

# Interpreting Vague Utterances in Context

David DeVault and Matthew Stone

Department of Computer Science

Rutgers University

Piscataway NJ 08854-8019

David.DeVault@rutgers.edu, Matthew.Stone@rutgers.edu

## Abstract

We use the interpretation of vague scalar predicates like *small* as an illustration of how systematic semantic models of dialogue context enable the derivation of useful, fine-grained utterance interpretations from radically underspecified semantic forms. Because dialogue context suffices to determine salient alternative scales and relevant distinctions along these scales, we can infer implicit standards of comparison for vague scalar predicates through completely general pragmatics, yet closely constrain the intended meaning to within a natural range.

## 1 Introduction

Modeling context and its effects on interpretation may once have seemed to call for an open-ended investigation of people’s knowledge of the common-sense world (Hobbs et al., 1993). But research on the semantics of practical dialogue (Allen et al., 2001) now approaches dimensions of context systematically, through increasingly lightweight, factored models. The evolving state of real-world activity proceeds predictably according to background plans and principles of coordination (Rich et al., 2001). The status of the dialogue itself is defined by circumscribed obligations to ground prior utterances, follow up open issues, and advance real-world negotiation (Larsson and Traum, 2000). Finally, the evolving state of the linguistic context is a direct outgrowth of the linguistic forms interlocutors use and the linguistic relationships among successive utterances (Ginzburg and Cooper, 2001; Asher and Lascarides, 2003). These compatible models combine directly to characterize an aggregate information state that provides a general background for interpretation (Bunt, 2000).

We argue in this paper that such integrated models enable systems to calculate useful, fine-grained utterance interpretations from radically underspecified semantic forms. We focus in particular on vague scalar predicates like *small* or *long*. These predicates typify *qualitative* linguistic expression of *quantitative* information, and are thus both chal-

lenging and commonplace. Building on a multidimensional treatment of dialogue context, we develop and implement a theoretically-motivated model of vagueness which is unique in treating vague predicates as genuinely vague and genuinely context-sensitive, yet amenable to general processes of contextual and interpretive inference.

### 1.1 Semantic insights

We pursue our argument in the context of an implemented drawing application, FIGLET, which allows users to give English instructions to draw a caricature of an expressive face. Figure 1 shows a representative interaction with FIGLET; the user gives the successive instructions in (1):

- (1) a. Make two small circles.
- b. Draw a long line underneath.

Like Di Eugenio and Webber (1996), we emphasize that understanding such instructions requires contextual inference combining linguistic, task and domain knowledge. For example, consider the response to (1a) of placing circles so as to form the eyes of a new face. To recognize the possibility of drawing eyes exploits knowledge of the ongoing drawing task. To put the eyes where they belong in the upper part of the new face exploits domain knowledge. The response to (1b) adds the linguistic context as another ingredient. To identify where the line goes, the user uses the objects mentioned recently in the interaction as the understood spatial landmark for *underneath*. Figure 1 highlights the importance of using multidimensional representations of dialogue context in understanding instructions for quantitative domains.

We leverage this background context in our computational approach to vagueness. We model a vague utterance like *draw a long line* as though it meant *draw a line with, you know, length*. In this approach, vague predicates are completely underspecified; linguistic knowledge says nothing about how long something *long* is. Instead, vague language explicitly draws on the background knowl-




Figure 1: Motivating interaction: Vague instructions to draw a face.

edge already being applied in utterance interpretation. The user’s motivation in using *long* is to differentiate an intended interpretation, here an intended action, from alternative possibilities in context. Background knowledge already sets out the relevant ways to *draw a line*; *drawing a long line* means singling out some of them by the length of that new line. This model recalls dynamic theories of vague scalar predicates, such as the semantics of Kyburg and Morreau (2000), Barker (2002), or Kennedy (2003), but it is exactly implemented in FIGLET. The implementation capitalizes on the richness of current models of context to recover content for the *you know* of vagueness.

## 1.2 Overview

In Section 2, we motivate approaches to the semantics of vague scalar predicates that associate them with a presupposed standard of comparison. We illustrate how context can be understood to supply possible standards, and how pragmatic reasoning from utterances allows interlocutors to infer them. In Section 3, we establish a bridge to the general treatment of practical dialogue, by showing how multiple dimensions of context generally contribute to recognizing possible interpretations for underspecified utterances. Section 4 builds on Sections 2 and 3 to show how FIGLET exploits a rich model of utterance context to respond cooperatively to vague utterances like (1a) and (1b), while Section 5 details FIGLET’s actual implementation. We conclude in Section 6 by suggesting further challenges that vagueness still poses for computational semantics.

## 2 Vague standards in context

We adopt a view of vague predicates motivated by linguistic theory, particularly Kennedy’s approach (1999; 2003). We assume that gradable adjectives are associated with measurement functions mapping *individuals* to *degrees* on a *scale*. In FIGLET’s drawing domain, the relevant measurements pertain to spatial properties. For *long*, for example, the measurement maps individuals to their spatial lengths; for *small*, it maps individuals to degrees on an inverted scale of size.

*Positive* gradable adjectives compare measured degrees against a *standard* on the scale which is derived from context. For example, *long* says that an object’s length exceeds the threshold set by the current standard for length. Other forms, such as comparative adjectives or adjectives with explicit measure phrases, compare degrees differently.

Importantly, grammar says *nothing* about how standards for positive gradable adjectives are derived. In other words, contra Cresswell (1977) and others, the interpretation of adjectives is not computed relative to a grammatically-specified “comparison class” of related objects. And, contra Oates et al. (2000) and Roy and Pentland (2002), the interpretation of adjectives need not require statistical knowledge about how objects with different measurements on a scale tend to be described. Instead, standards are derived directly from an evolving context by the general principles that govern pragmatic resolution of context dependence.

Kennedy synthesizes a range of evidence for this claim. Here we go further, and provide a formal, implemented model. We can sketch the evidence and our model by considering two key examples.

First, we illustrate that vagueness depends directly on specific contextually-relevant distinctions. Consider the session with FIGLET shown in Figure 2. The user has elected to draw two objects side-by-side. The initial context just contains a square. The user utters (2).

- (2) Make a small circle.

To interpret (2) it doesn’t seem to help to appeal to general knowledge about how big circles tend to be. (It’s quite convoluted to even frame the idea in a sensible way.) Graff (2000) observes that standards often implicitly appeal to what we expect about particular individuals, not just what we know about similar individuals. In context, here, the user just seems to be asking for a circle vaguely smaller than the square. This is the interpretation FIGLET builds; to comply, FIGLET draws the circle an arbitrary but representative possible size. The point is that salient objects and actions inevitably set up meaningful dis-




Figure 2: Taking standards from context in (2): *Make a small circle.*




Figure 3: Disambiguating contextual standards in (3): *Make the small square a circle.*

tinctions in the context. Interlocutors exploit these distinctions in using vague adjectives.

Figure 3 illustrates that understanding vagueness is part of a general problem of understanding utterances. Figure 3 shows FIGLET’s action in a more complex context, containing two squares of different sizes. We consider the user’s instruction (3):

(3) Make the small square a circle.

FIGLET’s action changes the smaller of the two squares. The standard behind this interpretation is implicitly set to differentiate the contextually-salient objects from one another; the natural resolution of (3) does not require that either square be definitely small (Kyburg and Morreau, 2000). In Figure 3, for example, there are different potential standards that would admit either both squares or neither square as small. However, we can rule out these candidate standards in interpreting (3). The user’s communicative intention must explain how a unique square from the context can be identified from (3) using a presupposed *small* standard. If that standard is too big, both squares will fit. If that standard is too small, neither square will fit. Only when that standard falls between the sizes of the squares does (3) identify a unique square.

The examples in Figures 2 and 3 show two ways new standards can be established. Once established, however, standards become part of the evolving context (Barker, 2002). Old standards serve as defaults in interpreting subsequent utterances. Only if no better interpretation is found will FIGLET go back and reconsider its standard. This too is general pragmatic reasoning (Stone and Thomason, 2003).

### 3 Dimensions of context in interpretation

To cash out our account of contextual reasoning with vagueness, we need to characterize the context for practical dialogue. Our account presupposes a context comprising *domain and situation knowledge*, *task context* and *linguistic context*. In this section, we survey each of these dimensions of context, and show how they converge in the resolution of underspecification across a wide range utterances.

Domain and situation knowledge describes the commonsense structure of the real-world objects and actions under discussion. Practical dialogue restricts this otherwise open-ended specification to the circumscribed facts that are directly relevant to an ongoing collaboration. For example, in our drawing domain, individuals are categorized by a few types: types of shape such as *circles* and *squares*; and types of depiction such as *eyes* and *heads*. These types come with corresponding constraints on individuals. For example, the shape of a mouth may be a line, an ellipse, or a rectangle, while the shape of a head can only be an ellipse. These constraints contribute to interpretation. For instance, a head can never be described as *a line*, for example, since heads cannot have this shape.

Task context tracks collaborators’ evolving commitment to shared goals and plans during joint activity. In FIGLET’s drawing domain, available actions allow users to build figure parts by introducing shapes and revising them. Our experience is that users’ domain plans organize these actions hierarchically into strategic patterns. For example, users tend to complete the structures they begin drawing before drawing elsewhere; and once they are satisfied with what they have, they proceed in natural sequence to a new part nearby. Task context plays a powerful role in finding natural utterance interpretations. By recording a plan representation and keeping track of progress in carrying it out, FIGLET has access to a set of candidate next actions at each point in an interaction. Matching the user’s utterance against this candidate set restricts the interpretation of instructions based on the drawing already created and the user’s focus of attention within it. For example, if the user has just drawn the right eye onto an empty face, they are likely to turn to the left eye next. This context suggests making a winking left eye in response to *draw a line*, an interpretation that might not otherwise be salient.

Linguistic context records the evolving status of pragmatic distinctions triggered by grammatical conventions. One role of the linguistic context is its contribution to distinguishing the prominent entities




Figure 4: Context in instructions.

that can serve as the referents of pronouns and other reduced expressions. To see this, note that, as far as domain knowledge and task context go, the instruction *make it bigger* could apply to any object currently being created. If the figure is hierarchical, there will be many possibilities. Yet we typically understand *it* to refer specifically to an object mentioned saliently in the previous utterance. The linguistic context helps disambiguate *it*.

Figure 4 illustrates how the three different dimensions of context work together. It illustrates an interaction with FIGLET where the user has just issued an instruction to create two eyes, resulting in the figure state shown at the left in Figure 4. The user’s next instruction is (4):

- (4) Draw a line underneath.

We focus on how the context constrains the position and orientation of the line.

Linguistic context indicates that *underneath* should be understood as *underneath the eyes*. This provides one constraint on the placement of the line. Task context makes drawing the mouth a plausible candidate next action. Domain knowledge shows that the mouth can be a line, but only if further constraints on position, orientation and length are met. In understanding the instruction, FIGLET applies all these contextual constraints simultaneously. The set of consistent solutions—drawing a horizontal line at a range of plausible mouth positions below the eyes—constitutes the utterance interpretation. FIGLET acts to create the result in Figure 4 by choosing a representative action from this set.

#### 4 Interpreting vague utterances in context

In our approach, the linguistic context stores agreed standards for vague predicates. Candidate standards are determined using information available from domain knowledge and the current task context. In FIGLET’s drawing domain, possibilities include the actual measurements of objects that have already been drawn. They also include the default domain measurements for new objects that task context says could be added. Setting standards by a measure-

ment is our shorthand for adopting an implicit range of compatible standards; these standards remain vague, especially since many options are normally available (Graff, 2000).

We treat the use of new candidate standards in interpretation as a case of presupposition accommodation (Bos, 2003). In presupposition accommodation, the interpretation of an utterance must be constructed using a context that differs from the actual context. When speakers use an utterance which requires accommodation, they typically expect that interlocutors will update the dialogue context to include the additional presumptions the utterance requires. We assume that all accommodation is subject to two Gricean constraints. First, we assume whenever possible that an utterance should have a uniquely identifiable intended interpretation in the context in which it is to be interpreted. Second, we assume that when interpretations in alternative contexts are available, the speaker is committed to the strongest one—compare Dalrymple et al. (1998). Inferring standards for vague predicates is a special case of this general Gricean reasoning.

The principles articulated thus far in Sections 2–4 allow us to offer a precise explanation of FIGLET’s behavior as depicted in Figure 1. The user starts drawing a face with an empty figure. In this domain and task context, *make two circles* fits a number of possible actions. For example, it fits the action of drawing a round head and its gaping mouth. However, in (1a), what the user actually says is *make two small circles*. The interpretation for (1a) must accommodate a standard for *small* and select from the continuum of size possibilities two new circles that meet this standard.

The standards in this context are associated with the size distinctions among potential new objects. The different qualitative behavior of these standards in interpretation can be illustrated by the standards set from possible new circular objects that are consistent with the face-drawing task. We can set the standard from the default size of an eye, from the default size of a mouth (larger), or from the default size of a head (larger still).<sup>1</sup> Because each standard allows all smaller objects to be created next, these standards lead to 1, 3, and 6 interpretations, respectively. So we recover the standard from the eye, which results in a unique interpretation.<sup>2</sup>

<sup>1</sup>Since the default sizes of new objects reflect the relative dimensions of any other objects already in the figure, FIGLET’s default sizes are not generally equivalent to static comparison classes.

<sup>2</sup>Note that there are many potential sources of standards for *small* that FIGLET does not currently pursue. E.g. the average size of all objects already in the figure. We believe that general

In tandem with its response, FIGLET tracks the changes to the context. The task context is updated to note that the user has drawn the eyes and must continue with the process of creating and revising the features of the face. The linguistic context is updated to include the new *small* standard, and to place the eyes in focus.

This updated context provides the background for (1b), the user’s next instruction *draw a long line underneath*. In this context, as we saw with Figure 4, context makes it clear that any response to *draw a line underneath* must draw the mouth. Thus, unlike in (1a), all the interpretations here have the same qualitative form. Nevertheless, FIGLET’s Gricean reasoning can still adjust the standard for length to differentiate interpretations quantitatively, and thereby motivate the user’s use of the word *long* in the instruction. FIGLET bases its possible standards for *length* on both actual and potential objects. It can set the standard from an actual eye or from the two eyes together; and it can set the standard from the default mouth or head. The mouth, of course, must fit inside the head; the largest standard is ruled out. All the other standards lead to unique interpretations. Since the length of the two eyes together is the strictest of the remaining standards, it is adopted. This interpretation leads FIGLET to the response illustrated at the right in Figure 1.

## 5 Implementation

We have implemented FIGLET in Prolog using CLP(R) real constraints (Jaffar and Lassez, 1987) for metric and spatial reasoning. This section presents a necessarily brief overview of this implementation; we highlight how FIGLET is able to exactly implement the semantic representations and pragmatic reasoning presented in Sections 2–4. We offer a detailed description of our system and discuss some of the challenges of building it in DeVault and Stone (2003).

### 5.1 Semantic representation

In FIGLET, we record the semantics of user instructions using *constraints*, or logical conjunctions of open atomic formulas, to represent the contextual requirements that utterances impose; we view these constraints as presuppositions that speakers make in using the utterance. We assume matches take the form of instances that supply particular domain representations as suitable values for variables. Stone (2003) motivates this framework in detail.

methods for specifying domain knowledge will help provide the meaningful task distinctions that serve as candidate standards for vague predicates on our approach, but pursuing this hypothesis is beyond the scope of this paper.

In (5a-d), we show the presuppositions FIGLET assigns to an utterance of *Make two small circles*, arranged to show the contributions of each individual word. In (5e), we show the contribution made by the utterance to an evolving dialogue; the effect is to propose that an action be carried out.

- (5) a.  $simple(A) \wedge target(A, X) \wedge fits\_plan(A) \wedge holds(result(A, now), visible(X)) \wedge holds(now, invisible(X)) \wedge$
- b.  $number(X, 2) \wedge$
- c.  $standard(small, S) \wedge holds(result(A, now), small(X, S)) \wedge$
- d.  $number(X, multiple) \wedge holds(result(A, now), shape(X, circle))$
- e.  $propose(A)$

We formulate these constraints in an expressive ontology. We have terms and variables for *actions*, such as  $A$ ; for *situations*, such as  $now$  and  $result(A, now)$ ; for *objects*, such as  $X$ ; for standards for gradable vague predicates (scale-threshold pairs), such as  $S$ ; and for quantitative points and intervals of varying dimensionality, as necessary.

### 5.2 Pragmatic reasoning

Constraint networks such as (5a-e) provide a uniform venue for describing the various contextual dependencies required to arrive at natural utterance interpretations. Thus, the contextual representation and reasoning outlined in Sections 3 and 4 is realized by a uniform mechanism in FIGLET: specifications of how to reason from context to find solutions to these constraints.

For example, Section 3 described domain knowledge that links particular object types like *eyes* and *heads* with type-specific constraints. In our implementation, we specify real and finite constraints that individuals of each type must satisfy. In order for an individual  $e$  of type  $t$  to serve as part of a solution to a constraint network like (5a-e),  $e$  must additionally meet the constraints associated with type  $t$ . In this way, FIGLET requires utterance interpretations to respect domain knowledge.

Solving many of the constraints appearing in (5a-e) requires contextual reasoning about domain actions and their consequences. Some constraints characterize actions directly; thus  $simple(A)$  means that  $A$  is a natural domain action rather than an abstruse one. Constraints can describe the effects of actions by reference to the state of the visual display in hypothetical situations; thus  $holds(result(A, now), shape(X, circle))$  means that the individual  $X$  has a circular shape once action  $A$  is carried out. Constraints can additionally char-

acterize causal relationships in the domain; thus  $target(A, X)$  means that action  $A$  directly affects  $X$ , and the constraints of (5a-d) together mean that carrying out action  $A$  in the current situation causes two small circles to become visible. These constraints are proved in FIGLET by what is in effect a planner that can find complex actions that achieve specified effects via a repertoire of basic domain actions.

Task context is brought to bear on interpretation through the  $fits\_plan(A)$  constraint of (5a). FIGLET uses a standard hierarchical, partially ordered plan representation to record the structure of a user’s task. We specify the solutions to  $fits\_plan(A)$  to be just those actions  $A$  that are possible next steps given the user’s current state in achieving the task. Since these task-appropriate actions can factor additional constraints into interpretation, enforcing the  $fits\_plan(A)$  constraint can help FIGLET identify a natural interpretation.

As discussed in Section 4, FIGLET records a list of current standards for vague scalar adjectives in the linguistic context. The constraint  $standard(small, S)$  of (5c) connects the overall utterance interpretation to the available standards for *small* in the linguistic context. FIGLET interprets utterances carrying semantic constraints of the form  $standard(vague-predicate, S)$  in one or two stages. In the first stage, the constraint is solved just in case  $S$  is the prevailing standard for *vague-predicate* in the linguistic context. If there is *no* prevailing standard for an evoked vague property, or if this stage does not yield a unique utterance interpretation, then FIGLET moves to a second stage in which the constraint is solved for any standard that captures a relevant distinction for *vague-predicate* in the context. If there is a strongest standard that results in a unique interpretation, it is adopted and integrated into the new linguistic context.

### 5.3 Parsing and Interpretation

Language understanding in FIGLET is mediated by a bottom-up chart parser written in Prolog. As usual, chart edges indicate the presence of recognized partial constituents within the input sequence. In addition, edges now carry constraint networks that specify the contextual reasoning required for understanding. In addition to finite instances (Schuler, 2001), these networks include real constraints that formalize metric and spatial relationships. Interpretation of these networks is carried out incrementally, during parsing; each edge thus records a set of associated candidate interpretations. Since domain reasoning can be somewhat time-intensive in our current implementation, we adopt a strategy

of delaying the solution of certain constraints until enough lexical material has accrued that the associated problem-solving is judged tractable (DeVault and Stone, 2003).

## 6 Assessment and Conclusion

In our approach, we specify a genuinely vague semantics: vague words evoke a domain-specific scale that can differentiate alternative domain individuals. To find a unique interpretation for a vague utterance, we leverage ordinary inference about the domain, task, and linguistic context to recover implicit thresholds on this scale.

We believe that further methodological advances will be required to evaluate treatments of vagueness in indefinite reference, such as that considered here. For example, obviously the very idea of a “gold standard” for resolution of vagueness is problematic. We believe that the best argument for a theory of vagueness in a language interface would show that naive users of the interface are, on the whole, likely to accept its vague interpretations and unlikely to renegotiate them through clarification. But the experiment would have to rule out confounding factors such as poorly-modeled lexical representation and context tracking as sources for system interpretations that users reject.

We intend to take up the methodological challenges necessary to construct such an argument in future work. In the meantime, while our current implementation of FIGLET exhibits the promising behavior discussed in this paper and illustrated in Figures 1–4, some minor engineering unrelated to language understanding remains before a fruitful evaluation can take place. As alluded to above, the tight integration of contextual reasoning and interpretation that FIGLET carries out can be expensive if not pursued efficiently. While our initial implementation achieves a level of performance that we accept as researchers (interpretation times of between one and a few tens of seconds), evaluation requires us to improve FIGLET’s performance to levels that experimental participants will accept as volunteers. Our analysis of FIGLET indicates that this performance can in fact be achieved with better-regimented domain problem-solving.

Nevertheless, we emphasize the empirical and computational arguments we already have in support of our model. Our close links with the linguistic literature mean that major empirical errors would be surprising and important across the language sciences. Indeed, limited evaluations of treatments of vague definite reference using standards of differentiation or very similar ideas have been promising

(Gorniak and Roy, In Press). The computational appeal is that all the expensive infrastructure required to pursue the account is independently necessary. Once this infrastructure is in place the account is readily implemented with small penalty of performance and development time. It is particularly attractive that the approach requires minimal lexical knowledge and training data. This means adding new vague words to an interface is a snap.

Overall, our new model offers three contributions. Most importantly, of course, we have developed a computational model of vagueness in terms of underspecified quantitative constraints. But we have also presented a new demonstration of the importance and the feasibility of using multidimensional representations of dialogue context in understanding descriptions of quantitative domains. And we have introduced an architecture for resolving underspecification through uniform pragmatic mechanisms based on context-dependent collaboration. Together, these developments allow us to circumscribe possible resolutions for underspecified utterances, to zero in on those that the speaker and hearer could adopt consistently and collaboratively, and so to constrain the speaker's intended meaning to within a natural range.

### Acknowledgments

We thank Kees van Deemter and our anonymous reviewers for valuable comments. This work was supported by NSF grant HLC 0308121.

### References

- J. Allen, D. Byron, M. Dzikovska, G. Ferguson, L. Galescu, and A. Stent. 2001. Towards conversational human-computer interaction. *AI Magazine*, 22(4):27–37.
- N. Asher and A. Lascarides. 2003. *Logics of Conversation*. Cambridge.
- C. Barker. 2002. The dynamics of vagueness. *Linguistics and Philosophy*, 25(1):1–36.
- J. Bos. 2003. Implementing the binding and accommodation theory for anaphora resolution and presupposition. *Computational Linguistics*, 29(2):179–210.
- H. Bunt. 2000. Dialogue pragmatics and context specification. In H. Bunt and W. Black, editors, *Abduction, Belief and Context in Dialogue*, pages 81–150. Benjamin.
- M. Cresswell. 1977. The semantics of degree. In B. H. Partee, editor, *Montague Grammar*, pages 261–292. Academic.
- M. Dalrymple, M. Kanazawa, Y. Kim, S. Mchombo, and S. Peters. 1998. Reciprocal expressions and the concept of reciprocity. *Linguistics and Philosophy*, 21(2):159–210.
- D. DeVault and M. Stone. 2003. Domain inference in incremental interpretation. In *Proc. ICoS*.
- B. Di Eugenio and B. Webber. 1996. Pragmatic overloading in natural language instructions. *Int. Journal of Expert Systems*, 9(2):53–84.
- J. Ginzburg and R. Cooper. 2001. Resolving ellipsis in clarification. In *Proc. ACL*.
- P. Gorniak and D. Roy. In Press. Grounded semantic composition for visual scenes. *Journal of Artificial Intelligence Research*.
- D. Graff. 2000. Shifting sands: An interest-relative theory of vagueness. *Philosophical Topics*, 28(1):45–81.
- J. Hobbs, M. Stickel, D. Appelt, and P. Martin. 1993. Interpretation as abduction. *Artificial Intelligence*, 63:69–142.
- J. Jaffar and J.-L. Lassez. 1987. Constraint logic programming. In *Proc. POPL*, pages 111–119.
- C. Kennedy. 1999. *Projecting the adjective: The syntax and semantics of gradability and comparison*. Garland.
- C. Kennedy. 2003. Towards a grammar of vagueness. Manuscript, Northwestern.
- A. Kyburg and M. Morreau. 2000. Fitting words: Vague words in context. *Linguistics and Philosophy*, 23(6):577–597.
- S. Larsson and D. Traum. 2000. Information state and dialogue management in the TRINDI dialogue move engine toolkit. *Natural Language Engineering*, 6:323–340.
- T. Oates, M. D. Schmill, and P. R. Cohen. 2000. Toward natural language interfaces for robotic agents. In *Proc. Agents*, pages 227–228.
- C. Rich, C. L. Sidner, and N. Lesh. 2001. COLLAGEN: applying collaborative discourse theory to human-computer interaction. *AI Magazine*, 22(4):15–26.
- D. Roy and A. Pentland. 2002. Learning words from sights and sounds: A computational model. *Cognitive Science*, 26(1):113–146.
- W. Schuler. 2001. Computational properties of environment-based disambiguation. In *Proc. ACL*, pages 466–473.
- M. Stone and R. H. Thomason. 2003. Coordinating understanding and generation in an abductive approach to interpretation. In *Proc. DiaBruck*, pages 131–138.
- M. Stone. 2003. Knowledge representation for language engineering. In A. Farghaly, editor, *A Handbook for Language Engineers*, pages 299–366. CSLI.