annotation.

³<http://ontotext.fbk.eu/icab.html>

⁴<https://sites.google.com/site/eventievalita2014/data-tools>

⁵<http://hlt-nlp.fbk.eu/technologies/fact-ita-bank>

⁶<http://hlt-nlp.fbk.eu/technologies/cripco>

¹<http://www.newsreader-project.eu/>

²Wikinews is a collection of multilingual online news articles written collaboratively in a wiki-like manner (<http://en.wikinews.org>).

The most well-known Dutch annotated corpus is the dataset developed for the CoNLL 2002 Language-Independent Named Entity Recognition shared task.⁷ It consists of four editions of a Belgian newspaper that were annotated with Person, Location, Organization and Miscellaneous entities (Tjong Kim Sang, 2002). Between 2008 and 2011, several Dutch and Flemish universities collaborated to create the SoNaR corpus⁸ (Oostdijk et al., 2013), containing several layers of annotations (such as parts-of-speech, named entities and semantic roles), which were carried out in part manually, in part semi-automatically, and in part automatically.

A number of initiatives to build large annotated corpora for Spanish have been carried out, with CORPES XXI⁹ (Real Academia Española, 2015) being the most prominent. This corpus is annotated with morphosyntactic information, multiword expressions, and named entities. AnCora (Taulé et al., 2008) is a multilingual corpus of Catalan and Spanish newspaper texts annotated at different linguistic levels: morphological, syntactic, and semantic (covering argument structures, thematic roles, semantic verb classes, named entities, and WordNet nominal senses). AnCora was later enriched with coreference links between pronouns (including elliptical subjects and clitics), full noun phrases (including proper nouns), and discourse segments (Recasens and Martí, 2010). Since the corpora for Catalan and Spanish are not parallel, no cross-lingual annotation was carried out.

Among parallel multilingual corpora, we focus on translation corpora, which are often aligned at the sentence level (Mouka et al., 2012; Gavrilidou et al., ; Koehn, 2005). The MultiSemCor corpus (the SemCor English corpus translated into Italian), on the other hand, is aligned at the word level (Bentivogli and Pianta, 2005). Exploiting the alignment between English and Italian, the WordNet sense annotation available in SemCor was projected onto the Italian translation.

Our method to project annotation across texts in different languages taking advantage of text alignment is also similar to other methods used, for example, to build annotated corpora with semantic roles (Padó and Lapata, 2009), temporal information (Spreyer and Frank, 2008; Forascu and Tufi, 2012), and coreference chains (Postolache et al.,). However, previous work is based on the use of corpora aligned at the token level, whereas our method envisages an alignment at the markable level, where each annotated element is aligned to English on a semantic rather than syntactic basis.

3. Levels of semantic annotation

3.1. Annotation at document level

The intra-document semantic annotation is based on the NewsReader guidelines (Tonelli et al., 2014). It revolves around markables (e.g. entity mentions, event mentions, time expressions, and numerical expressions), relations between markables, and entity/event co-reference (a relation

called REFERS_TO is used to mark that one or more markables refer to the same entity or event).

Markables Entity mentions are the textual spans in which entity instances are referenced (e.g. “the US president” and “Obama” refer to the same entity instance, i.e. *Barack Obama*). Given that the focus of the project is on the economic and financial domain, we defined two new types of entities (namely Products and Financial entities) in addition to the classic Person, Organization, and Location entities. Product entities include anything that can be offered to the market to satisfy a want or need (e.g. *iPhone 4*), while Financial entities are entities belonging specifically to the financial domain (e.g. *Dow Jones*). In the annotation process, each entity mention is described through a text span and two optional attributes, namely its syntactic head and syntactic type.

As Italian and Spanish, unlike English, are null-subject languages where clauses lacking an explicit subject are permitted, we devised specific guidelines. Null subjects having finite verb forms as predicates and referring to existing entity instances were marked through the creation of an empty (i.e. non text-consuming) tag, which was then linked to other markables following the guidelines for regular text consuming entity mentions. For instance, in the following Italian sentences “Obama fece un discorso. [Ø] Disse che [...]” (‘Obama gave a speech. [He] said that [...]’) the null subject of “disse” (‘said’) is annotated using a non text-consuming tag and linked to the instance *Obama*.

Event mentions, i.e. textual realizations of event instances, can be verbs, nouns, pronouns, adjectives, or prepositional constructions (for example, “this conference”, “LREC2016”, and “it”, all refer to the same event instance, i.e. the *10th edition of the Language Resources and Evaluation Conference*). Each event mention is annotated through its text span and a number of attributes, such as predicate (lemma) and part-of-speech, and is associated with a factuality value described through several attributes (van Son et al., 2014), including time, certainty and polarity.

The annotation of temporal expressions, based on the ISO-TimeML guidelines (ISO TimeML Working Group, 2008), includes durations (e.g. “three months”), dates (e.g. “March 10th, 2016” or the document creation date), times (e.g. “5.30 PM”), and sets of times (e.g. “every week”) and the following attributes: type, normalized value, anchorTimeID (for anchored temporal expressions), and beginPoint and endPoint (for durations).

Given their relevance in the economic and financial domain, numerical expressions are also annotated; they include percentages (e.g. “70%”), amounts described in terms of currencies (e.g. “10,000 Euros”), and general amounts (e.g. “more than 90”).

Temporal and causal relations can be made explicit by textual elements (e.g. “while” for temporal relations or “because of” for causal relations). In MEANTIME these temporal signals (SIGNAL), inherited from ISO-TimeML, and causal signals (C-SIGNAL) have been annotated.

Relations Temporal relations (such as “before” and “after”) are based on the TimeML approach (Pustejovsky et al., 2003); they can link event mentions and/or temporal

⁷<http://www.cnts.ua.ac.be/conll2002/ner/>

⁸<http://lands.let.ru.nl/projects/SoNaR/>

⁹<http://www.rae.es/recursos/banco-de-datos/corpes-xxi>

expressions.

Subordinating relations (relating two event mentions) are used for the annotation of reported speech (their annotation also leans on TimeML, although we have reduced their scope).

Causal relations are used to link causes and effects denoted by event mentions; we have annotated explicit causal relations taking into consideration the *cause*, *enable*, and *prevent* categories of causation.

Grammatical relations are created for events that are semantically dependent on another event (e.g. “make” in “make a call”) to link them to their governing content verb/noun (e.g. “call”).

Participant relations are one-to-one relations that link an event mention to an entity mention (or a numerical expression) that plays a role in the event. Participant relations are used to model semantic role labeling; PropBank (Bonial et al., 2010) is the reference framework for assigning semantic roles. For example in the sentence “Apple Inc. today has introduced the iPhone 4”, the event “introduced” has two participants: “Apple Inc.” as Arg0 (i.e. agent) and “iPhone 4” as Arg1 (i.e. patient).

Intra-document event and entity coreference The annotation of coreference chains that link different mentions to the same instance is based on the REFERS_TO relation.

Entity and event instances are described respectively through the non text-consuming ENTITY and EVENT tags and two attributes, i.e. tag descriptor and type.

Entity coreference annotation in MEANTIME follows the ACE 2008 guidelines (Linguistic Data Consortium,); for events, two mentions are considered as coreferring if their discourse elements (e.g. agents, location and time) are identical in all respects (Hovy et al., 2013), as far as one can tell from their occurrence in the text.

3.2. Annotation at corpus level

The cross-document annotation, based on the NewsReader cross-document annotation guidelines (Speranza and Minard, 2014), consists of linking event and entity instances annotated at the intra-document level.

More specifically, if two or more (entity or event) coreference chains annotated in different documents refer to the same instance, they are linked through a unique instance ID and the DBpedia URI (when available).

This annotation layer consists of:

- cross-document entity coreference for the coreference chains annotated at the document level;
- cross-document entity and event coreference in the whole document for a subset of 44 seed entities (i.e.,

annotation and coreference of all mentions referring to the seed entities and of the events of which the entities are participants).

4. The annotation process

The English documents have been annotated by six trained annotators using two different tools:

- CAT (Bartalesi Lenzi et al., 2012), for the annotation at the document level;
- CROMER (Girardi et al., 2014), for the annotation at the corpus level.

For the Italian, Spanish, and Dutch translations of the English articles, we devised a method to speed up the process and, at the same time, ensure cross-lingual annotation.

4.1. Cross-language projection for Italian and Spanish

The procedure we used for the annotation of the Spanish and Italian texts consists of the cross-lingual projection of annotations from the source text to the target text (Speranza and Minard, 2015). It involves five steps, starting with a file containing the source annotated text and its translation aligned at the sentence level.

Mention annotation The first step of the annotation is performed using the CAT tool and consists of annotating all markables (textual extent only).

Alignment The second step consists of the alignment between Italian/Spanish and English markables, which is done starting from files where the Italian/Spanish and the English texts have been aligned at the sentence level. Markable alignment is performed (using the CAT tool) by means of a new attribute associated with each markable, which takes as a value the corresponding English markable.

Automatic projection The automatic projection is performed using a Python script. It consists of importing the following English annotations into the Italian/Spanish texts: markable attributes (only non-language-specific attributes), event instances, entity instances, and relations (including the REFERS_TO relation which models intra-document coreference).

Manual revision Manual revision involves an overall check of the annotations imported automatically.

Projection of cross document coreference This consists of importing coreferences from the English section taking advantage of the alignment between English and Italian markables and extending the entity and event instances by importing English instance IDs and DBpedia URIs.

		First 5 sentences		Whole article	
		# sentences	# tokens	# sentences	# tokens
Airbus	(30 articles)	150	3,590	445	9,909
Apple	(30 articles)	149	3,423	462	10,343
GM	(30 articles)	148	3,636	428	10,063
Stock	(30 articles)	150	3,332	458	9,916
Total	(120 articles)	597	13,981	1,797	40,231

Table 1: Statistics about the 120 Wikinews articles in MEANTIME

	English	Dutch	Italian	Spanish
# files	120	120	120	120
# sentences	597	597	597	597
# tokens	13,981	14,647	15,676	15,843
EVENT_MENTION	2,096	1,510	2,208	2,223
EVENT_INSTANCE	1,717	1,210	1,773	1,744
ENTITY_MENTION	2,790	2,729	2,709	2,704
ENTITY_INSTANCE	1,292	1,325	1,281	1,271
TIMEX3	525	480	507	486
VALUE	418	412	415	404
SIGNAL	291	291	253	280
C-SIGNAL	29	61	35	40
REFERS_TO	2,983	2,516	3,054	3,015
TLINK	1,789	1,516	1,711	2,186
CLINK	50	48	61	61
GLINK	209	122	300	310
HAS_PARTICIPANT	1,978	1,930	1,865	2,152
SLINK	239	211	220	238

Table 2: Document level annotation in English, Dutch, Italian, and Spanish

4.2. Annotation and cross-language projection for Dutch

For Dutch, the complete document level annotation (i.e. markables, markable attributes and relations) has been performed manually using the CAT tool.

Then, in order to obtain a cross-lingually annotated corpus, the annotation in Dutch and English has been aligned and the cross-document annotation has been projected following a method similar to the one described above (Section 4.1.).

5. Corpus Description

5.1. Source Texts

The core of the corpus is composed of 120 English Wikinews articles written between 2004 and 2011. The articles were manually chosen to cover the following four topics: (i) Airbus and Boeing, (ii) Apple Inc., (iii) Stock market, and (iv) General Motors, Chrysler and Ford. Table 1 provides some general statistics about them.

The English articles were translated by professional translators into Spanish, Italian and Dutch. This ensured access to non-copyrighted articles in all project languages on the same topics, with the option to compare the results of the NewsReader pipeline in the different languages at a fine-grained level.

The annotation at the document level was performed for the headline and the first four sentences of each file, while the annotation at the corpus level spans over the whole document.

5.2. Data Format

The annotated texts are in the CAT labeled format, an XML stand-off format, where different annotation layers are stored in separate document sections and are related to each other and to source data through pointers. For each article we have an XML file containing the raw text and the annotations. In addition, a CVS file contains the list of instances shared by all sections of the corpus, with their type, DBpedia URI, time anchors and participants.

5.3. Corpus Statistics

The MEANTIME corpus is composed of 480 articles in four languages. In Table 2 we present statistics about the annotation done at the document level. For English, we annotated 2,096 event mentions and linked them to a total of 1,717 instances. For Italian and Spanish the number of event mentions is higher (over 2,200); this is partly due to the fact that in the Italian and Spanish sections modal verbs have been annotated as independent mentions.

In Table 3 we present some general statistics about the

	Stock	GM	Airbus	Apple	Total
# event instances cross-lingual	174	176	292	119	761
# entity instances cross-lingual	118	128	185	153	584
# event instances EN-NL	194	209	312	139	854
# event instances EN-IT	201	204	341	141	887
# event instances EN-ES	199	208	329	144	880
# event instances NL-IT	185	182	306	125	798
# event instances NL-ES	181	185	298	130	794
# event instances IT-ES	189	197	320	133	839
# event instances EN	215	237	350	156	958

Table 3: Cross-lingual annotation in the Wikinews corpus (EN: English, NL: Dutch, IT: Italian, ES: Spanish)

cross-lingual aspect of our data. In particular, we provide the number of event and entity instances annotated and linked in the four languages through a unique identifier. For example, in the “Stock Market” subcorpus we have 174 event instances and 118 entity instances shared by the four languages.

In the lower part of the table, we provide the number of event instances shared by language pairs and the number of event instances annotated in English. For example, in the “Airbus and Boeing” subcorpus, 312 event instances are annotated both in English and Dutch, while 320 are shared by Italian and Spanish.

To illustrate the cross-lingual coreference annotated in the corpus, we consider the following four aligned sentences in English, Italian, Dutch and Spanish:

EN: “The White House is considering an auto rescue plan”
 IT: “La Casa Bianca valuta il piano di salvataggio dell’auto”
 NL: “Het Witte Huis overweegt reddingsplan autosector”
 ES: “La Casa Blanca está considerando un plan de rescate del sector automovilístico”

The entity instance *White House*, which is represented through a unique ID and a DBpedia URI (http://dbpedia.org/page/White_House), is linked to the mentions “The White House”, “La Casa Bianca”, “Het Witte Huis” and “La Casa Blanca”. The event instance *considering*, which is represented through a unique ID and an *has-participant* relation with *White House*, is referenced by the mentions “considering”, “valuta”, “overweegt” and “considerando”.

6. Conclusions

We presented MEANTIME, a multilingual corpus of news articles annotated with entities and events at both the cross-document and cross-lingual level, which can serve as a benchmark dataset for English, Spanish, Italian and Dutch natural language processing tools. Recently, Basque translations of the English articles have been added and are currently being annotated with temporal information.

MEANTIME has been used within the NewsReader project to evaluate the NLP pipelines developed for the four languages as well as to conduct cross-lingual experiments.

The English section of the corpus was extended with timeline annotations and used in the SemEval 2015 TimeLine shared task (Minard et al., 2015a). The Dutch data have been used in a shared task at CLIN26.¹⁰ The Italian section will be used as test data in FactA,¹¹ an EVALITA 2016 shared task focusing on factuality annotation that is partially based on Minard et al. (2015b).

The English, Dutch and Spanish sections are freely available from <http://www.newsreader-project.eu/results/data/wikinews> and distributed under a CC-BY license, while the Italian section will be released in conjunction with the shared task.

Acknowledgments

This research was funded by European Union’s 7th Framework Programme via the NewsReader project (ICT-316404). We would like to thank the NewsReader team for its constant collaboration and support, and in particular Sara Tonelli and Rachele Sprugnoli for their work in the initial phase of the creation of the corpus.

References

- Artiles, J., Gonzalo, J., and Sekine, S. (2007). The SemEval-2007 WePS Evaluation: Establishing a Benchmark for the Web People Search Task. In *Proceedings of the 4th International Workshop on Semantic Evaluations*, SemEval ’07, pages 64–69. Association for Computational Linguistics.
- Bartalesi Lenzi, V., Moretti, G., and Sprugnoli, R. (2012). CAT: the CELCT Annotation Tool. In *Proceedings of the 8th International Conference on Language Resources and Evaluation (LREC’12)*, pages 333–338, Istanbul, Turkey, May. European Language Resources Association (ELRA).
- Bentivogli, L. and Pianta, E. (2005). Exploiting Parallel Texts in the Creation of Multilingual Semantically Annotated Resources: The MultiSemCor Corpus. *Nat. Lang. Eng.*, 11(3):247–261, September.
- Bentivogli, L., Girardi, C., and Pianta, E. (2008). Creating a gold standard for person cross-document coreference resolution in italian news. In *Proceedings of the LREC Workshop on Resources and Evaluation for Identity Matching, Entity Resolution and Entity Management*.
- Bonial, C., Babko-Malaya, O., Choi, J. D., Hwang, J., and Palmer, M. (2010). Propbank annotation guidelines, version 3.0. Technical report, Center for Computational Language and Education Research, Institute of Cognitive Science, University of Colorado at Boulder. http://clear.colorado.edu/compsem/documents/propbank_guidelines.pdf.
- Caselli, T., Sprugnoli, R., Speranza, M., and Monachini, M. (2014). EVENTI Evaluation of Events and Temporal Information at Evalita. In *Proceedings of the First Italian Conference on Computational Linguistic CLiC-it 2014 & the Fourth International Workshop EVALITA 2014*, pages 27–34.
- Cybulska, A. and Vossen, P. (2014). Using a sledgehammer to crack a nut? Lexical diversity and event coreference resolution. In *Proceedings of the 9th Language Resources and Evaluation Conference (LREC2014)*, Reykjavik, Iceland, May 26–31.
- Forascu, C. and Tufi, D. (2012). Romanian TimeBank: An Annotated Parallel Corpus for Temporal Information. In *Proceedings of the 8th International Conference on Language Resources and Evaluation (LREC’12)*, Istanbul, Turkey, May. European Language Resources Association (ELRA).
- Gavrilidou, M., Labropoulou, P., Piperidis, S., Giouli, V., Calzolari, N., Monachini, M., Soria, C., and Choukri, K.). Language resources production models: the case of INTERA multilingual corpus and terminology. In *Proceedings of the 5th International Conference on Lan-*

¹⁰<http://wordpress.let.vupr.nl/clin26/>

¹¹<http://facta-evalita2016.fbk.eu>

- guage Resources and Evaluation (LREC-2006), Genova, Italy.
- Girardi, C., Speranza, M., Sprugnoli, R., and Tonelli, S. (2014). CROMER: A Tool for Cross-Document Event and Entity Coreference. In *Proceedings of the 9th Language Resources and Evaluation Conference (LREC2014)*, Reykjavik, Iceland, May 26-31.
- Hovy, E., Mitamura, T., Verdejo, F., Araki, J., and Philpot, A. (2013). Events are Not Simple: Identity, Non-Identity, and Quasi-Identity. pages 21–28.
- ISO TimeML Working Group. (2008). ISO TC37 draft international standard DIS 24617-1, August 14. <http://semantic-annotation.uvt.nl/ISO-TimeML-08-13-2008-vankiyong.pdf>.
- Koehn, P. (2005). Europarl: A Parallel Corpus for Statistical Machine Translation. In *Proceedings of the tenth Machine Translation Summit*, pages 79–86, Phuket, Thailand. AAMT.
- Linguistic Data Consortium.). Technical report, June. http://projects.ldc.upenn.edu/ace/docs/English-Entities-Guidelines_v6.6.pdf.
- Magnini, B., Pianta, E., Girardi, C., Negri, M., Romano, L., Speranza, M., Lenzi, V. B., and Sprugnoli, R. (2006). I-CAB: the Italian Content Annotation Bank. In *Proceedings of the 5th Conference on Language Resources and Evaluation (LREC-2006)*.
- Minard, A.-L., Marchetti, A., and Speranza, M. (2014). Event Factuality in Italian: Annotation of News Stories from the Ita-TimeBank. In *Proceedings of the First Italian Conference on Computational Linguistic CLiC-it 2014*.
- Minard, A.-L., Speranza, M., Agirre, E., Aldabe, I., van Erp, M., Magnini, B., Rigau, G., and Urizar, R. (2015a). 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. Association for Computational Linguistics.
- Minard, A.-L., Speranza, M., Sprugnoli, R., and Caselli, T. (2015b). FacTA: Evaluation of Event Factuality and Temporal Anchoring. In *Proceedings of CLiC-it 2015, Second Italian Conference on Computational Linguistic*.
- Mouka, E., Giouli, V., Fotopoulou, A., and Saridakis, I. (2012). Opinion and emotion in movies: a modular perspective to annotation. In *Proceedings of the 4th International Workshop on Corpora for Research on Emotion, Sentiment & Social Signals (ES³ 2012)*, Istanbul, Turkey, May, 2012.
- Oostdijk, N., Reynaert, M., Hoste, V., and Schuurman, I. (2013). The Construction of a 500-Million-Word Reference Corpus of Contemporary Written Dutch. In Peter Spyns et al., editors, *Essential Speech and Language Technology for Dutch*, Theory and Applications of Natural Language Processing, pages 219–247. Springer Berlin Heidelberg.
- Padó, S. and Lapata, M. (2009). Cross-lingual Annotation Projection of Semantic Roles. *Journal of Artificial Intelligence Research*, 36(1):307–340, September.
- Postolache, O., Cristea, D., and Orasan, C.).
- Pustejovsky, J., Castaño, J. M., Ingria, R., Saurí, R., Gaizauskas, R. J., Setzer, A., Katz, G., and Radev, D. R. (2003). TimeML: Robust Specification of Event and Temporal Expressions in Text. In *New Directions in Question Answering*, pages 28–34.
- Real Academia Española. (2015). Banco de datos [online]. In *Corpus del español del siglo XXI (CORPES XXI)*. <http://www.rae.es/recursos/banco-de-datos/corpes-xxi>.
- Recasens, M. and Martí, M. A. (2010). AnCora-CO: Coreferentially annotated corpora for Spanish and Catalan. *Language Resources and Evaluation*, 44(4):315–345.
- Speranza, M. and Minard, A.-L. (2014). News-Reader Guidelines for Cross-Document Annotation. Technical Report NWR2014-9, Fondazione Bruno Kessler. <http://www.newsreader-project.eu/files/2014/12/NWR-2014-9.pdf>.
- Speranza, M. and Minard, A.-L. (2015). Cross-language projection of multilayer semantic annotation in the NewsReader Wikinews Italian Corpus (WItaC). In *Proceedings of CLiC-it 2015, Second Italian Conference on Computational Linguistic*.
- Spreyer, K. and Frank, A. (2008). Projection-based Acquisition of a Temporal Labeller. In *Proceedings of IJCNLP*, pages 489–496, Hyderabad, India, January.
- Surdeanu, M., Johansson, R., Meyers, A., Márquez, L., and Nivre, J. (2008). The CoNLL-2008 Shared Task on Joint Parsing of Syntactic and Semantic Dependencies. In *Proceedings of the Twelfth Conference on Computational Natural Language Learning*, CoNLL '08, pages 159–177.
- Taulé, M., Martí, M. A., and Recasens, M. (2008). AnCora: Multilevel Annotated Corpora for Catalan and Spanish. In *Proceedings of the 6th International Conference on Language Resources and Evaluation (LREC-2008)*.
- Tjong Kim Sang, E. F. (2002). Introduction to the CoNLL-2002 Shared Task: Language-Independent Named Entity Recognition. In *Proceedings of CoNLL-2002*, pages 155–158, Taipei, Taiwan.
- Tonelli, S., Sprugnoli, R., Speranza, M., and Minard, A.-L. (2014). NewsReader Guidelines for Annotation at Document Level. Technical Report NWR2014-2-2, Fondazione Bruno Kessler. <http://www.newsreader-project.eu/files/2014/12/NWR-2014-2-2.pdf>.
- UzZaman, N., Llorens, H., Derczynski, L., Allen, J., Verhagen, M., and Pustejovsky, J. (2013). SemEval-2013 Task 1: TempEval-3: Evaluating Time Expressions, Events, and Temporal Relations. In *Proceedings of the Seventh International Workshop on Semantic Evaluation, SemEval '13*, pages 1–9, Atlanta, Georgia, USA.
- van Son, C., van Erp, M., Fokkens, A., and Vossen, P. (2014). Hope and Fear: Interpreting Perspectives by Integrating Sentiment and Event Factuality. In *Proceedings of the 9th Language Resources and Evaluation Conference (LREC2014)*, Reykjavik, Iceland, May 26-31.