

# The Use-Mention Distinction and its Importance to HCI\*

Michael L. Anderson<sup>1</sup>, Yoshi A. Okamoto<sup>3</sup>, Darsana Josyula<sup>2</sup> and Don Perlis<sup>1,2</sup>

(1) Institute for Advanced Computer Studies

(2) Department of Computer Science

University of Maryland, College Park, MD 20742

(3) Microsoft Corporation, Redmond, WA 98052

(301) 405-1746

{mikeoda, darsana, perlis}@cs.umd.edu, yoshio@microsoft.com

The Use-Mention Distinction and its  
Importance to HCI

## Abstract

In this paper we contend that the ability to engage in meta-dialog is necessary for free and flexible conversation. Central to the possibility of meta-dialog is the ability to recognize and negotiate the distinction between the use and mention of a word. The paper surveys existing theoretical approaches to the use-mention distinction, and briefly describes some of our ongoing efforts to implement a system which represents the use-mention distinction in the service of simple meta-dialog.

## 1 Introduction

We use the term *conversational adequacy* to denote the ability to engage in free and flexible conversation. It is our contention that the ability to engage in meta-dialog is necessary for conversational adequacy, and more importantly, that a robust meta-dialogic ability can make up for weaknesses in other areas of linguistic ability (Perlis et al., 1998). For this reason, we think that time spent understanding and implementing meta-dialog in natural language HCI systems will be well rewarded; the ability to engage in even simple meta-dialog can be used to fruitfully enhance the performance of interactive systems, even those having relatively limited speech recognition and language processing abilities.

Central to the possibility of meta-dialog is the ability to recognize and negotiate the distinction between the use and mention of a word. Although this distinction has attracted some attention from philosophers of language, it has not gotten sufficient notice from those most directly involved in the theory and design of natural language HCI. We hope with this paper to begin to correct this lacuna, and stimulate further research in this important area. To this end,

\* This research was supported in part by the AFOSR and ONR

we will first try to establish the importance of meta-dialog, and of mastery of the use-mention distinction in particular, to conversational adequacy. We will then survey and evaluate the existing theoretical frameworks for understanding the use-mention distinction. Finally, we will briefly describe some of our efforts to implement a system which represents the use-mention distinction in the service of simple meta-dialog, and suggest some directions for future research.

## 2 The importance of meta-dialog

Conversation is not generally the exchange of fully formed, grammatically correct, and error-free utterances. Indeed, it is unlikely that there could ever be a fully fluent, error-free dialog; even putting aside problems of signal reception, and assuming perfect syntactic processing, the ability to understand a dialog partner involves such complicated and uncertain tasks as modeling their knowledge state and using context to disambiguate reference. In practice, conversation tends to be like this: a series of contributions in which a common context is constantly assessed and maintained with the use of inference, clarifications and confirmations. The main thread of the conversation starts and stops; there are pauses and hedges, questions are asked and sub-dialogs are spawned. Far from representing a break-down or failure of conversation, these processes, which monitor and establish *grounding*, are central to the nature of dialog. As was modeled by (Clark and Schaefer, 1987; Clark and Schaefer, 1989b) and has been confirmed by a great deal of research since (e.g. (Brennan, 1998; Brennan, 2000; Brennan and Hulstijn, 1995; Cahn and Brennan, 1999; Clark and Brennan, 1991; Krahmer et al., 1999a; Krahmer et al., 1999b; Paek and Horvitz, 1999)), conversation is composed of two-phase contributions—an utterance is first *presented* to a dialog partner, but it is not until evidence of understanding is accrued that the utterance is *accepted* and becomes part of the common ground of the two speakers.

Insofar as this process involves monitoring not just of the content of the conversation, but also of one's own *understanding* of the content, as well as making inferences about the other's state of understanding based on aspects of the conversation like timing, feedback, etc., this suggests to us the centrality of meta-linguistic skills to conversational ability. Clark (Clark and Schaefer, 1989a) presents empirical evidence for the use of meta-linguistic skills by young children. Among many other things, these skills include:

1. **Monitoring one's ongoing utterance:** An example of this was seen in a 2 year, 6 month child practicing parts of speech (in this case its pronunciation of "berries") on its own: "Back please/berries/*not* barries/barries, barries/*not* barries/berries /ba ba"
2. **Checking the result of an utterance:** Children at least as young as 5 years, 4 months comment on and correct the utterances of others. They also verify that the listener has understood their utterance and attempt a repair otherwise.
3. **Deliberately trying to learn:** For instance, a 4 year old will ask things like: "Mommy, is it AN A-dult or A NUH-dult?"
4. **Predicting the consequences of using inflections, words, phrases or sentences:** This includes judging the politeness of utterances, which is exhibited by children aged four and a half. Children can also correct word order in sentences judged "silly". Clark cites instances of this being done by two-year olds.

This variety of meta-linguistic skills exhibited by very young children further suggests the central importance these skills have for learning and processing language.

It seems fairly clear that these abilities could not exist apart from a facility with meta-reasoning (reasoning about the reasoning process, whether yours or your interlocutors') and meta-dialog, and in particular without a grasp of the distinction between the use and mention of a word. This can be seen most explicitly in item (3), above, but the general ability to take language or a linguistic item as an object of scrutiny—which ability is at the center of the use-mention distinction—is part of all the abilities listed above.

For current purposes, we would like to draw your attention to the following key conversational ability—reasoning about, discussing, and adapting word meanings and references. This can include reasoning about:

- word meaning—knowledge of the word itself, part of speech, lexical semantics, concept picked out, etc.
- specific reference of a term denoting a particular object, including both names and definite, indefinite and deictic expressions, using a concept to pick out an individual.
- the relation of words to the concept(s) or entities referred to

Clearly the exercise of this capacity—which can be understood, for instance, as the ability to understand sentences like: (a) What does 'felicitous' mean? (b) 'Felicitous' means happy, or well-formed. (c) Is 'shop' a verb as well as a noun? (d) What is a 'VCR'?—requires the ability to take words themselves as objects for reasoning and discussion and, more particularly, the ability to negotiate the distinction between the use and mention of a given term.

### 3 Survey of Theoretical Approaches to the Use-Mention Distinction

In this section, we briefly sketch the various theories of the use-mention distinction, and the related issue of understanding quotation. Extensive citations are provided for the reader who wishes to explore the issues in more detail.

#### 3.1 Classic theories

There are four main theoretical positions which define the early discussion of quotation: the Name theory, the Description theory, the Demonstrative theory and the Identity theory.<sup>1</sup>

In the Name theory, quotations are considered to be names, that is, quoted expressions name their referents. (Tarski, 1933) explores the frontier by elucidating the issue of quotation in general with this approach.

Quotation-mark names may be treated like single words of a language ... the single constituents of these names ... fulfill the same function as the letters and complexes of successive letters in single words. Hence they can possess no independent meaning. (Tarski, 1933) p. 159

(Quine, 1940) and (Cram, 1978) travel the same theoretical path. The main point seems to be that, as in a name, "[t]he meaning of the whole does not depend upon the meanings of the constituent words."

<sup>1</sup>The discussion is influenced throughout by (Saka, 1998).

(Quine, 1940) p. 26 “Chief Sitting Bull” refers to Sitting Bull, and not, for instance, to a sitting bull, because that linguistic item is a name. However, whether or not this is generally true of names, it does not seem to be a good model for quotation in general. For instance, I can say: “‘ $\alpha$ ’ is a Greek letter,” and be entirely understood. Had I been thereby told only the *name* for  $\alpha$ , I would not fully understand the sentence until I had determined what the name named. But this does not appear to be the case. More generally, quoted expressions and the linguistic items they refer to seem to be systematically related by the shape or nature of the expression itself. This and like considerations give rise to the Description theory.

The Description theory provides more structurally “descriptive power” to the quoted expressions. There are roughly two approaches. The formal approach, outlined in (Tarski, 1933; Quine, 1940; Richard, 1986), is based on formal elements, such as orthographical elements or phonological elements. The other is the lexical approach, exemplified by (Geach, 1957), which analyzes the quoted expression word-by-word. In general, the idea is that the quoted expression describes its referent, thus “It was the best of times ...” = the expression formed by ‘It’, plus ‘was’, etc. The descriptive theory thus appears to rely on having names for some base elements of the language, whether orthographic, phonemic, or lexical; insofar as this is so, it falls prey to the same criticisms as the Name theory.

(Davidson, 1969) explores the Demonstrative account of quotation, and this framework has convinced many linguists and philosophers over three decades. Just to cite a few: (Partee, 1973; Christensen, 1967; Goldstein, 1984; Garcia-Carpintero, 1994; Cappelen and Lepore, 1997). The Demonstrative theory holds that quote marks demonstrate, or point to, the quoted material, on one interpretation to its shape. Hence, (e) is understood as (f) in the following:

- e. “Cats” is a noun.
- f. Cats. That complex of shapes is a noun.

But taken as a formal transformational rule, this cannot account for the possibility of iterated quotation. It appears that (g) would be transformed into (h) and then something like (i):

- g. “‘Cats’ ” is a noun phrase.
- h. “Cats.” That is a noun phrase.
- i. Cats. That that is a noun phrase.

Opposed to the demonstrative theory is the Identity theory, according to which quoted expressions co-refer to themselves. The Identity theory is as widely discussed as the demonstrative theory. See (Frege, 1892/1980; Quine, 1940; Wittgenstein, 1953; Tajtelbaum, 1957; Whitely, 1957; Searle, 1969; Washington, 1992; Reimer, 1996).

Whereas the Demonstrative Theorist regards quote marks (or context) as referential and the quoted material as an inert adjunct, the Identity Theorist conversely regards the quoted material as (self-) referential and the quote marks as semantically empty. Despite these differences, the Identity and Demonstrative accounts can both be called picture theories, for both claim that quotation resembles its referent, the quoted material. (Saka, 1998)

It appears that the Identity theory also falls prey to the recursion problem, for it does not appear to be able to account for the difference between “‘sunset’” and “sunset”, which have distinct referents.

### 3.2 Syntactic and Semantic Treatments of Quoted Expressions

Although the title, *The Syntax and Semantics of Quotation* sounds promising, in her article Partee (Partee, 1973) only deals with the direct quotation of the whole sentence, thereby, unfortunately, excluding the cases in which we are most interested. She classifies quotation into *word quotation* (j) Should “pickup” ever be hyphenated?, *sentence quotation*, (k) “I am speaking now” is always true when spoken, and *direct quotation* (l) Tom said, “My grandfather was in Seattle.” Focusing her attention on the last, Partee basically supports Davidson’s claim (Davidson, 1969) that the quoted sentence is not syntactically or semantically a part of the sentence that contains it. Partee then argues that the quoted sentence represents nothing more than the surface structure, and it is that surface structure (and not the meaning of the quoted sentence) that contributes to the meaning of the sentence that contains it.

(A)ll the apparent evidence for deeper syntactic and semantic structure is a result of the main sentence speaker’s understanding and analyzing the noises he is quoting as a sentence, just as he understands and analyzes a sentence, a string of noises, that comes to him from someone else. It is in this process that the language/metalinguage distinction is crucially involved. (Partee, 1973) p. 418

Cram (Cram, 1978) considers Partee’s conclusion “absurd” (p. 43). He disagrees with Partee’s selective application of her constraint only to direct quotations. He thinks that the same constraint “would hold equally well for sentences in general” (p. 43). He instead argues that the quoted expression has the status of a lexical item—more specifically of a noun phrase—and the structure is specified at the level of the matrix sentence. Consider the following examples:

- m. “They are cooking apples” ist zweideutig
- n. Der Satz “They are cooking apples” ist zweideutig
- o. The sentence “They are cooking apples” ist zweideutig.

Cram presents the contrast between (m) and (n) (the tagging test (Reichenbach, 1947)) as a supportive case for the first part of his argument that the quoted expression has the status of a noun phrase. Then he contrasts (n) and (o) to argue that the NP tag “belongs both logically and semantically to the level of the matrix sentence” (p. 44). In order for his analysis to work, however, he has to make one assumption:

The base rules which generate the matrix sentence also generate a quotational tag under a Noun Phrase node, which may occur either in subject position or elsewhere. The quotation tag is always specified at the level of Deep Structure, but may be deleted by an optional transformation. (Cram, 1978) p.44

From the perspective of computational applications, it is not hard to implement the aforementioned rule as a system designer. In fact, treating a quoted expression as an NP can be a way to solve a simple parsing problem. For instance, following (Reichenbach, 1947), a simple lexical insertion rule can allow one to treat a quoted sentence as a syntactically opaque unit.

However, there are cases when expressions are mentioned without overt quote marking. If we were to design a system simply to look for such quote marks, then we would miss such cases. Yet, we do not want to design a parser or a system that automatically treats any syntactically dubious phrases as noun phrases, either.

### 3.3 Disambiguated Ostention Theory

In our judgment, the Disambiguated Ostention theory (Saka, 1998) is the most satisfactory of the var-

ious theoretical accounts which formalize the use-mention distinction. In general, according to Saka, the capacities for both use and mention stem from the same source, namely from the fact that the human mind associates a multiplicity of deferred ostensions with any exhibited token, thus giving rise to pragmatic ambiguity. Among the possible ostensions are:

- p. orthographic form: cat
- q. phonic form: /kaet/
- r. lexical entry: [cat, /kaet/, count noun, CAT]
- s. intension: CAT
- t. extension: x: x a cat

Saka then explains how these ostensibly deferred items interact in the following way:

- 5. Exposure to the written label cat (p) or the spoken label /kaet/ (q) evokes the corresponding lexeme (r) in every competent speaker of English, where a lexeme is an arbitrary ordered n-tuple including orthographic form, phonic form, syntactic category, meaning, register, etc.
- 6. The lexeme [cat, /kaet/, count noun, CAT] specifies the intension CAT (s) according to the pragmatic function.
- 7. CAT determines the extension x: x a cat (t) according to some function.

Unfortunately, he does not elaborate much on the nature of the “pragmatic function” (6) nor “some function” (7). However, he does provide a useful characterization of the use-mention distinction:

Speaker S **uses** an expression X iff:

- i. S exhibits a token of X;
- ii. S thereby ostends the multiple items associated with X (including X’s extension);
- iii. S intends to direct the thoughts of the audience to the extension of X.

Speaker S **mentions** an expression X iff:

- iv. S exhibits a token of X;
- v. S thereby ostends the multiple items associated with X;
- vi. S intends to direct the thoughts of the audience to some item associated with X *other than its extension*.

Saka’s approach is significant in that the use-mention distinction is characterized in terms of extension. And, more importantly, it “allows for the existence of mentioning without quote marks” (p. 127). This is definitely an advantage over the rest of the known approaches, because they assume that quote marks are required for an expression to be considered as mentioned. The Disambiguated Ostension theory has been criticized in (Cappelen and Lepore, 1999). However, that criticism seems to center around the question of whether words really have multiple *referents*, as Saka contends, or only one (perhaps complex) referent. In so far as the matter turns on technical issues in the philosophy of language and on the nature of reference, it does not affect the arguments offered here. It is enough for our purposes to hold that a speaker might, by mentioning (or with the use of quotation), mean to draw our attention to items other than the extension of a word, whether or not these other items are legitimately considered to be *referents* of that word.

## 4 Description of Current Research

The long-term goal we have set for ourselves is the design of a system modeled not on conversation with a fluent colleague, but rather, for example, on a task-oriented interaction with a stranger who doesn’t speak much of a common language. In these situations, e.g., buying a train ticket in a foreign country, speakers are often able to communicate to effectively solve a joint task, in spite of problems in word recognition, or use of unfamiliar words or syntactic structures. Despite the difficulties of understanding the language, interactive dialog behaviors and ongoing repairs allow humans to overcome some of these problems.

One major step toward this goal was the design and implementation of a model of action-directive exchanges (task oriented requests) (Traum et al., forthcoming; Traum and Andersen, 1999) based on an evolving-time (“active logic”) model of inference (Elgot-Drapkin et al., 1993; Elgot-Drapkin and Perlis, 1990; Purang et al., 1999). Our model works via a step-wise transformation of the literal request made by a user (e.g. “Send the Boston train to New York”) into a specific request for an action that can be performed by the system or domain. In the case of “the Boston train,” the system we have implemented is able to interpret this as “the train in Boston,” and then further disambiguate this into a specific train currently at Boston station, which it will send to New York. Information about each step in the transformation is maintained, to accommodate any repairs that might be required in the case of negative feedback (if for instance, the system picks

the wrong train, and the user says “No” in response to the action). (See (Traum et al., forthcoming) for more detailed information.) This implementation represents an advance not just in its ability to reason initially about the user’s intention (e.g., by “the Boston train” the user means . . .”) but in its ability to respond in a context-sensitive way to post-action user feedback, and use that feedback to aid in the interpretation of the user’s original and future intentions. (For instance, if the user says “No, send the Boston train to New York” the system is able to identify its mistake as its choice of referent for “the Boston train,” and choose a different train; if there are no other trains in Boston, it tells the user: “Please specify the train by name.”) However, although our current system is able to respond in this way to post-action user feedback, the actual transformation of the user’s utterance from its literal form into a performable request takes place without user feedback. We believe that a system which could request and utilize user feedback in this process—one, that is, which contained a more robust meta-dialogic component—would represent a significant advance.

Thus, the next step for us is to design and implement a model of Question-Answer exchanges which can be used not just independently to track ongoing interactions of this type, but also (more significantly for our purposes) to supplement the interpretation of task-oriented requests by allowing the system to request and use user feedback during the interpretive process. It is important not just that the dialog agent be able to respond to questions from the user, but also be able to ask questions *of* the user. For instance, consider the user request “Send the Bullet train to Bean Town.” Assuming that the system is unable to determine the meaning of “Bean Town” by itself, it should be able to ask the user for specific help. Understanding and implementing ways of representing the use-mention distinction in natural language HCI is vital to making this possible, for should the user respond: “‘Bean Town’ means Boston” the system will have to recognize this as an instance of mention rather than use, and deal with the utterance appropriately. It is important that the system recognizes this distinction in order to identify the referent of the word “Bean Town” correctly. In the first case, the system has to disambiguate the referent as a city whereas in the latter case, it has to interpret the referent as a word.

Our own general approach to this issue involves implementing a specialized domain, which (with apologies to Austin (Austin, 1962)) knows how to do things with words. Following Saka’s definition, if the system determines that a given sentence is meant to draw attention to something other than the extension of a word, processing of the sentence is passed

on to the “words” domain, where words are objects and are dealt with in terms of ostensions (p)-(s), above. In the simplest cases this means doing things like updating the lexicon and correcting spelling; more complex cases include adopting new conventions for word use, and learning new words.

However, such a system depends on being *able* to determine when a word is being mentioned; methods for doing this reliably is one area in which we would especially like to see more research. If the user were to characterize each instance of a mention with explicit quotation marks, then it would be easy for the system to determine when a word is being mentioned rather than being used. However, in spoken natural language HCI, such explicit quote marks are neither feasible nor reliable, since it is not natural to say “Open quote” before and “Close quote” after each mention, and in any case human-human conversation seems to proceed successfully in most cases without such explicit quotation marks.

Context seems to play a key role in determining whether a word is being used or mentioned. If the system comes across a word that it does not know about, then it could ask the user what it is, and generate the expectation that a future user response would be to provide the answer. Keeping track of such expectations can help the system to identify the mention of a word. For instance, in the example, since the system is not able to determine the meaning of “Bean Town,” it can create an expectation that the user is going to provide a definition for “Bean Town” and in that definition the word “Bean Town” would be mentioned rather than used. Still, this approach is clearly not going to cover a large percentage of cases, and other methods are needed. One method that we are exploring is to make the system sensitive to special words like “means” and “is”, often used to introduce definitions of unknown words. Such intension words can indicate that a word is being mentioned rather than used. For instance, if the user were to say “‘Bean Town’ means Boston,” “means” would signal that it is a mention of the word “Bean Town” rather than its use that the user intends. Lexical category words (like “noun” and “verb”), can also act as cues for recognizing the possibility of a mention, as in “‘Cats’ is a noun.” If the system misidentifies a mention to be a use, the user would engage in meta-dialog to correct the system. The fact that dialog is proceeding at the level of meta-dialog is itself a hint to the system to look for the possibility of a mention.

A forthcoming paper describes in more detail the progress of our research to date, including our proposed representation scheme and overall system architecture. However, as indicated, this paper is in-

tended primarily to draw attention to a problem, finding solutions (or exploring approaches) to which we believe will have a great impact on the future of natural language HCI.

## References

- J. L. Austin. 1962. *How to Do Things with Words*. Harvard University Press.
- Susan E. Brennan and Eric A. Hulstien. 1995. Interaction and feedback in a spoken language system: A theoretical framework. *Knowledge-Based Systems*, 8:143–151.
- Susan E. Brennan. 1998. The grounding problem in conversations with and through computers. In S.R. Fussell and R.J. Kreuz, editors, *Social and Cognitive Psychological Approaches to Interpersonal Communication*, pages 201–225. Lawrence Erlbaum.
- Susan E. Brennan. 2000. Processes that shape conversation and their implications for computational linguistics. In *Proceedings of the 38th Annual Meeting of the Association for Computational Linguistics*.
- Janet E. Cahn and Susan E. Brennan. 1999. A psychological model of grounding and repair in dialog. In *Proceedings of the AAAI Fall Symposium on Psychological Models of Communication in Collaborative Systems*, pages 25–33.
- Herman Cappelen and Ernest Lepore. 1997. Varieties of quotation. *Mind*, 106:429–50.
- Herman Cappelen and Ernest Lepore. 1999. Using, mentioning and quoting: A reply to saka. *Mind*, 108:741–50.
- Niels Christensen. 1967. The alleged distinction between use and mention. *Philosophical Review*, 76:358–67.
- H.H. Clark and Susan E. Brennan. 1991. Grounding in communication. In J. Levine L.B. Resnik and S.D. Teasley, editors, *Perspectives on Socially Shared Cognition*, pages 127–149.
- H.H. Clark and E.F. Schaefer. 1987. Collaborating on contributions to conversations. *Language and Cognitive Processes*, 2:19–41.
- Herbert H. Clark and Edward F. Schaefer. 1989a. Contributing to discourse. *Cognitive Science*, 13:259–294. Also appears as Chapter 5 in (Clark, 1992).
- H.H. Clark and E.F. Schaefer. 1989b. Contributing to discourse. *Cognitive Science*, 13:259–294.
- Herbert H. Clark. 1992. *Arenas of Language Use*. University of Chicago Press.
- David F. Cram. 1978. The syntax of direct quotation. *Cahiers de Lexicologie*, 33(2):41–52.
- Donald Davidson. 1969. On saying that. *Synthese*, 19:130–46.
- J. Elgot-Drapkin and D. Perlis. 1990. Reasoning situated in time I: Basic concepts. *Journal of Experimental and Theoretical Artificial Intelligence*, 2(1):75–98.
- J. Elgot-Drapkin, S. Kraus, M. Miller, M. Nirkhe, and D. Perlis. 1993. Active logics: A unified formal approach to episodic reasoning. Technical Report UMIACS TR # 99-65, CS-TR # 4072, Univ of Maryland, UMIACS and CSD.
- Gottlob Frege. 1892/1980. On sense and reference. In Peter Geach and Max Black, editors, *The Philosophical Writings of Gottlob Frege*, pages 56–78. Oxford: Basil Blackwell.
- Manuel Garcia-Carpintero. 1994. Ostensive signs. *Jphil*, 91:253–64.
- Peter T. Geach. 1957. On names of expressions. *Mind*, 59:388.
- Laurence Goldstein. 1984. Quotation of types and types of quotation. *Analysis*, 44:1–6.
- Emiel Krahmer, Marc Swerts, Mariet Theune, and Mieke Weegels. 1999a. Error detection in spoken human-machine interaction. In *Proceedings of Eurospeech'99*, Budapest, Hungary.
- Emiel Krahmer, Marc Swerts, Mariet Theune, and Mieke Weegels. 1999b. Problem spotting in human-machine interaction. In *Proceedings of Eurospeech'99*, Budapest, Hungary.
- Tim Paek and Eric Horvitz. 1999. Uncertainty, utility and misunderstanding: A decision-theoretic perspective on grounding in conversational systems. In *Proceedings, AAAI Fall Symposium on Psychological Models of Communication in Collaborative Systems*.
- Barbara Partee. 1973. The syntax and semantics of quotation. In Stephen Anderson and Paul Kiparsky, editors, *Festschrift for Morris Halle*, pages 410–18. New York: Holt, Rinehart, Winston.
- D. Perlis, K. Purang, and C. Andersen. 1998. Conversational adequacy: mistakes are the essence. *Int. J. Human-Computer Studies*, 48:553–575.
- K. Purang, D. Purushothaman, D. Traum, C. Andersen, D. Traum, and D. Perlis. 1999. Practical reasoning and plan execution with active logic. In *Proceedings of the IJCAI'99 Workshop on Practical Reasoning and Rationality*.
- W.V.O. Quine. 1940. *Mathematical Logic*. Harvard University Press, Cambridge, MA.
- Hans Reichenbach. 1947. *Elements of Symbolic Logic*. Macmillan, London.
- Marga Reimer. 1996. Quotation marks: demonstratives or demonstrations? *Analysis*, 56:131–42.
- Mark Richard. 1986. Quotation, grammar and opacity. *Linguistics and Philosophy*, 9:383–403.
- Paul Saka. 1998. Quotation and the use-mention distinction. *Mind*, 107:113–35.
- John Searle. 1969. *Speech Acts*. Cambridge University Press, Cambridge.
- Al Tajtelbaum. 1957. It is impossible to be told anyone's name. *Analysis*, 17:52.
- Alfred Tarski. 1933. The concept of truth in formalized languages. In J. Woodger, editor, *Logic, semantics, mathematics*. Oxford: Oxford University Press.
- D. Traum and C. Andersen. 1999. Representations of dialogue state for domain and task independent

meta-dialogue. In *Proceedings of the IJCAI99 workshop: Knowledge And Reasoning in Practical Dialogue Systems*, pages 113–120.

David Traum, Carl Andersen, Yuan Chong, Darsana Josyula, Michael O'Donovan-Anderson, Yoshi Okamoto, Khemdut Purang, and Don Perlis. forthcoming. Representations of dialogue state for domain and task independent meta-dialogue. *Electronic Transactions on Artificial Intelligence*.

Corey Washington. 1992. Quotation. *jphil*, 89:582–605.

C.H. Whitely. 1957. Names and words. *Analysis*, 17:119–20.

Ludwig Wittgenstein. 1953. *Philosophical Investigations*. Macmillan, New York.