Lexicon, Meaning Relations, and Semantic Networks

Prakash Mondal

Indian Institute of Technology Hyderabad
Kandi, Sangareddy, Telangana 502285 , India
prakashmondal@iith.ac.in

Abstract. This paper proposes a fresh formulation of conceptually grounded meaning relations by way of construction of certain well-defined relations over the lexicon of a natural language. These relations are constrained by the logical structures of linguistic meanings across sentence and discourse contexts. One of the biggest advantages of such meaning relations is that they are not defined over, or do not ride on, the syntactic structure of a given language. Nor do they turn on compositional relations for the computation of meaning values. This helps in the formulation of meaning relations to be defined on the symbolic elements of a lexicon on the one hand, and to be extracted from the surface structure of linguistic constructions on the other. This has consequences not merely for the nature of lexical meaning but also for the construction of a kind of (shallow) semantic networks that can be used for semantic processing in natural language understanding or machine translation systems that are driven by a kind of shallow processing of linguistic meanings. Thus, this paper aims to show the usefulness of a kind of conceptually based characterization of linguistic meaning for its relevance to computational language processing.

Keywords: Lexicon, Meaning Relations, Semantic Networks, Shallow Processing

1 Introduction

Lexicon is the central knowledge base of linguistic meanings as meanings are primarily grounded in words. Any expansions or extensions of linguistic meanings ride on the constructions of larger structures out of the elements of the lexicon. That is, linguistic meanings grow out of the constructions of more complex meanings as larger constructions are built from words. From this perspective, meanings are thus a superposition of formulas for syntactic constructions derived from words. The need for semantic processing in natural language processing arises from an in-depth of representation of various levels of natural language and the extraction of relevant information about the meaning-bearing elements from linguistic structures. Systems that rely heavily on such information such as machine translation systems or natural language understanding systems must have a rich formalism or system of (statistical) techniques for semantic analysis. Needless to say, semantic analysis requires decoding lexical relations from the text either in a stochastic manner or in a rule-based format. That is why various lexical resources such as WordNet, FrameNet, VerbNet are often utilized for semantic processing as an appropriate representation of lexical relations feeds into richer semantic descriptions. In this paper, we propose a new formulation of semantic description that can be extracted or read off from the surface structure of sentences. The semantic description defines meaning relations over the lexicon of a language, and since this description is not defined over compositional functions in syntax or semantics, such relations are amenable to extraction in both statistical or rule-based formats. Since the annotation and construction of deep semantic representations in any kind of semantic treebank is costly and heavily resource-consuming, the extraction of meaning relations for (shallow) semantic processing can be an optimal alternative to the extraction of semantic representations from semantic treebanks that incorporate rich semantic structures.

Before we proceed to formulate the new semantic description that can be exploited by various NLP applications, something about the nature of systematicity in relation to the lexicon needs to be said because this can help gauge the expressive power of the new formalism proposed here.

2 Lexicon and Systematicity

The lexicon of a natural language contains all lexical items, that is, words. In a certain sense, the lexicon of any natural language is the stock of idiosyncratic and irregular pieces of information [1]. However, this does not necessarily mean that the lexicon is not rule-governed. A number of morphological and syntactic regularities can be expressed even within the lexicon of a language. Idioms such as 'kick the bucket', 'beat about the bush' are whole chunks that are syntactically constructed but not semantically composed. Even though such idioms have to be listed in the lexicon anyway, they conform to syntactic rules since we do not have 'the bush about beat' instead of 'beat about the bush'. Similarly, the past and past participle forms of verbs such as 'sing', 'ring' have to be listed as such in the lexicon, but they have a rule-like commonality in their forms too. This is also true of the pair: 'breed' and 'bleed'.

Nevertheless, there is no denying that the lexicon is the least abstract system within the complex ensemble of linguistic systems including syntax, semantics, morphology and phonology. If any linguistic system that is closest to cultural conventions, contingencies of language use and the world out there, it is the lexicon of a natural language. Plus each person may know a few thousand words, but the entire lexicon of a language cannot be said to be located within the confinements of one's brain. Rather, the lexicon of a language resides in the collective inter-subjective memory of a linguistic community (see for a related discussion, [2]). Significantly, other linguistic systems (that is, syntax, semantics, morphology and phonology) tend to be systematic and hence axiomatic in nature due to the definability of rules over their domains. On the other hand, most of what a lexicon contains has to be learnt item-wise with a smaller number of rules that may help streamline the learning of the semi-regular properties of the lexicon (as discussed in the paragraph just above). What this indicates is that the lexicon of a language contains and specifies disparate pieces of information that incorporate and integrate features of phonological, syntactic, semantic and possibly pragmatic properties. In a sense, the structure of the lexicon is such that every lexical item in it is the source of and generates a certain amount of information. This makes the lexicon a kind of repository with a hugely diverse variety of lexical objects each containing an amount of information, insofar as each lexical item encapsulates a certain amount of uncertainty with an accompanying probability. It is this form of the lexicon that makes it, at least in part, insensitive to the formulation of rules by way of induction. That is, once we formulate a rule for a certain domain in syntax or semantics or even in phonology, we infer that the rule in question will apply to the whole range of items under the domain. For we do not have to check all items one by one to verify whether the rule formulated applies to each one of them. The axiomatic character of rules helps minimize the enormity of information by taming it only within the bounds of the symbols of the axiom(s) concerned. On the other hand, the lexicon of a language is full of surprises as one learns the vocabulary of a language.

2.1 A Formalism of Meaning Relations

Let's suppose that the lexicon of a language is a finite set of lexical items (L^i). Let this set be *Lex*= $\{L^i_{1,} \dots , L^i_n\}$[1]. Here, the indices *1…n* in the set *Lex* are indices of information contained in or generated from each of the lexical items in the lexicon of a language. The indices do not necessarily impose any order on the lexical items in *Lex* which are, by definition, unordered. Rather, the indices are countable indicators or *trackers*—stable or dynamic—of the amount of information each lexical item as a symbol or as a signal generates or carries. Since the lexical items in *Lex* are not just placeholders, the information contained in or generated from each of the lexical items in *Lex* may be rooted in phonological, syntactic, morphological and semantic features. This is not certainly to deny that such information may also have links to perceptions, actions, language use and properties of the world out there. Whatever the grounding for information in each lexical item is, this information must be understood in its information-theoretic sense. The informational indices of lexical items in all the examples throughout this paper ought to be interpreted to be present when specific relations are formed, although the indices have been omitted here for notational simplicity.

What about linguistic meaning—whether lexical or phrasal or even discoursal? Let's now assume that any linguistic meaning that can be constructed from a combination of the lexical items in *Lex* can be characterized in terms of some relation(s) drawn from among infinitely many relations defined on $\{Lex \cup R_1, \dots , R_k\}$, given that $R_1, \dots , R_k \subset Lex \times Lex$, where k is an arbitrary number. Hence these infinitely many relations have the form $R_1, \dots , R_k, R_{k+1}, \dots , R_\infty$, where $R_{k+1}, \dots , R_\infty$ are high-

[1] It appears that having a set for lexical items may create a problem for languages like Chinese, Japanese and Korean, since there does not exist in these languages any one-to-one correspondence between logographic characters and words, and such characters, often equivalent to single morphemes, seamlessly come together to form words and phrases. However, what matters for us is not how the characters in such languages can be defined to form words; rather, the possibility of having *discrete* word-like entities by imposing a certain organization—conceptual or otherwise—on the string of characters is all that matters.

er-order relations. Thus, whatever $R_1, \dots, R_k$ are constructed on *Lex* form a union with *Lex* itself. An example can make this much clearer. For instance, if we want to construct a meaning of the phrase 'an unusually brilliant scientist' from the lexicon of English, the lexical items 'an', 'unusually', 'brilliant' and 'scientist' from *Lex* can be related to one another in terms of meaning relations among the three lexical items[2]. Thus, one meaning relation obtains between 'an' and 'scientist'; a meaning relation between 'unusually' and 'brilliant'; one between the relation for 'unusually' and 'brilliant' together and 'scientist', and a *second-order* relation between 'an' and a meaning relation for 'unusually brilliant scientist'. Each of these relations R_i will have the form $R_i= \{(x_1, y_1), \dots, (x_n, y_n)\}$, where $n \geq 1$ and either x or y can be a relation. Meaning relations are thus distinct from lexical relations such as synonymy, polysemy etc. because they cover only one dimension such as that of semantic similarity or differentiated form-meaning mappings (one-to-many or many to one). Meaning relations can cover many dimensions of conceptual associations such as actor-action relations, theme-theme relation, agent-location relation, state-state relation, predicate-argument relations etc. etc. It needs to be stressed that the defining of meaning relations on *Lex* does not have anything to do, in a direct way, with the way syntactic relations, and for that matter, semantic compositions, are defined on the hierarchy of a tree, although in certain cases a meaning relation may well correspond to the way lexical items are syntactically combined. For example, the meaning relation between 'an' and 'scientist' does not form any syntactically defined constituent, yet it constitutes a meaning relation. In a nutshell, meaning relations are those relations that constitute conceptually viable elaborations or associations of linguistic contents expressed in words. In other words, meaning relations are those that are instantiated by conceptually constrained associations of a set of given words. The nature of the relevant conceptual constraints will be explicated below as we look at the illustration of meaning relation with a number of appropriate examples.

It is true that an *intrinsically* non-compositional approach to linguistic meaning can be specified in such a way that the elements of meanings have correlates in semantically composed expressions, as in models of linguistic meanings in which meanings are defined in terms of vectors defined over the features of words or combinations of words (see [4], [5]). Crucially, the present formulation of meaning relations, which may look similar to these models in having nothing to do with compositionality per se, is way distinct from these models of meaning, in that meaning relations are here conceptually constrained, regardless of whether features of words match or not. Hence 'an', and 'unusually' in the phrase 'an unusually brilliant scientist' do not form a meaning relation precisely because this relation is conceptually vacuous. One may now wonder what sense one can make of the notion of a relation being conceptually vacuous. One way of determining whether or not some relation is conceptually vacu-

[2] Note that this notion of relation is way different from the relations that can be constructed, as in model-theoretic syntax, for nodes in a tree (such as precedence or dominance relations) and for categories such as NP (Noun Phrase), VP (Verb Phrase), S (Sentence) etc. which are properties of nodes (see for details, [3]). In fact, the relations $R_1, \dots, R_k, R_{k+1}, \dots, R_\infty$ encompass many *dimensions* (such as string adjacency, precedence, dominance and parent-of relations etc.) in terms of which linguistic constructions can be characterized.

ous can be offered. In order to check this, one may look into the logic of the expressions concerned. In the present case, it may be noted that the determiner 'an' specifies the content of a nominal expression whereas an adverbial such as 'unusually' modifies an expression (an adjective or a sentence) and thus reduces the cardinality of the set of entities characterized by the expression concerned (for example, the set of individuals who are brilliant must be greater than the set of individuals who are unusually brilliant). When we form a relation between these two words, the resulting relation does not lead to a harmony in the logical structures of the words involved. Since logical structures of words can go beyond what the syntax of a language permits, meaning relations are ultimately grounded in logically possible *and yet* conceptually constrained relations between words. Similarly, in the sentence 'Rini knows that the man is sick' we cannot make a conceptually significant relation between 'that' and 'the' since their logical structures do not go on to make up a conceptually viable relation. Likewise, 'knows' and 'that' cannot form a conceptually viable relation because 'that' is a complementizer that logically acts as a function from propositional verbs to propositions, whereas 'knows' as a propositional verb relates individuals to things that are known. But 'knows' and 'sick' can form a non-compositional meaning relation specifying (someone's) knowing about the health status of someone (else). One simple way of constructing meaning relations is to think of them as *filler-gap* relations in which one entity in R must be/contain a gap that is filled in by a filler. Thus, in $R_i = \{(x_1, y_1), \ldots, (x_n, y_n)\}$ x or y must be/contain a gap and the other would be a filler. In the previous example, 'knows' contains a gap (requiring something which is known) which is filled in by 'sick', whereas 'knows' and 'that' cannot form a meaning relation because while 'knows' contains a gap, 'that' is not a filler because it contains gaps and hence is a gappy entity (in requiring propositions). Likewise, ‘an’, and ‘unusually’ in the phrase ‘an unusually brilliant scientist’ do not form a meaning relation because 'an' contains a gap but 'unusually' is also a gap (in being a modifier requiring an adjective or a noun). The notion of gap here is thus more general than is recognized in linguistic theory as it may encompass arguments, complements, predicates including verbs or even relations formed through a filler-gap link itself. Therefore, one rule of the thumb is that any meaning relation must contain a gap and (at least) a filler. Therefore, any arbitrary meaning relation R_i must be antisymmetric. We state this formally below.

The Antisymmetry Condition for Meaning Relations: Any arbitrary meaning relation R_i is antisymmetric because whenever $(x,y) \in R_i$ where *x* and *y* are either the gap and the filler or the filler and the gap, but since $(y, x) \notin R_i$ it follows that R_i is antisymmetric.

Therefore, relations that are not meaning relations on this generalization are usually symmetric. In sentences such as 'Think, plan and feel the difference' we may construct a meaning relation of co-occurring actions between 'think' and 'plan' where it may seem that both are actually gaps. But this is not the case because if the relevant relation is *action-action*, one action must be/contain a gap and hence in this context thinking as an action (gap) requires the action of planning (filler). In this sense, meaning relations are different from elementary predicates, as in *minimal recursion seman-*

tics [6], because each lexical item in minimal recursion semantics instantiates a relation, which is not the case in the present context since any meaning relation requires at least two lexical items even if they are identical in form. This can also be illustrated in the following manner.

$$(\text{think}(x), \text{plan}(x), \text{feel}(x, y)), (\text{the}(y), \text{difference}(y)) \quad (1)$$

Notice that each lexical item above introduces a predicate with its corresponding variable which is the argument of a number of elementary predicates. Also, the predicates 'think', 'plan' and 'feel' are syntactically coordinated in a group demarcated by the parentheses. Hence each conceptually viable relation within minimal recursion semantics is governed by predicate-argument relations, whereas in the current context a meaning relation has to merely fulfill the requirements of standing in an antisymmetric relation with or without any regard for syntactically governed compositions. A similar difference holds with respect to Discourse Representation Theory [7] as well. The following diagram shows this well for the same example sentence.

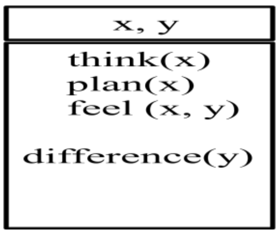

Fig. 1. A DRT representation of the sentence 'Think, plan and feel the difference'

The crucial difference between this type of representation and the current formulation of meaning relations is that inter-predicate relations is more expressively encoded in terms of meaning relations. In Discourse Representation Theory this has to be achieved only by means of linking via the same variable, which is *x* in Fig. 1. The relevant difference here is significant because the formalism of meaning relations is suitably grounded in the combinations of logically combinable concepts, whereas Discourse Representation Theory has arisen from the need to treat complex issues of variable binding in natural language. While variable binding has its own richness that is part and parcel of the system of links among clauses in natural language, concepts within and outside any clause bind together to make these resources available. In this sense, the motivation of Discourse Representation Theory and that of the current formulation are complementary with respect to each other. A corpus of Discourse Representation Theory-based semantic relations is thus particularly suited to the mapping of annotations to linguistic resources involving even event semantics that permit the use of efficient inference engines [8]. But the gamut of plausible conceptual linkages can be best exploited by the kind of meaning relations formulated in the current work, thereby providing a suitable ground for parallel annotations of semantic and conceptual possibilities offered by natural language texts.

From another related perspective, cases of idioms which seem to be built on syntactic combinations without any accompanying semantic composition can be cashed out in terms of a meaning relation that incorporates the relevant lexical items and builds the relation. So for an idiom 'face the music', for example, the lexical items 'face', 'the' and 'music' will be brought forward in order to construct a binary relation, say, $R_j = \{(X, (x, y, z))\}$, when X is the meaning of the whole idiom and x, y, z denote the individual lexical items ('face', 'the' and 'music'). Cases having syntactic items with no viable contribution to the meaning of a phrase/sentence can be dealt with in terms of the current notion of meaning relations, insofar as a sentence such as 'It rains' can be assigned the meaning relation which is constructed as $R_m= \{(X, (x, y))\}$, where X denotes the meaning of the whole sentence, and x, y designate the lexical items. It should also be noted that the relevant meaning relations are not simply dependency relations, as in Dependency Grammar [9]. For example, for a sentence like 'I have an enormous respect for this man' a meaning relation involving 'I' and a relation for 'an enormous respect', or 'I' and a relation for 'this man' can be constructed. Clearly these are not dependency relations because dependency relations themselves depend on constituency relations. Further, this account also differs from *frame semantics* [10] because the approach here is much more granular. For instance, a meaning relation between 'an' and 'scientist' will not be easily captured by a syntactic unit or by a frame, say, 'being a scientist' which is supposed to map the whole noun phrase and its grammatical function to the frame but 'an' and 'scientist' in 'an unusually brilliant scientist' are discontinuous frame elements. This issue is important also because some computational work involving abstract semantic representations has been done with the help of frame semantics [11]. In any case, what is important is that such relations as are proposed here in virtue of being independent of syntactic relations provide the advantage of granular and also discontinuous conceptual associations from the lexicon. In fact, any relation in $R_1, \ldots, R_k$ can be practically constructed by means of a relevant definition of a relation on the subset of *Lex*. This gives us the desired flexibility in having any possible collection of relations.

3 Meaning Relations and their Extraction from Texts

In this section, we specify a relevant procedure for extracting meaning relations from some sample texts. This will help understand how meaning relations can be extracted from real domains of language use and further utilized for natural language processing applications. To this end, we can use some real language data taken from newspapers, corpora, social media, texts on the Internet such as Wikipedia, or any other context of language use. For the purpose of our illustration, we provide below a sample text taken from the New York Times. The relevant portion is produced below.

'The study, published in June in Wildlife Monographs, suggests that when the Alaskan authorities were limiting wolf populations outside the Yukon-Charley preserve, survival rates of wolves within the preserve were lower than usual'.

Once parts-of-speech (POS) tagging of the text is completed, a number of relevant meaning relations can be extracted from the text above. Thus, we can have R1={(study, published)}, R2={(published, R3)}, R3={(in, June)}, R4={(study, Wildlife Monographs)}, R5= {(limiting, wolf populations)}, R6={(Alaska Authorities, R5)}, R7 = {(wolf populations, R8)}, R8={(outside, Yukon-Charley preserve)}, R9= {(when, R6, R7)}, R10={(survival rates, R11)}, R11={of, wolves)}, R12={(R10, R13)}, R13={(lower, R14)}, R14={(than, usual)}, R15={(when, R12)}, R16={(R1, suggests, R17)}, R17={(R9, R12)}. As can be noted above, in many cases meaning relations can allow for self-embedding. What is important to note here is that the exact names of the relations can be flexibly determined by finding out the logical structure of a given relation. For example, R3 can be named 'time of month' or even 'time'; R6 can be named 'agent-action' or something like 'actor-action'. A complex relation such as R17 can be named 'event co-occurrence' or 'event simultaneity'. Now that the POS tags can be made available, meaning relations that are not viable can be easily ruled out. No meaning relation can be constructed between 'when' (a subordinating conjunction) and the' (a determiner), for example, just as there cannot be any meaning relation between 'outside' (a preposition) and 'the' (a determiner). But one may now wonder how this can stop 'preserve' and 'lower' from being in a meaning relation since it is possible to have them in a conceptual association although the sentence above does not say anything about the preserve being lower. This is indeed a viable meaning relation. But, since this meaning relation is constrained by the syntactic structure of the sentence, this can be ruled out once the preposition phrase 'within the preserve' is analyzed as part of the noun phrase 'wolves'. In fact, a meaning relation between 'wolves' and 'preserve' may also be constructed so as to block, by being guided by syntactically parsed structures, such meaning relations from being extracted from certain texts. So it seems that in this particular case the compositional relations in syntax help extract the appropriate meaning relation, but note that this is not necessary for meaning relations to be defined. Although meaning relations can be governed by the compositional relations in syntax, they are not in themselves compositional relations. Consider the following sentence now. This is also an excerpt from the New York Times.

'The findings highlight the notion that managing wildlife within human-imposed boundaries requires communication and cooperation with the authorities beyond a preserve's boundaries, and could have implications for wildlife management programs elsewhere'.

The meaning relation between 'findings' and 'implications', for example, does not ride on syntactic composition since the two nouns are parts of different verb phrases joined by 'and' which share the same subject as far as the surface structure is concerned. This is so simply because even when we posit ' (the) findings' as the implicit subject for the second verb phrase 'could have implications for wildlife management programs elsewhere', the verb phrase is formed by combining 'could have' and 'implications' and the associated adjunct (the prepositional phrase 'for wildlife management programs elsewhere'). The postulated hidden subject engages in a compositional relation with the entire verb phrase thus formed, *but not with* the complement 'implications'. Similarly, a complex meaning relation between 'managing', 'requires' and a relation between 'communication' and 'cooperation' is also possible even though they are actually fragments lifted from the compositionally structured phrases.

It is now easy to show the meaning relations R1-R17 for the first example sentence in a graphical format in terms of semantic networks. These semantic networks can serve to mark the skeletal structure of basic semantic relations that can be read off from the surface structures of sentences in a text. Only in this sense can these networks be deemed to be shallow semantic networks.

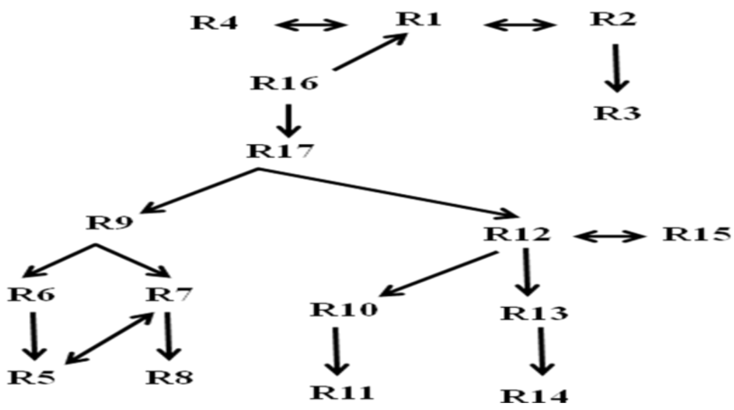

Fig. 2. A semantic network of meaning relations

Here, the bidirectional arrows designate relationships between two or more meaning relations that do not in themselves involve embedding of meaning relations, whereas the unidirectional arrows represent relationships from the embedding meaning relation to the meaning relation that is embedded inside. While major semantic treebanks such as FrameNet [12], PropBank [13] etc. extract schemas of event structures, predicate-argument relations and also semantic roles of arguments (especially in PropoBank), these semantic representations are either syntactically constrained or discourse-governed (especially in FrameNet). Syntactic constraints are useful for the characterization of the basic syntactic format of constructions, whereas discourse constraints are relevant to the construction of event types or scenes within which the syntactic structures come to mesh with the semantic structures. The underlying semantic representation in such frameworks is centered around the verb, and it is more so in a resource like VerbNet [14]. But there are many possible conceptual associations and links among words in a given sentence that cannot be captured by verb-centered relations such as predicate-argument relations. The non-compositional meaning relations mentioned in the paragraph above can be suitable examples of such conceptual associations. This is not, however, to deny that the relations among words that incorporate properties of conceptual associations cannot be compositionally governed.

The Universal Conceptual Cognitive Annotation (UCCA) [15] is exactly such a framework of semantic representations which import rich semantic structures of sentences. But even here the labels such as Connectors, Centers, Relators, Functions etc. operate over syntactically governed constituent relations, not over relations that go beyond syntactic constituent structures. Just for instance, UCCA in a sentence such as

'The professor we met yesterday was amazing' will label 'professor' as the Center that does not introduce an event or state on its own (while a verb does so) and 'amazing' as the Center that does the same job as 'professor'. But the conceptual association between 'professor' and 'amazing' cannot be shown by any of these relational labels. The concept associated with the professor here is not linked to the concept associated with the quality or property of being amazing. A useful way of augmenting any semantic treebanks that utilize syntactic information to build semantic representations, however rich or poor they may be, is to adjoin semantic networks of meaning relations to the semantic representations wherever rich syntactic resources are missing or not easy to access. An augmented treebank of semantic representations of this kind can integrate syntactic, discourse-related and non-syntactic/non-compositional properties of semantic structures. Many non-trivial conceptual links captured in meaning relations can thus be ensconced in semantic networks which can be built by importing the labels from the UCCA, for example, and then building the relevant meaning relations among these labels. Likewise, the labels used in FrameNet or PropBank can be re-deployed to build a parallel resource consisting of semantic networks of meaning relations.

On the basis of such considerations, we propose the following procedure for extracting meaning relations from texts.

Fig. 3. The procedure for the extraction of meaning relations

If semantic networks of meaning relations are adjoined to the UCCA or PropBank representations, the first stage and also the second are obviated. Hence the following figure reflects how an augmented treebank can be constructed.

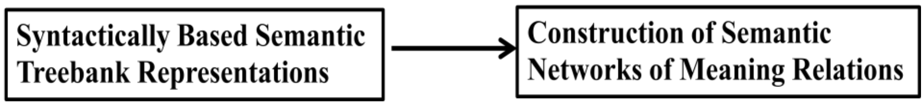

Fig. 4. The procedure for the construction of augmented treebank representations

One can thus formulate meaning relations freely so long as the ordering relation schematized above is obeyed. As a matter of fact, meaning relations can generally be represented in graphical formats as well, chiefly because meaning relations are, after all, conceptual units. The theory of conceptual graphs [16], [17] can help convert algebraic formats of meaning relations into graphical representations. An example can serve to illustrate this point further. Take, for instance, the meaning relation between 'findings' and 'notion'. The relevant relation can be named 'showing' or simply 'highlighting'.

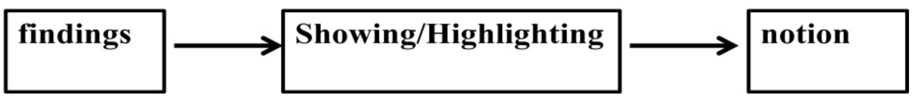

Fig. 5. A conceptual graph representation of a meaning relation

Thus, meaning relations can be reliably and appropriately organized in terms of graphs that have relations of partial order. The extracted meaning relations are in this sense associated with, and may be shown to be equivalent to, different kinds of lattice structures within which graphs instantiate relations of subsumption, homomorphism etc. We shall not try to show the details of how this can be done because of space considerations (and this is also outside the scope of this paper), but the underlying idea is that the relational properties of conceptual graphs can be shown to be orthogonal to the conceptual relations among meaning relations.

4 Meaning Relations and Shallow Semantic Processing: Implications for NLP Applications

The formalism of meaning relations elucidated here can have pivotal consequences for certain NLP systems such as machine translation systems, natural language understanding systems etc. The considerations laid out below are more or less general enough, though. The extraction of meaning relations can be very useful for NLP systems that do not dig too deep into linguistic structures. It is often the case that going deep into layers of semantic representation can be computationally costly. Besides, even when statistical techniques for machine learning from databases and corpora are used to draw out the required parameters of regularities in meaning construction, the resource bottleneck can hamper the effectiveness of any such techniques [18]. Resource-heavy operations in NLP systems have always been a big obstacle for machine learning whether supervised or semi-supervised. The problem is more acute especially for resource-poor languages, but with the proviso that some resources for even shallow syntactic parsing are required for the construction of meaning relations. The underlying idea is that the amount of resources required for both syntactic and semantic parsing strategies can be reduced only to the amount required for shallow syntactic parsing. Hence shallow semantic processing of linguistic structures can be more optimal in such situations. Since the extraction of meaning relations in a natural language turns on a sort of shallow parsing of linguistic structures, and meaning relations are themselves defined on the surface structure, NLP systems that can extract meaning relations from the text on the fly and then feed them into further levels of representation can avoid relying on lots of language resources of various kinds. Note that meaning relations entail shallow semantic processing also because meaning relations do not provide information about interpretative possibilities in vagueness and ambiguity, contextual salience effects, pragmatically governed constraints on meaning construction etc.

Beyond that, meaning relations being neutral with respect to statistical vs. rule-based systems, they can be easily accommodated within systems that are biased towards either type. The extraction of meaning relations is, of course, compatible with rule-based systems that apply parsing rules to decompose sentence structures. Additionally, the extraction of meaning relations can also be done on the basis of calculation of n-gram probabilities (especially bi-gram probabilities) of occurrences of words. One rule of the thumb is that if two words occur in a certain sequence more frequently, and if the word occurring second in the sequence has a higher bi-gram probability with respect to the first word, they will generally form a meaning relation. One example can illustrate this rule. For modifying adjectives appearing before or after a noun, if certain nouns appear more frequently after modifying adjectives (the word 'student' appearing in 'a/the brilliant student', for example), they can freely participate in a meaning relation with those modifying adjectives. This applies to post-nominal modifiers as well (such as 'the paper published in the journal...'). Even if two words are placed farther apart from each other and can potentially engage in a meaning relation (such as 'findings' and 'implications' in the preceding section), the joint probability of such words across phrases and/or clauses can be computed to predict the possibility of having a meaning relation between the two words. In fact, natural language has many such linguistic devices which can easily have us co-opt the policy of using a higher conditional probability of two words as a proxy for the prediction that they can entertain a meaning relation. Correlatives are a good example. Thus, whenever 'if' is found, the conditional probability of 'then' will be higher than that of any other conjunction. In such a case, the word 'if' along with the relation for the *if*-clause can always involve the relation for the *then*-clause in a complex meaning relation.

The advantages that machine translation systems and natural language understanding systems can accrue from the extraction of meaning relations are parallel. These systems can first extract meaning relations from the raw surface texts and then map them onto some level of ontology which can help re-code these meaning relations in another language, if necessary (as in machine translation), or which can feed into the construction of further meaning relations (as in natural language understanding). Any other post-processing level is not thus required if meaning relations are processed as part of the shallow semantic processing component in any NLP system. This needs to be further examined though.

5 Concluding Remarks

This paper has presented a new formalism of linguistic meanings, which is the formulation of meaning relations, and has also attempted to show its importance for semantic treebanks in natural language processing and beyond. Since it is based on the surface constituents of sentence structures, the hope is that it can be made part of shallow processing components of NLP systems that cannot heavily depend on scarcely available language resources. Further research can, of course, tell us whether meaning relations will really give us an edge. One limitation of the current proposal is that the construction of meaning relations cannot be done unless the text is parsed at least at a shallow level which requires a syntactic analysis to begin with, although meaning relations are not in themselves based on syntactic relations. But the proposal presented here emphasizes the role of semantic networks of meaning relations in augmenting, *rather than replacing*, the representations of popular semantic treebanks. This heavily weighs in favor of the viability of the construction of semantic networks of meaning relations to be attached to syntactically governed semantic representations in semantic treebanks.

References

1. Bloomfield, L.: Language. Henry Holt, New York (1933).
2. Mondal, P.: Can internalism and externalism be reconciled in a biological epistemology of language?. Biosemiotics, vol. 5, pp. 61-82 (2012).
3. Pullum, G. K.: The central question in comparative syntactic metatheory, Mind and Language, vol. 28(4), pp. 492-521 (2013).
4. Turney, P. D., Pantel, P. : From frequency to meaning: Vector space models of semantics. Journal of Artificial Intelligence Research, vol. 37, pp. 141-188 (2010).
5. Clark, S.: Vector space models of lexical meaning. In: Lappin, S. and Fox, C. (eds.) Handbook of Contemporary Semantic Theory. Oxford, Blackwell (2015).
6. Copestake, A., Flickinger, D., Sag, I., Pollard, C.: Minimal recursion semantics: An introduction. Journal of Research on Language and Computation, vol. 3, pp. 281-332 (2005).
7. Kamp, H, Reyle, U.: From discourse to logic: An introduction to model-theoretic semantics of natural language, formal logic and DRT. Kluwer, Dordrecht (1993).
8. Basile, V., Bos, J., Evang, K., Venhuizen, N.: Developing a large semantically annotated corpus. In: Proceedings of the Language Resources and Evaluation Conference, vol. 12, pp. 3196–3200 (2012).
9. Tesniére, L.: Eléments de syntaxe structurale. Klincksieck, Paris (1959).
10. Fillmore, C.: Frame semantics and the nature of language. In: Annals of the NewYork Academy of Sciences: Conference on the Origin and Development of Language and Speech, vol. 280, pp. 20-32 (1976).
11. Banarescu, L., Bonial, C., Cai, S., Georgescu, M., Griffitt, K., Hermjakob, K., Knight, K., Koehn, P., Palmer, M., Schneider, N.: Abstract meaning representation for sembanking. In: Proceedings of the 7th Linguistic Annotation Workshop and Interoperability with Discourse, pp. 178–186 (2013).
12. Ruppenhofer, J., Ellsworth, M., Petruck, M. R. L., Johnson, C. R., Baker, C. F., Scheffczyk, J.: FrameNet II: Extended Theory and Practice. The Berkeley FrameNet Project (2016).
13. Palmer, M., Gildea, D., Kingsbury, P.: The Proposition Bank: An annotated corpus of semantic roles. Computational Linguistics, vol. 31(1), pp.71–106 (2005).
14. Kipper, K., Korhonen, A., Ryant, N., Palmer, M.: A large-scale classification of English verbs. Language Resources and Evaluation, vol. 42, pp.21–40 (2008) .
15. Abend, O., Rappaport, A.: Universal conceptual cognitive annotation (UCCA). In: Proceedings of the Association for Computational Linguistics, http://aclweb.org/anthology/P/P13/ (2013).
16. Sowa, J. F.: Conceptual Structures: Information Processing in Mind and Machine. Addison-Wesley, Reading, MA (1984).
17. Sowa, J.F.: Knowledge Representation: Logical Philosophical, an Computational Foundations. :.Brooks Cole Publishing Co., Pacific Grove, CA (2000).
18. Manning, C. D., Schütze, H.: Foundations of Statistical Natural Language Processing, 6th edn.. MIT Press, Cambridge, MA. (1999).