Tagging and Labelling Portuguese Modal Verbs

Paulo Quaresma[1,4], Amália Mendes[2], Iris Hendrickx[2,3], and Teresa Gonçalves[1]

[1] Department of Informatics, University of Évora, Portugal
[2] Center for Linguistics of 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 in this paper an experiment in automatically tagging a set of Portuguese modal verbs with modal information. Modality is the expression of the speaker's (or the subject's) attitude towards the content of the sentences and may be marked with lexical clues such as verbs, adverbs, adjectives, but also by mood and tense. Here we focus exclusively on 9 verbal clues that are frequent in Portuguese and that may have more than one modal meaning. We use as our gold data set a corpus of 160.000 tokens manually annotated, according to a modality annotation scheme for Portuguese. We apply a machine learning approach to predict the modal meaning of a verb in context. This modality tagger takes into consideration all the features available from the parsed data (pos, syntactic and semantic). The results show that the tagger improved the baseline for all verbs, and reached macro-average F-measures between 35 and 81% depending on the modal verb and on the modal value.

1 Introduction

Nowadays we observe a growing interest in text mining applications that can automatically detect opinions, facts and sentiments in texts. Many of the current opinion or sentiment mining applications use a crude division between negative, neutral and positive sentiments. Modality (defined as the speaker's attitude towards the proposition in the text [19]) that has been studied for many decades in Linguistics, offers a theoretical framework to make a rather fine-grained distinction between different attitudes. For example, the speaker expresses his or someone else's commitment, hope, belief or knowledge about a certain proposition. We believe that such more detailed tagging of modal information will be helpful to improve opinion mining applications. Furthermore, a systematically tagged corpus with modal information will also be a rich resource for linguists who are interested in modality and how this influences the meaning of a text.

Modality can be expressed by different grammatical categories like modal adjectives, adverbs, nouns and verbs which we will denote as *modal trigger* in the rest of this paper. Here we focus on modality tagging for the Portuguese language and we concentrate on high frequent modal verbs. For our experiments we deliberately selected the highly ambiguous verbs that can express multiple modal meanings like for example the Portuguese verb *poder* that can express

"to be possible", "to be able" or "to give permission". The modality tagger that we developed has two objectives: the identification of modal triggers and the attribution of a modal value to this trigger.

Constructing an automatic modality tagger requires a data set with labelled examples to train and evaluate the tagger. As we are currently in need of a suitable data set, one of the goals of the current experiments is to develop an automatic modality tagger on a small manually labelled sample that can later be applied (semi-automatically) to generate a larger data set.

In this paper, we restrict our experiment to 9 modal verbs taken from a small sample of 160K tokens, manually annotated with a modality scheme for Portuguese [11]. The selected verbs are ambiguous and have at least two modal meanings. This polysemy increases the level of difficulty of the automatic annotation task. To create the modality tagger, we first automatically assign lemmas, 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 (manually annotated) labelling.

Next we first discuss related work on modality annotation and tagging in section 2. In section 3 we describe the modality scheme and manually annotated dataset for Portuguese that we used in the experiments. The modality tagging approach is presented in section 4 and the results of automatic attribution of modal value in 5. We conclude in section 6.

2   Related work

The annotation schemes covering modality differ greatly in their objectives and in the nature of the concepts that are labelled. Modality may be one aspect of the semantic information encoded in the properties of events, such as Time and Condition (Matsuyoshi et al. [14], Baker et al. [2] and Nirenburg and McShane [17]). Modal values may also be included in annotation schemes that cover both factuality and modality, as in Sauri et al. [24], or that are concerned with subjectivity, beliefs and opinions. Contrary to many of these approaches, Rubinstein et al. [22] use a restricted notion of modality and establish conditions for an expression to be considered modal, such as the requirement for a propositional argument.

The options regarding what textual elements to annotate also differ: the modal value may be attributed globally to the sentence/event or it can be encoded on specific triggers (and other components such as target and source). The nature of the data over which to apply the annotation also varies: some work has been developed specifically on spoken data, as the scheme described for Brazilian Portuguese in Avila and Melo (2013) [26], while modality has enriched the semantic annotation of events in biomedical texts. For more contrastive information on the existing annotation schemes, see an overview in Nissim et al. [18].

Some of these modality schemes have led to experiments in the automatic annotation of modal values, mainly for English. For instance, in BioNLP Miwa

et al. [16] annotated pre-recognised events with the epistemic value "level of certainty" and attain F-measures of 74,9 for "low confidence" and 66,5 for "high but not complete confidence". Another system is described in Battistelli et al. [3] and plans to attribute enunciative and modal features (E_M) to textual segments. Modal triggers are to be identified through semantic clues and a syntactic parser will be used to calculate the length of each segment. The design of the system is well described but has not been fully implemented yet. The modal scheme presented in [24] focuses on factuality and has been applied to an experiment in the automatic identification of events in text, together with the characterization of their modality features. The authors report results with accuracy values of 97.04 with the EviTA tool. Also, a specific task for detecting uncertainty through the use of hedging clues was organized at CoNLL2010 [7].

Baker et al. [2] go a little further than our own experiment in this paper: they identify both the modal trigger and its target and report results of 86% precision with two rule-based modality taggers over a standard LDC data set. The approach of Diab et al [6] covers modality but is essentially geared towards the identification of belief. Contrary to our own approach, the authors do not take into consideration the polysemy of the auxiliary verbs and encode a single value (epistemic), although they do report that the verbs may be deontic in some contexts. This experiment has been extended to other modality values (Ability, Effort, Intention, Success and Want), as reported in Prabhakaran et al [20]. The authors applied their system to their Gold data (containing emails, sentences from newswire, letters and blogs) and on a corpus of emails annotated on MTurk. The two experiments attained quite different overall measures: 79.1 F-measure for the MTurk experiment and 41.9 on the Gold data, that the authors attribute to the difference in the corpora design and to the higher complexity of the experts annotations.

Our objective to identify and label Portuguese modal verbs with modal values is closely related to the work of Ruppenhofer and Rehbein [23] on English modals. The five English modal verbs (*can/could, may/might, must, ought, shall/should*) were first identified in texts and their modal value was predicted by training a maximum entropy classifier on features extracted from the training set. The experiment achieved an improvement of the baseline for all verbs but *must*, and accuracy numbers between 68.7 and 93.5.

The diversity of the concept of modality and of the textual scope of the annotation makes it difficult to fully contrast the results described in the experiments, since no standard has yet emerged in this domain.

3   Data set

Our experiment in the automatic annotation of modality is based on the annotation scheme for Portuguese presented in [12] and [11]. The annotation is based on the identification of lexical clues, called triggers, which include verbs, but also other POS classes, such as nouns, adjectives and adverbs. Each modal trigger is labelled with a modal value.

The scheme covers seven main modal values (epistemic, deontic, participant-internal, volition, evaluation, effort and success) and several sub-values. The first five modal values follow linguistic typologies ([19][1]), while the last two are influenced by typologies used in annotation schemes for modality (e.g. [2]). Epistemic modality is the more complex value (and the more stable across typologies) and is subdivided into five sub-values: knowledge, belief, doubt, possibility and interrogative. Evidential modality (i.e., belief supported by evidence) is not considered an independent value and is rather annotated as epistemic belief. Deontic modality is divided into obligation and permission and includes values that are described as participant-external modality in van der Auwera and Plungian [1] (obligation and possibility that are not dependent on the participants but rather on external conditions that make something required or possible). However, the scheme does include what these authors describe as Participant-internal modality and its two subvalues: necessity and capacity (an internal capacity or an internal necessity of the participant, usually the subject). Finally, four other values are included: evaluation (of the proposition), volition (desires, wishes and fears), effort and success.

The annotation scheme components comprise the trigger (which is the lexical clue conveying the modal value), its target, the source of the event mention (speaker or writer) and the source of the modality (for instance, the entity that considers the proposition to be possible, that establishes an obligation or that has an internal capacity to do something). The source of the event mention and the source of the modality are in many cases the same entity: the speaker/writer produces a discourse/text unit where it states its belief or doubt or the possibility that something may happen. However, they may also be different entities when the text presents the views of someone else than its producer. The trigger receives an attribute *modal value*, while both trigger and target are marked for polarity. Modal verbs may have more than one meaning and it is sometimes difficult to distinguish between those modal values, even when the annotator takes into consideration a larger context. To address this issue, the scheme includes the Ambiguity component, where the annotators can write down secondary meanings that would also be available in that specific context. This annotation scheme was further enriched with Focus information to address the interaction betwen exclusive adverbs and modal triggers [15], although we will not take these components into consideration in our current experiment.

The verbs that are the subject of our experiment are good examples of the polysemy of modal verbs. Two of them may also have non modal meanings, as it is the case with *dever* and *poder*. We give in (1) an example of a non modal use of *dever* and in (2) and (3) examples of modal uses of the verb. In (2), *dever* has an epistemic reading, stating that the proposition is probable, while in (3), it has a Deontic obligation reading (the adverb *obrigatoriamente* 'obligatorily' is also a modal trigger). In (3) we also provide the total description of the components of the sentence annotation. The obligation has scope over the whole proposition [devem [faltas sucessivas aos julgamentos ser contrariadas]] and there is no overt source of the modality (in these cases the verb itself is

marked as the source)[5]. Cases marked as ambiguous are illustrated in (4): the context is ambiguous between an Epistemic possibility and a Deontic obligation reading (it is probable vs it is required that the money be spent in such a way). The annotator selects what he considers to be the primary value and marks the ambiguity in the Ambiguity component.

(1) O Governo *deve* explicações muito mais claras e completas do que as que deu (...).
'The Government owes much more clear and complete explanations than the ones that it gave'

(2) Entre copos de vinho, muitos cigarros e piadas trocadas com os músicos, Chico Buarque de Hollanda ensaia o seu próximo "show", que *deve* estrear em Janeiro.
'Among wine, cigarettes and jokes with the musicians, Chico Buarque de Hollanda rehearses his next show, which should start in January.'

(3) Faltas sucessivas aos julgamentos contribuem para a morosidade da justiça e, obrigatoriamente, *devem* ser contrariadas.
'Repeated absence to trials contribute to the slowdown of justice and must obligatorily be opposed.'
Trigger: devem
    Modal value: deontic_obligation
    Polarity: positive
Trigger: Faltas sucessivas@ser contrariadas
Source of the modality: devem
Source of the event: writer
Ambiguity: none

(4) O seu orçamento global é de 12,5 milhões de euros, sendo que uma grande maquia deverá ser aplicada na construção de infraestruturas rodoviárias.
'Its global funding is 12,5 million euros, and a large part (might/has to) be applied in the construction of road infrastructures.'
Trigger: deverá
    Modal value: Epistemic_possibility
    Polarity: positive
Trigger: uma grande maquia@ser aplicada na construção de infraestruturas rodoviárias
Source of the modality: deverá
Source of the event: writer
Ambiguity: Deontic_obligation

This annotation scheme was applied to a corpus sample of 2000 sentences extracted from the written subpart of the Reference Corpus of Contemporary Portuguese (CRPC)[9]. The sentences were selected on the basis of a list of

---

[5] Notice that the discontinuity of the target is marked with the symbol @ in this example, but is encoded in XML in the data set.

potential modal verbs mentioned in the literature. In the annotation however, all modal triggers including nouns, adjectives and adverbs were annotated. In total the data set contains approximately 3200 modal triggers.

In this experiment however, we only evaluate on a set of highly ambiguous verbs that can have multiple modal values. Many of the verbs in the data set, such as high frequent modal verbs like *querer 'to want'* and *tentar 'to try'*, only have one modal value and for those assigning the correct modal value becomes a trivial task. We aim to only study the hard cases that involve true ambiguous verbs. We focus on modal verbs with multiple modal meanings that each occur at least 5 times in the small annotated corpus sample. Only a handful of verbs met this criteria, giving us a list of 9 verbs to work with. In Table 1 we show the 9 verbs and their distribution between the different main modal values. We see that all different modal values of the data set are covered but most verbs only have two different modal values resulting in a sparse matrix. Two of the verbs (*poder* and *dever*) are polysemous verbs that can be used as a semi-auxiliary verb where it has a modal meaning or as a main verb without a modal meaning.

The current experiment is a follow-up of our experience with tagging 3 Portuguese verbs (*poder*, *dever* and *conseguir*) with modality values, as reported in [21]. Here, we extend our work, so as to cover a larger list of verbs and provide a better testing set for our system.

Table 1. Data set used in these experiments. Abbreviations used: non(non-modal), EB (epistemic belief), EP (epistemic possibility), EF (effort), EV (evaluation), OB (deontic obligation), NE (participant-internal necessity), PC (participant-internal capacity), PE (deontic permission), SE (success) and VO (volition)

verb	English	total	non	EB	EP	EF	EV	OB	NE	PC	PE	SE	VO
arriscar	risk	44			25	19							
aspirar	aspire	50		19									31
considerar	consider	29		18			11						
conseguir	succeed	84								41		43	
dever	must/might	120	12		37			71					
esperar	hope	52		26									26
permitir	allow	78			60						18		
poder	can/may/be able	258	22		154					40	42		
precisar	need	54						9	45				
total		769	34	63	276	19	11	80	45	81	60	43	57

4 Modality tagging

The modality tagging is done in three different steps. First we preprocess the data set with a syntactic parser. Next, the parser output is used to detect the

modal triggers. And finally we label each modal trigger with the appropriate modal value in context.

We used the PALAVRAS parser [5] for the syntactic analysis, and rewrote the result to an XML format with logical terms using the tool Xtractor [8]. We use a list of the selected verbs described in section 3 to detect the modal triggers. Some of the verbs can be used as modal or non-modal like *dever* and *poder*. The parser labels these verbs as auxiliary verbs when they are used as modals and we exploit these parser predictions to detect the modal usage.

For the final step of labelling the modal triggers with their appropriate modal value we apply a machine learning approach. We adopted the "word expert" strategy that is often used in automatic word sense disambiguation: we train a specialised classifier for each verb. This strategy is ideal for classification tasks where every verb has a different type of meanings or modal values and a different distribution in usage [4].

After experimenting with several machine learning algorithms with Weka [10], we achieved the best results with SVM, Support Vector Machines [25] using the subsequence string kernel [13]. We performed two experiments with different feature representations (using default algorithm values). In the first one we only use the words in the original sentences and in the second experiment we used 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. Remark that this is a challenging task as we only have a few examples for every modal verb word expert to train and test. In the next section we present the results of our experiments.

5 Experiments: attribution of the modal value

In the context of this work, and as presented in table 1, we analysed 9 modal verbs and created automatic classifiers able to attribute a modal value to each occurrence of these verbs in sentences written in the Portuguese language.

Occurrences of these 9 verbs were automatically detected from the analysis of the output of the syntactical parser PALAVRAS [5], which associates a lemma to each word in a sentence. As verbs *poder* and *dever* may occur with non-modal uses, it was necessary, for these verbs, to distinguish the two kind of uses (modal and non-modal). We used the sentences' parse trees to identify the semi-auxiliary situations, which are modal. Based on the parsers' output we were able to detect the modal usage of *poder* and *dever* with and F-score of 98%.

To identify the modal value, we applied a machine learning approach to the sentences with occcurrences of the 9 modal verbs analysed in this work. 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, [±human] nature of arguments. The results for both experiments (using the word sentences and a text linearized format of the parse tree within a window around the verb) are presented in Table 2. 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. As baseline we used a system that always assigns the most frequent modal value for each verb.

As summary we can refer the following points:

– Verb *arriscar*: the classifier which uses as input the partial parse tree was able to improve the baseline F-measure from 36.2 to 63.0.
– Verb *aspirar*: the classifier which uses as input the partial parse tree was able to improve the baseline F-measure from 38.3 to 81.1.
– Verb *conseguir*: the classifier which uses as input the partial parse tree was able to improve the baseline F-measure from 33.9 to 76.2.
– Verb *considerar*: the classifier which uses as input the sentences was able to improve the baseline F-measure from 38.3 to 41.6.
– Verb *dever*: the classifier which uses as input the partial parse tree was able to improve the baseline F-measure from 39.7 to 63.8.
– Verb *esperar*: the classifier which uses as input the sentences was able to improve the baseline F-measure from 33.3 to 53.6.
– Verb *permitir*: the classifier which uses as input the partial parse tree was able to improve the baseline F-measure from 43.5 to 55.8.
– Verb *poder*: the classifier which uses as input the sentences was able to improve the baseline F-measure from 33.3 to 35.3.
– Verb *precisar*: the classifier which uses as input the sentences was able to improve the baseline F-measure from 45.7 to 62.5.

So, our system was able to clearly improve the F-measure baseline approach for all 9 modal verbs, with more than 40 points improvement for the verbs *aspirar* and *conseguir*, more than 20 points for the verb *arriscar*, and more than 10 points for the verbs *dever*, *esperar*, *permitir*, and *precisar*. The values for the verb *poder* are lower, probably because this verb has three modal meanings while the other verbs have two.

We also observed that for 5 modal verbs the partial parse tree input was able to produce a better classifier and for the other 4 verbs the list of words of the sentences were enough to obtain the best classifier.

We plan to follow up this analysis with a detailed study identifying the individual role of the syntactic and semantic features that are used for the automatic attribution of the modal value in our system. More specifically, for each parse tree we intend to evaluate the relevance of the following features in the modal value classification: word, lemma, POS tag, syntactic tag, semantic information, role label. The analysis will be performed over the partial parse trees that include the modal verbs and their parents and grand-parents and also over the nodes in the path from the root of the sentence parse trees to the modal verbs. We will also make a comparative study of the relevant features for each of the studied modal verbs. In fact, from the analysis of the already obtained results, we foresee the existence of great differences between the 9 modal verbs. The fact that, for some verbs, the parse tree input obtains better results and for others it is better to use the list of words of the sentences, suggests that syntactic and semantics features might not be equally relevant for all verbs.

Table 2. Results of all verbs

arriscar		baseline			sentences			tree		
		P	R	F	P	R	F	P	R	F
effort	19	0	0	0	47.4	47.4	47.4	57.9	57.9	57.9
possibility	25	56.8	100	72.5	60.0	60.0	60.0	68.0	68.0	68.0
average	44	28.4	50.0	36.2	53.7	53.7	53.7	63.0	63.0	**63.0**

aspirar		baseline			sentences			tree		
		P	R	F	P	R	F	P	R	F
volition	31	62.0	100	76.5	65.7	74.2	69.7	86.7	83.9	85.2
belief	19	0	0	0	46.7	36.8	41.2	75.0	78.9	76.9
average	50	31.0	50.0	38.3	56.2	55.5	55.5	80.9	81.4	**81.1**

conseguir		baseline			sentences			tree		
		P	R	F	P	R	F	P	R	F
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
average	84	25.6	50.0	33.9	57.1	57.0	56.8	76.3	76.2	**76.2**

considerar		baseline			sentences			tree		
		P	R	F	P	R	F	P	R	F
belief	18	62.1	100	76.6	60.0	83.3	69.8	55.0	61.1	57.9
evaluation	11	0	0	0	25.0	9.1	13.3	22.2	18.2	20.0
average	29	31.0	50.0	38.3	42.5	46.2	**41.6**	38.6	39.7	39.0

dever		baseline			sentences			tree		
		P	R	F	P	R	F	P	R	F
obligation	71	65.7	100	79.3	74.0	80.3	77.0	75.0	76.1	75.5
possibility	37	0	0	0	54.8	45.9	50.0	52.8	51.4	52.1
average	108	32.9	50.0	39.7	64.4	63.1	63.5	63.9	63.8	**63.8**

esperar		baseline			sentences			tree		
		P	R	F	P	R	F	P	R	F
belief	26	50.0	100	66.7	57.1	46.2	51.1	54.5	46.2	50.0
volition	26	0	0	0	54.8	65.4	59.6	53.3	61.5	57.1
average	52	25.0	50.0	33.3	56.0	55.8	**55.4**	53.9	53.9	53.6

permitir		baseline			sentences			tree		
		P	R	F	P	R	F	P	R	F
possibility	60	76.9	100	87.0	75.7	88.3	81.5	79.4	83.3	81.3
permission	18	0	0	0	12.5	5.6	7.7	33.3	27.8	30.3
average	78	38.5	50.0	43.5	44.1	47.0	44.6	56.4	55.6	**55.8**

poder		baseline			sentences			tree		
		P	R	F	P	R	F	P	R	F
permission	42	0	0	0	33.3	9.5	14.8	12.5	11.9	12.2
possibility	154	65.3	100	79.0	65.8	83.8	73.7	63.5	68.8	66.0
capacity	40	0	0	0	17.9	12.5	14.7	24.1	17.5	20.3
average	236	21.8	33.3	26.3	39.0	**35.3**	34.4	33.4	32.7	32.8

precisar		baseline			sentences			tree		
		P	R	F	P	R	F	P	R	F
necessity	45	84.0	100	91.3	86.3	97.8	91.7	83.3	88.9	86.0
obligation	9	0	0	0	66.7	22.2	33.3	16.7	11.1	13.3
average	54	42.0	50.0	45.7	76.5	60.0	**62.5**	50.0	50.0	49.7

6 Conclusion

We have presented our experiment to automatically annotate modality for 9 Portuguese modal verbs, using a corpus sample of 160K tokens, manually tagged with modal values. For this experiment, we selected verbs that had more than one modal meaning, which occurred at least 5 times in the corpus. We used SVM and performed two experiments: one uses the original sentences and another uses the information available from the parser's output. The results of the attribution of modal value improved the baseline for all verbs, with macro-average F-measures between 35.3 and 81.1%, depending on the modal verb and on the modal value. Considering that our training corpus was relatively small and that we selected challenging verbs in our experiment, we believe that our goal, of creating a larger corpus with modal information by a (semi) automatic tagging process, could lead to positive results in the future. We plan to study the role played by each feature in our system and to observe in more detail the reason why some verbs reach higher scores with the experiment that uses sentences instead of the parse tree. 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. Also, for our system to be able to label new verbs that didn't occur in the initial data set, we plan to train a general modal trigger classifier that is not dependent on the verb itself.

Acknowledgements

References

1. der Auwera, J.V., Plungian, V.A.: Modality's semantic map. Linguistic Typology 1(2), 79–124 (1998)
2. Baker, K., Bloodgood, M., Dorr, B., Filardo, N.W., Levin, L., Piatko, C.: A modality lexicon and its use in automatic tagging. In: Chair), N.C.C., Choukri, K., Maegaard, B., Mariani, J., Odijk, J., Piperidis, S., Rosner, M., Tapias, D. (eds.) Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC'10). European Language Resources Association (ELRA), Valletta, Malta (may 2010)
3. Battistelli, D., Damiani, M.: Analyzing modal and enunciative discursive heterogeneity: how to combine semantic resources and a syntactic parser analysis. In: Proceedings of the IWCS 2013 Workshop on Annotation of Modal Meanings in Natural Language (WAMM). pp. 7–15. Association for Computational Linguistics, Potsdam, Germany (March 2013)
4. Berleant, D.: Engineering "word experts" for word disambiguation. Natural Language Engineering 1(4), 339–362 (1995)
5. Bick, E.: The parsing system PALAVRAS. Aarhus University Press (1999)

6. Diab, M.T., Levin, L.S., Mitamura, T., Rambow, O., Prabhakaran, V., Guo, W.: Committed belief annotation and tagging. In: Third Linguistic Annotation Workshop. pp. 68–73. The Association for Computer Linguistics, Singapore (August 2009)

7. Farkas, R., Vincze, V., Móra, G., Csirik, J., Szarvas, G.: 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. pp. 1–12. Association for Computational Linguistics, Uppsala, Sweden (July 2010)

8. Gasperin, C., Vieira, R., Goulart, R., Quaresma, P.: Extracting xml syntactic chunks from portuguese corpora. Proc. of the TALN Workshop on Natural Language Processing of Minority Languages and Small Languages pp. 223–232 (2003)

9. Généreux, M., Hendrickx, I., , Mendes, A.: Introducing the reference corpus of contemporary portuguese on-line. In: Calzolari, N., Choukri, K., Declerck, T., Dogan, M.U., Maegaard, B., Mariani, J., Odijk, J., Piperidis, S. (eds.) LREC'2012 – Eighth International Conference on Language Resources and Evaluation. pp. 2237–2244. European Language Resources Association (ELRA), Istanbul, Turkey (May 2012)

10. Hall, M., Frank, E., Holmes, G., Pfahringer, B., Reutemann, P., Witten, I.H.: The weka data mining software: An update. SIGKDD Explor. Newsl. 11(1), 10–18 (Nov 2009)

11. Hendrickx, I., Mendes, A., Mencarelli, S.: Modality in text: a proposal for corpus annotation. In: Calzolari, N., Choukri, K., Declerck, T., Dogan, M.U., Maegaard, B., Mariani, J., Odijk, J., Piperidis, S. (eds.) LREC'2012 – Eighth International Conference on Language Resources and Evaluation. pp. 1805–1812. European Language Resources Association (ELRA), Istanbul, Turkey (May 2012)

12. Hendrickx, I., Mendes, A., Mencarelli, S., Salgueiro, A.: Modality Annotation Manual. Centro de Linguística da Universidade de Lisboa, Lisboa, Portugal, v1 edn.

13. Lodhi, H., Saunders, C., Shawe-Taylor, J., Cristianini, N., Watkins, C.: Text classification using string kernels. Journal of Machine Learning Research 2, 419–444 (Mar 2002)

14. Matsuyoshi, S., Eguchi, M., Sao, C., Murakami, K., Inui, K., Matsumoto, Y.: Annotating event mentions in text with modality, focus, and source information. In: Chair), N.C.C., Choukri, K., Maegaard, B., Mariani, J., Odijk, J., Piperidis, S., Rosner, M., Tapias, D. (eds.) Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC'10). European Language Resources Association (ELRA), Valletta, Malta (may 2010)

15. Mendes, A., Hendrickx, I., Salgueiro, A., Ávila, L.: Annotating the interaction between focus and modality: the case of exclusive particles. In: Proceedings of the 7th Linguistic Annotation Workshop and Interoperability with Discourse. pp. 228–237. Association for Computational Linguistics, Sofia, Bulgaria (August 2013)

16. Miwa, M., Thompson, P., McNaught, J., Kell, D.B., Ananiadou, S.: Extracting semantically enriched events from biomedical literature. BMC Bioinformatics 13, 108 (2012)

17. Nirenburg, S., McShane, M.: Annotating modality. Tech. rep., University of Maryland, Baltimore County, USA (March 2008)

18. Nissim, M., Pietrandrea, P., Sanso, A., Mauri, C.: Cross-linguistic annotation of modality: a data-driven hierarchical model. In: Proceedings of IWCS 2013 WAMM Workshop on the Annotation of Modal Meaning in Natural Language. pp. 7–14. Association for Computational Linguistics, Postam, Germany (2013)

19. Palmer, F.R.: Mood and Modality. Cambridge textbooks in linguistics, Cambridge University Press (1986)

20. Prabhakaran, V., Bloodgood, M., Diab, M., Dorr, B., Levin, L., Piatko, C.D., Rambow, O., Van Durme, B.: Statistical modality tagging from rule-based annotations and crowdsourcing. In: Proceedings of the Workshop on Extra-Propositional Aspects of Meaning in Computational Linguistics. pp. 57–64. ExProM '12, Association for Computational Linguistics, Stroudsburg, PA, USA (2012)
21. Quaresma, P., Mendes, A., Hendrickx, I., Gonçalves, T.: Automatic tagging of modality: identifying triggers and modal values. In: Proceedings of ISA-10 - 10th Joint ACL-ISO Workshop on Interoperable Semantic Annotation. European Language Resources Association (ELRA), Reykjavik, Iceland (May 2014)
22. Rubinstein, A., Harner, H., Krawczyk, E., Simonson, D., Katz, G., Portner, P.: Toward fine-grained annotation of modality in text. In: Proceedings of the Tenth International Conference for Computational Semantics (IWCS 2013) (2013)
23. Ruppenhofer, J., Rehbein, I.: Yes we can!? annotating english modal verbs. In: Chair), N.C.C., Choukri, K., Declerck, T., Doğan, M.U., Maegaard, B., Mariani, J., Odijk, J., Piperidis, S. (eds.) Proceedings of the Eight International Conference on Language Resources and Evaluation (LREC'12). European Language Resources Association (ELRA), Istanbul, Turkey (may 2012)
24. Sauri, R., Verhagen, M., Pustejovsky, J.: Annotating and recognizing event modality in text. In: FLAIRS Conference. pp. 333–339 (2006)
25. Vapnik, V.N.: Statistical Learning Theory. Wiley-Interscience (1998)
26. Ávila, L., Melo, H.: Challenges in modality annotation in a brazilian portuguese spontaneous speech corpus. In: Proceedings of IWCS 2013 WAMM Workshop on the Annotation of Modal Meaning in Natural Language. Association for Computational Linguistics, Postam, Germany (2013)