An Integrative, Layered Approach to Lexical Semantics and its Application to Machine Translation

Inderjeet Mani

The MITRE Corporation, Z401
7525 Colshire Drive
McLean, VA 22102-3481
Phone: (703) 883-6149
mani@starbase.mitre.org

1. Introduction

One of the practical problems which stands in the way of the goal of developing a common lexicon for a language is the variation in granularities of meaning representation. These varying granularities have developed out of ecological needs. For example, in the field of machine translation (MT), researchers working on problems of translation divergences at the semantic level [Dorr 93a] [Dorr 93b], [Dorr and Voss 93] have tended to favor fairly abstract representations of meaning, such as Jackendoffian Conceptual Structures (CS) [Jackendoff 83], [Jackendoff 90]. Others (whether working in an "interlingual" paradigm, e.g., [Carbonell and Tomita 87], [Carlson and Nirenburg 90], [Meyer et al. 90], [Barnett et al. 94], or semantic transfer paradigm, e.g., [Nagao 87]) have depended on more concrete specifications of meaning, for example, based on microtheories of different domains, e.g., [Levin and Nirenburg 94]. The question then arises as to (1) whether these different granularities of meaning representation can be related to each other. In addition, there is (2) the problem of framework compatibility; in particular, between meaning representations more or less based on logic, with its emphasis on the semantics of reference and quantification, and the more conceptual structure-oriented approaches (like Jackendoff-inspired approaches) which focus mainly on a rich ontology of primitives and structures. In general, the logic-based representations appear to be especially useful in application to "functional" (or closed-class) categories associated with modality, mood, tense, aspect, conjunctions, determiners, etc., while the Jackendoffian approaches are particularly useful in connection with "lexical" (or open-class) categories.

This paper addresses these two problems by describing a representation for meaning which combines the advantages of both the logic-oriented and the conceptual structure oriented frameworks. The approach is layered, allowing different aspects or distinctions of meaning to appear at different levels. As will be illustrated here, in accommodating different granularities of meaning, it offers a basis for the development of a common lexicon.

2. The Approach: Integration

2.1 Representational Assumptions

We will begin with a semantic representation that is logic-based. We will assume a standard Davidsonian [Davidson 67] representation for reference to events. To simplify the presentation, we will use for the most part a uniform semantic treatment of modification within the Davidsonian approach, as sketched by [Parsons 90]. Our overall framework will be that of Discourse Representation Theory [Kamp 84], but this is merely a matter of convenience, related in part to its theoretical interest and in part to experience with using this framework in various NLP areas, including MT ([Barnett et al. 90], [Barnett et al. 94]).

To illustrate the approach, consider the sentence:

(1) A boy swam across a river.

In the Discourse Representation Structure (DRS) for this sentence, shown in Figure 1, certain letter variables stand for certain basic sorts: e1..en: events, x1..xn: things. The sorts indicate the ontological class of the object denoted by the variable. The terms contributed by the morpheme "swim" are in boldface.

e1 x1 x2
swimming(e1, x1, x2), past(e1),
boy(x1), river(x2), **water(x2)**, across(e1, x2)

Figure 1: DRS for "A boy swam across a river."

2.2 Decomposition

Now, assume we wish to represent the lexical semantics of "swim" at a slightly finer level of granularity, in terms of "<u>movement, across water, by means of limbs, caused by an animate agent</u>". The new terms contributed by this new

view of the morpheme "swim" are shown in boldface in Figure 2.

e1 e2 x1 x2 x3 **cause(e2, x1, e1), move(e1, x1, x2),** past(e1), boy(x1), **animate(x1),** water(x2), river(x2), **across(e1, x2),** limbs(x3), **with(e1, x3)**

Figure 2: DRS for "A boy swam across a river.", where "swim" is decomposed in terms of "move".

It is important to note that these DRS's can encode a decompositional representation such as Jackendoff's CS. This follows from the formalization of CS described by [Zwarts and Verkuyl 94], where the authors show how CS can be interpreted as a (many-sorted) first-order logic. Viewed as representational languages, the main difference between CS and typical logical representations for meaning is that CS has many different sorts of entities, with the variables and existential quantifiers which are explicit in logic made invisible in CS. To see the intuition behind this encoding, assume we view **moving** as a kind of **going**. Let sort variables j stand for spatial events, sort variables k stand for spatial paths, and sort variables w stand for instrumental positions. (Here we are using an extended set of Lexical Conceptual Structure (LCS) primitives inventoried in [Dorr 93a, p.170].) We then have the DRS in Figure 3:

j1 e1 x1 x2 x3 k1 w1 cause(e1, x1, j1), go(j1, x1, k1, w1), past(j1), boy(x1), water(x2), river(x2), via(k1, x2), limbs(x3), with(w1, x3)

Figure 3: DRS for CS-level decomposition for "A boy swam across a river."

A simple syntactic transformation on the DRS in Figure 3 yields the equivalent CS shown in Figure 4:

$[_{Event}$ CAUSE $([_{Thing\text{-}Animate}$ X1],
$[_{Event}$ $GO_{Spatial}$ $([_{Thing\text{-}Animate}$ X1],
$[_{Path}$ VIA$([_{Water}$ X2])])
$[_{Position}$ $WITH_{Instr}$ $([_{Limbs}$ X3])]])]

Figure 4: CS for "A boy swam across a river.", from DRS of Figure 3.

Note that the DRS in Figure 3 is also equivalent to the following expression in standard first-order logic:

∃j1 ∃e1 ∃x1 ∃x2 ∃x3 ∃k1 ∃w1 [cause(e1, x1, j1) & go(j1, x1, k1, w1) & past(j1) & boy(x1) & river(x2) & via(k1, x2) & limbs(x3) & with(w1, x3)]

Figure 5: First-order logic representation for "A boy swam across a river.", from DRS of Figure 3.

In general, for any sentence, a meaning representation in any one of these three formalisms (CS, DRT, and first-order logic) can be transformed into an equivalent representation in any of the others in which the sentence's meaning can be expressed. It is worth pointing out that while existential quantification is expressed in a straightforward manner in CS, the extent to which universal quantification can be expressed in CS is the subject of further research. Thus the DRS representation allows us to capture the quantifier scoping in sentences like (2), while the CS representation (at present) does not.

(2) Every boy who knew where the river was swam across it.

2.3 Using different ontologies

By using a many-sorted DRS in the manner illustrated, we have demonstrated how to encode the semantics of sentences like (1) at a variety of different levels of granularity. The CS approach turns out to be simply a particular level of granularity, with variables in the representation associated with a specific ontology, with sorts like spatial events (j variables), spatial paths (k variables), and instrumental positions (w variables). For a particular application, we could consider substituting a different ontology, for example, the PENMAN Upper Model [Bateman et al. 1990], an ontology for NLP used widely in generation and MT. For a sample representation in the PENMAN ontology, let the e variables be of the sort relational process, j variables of the sort motion process, x variables of the sort object, k variables of the sort spatial extent, and w variables of the sort instrumental, where all these sorts are appropriate concepts in the Penman ontology. We then have the same DRS as in the "CS" view in Figure 3, but now in

terms of the Penman upper model. Of course, in this case the mapping from the CS ontology to a different ontology was apparently not problematic; in practice we would expect conflicts[1].

2.4 Application

In the examples in Section 2.2, we showed how a given meaning could be decomposed into and expressed in rather different representations, without loss of meaning. In Section 2.3, we showed how a common representation of meaning could be tied to different ontologies. These suggest practical applications, in allowing MT systems to exploit the advantages of both logic-based representations, e.g., [Barnett et al. 94], and conceptual structure oriented representations, e.g., [Dorr 93a] [Dorr 93b], [Dorr and Voss 93]. For example, we can represent precisely the semantics of (3a) and (3b) in DRT (but not in CS).

(3a) Every boy swam across the Potomac.

(3b) Every boy crossed the Potomac by swimming.

Note that the DRS for (1) in Figure 2 could also serve as the semantics for sentences like (4)[2]:

(4) A boy crossed a river by swimming.

Similarly, by decomposing "swim" to a CS-level **go** in the DRS for (3a), we will end up with the decomposed DRS for sentences like (3b), which happens to be a form of the natural French translation of (3a). The difference between this pair of sentences (as well as (1) and (4)), termed translation divergences [Dorr 93a], [Dorr 93b], [Barnett et al. 94], reflects the fact that English typically allows incorporation (conflation) of manner (which has been decomposed out in our representations) into the verb, whereas French doesn't [Talmy 85][3]. By integrating these different lexical approaches in a common formalism, which can, further, be tied to different ontologies (e.g., CS, PENMAN), the MT system gains considerable power.

3 The Approach: Layering

Thus far, our discussion has dealt in general with sentential semantics. We now turn to the problem of lexical representation, in relation to different granularities of meaning. We will first introduce the idea of **views** in relation to decomposition, as a device for information hiding. We will then discuss the applications of this idea.

3.1 Decomposition and Views

The different levels of decomposition are analogous to different **views** of a concept in terms of its position in a subsumption hierarchy. Figure 6a shows the "minimal" view for the morpheme "swim" (i.e., the view corresponding to the terms for "swim" in Figure 1). For clarity of presentation, we used a more structured variant of a DRS encoded in a feature-structure formalism instead of a term-based formalism. This representation is in fact the language InL [Calder et al. 89], where "(i) every expression has a privileged variable (its index) and (ii) every variable is sorted, so indicating the ontological category of the object denoted by the variable" [Calder et al. 89, p. 234].

As shown in Figure 6a, the meaning of "swim" is expressed in the DRS of Figure 1 by means of a minimal view of the concept of swimming - with only certain structures being considered salient. These structures include required arguments. For sorts, only the most general sorts (event and thing) and important selectional restrictions are considered salient at this level.

[pred: swimming
index: $e1_{events}$
arg1: [index: $x1_{thing}$]
arg2: [index: $x2_{water}$]]

Figure 6a: View of "swim" (minimal representation)

Figure 6b shows an initial decomposition view of the same morpheme, where "swim" is decomposed to movement (as before). Here, we

[1]It is worth noting that the PENMAN ontology has since been merged (by hand) with the ONTOS [Carlson and Nirenburg 90] ontology, as reported by [Knight and Luk 94].

[2]We briefly discuss problems with claiming equivalence via decomposition in Section 4.

[3]Note that such language dependencies in meaning alternations, e.g., [Talmy 85], [Levin 93], etc., seem to support the idea of distinguishing language-dependent aspects of lexical meaning from world knowledge. See [Hobbs 87], [Kegl 87], [Pustejovsky 90] for various points of view on this.

augment the minimal view with other structures. In addition to the emergent cause structure, the "agent" argument (common to both events) becomes an animate being, a swimming event becomes a "move" event, and the optional "instrument" argument and direction feature are included.

[pred: cause
index: $e2_{causing\text{-}events}$
arg1: [index: $x1_{animate\text{-}being}$]
arg2: [pred: move
index: $e1_{events}$
arg1: [index: $x1_{animate\text{-}being}$]
arg2: [index: $x2_{water}$
direction: across]
arg3: [index: $x3_{limbs}$
means: with]]]

Figure 6b: View of "swim" decomposed to "move"

In Figure 6c, we show the LCS representation. Here the "move" event becomes a "go" spatial event, whose second and third arguments are "via" spatial paths and instrumental positions, respectively.

[pred: cause
index: $e_{causing\text{-}events}$
arg1: [index: $x1_{animate\text{-}being}$]
arg2: [pred: go
index: $j1_{go\text{-}spatial\text{-}events}$
arg1: [index: $x1_{animate\text{-}being}$]
arg2: [index: $k1_{spatial\text{-}path}$
pred: via
arg1: [index: $x2_{water}$]]
arg3: [index: $w1_{instr\text{-}pos}$
pred: with
arg1: [index: $x3_{limbs}$]]]]

Figure 6c: View of "swim" decomposed to LCS

Thus, at each level, different aspects of structure, including sortal changes, emerge and become salient.

3.2 Application

The layered representation of lexical meaning has several applications. First, the minimal representation helps to distinguish verbs which may have the same CS decomposition. Second, layering obviously provides for a more perspicuous representation (analogous to the use of a hierarchy of lexical templates). Third, the different layers could be associated with different lexical modules: for example, in MT applications, a minimal representation of meaning of "swim" used in a semantic-transfer oriented lexicon (with certain associated sortal restrictions on arguments) may be decomposed further into more elaborated representations of meaning used in more "interlingual" lexicons (with sortal restrictions associated with more elaborated ontologies). Fourth, it is likely that different levels of decomposition may be associated with different linguistic processes. For example, processes like ellipsis, as in (5), might look only at the minimal representation. On the other hand, an item in "implicit focus", like the limbs in (6), may be associated with a more decomposed representation.

(5) John swam across the river and Bill did so too.

(6) John swam across the river, even though his limbs were aching.

Finally, one might speculate that such a layered representation of morpheme meaning may be necessary for efficiency in natural language communication.

4 Decomposition and Problems of Generation and Translation

In practice, some of the terms in a more decomposed DRS may correspond to new morphemes which were not present in the minimal DRS. The DRS in Figure 2 could serve as the semantics (for the generation of) sentences like (7) and (8), in addition to (1) (repeated here):

(1) A boy swam across a river.

(7) A boy crossed a river by swimming.

(8) A boy caused his limbs to move through water in a river.

If the decomposition is meaning preserving (as it should be in principle), then the sentences corresponding to the original and decomposed DRS should be equivalent in meaning. As with many instances of decomposition (cf. the discussion of kill = cause-to-die [Lakoff 71], [Fodor 73], [Jackendoff 84], [Dowty 90], [Parsons 90]), it is unclear if these sentences are

in fact equivalent. However, this can, in principle, be tested by checking entailments.

Whether or not certain decompositions preserve meaning, a layered view of meaning suggests that different aspects of meaning become salient at different levels of decomposition. Indeed, sentences which are related by decomposition may differ in pragmatic properties, such as focus. For example, consider the alternations involved in a causative/inchoative decomposition. A given DRS could generate, depending on the level of decomposition:

(9a) Mary broke the vase.

(9b) Mary caused the vase to become broken.

(9c) Mary did something which caused the vase to get into a broken state.

The choice between these forms may depend on factors like focus, with the focus in (9b) being in general more likely to fall on the vase than on Mary.

4 Conclusion

This paper described a representation for lexical semantics which is able to integrate in a uniform way both logic-oriented aspects as well as structure-oriented aspects of meaning, which have traditionally been considered separately. Our representational tool was the language of many-sorted, indexed, first-order logic, embedded in a DRT framework. The representation has the potential to map to different ontologies. It is also layered, allowing for different linguistic processes like ellipsis and implicit focus to operate at different levels. In terms of these layers, minimal representations of lexical meaning appear to preserve what is explicitly mentioned in the utterance. Decomposed representations appear to differ (at least) pragmatically from their originals, but offer considerable advantages for grouping together divergences in form (whether across or within languages). We believe the framework described here helps to bridge the gap between various traditions in lexical semantics, and between various lexicographic approaches to MT.

Acknowledgements

I am grateful to Paul Portner for many insightful comments and suggestions on draft versions of this paper.

References

[Barnett et al. 90] Jim Barnett, Kevin Knight, Inderjeet Mani, and Elaine Rich. "Knowledge and Natural Language Processing", Communications of the ACM, pp. 50-71, August 1990.

[Barnett et al. 94] Jim Barnett, Inderjeet Mani, and Elaine Rich, "Reversible MT: What To Do When the Languages Don't Line Up". In T. Strzalkowski, ed., Reversible Grammar in Natural Language Processing, pp. 321-364, Kluwer, 1994.

[Bateman et al. 90] John Bateman, Robert Kasper, Johanna Moore, and Richard Whitney, "A General Organization of Knowledge for Natural Language Processing", ISI Technical Report, March 1990.

[Calder et al. 89] J. Calder, M. Reape, and H. Zeevat, "An Algorithm for Generation in Unification Categorial Grammar", Proceedings of the 4th Conference of the European Chapter of the ACL, pp. 233-240, Manchester, 10-12 April, 1989.

[Carbonell and Tomita 87] Jaime Carbonell and Masaru Tomita, "Knowledge Based Machine Translation: the CMU Approach". In Sergei Nirenburg, ed., "Machine Translation, Theoretical and Methodological Issues", Cambridge University Press, 1987.

[Carlson and Nirenburg 90] Lynn Carlson and Sergei Nirenburg, "World Modeling for NLP", Center for Machine Translation, Carnegie Mellon University, Tech Report CMU-CMT-90-121.

[Davidson 67] Donald Davidson, "The Logical Form of Action Sentences". In N. Rescher, ed. "The Logic of Decision and Action", University of Pittsburgh Press, Pittsburgh, 1967.

[Dorr 93a] Bonnie Dorr, "The Use of Lexical Semantics in Interlingual Machine Translation", Machine Translation 7, 135-193, 1992/3.

[Dorr 93b] Bonnie Dorr, "Interlingual Machine Translation: a parameterized approach", Artificial Intelligence 63 (1993) 429-492.

[Dorr and Voss 93] Bonnie Dorr and Clare Voss, "Machine Translation of Spatial Expressions: Defining the Relation between an Interlingua and a Knowledge Representation System", in Proceeedings of AAAI-93, Washington, DC, 1993.

[Dowty 90] David Dowty, "Word Meaning and Montague Grammar", D. Reidel, Boston, 1990.

[Hobbs 87] Jerry Hobbs, "World Knowledge and Word Meaning", Proceedings of TINLAP-3, Las Cruces, NM. 1987.

[Jackendoff 83] Ray Jackendoff, "Semantics and Cognition", MIT Press, Cambridge, MA, 1983.

[Jackendoff 90] Ray Jackendoff, "Semantic Structures", MIT Press, Cambridge, MA, 1990.

[Kamp 84] Hans Kamp, "A Theory of Truth and Semantic Representation", in J.A. G. Groenendijk, T.M. V. Janssen and M.J.B. Stokhof, eds., "Formal Methods in the Study of Language", Dordrecht, Foris, 1984.

[Kegl 87] Judy Kegl, "The Boundary Between Word Knowledge and World Knowledge", Proceedings of TINLAP-3, Las Cruces, NM, 1987.

[Knight and Luk 94] Kevin Knight and Steve K. Luk, "Building a Large-Scale Knowledge Base for Machine Translation", Proceeedings of AAAI-94, Seattle, 1994.

[Lakoff 71] George Lakoff, "On Generative Semantics". In D. Steinberg and L. Jakobovits, "Semantics: An Interdisciplinary Reader", Cambridge University Press, 1971.

[Levin 93] Beth Levin, "English Verb Classes and Alternations", University of Chicago Press, 1993.

[Levin and Nirenburg 94] Lori Levin and Sergei Nirenburg, "The Correct Place of Lexical Semantics in Interlingual MT", Proceedings of COLING-94, Kyoto, 1994.

[Meyer et al. 90] Ingrid Meyer, Boyan Onyshkevych, Lynn Carlson, "Lexicographic Principles and Design for Knowledge-Based Machine Translation", CMU-CMT-90-118.

[Nagao 87] Makoto Nagao, "Role of Structural Transformation in a Machine Translation System". In Sergei Nirenburg, ed., "Machine Translation, Theoretical and Methodological Issues", Cambridge University Press, 1987.

[Parsons 90] Terence Parsons, "Events in the Semantics of English", MIT Press, Cambridge, 1990.

[Pustejovsky 90] James Pustejovsky, "The Generative Lexicon", Computational Linguistics 17:4, 409-441.

[Talmy 85] Leonard Talmy, "Lexicalization Patterns". In T. Shopen, ed., "Language Typology and Syntactic Description", Volume III, Cambridge University Press, 1985.

[Zwarts and Verkuyl 94] "An Algebra of Conceptual Structure: An Investigation into Jackendoff's Conceptual Semantics". Linguistics and Philosophy 17: 1-28, 1994.