AperTO - Archivio Istituzionale Open Access dell'Università di Torino

Semantic Annotation of Legal Texts through a FrameNet-Based Approach

This is a pre print version of the following article:

Original Citation:

Availability:

This version is available <http://hdl.handle.net/2318/117691> since 2017-05-23T10:22:32Z

Publisher:

Springer

Published version:

DOI:10.1007/978-3-642-35731-2_17

Terms of use:

Open Access

Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.

(Article begins on next page)

Semantic Annotation of Legal Texts through a FrameNet-Based Approach

Marcello Ceci¹, Leonardo Lesmo², Alessandro Mazzei²,
Monica Palmirani¹, and Daniele P. Radicioni²

¹ University of Bologna, CIRSFD, Italy

{marcello.cecimonica.palmirani}@unibo.it

² University of Turin, Computer Science Department, Italy

{lesmo, mazzei, radicioni}@di.unito.it

Abstract. In this work we illustrate a novel approach for solving an information extraction problem on legal texts. It is based on Natural Language Processing techniques and on the adoption of a formalization that allows coupling domain knowledge and syntactic information. The proposed approach is applied to extend an existing system to assist human annotators in handling normative modificatory provisions –that are the changes to other normative texts–. Such laws ‘versioning’ problem is a hard and relevant one. We provide a linguistic and legal analysis of a particular case of modificatory provision (the *efficacy suspension*), show how such knowledge can be formalized in a linguistic resource such as FrameNet, and used by the semantic interpreter.

Keywords: Knowledge modelling, semantic interpretation, NLP.

1 Introduction

Legal systems are dynamic by nature, since they change over time. Modifications affect legal texts, their temporal properties, and even the meaning of the norms expressed in those texts. Many efforts have been invested in the last ten years towards the digitalization in the legal domain. Researches to produce updated collections of legal documents on the Web are being conducted with multiple aims, such as intelligent indexing, querying, searching, filtering, retrieval of documents or of meaningful parts, and to help managing changes in the legal content, through the so-called *consolidation process*. The digitalization process requires solving two sorts of problems: defining (XML) file formats to conveniently encode the texts, and designing systems to assist human experts in the annotation of the legal texts according to a format devised at the previous step. Much work has been done in both directions. Various initiatives have been established at the national and international levels to devise XML standards for describing legal sources and schemas to identify legal documents [8]. Also, systems have built that automatically identify and classify structural portions of legal documents and their intra- and inter-references [2,12]; the problem of semantic analysis is currently being investigated[15]. Unfortunately, due to the

‘natural language barrier’ (i.e., the problem of translating a sentence into some form of semantic interpretation [11]), this is still an open problem. This paper is concerned with this problem, that is the extraction of modificatory provisions and their annotation.

One main aspect makes legal texts suitable for applying information technologies commonly used to deal with hypertexts. Legal texts contain references to other legal texts or to other parts within the same document, so that a legal text can be naturally considered as a particular case of hypertext.¹ To have a sound example, let us consider the following example of *consolidation* problem: a legal document *A* contains a reference to document *B*. Say, e.g., that *A* contains a locution such as ‘the first article of the law number 9 in the document *B* is suspended until January 29, 2011’. Unfortunately, a person interested in inspecting the validity of the norms in *B* could encounter some problems to figure out whether the norms in *B* are still valid, in force, etc., because *B* contains no reference to *A* (that is, it is not possible to add backward pointers to existing legal documents). Under this perspective, legal systems can be seen as tangled webs. It should be noted that at least for some normative systems –such as the Dutch and the Italian ones (see [9] and [15,3,4], respectively)–, the consolidation problem is a relevant one. In fact, the uncertainty on the effects of normative modifications would undermine the certainty of the law, making it hard to clearly understand which one of several versions of a provision counts as law. Automating the process of semantically annotating modificatory clauses and provisions would be of great help in simplifying the legal system and in consolidating texts of law,² because the human annotation process is expensive and error-prone. From a practical perspective, the consolidation process involves identifying the main elements of the modificatory provisions, annotating them in the legal text according to a given DTD (we adopt the NIR standard, but in principle in a different context another standard could be adopted), and generating a set of metadata that compactly describes the considered modification.

In past works we detected some regularity in the linguistic structure of modificatory provisions [14], and we showed how this regularity, coupled with a XML markup [12] can be used by automated tools to qualify a modificatory provision [13]. In particular, our approach relies on a tree-matching technique to put together deep parsing and shallow semantic interpretation [7]. In the present work we extend our approach by devising a specialised version of FrameNet [1] to cope with modificatory provisions. In particular, we model the efficacy suspension modificatory provisions. The paper is structured as follows: we first illustrate the considered problem of automatically annotating XML files with information describing the modificatory provisions. We then consider the case of the efficacy suspension, which is by far more complex on a linguistic perspective, and argue that it requires enhanced modelling efforts with respect to integration,

¹ According to WordNet, hypertext is a “machine-readable text that is not sequential but is organized so that related items of information are connected”, <http://wordnetweb.princeton.edu/perl/webwn?s=hypertext>

² *Consolidated text* is the updated version of a normative text, embodying the changes.

substitution and repeal modifications, that have been previously considered. Finally we illustrate how the novel approach can be integrated in the implemented system by considering how a sentence containing a suspension is extracted from the XML input format, processed in the syntactic analysis phase, and then how the main elements of the suspending modification are individuated by the semantic interpreter, based on the FrameNet formalization.

2 Suspension Analysis

The Suspension is the action by which a textual provision interrupts the efficacy of a legal text (or fragment thereof) for a given period. It is important to deal with suspensive modifications because from a linguistic perspective they are more complex and rich than other temporal modifications. Furthermore, suspension is a relevant modification in that it is often used as a *legislative drafting technique* for introducing a *temporary* law. This need stems from for two main reasons: *i*) when the topic is so complex but urgent that it is necessary to have a temporary solution (e.g. Genetic Law); and *ii*) when some time is needed to fully apply the new dispositions (e.g. Euro Law in 1999). The rationale underlying the Suspension of Efficacy is that some norms so strongly affect their addressees (citizens, businesses, social actors) that some time is needed to tune them up. One ambitious and long-term goal is to track this rationale over time. Recognizing the suspension process even if it is fragmented across several intervals of efficacy will allow unveiling that each macro-suspension is driven by a normative principle.³

The suspension can be either *explicit* or *implicit*, depending on the language of the provision in question. And, temporally, it can be either *defined* or *undefined*. A suspension is *defined* when the period during which a norm efficacy is interrupted is explicitly stated in the text, with the suspending provision clearly indicating a beginning and an end (or an initial and a final event). By converse, a suspension is *undefined* when this time interval is not explicitly set out in any part of the suspending provision. This class of suspensions includes at least three subclasses as follows: *(i)* *sine die* (that is, with no ending date) suspension; *(ii)* suspension conditioned by an external event (e.g., “Article 5 is suspended for a six-month period starting from entry into force of the Treaty [...]”); and *(iii)* suspension intervals described with a set of other parameters such as the duration (e.g., “Article 5 is suspended for four months starting from January 31, 2011”). In these cases, extracting the correct values may be a complex task.

Suspension modificatory provisions are themselves subject to *modification*. A suspension can be reflexive, with the law introducing the suspension being the same as that affected by the suspension. However, it is rather habitual that later provisions are compiled to modify that suspension for the same reasons that led to its introduction. For example, the Decision 2000/185/EC (Article 3) said

³ Such as the principle that all norms on the use of human embryonic material will remain suspended until a coordinated regulatory framework is in place. Ordinanza 30 maggio 2003 (GU n. 158 del 10/07/2003).

that the decision itself “shall apply from January 1, 2000 to December 31, 2002”, thus limiting the document efficacy. Later, the Decision 2002/954/EC modified the second subparagraph of Article 3 by replacing “December 31, 2002” with “December 31, 2003”. Then, finally, a third Directive again changed the term, from ‘2003’ to ‘2005’. The rationale guiding this suspension remains the same, and it is important to grasp it by first detecting the arguments that characterize the suspension modification—so to identify and adjust the main suspension of efficacy—and then describing the phenomena in their atomicity. Another type of suspension provision is the *disapplication*. When a document “disapplies” another document, the latter is “frozen,” its efficacy being suspended.⁴

3 Extraction of Suspension Modifications

The annotation of modificatory provisions is a three steps process. Although this process has been illustrated in previous work (full details are provided in [10]), we briefly recall them in order to make the paper more complete and readable. We then show how the FrameNet formalization is used in the semantic interpretation process, pointing out the benefits due to encoding the knowledge about modifications in declarative form.

3.1 System Architecture

In the first step we look for the possible location of a modificatory provision within the document, and we simplify the input sentences, so to prune text fragments that do not convey relevant pieces of information (*input preprocessing*). In the second step we perform the syntactic analysis (*parsing*) of the retrieved sentences; in the third step (*semantic interpretation*) we semantically annotate the retrieved provisions through a tree matching approach. We briefly recall the first two steps and then focus on the annotation phase and on the semantic interpreter design.

The input to the system is encoded in the NormeInRete (NIR) XML standard format for Italian Legal Text. The NIR format encodes the structural elements used to mark up the main partitions of legal texts, as well as its atomic parts (such as articles, paragraphs, subparagraphs, and lettered and numbered items) and any non-structured text fragment. Additionally, the NIR standard includes in its Document Type Definitions a part describing modifications, to implement this model in XML. Based on the XML structure, we retain the text excerpts contained between some meaningful tags (e.g., `<corpo>`, which is the Italian word for body, where the modifications may be found). The text tagged by `<rif>` (Italian abbreviation for reference) and `<virgolette>` (Italian word for quotes) is then rewritten with the IDs of the corresponding tags. For example, given the XML encoding of a sentence such as “L’efficacia del decreto ministeriale 17 novembre

⁴ Disapplication may be motivated by various legal phenomena, such as the aim at resolving conflicts of laws between regional and national law or between national law and European regulations.

2006 è sospesa fino alla data del 30 aprile 2007” (*The efficacy of the Ministerial Decree is suspended until the date of April 30th, 2007*), we rewrite the sentence we rewrite the sentence like “L’efficacia del RIF12 è sospesa fino alla data del 30 aprile 2007”. This sentence, which is much simpler to analyze with no loss of information, is then given in input to the parser.

The TULE parser is a broad coverage rule based parser for Italian [6]: it returns a dependency tree that represents the syntactic analysis of the source Italian sentence. It relies on a morphological dictionary of Italian (about 25,000 lemmata) and on a rule-based grammar that describes dependency structures. Let us consider again the sentence: “L’efficacia del RIF12 è sospesa fino alla data del 30 aprile 2007”. After two preliminary steps (the morphological analysis and part of speech tagging), necessary to recover the lemma and the part of speech (PoS) tag of the words, the words sequence goes through three phases: chunking, syntactic analysis of the coordination, and verbal subcategorization. The parser produces in output a dependency tree that makes explicit the structural syntactic relationships occurring between the words of the sentence. Each word in the sentence is associated with a node of the tree, as depicted in Figure 1.⁵ The nodes are linked via labeled arcs that specify the role of the dependents with respect to their governor (the parent). In the considered example, “efficacia” (*efficacy*) is the subject of the verb “(è) sospesa” (*(is) suspended*), while “è” (*is*) is the auxiliar, marked with *aux*. A special node “trace” is framed by a dashed line and labeled *t*: it specifies that the deep subject of the suspension (the agent, in terms of roles) is not expressed. Finally, the temporal argument is in a dependent that is labeled as a *modifier*, tagged as *RMOD* in Figure 1.

3.2 The Interpretation Process

Modifications are represented by means of semantic frames, composed by slots [5]. Retrieving a modificatory provision amounts to choosing the frame describing that modification, and to filling its slots with the correct arguments. Alternatively, annotating a modificatory provision means that given a modification description we are able to recognize it in a sentence. The task of the semantic interpreter is twofold. First it consists in inspecting the dependents of the verb on the one hand, and in inspecting the frames and the available syntactic and semantic information on the other hand. Then the semantic interpreter is charged to find the frame that best fits to current setting. Secondly, once the appropriate frame has been individuated, the related set of rules is applied to retrieve the fillers for the frame slots. The information stored in the FrameNet formalization is thereby fundamental, since it provides a necessary interface between the syntactic and the semantic levels. Additionally, it allows formalizing syntactic and semantic knowledge about modificatory provisions in a *declarative* (as opposed to *procedural*) manner. That is, the FrameNet formalization allows illustrating the rationale underlying and governing the application of rules, since it puts

⁵ Actually, the nodes include further data (e.g., the gender and number for nouns and adjectives and verb tenses) which do not appear in the figure for space reasons.

Fig. 1. The (simplified) *dependency tree* structure for the input sentence “L’efficacia del RIF12 è sospesa fino alla data del 30 aprile 2007” (*The efficacy of the rif12 is suspended until the date of April 30th, 2007*)

together both the information about the modification, and their grammatical and syntactical possible realizations.

FrameNet Encoding. FrameNet is a lexical database that represents concepts related to events, relations and states in terms of *semantic frames* [1]. Some features make FrameNet particularly well-suited to our modeling purposes. Frames can be thought of as concepts, composed by sub-elements (called *frame elements*, FEs), that act as semantic roles. Words meaning is encoded through *lexical units* (LUs) that are the FrameNet counterpart of words *senses* in a traditional dictionary.⁶ Moreover, which is perhaps more relevant to our present ends, for each such lexical unit an annotation is provided, where the possible realizations of that LU are mapped onto a syntactic structure. The annotated component of FrameNet is of the highest relevance to computational approaches to linguistics (be them based on hand-crafted rules, or acquired through machine learning techniques), in that it provides fully analyzed working examples for each lexical unit. FrameNet retains also information on parts of speech (PoS) such as verbs, adjectives, nouns, etc., so that it can be exploited at these levels.

⁶ This implies, e.g., that polysemous words are represented by different lexical units.

We are currently developing some FrameNet-like frames (ideally extending the original FrameNet) to encode the legal knowledge needed for recognizing the main features of suspensive provisions (and their variants, such as the modified suspension and the disapplication). To provide an account for suspension modifications it is possible to start by devising two frames, the *Efficacy_Inclusion* and *Efficacy_Exclusion* frames. Such frames are composed of the elements illustrated below and are part of a *Main_Suspension* frame, which can be thought of as an *abstract class* to be implemented through the *Efficacy_Inclusion/Exclusion* frames (Table 1).

Table 1. The basic frames for *Efficacy_Inclusion* and *Efficacy_Exclusion*

frame (Efficacy_Inclusion) frame_elements (Passive_Norm, Period_Start, Period_End) scene (<i>Passive_Norm</i> has efficacy from <i>Period_Start</i> to <i>Period_End</i>)
frame (Efficacy_Exclusion) frame_elements (Passive_Norm, Period_Start, Period_End) scenes (Passive_Norm is suspended from <i>Period_Start</i> to <i>Period_End</i> , <i>Passive_Norm</i> has not efficacy from <i>Period_Start</i> to <i>Period_End</i> , <i>Passive_Norm</i> has efficacy until <i>Period_End</i>)
frame (Main_Suspension) frame_elements (Passive_Norm, Suspension_Start, Suspension_End) scenes (<i>Passive_Norm</i> is suspended from <i>Suspension_Start</i> to <i>Suspension_End</i> , <i>Passive_Norm</i> has efficacy from <i>Suspension_End</i> to <i>Suspension_Start</i>)

Also, a *Suspension_Modification* frame can be used to describe provisions modifying a suspension previously introduced by another norm (please refer to the analysis of suspensions, Section 2). It is fairly easy to distinguish between the two kinds of provisions, since they are textual modifications lacking a term that evokes an *Efficacy* frame and contains some *Change_event_time* frame. In order to properly interpret the modification, there needs to be a comparison between the *Suspension_Modification* and the *Main_Suspension* (contained elsewhere). For this reason, the *Suspension_Modification* element is presented without any semantic specification of its content, since the exact interpretation of the provision is entrusted to the semantic interpreter.

We have collected a set of relevant terms, that evoke either the *Efficacy_Inclusion* or the *Efficacy_Exclusion* frame. The ‘.n’, ‘.v’, etc. notation reports about PoS information for nouns, verbs, adjectives, and so forth (Table 2). The *TemporalArguments* of the shift in efficacy is captured by the *Period_Start* and *Period_end* Frame Elements (FEs), and the target norm is marked as *Passive_Norm*. Frame Element Groups (FEGs) represent the occurrence of FEs in the examined provisions (P=*Passive_Norm*, S=*Period_Start*, E=*Period_End*). Some typical examples of annotated suspensions are provided in Table 3.

The *Main_Suspension* frame is modelled by inheriting the *Process* frame. The *Suspension* is therefore treated as a process, with a “target” represented by the *Passive_Norm* and whose state is affected by one or more events: it starts with the *Suspension_Start* event and/or ends with the *Suspension_End* event.

Table 2. Terms relevant, *evoking* the efficacy of suspension

Efficacy_Inclusion (efficacia. <i>n</i> , efficace. <i>adj</i> , applicarsi. <i>v</i> , valido. <i>adj</i> , validità. <i>n</i> , effetto. <i>n</i> , applicazione. <i>n</i> , vigore. <i>n</i>)
Efficacy_Exclusion (sospendere. <i>v</i> , disapplicare. <i>v</i> , cessare. <i>v</i> +efficacia. <i>n</i> , non. <i>adv</i> +Efficacy_Inclusion)

Table 3. Example of annotated sentences containing efficacy suspensions

FEG	Annotated Example
P, S, E	[L'obbligo di cui all'articolo 51, comma 1, della legge 27 dicembre 2002, n. 289] _P , è sospeso [dalla data di entrata in vigore del presente decreto] _S [fino al 31 dicembre 2006] _E .
P, E	[Le disposizioni del presente provvedimento] _P hanno efficacia [sino a tutto il 7 maggio 2007] _E .
P, S+E	[Le disposizioni della legge 29 dicembre 1988 n. 554] _P si applicano [negli anni 1989 e 1990] _{S+E}

Moreover, the start of the process can be advanced or postponed by another norm, and the same can be done to its end. These events are represented by specific frames, subclasses of the *Suspension_Modification* frame that are presently not reported for lack of space.

The FrameNet model described above is designed to deal with legislative texts encoded in XML format, with some elements already annotated, in a supervised manner. A parser called Norma-Editor automatically detects *references*, *dates*, and allows adding metadata in legislative texts [14]. Norma-Editor is employed to convert legal texts in a XML format based on Legal XML standards (such as Akoma Ntoso and NiR, [2]). The XML file is then given in input to the TULE parser. The FrameNet modelling helps us clearly investigate and understand the possible linguistic realizations of suspensions and how such information can be exploited by a syntactic interpreter. Efficacy-evoking terms help us formulate an hypothesis on the type of provision being examined: for example, if the evoking word occurs as the subject, then the prepositional phrase is marked as *Passive-Norm* (as in “Efficacy of law X”). Also, if the evoking word occurs as the predicate, the *Passive-Norm* element will be represented by the subject (“Law X is suspended”). Words and locutions expressing (the beginning or the end of) a time span are marked as *Period_Start* and *Period_End*.

Arguments Extraction. After describing how legal and linguistic knowledge is represented in FrameNet terms, we show how such knowledge is used by the semantic interpreter.

The semantic interpreter is charged to test whether the root node of the syntactic tree is a verb, and if it belongs to the taxonomy of the verbs relevant to modificatory provisions (see [10]). For example, given the parse tree in Figure 1, we take the verb lemma *sospendere* (*suspend*), search for it in the knowledge

base, and find that it is a possible instantiation of a modification whose *legal-Category* property is *suspension*, together with the verbs *disapplicare* (*to cease to apply*), *applicarsi* (*enforce*), etc.. In this case we have a fundamental cue that the sentence being analyzed contains a modificatory provision and the semantic interpreter is triggered. We note that there is potentially a terminological collision, in that the frame allocated (and to be filled) is a data structure that can be thought of as an object, and has nothing to do with the frames of FrameNet. Once the frame is allocated, the main task of the semantic interpreter consists in filling its slots. To identify the main elements of the *efficacy inclusion* and *efficacy exclusion* modifications we have to retrieve the information needed to fill the following slots: *Passive_Norm*, [*Position*], *Start* and *End*.

Once discovered that the modification is probably a suspension, the appropriate set of rules is executed so to exploit the information grasped through the FrameNet formalization to retrieve the correct slot fillers from the parse tree. Filling a modification frame amounts to finding an appropriate mapping between tree dependents and frame slots. To carry on with the sentence under consideration, let us consider a typical realization for the *Efficacy_Exclusion* frame:

Passive_Norm is suspended from *Start* to *End*. (1)

By introducing the terms used above, we can rewrite the previous sentence as:

P is suspended from *S* to *E*. (2)

In practical cases it may happen that either the *Start* or the *End* argument is lacking, therefore determining an *open* time span, where one of the two temporal arguments may be absent. Among many possible variants of the sentence in (3.2), a slightly different linguistic construction can be

The efficacy of *P* is suspended from *S* to *E*. (3)

Once the semantic interpreter recognizes a particular surface realization, further relevant information can be made available and exploited, that is directly related to the syntactic structure:

[The efficacy of *P*]_{subj} is suspended [from *S*]_{rmod} [to *E*]_{rmod}. (4)

Like it is apparent from this simple example, the FrameNet formalization provides a compact description for (some of) the possible syntactic realizations of the modificatory provisions. That is, the locution “The efficacy of *P*” is expected to occur in a branch of the parse tree rooted under the main verb. Namely, the semantic interpreter inspects the branch containing the subject of the sentence, labeled *verb-subj*. The processing of such tree branch allows extracting the reference to the *passive norm*. Similarly, extracting both the *Start* and the *End* time will imply traversing the tree branches labeled *RMOD* (see Figure 1). As suggested in the description of the frame, the presence of words/locutions such as “a partire da” (*starting from*), “a decorrere da” (*starting day will be*) or “fino a”, “sino a” (*until*) will provide precious cues about where to find the starting and ending time of the suspension time span.

Triggered by the recognition of the root verb, the set of rules related to each modification are executed to test the content of the verb arguments and the verb modifiers to fill the slots of current frame. The rules are charged to discover whether in the syntactic arguments like *subject*, *object* or in any modifier are present any meaningful locutions or constants, such as *RIF*. In this way we can conveniently map the syntactic pattern described in the FrameNet formalization onto the set of slots of a semantic frame.

4 Conclusions

In this paper we provided a linguistic account and syntactic analysis for a particular type of modificatory provision, that is efficacy suspension. A system aimed at automating the consolidation process is being developed, that extends an existing one in dealing with further sorts of modifications (in its first release we only accounted for integration, substitution and repeal). The system is designed to extract modificatory provisions from a large database consisting of about 29,000 normative documents. The system is grounded on a description of modifications paired with a full-fledged syntactic annotation of such modifications.

In this paper we described a methodology for approaching legal texts analysis, with special focus on temporal modifications. We showed how the adoption of the FrameNet approach allows to use a wealth of information about legal language phenomena, that span over different layers, such as the legal one, the grammatical one and the syntactic one.

In our view, the proposed approach benefits from a declarative description of modifications. Decoupling declarative knowledge from procedural components of the system is helpful in separating legal knowledge from its use, which is not only more convenient on a software engineering perspective, but is also helpful in extending the system coverage. Further, from preliminary tests, we are confident to be able to improve the system accuracy, that over simpler modifications (substitution, integration and repeal) is around 70% recall and over 80% accuracy. The results of the first experiments of the system seem to corroborate the approach undertaken; however an extensive experimentation is necessary to assess the approach.

Future works will involve investigating the related –though different– modification of exceptions in its connections to suspensions, in order to yield a broader coverage of the modifications handled and a deeper comprehension of legal and linguistic phenomena.

Acknowledgments. The authors would like to thank the ICT4Law project (<http://www.ict4law.org/>), funded by Piedmont Region, for partly supporting current research.

References

1. Baker, C.F., Fillmore, C.J., Lowe, J.B.: The Berkeley FrameNet Project. In: Proceedings of the 17th International Conference on Computational Linguistics, vol. 1, pp. 86–90. Association for Computational Linguistics, Montreal (1998), <http://portal.acm.org/citation.cfm?id=980860>

2. Biagioli, C., Francesconi, E., Spinosa, P., Taddei, M.: The NIR project: Standards and tools for legislative drafting and legal document web publication. In: *Proceedings of ICAIL Workshop on e-Government:Modelling Norms and Concepts as Key Issues*, pp. 69–78 (2003)
3. Brighi, R., Lesmo, L., Mazzei, A., Palmirani, M., Radicioni, D.P.: Towards Semantic Interpretation of Legal Modifications through Deep Syntactic Analysis. In: *Jurix 2008: The 21st Annual Conference. Frontiers in Artificial Intelligence and Applications*, vol. 189, pp. 202–206. IOS Press (2008),
<http://www.di.unito.it/~radicion/papers/brighi08towards.pdf>
4. Cherubini, M., Giardiello, G., Marchi, S., Montemagni, S., Spinosa, P., Venturi, G.: NLP-based metadata annotation of textual amendments. In: *Proc. of Workshop On Legislative Xml, JURIX 2008* (2008)
5. Fillmore, C.J.: Scenes-and-frames semantics. In: Zampolli, A. (ed.) *Linguistic Structures Processing*, pp. 55–79. North-Holland, Amsterdam (1977)
6. Lesmo, L.: The Rule-Based Parser of the NLP Group of the University of Torino. *Intelligenza Artificiale* 2(4), 46–47 (2007)
7. Lesmo, L., Mazzei, A., Radicioni, D.P.: Semantic Annotation of Legal Modificatory Provisions. In: Serra, R., Cucchiara, R. (eds.) *AI*IA 2009. LNCS*, vol. 5883, pp. 304–313. Springer, Heidelberg (2009),
<http://dblp.uni-trier.de/db/conf/aiia/aiia2009.html#LesmoMR09>
8. Lupo, C., Vitali, F., Francesconi, E., Palmirani, M., Winkels, R., de Maat, E., Boer, A., Mascellani, P.: General XML format(s) for legal Sources - ESTRELLA European Project. Deliverable 3.1, Faculty of Law, University of Amsterdam, Amsterdam, The Netherlands (2007), <http://www.estrellaproject.org>
9. de Maat, E., Winkels, R.: Automatic Classification of Sentences in Dutch Laws. In: *Legal Knowledge and Information Systems, Jurix 2008: The 21st Annual Conference. Frontiers in Artificial Intelligence and Applications*, pp. 207–216. IOS Press (December 2008)
10. Mazzei, A., Radicioni, D., Brighi, R.: NLP-based Extraction of Modificatory Provisions Semantics. In: *Proceedings of the International Conference on Artificial Intelligence and Law, ICAIL 2009*, pp. 50–57. ACM, Barcelona (2009)
11. McCarty, L.T.: Deep semantic interpretations of legal texts. In: *ICAIL 2007: Proceedings of the 11th International Conference on Artificial Intelligence and Law*, pp. 217–224. ACM, New York (2007)
12. Palmirani, M., Brighi, R.: An XML Editor for Legal Information Management. In: Traunmüller, R. (ed.) *EGOV 2003. LNCS*, vol. 2739, pp. 421–429. Springer, Heidelberg (2003)
13. Palmirani, M., Brighi, R.: Model Regularity of Legal Language in Active Modifications. In: Casanovas, P., Pagallo, U., Sartor, G., Ajani, G. (eds.) *AICOL-II/JURIX 2009. LNCS*, vol. 6237, pp. 54–73. Springer, Heidelberg (2010)
14. Palmirani, M., Brighi, R., Massini, M.: Automated Extraction of Normative References in Legal Texts. In: *ICAIL 2003*, pp. 105–106 (2003),
<http://dblp.uni-trier.de/db/conf/icail/icail2003.html#PalmiraniBM03>
15. Soria, C., Bartolini, R., Lenci, A., Montemagni, S., Pirrelli, V.: Automatic Extraction of Semantics in Law Documents. In: Biagioli, C., Francesconi, E., Sartor, G. (eds.) *Proceedings of the V Legislative XML Workshop*, pp. 253–266. European Press Academic Publishing (2007)