


→ Summary of the paper

Title	T-PAS; A resource of Typed Predicate Argument Structures for linguistic analysis and semantic processing
Authors	Elisabetta Jezeq, Bernardo Magnini, Anna Feltracco, Alessia Bianchini and Octavian Popescu
Abstract	The goal of this paper is to introduce T-PAS, a resource of typed predicate argument structures for Italian, acquired from corpora by manual clustering of distributional information about Italian verbs, to be used for linguistic analysis and semantic processing tasks. T-PAS is the first resource for Italian in which semantic selection properties and sense-in-context distinctions of verbs are characterized fully on empirical ground. In the paper, we first describe the process of pattern acquisition and corpus annotation (section 2) and its ongoing evaluation (section 3). We then demonstrate the benefits of pattern tagging for NLP purposes (section 4), and discuss current effort to improve the annotation of the corpus (section 5). We conclude by reporting on ongoing experiments using semiautomatic techniques for extending coverage (section 6).
Topics	Lexicon, Lexical Database, Ontologies
Full paper 	T-PAS; A resource of Typed Predicate Argument Structures for linguistic analysis and semantic processing
Bibtex	<pre>@InProceedings{JEZEK14.465, author = {Elisabetta Jezeq and Bernardo Magnini and Anna Feltracco and Alessia Bianchini and Octavian Popescu}, title = {T-PAS; A resource of Typed Predicate Argument Structures for linguistic analysis and semantic processing}, booktitle = {Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)}, year = {2014}, month = {may}, date = {26-31}, address = {Reykjavik, Iceland}, editor = {Nicoletta Calzolari (Conference Chair) and Khalid Choukri and Thierry Declerck and Hrafn Loftsson and Bente Maegaard and Joseph Mariani and Asuncion Moreno and Jan Odiijk and Stelios Piperidis}, publisher = {European Language Resources Association (ELRA)}, isbn = {978-2-9517408-8-4}, language = {english} }</pre>

T-PAS: A resource of corpus-derived Typed Predicate Argument Structures for linguistic analysis and semantic processing

Elisabetta Jezeq¹, Bernardo Magnini², Anna Feltracco²,
Alessia Bianchini¹, Octavian Popescu²

¹Università di Pavia, Strada Nuova 65, 27100 Pavia, Italy

²Fondazione Bruno Kessler, Via Sommarive 18, 38100 Povo-Trento, Italy

jezek@unipv.it, magnini@fbk.eu, feltracco@fbk.eu, alessia.bianchini01@universitadipavia.it, popescu@fbk.eu

Abstract

The goal of this paper is to introduce T-PAS, a resource of typed predicate argument structures for Italian, acquired from corpora by manual clustering of distributional information about Italian verbs, to be used for linguistic analysis and semantic processing tasks. T-PAS is the first resource for Italian in which semantic selection properties and sense-in-context distinctions of verbs are characterized fully on empirical ground. In the paper, we first describe the process of pattern acquisition and corpus annotation (section 2) and its ongoing evaluation (section 3). We then demonstrate the benefits of pattern tagging for NLP purposes (section 4), and discuss current effort to improve the annotation of the corpus (section 5). We conclude by reporting on ongoing experiments using semiautomatic techniques for extending coverage (section 6).

Keywords: predicate argument structure, corpus annotation, semantic type, ontology, evaluation of lexical resources

1. Introduction

This paper introduces T-PAS, a repository of typed predicate argument structures (T-PAS) for Italian acquired from corpora by manual clustering of distributional information about Italian verbs, freely available under a Creative Commons Attribution 3.0 license¹. T-PAS are corpus-derived verb patterns with specification of the expected semantic type (ST) for each argument slot, such as [Human] *guida* [[Vehicle]]. T-PAS is the first resource for Italian in which semantic selection properties and sense-in context distinctions of verbal predicates are characterized fully on empirical ground. In the resource, the acquisition of T-PAS is totally corpus-driven. We discover the most salient verbal patterns using a lexicographic procedure called Corpus Pattern Analysis (CPA, Hanks 2004), which relies on the analysis of co-occurrence statistics of syntactic slots in concrete examples found in corpora.

An important feature of T-PAS is that they are semantically motivated; different syntactic realizations are encoded as alternating subcategorization frames within the same T-PAS. For example, the T-PAS in (1) subsumes two distinct syntactic realizations (object and clausal object) for the same ST of argument of *finire* ‘finish’:

- (1) [[Human]-subj] *finisce* [[Event]-obj | di INF [V]]
 - a. *Finisce l'allenamento.*
 - b. *Non faccio in tempo a finire di bere la mia birra.*

Complements are included in T-PAS if they contribute to the way the verb is interpreted in the context of use. This

is how we define what counts as an argument, including adverbials. In this way, we offer an empirically grounded criterion to approach the traditional distinction between argument and adjunct, which is often questionable and hard to turn into robust generalizations. T-PAS are sense-stable objects, i.e. phrases where all the words are disambiguated; they provide the exact context carrying the relevant information for word senses. This has important consequences for the use of T-PAS in NLP tasks, as we will explain below. Moreover, in T-PAS the STs which are responsible for the sense of the verb in the context of the pattern (for example *arrestare*_{s1=arrest} [[Human]] vs *arrestare*_{s2=stop} [[Process]]) are not abstract categories but semantic classes discovered by generalizing over the statistically relevant list of collocates that fill each position; for example *arrestare*_{s2=stop} [[Process]]: {*emorragia, declino, corsa, marcia, caduta, flusso, crescita, desertificazione, erosione, epidemia ...*}.

Important reference points for the T-PAS project are FrameNet (Ruppenhofer et al. 2010) and VerbNet (Kipper-Schuler 2005). They differ from T-PAS because the structures they identify are not acquired from corpora following a systematic procedure (see, however, Bonial et al. 2013). Another important resource is PDEV (Hanks and Pustejovsky 2005), a pattern dictionary of English verbs which is the main product of the CPA procedure applied to English. As for Italian, a complementary project is LexIt (Lenci et al. 2012), a resource providing automatically acquired distributional information about verbs, adjectives and nouns. Differently from T-PAS, LexIt does not convey an inventory of patterns and the categories used for classifying the semantics of arguments are not corpus-driven. Inventory of senses such as MultiWordNet (Pianta et al. 2002) and Senso

¹ tpas.fbk.eu.

Comune (Oltramari et al. 2013) are resources to which T-PAS can be successfully linked with the goal of populating the former with corpus-driven pattern-based sense distinctions for verbs.

2. Resource Overview

T-PAS is being developed at the Dept. of Humanities of the University of Pavia, in collaboration with the Human Language Technology group of Fondazione Bruno Kessler (FBK), Trento and the technical support of the Faculty of Informatics at Masaryk University in Brno (CZ). The first release contains 1000 analyzed average polysemy verbs, selected on the basis of random extraction of 1000 lemmas out of the total set of fundamental lemmas of Sabatini Coletti 2008, according to the following proportions: 10 % 2-sense verbs, 60 % 3-5-sense verbs, 30 % 6-11-sense verbs.

The resource consists of three components:

- 1) a repository of corpus-derived T-PAS linked to lexical units (verbs);
- 2) an inventory of about 230 corpus-derived semantic classes for nouns, relevant for disambiguation of the verb in context;
- 3) a corpus of sentences that instantiate T-PAS, tagged with lexical unit (verb) and pattern number.

The reference corpus is a reduced version of ItWAC (Baroni & Kilgarriff, 2006), which was prepared at the the Laboratory of Natural Language Processing at Masaryk University (CZ) by J. Pomikalek by removing 7-grams duplicates from ItWac. Finally, we use a suite of corpus tools: Manatee, Bonito, and the Sketch Engine (Kilgarriff, Rychly, Smrz, Tugwell, 2004).

As referenced above, T-PAS specify the expected semantic type (ST) for each argument slot in the structure; in ST annotation, the analyst employs a “shallow” list of semantic type labels (HUMAN, ARTIFACT, EVENT, ecc.) which was obtained by applying the CPA procedure to the analysis of concordances for ca 1500 English and Italian verbs². These types look very much like conceptual / ontological categories for nouns but should instead be conceived as semantic classes, as they are induced by the analysis of selectional properties of verbs. They are derived by manual clustering and generalization over sets of lexical items found in the argument positions in the corpus. They are language-driven, and reflect how we talk about entities in the world. Despite the obvious correlations, they differ from categories of entities defined on the basis of ontological axioms, such as those of DOLCE (Descriptive Ontology for Linguistic and Cognitive Engineering, cf. Masolo et al., 2003).

Pattern acquisition and ST tagging involves the following steps:

1) choose a target verb and create a sample of 250 concordances in the corpus;

2) while browsing the corpus lines, identify the variety of relevant syntagmatic structures corresponding to the minimal contexts where all words are disambiguated;

3) identify the typing constraint of each argument slot of the structure by inspecting the lexical set of fillers: such constraints are crucial to distinguish among the different senses of the target verb in context. Each semantic class of fillers corresponds to a category from the inventory the analyst is provided with. If none of the existing ones captures the selectional properties of the predicate, the analyst can propose a new ST or list a lexical set, in case no generalization can be done;

4) when the structures and the typing constraints are identified, registration of the patterns in the Resource using the Pattern Editor (see Fig. 1). Each pattern has a unique identification number, and a description of its sense, expressed in the form of an implicature linked to the typing constrains of the pattern, for example the T-PAS in Fig. 1. has the implicature `[[Human]] legge [[Document]] con grande interesse`:

2. 11.6% `[[Human]] divorare [[Document]]`
`[[Human]] legge [[Document]] con grande interesse`

Fig. 1: Selected pattern for verb *divorare*

5) assignment of the 250 instances of the sample to the corresponding patterns, as shown in Fig. 2.

ID	Text	Label	Count
#3641905	coinvolgere dalla storia e l'ho letteralmente	DIVORATO	2
#96435540	coinvolto e lo consiglio a chi ha voglia di	divorare	2
#198128862	scrittura è nato nell'adolescenza, quando	divorare	2
#205411108	sono barricata in casa, mangio e studio.	divoro	2
#226711567	quotidianamente dai giornali, che venivano	divorati	2
#227271345	carina. Argomento: MANGA (e, visto che ne	divoro	2
#398209327	sfigato "quattrocchi" sempre immerso a	divorare	2
#422433394	poi gli avrei reso la cortesia! Mentre	divoravamo	2
#528732547	a chi ancora non lo ha letto, è di non	divorare	2
#642581344	aveva profondamente studiato, compreso, e	divorato	2
#646598005	infioati di Romero, mi butto su un libro che	divoro	2
#676881330	Non le importava di balli o teatri, ma	divorava	2
#742321485	L'anima cerca calore tra gli antichi e	divorati	2

Fig. 2: example of sample annotation for pattern 2 of *divorare* - SkE

In this phase, the analyst annotates the corpus line by assigning it the same number associated with the pattern. Concordances containing tagging errors are annotated as *x* and verb uses that do not come close to matching any of the normal patterns are tagged *u* (unclassifiable);

All above mentioned steps are explained in details in Guidelines, which are provided to analysts before starting the annotation.

3. Evaluation

3.1 Inter Annotator Agreement

In order to determine how much the annotation procedure is reliable and how much the task can be

² <http://corpora.fi.muni.cz/cpaeditor/onto.html>.

reproduced, we estimated the degree of inter annotator agreement among two annotators. We selected a sample of 50 verbs whose characteristics are representative of the whole T-PAS resources (i.e. one thousand verbs), and asked a second annotator, a persons with annotation experience and linguistic background, to develop a T-PAS for that sample. We adopted exactly the same procedure and annotation conditions, giving the second annotator the same set of 250 verbs used by the first annotator.

As explained in Section 2, the annotation task is quite complex, and includes two main interrelated steps: (i) defining a set of relevant patterns for a given verb, (ii) assigning examples (i.e. short sentences) to each of the pattern. In order to calculate the agreement we have considered the annotation as a clustering task, where the initial set of examples corresponds to the objects to be clustered, and each pattern defined by the annotator corresponds to a single cluster. Under this interpretation, the agreement between the two annotators corresponds to the degree of similarity among two distributions of the initial set of examples. To this aim, we have adopted two measures largely used in clustering: Purity and BCubed. Both the measures originate from Precision (purity), Recall (inverse purity) and their harmonic mean (F-measure), as used in Information Retrieval. Purity (Amigò et al., 2009) focuses on the frequency of the most common category into each cluster: it penalizes the noise in a cluster but it does not reward grouping items from the same category together. On the other side, Inverse Purity rewards grouping items together, but it does not penalise mixing items from different categories. Unlike Purity which computes independently the quality of each cluster and category, BCubed (Bagga & Baldwin, 1998) estimates the precision and recall associated to each object in the distribution. According to (Amigò et al., 2009) BCubed is the only measure sensible to several phenomena affecting clustering, including cluster homogeneity, cluster completeness, the “rag bag” phenomenon and cluster size versus quantity. In our case, cluster homogeneity rewards the situation where the two annotators do not mix examples belonging to different patterns. Cluster completeness rewards the case where annotators assign examples of a pattern to the same cluster (e.g. without splitting). The third aspect, “rag bag”, penalizes the situation where a non appropriate example is assigned to a well formed pattern (i.e. a pure cluster) rather than in a noisy one. Finally, the fourth aspect states that a small error in a big pattern should be preferable to a large number of small errors in small clusters. While BCubed is sensible to all the above mentioned aspects, Purity and Inverse Purity do not consider neither completeness and “rag bag”.

The selected examples of each verb which were considered either as “not pertinent” or as “undecided” by at least one of the two annotators were removed from the initial set. After this operation we got an average of 158 examples per verb, with a range from 120 to 190, which

have been used for clustering. The average number of patterns (i.e. clusters) per verb created by the two annotators was 3.82 for annotator 1 and 6.68 for annotator 2, showing a quite high variability. As an example, for the verb *crescere* (to grow), Annotator 1 has defined four patterns, while annotator 2 has defined eleven patterns. The difference is due a different interpretation of specific usages of the verb (e.g. “far crescere tutto il Movimento” – “make the Movement to grow”), which Annotator 1 decided to cluster together to one of the main sense of the verb:

[[Animate | Body Part | Plant]] *crescere* [NO OBJ]

while Annotator 2 has decided to assign to a specific pattern:

[[Human | Activity]] *fare crescere* [[Human Group]]

with very few examples.

The averaged BCubed (F1) per example is 0.77, which indicates that for big clusters (usually corresponding to frequent usages) there is a high overlap. For instance, despite the high variability in the number of patterns, the element BCubed for *comprare* is 0.90, (precision 0.83, Recall 0.98) showing that the distribution of the examples among the two annotators has a very high degree of overlap. In fact, most of the examples are assigned to four very populated patterns, while just a few examples are assigned to sparse clusters by Annotator 2. The average Purity (F-measure) is 0.60, as Purity, unlike BCubed, does not reward the cases, quite frequent in our sample, when all the examples for a certain pattern are grouped in a single cluster rather than split in several clusters.

Overall, given the complexity of the task (i.e. defining patterns and assign examples), we think that the inter annotator agreement show a very good reliability of the T-PAS resource.

3.2 Ongoing Evaluation

Ongoing experiments of evaluation focus on measuring inter-annotator agreement on pattern structure, and in particular on the following parameters: a) agreement on the span of context considered as pattern; b) agreement on ST tagging (cf. Cinková et al., 2012). As an example of disagreement on b) consider mismatches in ST tagging, which often arise due to metonymies or systematic polysemy of argument fillers (cf. Jezek and Quochi, 2010). According to the Guidelines, when corpus investigation reveals a mismatch between the ST specified in the pattern (*pattern type*) and the ST associated with the terms that populates it in the corpus line (*instance type*), annotators have two options: they can either tag non-canonical fillers as anomalous arguments (tag .a) at the level of concordance, or tag regular choices of STs as type alternations at the pattern level (tag |). As shown in Table 1, this may lead to disagreement; annotator 1 tagged the concordance B as

including an anomalous argument, while annotator 2 encoded an alternation in the object position of the pattern.

Corpus lines	Annotator	Pattern and Corpus Tagging
A. <i>Abbiamo raggiunto l'isola alle 5.</i> 'We reached the island at 5'	annotator 1	[[Human]] raggiunge [[Location]] Tagging: A. <i>Abbiamo raggiunto l'isola alle 5.</i> -> pattern1 B. <i>Ho raggiunto il semaforo e ho svoltato a destra</i> -> .a of pattern1
B. <i>Ho raggiunto il semaforo e ho svoltato a destra.</i> 'I reached the traffic light and turned right'	annotator 2	[[Human]] raggiunge [[Location Physical Object]] Tagging: A. <i>Abbiamo raggiunto l'isola alle 5.</i> -> pattern1 B. <i>Ho raggiunto il semaforo e ho svoltato a destra</i> -> pattern1

Table 1. Pattern 1 of *raggiungere* 'reach'

We improve this aspect by adding examples in the Guidelines and providing the annotators with a list of the most frequent type alternations found so far. Moreover, clustering techniques being developed within SkE will help the lexicographer to gain more insight about this phenomenon.

4. Computational Applications

T-PAS is an attractive resource to be exploited in several NLP tasks. The most natural application of T-PAS structures is for a wide range of applications where the participants of a certain event, represented by a predicate, need to be automatically detected. This is the case of Semantic Role Labeling (Gildea and Jurafsky, 2002) where the semantic arguments of a predicate are recognized with their role on the base of the syntactic structure of the sentence. In turn, semantic roles have shown to play a crucial role in more complex tasks, including Question Answering (Shen and Lapata, 2007), and the recognition of Textual Entailment (Dagan et al., 2009), where the semantic relation between two sentences often depends on the logical relation between predicates. In addition, T-PAS structures are relevant when the correct sense of a predicate-argument structure plays a crucial role, as in Machine Translation. For the above mentioned tasks T-PAS, and T-PAS resources in general, are an excellent source of features for approaches based on machine learning algorithms, as

both the predicate patterns and the example associated with them extracted from corpora provide a rich set of discriminative information, not easily available from other resources.

5. The Annotated Corpus

At present, the corpus is tagged with lexical item, pattern number and anomalous arguments (see section 3). Ongoing work (Bianchini, 2013) focuses on improving the annotation of the elements of the patterns onto the corpus, including type mismatches (cf. Pustejovsky, 2006). An example of type mismatch is reported in Table 2, where being the appropriate pattern for *annunciare* 'announce' [[Human1]] *annuncia* [[Event]] (a [[Human2]]), there is a mismatch between the pattern type [[Human1]] and the instance type, because the argument filler *altoparlante* 'loudspeaker' does not match the typing constraints specified in the pattern.

Pattern 1 <i>annunciare</i>	Corpus lines
[[Human1]] annuncia [[Event]] a [[Human2]]	<i>La segretaria annuncia l'arrivo del direttore</i> (matching). <i>L'altoparlante annunciava l'arrivo del treno</i> (type mismatch).

Table 2. Example of type mismatch

Preliminary work carried out for SemEval task 7 Argument Selection and Coercion (Pustejovsky et al., 2010)³ has shown that the phenomenon is pervasive (Jezek and Quochi, 2010) and spread over several of Levin's 1993 verb classes (aspectual verbs, communication verbs, perception verbs, directed motion verbs) as well as shift types (Artifact as Event, Artifact as Human, Artifact as Sound, Event as Location etc., cf. Jezek, 2012). The goal of the ongoing annotation is twofold: building an inventory of type mismatches to be used for linguistic analysis and metonymy resolution tasks and manually populating the ontology of STs, to be used in experiments for extending the resource (see section 6).

6. Techniques for extending coverage

Building manually a T-PAS resource is a time consuming process. Moreover, the errors in annotation due to human fatigue represent a big problem especially for pattern construction. To cope with these problems, we are implementing specific machine learning techniques in order to automatically extend the coverage of the resource. The main problem that needs to be addressed is the fact there is no dictionary that associates

³ In SemEval Task 7 we used patterns from T-PAS for 26 verbs as a reference for the Italian dataset (ca 4000 sentences of the PAROLE corpus) to be annotated.

each Italian word with its possible STs. While populating our own ontology (Section 5), we have resolved this problem by mapping the T-PAS semantic types to DOLCE's categories (Gangemi et al., 2002). We use the examples associated with each pattern to compute the confusion matrix between the set of semantic features on specific syntactic slots and patterns. Specifically, we have adopted an approach (Popescu, 2013) based on the implementation of the Angluin algorithm for learning regular languages (Angluin, 1987), through which we generate T-PAS patterns with semantic features.

In parallel, we addressed the issue of automatically recognize the occurrence of a certain T-PAS pattern in a text. This activity will potentially result in a T-PAS matcher, able to correctly individuate and disambiguate occurrences of verbs and their semantic participants and to link them to T-PAS. In this perspective we started developing a statistical model to match a pattern against raw text. The model is being applied to match patterns against raw text, with three main purposes:

- *Classify examples based on pre-existing patterns*: this will potentially allow to extend the set of sentences currently associated to each pattern, with an evident benefit for the whole resource.
- *Perform supervised semantic parsing*: the presence of sentences associate to patterns allows to apply supervised methods for training a semantic matcher (the T-PAS matcher). This will be crucial to exploit T-PAS in applications settings (see Section 4), particularly question answering and textual entailment.
- *Predict new patterns*: this is the most difficult task, where new patterns, with respect to those already manually defined in T-PAS, are automatically induced from text, either for a new verb or for a verb already present in the resource.

Finally, we believe that the statistical model can be extended to bilingual corpora, opening the possibility to cross-language alignments. In fact, using a bilingual English-Italian corpus we can infer the patterns in the target language by aligning the verb dependencies and the matched patterns in the source language (Popescu and Jezek, 2013).

7. Conclusions and future work

We have introduced T-PAS, a resource of corpus-derived typed predicate argument structures for Italian language. The current T-PAS release includes typed patterns for one thousand Italian verbs, and consists of three components: 1) a repository of corpus-derived T-PAS linked to lexical units (verbs); 2) an inventory of about 200 corpus-derived semantic classes for nouns, relevant for disambiguation of the verb in context; 3) a corpus of sentences that instantiate T-PAS, tagged with lexical unit and pattern number.

We have described the process of manual acquisition of

T-PAS, for which we have incrementally produced a rich document of guidelines for annotators. In addition, we have reported both the procedure and the result for the inter annotator agreement (F-BCubed of 0.77), showing how T-PAS is a reliable lexical resource, opening the way for computational usages in the NLP area.

T-PAS is distributed under a Creative Common license and is available through the META-SHARE catalogue.

Ongoing work includes the manual population of the ontology of semantic types and the creation of a gold standard of type mismatches. In addition, we are carrying on experiments for extending the coverage of the resource by finding new examples of existing patterns through a T-PAS pattern matcher.

8. Acknowledgements

The first author has been partially supported by Pavia University FAR (Finanziamento Ateneo alla Ricerca) 2012. The second and third authors have been partially supported by the EU project Excitement (Exploring Customer Interaction Through Textual Entailment) under the grant FP7 ICT-287923. The fourth author has been supported by a PhD grant of the University of Pavia. The last author has been partially supported by the BCROCE project, under a grant of the Province of Trento. We thank Patrick Hanks for his scientific support for the Italian project. We thank the Faculty of Informatics at Masaryk University in Brno (CZ) for the technical support, whose work has been partly supported by the Ministry of Education of CR within the LINDAT-Clarín project LM2010013. We are grateful to An Vo Ngoc for his support during the evaluation phase.

9. References

- E. Amigò, J. Gonzalo, J. Artiles and F. Verdejo. (2009). A comparison of extrinsic clustering evaluation metrics based on formal constraints. *Information Retrieval*, Vol. 12, pp. 461-486.
- D. Angluin. (1987). Learning regular sets from queries and counterexamples. *Inf. Comp.*, 75(2), pp. 87-106.
- A. Bagga and B. Baldwin. (1998). Entity-Based Cross-Document Coreferencing using the Vector Space Model. In *Proceeding of Coling-ACL 1998*, pp. 79-85.
- M. Baroni and A. Kilgarrieff. (2006). Large Linguistically-Processed Web Corpora for Multiple Languages. In *EACL 2006 Proceedings*, pp. 87-90.
- A. Bianchini. (2013). Towards a corpus annotated with syntactic and semantic verbal patterns for Italian. *Wolverhampton CPA Analysis Workshop*, University of Wolverhampton, Aug. 28-29, 2013.
- C. Bonial, O. Hargraves and M. Palmer. (2013). Expanding VerbNet with Sketch Engine. In *Proceedings of the 6th International Conference on Generative Approaches to the Lexicon*. Pisa, ILC - CNR, Sept. 24-25, 2013.
- S. Cinková, M. Holub and V. Kříž. (2012). Managing Uncertainty in Semantic Tagging. In *Proceedings of 13th Conference of the European Chapter of the*

- Association for Computational Linguistics*, Association for Computational Linguistics, Avignon, France, pp. 840-850.
- I. Dagan, B. Dolan, B. Magnini and D. Roth. (2009). Recognizing textual entailment: Rational, evaluation and approaches. In *Journal of Natural Language Engineering*: Volume 15, issue 04: Cambridge University Press, pp. 1-17.
- D. Gildea and D. Jurafsky. (2002). Automatic Labeling of Semantic Roles, *Computational Linguistics*, Volume 28, Number 3.
- A. Gangemi, N. Guarino, C. Masolo, A. Oltramari, L. Schneider L. (2002). Sweetening Ontologies with DOLCE. In A. Gómez-Pérez and V.R. Benjamins (eds.) *Proceedings of the 13th International Conference on Knowledge Engineering and Knowledge Management (EKAW'02)*, Ontologies and the Semantic Web, Berlin, Springer-Verlag, pp. 166-181.
- P. Hanks. (2004). Corpus Pattern Analysis. In G. Williams and S. Vessier (eds.). *Proceedings of the XI EURALEX International Congress*. Lorient. France (July 6-10, 2004), pp. 87-98.
- P. Hanks and J. Pustejovsky. (2005). A Pattern Dictionary for Natural Language Processing. In *Revue française de linguistique appliquée*, 10 (2).
- A. Kilgarriff, P. Rychly, P. Smrz and D. Tugwell (2004). The Sketch Engine. In G. Williams and S. Vessier (eds.). In *Proceedings of the XI EURALEX International Congress*. Lorient. France (July 6-10, 2004), pp. 105-111.
- K. Kipper-Schuler. (2005). *VerbNet: A broad coverage, comprehensive verb lexicon*. Ph.D. Thesis. University of Pennsylvania, USA.
- E. Jezeq. (2012). Acquiring typed predicate-argument structures from corpora. In Bunt, H. (ed.) *Proceedings of the Eighth Joint ISO - ACL SIGSEM Workshop on Interoperable Semantic Annotation ISA-8*, Pisa, Università di Pisa - CNR, (October 3-5, 2012), pp. 28-33.
- E. Jezeq and V. Quochi. (2010). Capturing Coercions in Texts: a First Annotation Exercise. In *Proceedings of the Seventh conference on International Language Resources and Evaluation (LREC'10)*. Valletta, Malta (May 19-21, 2010), pp. 1464-1471.
- P. Lauwers and D. Willems. (2011). Coercion: Definition and Challenges, current approaches and new trends. In *Linguistics* 49:6, pp. 1219-1235.
- A. Lenci, G. Lapesa and G. Bonansinga. (2012). LexIt: A Computational Resource on Italian Argument Structure, in *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC'12)*, Istanbul (May 23-25, 2012), pp. 3712-3718.
- B. Levin. (1993). *English Verb Classes and Alternations: A Preliminary Investigation*, Chicago, The University of Chicago Press.
- C. Masolo, S. Borgo, A. Gangemi, N. Guarino, A. Oltramari. (2003). *Wonderweb deliverable D18: Ontology Library (final)*, IST Project 2001-33052 WonderWeb: Ontology Infrastructure for the Semantic Web.
- A. Oltramari, G. Vetere, I. Chiari, E. Jezeq, F. M. Zanzotto, M. Nissim and A. Gangemi. (2013). Senso Comune: A collaborative Knowledge Resource for Italian. In Iryna Gurevych and Jungi Kim (eds.). *The People's Web Meets NLP: Collaboratively Constructed Language Resources*, Berlin-Heidelberg, Springer, pp. 45-68.
- E. Pianta, L. Bentivogli and C. Girardi. (2002). Multiwordnet: developing an aligned multilingual database. In *First International Conference on Global WordNet*, Mysore, India.
- O. Popescu. (2013). Learning Corpus Pattern Using Finite State Automata. In *Proceeding of IWCS*, Potsdam.
- O. Popescu and E. Jezeq. (2013). Phrase based Translation. In *Proceedings of Tralogy – the quest for meaning*, Paris.
- J. Pustejovsky. (2006). Type theory and lexical decomposition. In *Journal of Cognitive Science* 6, pp. 39-76.
- J. Pustejovsky, A. Rumshisky, A. Plotnick, E. Jezeq, O. Batiukova and V. Quochi. (2010). SemEval Task 7: Argument Selection and Coercion. In *Proceedings of the Fifth International Workshop on Semantic Evaluation*. Uppsala University, Sweden, July 11-16 2010.
- J. Ruppenhofer, M. Ellsworth, M. R. L. Petruck, C. R. Johnson and J. Scheffczyk. (2010). *FrameNet II: Extended theory and practice*. International Computer Science Institute, University of Berkeley. (Manuscript, Version of September 14, 2010).
- F. Sabatini and V. Coletti (2007). *Il Sabatini-Coletti. Dizionario della lingua italiana 2008*, Milano, Rizzoli-Larousse.
- D. Shen and M. Lapata. (2007). Using Semantic Roles to Improve Question Answering. In *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, Prague (June 2007), pp. 12-21.