
A Theory of Granularity and its Application to Problems of Polysemy and Underspecification of Meaning

Inderjeet Mani

The MITRE Corporation, W640
11493 Sunset Hills Road
Reston, Virginia 22090, USA
imani@mitre.org

Abstract

Communication using natural language is remarkably efficient, by allowing reuse (through the use of generative devices) of a finite vocabulary to describe a potentially infinite set of situations. This vocabulary reuse contributes to words having many related senses (polysemy). Further, meanings can be relatively vague or precise; in other words, varying in their degree of specification of meaning. I suggest that these problems can be addressed by developing a knowledge representation which makes explicit the notion of granularity. As the grain size changes, we may fold certain distinctions, or split meanings more finely. In this paper, I formalize a theory of granularity and demonstrate how it can be applied to problems of meaning representation. Such a theory requires a world model which provides a rich sortal differentiation of entities based on the distinctions made by natural language, including the representation of meronymic structure and reification. Granularity will be represented in terms of structural operations defined as abstractions. I illustrate how this applies to problems of polysemy and vagueness in nominalizations, where splitting and folding of meanings are particularly evident.

1 INTRODUCTION

Natural language provides us humans with a very powerful means of describing the world as we experience it. Communication using natural language is remarkably efficient, by allowing reuse (through the use of generative devices) of a finite vocabulary to describe a potentially infinite set of situations. This vocab-

ulary reuse contributes to words having many related senses (polysemy). Further, meanings can be relatively vague or precise; in other words, varying in their degree of specification of meaning. If every aspect of a word's meaning were made explicit in everyday communication, communication would become very inefficient. These properties of natural language communication, polysemy and vagueness/underspecification, pose challenges for knowledge representation, in particular for lexical semantics, where the meanings of words have to be represented.

Lexical ambiguity in natural language can be viewed as falling into two disjoint categories: homonymy and polysemy. A lexically ambiguous word is polysemous iff it has multiple semantically related meanings. If a word has meanings which are semantically unrelated, those unrelated meanings are called homonymies. For an examples of polysemy, consider (1a-b). In (1a), *the bottle* refers to bottle as a container; in (1b) it refers to the contents.

(1a) Mary broke the bottle.

(1b) Mary finished the bottle.

For another example of polysemy, consider (2a), in which the nominalization is ambiguous at least between a process reading (2b) and a result reading of an object created (2c). (The term "nominalization" as used here means a noun or noun-headed constituent - such as a noun phrase - which can denote something event-like.) The result reading can be replaced by an ordinary count noun where one exists, but a count noun can't be used in the process reading.

(2a) The painting lasted three hours.

(2b) The painting/*picture lasted three hours, because they had to stop work that evening.

(2c) The painting/picture lasted three hours, because it was so flimsy.

In both (1) and (2), the shift in meaning involves some inferential path which relates the default meaning to the meaning expected by the predicate (e.g., “break”/“finish” in (1)). Clearly, the inferential paths involved in these examples encode some notion of perspective shift.

It is worth noting that perspective shifts can also result in information being lost. This is seen in each of (3a-c) (adapted from [Pustejovsky 95]), where the nominalization appears to denote both a process and a result object:

(3a) The house’s construction was finished in two months.

(3b) John’s construction of the roof frame for the house was done yesterday.

(3c) John’s construction of the roof frame for the house was done yesterday, and even though it took seven hours to get done, when it was done, it looked real good.

Such readings are underspecified with respect to whether we are dealing with a process or a result object.

There are numerous other cases where perspective shifts across basic ontological categories can occur, as when we shift from describing a *meeting* in terms of (say) a time interval to a description in terms of specific instants, as in (4a) (from [Pianesi and Varzi 96]), or from a description of a *pencil point* as a point to its being a surface, as in (4b) (from [Asher and Vieu 95]):

(4a) That’s how they met: *at a certain point*, John asked the waiter to invite her at his table; *the next moment* she was sitting in front of him.

(4b) The point of this pencil is actually an irregular surface with several peaks.

It is clear, then, that the examples in (1-4) all involve a shift in perspective. Shifts in perspective, in turn, reflect efficiencies in reasoning processes. Collectively, these examples suggest that it may be helpful to view efficiencies in language as a reflection of efficiencies in reasoning processes. Such efficiencies are very closely related to Hobbs’ notion of “granularity”: [Hobbs 85] suggests that in the course of reasoning we conceptualize the world at different levels of granularity, and that in a particular reasoning process we distinguish only those things that are relevant to that process. As [Hobbs 85] puts it, “our knowledge consists of a global theory together with a large number of relatively simple, idealized, grain-dependent, local theories, interrelated by articulation axioms.” In any given situation, a granularity is determined, allowing the local theory to be selected. Since the local theory has a smaller,

less complex ontology than one which would obtain if all predicates were always relevant, there is more hope of exploiting computationally tractable reasoning processes.

In what follows, I will show how problems of polysemy and underspecification can be addressed by a theory of reasoning which makes explicit the notion of granularity. The overall idea is that as the grain size changes, we may fold certain distinctions, or split meanings more finely. This idea complements existing work, providing a foundation for looking at problems in natural language semantics in terms of a theory of reasoning processes. In current approaches to semantics for natural language processing, cases of meaning shift like (1-2) are handled by some form of type coercion, involving, e.g., lexical rules defined over typed feature structures [Briscoe and Copestake 91], or [Pustejovsky 95] building it into the combination operations (e.g., function application) used to assemble phrases and sentences out of words. Thus, a system would have “container/contents” rules. The account I propose here does not seek to supplant these traditional rules; rather, I suggest that such rules are part of a more general inferential process involving granularity shifts. However, I go beyond previous approaches to address cases like (3) as well in terms of granularity shifts. Thus, in applying a theory of granularity to problems of lexical semantics, I will characterize the representational requirements in more abstract ways, providing a relatively problem-independent classification of the inferential procedures needed.

For reasons of space, I confine myself here to illustration of this idea with respect to problems of polysemy and underspecification in nominalizations, such as (2) and (3). This allows us to focus on the representation of event structure in the ontology, which allows us to make relatively precise certain key distinctions we have described informally, such as “event”, “process”, “result”, “process or result”, etc. The extension to (4a) will follow from this; the extension to (4b) requires some analysis of mereotopology (cf. [Asher and Vieu 95]), which we leave to future work.

In order to analyze the above phenomena, we have to first represent some of the distinctions which natural languages make of the world, which we take up in the next section.

2 ONTOLOGICAL DISTINCTIONS FOR NATURAL LANGUAGES

Following [Bach 86], let us define a join relation $\cup$ in R over the **entities** in D , so that if $q1$ and $q2$ are entities, $q1 \cup q2$ is also an entity. (For convenience, terms given

a definition in our ontology will be boldfaced, except when it's obvious that it's the technical sense that's intended.) The inclusion ordering relation $\preceq$, which defines a join semilattice over the set of entities, is:

Definition 1 (inclusion) $\forall(q1, q2) \ q1 \preceq q2 \equiv q1 \cup q2 = q2$

Among the sorts of entities are time **periods** and **individuals**. Instants will approximated as very short time periods (of course, alternative approaches, e.g., treating instants as primitive and defining intervals in terms of them, [Dowty 79], are clearly possible). Both **objects** and **eventualities** are subsorts of **individuals**.

Eventualities are further subdivided into the Vendler [Vendler 67] sorts of **states**, **processes** (activities), **accomplishments**, and **achievements** (the latter three subsorts will be called **events**). As a notational device, we will use sortal variables $q1, q2$ to stand for entities, $t1, t2$ for time periods, x, y for individuals, $v1, v2$ for eventualities, e for events, p for processes, s for states. **Processes** differ from other **events** in that if a process p (like John's walking) occurs in period t , it must also occur for every proper sub-period of t (this definition excludes gapped processes). **Accomplishments** are events e of which the primitive predicate *culminate*(e, t) can hold. (Intuitively, an event culminates in t if at the end of t the event has finished.) **Achievements** are instantaneous accomplishments (instants here are very short periods). Both accomplishments and achievements are defined so that if they occur over a time period t they may not occur in any other time period $t1$ such that $t1 \preceq t$.

In addition to representing **individuals** and time **periods** as subsorts of **entities**, we will also explicitly represent **sets** as another subsort of **entities**. This is needed, at the very least, to represent **propositions**. A proposition p is usually conceived of as the set of possible worlds where p is true. Here, a proposition will be represented as denoting a set of times, namely, the set of times in which the content of the proposition is true. As such, propositions can be viewed as abstract individuals representing sets, and we therefore include **abstract individuals** as a further primitive category. Facts (found in factive expressions such as *John admits that Mary came*, or *John admits to the fact that Mary came*) are simply true propositions (following [Zucchi 93]).

Associated with a set S will be a special kind of abstract individual called an **individual correlate** ${}^n S$ of the set, which is essentially an individual standing for that set. The concept of an abstract individual is relatively well established (e.g., [Chierchia 82],

[Link 83], [Hwang and Schubert 93], [Krifka 87], etc.). Summation operators have of course been discussed elsewhere in the literature, e.g., [Link 83], [Krifka 87], [Landman 89]. For reasons of space, I forego a comparison here with these earlier works (see [Mani 97b] for more discussion).

I now introduce one specialized inclusion relation:

Definition 2 (aspectual inclusion)

$\forall(v1, v2) \ v1 \preceq_w v2 \equiv v1 \preceq v2$ and $v1$ is an *aspectually salient part* of $v2$

By introducing the term “aspectually salient”, I am recognizing that the inclusion relation might be dense, in that an event may be broken down into an infinite number of parts, corresponding to the dense linear ordering of the real numbers. In some instances, we only want to consider subparts which relate to aspectual substructure of the event. Thus for example, an accomplishment event which is an event of building a house can be broken down into two aspectually salient parts, a preparatory process and a resulting state. The resulting state overlaps with the event, and can of course extend beyond it.

Now, I will define a ternary relation first described by [Pustejovsky 95]:

Definition 3 (exhaustive ordered part-of (*eop*))

$\forall(v1, v2, v3) \ eop(v1, v2, v3) \equiv$
 $v1 \preceq_w v3 \wedge v2 \preceq_w v3 \wedge event_meets(v1, v2)$
 $\wedge v3 = v1 \cup v2$
 $\wedge (\forall v4 \text{ if } v4 \preceq_w v3, \text{ then } v4 \preceq v1 \vee v4 \preceq v2)$

Thus for example, an accomplishment event $a1$ which is an event of building a house will have a preparatory process $p1$ and a resulting state $s1$ related by *eop*($p1, s1, a1$). For another example (from [Pustejovsky 95]) the event $e1$ of killing can be viewed as having a process (call it **killing-act**) $p1$ and a resulting state (call it **being-dead**) $s1$, related to $e1$ by *eop*($p1, s1, e1$).

The *eop* definition relies on the following auxiliary definitions (here **occurs** - for events - and **holds** -for states - are treated as primitives):

Definition 4 (happens) $\forall(v, t) \ happens(v, t) \equiv$

(i) v is an **event** and t is the maximal interval over which *occurs*(v, t) or (ii) v is a **state** and t is the maximal interval over which *holds*(v, t)

Definition 5 (event_meets) $\forall(v1, v2, t1, t2)$

$event_meets(v1, v2) \equiv$
 $happens(v1, t1) \wedge happens(v2, t2) \wedge meets(t1, t2)$




Figure 1: Ontological Distinctions

Here, for any $t1, t2$, $meets(t1, t2)$ iff $t1$ ends before $t2$ starts, with no other period in between. In somewhat more precise terms (derived in part from [Allen 81], [Allen 83], [Allen 84]):

Definition 6 (meets) $\forall(t1, t2) \text{ meets}(t1, t2) \equiv \exists(t1e, t2s) \text{ such that } t1 \text{ ends at } t1e \text{ and } t2 \text{ starts at } t2s \text{ and } t1e \text{ precedes } t2s \text{ and } \forall t3, (t3 \text{ precedes } t1e) \vee (t2s \text{ precedes } t3)$

Finally, we introduce an informal definition which says that x is the result object associated with an eventuality if it takes the patient role in the eventuality and it is around during the time of the result state associated with the eventuality:

Definition 7 (eventuality result-object (ero))
 $\forall(v, x) \text{ ero}(v, x) \equiv \exists(s, p, t2, t1) \text{ such that } happens(v, t2) \wedge eop(p, s, v) \wedge patient(v)(x) \wedge holds(s, t1) \wedge exists_at(x, t1)$

The sorts we have enumerated here are ordered by the sortal ordering relation $\subseteq$. All the sorts, along with $\subseteq$, are shown in Figure 1.

3 ANALYSIS USING TYPE-SHIFTING OPERATORS

Now that we have represented certain key distinctions in our ontology, we are in a position to analyze the linguistic examples we began with. As mentioned earlier, cases of meaning shift are treated in contemporary approaches as involving some form of type coercion. Prior to characterizing these in terms of granularity, however, I will first provide a type-shifting analysis based on the ontological distinctions introduced in the previous section.

Before launching into this account, let me state certain assumptions. I assume that the meanings of natural language sentences can be represented by logical expressions, each of which has a denotation in a world model. These logical expressions representing the meanings of sentences will be called “logical forms”. I will have to assume that the reader has a basic familiarity with the lambda calculus [Church 41], as well as the use of categorial grammars and type-theoretic representations for the semantics of natural languages (see [Dowty 79]). Finally, it may be helpful if the reader has some familiarity with Davidsonian semantics [Davidson 67], [Davidson 69], [Parsons 90], though it is not essential.

Now, let’s return to (2) above. I will assume that there

is a basic reading of a nominalization corresponding to a $\overline{N}$ constituent (i.e., a noun-headed constituent which combines with a determiner like *the*), as denoting a set of **events**. Thus, the noun *construction* has the logical form (5a), and the $\overline{N}$ *construction* has the logical form (5b), where the arguments are filled (i.e., “saturated”) from context.

$$(5a) \lambda e \lambda x \lambda y [constructing(e) \wedge agent(e)(y) \wedge patient(e)(x)]$$

$$(5b) \lambda e \exists(x, y) [constructing(e) \wedge agent(e)(y) \wedge patient(e)(x)]$$

Now, in many situations, (5b) is a perfectly satisfactory representation of the meaning of (the $\overline{N}$) *construction*. Consider what happens when the grain size shifts, causing the meanings to be split more finely. We can define a type shifting operator which shifts the meaning from an event to a process:

$$(6) \chi \equiv \lambda Q \lambda p \exists(e, s)[Q(e) \wedge eop(p, s, e)]$$

Applying this operator to (5b), we get the logical form for the process reading (which we illustrated in (2b)), where the nominalization denotes a set of **processes**:

$$(7) \lambda p \exists(e, x, y)[constructing(e) \wedge agent(e)(y) \wedge patient(e)(x) \wedge eop(p, s, e)]$$

Likewise, we have a type-shifting operator θ :

$$(8) \theta \equiv \lambda Q \lambda z \exists e[Q(e) \wedge ero(e, z)]$$

θ extracts the result **object** meaning (as was illustrated in (2c)) when applied to (5b), and this is shown in (9):

$$(9) \lambda z \exists(e, x, y)[constructing(e) \wedge agent(e)(y) \wedge patient(e)(x) \wedge ero(e, z)]$$

Note that although there are three possible polysemies: **process**, result **object**, and result **state**, we never get a three-way polysemy in a single use: we either have process/result object (*the construction*), or we have process/result state (*the destruction*). Given a $\overline{N}$ expression like *destruction*, ω extracts the result state meaning when applied to (5b).

$$(10) \omega \equiv \lambda Q \lambda s \exists(e, p)[Q(e) \wedge eop(p, s, e)]$$

Next, we turn to the cases of underspecified readings of “process or result” in (3). Here, we define an underspecification operator:

$$(11) \mu \equiv \lambda Q \lambda e[e = {}^\cap \chi Q \cup {}^\cap \theta Q]$$

The result of applying μ to (5b) is (12a), which when expanded to (12b), reveals that it denotes a set of **individual correlates** of the set of **processes** and result **objects** denoted by the process and result object readings, respectively:

$$(12a) \mu \lambda e \exists(x, y)[constructing(e) \wedge agent(e)(y) \wedge patient(e)(x)]$$

$$(12b) \lambda e 1[e 1 = \lambda p \exists(e, s, x, y)[constructing(e) \wedge agent(e)(y) \wedge patient(e)(x) \wedge eop(p, s, e)] \cup \lambda z \exists(e, x, y)[constructing(e) \wedge agent(e)(y) \wedge patient(e)(x) \wedge ero(e, z)]]$$

Likewise, there is another underspecification operation which folds **processes** and result **states**:

$$(13) \vartheta \equiv \lambda Q \lambda e[e = {}^\cap \chi Q \cup {}^\cap \omega Q]$$

Finally, we have an underspecification operator which folds **events** and **propositions**:

$$(14) \xi \equiv \lambda Q \lambda e[e = {}^\cap the Q \cup {}^\cap comp Q]$$

where *the* $\equiv \lambda P \iota e[P(e)]$ maps one-place predicates to individuals, ι is the usual iota operator, and *comp* $\equiv \lambda P \exists e[P(e)]$ maps one place predicates to times.

It is worth noting that underspecification over individuals (such as individual processes or results) is but one of a variety of different forms of underspecification that have been studied in the literature, including underspecified scope representations [Reyle 93]. The approach to underspecification described here applies only to underspecification over individuals.

4 GRANULARITY

Now, I will try to make good on my promise to characterize type shifting operators in terms of more general inferential processes involving changes in granularity. In order to do so, however, I must first flesh out a theory of granularity. Once I have done that, I will return in the next section to the linguistic phenomena.

As mentioned earlier, [Hobbs 85] suggests that in the course of reasoning we conceptualize the world at different levels of granularity, and that in a particular reasoning process we distinguish only those things that are relevant to that process, making other things indistinguishable for all practical purposes. In particular, Hobbs defines an indistinguishability relation such that two domain entities x and y are indistinguishable iff for all relevant predicates p , $p(x)$ is true iff $p(y)$ is true. The particular level of granularity determines which predicates are relevant. Thus, a theory in which temperatures are distinguished to the nearest degree, can, via an indistinguishability relation, be mapped to one in which temperatures are distinguished to the nearest 10 degrees (e.g., *in the fifties*). In any given situation, a granularity is determined, allowing the local theory to be selected. When the grain size shifts, certain “articulation axioms” are used to map to another local theory.

In general, a mapping which induces a change in granularity can be considered a special case of an **abstraction** [Giunchiglia and Walsh 92], which we will consider here to be a total, surjective function from formal languages $\mathcal{L}1$ to $\mathcal{L}2$. I restrict myself here to abstractions where the source and target language are the same language $\mathcal{L}$ (though there will be lots of applications - outside the scope of this paper - where the languages are different formal languages). I will write $\mathcal{F}_\gamma^\Rightarrow$ to mean an abstraction which uses a function γ to map a source logical form to an equivalent (denotationally identical) target logical form, both drawn from the same language $\mathcal{L}$. Likewise, let $\mathcal{F}_\gamma^\Rightarrow$ (or $\mathcal{F}_\gamma^\Leftarrow$) mean an abstraction which maps a source logical form to an entailed (or entailing) target logical form. As a notational convention, I will drop the entailment superscript if it isn't specified. The notions of equality of abstractions (based on extensional equality of operators), identity, composition and inverse are similar to those defined in [Giunchiglia and Walsh 92] (cf. pp. 370-71), for reasons of space, these are excluded here.

As [Giunchiglia and Walsh 92] also point out, Hobbs' notion of granularity can be represented in terms of a **granularity abstraction** which we define here as:

Definition 8 (granularity abstraction) *An abstraction $\mathcal{F}_\gamma^\Rightarrow$ is a granularity abstraction iff*

- (i) $\mathcal{F}_\gamma^\Rightarrow$ maps individual constants in $\mathcal{L}$ to their equivalence class under the indistinguishability relation $\sim$ in $\mathcal{L}$. (Thus, for any individual x in $\mathcal{L}$, $\mathcal{F}_\gamma^\Rightarrow(x) = \kappa(x)$ where $\kappa(x) = \{y \text{ such that } x \sim y\}$.)
- and
- (ii) $\mathcal{F}_\gamma^\Rightarrow$ maps everything else, including the predicates in $\mathcal{L}$, to itself.

Note that granularity abstractions may be particularly prone to lose information, since elements which are indistinguishable (given the context) are collapsed.

For problems of natural language semantics (which [Giunchiglia and Walsh 92] do not address), it is important to establish whether applying an abstraction to a logical form which is then combined with another logical form is the same as applying the abstraction to the result of the combination of the two logical forms. If it is the same, then the abstraction preserves compositionality. I will call such an abstraction **endocentric**, defined as follows:

Definition 9 (endocentric abstraction)

Given logical forms α, β in $\mathcal{L}$, and a binary combination operator $$ which forms new logical forms, an abstraction $\mathcal{F}_\gamma$ is endocentric iff*

- (i) $\mathcal{F}_\gamma(\alpha * \beta) = \mathcal{F}_\gamma(\alpha) * \beta$
- or

- (ii) $\mathcal{F}_\gamma(\alpha * \beta) = \mathcal{F}_\gamma(\beta) * \alpha$

Among the abstractions we will introduce and exploit are **meronymic abstractions**:

Definition 10 (meronymic abstraction) *An*

abstraction $\mathcal{F}_\gamma$ is a meronymic abstraction iff $\forall \alpha$ in $\mathcal{L}$,

- (i) α denotes an entity $q1$, and $\mathcal{F}_\gamma(\alpha)$ denotes an entity $q2$ such that $q2 \preceq q1$

or

- (ii) α denotes a set of entities $Q1$, and $\mathcal{F}_\gamma(\alpha)$ denotes a set of entities $Q2$ such that $\forall q2$ in $Q2 \exists q1$ in $Q1$ such that $q2 \preceq q1$

Similarly, we can represent **sortal abstractions**, whose definition is identical to Definition 10, except that $\subseteq$ is used instead of $\preceq$. To distinguish meronymic from sortal abstractions, and to indicate their direction, we will introduce another subscript, e.g., $\preceq \mathcal{F}_\gamma$.

It is worth noting that a meronymic abstraction which maps an individual to its equivalence class under the indistinguishability relation $\sim_1$, defined below, is a granularity abstraction.

Definition 11 (indistinguishability under $\cap$)

$\forall (x, y) x \sim_1 y \equiv x \preceq^\cap \{x, y\} \text{ and } y \preceq^\cap \{x, y\}$

The nature of the entailment varies with the abstraction. For granularity abstractions, we have upward entailment $\Rightarrow$; for sortal abstractions $\subseteq \mathcal{F}$, we have downward entailment $\Leftarrow$. In other cases, the entailment, if any, depends on the specifics of the logical form and the abstraction operator.

Finally, we introduce a notion of **grain size**:

Definition 12 (relative grain size) $\forall (\alpha, \beta)$ in $\mathcal{L}$, α is **finer-grained** than β iff there is an abstraction $R\mathcal{F}_\gamma$ such that $R\mathcal{F}_\gamma(\alpha) = \beta$ and either

- (i) $R = \preceq$

or

- (ii) $R = \subseteq$

5 TYPE SHIFTING OPERATORS AND ABSTRACTIONS

We now return to the analyses of Section 3, but this time with a characterization of meaning shifts in terms of abstractions. We can define a family of meronymic abstractions $\preceq \mathcal{F}_\chi$, $\preceq \mathcal{F}_\theta$, and $\preceq \mathcal{F}_\omega$, where the operators χ , θ , and ω are as defined earlier in (6), (8) and (10). None of these abstractions are endocentric, as the associated operator takes an event-selecting predicate and

returns a process- (or result-object or result-state) selecting predicate. In other words, each of these three operators returns a functor which takes a more specific argument; and so, an abstraction using another operator which expects the same sort of argument as χ , θ , or ω won't be able to combine with any of the results yielded by these three abstractions.

The meronymic abstraction $\preceq \mathcal{F}_\mu^\rightarrow$ is endocentric, since the operator μ (defined in (11)) yields an individual which is more general than its input individual and so can combine with a wider variety of functors (e.g., both process and result-object selecting functors) than the input individual. Further, it is a granularity abstraction under the indistinguishability relation $\sim_1$. $\preceq \mathcal{F}_\mu^\rightarrow$ is part of a family of endocentric meronymic granularity abstractions which apply to nominalizations, including $\preceq \mathcal{F}_\theta^\rightarrow$ and $\preceq \mathcal{F}_\xi^\rightarrow$.

At this point, one may ask what has been gained by characterizing the type-shifting operators above as abstractions. We first address underspecification. It is useful to characterize underspecification over individuals as an abstraction instead of merely a type shift, because granularity seems clearly involved - in particular, the concept of **indistinguishability** (e.g., Definition 11), which is absent from type shifting, obviously applies.

The case for characterizing polysemy in terms of granularity is not very striking, but nevertheless there are intuitively appealing grounds for doing so. Let's say the grain size shifts, so that the discourse shifts from talking about an event to talking about a part of the event (causing, in our view, the meaning of polysemous nominalizations describing the event to become more fine-grained). Type shifts in themselves have no representation of the current "level of abstraction" of the discourse (and consequently of the reasoning processes), whereas the granularity characterization offers the concept of **grain size**. Tracking the grain size, in turn, can help constrain reasoning processes, as revealed in investigations of research on "task-oriented" dialog (e.g., [Grosz 77]).

The final reason for viewing these operations as abstractions is methodological. In thinking of them as abstractions and formalizing them, we are forced to specify a meta-theory of these operations, which can provide for an improved modularity in our reasoning systems. Such a meta-theory tells us how expressive the language for representing abstractions needs to be, provides a classification of abstraction, specifies what particular form abstractions can take, what their truth-conditional properties are, and what their effect on compositional structure is. Except for the first question, the others have been directly addressed

in this paper.

6 RELATED WORK

This formalization of abstractions originates in the work of [Giunchiglia and Walsh 92], who consider abstractions to be mappings between formal systems, which consist of a language L , inference rules D , and axioms W . The abstractions defined here correspond roughly to what they call " L/W invariant" abstractions, where the axioms are not distinguished from the well-formed formulas of the language. There are several major points of difference, however. First of all, the discussion in [Giunchiglia and Walsh 92] includes a syntactic classification of abstractions based on the structure of terms in the logical form language. For example, they distinguish Predicate Abstractions, which map predicate names, Domain Abstractions which map constants and function symbols, and Propositional Abstractions, where the number of arguments of a predicate can be decreased or increased. Such a syntactic classification is not useful for natural language logical forms unless the logical forms are canonicalized in some form, e.g., in the form of a Davidsonian representation where event predicate symbols are monadic. Second, our abstractions are of special interest to phenomena like natural language, where composition plays a critical role. Finally, we classify the set of abstractions in terms of particular types of inferential paths, including those involving $\subseteq$ and $\preceq$, making such abstractions very appropriate as inference mechanisms in ontologies.

A formalization of granularity was proposed by [Hobbs 85]. As we have pointed out earlier, following Giunchiglia and Walsh, Hobbs' notion of granularity shifts corresponds to a special class of granularity abstractions, some of which have been exemplified in this paper. There are also several other accounts of granularity shifts. [Euzenat 95] discusses granularity operators, which are mappings on relations in a ontology. In our case, abstractions have been defined as mappings among logical forms, with different kinds of abstractions being characterized in terms of their denotations. His account exploits complex relations which are disjunctions of other relations, which in our framework would correspond to **individual correlates** of the sets corresponding to the complex relations. Finally, in contrast to Euzenat's approach, which postulates a number of fundamental algebraic properties of granularity operators, our approach is based on having abstractions specify explicitly their truth-preserving properties.

Finally, there are references to the idea of granularity in other work. For example, [Pianesi and Varzi 96]

discuss degrees of temporal granularity in event structure. The abstraction which shifts the perspective of *meeting* from (say) time intervals to specific instants, as illustrated in (4a), is characterized in their work in terms of a “minimal divisor” on structures corresponding to sets of events, where temporal differences within the divisor are neglected. It may be characterized in our approach in terms of a **sortal abstraction**. [Asher and Vieu 95] discuss perspective shifts in spatial domains; the shift of perspective in (4b) of “a pencil point” from a point to a surface could be characterized in terms of abstractions, but it first requires some representation of mereotopology (cf. [Asher and Vieu 95]), which we will examine in future work.

7 PRACTICAL APPLICATIONS

The motivation for this work is not entirely theoretical. In binding a NLP system to an application, the meanings of natural language expressions have to be mapped in a variety of different ways. A particular module may choose to collapse certain distinctions or to fill in and elaborate more detail for a particular application program. Providing a theory characterizing these mappings can remove a degree of adhocness and task-dependence in the design of the semantics/pragmatics interface in a wide variety of different tasks involving NLP, e.g., the business of constructing database queries, filling templates, generating a translation into a different natural language, and translating commands into directives to a simulation [Moore et al 96]).

The goal of developing robust ontologies capable of being used for inference in a variety of applications remains a challenge, since the concepts which need to be represented in different domains vary greatly, and the merging of disparate ontologies can be a formidable task. Nevertheless, several trends suggest that we are making progress towards this goal. First, standards have been evolving for knowledge exchange (e.g., [Gruber 93]) among different ontologies, as well as layered architectures for developing ontologies (e.g., [HPKB]). Second, there has also been progress on standardization of ontologies [ANSI]. Third, progress has been made in terms of a meta-theory of primitives used in knowledge representation systems [Guarino 94a], [Guarino 94b]. Fourth, there have emerged several large-scale ontologies [Mikrokosmos], [CYC], where a significant degree of reuse across applications is the norm rather than the exception. Fifth, relatively fine-grained thesauri such as [WordNet] (see also [Euro-WordNet]) continue to be used in different applications, although their use is not

entirely problem-free. Finally, interesting techniques have been developed (e.g., [Knight and Luk 94]) for semi-automatic merging of diverse ontologies. All this suggests that some day in the foreseeable future we will have rich large-scale semantic resources available, that can be reused across applications.

When used for lexical semantics, such a resource has the potential to give rise to massive ambiguity, unless we can develop generative theories which allow for dynamic determination of the set of relevant meanings by constraining our inferential processes to certain sub-ontologies. Abstractions, defined as they are as ontological operations, provide a generative mechanism for deriving new granularity-shift-related meanings

This approach to granularity may be particularly useful in tying more fundamental concepts in an ontology to more domain-specific ones. For example, a natural language front-end to a battle simulation system (e.g., [Moore et al 96]) may map words like *attack* or *approach* to complex sequences of actions in the target simulation system (e.g., *Attack Checkpoint Charlie at 1500*). To reuse lexical information across multiple applications, it may be desirable to integrate a general English lexical semantic ontology (where *attack* may have several meanings, e.g., related to military or medical situations, etc.) with an ontology for the application (where *attack* maps to a particular battle simulation action). If the attack is to begin at a particular time (e.g., *Platoon 6, launch an attack at 0900 hours*), the event structure of the general meaning (i.e., that attacks are processes) is relevant. A simulation may treat attacks as instantaneous events, akin to achievements, and military units like Platoon 6 as primitive (point) objects, simplifying the specification of an attack, or it may be more fine-grained, treating attacks as involving a non-infinitesimal time interval, and units as having component objects, including soldiers, vehicles, etc. (which may in turn be realized as geometric shapes on a map). These mappings can be characterized as abstractions, which may be reused in different ways for an NL system which must talk to several different simulations (a not uncommon situation).

There are also other cases where this approach may bear particular fruit. For example, a particular (transfer-based) machine translation system (e.g., [Nagao 87]) may represent just enough word meaning in the source language to disambiguate its possible translations in a target language; a more interlingual approach focused on problems of translation divergences at the semantic level [Dorr 93a], [Dorr 93b], [Dorr et al 94] may use fairly abstract representations of meaning, such as Conceptual Structures (CS) [Jackendoff 83], [Jackendoff 90], [Jackendoff 91] (see

also [Zwarts and Verkuyl 94]). In applying my approach, a particular semantic representation of *swim* used in a semantic-transfer oriented lexicon might be decomposed into more “abstract” semantic representation used in more “interlingual” lexicons. We may define (as in [Mani 97b]) an endocentric abstraction which maps swimming into a particular kind of going event, which in turn allows certain path and manner incorporation translation divergences [Talmy 85], [Barnett et al. 94] (e.g., *swim across the river* ~ (French) *traverser la riviere a la nage*) to be resolved. Thus, the same semantic lexicon could be used for multiple theoretical approaches.

8 CONCLUSIONS

This paper has illustrated how a theory of reasoning processes based on the notion of granularity can be used as a foundation to analyze certain problems in natural language semantics. Efficiencies in language use, reflected in phenomena of polysemy and underspecification, are viewed as mirroring efficiencies in underlying reasoning processes. The overall idea is that we may fold certain distinctions, or split meanings more finely, as the grain size changes, in order to carry out the inferences needed to communicate. As such, the paper shows one way in which further synergy between the fields of knowledge representation and natural language semantics can be achieved. The synthesis of ideas presented here is intended to stimulate further discussion, and is fairly open-ended in scope. In the future, we expect to continue to develop the theory, and applying it to problems of lexical semantics in the domain of spatial relations, in particular, exploiting current work in mereotopological representations.

Acknowledgements

I am grateful to Leo Obrst and two anonymous referees for helpful comments on earlier drafts of this paper.

References

- [Allen 81] James Allen, “What’s Necessary to Hide?: Modeling Action Verbs”, Proceedings of the 19th ACL, Stanford, June 29-July 1, 1981.
- [Allen 83] James Allen, “Maintaining Knowledge About Temporal Intervals”, Communications of the ACM, Volume 26, Number 11, 1983.
- [Allen 84] James Allen, “Towards a General Theory of Action and Time”, Artificial Intelligence 23 (1984), 123-154.
- [ANSI] ANSI Ad Hoc Group on Ontology Standards, see <http://WWW-KSL.Stanford.edu/ontology/index.html>.
- [Asher and Vieu 95] N. Asher and L. Vieu, “Towards a Geometry of Common Sense: A Semantics and Complete Axiomatization of Mereotopology”, Proceedings of IJCAI-95, Montreal, 1995, pp. 846-852.
- [Bach 86] Emmon Bach, “The Algebra of Events”, Linguistics and Philosophy 9, 5-16.
- [Barnett et al. 94] J. Barnett, I. Mani, and E. Rich, “Reversible Machine Translation”, in T. Strzalkowski, ed., Reversible Grammar in Natural Language Processing, pp. 321-364, Kluwer Academic Publishers, 1994.
- [Briscoe and Copestake 91] E. H. Briscoe and A. Copestake, “Sense Extensions as Lexical Rules”, in D. Fass, E. Hinkelman and J. Martin, eds., Proceedings of the IJCAI Workshop on Computational Approaches to Non-Literary Language, Sydney, Australia, pp. 12-20.
- [Chierchia 82] Gennaro Chierchia, “Nominalization and Montague Grammar: A Semantics Without Types for Natural Languages”, Linguistics and Philosophy 5 (1982), 303-354.
- [Chierchia 84] Gennaro Chierchia, Topics in the Syntax and Semantics of Infinitives and Gerunds, Ph. Dissertation, University of Massachusetts, 1984.
- [Church 41] Alonzo Church, “The Calculi of Lambda-Conversion”, Annals of Mathematical Studies, Vol. 6, Princeton University Press, 1941.
- [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.
- [Davidson 69] Donald Davidson, “The Individuation of Events”. In N. Rescher, et al, eds. Essays in Honor of Carl G. Hempel, D. Reidel, Dordrecht, 1969.
- [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 Proceedings of AAAI-93, Washington, DC, 1993.
- [Dorr et al 94] Bonnie Dorr, Clare R. Voss, Eric Peterson, and Mike Kiker, “Concept- Based Lexical Selection”, Proceedings of the AAAI Fall Symposium on MT Lexicons, 1994.

- [Dowty 79] David Dowty, "Word Meaning and Montague Grammar", D. Reidel, Boston, 1979.
- [Euro-WordNet] Report on Euro WordNet, see <http://WWW-KSL.Stanford.EDU/ontology/eurowordnet.pdf>
- [Euzenat 95] Jerome Euzenat, An Algebraic Approach to Granularity in Qualitative Time and Space Representation, Proceedings of IJCAI-95.
- [Giunchiglia and Walsh 90] Fausto Giunchiglia and Toby Walsh, "Abstraction in AI", Proceedings of the Conference for the Society for the Study of Artificial Intelligence and Simulation of Behavior (AISB-90), 22-26.
- [Giunchiglia and Walsh 92] Fausto Giunchiglia and Toby Walsh, "A Theory of Abstraction", Artificial Intelligence, 57, 2-3, 323-390.
- [Grosz 77] B. Grosz, "The Representation and Use of Focus in Dialog Understanding", Stanford Research International Technical Note 151, Menlo Park, CA, July 1977.
- [Gruber 93] T. R. Gruber, A Translation Approach to Portable Ontology Specifications, Knowledge Acquisition, 5(2):199-220, 1993. Available at <ftp://ksl.stanford.edu/pub/knowledge-sharing/papers/>.
- [Guarino 94a] Guarino N. The Ontological Level. Invited paper Presented at IV Wittgenstein Symposium, Kirchberg, Austria, 1993. In R. Casati, B. Smith and G. White (eds.), Philosophy and the Cognitive Sciences, Vienna, Hlder-Pichler-Tempsky 1994.
- [Guarino 94b] Guarino N., Carrara M., Giarretta P. An Ontology of Meta-Level Categories. KR'94, Bonn, 1994.
- [Hobbs 85] Jerry Hobbs, "Granularity", Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI-85), 1985.
- [HPKB] DARPA High-Performance Knowledge Base Program, <http://www.teknowledge.com/HPKB/about/HPKBpipGoals.html>.
- [Hwang and Schubert 93] C. H. Hwang and L. K. Schubert, Episodic Logic: A Situational Logic for Natural Language Processing, in P. Aczel, D. Israel, Y. Katagirm and S. Peters, eds., Situation Theory and its Applications, Vol. 3, CSLI Lecture Notes No. 37, Center for the Study of Language and Information, Stanford University, 1993.
- [Jackendoff 83] Ray Jackendoff, "Semantics and Cognition", MIT Press, Cambridge, MA, 1983.
- [Jackendoff 90] Ray Jackendoff, "Semantic Structures", MIT Press, Cambridge, MA, 1990.
- [Jackendoff 91] Ray Jackendoff, "Parts and Boundaries", Cognition, 41, (1991) 9-45
- [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.
- [Knight and Luk 94] Kevin Knight and Steve K. Luk, "Building a Large-Scale Knowledge Base for Machine Translation", Proceedings of AAAI-94, Seattle, 1994.
- [Krifka 87] Manfred Krifka, "Nominal Reference and Temporal Constitution: Towards a Semantics of Quantity", in J.A. G. Groenendijk, M.J.B. Stokhof, and F. Veltman, eds., Proceedings of the 6th Amsterdam Colloquium, 1987, 153-173.
- [Landman 89] F. Landman, "Groups", Linguistics and Philosophy, 12, 1:559-605, II: 723-744.
- [CYC] Lenat, D. and Guha, R.V. Building Large Knowledge-based Systems: Representation and Inference in the CYC Project. Addison Wesley, 1990.
- [Link 83] G. Link, The Logical Analysis of Plurals and Mass Terms: A Lattice-Theoretic Approach, in R. Bauerle, C. Schwarze, and A. von Stechow, eds., Meaning, Use, and Interpretation of Language, Mouton, Berlin, 1983.
- [Mani 97a] Inderjeet Mani, "Underspecification in the Semantics of Nominalizations", Lexicalische Ambiguitaet und semantische Unterspezifikation, 1997 Annual Conference of the German Linguistic Society, Dusseldorf, February 26-28, 1997.
- [Mani 97b] Inderjeet Mani, "The Semantics of Nominalizations: A Study in Decompositional Semantics", Ph.D. Dissertation, Georgetown University, Washington, D.C., February 1997.
- [Mikrokosmos] Stephen Beale, Sergei Nirenburg and Kavi Mahesh, Semantic Analysis in The Mikrokosmos Machine Translation Project, Proc. Symposium on NLP, Kaset Sart University, Bangkok, Thailand, 1995.
- [WordNet] Miller, G., "WordNet: A Lexical Database for English", Communications of the ACM, 38, 11, pp. 39-41, 1995.
- [Moore et al 96] Robert Moore, John Dowding, Harry Bratt, J. Mark Gawron, Yonael Gorf, and Adam Cheyer, "CommandTalk: A Spoken-Language Interface for Battlefield Simulations", Proceedings of the 5th Conference on Applied Natural Language Processing, 31 March -3 April 1997.
- [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.
- [Pianesi and Varzi 96] Fabio Pianesi and Achille C. Varzi, Refining Temporal Reference in Event Structures, *Notre Dame Journal of Formal Logic*, 37, 1, Winter 1996.
- [Pustejovsky 95] James Pustejovsky, *The Generative Lexicon*, MIT Press, Cambridge, MA, 1995.
- [Reyle 93] Uwe Reyle, “Dealing with Ambiguities by Underspecification: Construction, Representation, and Deduction”, *Journal of Semantics*, 10, 123-179, 1993.
- [Talmy 85] Leonard Talmy, “Lexicalization Patterns”, in T. Shopen, ed., *Language Typology and Syntactic Description*, Volume III, Cambridge University Press, 1985.
- [Vendler 67] Zeno Vendler, *Linguistics in Philosophy*, Cornell University Press, Ithaca, 1967.
- [Zwarts and Verkuyl 94] “An Algebra of Conceptual Structure: An Investigation into Jackendoff’s Conceptual Semantics”. *Linguistics and Philosophy* 17: 1-28, 1994.
- [Zucchi 93] Alessandro Zucchi, *The Language of Propositions and Events*, Kluwer, 1993.