

Automatic Tagging of Modality: identifying triggers and modal values

Paulo Quaresma^{1,4}, Amália Mendes², Iris Hendrickx^{2,3}, Teresa Gonçalves¹

¹ Department of Informatics, University of Évora, Portugal

² Center for Linguistics at the University of Lisbon, Portugal

³ Center for Language Studies, Radboud University Nijmegen, The Netherlands

⁴ L2F – Spoken Language Systems Laboratory, INESC-ID, Portugal

Abstract

We present an experiment in the automatic tagging of modality in Portuguese. As we are currently lacking a suitable resource with detailed modal information for Portuguese, we experiment with small sample of 160.000 tokens, manually annotated according to the modality scheme that we previously developed for European Portuguese (Hendrickx et al., 2012). We consider modality as the expression of the speaker (or subject)’s attitude towards the proposition and our modality scheme accounts for seven major modal values, and nine sub values. This experiment focuses on three modal verbs, *poder* ‘may/can’, *dever* ‘shall/might’ and *conseguir* ‘manage to/ succeed in/ be able to’, which may all have more than one modal value. We first report on the task of correctly detecting the modal uses of *poder* and *dever*, since these two verbs may have non modal meanings. For the identification of the modal value of each occurrence of those three verbs, we applied a machine learning approach that takes into consideration all the features available from a syntactic parser’s output. We obtained the best performance using SVM with a string kernel and the system improved the baseline for all three verbs, with a maximum F-score of 76.2.

Keywords: modality, annotation scheme, automatic tagging

1. Introduction

As the vast amount of digitally available data keeps growing, so does the demand to automatically extract relevant information. A clear problem for automatic extraction tools is to recognize the factual or non-factual nature of events, and the subjective perspective underlying the texts. In this paper we focus on modality: an important indicator of subjectivity and factuality in text. Modality is usually defined as the expression of the speaker’s opinion and of his attitude towards the proposition (Palmer, 1986). It traditionally covers epistemic modality, which is related to the degree of commitment of the speaker to the truth of the proposition (whether the event is perceived as possible, probable or certain), but also deontic modality (obligation or permission), capacity and volition. Modality detection is therefore also clearly linked to the current trend in NLP on sentiment analysis and opinion mining.

This paper presents an experiment in the automatic tagging of modality in Portuguese. Not much related work has been done in this area, certainly not for languages other than English. A prerequisite for building an automatic modality tagger is to have a corpus with labeled examples to train and evaluate such tool. As we are currently lacking a large and suitable corpus, one of the main aims of the study presented here is to create a tagger on a small corpus sample in order to (semi) automatically tag a larger corpus with modality information. For this purpose, we use a corpus of 158.553 tokens, manually annotated with a modality scheme for Portuguese (Hendrickx et al., 2012b). In this paper, we restrict our experiment to three modal verbs: *poder* ‘may/can’, *dever* ‘shall/might’ and *conseguir* ‘manage to/ succeed in/ be able to’. These three verbs are high frequent words in Portuguese and have different modal meanings, what makes them an excellent study object for our experiments.

The automatic modality tagger that we devised has two objectives: the identification of modal verbs (which we call the modal trigger) and the attribution of a modal value to this trigger. All three verbs have two or more modal meanings: for example, *poder* may be Epistemic, stating that something is possible, as in example (1); Deontic, denoting a permission, as in (2); or it may express an Internal capacity, the fact that someone is able to do something, as in (3). And frequently, a single context may be ambiguous between one and more of these readings.

- (1) E é evidente que um jogador que arrisque **pode** vir a ser apanhado mas, sem a certeza do controlo, a minha opinião é de que vai ter tendência para arriscar mais.

‘It is obvious that a player that takes risks might be caught but, without the certainty that there will be a control, in my opinion he will tend to take more risks.’

- (2) Segundo Cândida Almeida, "os jornalistas não **podem** usar meios que a própria lei veda a polícias e magistrados em nome dos direitos, liberdades e garantias dos cidadãos".

‘According to Cândida Almeida, “the journalists can not use means that the law itself forbids to the police and to prosecutors in the name of the citizen’s rights, liberties and warranties.

- (3) Os deputados portugueses, para serem ouvidos e terem influência, precisam de **poder** comunicar facilmente com os seus colegas, o que implica, num ambiente genuinamente multilinguístico, o domínio de várias línguas estrangeiras (...).

‘The Portuguese representatives to the European Parliament, to be heard and to have influence, need to be able to communicate easily with their

colleagues, what implies, in a genuinely multilingual environment, the mastery of several foreign languages.’

This polysemy increases the level of difficulty of the automatic annotation task. To create the modality tagger, we first automatically assign POS and syntactic tags, we then automatically identify modal triggers and apply a machine learning approach to attribute a modal value to the triggers, comparing the results with our gold dataset of 158.553 tokens.

The paper is structured as follows: we first revise related work in section 2, before briefly presenting our modality scheme and golden dataset in 3. Our automatic annotation system is described in section 4, the results of trigger identification are presented in 5.1 and the results of automatic attribution of modal value in 5.2, followed by a conclusion in 6.

2. Related work

Several annotation schemes of modality have been proposed in recent years, such as Baker et al. (2010), Matsuyoshi et al. (2010); Saurí et al. (2006), Nirenburg and McShane (2008) and, for Brazilian Portuguese, Ávila and Melo (2012). We will not discuss here in detail the differences between those annotation schemes (see Hendrickx et al. (2012b) and Nissim et al. (2013)) but rather focus on some experiments in the automatic annotation of modality that have been reported, mainly for English. Baker et al. (2010) tested two rule-based modality taggers to identify the modal trigger and its target and report results of 86% precision for tagging of a standard LDC data set. Also, Saurí et al. (2006) report on the automatic identification of events in text, and their characterization with modality features, achieving accuracy values of 97.04 with the EviTA tool. Battistelli and Damiani (2012) aim to annotate textual segments that have enunciative and modal (E_M) features. They use semantic clues to identify modal triggers and a syntactic parser to calculate the length of the E_M segment. However, the implementation of the system is an upcoming work. A specific system for the annotation of belief is reported by Diab et al. (2009). The authors mention that they treat all auxiliary verbs as epistemic, although they are aware of the fact that they may be deontic, and consider that this might be a source of noise in their system (an aspect that we also have to deal with). An extension of this experiment is reported in Prabhakaran et al. (2012), testing the tagging of different modality values (Ability, Effort, Intention, Success and Want). The authors report experiments on MTurk annotations (using only those examples for which at least two Turkers agreed on the modality and the target of the modality) and on a gold dataset, with respectively an overall 79.1 and 41.9 F-measure. It is important to mention that the corpora for both experiments differ greatly: MTurk data is entirely from email threads, whereas Gold data contains sentences from newswire, letters and blogs in addition to emails.

The work of Ruppenhofer and Rehbein (2012) is close to our own objectives in this paper. The authors report an experience to automatically identify five English modal verbs (*can/could*, *may/might*, *must*, *ought*, *shall/should*) in texts and predict their modal value, by training a maximum entropy classifier on features extracted from the training set. The authors manage to improve the baseline for all verbs but *must*, and achieve accuracy numbers between 68.7 and 93.5.

The detection of uncertainty and its linguistic scope was the subject of a shared task at CoNLL2010 (Farkas et al., 2010) focusing on hedging clues, which includes a broader set of lexical and syntactic clues than modality as we contemplate it in this paper. The area of BioNLP includes modality and factuality in the annotation of events: the dimension “level of certainty” is part of the system of meta-knowledge assignment to pre-recognised events described in Miwa et al. (2012), which attains F-measures of 74,9 for “low confidence” and 66,5 for “high but not complete confidence”.

3. Annotation Scheme and Corpus

The annotation scheme for Portuguese presented in Hendrickx et al. (2012a) is not restricted to modal verbs and also covers nouns, adjectives and adverbs. Modality is understood as the expression of the speaker’s attitude towards the proposition. So, the concept of factuality is not included, contrary to approaches such as Nissim et al. (2013), who accounts for both values but in different layers of the annotation scheme. Furthermore, our annotation scheme does not account for verb tense and mood, although this category is related to modality. The approach is very similar to the OntoSem (Mcshane et al., 2005) annotation scheme for modality (Nirenburg and McShane, 2008).

We include several modal values, based on the modality literature, but also on studies focused on annotation and information extraction (e.g. Palmer (1986); van der Auwera and Plungian (1998); Baker et al. (2010)). Seven main modal values are considered (Epistemic, Deontic, Participant-internal, Volition, Evaluation, Effort and Success), and several sub-values. There are five sub-values for epistemic modality: Knowledge, Belief, Doubt, Possibility and Interrogative. Contexts traditionally considered of the modal type “evidentials” (i.e., supported by evidence) are annotated as Epistemic belief. Two sub values are identified for deontic modality: Deontic obligation and Deontic permission (this includes what is sometimes considered Participant-external modality, as in van der Auwera and Plungian (1998)). Participant-internal modality is subdivided into Necessity and Capacity. Four other values are included: Evaluation, Volition and, following Baker et al. (2010), Effort and Success. We present the list of values and sub values in Table 1, together with their frequency in our golden set.

Nunca *me esqueço da ironia arrasadora de Churchill*, que *defendia que o político devia ser capaz de prever o que se vai passar amanhã*, no próximo mês e no próximo ano e de explicar depois por que é que aquilo que previu não aconteceu.

Figure 1: Screenshot of MMAX2 annotation tool

Main modal values	Sub values	Freq	%
Epistemic			
	knowledge	183	7,1
	belief	161	6,3
	doubt	29	1,1
	possibility	279	10,9
	interrogative	87	3,4
Deontic			
	obligation	581	22,7
	permission	159	6,2
Participant-internal			
	capacity	126	4,9
	necessity	122	4,8
Evaluation		159	6,2
Volition		396	15,4
Effort		110	4,3
Success		119	4,6

Table 1: Modal values and frequencies in our golden set

The annotation scheme comprises several components: (a) the trigger, which is the lexical element conveying the modal value; (b) the target; (c) the source of the event mention (speaker or writer) and (d) the source of the modality (agent or experiencer). The trigger receives an attribute *modal value*, while both trigger and target are marked for polarity. An example with the verb *dever* is given in (4)¹. In fact, the example sentence in (4) contains three other triggers as well. In this particular context, the trigger *esqueço* ‘I forget’ expresses the modal value Epistemic knowledge, the trigger *defendia* ‘argued’ expresses Epistemic belief, and the trigger *capaz* ‘be able’ expresses Participant-internal capacity. In example (4) however we focus on the annotation of the trigger *dever* in more detail.

- (1) Nunca me esqueço da ironia arrasadora de Churchill, que defendia que o político devia ser capaz de prever o que se vai passar amanhã, no próximo mês e no próximo ano e de explicar depois por que é que aquilo que previu não aconteceu.

‘I never forget the devastating irony of Churchill, who argued that a politician should be able of predicting what is going to happen tomorrow, next

¹ Notice that the discontinuity of the target is marked with the symbol @ in the example, but is encoded in XML in our data set.

month and next year and then explain why what he had predicted didn’t happen.’

Trigger: *devia*

Modal value: *deontic_obligation*

Polarity: positive

Target: o político@ ser capaz de prever o que se vai passar amanhã, no próximo mês e no próximo ano e de explicar depois por que é que aquilo que previu não aconteceu

Source of the modality: Churchill

Source of the event: writer

Ambiguity: none

This annotation scheme was applied to a corpus sample extracted from the written subpart of the Reference Corpus of Contemporary Portuguese (CRPC) (Généreux et al, 2012). Details about the selection of the sample are provided in Hendrickx et al (2012b). We used the MMAX2 annotation software tool (Müller and Strube, 2006) for our manual annotation task. The MMAX2 software is platform-independent, written in java and can freely be downloaded from <http://mmax2.sourceforge.net/>. The elements of our annotation consist of markables that are linked to the same modal event, which we call a "set". We present a screenshot of the results in Figure 1. The trigger *devia* and related markables are connected under a single set and are highlighted.

Full details on our annotation scheme and on the results of an inter-annotator experiment are provided in Hendrickx et al. (2012b). An enriched version with the interaction between Focus and Modality, specifically the case of exclusive adverbs, is presented in Mendes et al. (2013).

In the experiments that we present here, we focus on the Trigger component and its attribute *modal value*, and specifically on three semi-auxiliary modal verbs. The frequency of the modal verbs in our data set and their values are presented in Table 2.

The verb *dever* has two modal values in our golden set: Deontic obligation and Epistemic possibility. The value Participant-internal capacity is also possible with this verb but was never selected in our data as the primary meaning, although manual annotators have marked it in the ‘Ambiguity’ field of our annotation system in several cases. For this experiment, we didn’t take into consideration cases marked as ambiguous but this is certainly an important aspect to tackle in future research. Our experiments will therefore focus on five modal values: Deontic obligation, Deontic permission, Epistemic possibility, Participant-internal capacity and Success.

Main values	Sub values	Freq.
dever		113
	Deontic obligation	74
	Epistemic possibility	39
poder		244
	Deontic permission	43
	Epistemic possibility	158
	Participant-internal capacity	44
conseguir		84
	Participant-internal capacity	41
	Success	43

Table 2: Frequency of *dever*, *poder* and *conseguir* in our gold dataset.

4. Modality tagging

Our automatic modality tagger is composed by three modules:

- Syntactic analysis of the corpus;
- Identification of the modal verbs *poder*, *dever*, *conseguir*;
- Labeling of each verb with the appropriate modal value in its specific context.

The syntactic analysis was performed by the PALAVRAS parser (Bick, 1999), and the results were transformed into XML and logical terms (Prolog format) using the tool Xtractor (Gasperin et al., 2003). We then selected the set of parsed sentences that included the modal verbs and distinguished the modal uses of the verbs from the non-modal ones. As we aim to use this tagger to create a larger corpus, this first step of finding the modal triggers needs to be performed with very high accuracy.

We then used SVM, Support Vector Machines (Vapnik, 1998), to classify the modal value of each verb. We evaluated several machine learning algorithms and SVM kernel types with Weka (Hall et al., 2009), and obtained the best performance using SVM with a string kernel (Lodhi et al., 2002). We report the results obtained in two experiments: one using just the original sentences and another using the POS tags and functional and syntactic information extracted from the sentence’s parse tree, in a window of 70 characters around the verb. For the evaluation we used a 10-fold stratified cross-validation procedure. Note that this is a challenging task as we only have a few hundred examples to train and test the automatic tagger. We analyze the results in the next section.

5. Results

5.1 Modal verb detection

Here we first discuss to what extent we were able to correctly detect the modal verbs based on the output of the automatic syntactic parser. The verbs *poder* and *dever* may occur with non-modal uses, therefore the task

involves the correct identification of contexts that are indeed modal. The case of the verb *conseguir* is different because it always involves one of the modal values contemplated in our annotation system. For this specific verb, the system has to correctly identify sentences containing the lemma in the results of the parser, a much simpler task. Taking this into consideration, we will only discuss the results obtained for the verbs *poder* and *dever*, and compare our system’s output with the manually tagged information. This is summarized in Table 3.

	<i>poder</i>	<i>dever</i>
total verb occurrences	258	120
modal occurrences	244	113
automatic identification	236	108
false positives	0	0
error rate	3.1	4.2
precision	100	100
recall	96.7	95.6
F-measure	98.3	97.7

Table 3: Results of modal verb detection

Data from Table 3 show that the error rate in the identification of the modal occurrences is quite low: 3.1 for *poder* and 4.2 for *dever*. Precision receives the maximum value and Recall is above 95 for the two verbs. Errors are due to complex Portuguese sentences causing parsing problems, especially contexts where the semi-auxiliary modal verbs and the main verb are distant in the sentence. Another difficulty of the parser is to deal with cases where the semi-auxiliary modal is followed by a pronominal clitic. These issues could be partially dealt with in an additional post-processing step and would possibly result in an improvement of our performance in the future. However, syntactic complexity will remain a difficult challenge for semi-auxiliary detection.

5.2 Attribution of modal value

To identify the modal value, we applied a machine learning approach to the sentences detected by the previous module. Our system takes into consideration all the features available from the PALAVRAS output: lemma and POS of the trigger, left and right syntactic context, and semantic features: predicate argument structure, [\pm human] nature of arguments. We also computed scores for a baseline system that always assigns the most frequent modal value for each verb.

The results for both experiments (using the sentences and a text linearized format of the parse tree within a window around the verb) are presented in Table 4 (for *dever*), Table 5 (for *poder*) and Table 6 (for *conseguir*). We give results for a baseline and for both experiments (sentences and window parse tree), computing Precision (P), Recall (R) and F-value (F) and the macro-average over the different modal values.

dever	count	baseline			sentences			window parse tree		
		P	R	F	P	R	F	P	R	F
Total/macro-average	108	32.9	50.0	39.7	65.6	63.8	64.3	65.7	64.5	64.9
deontic obligation	71	65.7	100	79.3	74.4	81.7	77.9	75.0	80.3	77.6
epistemic possibility	37	0	0	0	56.7	45.9	50.7	56.3	48.6	52.2

Table 4: Results of the automatic modal value attribution for *dever*

poder	count	baseline			sentences			window parse tree		
		P	R	F	P	R	F	P	R	F
total/macro-average	236	21.8	33.3	26.3	34.6	33.4	32.2	34.3	34.0	33.7
deontic permission	42	0	0	0	23.1	7.1	10.9	18.8	14.3	16.2
epistemic possibility	154	65.3	100	79.0	64.6	80.5	71.7	65.5	75.3	70.1
participant internal capacity	40	0	0	0	16.1	12.5	14.1	18.5	12.5	14.9

Table 5: Results of the automatic modal value attribution for *poder*

conseguir	count	baseline			sentences			window parse tree		
		P	R	F	P	R	F	P	R	F
total/macro-average	84	25.6	50.0	33.9	57.1	57.0	56.8	76.3	0,762	76.2
participant internal capacity	41	0	0	0	57.1	48.8	52.6	76.9	73.2	75.0
success	43	51.2	100	67.7	57.1	65.1	60.9	75.6	79.1	77.3

Table 6: Results of the automatic modal value attribution for *conseguir*

The results in Tables 4-6 show that our system was able to improve the baseline for all three verbs: for *dever* it improves the baseline from 39.7 to 64.7 macro-average F-value, for *poder* from 26.3 to 33.7 and for *conseguir* from 33.9 to 76.2. The higher values attained for *conseguir* are tied to the fact that its two modal values have similar frequencies in our gold dataset, making it easier to improve the baseline.

With these experiments we obtained macro-average F-values between 33.7 and 76.2. We obtain better performance measures for *conseguir* and *dever* than for *poder*, possibly because *poder* has three modal values. Obviously, the automatic tagger obtains the best results for the most frequent values.

Comparing the experiments using the sentences and the window parse tree, the results show no significant differences, although the window parse tree experiment generally presents higher results, especially with *conseguir* (F-value 76.2 vs. 56.5).

6. Conclusion

We have presented a system for the automatic tagging of modality in Portuguese, using a manually annotated corpus as training data. The identification of the modal instances of the three auxiliary verbs receives high recall and precision values and could be further improved at the parsing level. The results of the attribution of the modal

value reach macro-average F-measures between 33 and 76 % F-value depending on the modal verb and on the modal value. The results are promising, considering that we trained our system on a tiny data set, and suggest that our aim: creating a larger corpus with modal information by a (semi) automatic tagging process based on a small sample seems to be a feasible next step.

In future work we plan to provide a detailed study identifying the individual role of the syntactic and semantic features that play a role in the automatic attribution of the modal value in our system. Another goal is to apply the modality tagger to a larger set of verbs to see whether we can keep a reasonable performance for a more diverse set of verbal triggers. We also aim to compute a learning curve to estimate the amount of manually annotated examples that are needed to get a good performance from the modality tagger.

As we are currently applying a 'word expert' approach and training separate classifiers for different verbal triggers, it is clear that this approach will not be able to handle modal triggers that it has not seen before. As a next step we will study this problem and for example try to train a general modal trigger classifier that is not dependent on the verb itself.

Acknowledgements

This work was partially supported by national funds

through FCT – Fundação para a Ciência e Tecnologia, under project Pest-OE/EEI/LA0021/2013 and project PEst-OE/LIN/UI0214/2013. We would like to thank the anonymous reviewers for their helpful comments and suggestions.

References

- Luciana Ávila and Heliana Melo (2013) Challenges in modality annotation in a Brazilian Portuguese Spontaneous Speech Corpus, *Proceedings of IWCS 2013 WAMM Workshop on the Annotation of Modal Meaning in Natural Language*, March 19-20, 2013, Postdam, Germany.
- Kathrin Baker, Michael Bloodgood, Bonnie Dorr, Nathaniel W. Filardo, Lori Levin, and Christine Piatko. 2010. A modality lexicon and its use in automatic tagging. In *Proceedings of LREC'10*, Valletta, Malta. ELRA, 1402-1407.
- Delphine Battistelli and Marine Damiani. 2013. Analyzing modal and enunciative discursive heterogeneity: how to combine semantic resources and a syntactic parser analysis. In *Proceedings of WAMM-IWCS2013*, Potsdam, Germany.
- Eckhard Bick. 1999. The parsing system PALAVRAS. *Aarhus University Press*.
- Mona Diab, Lori Levin, Teruko Mitamura, Owen Rambow, Vinodkumar Prabhakaran, and Weiwei Guo. 2009. Committed belief annotation and tagging. In *Proceedings of the Third Linguistic Annotation Workshop*, Suntec, Singapore, August. Association for Computational Linguistics, 68–73.
- Richárd Farkas, Veronika Vincze, György Móra, János Csirik, and György Szarvas. 2010. The CoNLL-2010 shared task: Learning to detect hedges and their scope in natural language text. In *Proceedings of the Fourteenth Conference on Computational Natural Language Learning*, ACL, 1-12.
- Caroline Gasperin, Renata Vieira, Rodrigo Goulart, and Paulo Quaresma. 2003. Extracting XML syntactic chunks from portuguese corpora. In *TALN'2003 - Workshop on Natural Language Processing of Minority Languages and Small Languages of the Conference on "Traitement Automatique des Langues Naturelles"*, Batz-sur-Mer, France, June 2003.
- Michel Génèreux, Iris Hendrickx, and Amália Mendes. 2012. Introducing the Reference Corpus of Contemporary Portuguese On-Line". In *Proceedings of the Eighth International Conference on Language Resources and Evaluation - LREC 2012*, Istanbul, May 21-27, 2012, 2237-2244.
- Mark Hall, Eibe Frank, Geoffrey Holmes, Bernhard Pfahringer, Peter Reutemann, and Ian H. Witten. 2009. The WEKA Data Mining Software: An Update; *SIGKDD Explorations*, Volume 11, Issue 1.
- Iris Hendrickx, Amália Mendes, Silvia Mencarelli, and Agostinho Salgueiro. 2012a. *Modality Annotation Manual*, version 1.0. Centro de Linguística da Universidade de Lisboa, Lisboa, Portugal.
- Iris Hendrickx, Amália Mendes, and Silvia Mencarelli. 2012b. Modality in Text: a Proposal for Corpus Annotation. In *Proceedings of the Eighth International Conference on Language Resources and Evaluation - LREC 2012*, Istanbul, May 21-27, 2012, 1805-1812.
- Huma Lodhi, Craig Saunders, John Shawe-Taylor, Nello Cristianini, and Christopher J. C. H. Watkins. 2002. Text Classification using String Kernels. *Journal of Machine Learning Research*. 2:419-444.
- Suguru Matsuyoshi, Megumi Eguchi, Chitose Sao, Koji Murakami, Kentaro Inui, and Yuji Matsumoto. 2010. Annotating event mentions in text with modality, focus, and source information. In *Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC'10)*, Valletta, Malta. ELRA.
- Marjorie McShane, Sergei Nirenburg, Stephen Beale, and Thomas O'Hara. 2005. Semantically rich human-aided machine annotation. In *Proceedings of the Workshop on Frontiers in Corpus Annotations II: Pie in the Sky*. ACL, 68-75.
- Amália Mendes, Iris Hendrickx, Agostinho Salgueiro, and Luciana Ávila. 2013. Annotating the Interaction between Modality and Focus: the case of exclusive particles. In *Proceedings of the 7th Linguistic Annotation Workshop & Interoperability with Discourse (LAW VII)*. Association for Computational Linguistics, Sofia, Bulgaria, August 8-9 2013, 228-237.
- Makoto Miwa, Paul Thompson, John McNaught, Douglas B Kell and Sophia Ananiadou. 2012. Extracting semantically enriched events from biomedical literature. *BMC Bioinformatics* 13:108.
- Christoph Müller and Michael Strube. 2006. Multi-level annotation of linguistic data with MMAX2. In *Corpus Technology and Language Pedagogy: New Resources, New Tools, New Methods*, 197-214. Peter Lang.
- Malvina Nissim, Paola Pietrandrea, Andrea Sansò and Caterina Mauri. 2013. Cross-linguistic annotation of modality: a data-driven hierarchical model. In Harry Bunt (ed.) *Proceedings of the 9th Joint ISO - ACL SIGSEM Workshop on Interoperable Semantic Annotation (isa-9)*, March 19-20, 2013, Postdam, Germany, 7-14.
- Sergei Nirenburg and Marjorie McShane. 2008. Annotating modality. Technical report, University of Maryland, Baltimore County, March 19, 2008.
- Frank R. Palmer. 1986. *Mood and Modality*. Cambridge textbooks in linguistics. Cambridge University Press.
- Vinodkumar Prabhakaran, Michael Bloodgood, Mona Diab, Bonnie Dorr, Lori Levin, Christine D. Piatko, Owen Rambow, and Benjamin Van Durme. 2012. Statistical modality tagging from rule-based annotations and crowdsourcing. In *Proceedings of the Workshop on Extra-Propositional Aspects of Meaning in Computational Linguistics (ExProM '12)*. Association for Computational Linguistics, Stroudsburg, PA, USA, 57-64.
- Josef Ruppenhofer and Ines Rehbein. 2012. Yes we can!?! Annotating the senses of English modal verbs. In *Proceedings of the 8th International Conference on*

Language Resources and Evaluation (LREC),
May 24-26, 2012, Istanbul, Turkey, 1538-1545.

Roser Saurí, Marc Verhagen, and James Pustejovsky.
2006. Annotating and recognizing event modality in
text. In *Proceedings of the 19th International FLAIRS
Conference*.

Johan Van der Auwera and Vladimir A. Plungian. 1998.
Modality's semantic map. *Linguistic Typology*, 2(1):
79-124.

Vladimir Vapnik. 1998. *Statistical learning theory*. Wiley,
NY.