

Modal Verbs in the Common Ground: Discriminating among “Actual” and “Non-Actual” Uses of *Could* and *Would* for Improved Text Interpretation

Lori Moon

(aralluna@gmail.com)

University of Illinois Urbana-Champaign
Urbana, IL 61801 USA

Abstract

Modal verbs occur in contexts which convey information about non-actual states of affairs as well as in contexts which convey information about the actual world of the discourse. Modeling the semantic interpretation of non-actual states of affairs is notoriously complicated, sometimes requiring modal logic, belief revision, non-monotonic reasoning, and multi-agent autoepistemic models. This work presents linguistic features which disambiguate those instances of the past tense modal verbs *could* and *would* which occur in contexts where the proposition in the scope of the modal is not true in the actual world of the discourse from those instances which presuppose or entail that an event in their scope occurred in the actual world of the discourse. It also illustrates the complexity of the role of modal verbs in semantic interpretation and, consequently, the limitations of state of the art inference systems with respect to modal verbs. The features suggested for improving modal verb interpretation are based on the analysis of corpus data and insights from the linguistic literature.

Introduction

Modal verbs such as *could* and *would* allow various interpretations depending on the surrounding linguistic context. The clause, ‘*he would walk to school*’ appears in the four sentences below with four different meanings:

1. When Jeron_{*i*} was five, *he_{*i*} would walk to school* (every day).
2. Jeron_{*i*} stood by the door waiting for hours, later *he_{*i*} would walk to school*.
3. If Jeron_{*i*} were you, *he_{*i*} would walk to school*.
4. Jeron_{*i*} bought a house close to campus so that *he_{*i*} would walk to school*.

In sentences (1) and (2), the modal verb occurs in a context which entails that the proposition in the scope of the modal (corresponding to the constituents underlined in (1) above) refers to an event which actually occurred: the sentences entail that there was a past event of Jeron walking to school in the actual world of the discourse.

In contrast, the modal verbs in sentences (3) and (4) are in non-actual contexts and do not have the same entailment

properties. Rather, sentence (3) is counterfactual and *presupposes* the negation of the antecedent proposition, which, if the presupposition is retained, requires the consequent clause to be evaluated in a world other than the actual world of the discourse (Lewis 1973).

In sentence (4), the italicized clause must be interpreted relative to the speaker’s beliefs about another agent’s intention rather than an actual event (Verstraete 2008). Sentence (4) can be followed with the clause, ‘*but he_{*i*} doesn’t [walk to school]*’, without a resulting contradiction, as can sentence (3). However, if the same clause is appended to the end of sentence (1) or sentence (2), a contradiction results. Negation of the proposition that Jeron walks to school is not possible in sentences (1) and (2) because the use of the clause containing the modal constitutes an assertion that the proposition in the scope of the modal actually occurred. In sentences (3) and (4), the speaker does not make an assertion that a ‘Jeron-walking-to-school’ event ever occurred, and the corresponding proposition, therefore, is not entailed.

Modal verbs pose special problems for implementation tasks which are concerned with meaning representation. Consider an agent which encounters one of the past tense modal verbs above while processing a text. In the first two uses, the agent needs to incorporate into the common ground the fact that the event in the scope of the modal verb occurred repeatedly during a time span which precedes the speech time (for sentence (1)) and at a time which is future relative to a past reference point, but past relative to the speech time for sentence (2).

The interpretation of the non-actual uses is more complicated to model in the common ground. Lewis (1973:9-13) demonstrated that counterfactual conditionals are *variably strict conditionals* and require interpretation in a non-monotonic logic. *Purpose* uses of past tense modal verbs, such as that in sentence (4), describe the speaker’s belief about the another agent’s intentions (Verstraete 2008) and would require interpretation in a multi-agent expansion of an auto-epistemic logic with beliefs and intentions (e.g., (Rao and Georgeff 1991)).

As part of a larger project determining which interpretation models are required for modal verb interpretation and how to link them to natural language expressions, this study describes the linguistic features distinguishing actual and non-actual uses of *could* and *would*. This distinction is of

special importance for systems which want to determine which actions some agent actually performed or avoid deriving real world entailments from non-actual contexts (e.g., the PASCAL RTE Challenge), for Natural Language Inference systems (see Schubert et. al 2010 for the presentation of such a system and for a recent concise summary of previous work), as well as for tasks extracting facts from report documents or encyclopedic sources in order to build a general-purpose knowledge base.

Project Design and Goals

The primary goal of this work is to determine what linguistic expressions co-occur with the past tense modal verbs in actual versus non-actual contexts. Generally, theoretical linguistic studies work with sets of linguistic expressions or morphological markings in order to determine what meanings they systematically express. The approach in this project differs slightly, in that it first classifies naturally occurring data with respect to its entailment properties, and second looks for evidence of the expected linguistic patterns as well as phenomena which have previously gone unnoticed.

Linguistic descriptions of the entailment patterns of modal verbs (see (Portner 2009) for references) and their association with cross-linguistic patterns has been an active area of investigation in the linguistics literature (e.g., (Iatridou 2000), (Giannakidou 2009), (von Stechow and Grønn 2008), and others). There has not, however, been a link made between this theoretical research and the practical task of mapping the generalizations to features which can be straightforwardly encoded in broad-coverage tasks. The gap is largely due to the complexity and variety of interpretations modal verbs require, and the high-level nature of the relevant linguistic features.

Previous implementations which attempted to characterize whether modal verbs were (non)actual overgeneralized by associating lexical modal verbs uniformly with an entailment pattern (e.g., MacCartney et. al 2006, MacCartney(2009:57)). The importance of expressions which signal non-actual contexts continues to be recognized in the textual entailment community, but the work has focused more on categories which can be associated with non-functional lexical expressions. For example, recent work has been done on Negative Polarity Items and downward entailing expressions (Danescu-Niculescu-Mizil and Lee 2010), and on embedding verbs (non-factive, implicative, etc.) and their relations to presuppositions (e.g., (Clausen and Manning 2009)). Modal verbs, however, constitute a significant portion of non-actual contexts in open domain corpora and interact with other expressions signaling (non)actual contexts.

Past tense modal verbs are also under-represented in test sets for state of the art tools, for example, the FraCas set (Cooper et al. 1996) contains only two examples with *would* and none with *could*, making the performance of systems with respect to modal interpretation difficult to measure. As inference systems improve, a more thorough understanding of the various interpretations of modal verbs is required for accurate broad-coverage performance.

Terms and Approach

Some non-actual uses of modal verbs have been associated with linguistic expressions such as clauses with *irrealis* mood or *non-veridical* operators (Giannakidou 2009) (for a concise summary of the complexity of these mappings, see Marques (2009:182-190)). The *non-actual* contexts here refer more broadly to contexts in which the proposition in the scope of the modal has some morphological marking indicating that it describes non-actual information (e.g., the presence of tense marking indicating the subjunctive mood in English) or which, due to some clausal or extra-clausal information, describes a non-actual state of affairs.

The notion of *actual* refers to uses of sentences corresponding to propositions which are either presupposed or entailed to be true in the actual world of the discourse. These propositions are assumed to enter the agent's common ground as facts about the world of the discourse. Non-actual uses are a much broader class and include sentences corresponding to propositions which are only true relative to a participant's belief model, for example, or only true relative to a non-actual possible world. These propositions do not enter the common ground as facts about the actual world of the discourse, but there are fairly reliable presuppositions or inferences associated with them which do provide information about the actual world.

The actual/non-actual distinction differs from (non)veridicality as follows: An operator (e.g., a modal verb, a propositional attitude verb) is considered to be *veridical* if the proposition in its scope is presupposed or entailed to be true in at least one belief world of the speaker's (or subject's) epistemic model (Giannakidou (2009:1889)). Therefore (non)veridicality is a cross-cutting notion with respect to (non)actuality in that it includes propositional attitude verbs which enter the shared common ground as beliefs about another's beliefs or, more generally, beliefs about a proposition under the category 'veridical', making it not equivalent to our notion of 'actual'.

Modal verbs have traditionally been claimed to have a single meaning which takes two arguments: a proposition and a context (Kratzer 1977). The proposition corresponds to the sentence in which the modal occurs (underlined in (5) below), and the context refers to accessible worlds and an ordering relation on the accessible worlds (Kratzer 1981).

5. He could do more to help.

The seeming ambiguities, such as that between epistemic and deontic readings, are said to arise from a difference in accessible worlds and ordering relations.

Non-actual states of affairs relative to which past tense modal verbs are sometimes evaluated can be expressed by an antecedent proposition or by some other linguistic means. The non-actual state of affairs can also be contextually specified. In the descriptions that follow, the term 'non-actual state of affairs' includes all means of specifying the restriction on the clause containing the past tense modal, unless the description mentions linguistic properties of the antecedent clause or the embedding verb.

The letter ϕ and γ will be used as abbreviations for propositions in the discussion which follows.

Data and Classification

In this project, sentences containing *could* and *would* were extracted from an 80,000 line corpus of weblog texts. Each instance of *could* or *would* was indexed and hand-classified according to whether the instance was actual or non-actual, then according to its entailment class as shown in the table below:

modal	(non)actual	entailment class
<i>could</i>	actual	
	non-actual	non-past counterfactual past counterfactual past perfect counterfactual
<i>would</i>	actual	habitual in the past future in the past 'dispositional-indicatives'
	non-actual	hypothetical non-past counterfactual past perfect counterfactual

The three columns in the table above, *modal*, *(non)actual*, and *entailment class* will be referred to as 'levels of classification'. The amount of data classified were approximately 10% of the instances of *could* and *would* extracted from the corpus, in anticipation of using cross-fold validation when automatically classifying the remaining instances with semi-supervised machine learning methods.

This section describes the semantic entailments and presuppositions which enter the common ground and how they are associated with the proposed entailment classes. The following section discusses the probabilistic correlation between the entailments and easily-identifiable syntactic features.

All actual uses of *could* were put into one group, partly due to the small sample size (28 instances). The following distinctions, however, were relevant: (i) uses with a reference time to a past possibility which was realized either simultaneously (with respect to reference time) or subsequently (but before the speech time). In this case, *could*(ϕ) entails that the proposition ϕ enters the common ground at the reference time but preceding speech time. (ii) uses which describe a past impossibility being removed at the reference time. In these uses, *could*(ϕ) entails $\diamond\phi$ and the presupposition holds that $\neg\diamond\phi$ held before the reference time. Instances in the data imply that ϕ occurred, therefore, they were classified with actual uses. And (iii), uses which are in the scope of a veridical verb (e.g., a verb of perception). Often, 'A *could* (*see/hear/discern*) X' entails 'A (*saw/heard/discerned*) X'.

The 85 non-actual instances of *could* were divided into three categories. *Non-past counterfactuals*, *past counterfactuals*, and *past perfect counterfactuals*.

Non-past counterfactuals included present counterfactual and hypothetical instances. Following Iatridou (2000), it holds in the data that the distinction between the hypothetical (or Future Less Vivid) and present counterfactual uses is disambiguated by the Aktionsart of the antecedent predicate, however, many instances of non-past *could* uses were not in conditional constructions, and, therefore, lacked a temporal

anchor at the sentence level. Because this ambiguity is likely to be resolved at the discourse level, it was considered beyond the scope of the current project and the two forms were put into a single category. Non-past counterfactual uses of *could* enter the common ground as information regarding a non-actual state of affairs which, in general, does not hold at speech time, allowing the actual world inference that the proposition in the scope of the modal does not hold.

Past counterfactual uses of *could* are well-known to exhibit 'sequence of tense' phenomena (Abusch 1997) (Portner (2009:223-230)) when embedded under a past tense verb. In these cases, the past tense marking associated with the modal verb *could* can either refer to an event located at the same temporal point as the embedding verb or at a time following. Because a significant number of the past counterfactual uses occurred as the complement of a verb, the ambiguity was frequently exhibited in the data. Because either reading is past relative to the speech time, they were merged into the category of past counterfactuals. Past counterfactuals enter the common ground as information about a state of affairs which (in general) did not hold at a time preceding the speech time. In these cases *could* (ϕ), where the sentence corresponding to ϕ contains a past tense verb, holds relative to a past non-actual state of affairs and, it is inferred that ϕ did not hold in the actual world at that time.

Past perfect counterfactuals were put into their own class in order to keep some parallel between the *would* and *could* non-actual uses and because they never show ambiguity with actual uses. This distinction is more semantically motivated with *would* than *could*, as some examples of *could* without the past perfect marker *have* share a temporal structure with past perfect uses.

The 48 actual instances with *would* seemed to form three major groups. The *habitual in the past* uses involve reference to an event which occurred repeatedly in the past. In uses with habitual past tense, *would*(ϕ) updates the common ground with the information that ϕ occurred repeatedly at some time span which precedes the speech time. The *future in the past* uses refer to an event which occurred after the reference time but before the speech time. In the case of future in the past uses, *would*(ϕ) updates the common ground with information that ϕ occurred in the future relative to the reference time but preceding the speech time.

A third category of uses which naturally emerged are here labeled '*dispositional-indicative*', for lack of a better mnemonic. They include not only examples that would be considered dispositional (e.g., *What man wants, they would do anything to get it. (C1_27)*), but also examples which correspond to similar instances labeled *gnomic* or *habitual future* (*An SUV would get good gas mileage sitting in the drive. (C1_239)*) in other taxonomies (e.g.,(Giomi 2010)). While these examples do not exhibit the same clear-cut entailment properties of the genuine past uses, they have a habitual flavor which makes them seem to support at least a valid common-sense inference that the event has held at some time point or currently holds, that is, that an SUV does get (or has gotten) good gas mileage sitting in the drive and men do indeed do anything (or have done anything) to get what they want (if the speaker's authority to say this can be

trusted). They have a similarity to indicative conditionals in that, whenever the condition is met (e.g., the SUV is in the drive), the proposition in the scope of *would* holds. These forms might be treated as holding whenever the restrictor holds, regardless of the temporal reference point relative to speech time.

Among the non-actual uses of *would*, three categories were identified roughly similar to those for *could*: *hypotheticals*, *non-past counterfactuals*, and *past perfect counterfactuals*. In the case of *would*, the hypothetical and present tense counterfactuals were separated into their own categories. The main reason for this separation is that non-actual uses of *would* were most frequently expressed in the form of a conditional. Therefore, it was possible to disambiguate hypotheticals and present tense counterfactuals based on the Aktionsart of the antecedent verb. Unlike *could*, all of the past tense counterfactuals include the past perfect auxiliary *have*; therefore, there was no motivation to have a past label as well. None of the past uses were ambiguous with actual uses. For hypothetical uses of *would*, the proposition in the scope of *would* enters the common ground relative to a condition which is future with respect to speech time. In present counterfactual uses, the condition does not hold at speech time.

Past counterfactuals have a proposition in the scope of *would* which enters the common ground relative to a past condition which did not hold, as with *could*.

Correlation between Semantic Categories and Syntactic Features

Within each semantic category, the data was organized according to sentence-level features which helped distinguish the uses. The first division was according to syntactic structure. Each final category of the semantic classification was sorted according to the following syntactic division:

Conditional (Cond): instances in the consequent clause of a conditional sentence

Free-Standing (FS): instances in a main clause, possibly with adjuncts or coordination

Relative Clause (RC): instances in a clause headed by a noun phrase

Complement Clause (CC): instances in a clause which is the complement of a verb

Other (OT): instances in adjunct or coordinating constructions which seemed to exhibit exceptional properties

Syntactic constructions were chosen as a level of classification after the semantic categories were formed because the data exhibited some systematic use of constructions and because correlations between syntactic construction and meaning can be encoded relatively easily. Using syntactic properties as a feature correlated with meaning can make improved interpretation easier to achieve and reduce the need for hand annotation (cf. Fernandez et. al 2007). Correlations between syntactic category and semantic classification were measured at the modal verb level, the (non)actual level, and the more detailed entailment class level of the classification.




Figure 1: The distribution of ‘actual’ uses of *could* (gray) and *would* (black).

At the classification level of the (non)actual distinction, no actual instances with either modal verb occurred in the consequent of a conditional.¹ Although such uses are attested in the linguistic literature (Crouch 1993:176), they did not occur in the sample, indicating that they are not frequent. Two highly probable rules, then, are:

Rule i: If a clause contains a past tense modal verb (*could* or *would*), some form of the past perfect marker *have*, and a verb in the past perfect tense, the proposition in the scope of the modal is non-actual.

Rule ii: Given a conditional sentence with a past tense modal (*could* or *would*) in the consequent clause, the use is non-actual.

Rules (i) and (ii) account for the classification of 40.9% of the total data.

Within this distinction, each example contained the temporal sequence in which the antecedent preceded the consequent, that is, there were no instances of *backtracking* (Lewis 1979). The reason for the lack of backtracking was that the Aktionsart of the predicates in the sample was such that the antecedent preceding the consequent was the order predicted by standard theories. So we can also add the highly probable temporal rule to the common ground that, given a conditional sentence with the antecedent γ and consequent ϕ , the reference time of $\gamma \leq$ the reference time of ϕ in the world(s) of evaluation.

With respect to other syntactic categories, there was variation between the *could* and *would* examples: For actual uses of *would*, 54% occurred in free-standing form, 27% as the complement of verbs with relative clause and other uses comprising 10% and 8% respectively, as shown in Figure 1. In contrast, actual uses with *could* were fairly evenly distributed over all non-conditional construction types in the 20%-30% range.

For the non-actual uses, conditionals were much more common for *would* uses, comprising 43% of the instances.

¹Instances with *even if* in the antecedent or *still* in the consequent were difficult to classify. Of the ten examples present in the corpus, five seemed arguably actual, however, since the antecedent is hypothetical, the semantics would still require evaluation of the conditional relative to a non-actual world.




Figure 2: The distribution of ‘non-actual’ uses of *could* (gray) and *would* (black).

Free-standing uses were 32%, making free-standing instances of *would* highly ambiguous without finer-grained analysis. As a working hypothesis for feature design, free-standing actual uses of *would* typically occur in narrative sequences, where the surrounding clauses are in the past tense or as the complement of a verb which is generally considered to be veridical. Relative clauses and complement clauses were a more common strategy for encoding actual uses than non-actual.

Non-actual uses of *could* were much more rare in conditionals (only 11%), with free-standing uses at 33%, comparable with *would*. Relative clause and complement clause, uses, however, were much more prevalent than with *would*, as shown in Figure 2.

Looking at the question of by what linguistic methods actual versus non-actual uses are encoded, it is clear that construction-type plays a much more prominent role in disambiguation with *would* than with *could*. Also, propositional attitude verbs and ‘verbs of saying or telling’ (Levin 1993) form an ambiguous class in naturally occurring data, being present as heads of complement clauses in both actual and non-actual instances.

Of the complement clause examples of both modal types with very strong entailments, 80% of actual and 20% of non-actual uses were headed by verbs of speaking, telling, or transfer of message (Levin 1993), and 40% actual and non-actual examples of both were headed by verbs of ‘declaring’ or ‘conjecturing’ (note that a verb can belong to more than one Levin class).

Some verbs were expected to be heads of non-actual complement clauses (e.g., *wish*, *deny*, *wager*, *predict*, *think*) and actual complement clauses (e.g., *know*, *be glad*, *understand*). But other verbs crossed categories, including *admit*, and *said*.

In current work, we are testing when human subjects infer the proposition in the scope of *could* given the embedding verb’s class and the presence of temporal or descriptive adverbs in the complement clause. Similar testing will be conducted in order to isolate the role of additional syntactic features within the constructions examined in this work.

At the more fine-grained entailment class level, we found that future in the past uses of *would* were usually encoded as complement clauses and that past counterfactual *could*




Figure 3: Given an actual occurrence of the modal *could* (black) or a non-actual occurrence (gray), this chart shows the likelihood that a particular construction will be used.

was usually in conditional form. For reasons previously mentioned, present counterfactual and hypothetical uses of *would* were usually in conditional form, however, hypothetical uses were more frequently in free-standing constructions than present counterfactuals, indicating that, in the absence of antecedent cues, uses of *would* are often taken to be hypothetical. For non-actual uses of *could*, the choice of relative clauses for encoding was prevalent for all entailment classes.

Summary

In conclusion, two easily encoded rules were given which account for the (non)actual classification of 40.9% of the data. For the remaining 60% of the data, the following features are relevant: (i) For complement clauses, inventories of verb classes for the embedding verb (e.g., (Karttunen 1973)), and their (non)veridical properties play an important role. (ii) For free-standing uses, important features are sequences of modal verbs or past tense verbs in coordinated clauses and the presence of temporal adverbs (iii) For relative clauses, the nature of the antecedent matters: If it is non-actual, so is the proposition in the relative clause. The type of verb in the scope of the modal also plays a role in the (non)actual distinction. (iv) In other uses, the type of coordinating conjunction is important (e.g., modal verbs following *so that* were always non-actual, but following *so*, in contrast, were either). Suggestions for finer-grained classification were also suggested.

By proposing an entailment-based taxonomy then correlating it with easily encoded syntactic features, the study makes two significant innovations: First, it elucidates the entailments, presuppositions, and inferences which are systematically associated with various uses of past tense modal verbs. Second, it renders the problem of associating a meaning with natural language data tractable by uncovering the correlations between entailment classes and syntactic features which are easily identified by state of the art tools such as syntactic parsers or lexical ontologies.

Future Work

Because this corpus was built from weblogs which were returned by search results with *could* and *would*, there is a possibility that the sample is biased towards documents with




Figure 4: Given an actual occurrence of the modal *would* (black) or a non-actual occurrence (gray), this chart shows the likelihood that a particular construction will be used.

a high frequency of modal verbs: there were approximately four times as many non-actual instances of *would* as actual and approximately three times as many non-actual instances of *could* as actual. Such documents would, more likely, contain sequences of modal verbs in a description of non-actual events, which is what we found looking at the full weblog texts. For this reason, future work is being done using the English Gigaword Corpus (4th ed.) and will proceed by extracting instances of past tense modal verbs from that corpus. Furthermore, the two past tense English modal verbs, *might* and *should* will be added to the target data, making the set of English past tense modal verbs covered complete.

For cases which were hard to classify, similar examples are being constructed for an experimental study with human subjects to see what inferences are associated with variations in syntactic form and verb classes. A formal theoretical model of the entailment classes is also being further developed for implementation in a logical model.

Acknowledgements

This work benefited from the support of my adviser, Peter Lasersohn, and committee members Tania Ionin, Shalom Lappin, and Jose Meseguer as well as comments from the audience at TMEI 2011 on slides over this topic. All errors are my own.

References

Abusch, D. 1997. Sequence of tense and temporal de re. *Linguistics and Philosophy* 20(1):1–50.

Clausen, D., and Manning, C. D. 2009. Presupposed content and entailments in natural language inference. In *Proceedings of the 2009 Workshop on Applied Textual Inference, ACL-IJCNLP 2009*, 70–73. ACL and AFNLP.

Cooper, R.; Crouch, D.; van Eijck, J.; Fox, C.; van Genabith, J.; Jaspars, J.; Kamp, H.; Milward, D.; Pinkal, M.; Poesio, M.; and Pulman, S. 1996. Using the framework. The FraCaS Consortium.

Crouch, R. 1993. *The Temporal Properties of English Conditionals and Modals*. Ph.D. Dissertation, University of Cambridge.

Danescu-Niculescu-Mizil, C., and Lee, L. 2010. Don't 'have a clue'? Unsupervised co-learning of downward-entailing operators. In *Proceedings of the ACL Short Papers*, 247–252.

Fernandez, R.; Ginzburg, J.; and Lappin, S. 2007. Classifying non-sentential utterances in dialogue: a machine learning approach. *Computational Linguistics* 33(3):397–427.

Giannakidou, A. 2009. the dependency of the subjunctive revisited: Temporal semantics and polarity. *Lingua* 119:1883–1908.

Giomi, R. 2010. Para uma caracterização semântica do futuro sintético românico: descrição e análise dos valores do futuro do indicativo em Português e em Italiano. Master's thesis, Universidade de Lisboa Faculdade de Letras.

Iatridou, S. 2000. The grammatical ingredients of counterfactuality. *Linguistic Inquiry* 31(2):231–270.

Karttunen, L. 1973. Presuppositions of compound sentences. *Linguistic Inquiry* 4(2):169–193.

Kratzer, A. 1977. What 'must' and 'can' must and can mean. *Linguistics & Philosophy* 1:337–355.

Kratzer, A. 1981. *Words, Worlds and Contexts*. Berlin: de Gruyter. chapter The Notional Category of Modality.

Levin, B. 1993. *English Verb Classes and Alternations*. University of Chicago Press.

Lewis, D. 1973. *Counterfactuals*. Blackwell Publishers.

Lewis, D. 1979. Counterfactual dependence and time's arrow. *Nous* 13:455–476. Reprinted in Lewis (1986) *Philosophical Papers* (Vol. 2). Oxford: Oxford University Press.

MacCartney, B.; Grenager, T.; de Marneffe, M.-C.; Cer, D.; and Manning, C. D. 2006. Learning to recognize features of valid textual entailments. In *Proceedings of the Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics (HLT-NAACL 2006)*, 41–48.

MacCartney, B. 2009. *Natural Language Inference*. Ph.D. Dissertation, Stanford University.

Marques, R. 2009. *Cross-linguistic Semantics of Tense, Aspect, and Modality*. John Benjamins Publishing Company. chapter On the selection of mood in complement clauses, 179–204.

Portner, P. 2009. *Modality*. Oxford Surveys in Semantics and Pragmatics. Oxford University Press.

Rao, A. S., and Georgeff, M. P. 1991. Modeling rational agents within a BDI-architecture.

Schubert, L. K.; Durme, B. V.; and Bazrafshan, M. 2010. Entailment inference in a natural logic-like general reasoner. In *AAAI 2010 Fall Symposium Series*, volume (CSK-10).

Verstraete, J.-C. 2008. The status of purpose, reason, and intended endpoint in the typology of complex sentences: implications for layered models of clause structure. *Linguistics* 46(4):757–788.

von Stechow, A., and Grønn, A. 2008. Tense and presupposition in counterfactuals. Handout from a talk at New York University.