

Modality: logic, semantics, annotation and machine learning

ANNIE ZAENEN, STANFORD UNIVERSITY

Up to rather recently Natural Language Processing has not given much attention to modality. As long as the main task was to determine what a text was about (Information Retrieval) or who the participants in an eventuality were (Information Extraction), this neglect was understandable. With the focus moving to questions of natural language understanding and inferencing as well as to sentiment and opinion analysis, it becomes necessary to distinguish between actual and envisioned eventualities and to draw conclusions about the attitude of the writer or speaker towards the eventualities referred to. This means, i.a., to be able to distinguish ‘John went to Paris’ and ‘John wanted to go to Paris’. To do this one has to calculate the effect of different linguistic operators on the eventuality predication.¹

Modality has different shades of meaning that are subtle, and often difficult to distinguish, being able to express hypothetical situations (he could/may come in), desired or undesired (permitted or non-permitted situations (he can/may come in/enter), or (physical) abilities: he can enter. The study of modality often focusses on the semantics and pragmatics of the modal auxiliaries because of their notorious ambiguity but modality can also be expressed through other means than auxiliaries, such as adverbial modification and non-auxiliary verbs such as *want* or *believe*. In fact, the same modality can be expressed by different linguistic means, e.g. ‘Maybe he is already home’ or ‘He may already

¹It also means to be able to distinguish among non-propositional aspects of meaning (e.g. among ‘Jane is strong-minded’, ‘Jane is stubborn’ and ‘Jane is pigheaded’) but these are not the topic of this volume.

be home'. These paraphrases can provide a means to disambiguate the modals auxiliaries as is illustrated in one of the contributions to this volume (Modal Sense Classification at Large).

Huddleston and Pullum (2002) distinguish between three dimensions of modal expressions, the KIND, THE STRENGTH and the DEGREE. Linguistic research (Palmer (1976), Kratzer (1981), Kratzer (1991)) has focussed mainly on the many KINDS of modals. The most ambitious and influential proposal to date to structure the domain is that of Kratzer (1981,1991) who developed a possible world semantics. In this theory all KINDS of modals have one pole that corresponds to existential quantification (i.e. possibility) and one that corresponds to universal quantification (i.e. necessity). This force (STRENGTH) of the modal is lexically determined. The difference among the KINDS of modals arises from the different choices in sets of possible worlds that the quantification ranges over, the *modal base*. This aspect is determined by the context, the situation in which the modal is used. Moreover, an *ordering* on the possible worlds is assumed and this ordering too depends on the context.

The main KINDS of modal meanings that have been distinguished are *epistemic*, *deontic*, *dynamic/circumstantial*, *bouletic* and *teleological* modality. *Epistemic* modality is about what is necessary or possible given what is known; *deontic* is about what is necessary or possibly given what is moral or legal; *dynamic* or *circumstantial* modality about what is necessary or possible given a set of circumstances and *teleological* about what is necessary of possible to achieve a given goal. The *modal base* for the latter three categories can be seen as the same, it is circumstantial, but the textitordering relation is different, depending on legal and moral principles in the case of deontic modality, on desires in the case of bouletic modality and on goals in the case of teleological modality. In this approach conditional statements are closely related to modals. Kratzer's proposal is that if-clauses are modifiers of the modal base: they restrict the set of accessible worlds.

Differences in STRENGTH, as illustrated in the difference among the adverbs *necessary*, *probable*, *possible* (and their negations) can be modeled through the ordering on the set of worlds in the Kratzer model. Especially from a computational point of view, this is not desirable as discussed in 'Goal-Oriented Modality in Type Theory with Records (TTR)' below. Alternatives have been presented e.g. in Frank (1997) and in Lassinter (forthcoming). In most of the annotation proposals reported on in this volume either STRENGTH is not addressed or it is seen as a subdivision in the kinds of modals that are distinguished, e.g. distinctions between knowledge and belief, permission and prohibition.

The DEGREE dimension ‘has to do with the extent to which there is a clearly identifiable and separable element of modal meaning’ (Huddleston and Pullum, p.179). It is not addressed by any of the contributions here.

The volume addresses two other phenomena that fall outside of the taxonomy proposed in Huddleston and Pullum. One is that of modal subordination, first discussed in Roberts (1989) and illustrated in the contrast between (1) and (2):

(1) # Bill can make a kite. The kite has a long string (Karttunen (1969))

(2) A thief might brake into the house. He would take the silver. (Roberts (1989))

The first sentence shows that normally an indefinite introduced in a modal context cannot serve as a referential anchor in the subsequent discourse. (2) shows that it is possible to create a modal context in which these references are possible, namely, when there is a modal in the subsequent context that indicates that the non-factual situation extends to it.

The other phenomenon that we include might come as a surprise: *sarcasm*. We see this as a type of modality in the sense that it involves two worlds: one in which the literal proposition is evaluated and the actual one in which the intended proposition which is the negation of the literal one is supposed to hold.

In 2012 *Computational Linguistics* devoted an issue to Modality and Negation. This volume can be seen as an update. It includes two papers that extend the coverage of precision grammar systems to incorporate some modal phenomena (precision grammar approaches were not represented in the 2012 volume). The other papers are devoted to annotation efforts. In those, the focus is on distinguishing the various KINDS of modal expressions, DEGREE is not focussed on and STRENGTH is handled as a subdivision in the kinds of modals. The contributions all follow the linguistic conceptualization of modality and have the same core kinds of modalities but they differ in the subdivisions and the additions to the core.

The first two papers build on type-theoretical logics to extend discourse semantics to account for some modal phenomena. The first one, **Modal Subordination in Type Theoretic Dynamic Logic (TTDL)**, extends the TTDL, a Montagovian account of discourse se-

mantics developed by de Groot (2001), to handle epistemic modality and modal subordination. In contrast with some other dynamic semantics theories, TTDL is completely compositional. It has been implemented in the Abstract Categorical Grammar Toolkit of de Groot. The extension to modals follows Kratzer's theory. A part of this extension has also been implemented but, as the authors observe, the system uses a full possible world semantics to model propositions which leads to efficiency and complexity problems.

The second paper, **Goal-Oriented Modality in Type Theory with Records (TTR)**, uses a different incarnation of Type Theory that addresses this implementation problem by using types of situations instead of possible worlds. Whereas the number of possible worlds and their orderings needed to represent anankastic modalities is unworkable, an ordering on a limited number of types of situations might be computationally acceptable. The paper develops a treatment of anankastic (teleological) modality in the TTR and show how it can be extended to other goal-oriented statements such as purpose clauses. An adequate treatment of this type of modality is of obvious importance to planning.

The next four papers address different aspects of annotating corpora with modal information as input for a statistical learning approach to identify the different kinds of modals. **Modal Sense Classification at Large** is the most ambitious of the four in that it aims, first, at creating a large annotated corpus because the usual small corpora are not sufficient to test whether the different senses can be distinguished reliably and, second, to apply statistical learning methods to automatically annotate it. The paper also proposes a semantically grounded feature space for this classification and investigates whether there are genre differences that influence the performance of the classification models. The corpus construction scales up by exploiting paraphrases. As noted above, modality can be expressed by various linguistic means. Whereas modal verbs tend to have many senses, adverbial modal expressions tend to be less ambiguous. This makes it possible to exploit translations as a means to disambiguate modal verbs in one language via the sense of a modal expression in another language for which there are parallel corpora. Here a translation between German and English is used to achieve this. Adapting Ruppenhofer and Rehbein (2012), the system is geared to distinguish between possibility (*epistemic*), request and permission (*deontic*) and ability (*dynamic*). The range of features used is different from those used in most work on modality where mainly POS and simple lexical features are used. Here, the features are morpho-syntactic reflexes of semantic notions such as voice, tense, aspect, person (and other conceptual types), and WordNet fea-

tures. The addition of semantic features showed strong improvements over prior work. The project also hand-annotated a subset of MASC (Ide et al., 2008) to study how modal sense distributions (that are heavily skewed) can also differ across genres showing one cannot rely on a uniform ‘most frequent sense’, but really needs good (semantic) classification models for the modal senses.

The two next annotation papers aim to develop annotation schemes for languages other than English, specifically for Spanish and Portuguese.

A linguistically-motivated annotation model for modality in English and Spanish is an interim report on the elaboration of an annotation schema for these two languages. It starts from a four-way distinction among modals: *epistemic*, *deontic*, *dynamic* and *volitional modality*. These are further subdivided in several subclasses. The annotation experiment run in English and in Spanish, not only annotates (the Spanish equivalents of) modal auxiliaries but modal expressions in all grammatical classes. It uses a bilingual corpus, MULTINOT, consisting of both original and translated texts. The inter-annotator agreement for the experiments was judged high for both the coarse and the more refined classes. The experiment allowed to identify some lexical items that can be annotated automatically but also some difficult cases of ambiguous modals and of non-auxiliary triggers where annotators differ in their appreciation of the term as modal or not.

The contribution on Portuguese, **Modality annotation for Portuguese: from manual annotation to automatic labeling**, aims at developing an annotation scheme for both written and spoken Portuguese and for the European and Brazilian variety of the language. It evaluates two existing proposals that have been developed separately, one for written European Portuguese and one for spoken Brazilian. They consider different types of expressions of modality, not just modal auxiliaries, and aim to identify not only the modal but also the target of the modal expression and the source of the modality. The main modal categories are *epistemic* (with subdivisions), *deontic* (with subdivisions) but the European Portuguese scheme had also *participant-internal* (with subdivisions), *volition*, *evaluation*, *effort* and *success* as categories. The unified scheme groups some of the latter under a general *dynamic* category. The European corpus is also used to train an automatic tagger using different sets of features extracted from the parse tree for the ambiguous items. As in Modal Sense Classification at Large, the features were inspired by Ruppenhofer and Rehbein (2012). The results of various feature combinations and a bag-of-words approach were compared. Most results show an improvement over the baseline (most

common interpretation) but only rarely improvement due to parsing information. The paper presents a discussion of some issues that the annotation effort highlighted: the difficulty for annotators to decide on the scope of negation, the influence of markers that influence the strength of the modality, the importance of some syntactic features for some verbs, the difficulty to distinguish between ambiguity and indeterminacy. The unified scheme will be used for new experiments that will be the basis of a further study of the relevant syntactic and semantic features.

The last paper about annotation, **Selective annotation of Modal readings**, focuses on a very specific problem: how to distinguish the *epistemic* and the *ability* reading of *could*. The authors observe that these two readings correlate with the eventuality described being in the past (ability) or in the future (epistemic) and hypothesize that by asking annotators to decide about the temporal relation of the modal sentence to speech time, one can improve inter-annotator agreement on *could*. They show that this is indeed the case. The finding confirms that of Marasovic et al. (Modal Sense Classification at Large) who found *could* difficult and observed that the difference between past and non-past help distinguish epistemic readings from deontic/dynamic ones.

The last paper in the volume, **Sarcastic Soulmates**, reports on an experiment to train a machine-learning classifier on Twitter messages with the aim of distinguishing between sarcastic messages that were addressed to specific users (known to the sender) and those that were not addressed to specific users. They found that the Tweets used more ‘markers of irony’ when the tweet is sent to unidentified addressees than when it is sent to a known addressee. In a certain sense, this result is surprising because one can argue that there is no such thing as sarcastic or ironic language, there is only the sarcastic or ironical use of language. If one looks at what the authors have in mind, one sees that, indeed, the indicators of irony are intensifiers, hyperbole, exclamation marks and the like, all linguistic devices that, as the authors point out, are not specific to irony. But the experiments reported confirm previous experiments that these textual features allow to distinguish ironic messages from non-ironic ones and show that ironical messages sent to known users differ from those sent to non-known users. To obtain reliable results in detecting irony or sarcasm, however, it is necessary to have an explicit model of the communication situation.

From the submissions it is clear that most of the action is still on the level of annotation. In their introduction to the Computational Linguistics 2012 volume, Morante and Sporleder (2012) write that

‘Currently, many annotation schemes exist in parallel (...). As a consequence, the existing annotated corpora are all relatively small. ... Ideally, any larger scale resource creation project should be preceded by a discussion in the computational community about which aspects of negation and modality should be annotated and how this should be done.’

As said above, there still is no common classification for the modality triggers but Ruppenhofer and Rehbein (2012), based on Kratzer, seems to have inspired some commonalities. Although several different annotation schemes are still deployed on small corpora, at least for English, Marasovic et al. have created a corpus of adequate seize. As Nissim et al. (2013) observe, apart from agreeing on a taxonomy, there should be more uniformity in the actual annotation procedure. They propose the TRIGGER, the SCOPE and the SOURCE as the relevant categories (a proposal that is followed in ‘Modality annotation for Portuguese: from manual annotation to automatic labeling’ in this volume) but observe that various efforts have different ways to annotate those and in most cases do not annotate all of them. The same lack of uniformity can be observed in the papers in this volume. It is most likely easier to make headway on the annotation procedure issue than on coming to a common understanding of that modality taxonomy. Different annotation projects will continue to have different aims so in certain cases a distinction between veridical and non-veridical might be enough, in other cases one might want to know about several degrees of desirability for e.g. some products. In fact, the proposal in Nissim et al. (2013) itself is very incomplete. But this difference in aim should not prevent researchers from having common definitions. At this point we don’t even know whether what is called *deontic* in project A is the same thing as what is called *deontic* in project B. In the realm of *epistemic* modality, there seems to be a general common understanding but for the other kinds of modals the fact that they basically have the same modal base seems to leave the door open for important differences in classification.

For reviewing and advice, thanks to:

Michael Bloodgood

Sergios Chatzikyriadis

Wesley H. Holliday

Justin Kao

Katja Markert

Roser Morante

Larry Moss

Malvina Nissim
Valeria de Paiva
Ines Rehbein
Aynat Rubenstein
Joseph Ruppenhofer
Roser Saurí
Caroline Sporleder
 and especially **Cleo Condoravdi** and **Anette Frank**.

References

- de Groote, Philippe. 2001. Towards abstract categorial grammars. In *Proceedings of the 39th Annual Meeting of the Association for Computational Linguistics*.
- Frank, Anette. 1997. *Context dependence and modal constructions*. Ph.D. thesis, University of Stuttgart.
- Huddleston, Rodney and Geoffrey Pullum. 2002. *The Cambridge Grammar of the English Language*. Cambridge University Press.
- Ide, Nancy, Colin Baker, Christiane Fellbaum, and Charles Fillmore. 2008. Masc: The manually annotated sub-corpus of american english. In *Proceedings of the Sixth International Conference on Language Resources and Evaluation (LREC-2008)*.
- Karttunen, Lauri. 1969. Discourse referents. In *Proceedings of the 1969 Conference on Computational Linguistics (COLING)*.
- Kratzer, Angelika. 1981. The notional category of modality. In H. Eikmeyer and H. Rieser, eds., *Words, Worlds, and Contexts: New Approaches to Word Semantics*. de Gruyter.
- Kratzer, Angelika. 1991. Modality. In v. S. A. Eikmeyer and D. Wunderlich, eds., *Semantics: An International Handbook of Contemporary Research*. de Gruyter.
- Lassinter, Dan. forthcoming. *Graded Modality*. Oxford University Press.
- Morante, Roser and Caroline Sporleder. 2012. Introduction to the special issue. *Computational Linguistics* 38.
- Nissim, M., P. Pietandrea, A. Sansò, and C. Mauri. 2013. Cross-linguistic annotation of modality: a data-driven hierarchical model. In *Proceedings of the 9th Joint ISO-ACL SIGSEM Workshop on Interoperable Semantic Annotation*.
- Palmer, Frank. 1976. *Mood and Modality*. Cambridge University Press.
- Roberts, Craige. 1989. Modal subordination and pronominal anaphora in discourse. *Linguistics and Philosophy* 12.
- Ruppenhofer, Joseph and Ines Rehbein. 2012. Yes, we can. In *Proceedings of the Eight International Conference on Language Resources and Evaluation (LREC-2012)*.