Sense-Specific Implicative Commitments

Gerard de Melo[1,⋆] and Valeria de Paiva[2]

[1] IIIS, Tsinghua University, Beijing
gdm@demelo.org
[2] Nuance Communications, Sunnyvale, CA
Valeria.dePaiva@nuance.com

Abstract. Natural language processing systems, even when given proper syntactic and semantic interpretations, still lack the common sense inference capabilities required for genuinely understanding a sentence. Recently, there have been several studies developing a semantic classification of verbs and their sentential complements, aiming at determining which inferences people draw from them. Such constructions may give rise to implied commitments that the author normally cannot disavow without being incoherent or without contradicting herself, as described for instance in the work of Kartunnen. In this paper, we model such knowledge at the semantic level by attempting to associate such inferences with specific word senses, drawing on WordNet and VerbNet. This allows us to investigate to what extent the inferences apply to semantically equivalent words within and across languages.

1 Introduction

Understanding a sentence requires more than just decoding its syntactic and semantic structure. Even when supplied with proper syntactic and semantic interpretations of a sentence, current natural language processing systems still lack the common sense inference capabilities required to interpret it in the way humans do. Given a sentence like "John missed that Mary had left", we are inclined to infer that Mary had indeed left. In contrast, given "John pretended that Mary had left", we are inclined to presuppose the opposite. Although both sentences share a common structure, involving verbs with sentential complements, there are clear differences in the types of implicative commitments the author is making. Such differences have been studied in detail by recent studies that have attempted to develop a classification of verbs (and verb-noun collections) that take sentential complements [1]. Similar analyses can be made with respect to other constructions, e.g. to study adjective constructions like "It is confusing that Mary has left" vs. "It is improbable that Mary has left" [2]. Likewise, one can also study implicative commitments regarding the existence of entities. For

---

instance, "She cancelled the meeting" generally leads us to believe that there was no meeting, while "They caused a strike" means that a strike did occur [3].

Certain sentences give rise to implied commitments that the author normally cannot disavow without being incoherent or without contradicting herself. Previous work has developed detailed classifications of words and their implicative commitments at the lexical level [4]. In this paper, we build on this work but attempt to model such knowledge about commitments at the semantic level instead. For this, the commitments are tied to specific word senses, drawing on WordNet and VerbNet as sense inventories. We rely on automatic disambiguation methods to produce a lexicon of sense-specific implicative commitments. Among other things, this allows us to investigate to what extent such inferences also apply to semantically equivalent words within and across languages. Our findings indicate that they are for the most part preserved when transitioning from words to synonym sets, and at least most of the English classifications collected by Nairn et al. [4] do seem to transfer to other languages such as Portuguese as well.

2   Implicative Commitments

Karttunen's seminal "The Logic of English Predicate Complement Constructions" [5] makes the point that while it is valid and helpful to classify verbal constructions according to their syntactic characteristics, e.g. whether they take propositional complements, it is also valid and even more useful to classify them according to their semantic characteristics, such as the *factivity* of complements. The term *factive verb* introduced by Kiparsky and Kiparsky [6] refers to the notion that any simple assertion using one of these verbs (e.g. "John knew that Mary had left") commits the speaker to the belief that the complement sentence ("Mary had left"), just by itself, is also true.

Karttunen observes that it is sometimes "possible to show that there is a definite connection between the semantic properties of a verb and certain syntactic characteristics", and thus goes on to investigate parts of this connection. The syntactic characteristics referred to are that these are propositional complement constructions, while the semantic characteristics are what he refers to as the "logic" in the title of his book. The latter is more general than simply factivity. In fact, 35 years later, as part of a project to extract information from text for question answering, Karttunen, together with Nairn and Condoravdi [4] presented an algorithm for detecting author commitment to the truth or falsity of complement clauses based on their syntactic type and on the meaning of their embedding predicate. They also created a small lexicon of around 300 verbs spelling out the implicative commitments that these verbs indicate. This was the starting point of our work. We contend that such implicative commitments are best described at the semantic level, as one would hope given Karttunen's choice of the word "logic" in the title of his seminal work.

The same applies to other sources and forms of implicative commitments. For instance, in some cases, nouns can induce factivity commitments. As an example, "John took the opportunity to sing" entails that he did sing [7]. Similarly,

adjectives, too, can be marked for factivity and implicativity, as in a recent study [2] that collected such annotations using Amazon Mechanical Turk. Additionally, one can study implicative commitments relating to existential and temporal aspects [3].

3   Veridicity, Veracity, or Veridicality?

Throughout this paper, we are adopting a generalized, somewhat underspecified notion of implicativity that subsumes factivity, entailment, and presupposition, calling them "implicative commitments" for want of a better name. This notion seems more robust and clear-cut, as described by Nairn et al. [4], and exemplified in their small lexicon.

It remains clear that our lexicon would have to be just one of several ingredients in any natural language understanding systems aimed at fully assessing the *veridicity* of textual content. Has an event mentioned in the text really occurred? Who is the source of the information? What is the stance of the author of the text? Does the author indicate whether he or she believes the source? This more encompassing problem is discussed by Karttunen and Zaenen [8] and in Saurí and Pustejovsky's FactBank corpus and factuality studies [9,10,11].

The work in FactBank is an attempt to mark the implicative commitments of both the author and some participants of the events in some given text. As Saurí and Pustejovsky explain

> Identifying the veracity, or factuality, of event mentions in text is fundamental for reasoning about eventualities in discourse. Inferences derived from events judged as not having happened, or as being only possible, are different from those derived from events evaluated as factual. Event factuality involves two separate levels of information. On the one hand, it deals with polarity, which distinguishes between positive and negative instantiations of events. On the other, it has to do with degrees of certainty (e.g., possible, probable), an information level generally subsumed under the category of epistemic modality.

The work on FactBank has been extended by the Stanford NLP group [12], who re-annotated the same sentences using workers on Amazon Mechanical Turk, given their stance that textual entailment is to be measured in terms of the understanding of text in ordinary people, not logicians, philosophers, or linguists.

While veridicality or factuality is a possible long term goal, we instead start from the less lofty goal of *veridicity*, a level of assessment that we believe can be coded up in the lexicon. As Karttunen and Zaenen put it:

> It is useful to distinguish between two ingredients that go into determining the truth value of utterance, one is the trustworthiness of the utterer and the other is the stance of the utterer vis-à-vis the truth of the content. The latter we will call the veridicity of the content [...].

The classification chosen by Nairn et al. [4] focuses on cases in which the author's commitment to the truth of a complement clause arises solely from the larger sentence it belongs to, leaving aside other sources of information about the beliefs of the author. Note that it can be difficult to decide between *entailments*, that is, what the author is actually committed to, and the more pragmatic notion of *conversational implicatures*, that is, what a reader/hearer may feel entitled to infer. For example, "Ed did not refuse to participate" might lead the hearer to conclude that Ed participated. But the speaker could continue with "He was not even eligible" indicating the opposite.

The tags used to describe the implicative commitments of lexical items need to distinguish several types of context-specific behaviors. Some words yield an entailment in both affirmative and negative environments but there are others, "one-way implicatives", that yield entailments only in one or the other environment. Furthermore, the entailment may be either positive or negative depending on the polarity of the environment. Altogether Nairn et al. provide us with a table of "implication signatures" for a large class of complement-taking constructions. These implication signatures can be stacked together, which gives rise to their algorithmic solution for dealing with more complex phrases [4].

4   Describing Sense-Specific Implicative Commitments

In this paper we advance the state-of-the-art in that, instead of associating implicative commitments just with lexical items, we associate them with word senses, as represented by WordNet synsets. Our input is a set of resources that describe implicative commitments at the lexical level (included the Nairn et al. lexicon and further resources listed later on in Section 5.1). Our output will contain descriptions of implicative commitments at the sense level.

Generally speaking, an input entry at the lexical level consists of a partially instantiated construction and a set of tags describing the implicative commitments of that construction in particular contexts. The construction may involve one or more lexical items to be disambiguated. For instance, often, there is simply a single verb like "to find" in a syntactic frame expecting a subject and a that-complement ("John found that Mary had arrived"). Instead of describing the construction for a specific lexical item like "to find", we can instead characterize it at a more abstract semantic level, allowing an instantiation with several additional words ("notice", "discover", etc.). We thus have to disambiguate the word "to find" in order to establish which senses are relevant and know which sense-specific synonyms can be assumed to give rise to the same commitments.

In some cases, the construction may also involve two lexical items that need to be disambiguated. For instance, the construction may involve both the verb "avoid" and a description of its possible arguments, described by the word "event". These will then both be mapped to WordNet synsets, and WordNet's synset hierarchy will allow us to to recognize specific instantiations, e.g. "avoiding a strike" vs. "avoiding a road bump" – In the first case, there is no strike, whereas in the second case, the road bump remains.

4.1   Sense Mappings and Disambiguation

Our goal is thus to map one or more input lexical items $e = (l, p, f)$ on the one side, each consisting of a lemma l, part-of-speech tag p, and syntactic frame marker f, to one or more candidate WordNet synsets s on the other side. For example, our input entry could be for the lemma "admit" with part-of-speech tag `verb`, and a marker for a syntactic frame with infinitival complement (as in "The report was admitted to be incorrect"). WordNet lists a total of 8 verb senses for the input lemma "admit". For instance, the first WordNet sense for "admit" refers to the process of allowing people to enter a building, and the second one refers to admittance into a group or a community. These do not reflect the meaning that our input entry is targeting. Our system instead needs to choose the WordNet sense for *declaring to be true or admitting the existence or reality or truth of something.*

In order to achieve such disambiguated mappings in the case of verbs, we draw on another traditional lexical resource VerbNet [13]. VerbNet is a lexicon of Levin-style verb classes based on the syntactic frames of English verbs. Since the input lexicons often refer to specific syntactic frames, e.g. for sentential complements, the syntactic characterizations present in VerbNet can help us pick relevant senses.

For a given input entry e with lemma l, let $W(l)$ denote the set of WordNet senses and $V(l)$ denote the set of VerbNet senses for l. We first use the following similarity function to score the candidate VerbNet senses $v \in V(l)$:

$$\text{sim}(e, v) = \frac{1}{|V(l)|} \left(\epsilon + \max_{f_v \in F(v)} \text{sim}(f_v, f) \right) \tag{1}$$

Here, the strength of association between e and one of the verb senses $v \in V(l)$ first of all depends on the number of candidates $|V(l)|$. Additionally, we obtain the relevant VerbNet class and its syntactic frame information from VerbNet, using $F(v)$ to denote the set of syntactic frames associated with the VerbNet class of v.

We rely on a simple rule-based similarity measure between syntactic descriptors f_v from VerbNet and descriptors f from the original lexicon, which in our case are based on PARC's Bridge system [14]. For instance, VerbNet's `NP V that S` would match the `V-SUBJ-COMPEXthat` descriptor from the Bridge system but not `V-SUBJ-OBJexpl-XCOMPinf`. While this similarity function clearly prefers VerbNet senses that have at least one matching syntactic frame, we still retain the remaining VerbNet senses as possible candidates by setting ϵ to a small constant greater than 0, because VerbNet's syntactic frame descriptions are not always complete.

We then use these VerbNet similarities to help us in scoring the candidate WordNet synsets $s \in W(l)$:

$$\text{sim}(e, s) = \begin{cases} \epsilon + \sum\limits_{v \in V(l)} \dfrac{\max\limits_{s_v \in V_w(v)} \text{sim}(s_v, s)}{1 + |V_w(v)|} \, \text{sim}(e, v) & s \in W_p \\ 0 & s \notin W_p \end{cases} \tag{2}$$

Here, first of all, for a given part-of-speech tag p, W_p denotes the global set of all WordNet senses for that part-of-speech tag. If we are attempting to disambiguate an entry e for which the part-of-speech tag p denotes adjectives, for instance, then only adjective synsets $s \in W_p$ receive a non-zero similarity score.

For a given VerbNet sense v, we use $V_w(v)$ to denote the set of WordNet synsets mapped to v. For this, we make use of the existing disambiguated mappings between VerbNet and WordNet provided with the former. VerbNet's verb senses tend to be more coarse-grained, as their induction was mainly guided by syntactic considerations. Thus for a given VerbNet sense v, the set $V_w(v)$ may contain multiple WordNet synsets. We compare these synsets with our current synset of interest s, simply by using the identity function as the similarity measure $\mathrm{sim}(s_v, s)$, although there are also many existing WordNet sense similarity measures that could be plugged in here. If none of the synsets $s_v \in V_w(v)$ are similar, then the overall value of $\mathrm{sim}(e, s)$ will turn out to be just a small $\epsilon > 0$. In contrast, if there are similar synsets, we still make sure to discount using the set cardinality $|V_w(v)|$, based on the intuition that the score should be higher if there are fewer possible synset matches and thus less ambiguity.

These scoring functions thus allow us to select likely WordNet synsets for a given input entry.

4.2   Lexicon

Once we have disambiguated all the lexical items involved in the construction based on the disambiguation procedure described above, we can assign the implicative commitments tags associated with the original construction to our newly sense-disambiguated version of the construction. The latter refers to specific senses of lexical items, thus eliminating possible confusion about which senses are meant in the case of ambiguous words. At the same time, the latter is more general by enabling us to instantiate the construction with alternative synonymous words, rather than forcing us to select the original lexical item.

In addition to the new sense-specific implicative commitment markings, our final lexicon incorporates information from WordNet, VerbNet, and other lexical resources. For a given word, we list a number of (potentially overlapping) senses. Each sense corresponds either to a WordNet synset or to VerbNet sense or to both, based on the mappings between the two resources provided with VerbNet.

For some of these word senses, we now have information about constructions they can be involved in and the implicative commitments that are entailed when relevant words with those senses (or hyponym senses) are used in those specific constructions. The implicative commitments are provided as properties of the constructions, marked using tags like `fact_p` (based on the ones in the original lexicons). The constructions are linked to specific word senses and thus our lexicon describe the implicative commitments that words with a corresponding word sense entail when used in the corresponding syntactic constellation.

Additionally, our lexicon also provides references to other resources. For every sense corresponding to a VerbNet entry, we list the corresponding FrameNet

[15] frames, relying on the SemLink project's mappings between VerbNet and FrameNet [16]. For every sense corresponding to a WordNet entry, we also include the relevant SUMO [17] concepts based on the pre-existing WordNet-SUMO mappings [18].

We make use of the RDF standard for information interchange, but our data can be converted to any number of other formats, including simple tab-separated-value files. It is clear that this is just a first step towards a freely available open "unified" lexicon in the spirit of Crouch & King [19]. In future work, we additionally hope to add presupposition relationships between different verbs, as studied by Temper and Frank [20] (e.g., given "Spain *won* the tournament", one may presuppose that Spain *played* in the tournament). Another possible extension is to mine corpora for information about downward entailment as proposed by Danescu-Niculescu-Mizil et al. [21].

5   Results

5.1   Data

For our input, we relied on the lexicons from Nairn et al. [4] and related papers, as distributed by CSLI[1], the creation of which was largely DARPA-funded, in the context of the ACQUAINT project.

A lexicon entry consists of a lemma, a syntactic category descriptor, one or more tags to describe the implicative commitment behavior, and frequently also an example sentence for the lemma in the respective syntactic frame highlighting such a possible implicative commitment.

In addition to the simple factives data and the simple implicatives data, we also included the more recent phrasal implicatives data [1], which captures implicative behaviors that depend on the specific arguments of a verb. We likewise included the CSLI data about the factivity of adjectival statements, e.g. "It is accurate that John informed the president" vs. "It is untrue that John informed the president, as well as the reduced set of markings on nouns from the appendix of Price et al. [7].

Moreover, we incorporated the CSLI infinite temporal markings data, which provides information about possible temporal shifts in addition to factivity information. For instance, for a sentence like "Mary persuaded Ed to cook dinner", the cooking can happen after the persuasion, while for "Mary let Ed cook dinner" the cooking typically is simultaneous with the letting. Finally, we created a lexicon of entries about existential commitments, drawing on analyses in the Amaral et al. paper [3], among other sources.

To disambiguate all of this data with respect to WordNet, we made use of the algorithm described in Section 4.1, choosing the highest-ranked synset for each lexical item in the input. In the case of ties (in particular for the nouns and adjectives, where the VerbNet-based heuristics do not apply), we chose the top WordNet synset in terms of WordNet's original frequency-based sense ranking.

---

[1] `http://www.stanford.edu/group/csli_lnr/Lexical_Resources/`

Table 1 summarizes the results of this process. For each category of information, we describe the number of lexical items in the input and the relevant number of candidate senses in VerbNet and WordNet. We then list the number of annotated entries in the input and in our disambiguated output. Since for many data types, annotated entries only refer to a single lexical item, these numbers are often quite correlated. On a random sample of 20 simple implicatives, our automatic disambiguation had a precision of 75.0%. This is a reasonable result given that we are using fine-grained WordNet word senses and that particularly for verbs we have to choose from a very large set of candidates, as can be seen in Table 1. We could not reliably assess the recall of our disambiguation, because for many of the rarer word senses, it is in fact more difficult to establish whether the same implicatives hold as for the primary senses of those lexical items.

In any case, we are currently in the process of relying on human annotators to correct and extend our data, giving us a more complete set of reliable WordNet sense annotations.

Table 1. Implicative Commitment Data

Data Type	Lexical Items			Annotated Entries	
	Input	VN Cand.	WN Cand.	Input	Disambiguated
Simple factives	108	188	486	108	108
Simple implicatives	114	223	950	114	114
Phrasal implicatives	17	67	584	14	14
Factivity of adj. comp.	278	0	449	278	225
Factivity of extraposed adj. comp.	695	0	1534	695	623
Factivity of noun comp.	112	0	304	112	108
Temporal implicatives	86	160	811	86	86
Existential commitments	74	101	428	62	62

5.2   Discussion

One of the low hanging fruits of using our sense-specific implicative markings is the extension of coverage it affords. For example the original lexical markings repository provides the marking for "to bother" but not for its WordNet synonym "to trouble", as in "It troubled Ed that Mary was about to leave". Similarly there is a marking for "to acknowledge" as in "They acknowledged that the report was correct" but not for "to concede" in the same frame.

Obviously, not all synonyms can be used in the syntactic frame specified by the annotation entry. WordNet's synonym sets are semantically motivated, and as such may include words with different syntactic behavior. Thus the extension only applies to words with compatible syntactic behavior. For example, "to hate" is marked as factive as in "Ed hated to leave the party" or "Ed hated that Mary went home". The synonymous verb "to detest" was not marked in the original lexicon, but is a possible synonym listed by our lexicon.

Another advantage of a semantic approach is that we can more accurately describe the behavior of words when this behavior depends on the semantic context. For example, if you have "wasted your chance to go to Paris" , then you probably did not go, whereas if you "wasted your money to go to Paris", then you probably did go [1]. Similarly, "He withdrew his hand" doesn't affect the existence of the hand, while "He withdrew his offer" does reflect a change regarding the existence of the offer. Our lexicon distinguishes the two cases based on the semantic type of the complement. Given a novel example, an NLP system can make use of WordNet's hypernym hierarchy or of WordNet-based semantic relatedness measures to determine which of the two cases is more likely.

Of course, our lexicon nevertheless lacks descriptions to account for certain more involved contexts. As described in Section 3, it seems difficult to account for all possible context-specific behaviors. For example, while "to cause something" generally implies that that something comes to be, there are also examples like "The decree was causing a revolution when it was revoked", from which one is likely to conclude that a revolution was ultimately avoided [3].

5.3   Cross-Lingual Applicability

Our WordNet sense markings not only enable us to find new synonyms, but also allow us to look up non-English equivalents in sense-aligned non-English versions of WordNet, such as EuroWordNet [22] and UWN [23].

We have additionally arranged for a human-created Portuguese translation of the Nairn et al. lexicon and examples[2], and checked that most of the inferential behavior is preserved under direct translations. Some anecdotal observations can be made about this translation. First, as expected, the translation tends to indicate more English verbs, with subtle variations of meaning, going to a single Portuguese verb, for example "abhor" and "abominate" mapping to "abominar", or "acknowledge" and "recognize" to "reconhecer", or, more telling, "amaze"/"astonish"/"surprise" going to "surpreender". All the factive verbs and their examples provided by Nairn et al. seem to work in Portuguese, very much like in English. This is despite some non-direct translations, for example sometimes a single work in English like "affect" becomes a phrasal verb in Portuguese "fazer de conta" (make belief), or "to perplex" becomes "deixar perplexo". Things are not always as clear-cut on the implicatives and their examples. More work is needed here.

6   Conclusion

In this paper, we have described the creation of a freely available lexical resource that encodes sense-specific implicative commitments.

Our first contribution is to bring the remarkable information distributed by CSLI to a wider audience, so that they can be improved, and lexicographical

---

gaps plugged, as suggested by Nairn, Condoravdi, and Karttunen themselves. Their data, together with further data collected from other papers mentioned earlier forms the basis of our resource.

Our second contribution is a first step towards an open source "Unified Lexicon" that aligns and combines this information with main-stream semantic resources. We transfer the original information to the level of word senses, enabling applications to benefit from greater coverage as well as from finer-grained sense-specific entailment information. Our lexicon integrates this information with WordNet, VerbNet, FrameNet, and SUMO, leading to a single one-stop resource that has enormous potential as a backbone for semantic and pragmatic inference as well as for linguistically oriented ontologies.

Our third contribution is some evidence to the fact that lexical resources in English can possibly (and profitably) be used to induce lexical resources in other languages. Given that many years of work have been spent in producing a variety of lexical resources for English, we would like to channel all this effort into other useful projects for several different languages to the extent this is feasible. We appreciate that these resources will not work in multiple languages "out-of-the-box" in a satisfactory manner – completions and adaptations will be necessary – but they do seem to provide a baseline to bootstrap your work from. It is encouraging to see measurable signs that some things work the same way, if we discuss them at the level of concepts.

While some of the predictions in this sort of resource may be more clear-cut than others, individual researchers can tailor it to their applications. For example, if logicians insist that the expression "X says that Y" is always agnostic on the truth-value of "Y" no matter who or what "X" is, a more pragmatic system may decide that if "X" is "The New York Times" or the FDA (Food and Drug Administration of the US) and the author of the sentence is also a reputable source, then "Y" is to be considered true. We thus hope that our resource will be adopted for use in many systems, as the focus in natural language processing shifts from more fundamental operations to higher-level tasks requiring advanced pragmatic inferences.

Our lexicon is available from `http://lexvo.org/implicative-lexicon/`.

References

1. Karttunen, L.: Simple and phrasal implicatives. In: Proceedings of the First Joint Conference on Lexical and Computational Semantics: Proceedings of the Main Conference and the Shared Task, Proceedings of the Sixth International Workshop on Semantic Evaluation, SemEval 2012, vol. 1&2, pp. 124–131. Association for Computational Linguistics, Stroudsburg (2012)
2. Zaenen, A., Karttunen, L.: Veridicity annotation in the lexicon? a look at factive adjectives. In: Proceedings of the 9th Joint ISO - ACL SIGSEM Workshop on Interoperable Semantic Annotation, pp. 51–58. Association for Computational Linguistics, Potsdam (March 2013)

3. Amaral, P., de Paiva, V., Condoravdi, C., Zaenen, A.: Where's the meeting that was cancelled? existential implications of transitive verbs. In: Proceedings of the 3rd Workshop on Cognitive Aspects of the Lexicon, Mumbai, India, The COLING 2012 Organizing Committee, pp. 183–194 (December 2012)

4. Nairn, R., Condoravdi, C., Karttunen, L.: Computing relative polarity for textual inference. In: Proceedings of ICoS-5 (Inference in Computational Semantics) (2006)

5. Karttunen, L.: The logic of English predicate complement constructions. Publications of the Indiana University Linguistics Club, Bloomington (1971)

6. Kiparsky, P., Kiparsky, C.: Fact. In: Bierwisch, M., Heidolph, K.E. (eds.) Progress in Linguistics: A Collection of Papers, Mouton, The Hague, pp. 143–73 (1970)

7. Price, C., de Paiva, V., Holloway King, T.: Context inducing nouns. In: COLING 2008: Proceedings of the Workshop on Knowledge and Reasoning for Answering Questions, Coling 2008 Organizing Committee, Manchester, UK, pp. 9–16 (August 2008)

8. Karttunen, L., Zaenen, A.: Veridicity. In: Katz, G., Pustejovsky, J., Schilder, F. (eds.) Annotating, Extracting and Reasoning about Time and Events. Dagstuhl Seminar Proceedings 05151, Dagstuhl, Germany, Dagstuhl, Germany (2005)

9. Saurí, R., Pustejovsky, J.: Factbank: A corpus annotated with event factuality. Language Resources and Evaluation 43(3), 227–268 (2009)

10. Saurí, R., Pustejovsky, J.: From structure to interpretation: A double-layered annotation for event factuality. In: The 2nd Linguistic Annotation Workshop, LREC, Marrakech, August, May 26-27 (2008)

11. Saurí, R., Pustejovsky, J.: Are you sure that this happened? Assessing the factuality degree of events in text. Computational Linguistics 38(2) (2012)

12. de Marneffe, M.-C., Manning, C.D., Potts, C.: Veridicality and utterance meaning. In: Proceedings of the Fifth IEEE International Conference on Semantic Computing: Workshop on Semantic Annotation for Computational Linguistic Resources. IEEE Computer Society Press, Stanford (2011)

13. Kipper-Schuler, K.: VerbNet: A Broad-Coverage, Comprehensive Verb Lexicon. PhD thesis, University of Pennsylvania (2005)

14. Bobrow, D.G., Cheslow, B., Condoravdi, C., Karttunen, L., King, T.H., Nairn, R., de Paiva, V., Price, C., Zaenen, A.: PARC's Bridge and Question Answering system. In: King, T.H., Bender, E.M. (eds.) Proceedings of the GEAF 2007 Workshop, pp. 46–66. CSLI, Stanford (2007)

15. Baker, C.F., Fillmore, C.J., Lowe, J.B.: The Berkeley FrameNet project. In: Proc. COLING-ACL 1998, pp. 86–90 (1998)

16. Bonial, C., Stowe, K., Palmer, M.: Renewing and revising SemLink. In: Proc. 2nd Workshop on Linked Data in Linguistics, Pisa, Italy (2013)

17. Niles, I., Pease, A.: Toward a Standard Upper Ontology. In: Welty, C., Smith, B. (eds.) Proceedings of the 2nd International Conference on Formal Ontology in Information Systems, FOIS 2001 (2001)

18. Niles, I., Pease, A.: Linking lexicons and ontologies: Mapping WordNet to the Suggested Upper Merged Ontology. In: Proceedings of the IEEE International Conference on Information and Knowledge Engineering, pp. 412–416 (2003)

19. Crouch, D., King, T.: Unifying lexical resources. In: Proceedings of the Interdisciplinary Workshop on the Identification and Representation of Verb Features and Verb Classes, Saarbruecken, Germany (2005)

20. Tremper, G., Frank, A.: A Discriminative Analysis of Fine-Grained Semantic Relations including Presupposition: Annotation and Classification. In: Dipper, S., Zinsmeister, H., Webber, B. (eds.) Dialogue and Discourse, Special Issue on Annotating Pragmatic and Discourse Phenomena, pp. 282–322 (2013)

21. Danescu-Niculescu-Mizil, C., Lee, L., Ducott, R.: Without a "doubt"? Unsupervised discovery of downward-entailing operators. In: HLT-NAACL, pp. 137–145. The Association for Computational Linguistics (2009)
22. Vossen, P. (ed.): EuroWordNet: A Multilingual Database with Lexical Semantic Networks. Springer (1998)
23. de Melo, G., Weikum, G.: Towards a Universal Wordnet by learning from combined evidence. In: Proceedings of the 18th ACM Conference on Information and Knowledge Management, CIKM 2009, pp. 513–522. ACM, New York (2009)