Ontology and Description in Computational Semantics: Models and Processes

Matthew Stone
Computer Science and Cognitive Science
Rutgers University
Piscataway NJ 08854-8019
mdstone@cs.rutgers.edu

Abstract

This paper argues that natural language meanings should be modeled as DESCRIBING an underlying domain ontology rather than CORRESPONDING to it. Theoretically, description provides an attractive explanation for the flexibility and context-sensitivity of language use in complex domains. Practically, description offers a convenient and general way to relax algorithmic assumptions so that natural language semantics can be connected to domain ontologies in less constrained ways. I use case studies in natural language generation and descriptions of categories to suggest how this general framework provides an attractive computational setting in which to pursue applied models of language interpretation.

1 Introduction

In linguistics and philosophy, it is easy enough to assume that the correspondence between words and concepts is straightforward, and fixed in advance by our knowledge of language. According to one celebrated anecdote, semanticists can model the meaning of life once and for all just as *life′*. Many computational processes also implicitly appeal to a simple, pre-established correspondence between domain concepts and linguistic meanings. Pipelined natural language generators assume that input representations can be translated to linguistic representations by staged, more or less deterministic processes, with little input from context [Reiter and Dale, 2000]. Conversely, pipelined understanding modules generally assume that linking language to context simply amounts to resolving the reference of pronouns and definite noun phrases, leaving the constituency of semantic representations otherwise unchanged for application reasoning [Allen *et al.*, 2001].

Such models break down when agents come to language interaction with the kinds of rich, language-independent domain knowledge now being developed in the knowledge representation community. Take any important concept, like the relationships of parts and wholes, for example, and you find, on closer inspection, a family of closely related but clearly distinguished concepts, each of which has a contribution to make to application reasoning [Artale *et al.*, 1996]. The richness of current domain ontologies calls for flexible links between language and concepts. In fact, human speakers don't come to language with a fixed correspondence between, say, verbs and actions, either; people use different verbs to describe the same action and the same verbs to describe different actions [Furnas *et al.*, 1987]. Instead, it seems that people link language to concepts creatively based on the task at hand, and interactively come to agree on this conceptualization over the course of each conversation [Brennan, 2000]. This flexibility has been emphasized in many linguistic models, such as Jackendoff's PREFERENCE SEMANTICS [1983] or Lakoff and colleagues' models of METAPHOR in language interpretation [1980]. It seems to originate in the genuine complexity and flexibility with which agents must understand the world to act successfully in it.

This paper motivates, proposes and explores an architecture for linking linguistic meanings to a domain ontology based on the framework of DESCRIPTION. A description is a linguistic expression whose interpretation accesses a domain representation drawn from context, the OBJECT of the description. While the semantics of a description is a linguistic representation, the interpretation of a description involves a resolution that can link a speaker's meaning up with domain representations ARBITRARILY—the only constraint is that the object of the description is a term in the domain representation. With a flexible ontology, including things like actions, events, kinds, possibilities, places, paths, and so forth, almost any linguistic constituent can be modeled as a description. Thus, within the framework of description, language interpretations can link up with the underlying domain in more general ways, and linguistic modules can be deployed more flexibly within dialogue systems. The framework of description thus serves as an attractive setting in which to draw on a sophisticated domain ontology to pursue natural dialogue with human users.

This paper begins by showing how description accounts for the context-senstivity of language interpretation. I propose a simple linguistic representation for descriptions, and show how such representations can contribute to our theoretical and practical expertise in connecting language use to domain ontologies. I close with some suggestions about ongoing and future work within this paradigm. The presentation here draws heavily on two longer papers of mine that will be published soon [Stone *et al.*, 2001; Stone, 2003], and I refer the interested reader there for further details.

2 Language, Context and Representation

2.1 The problem of interpretation

Language is general. For the countless open-ended situations in which we might find ourselves, language offers a rather small stock of words and constructions, with rough-and-ready meanings which we must fit together creatively to suit our needs. Yet the abstract information we can convey in language allows us almost effortlessly to advance our specific projects in specific contexts. Thus, consider example (1).

(1) I would like coffee.

By uttering (1), a speaker can request an action by the hearer and thereby coordinate with the hearer to advance goals in the world. In particular, you have no doubt imagined (1) as an instruction to surrender a mug of steaming liquid to the speaker's physical control—thinking of (1) as a response to waitstaff's question (2) perhaps.

(2) Would you like a drink with dessert?

The specific content that an utterance carries when used by a speaker on a particular occasion, as with (1) in answer to (2), is its INTERPRETATION. Interpretation in this sense is part of linguistic pragmatics. By contrast, I will reserve MEANING for the general semantic description of the utterance, as provided by the grammar, abstracting away from any use of the utterance in a particular context.

The meaning of (1) in this sense is much more general than this one interpretation at first suggests. Too see this, note that (1) works just as well as a request for a clerk to scoop out an ice-cream cone in the context established by (3a), as a request for a wealthy host to make certain arrangements with the household cook on the speaker's behalf in the context established by (3b), or as a request for a coach to assign the speaker to affiliation on a particular team in the context established by (3c).

(3) a. Which of these flavors do you want?
b. What will you have to drink tomorrow morning?
c. Will you program for Team Coke or Team Coffee?

Semantically, (1) specifies the desired action just by its result: through it, the speaker must come in some sense to possess something identified as coffee. (1) looks to the context for the object the speaker intends to get, the kind of possession the speaker requires, and the kind of action the speaker expects. Here the possibilities are as limitless as the world itself, and when we link language to context, we seem quite simply to be linking language up with our full shared understanding of our present situation. In other words, interpretation may be FORMALIZED in terms of parameters from the context, describing speaker, time, place and so forth as is sometimes done in formal semantics, following [Kaplan, 1989], for example; however, interlocutors' pragmatic reasoning must in fact CONSTRUCT this formal context [Thomason, 1999], drawing on interlocutors' shared physical environment, their information about one another, and their expectation for the interaction. See [Clark, 1996; Stalnaker, 1998; Bunt, 2000].

Overall, language engineering brings a natural emphasis on implementing constrained processes that derive utterance interpretation from meaning in context. I will refer to this as the problem of interpretation. Its key practical and methodological difficulty is to capture the connection between natural language utterances and the underlying ontology of our domain of discourse. This domain ontology is typically specified in advance of any linguistic application, so the further burden of knowledge representation is to identify, describe and formalize any additional background knowledge about this ontology that figures in the meanings and interpretations of utterances in the domain.

2.2 Contrasts: WSD and NLKR

The problem of interpretation might seem simply to recast the familiar problem of word-sense disambiguation (WSD) in light of the general challenge of specifying the content of natural language discourse in a formal knowledge representation (NLKR). The three problems are in fact closely interrelated. A specific solution to the problem of interpretation must help construct the domain representations that specify discourse content; at the same time, it can presuppose the results of a word-sense disambiguation module. However, there is good reason, both theoretical and practical, to treat the problem of interpretation independently from WSD and NLKR.

Word-sense disambiguation refers to the process of resolving grammatical ambiguity in the uses of words. The ambiguity that pervades language extends to the lexicon, and many words appear to have multiple meanings. A WSD module determines which of these meanings fit an utterance. WSD is a coherent and active field of computational linguistics; see [Edmonds, 2002] for a current view. Most practical techniques use simple surface features of sentences to classify tokens of words into one of a relatively small number of senses, as defined by a general resource, such as a dictionary or thesaurus. [Yarowsky, 1995] epitomizes the approach.

Such WSD modules do not connect application language with indepedent domain models. Rather, they aim to inform and guide deeper processes of interpretation. For example, in (1), a WSD module would, at best, indicate that *like* is used in its sense meaning **wish-for**, not its sense meaning **find-agreeable**, and that *coffee* is used in the sense meaning **coffee** the beverage, rather than its sense meaning the beans or the shrub. This gives the meaning that (1) shares across the contexts in (3), not the specific interpretation that any one context requires. For this reason, it is natural to view practical WSD as a tool that provides the input meaning which we must then interpret.

Ultimately, a focus on the problem of interpretation may lead to critiques of WSD research. Once we model the inference that derives specific interpretations thoroughly, we are free to start from abstract or underspecified meanings. Such an architecture may no longer analyze individual lexical items as having multiple meanings, so it may not require an explicit step of WSD. Such an architecture might even respond better to what we know about word meaning. For example, Kilgarriff [1997] has argued that the fundamental data for lexicology is uses of words in context—in effect, interpretations. Kilgarriff suggests, following [Sweetser, 1990], that our knowledge of language links words directly to these interpretations. On this hypothesis, word senses need not have any linguistic or psychological reality; they would just be analytic abstractions that theorists use to identify particular clusters of interpretation that are usefully treated together within the context of a particular task. It seems too premature to endorse this tendentious argument (or to reject it). But regardless of the role that WSD may play in future NLP systems, Kilgarriff's proposal will prove instructive because it suggests that we may leverage the methodology of lexicography in building resources that characterize interpretation.

For researchers who aim to characterize our knowledge of the commonsense world, natural language discourse has always provided a rich set of data. Natural language interpretations effortlessly draw on subtle and intricate characterizations of physical action in the world—see particularly [Crangle and Suppes, 1994; Webber *et al.*, 1995]. Interpretations highlight the assumptions about time, action and causality that we must also use to predict, explain and plan for our changing circumstances; see [Steedman, 1997]. And interpretations connect with enduring regularities in our interactions with one another: our goals, beliefs, relationships and choices; for seminal studies, see [Charniak, 1973; Hobbs, 1978; Schank, 1980]. It increasingly possible to explore formal models of inference in interpretation within uniform and expressive frameworks, such as simulation semantics [Narayanan, 1997]. (Or see the alternative approaches represented in [Iwańska and Shapiro, 2000].)

At heart, this tradition of NLKR explores domain representation. It forms its own area of research because the domain knowledge is rarely treated elsewhere to the depth required to analyze discourse. Each NLKR formalism offers a precise way to characterize the content of interpretations; thus each sets up a specific problem of reasoning from meaning to interpretation. However, we will explore techniques for linking meaning and interpretation that do not depend on the specific formalism for interpretation, but only on its general structure. Such models of interpretation crosscut formalisms for NLKR research. Of course, where researchers have given analyses of context-dependence within specific formalisms, we expect to find instances of more general techniques. Thus Steedman [1997] complements a formal analysis of temporal ontology and temporal relations in the situation calculus with an analysis of temporal reference that constructs temporal interpretations as descriptions of events. Similarly, Chang and colleagues [2002] complement a formal analysis of event interpretation in terms of dynamic simulation with a model of interpretation that analyzes FrameNet frames as evoking—or indeed, describing—schematic simulation models drawn from general background knowledge.

3 Description, Meaning and Ontology

I recognize DESCRIPTION as a fundamental tool for documenting, characterizing and reasoning about the domain-specific interpretations of application utterances. In general, we speakers use utterances to describe or portray the things we are interested in—things as we conceive of them, individuated abstractly through our ideas about real-world causation, function and intention. I will say that an utterance DESCRIBES the things that it links up with and says something about.[1] In an application domain, the best way to characterize the interpretation of an utterance is in terms of the elements from the domain that the utterance describes.

Description is so useful because we intuitively understand not just what an utterance describes in general, but also the individual things that specific words describe. For example, the word *coffee*, as it is used after (2), describes something, namely **coffee**: a kind of beverage, made from certain roasted and ground seeds and known for its stimulant qualities. (I will continue to use italics when typesetting words under analysis and boldface when typesetting references to things.) And after (3), the word *coffee* describes other things: **coffee-ice-cream**, a kind of frozen confection flavored with the beverage coffee; or **team-coffee**, a group that represents its collective identity in the beverage coffee and its stimulant effects. By allowing us to factor the links between language and the world down to the atoms of linguistic structure, description naturally sets the stage for models of interpretation that generalize productively to new utterances and new situations.

3.1 Documenting interpretation

The ubiquity of description means that documenting the interpretation of utterances principally involves annotating words and phrases with the things in domain that they are understood to describe. In practice, documenting the interpretation of utterances this way builds closely on given domain representations, as informed by more general investigations in system-building. To act or reason effectively, any system needs an ontology of individuals and concepts which frames its information about the world. When we understand utterances as descriptions, we take these same individuals and concepts as constituents of utterance interpretation. By allowing a system to link utterances directly to its world, we collapse the tasks of delineating the individuals and concepts the system will reason about and those it will speak about. This gives us less work to do. It also makes the work easier, because specifying interpretations directly in domain terms alleviates many of the complexities and ambiguities of more generic levels of annotation.

This summary may seem a bit glib. After all, methodology for semantic and pragmatic annotation is a contentious and active area of research, and is sure to remain one. Still, annotation projects that cover constrained domains with comprehensive and specific guidelines have been able to achieve high reliability. The HCRC map task methodology is a good example [Anderson *et al.*, 1991]; so is the annotation methodology that informs SENSEVAL evaluations [Kilgarriff, 1999]. By contrast, annotation tasks that call for abstract judgments,

[1] Note here that, following philosophers, I am distinguishing description from REFERENCE; philosophers define reference as an objective real-world connection that links our thoughts or words to some constituent of the real world. See [Neale, 1990] for reference or [van Deemter and Kibble, 2000] for the concomitant difficulties of coreference. Reference is thus much rarer than description; we can describe Santa Claus, but cannot refer to him!

not clearly informed by specific relevant guidelines, see much more frequent disagreement. For example, Vieira and Poesio [2000] enlisted annotators to judge the antecedents of referring expressions in unrestricted text. With no domain model, annotators judged many bridging (or inferrable) references to have indeterminate antecedents. Kilgarriff found similar indeterminacy when annotators assigned abstract word senses to capture the specific interpretations that adjectives get in combination with particular nouns [2001].

An intuitive explanation for the advantages of specificity in these experiences is not hard to find. Specific applications can provide a meaningful array of individuals and categories that match the interpretations language users must construct in context; and specific applications can establish functional standards for what counts as a correct annotation. In fact, I expect these advantages to increase with further research, not to diminish, as modeling interpretation increasingly calls for obtaining large amounts of training data and minimizing the effort of annotation.

3.2 Documenting meaning

At the same time, description provides an attractive infrastructure for building general models of application language. If we characterize interpretation in terms of description, we can represent semantics for our application through CONSTRAINTS on what words in utterances can describe. This allows us to formulate semantic generalizations that correspond one-to-one to the elements of interpretation and generalize directly over them. Developing these generalizations brings the typical challenge of knowledge representation—we must identify categories of interpretation, regiment these concepts into a single consistent perspective that determines the granularity and organization of our semantic vocabulary, and iteratively refine regularities in interpretation in these semantic terms. However, we can readily carry out this knowledge representation effort in close contact both with the linguistic literature and with hypothesis-testing and data-mining tools that help to frame, refine and validate candidate generalizations about application language.

One close analogue to this enterprise is the corpus-based methodology already widely practiced by lexicographers to cluster attested uses of words [Kilgarriff, 1997; 2001]. Lexicographers aim to describe the uses of a word concisely, correctly and usefully. This requires regimenting an extensive body of interpretations and a correspondingly extensive set of communicative conventions that govern that word. However, this does not necessarily require capturing generative principles that go beyond attested patterns of use; such uses are inevitably rare—even unpredictable. Thus, lexicographers work by examining attested uses of words, and iteratively organizing what they see into a constrained and comprehensive framework. General concepts and definitions emerge from this effort. This sort of procedure, and this sort of result, is characteristic of applied methodology for knowledge representation generally [Brachman *et al.*, 1990].

3.3 An example analysis

Consider our case with example (1). A theoretically-inspired model of this case might describe the use of *coffee* by the principle in (4a). And it might offer the principle in (4b) for *I* and those in (4c) and (4d) *would like*.

(4) a. *Coffee* can describe any kind of thing connected with the beverage **coffee**, and always does so.

b. *I* can and does describe the speaker of an utterance.

c. *Would like* can describe any kind of event that brings its grammatical subject into a certain kind of possession relation with its grammatical object, and *would like* always does so.

d. *Would like* expresses the subject's preference for such an event, and so describes a particular **state**.

These suggestions to some degree encode a range of insights from the linguistic literature. (4a) suggests Nunberg's theory of deferred reference [1979]; what definite noun phrases and other descriptive expressions contribute to interpretation is often an entity related to what they literally evoke. (4b) treats *I* as an indexical much as Kaplan might [1989]; *I* is interpreted by accessing specific entities from the utterance context. (Note that in wider domains deferred reference might also be appropriate for *I*!) Finally, (4c) and (4d) abstract away from any predefined link to the world in a way that recalls accounts of productive polysemy and metaphor [Lakoff and Johnson, 1980; Pustejovsky, 1991]. Observe that the meaning given in (4c) does not simply say that the subject will have the object; rather it suggests that the context will supply the relationship that results between subject and object, and uses possession as an abstract concept to characterize these possible interpretations.

Such intuitions about the interpretation and meanings of application language can be formalized straightforwardly. To formalize description, we can INDEX each syntactic constituent to link it with the domain elements that this constituent is intended to describe. We can then use variables to abstract away from the things any particular utterance of a sentence might describe, and construct a general schema for utterance interpretation:

(5)
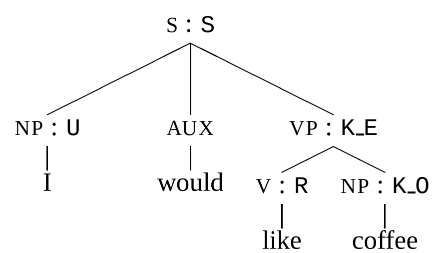

Here we implement the proposal in (4) by using a variable K_O for the kind of thing described by *coffee*, U for the speaker described by *I*, R for the possession described by *would like*, and K_E and S for the corresponding kind of event and state of preference.

Descriptive representations for computational semantics and pragmatics in some sense originate with the work of Hobbs and colleagues [1985; 1993]. Hobbs represents utterance meaning through "ontologically promiscuous" descriptions which represent a wide range of abstract individuals. Hobbs emphasizes that these representations make it easier to formulate sophisticated understanding processes, but Hobbs does not explore the possibility of using them to connect language to independent domain ontologies in a general but efficient way.

However, we can do this using constraints. The problem is to classify which values the variables in a structure like (5) should take. Constraints provide a way to describe this classification problem using formally-specified semantic concepts and generalizations. For example, the generalizations about meaning presented in (4) depend on a few key relationships. We can formalize these relationships through constraints on the variables in (5), as in (6).

(6) a. `connected(K_O, k_c)`, using a predicate `connected` to indicate that the thing described by `K_O` is connected by a close association with the beverage **coffee**.

b. `speaker(U)`, using a predicate `speaker` to indicate that `U` is the speaker of the utterance.

c. `possession(R)`, using a predicate `possession` to indicate that `R` is a kind of possession relation.

d. `result(K_E, holds(R, U, K_O))`, using an operator `result` to describe a view of the world which restricts attention just to the effects of `K_E`, and using a predicate `holds` to indicate that this view of the world finds `U` in `R` with `K_O`.

e. `preference(S, U, K_E)`, using a predicate `preference` to indicate that `S` represents `U`'s interest in seeing an event of kind `K_E`.

To formalize the semantics of (5), we pair it with the constraints on values of variables in (6).

With suitable additional representations, we should now be able to classify domain elements relevant to the context of utterance in semantically-meaningful terms. This gives a constraint-satisfaction model of interpretation. Such models have two crucial properties.

First and most importantly here, constraint-satisfaction models allow language to link up with a domain ontology flexibly. We can instantiate semantic variables to elements of our ontology arbitrarily as needed, so we can formulate general understanding and generation algorithms that abstract away from the specific formulation of domain knowledge, and so can be reused more generally to model language use in specific applications. For the theory of constraint-satisfaction see [Mackworth, 1987]; for applications in natural language, see [Mellish, 1985; Haddock, 1989; Stone and Webber, 1998; Schuler, 2001].

Second, constraint-satisfaction models provide a natural setting to resolve ambiguity. In context, each candidate instance of a constraint can be associated with a *weight* that quantifies the preference for resolving the constraint to that instance in context. (It may perhaps be correct to model this as a *probability*.) Such preferences are a standard ingredient of reference-resolution models [Grosz and Sidner, 1986; Lappin and Leass, 1994]; the generalization of this notion of preference to account for descriptive connections and pragmatic phenomena generally is due to [Hobbs *et al.*, 1993]. Describing the dynamics of these preferences in discourse remains a challenge, especially for formal models of dialogue [Stone and Thomason, 2002]. But a basic model is easy enough to implement—these preferences can simply privilege constraint-instances that have been used previously, and otherwise privilege constraint-instances that are seen frequently across application language. That basic model explains why the beverage interpretation of (1) is the default, while the interpretation changes in the other contexts of (3) to reuse the entities and instances introduced by the question.

4 Description in Models of Language

The significance of description is that it is not just limited to the familiar middle-sized physical objects of everyday experience. We can take linguistic descriptions to link up with domain representations of any type. Considering descriptions of KINDS provides a striking illustration of the flexible links between linguistic meaning and an underlying domain ontology. Complex as they are, descriptions of kinds can nevertheless be modeled in a descriptive framework, as naturally as descriptions of ordinary objects are.

In natural language, kind terms do not just supply predicates, and are not just used to characterize other things. Kinds are entities with their own distinctive properties. (7) illustrates this.

(7) a. Dodos are extinct.

b. Paper clips were invented in the early 1900s.

For (7a), the individual dodos are merely dead; only the species, **the-dodo**, is extinct. Likewise for (7b), individual paper clips are merely made; it is the kind of office supply, **the-paperclip**, that is invented. A classic source on descriptions of kinds is [Carlson, 1980]; see also [Carlson, 1992; Krifka *et al.*, 1995].

Once we represent kinds as things, to interpret sentences such as those in (7), we are free to exploit kinds elsewhere as well. Indeed, it seems that language always represents that domain individual **x** has domain property **p** as a relationship between autonomous elements, **x** and **p**, drawn from a background model of the world. So to say something has **p**, an utterance would always include a constituent that describes **p**. Using contextual parameters in place of predefined semantic translations is another common ingredient of natural language technology. For example, it is used in underspecification approaches to postpone resolutions of lexical ambiguity [Reyle, 1993; Pinkal, 1999]. But here we consider it as a general tool to mediate between semantic representations and a genuinely nonlinguistic ontology.

One reason for the approach is the general context-sensitivity of language—the properties by which we characterize domain individuals belong to that domain, not directly to language. For example, in different domains we might use *tape* in (8) to instruct our audience to use any of a variety of kinds of material.

(8) Cover the edge with some tape.

In office work, we might expect thin transparent tape, for the edge of a piece of paper; in craft work, we might expect masking tape, for the edge of a surface to be painted; or, in repair, we might expect duct tape, for the edge of a piece of metal or plastic. (The alternative interpretations of (1) make a similar point.)

It is easy enough to represent the facts about the world that make a particular kind of tape appropriate to a particular task. What is problematic is to draw these facts into the processes of language use in a regime that keeps understanding and generation symmetric and avoids open-ended reasoning in either. Modeling utterances as descriptions of domain kinds as well as domain objects solves this problem in a direct and straightforward way; generation and understanding processes then recover the described kinds by matching linguistic constraints against the context. In (8), then, we take *tape* to describe some kind of material from the context: specifically, the material is required to be a long thin fabric with a sticky side. Each context supplies its own kind of tape to satisfy this requirement.

Discourse connectivity provides an independent argument. Read (9) in the context established by (8).

(9) Isn't it the most versatile fastener ever invented?

As in (7b), the word *invented* in (9) signals that the utterance describes some kind, rather than some individual thing. Here that kind is picked out by the pronoun *it*. In the context of (8), *it* is naturally understood as whatever kind of tape we have just been instructed to use. On an analysis that implicates kind reference in attributions like those in (8)—but only on such an analysis—this interpretation of *it* in (9) is completely consistent with other uses of *it* [Webber, 1983]. In general, *it* can describe anything explicitly described in an earlier utterance, and little else.

Although language impels us to represent kinds pervasively in the interpretation of utterances, language provides little guidance about what kinds there are. In (8), for instance, language tells us that some kinds of things are *tape*. But domain knowledge actually supplies the specifics: there is this particular kind of transparent tape, this kind of masking tape, that kind of duct tape. Formalizing kinds in interpretation therefore involves investigating meaningful categories as constituents of the real world; this investigation draws on judgments about such real-world matters as causation and function, not judgments about language.

The literature of cognitive science calls our attention to three sorts of kinds in particular: NATURAL KINDS, ARTIFACT CATEGORIES and SOCIAL CATEGORIES. All of them crucially involve an explanatory understanding that puts us into contact with meaningful real-world categories. What distinguishes the different sorts of kinds is the character of the underlying causal properties and relationships that determine category membership. Things form a NATURAL KIND when they are created by the same general causal regularities in the natural world, and when they derive certain common characteristics from those regularities; see [Kripke, 1982; Putnam, 1975; Carlson, 1992]. Things form an ARTIFACT CATEGORY when they recognizably support a designated use in virtue of some effective structure, composition or design; see [Keil, 1989]. And SOCIAL CATEGORIES classify things in terms of the roles, powers and responsibilities people give them in specific culturally-sanctioned practices; see [Searle, 1997]. Proper inference about categories depends on rich reasoning that embraces defeasible inference, probability and causality. None of these categories are reducible to a boolean combinations of, say, perceptual features or other primitives, and none can be described purely in terms of logical definitions, theories and inferences in isolation from our connection to the world. See [Fodor, 1998].

A descriptive approach allow us to embrace this diversity in the world without reproducing it in the lexicon. It allows us, quite rightly, to represent interpretation while deferring the difficulties of domain reasoning. In fact, for many words, we may be able to say quite little about the categories they describe or the inferences they support in context in these taxonomic terms! For example, depending on the what matters in a particular case, we may take *vinegar* to describe a kind of result of fermentation (a natural kind, identified just by the preponderance of certain chemicals), as a kind of ingredient for imparting sourness to food (an artifact category, requiring a wholesome present state), or as a category of import, subject to particular taxes or controls (a social category, perhaps motivated by, but not identified with, specific methods of manufacture).

5 Description in Generation

Descriptive models thus promise offer an elegant computational analysis of the connection between application talk and application reasoning. More than that, description can in fact provide the basis for practical computational processes that connect a system's language use flexibly to its underlying domain ontology. I focus here on the case of natural language generation for dialogue. In generation, we start from a contribution that might usefully be made to an ongoing conversation. This contribution is expressed entirely in domain terms. But what we need to construct is a specific utterance, whose interpretation will make this contribution. In particular, the generator should expect that the hearer can use shared information, the utterance, the grammar, and the attentional and intentional state of the discourse, to reconstruct this interpretation.

This is the formulation of the language production problem that I and my colleagues arrived at in the SPUD generation system [Stone *et al.*, 2001]. SPUD can generate concise, contextually-appropriate utterances, including both speech and concurrent nonverbal behavior, by applying a simple, uniform and efficient decision-making strategy. This strategy exploits the lexicalized tree-adjoining grammar (LTAG) formalism in which SPUD's grammar is represented [Joshi *et al.*, 1975; Schabes, 1990]. LTAG grammars derive sentences by incorporating meaningful elements one-by-one into a provisional syntactic structure. SPUD makes these choices head-first and incrementally, in the order its grammar provides.

At each stage of derivation, SPUD determines both the intended interpretation for a provisional utterance and the interpretation that the hearer would recognize from it. SPUD implements this interpretation process directly in computational logic, using a constraint-satisfaction model of interpretation and an explicitly descriptive account of the meaning of utterances in context.

SPUD's choices of what elements to add to an incomplete sentence follow directly from SPUD's models of grammar and interpretation. The structure of the utterance suggests ways the sentence may be elaborated with further meaningful elements. The intended interpretation of each elaboration makes explicit the specific information that the utterance could contribute, and the specific links with the context that the utterance establishes. Meanwhile the recognized interpretation records SPUD's progress towards unambiguous formulation of descriptions. One measure of SPUD's success is that SPUD is able to use input representations that also suffice for other tasks—for animation for example—yet we can nevertheless design linguistic resources that allow SPUD to reproduce desired application language with high reliability. See [Cassell *et al.*, 2000] for further details.

6 Conclusion

In this paper, I have suggested that natural language meanings should be modeled as DESCRIBING an underlying domain ontology rather than CORRESPONDING to it. Theoretically, description provides an attractive explanation for the flexibility and context-sensitivity of language use in complex domains. Practically, description offers a convenient and general way to relax algorithmic assumptions so that natural language semantics can be connected to domain ontologies in less constrained ways.

Models may be good. But they are useless without resources to drive them. If we now know that natural language meanings allow us to connect language flexibly to the world using abstractions like possession to pick up specific relationships in context, and we now have a range of detailed domain ontologies that formalize that world knowledge for applications such as the semantic web [Fensel *et al.*, 2002], we must conclude that we are in urgent need of a new kind of language resource: a META-ONTOLOGY that formalizes the abstractions of semantics and allows us to classify elements of new domain ontologies in these linguistic terms.

References

[Allen *et al.*, 2001] James Allen, George Ferguson, and Amanda Stent. An architecture for more realistic conversational systems. In *Proceedings of Intelligent User Interfaces (IUI)*, 2001.

[Anderson *et al.*, 1991] A. Anderson, M. Bader, E. Bard, E. Boyle, G. M. Doherty, S. Garrod, S. Isard, J. Kowtko, J. McAllister, J. Miller, C. Sotillo, H. Thompson, and R. Weinert. The HCRC map task corpus. *Language and Speech*, 34:351–366, 1991.

[Artale *et al.*, 1996] Alessandro Artale, Enrico Franconi, Nicola Guarino, and Luca Pazzi. Part-whole relations in object-centered systems: An overview. *Data and Knowledge Engineering*, 20:347–383, 1996.

[Brachman *et al.*, 1990] Ronald Brachman, Deborah McGuinness, Peter Patel Schneider, Lori Alperin Resnick, and Alexander Borgida. Living with CLASSIC: when and how to use a KL-ONE-like language. In J. Sowa, editor, *Principles of Semantic Networks*. Morgan Kaufmann, 1990.

[Brennan, 2000] Susan E. Brennan. The vocabulary problem in spoken language systems. In Susan Luperfoy, editor, *Automatic Spoken Dialog Systems*. MIT, 2000. To appear.

[Bunt, 2000] Harry Bunt. Dialogue pragmatics and context specification. In Harry Bunt and William Black, editors, *Abduction, Belief and Context in Dialogue: Studies in Computational Pragmatics*, pages 81–150. John Benjamin, 2000.

[Carlson, 1980] Gregory N. Carlson. *Reference to Kinds in English*. Garland, 1980.

[Carlson, 1992] Gregory Carlson. Natural kinds and common nouns. In A. von Stechow and D. Wunderlich, editors, *Handbook of Semantics*, pages 370–398. de Gruyter, 1992.

[Cassell *et al.*, 2000] Justine Cassell, Matthew Stone, and Hao Yan. Coordination and context-dependence in the generation of embodied conversation. In *First International Conference on Natural Language Generation*, pages 171–178, 2000.

[Chang *et al.*, 2002] Nancy Chang, Srini Narayanan, and Miriam R. L. Petruck. Putting frames in perspective. In *COLING*, 2002.

[Charniak, 1973] Eugene Charniak. Jack and Janet in search of a theory of knowledge. In *Proceedings of IJCAI*, pages 337–343, 1973.

[Clark, 1996] Herbert H. Clark. *Using Language*. Cambridge University Press, Cambridge, UK, 1996.

[Crangle and Suppes, 1994] Colleen Crangle and Patrick Suppes. *Language and Learning for Robots.* CSLI, 1994.

[Edmonds, 2002] Philip Edmonds. SENSEVAL: the evaluation of word sense disambiguation systems. *ELRA Newsletter*, 7(3), 2002.

[Fensel *et al.*, 2002] Dieter Fensel, James A. Hendler, Henry Lieberman, and Wolfgang Wahlster, editors. *Spinning the Semantic Web: Bringing the World Wide Web to its Full Potential*. MIT, 2002.

[Fodor, 1998] Jerry A. Fodor. *Concepts: Where Cognitive Science Went Wrong*. Oxford, 1998.

[Furnas *et al.*, 1987] G. W. Furnas, T. K. Landauer, L. M. Gomez, and S. T. Dumais. The vocabulary problem in human-system communications. *Communications of the ACM*, 30:964–971, 1987.

[Grosz and Sidner, 1986] Barbara Grosz and Candace Sidner. Attention, intentions, and the structure of discourse. *Computational Linguistics*, 12:175–204, 1986.

[Haddock, 1989] Nicholas Haddock. *Incremental Semantics and Interactive Syntactic Processing.* PhD thesis, University of Edinburgh, 1989.

[Hobbs *et al.*, 1993] Jerry Hobbs, Mark Stickel, Douglas Appelt, and Paul Martin. Interpretation as abduction. *Artificial Intelligence*, 63:69–142, 1993.

[Hobbs, 1978] Jerry R. Hobbs. Resolving pronoun references. *Lingua*, 44:311–338, 1978.

[Hobbs, 1985] Jerry R. Hobbs. Ontological promiscuity. In *Proceedings of ACL*, pages 61–69, 1985.

[Iwańska and Shapiro, 2000] Łucja M. Iwańska and Stuart C. Shapiro, editors. *Natural Language Processing and Knowledge Representation: Language for Knowledge and Knowledge for Language*. MIT, 2000.

[Jackendoff, 1983] Ray S. Jackendoff. *Semantics and Cognition*. MIT Press, Cambridge, MA, 1983.

[Joshi *et al.*, 1975] Aravind K. Joshi, L. Levy, and M. Takahashi. Tree adjunct grammars. *Journal of the Computer and System Sciences*, 10:136–163, 1975.

[Kaplan, 1989] David Kaplan. Demonstratives. In J. Almog, J. Perry, and H. Wettstein, editors, *Themes from Kaplan*. Oxford, 1989.

[Keil, 1989] Frank C. Keil. *Concepts, Kinds and Cognitive Development*. MIT, 1989.

[Kilgarriff, 1997] Adam Kilgarriff. "i don't believe in word senses". *Computers and the Humanities*, 31(2):91–113, 1997.

[Kilgarriff, 1999] Adam Kilgarriff. 95% replicability for manual word sense tagging. In *EACL*, pages 277–278, 1999.

[Kilgarriff, 2001] Adam Kilgarriff. Generative lexicon meets corpus data: the case of non-standard word uses. In Pierrette Bouillon and Frederica Busa, editors, *The Language of Word Meaning*, pages 312–328. Cambridge, 2001.

[Krifka *et al.*, 1995] Manfred Krifka, Francis Jeffrey Pelletier, and Gregory Carlson. Genericity: An introduction. In G. N. Carlson and F. J. Pelletier, editors, *The Generic Book*, pages 1–124. University of Chicago, 1995.

[Kripke, 1982] Saul Kripke. *Naming and Necessity*. Harvard, 1982.

[Lakoff and Johnson, 1980] George Lakoff and Mark Johnson. *Metaphors We Live By*. Chicago, 1980.

[Lappin and Leass, 1994] Shalom Lappin and Herbert J. Leass. An algorithm for pronominal anaphora resolution. *Computational Linguistics*, 20(4):535–561, 1994.

[Mackworth, 1987] Alan Mackworth. Constraint satisfaction. In S.C. Shapiro, editor, *Encyclopedia of Artificial Intelligence*, pages 205–211. John Wiley and Sons, 1987.

[Mellish, 1985] Christopher S. Mellish. *Computer Interpretation of Natural Language Descriptions*. Ellis Horwood, Chichester, UK, 1985.

[Narayanan, 1997] Srini Narayanan. *KARMA: Knowledge-based Action Representations for Metaphor and Aspect*. PhD thesis, Berkeley, 1997.

[Neale, 1990] Stephen Neale. *Descriptions*. MIT, 1990.

[Nunberg, 1979] Geoffrey Nunberg. The non-uniqueness of semantic solutions: Polysemy. *Linguistics and Philosophy*, 3(2):143–184, 1979.

[Pinkal, 1999] Manfred Pinkal. On semantic underspecification. In Harry Bunt and Reinhard Muskens, editors, *Computing Meaning*, pages 33–56. Kluwer, 1999.

[Pustejovsky, 1991] James Pustejovsky. The generative lexicon. *Computational Linguistics*, 17(3):409–441, 1991.

[Putnam, 1975] Hilary Putnam. *Mind, Language and Reality*. Cambridge, 1975.

[Reiter and Dale, 2000] Ehud Reiter and Robert Dale. *Building Natural Language Generation Systems*. Cambridge, 2000.

[Reyle, 1993] Uwe Reyle. Dealing with ambiguity by underspecification: construction, representation and deduction. *Journal of Semantics*, 10:123–179, 1993.

[Schabes, 1990] Yves Schabes. *Mathematical and Computational Aspects of Lexicalized Grammars*. PhD thesis, Computer Science Department, University of Pennsylvania, 1990.

[Schank, 1980] Roger Schank. Language and memory. *Cognitive Science*, 4(3):243–284, 1980.

[Schuler, 2001] William Schuler. Computational properties of environment-based disambiguation. In *Proceedings of ACL*, pages 466–473, 2001.

[Searle, 1997] John R. Searle. *The Construction of Social Reality*. Simon and Schuster, 1997.

[Stalnaker, 1998] Robert Stalnaker. On the representation of context. *Journal of Logic, Language and Information*, 7:3–19, 1998.

[Steedman, 1997] Mark Steedman. Temporality. In Johan van Benthem and Alice ter Meulen, editors, *Handbook of Logic and Language*, pages 895–935. Elsevier, 1997.

[Stone and Thomason, 2002] Matthew Stone and Richmond H. Thomason. Context in abductive interpretation. In *EDILOG 2002: Proceedings of the sixth workshop on the semantics and pragmatics of dialogue*, pages 169–176, 2002.

[Stone and Webber, 1998] Matthew Stone and Bonnie Webber. Textual economy through close coupling of syntax and semantics. In *Proceedings of International Natural Language Generation Workshop*, pages 178–187, 1998.

[Stone *et al.*, 2001] Matthew Stone, Christine Doran, Bonnie Webber, Tonia Bleam, and Martha Palmer. Microplanning with communicative intentions: The SPUD system. *http://www.arxiv.org/abs/cs.CL/0104022*, 2001.

[Stone, 2003] Matthew Stone. Knowledge representation for language engineering. In Ali Farghaly, editor, *A Handbook for Language Engineers*. CSLI, 2003.

[Sweetser, 1990] Eve Sweetser. *From Etymology to Pragmatics: Metaphorical and Cultural Aspects of Semantic Structure*. Cambridge, 1990.

[Thomason, 1999] Richmond H. Thomason. Type-theoretic foundations for context, part i: Contexts as complex type-theoretic objects. In P. Bouquet, L. Serafini, P. Brezillon, M. Benerecetti, and F. Castellani, editors, *Modeling and Using Context*, LNAI 1688, pages 352–374. Springer, 1999.

[van Deemter and Kibble, 2000] Kees van Deemter and Rodger Kibble. On coreferring: Coreference in MUC and related annotation standards. *Computational Linguistics*, 26(4):629–637, 2000.

[Vieira and Poesio, 2000] Renata Vieira and Massimo Poesio. An empirically based system for processing definite descriptions. *Computational Linguistics*, 26(4):539–594, 2000.

[Webber *et al.*, 1995] Bonnie Webber, Norm Badler, Barbara Di Eugenio, Chris Geib, Libby Levison, and Michael Moore. Instructions, intentions and expectations. *Artificial Intelligence*, 73:253–259, 1995.

[Webber, 1983] Bonnie Lynn Webber. So what can we talk about now? In Michael Brady and Robert C. Berwick, editors, *Computational Models of Discourse*, pages 331–371. MIT, 1983.

[Yarowsky, 1995] David Yarowsky. Unsupervised word sense disambiguation rivalling supervised methods. In *ACL*, pages 189–196, 1995.