Linguistic Representation and Gricean Inference

Matthew Stone*
Computer Science and Cognitive Science
Rutgers, the State University of New Jersey
110 Frelinghuysen Road Piscataway NJ 08854-0819 USA
mdstone@cs.rutgers.edu
http://www.cs.rutgers.edu/~mdstone

Abstract

An essential ingredient of language use is our ability to reason about utterances as intentional actions. Linguistic representations are the natural substrate for such reasoning, and models from computational semantics can often be seen as providing an infrastructure to carry out such inferences from rich and accurate grammatical descriptions. Exploring such inferences offers a productive pragmatic perspective on problems of interpretation, and promises to leverage semantic representations in more flexible and more general tools that compute with meaning.

1 Introduction

In the philosophy of language, theorists have often attributed much of the flexibility with which we use language not to the complexity of our linguistic system, but rather to our robust abilities to reason about other agents' actions, choices and motivations. This insight is originally due to Grice [19, 20], and so I will refer to such reasoning as Gricean. The examples in (1) illustrate Gricean reasoning. Each has a natural interpretation that departs from the literal meaning that we would expect our grammar to assign it.

(1) a. Can you pass the salt?

b. You may now go.

c. Chris is a tiger.

*I owe the organizers of IWCS and a sabbatical leave from Rutgers for the opportunity to prepare this paper. This presentation benefits from a number of collaborations, especially with Rich Thomason and Bonnie Webber, and with David DeVault, Paul Tepper and Jesse Fischer at Rutgers. Thanks all.

(1a) is a yes-no question, but can function as a request. (1b) is a statement of fact, but can function to give the permission it describes. (1c) locates an animal in the natural order, but can function to characterize a human personality. These apparently unexpected functions may actually follow from our abilities to understand one another by reasoning cooperatively.

For example, Searle [37] explains utterances like (1a) by recognizing that observers can, in clear cases, anticipate an agent's planned future actions from an agent's present behavior. In particular, Searle suggests that (1a) functions as a request because it so clearly triggers the mutual expectation that the speaker plans a request eventually. Similarly, Lewis [30] explains utterances like (1b) by observing that collaborative social agents can take up and accomplish the goals that they recognize behind other agent's plans, even when the plans they recognize are themselves defective. Lewis suggests that what is permitted is defined by mutual understanding, within the context of agents' social relationships, and from this background argues that (1b) functions to give permission because the utterance so clearly communicates the goal of registering the speaker's newly-defined permission on the common ground. Finally, Grice [21] explains utterances like (1c) by observing that agents can rely on mutual assumptions of rationality to coordinate not just on literal interpretations of their utterances, but also on salient reinterpretations. (At least with suitable models of interpretation [32, 38].) In particular, with the word *tiger* in (1c), speaker and audience can coordinate to evoke a concept of aggression that is saliently associated with tigers' behavior. The Gricean reasoning we see in cases such as these reminds us that conventionalized literal meanings cannot always do justice to our intentions. In using language, we are improvising in an open-ended world, and we must act creatively if we are to meet our goals for communication, or balance them against the other significant goals we have for our interactions, including goals for self-presentation, affiliation and autonomy [8].

Utterances like (1) are not the sort that we typically aim to account for in principled ways in computational systems—we find enough in much more straightforward examples to occupy ourselves fully. Nevertheless, I believe that reconciling formal models of agency with formal models of meaning, in the ways utterances like (1) appear to require, is one of the most vital challenges for the cognitive science of language in general, and for computational semantics in particular. It's not just because our systems will someday have to exhibit the flexibility of language and action we see in (1), if we are correct that this flexibility is part and parcel of successful communication in the real world. More immediately, it's because general models of agency provide substantive high-level constraints we can draw on to flesh out semantic theories in more detail and to implement such ideas in extensible and elegant architectures. Few other frameworks provide such a clear guide to what to do next, and how to do it. My goal in this paper is to justify this contention.

I begin in Section 2, by exploring the close relationship between dynamic approaches to semantic interpretation [25, 40] and philosophical accounts of interpretation as intention recognition such as those [19, 20, 37, 21, 30] introduced in connection with (1). It seems a short and natural step to move from a description of meaning in terms of CHANGE, as in dynamic semantics, to the description of meaning in terms of ACTIONS THAT CAUSE CHANGE that the pragmatic theories require. But to take that step, we need the confidence to scrap those parts of the received wisdom that no longer fit our present motivations. Taking the step opens up a number of possibilities for applying ideas from computational semantics in new ways—I conclude by discussing two specific cases in which we are required to infer a speaker's assumptions about the context from their utterances. In Section 3.1, I describe a computational model of the interpretation of vague utterances. In Section 3.2, I describe a reversible model of referring-expression generation. Both models combine dynamic semantic representations with genuinely pragmatic processes of intention-recognition. Overall, the presentation in this paper explores the semantic implications of work that I have presented to rather different audiences in [41, 42], and sketches ongoing investigations that I and my colleagues aim to present more fully in the future.

2 Intentionality and Dialogue

The basic perspective of Gricean pragmatics is to analyze utterances as actions. This analysis frames knowledge of meaning in terms of the potential communicative effects of utterances, and locates a speaker's meaning in the choice of a particular utterance for an assortment of its effects. Adopting this perspective in systems that apply sophisticated knowledge of linguistic semantics should make computing meaning more robust and more useful. For example, on this perspective, understanding is a problem of recognizing the speaker's rationale for using an utterance in context. Such a rationale will link the utterance to the agreed context and purposes of the conversation, as understood by the speaker. Accordingly, it will provide a broad gauge of what makes sense in context, which the system can draw on in processes such as ambiguity resolution. Planning a response, meanwhile, means acting cooperatively to meet the requirements and objectives of the ongoing conversation. Models of cooperation provide flexible guides to action not just when a collaboration proceeds normally, but also in cases of unexpected outcomes, misunderstandings and disagreements. Finally, in generation, this perspective calls for representing the rationale behind a system's own actions. This can help a system keep its utterances concise by anticipating whether its utterances will be understood as intended, and can provide the system with the wherewithal to follow its utterances up through further interaction with the user.

This Gricean model represents a widespread view, particularly within the American AI tradition. The work of Allen and colleagues [16, 1, 5] provides a prominent example of the construction of conversational systems through general models of agency; and there are many other examples (too many to survey in this short space). In the past, systems in this tradition have tended to abstract intentionality away from linguistic representations and processes. For example, they have tended to encapsulate any Gricean reasoning in autonomous modules which presuppose fully disambiguated logical forms as inputs or outputs. This has left little opportunity to leverage computational semantics for conversational agency (or conversely, to apply agency in computing meaning). At present, however, this state of affairs is starting to change.

I see two reasons for this. First, computational semantics increasingly offers ready abstractions with which to precisely characterize the intended effects of an utterance. Take (2), after Donellan [15], which we imagine as the utterance of a speaker A on some particular occasion.

(2) The man with the champagne is happy.

One aspect of the speaker's meaning here is to use the definite description *the man with the champagne* ANAPHORICALLY, to EVOKE some DISCOURSE REFERENT m from the context [31, 46, 25]. Another aspect of meaning is to contribute the information that m is happy to the CONVERSATIONAL RECORD, an EVOLVING abstract model of the agreed content of the conversation [45, 40].

Such abstractions are enough to explain why A would choose (2) and commit to utter it. In this sense, they determine A's intention. The content of this intention can be spelled out as follows: A utters (2); in the context, some referent m is the man with the champagne; thus, as a result of uttering (2) in the context, it becomes part of the conversational record that m is happy. I'll refer to this content as i_1. I understand i_1 as a PRAGMATIC INTERPRETATION, a privileged symbolic structure through which we manage our efforts to understand, be understood, and play our part in the conversation.

Such abstractions really are a close fit to the theory of intention, as independently developed by Bratman [6], Pollack [34, 35] and others. In this theory, any plan or intention is a complex mental representation summarizing an agent's reason to act. It lays out a course of related actions, identifies some key circumstances that now hold in the world, and shows why the agent might expect the planned actions to lead to a desirable outcome in the current circumstances. Thus, when we recognize someone's intention, we know what they think they are doing, why they are doing it, and how they must think it will work. In applying this model to utterances, we require intentions that associate a concrete linguistic structure such as (2) (the course of action) with inferences that show how, in the current discourse context (the key circumstances), the meaning of that structure brings specific information to the conversation (a desired outcome). This is indeed how we characterized (2) with i_1 above.

A second reason Gricean inference can increasingly interact with computational semantics is our general experience with complex symbolic representations that spell out links between an utterance and the context and goals of language use. These representations include highly-structured objects such as feature structures [17] and proof-terms [24, 33]. Our experience with such representations paves the way to recast formalisms for interpretation in explicitly pragmatic terms. Recall that plans and intentions don't just describe changes to the context, they EXPLAIN how those changes come about. Plans and intentions such as i_1 must be represented as ARGUMENTS, inferences that track the causal connections at play in an utterance, and use those causal connections to predict the effects those utterances might hypothetically achieve.

Consider our example (2) and its pragmatic interpretation i_1. (3) reports i_1 as an axiomatic deduction in a modal logic of knowledge and time, using $[\textsc{cr}]p$ meaning that p follows from the conversational record and using $[\textsc{n}]p$ meaning that p holds after the current action takes place.

(3) a. $[\textsc{cr}](\mathit{man}(m) \wedge \mathit{with}(m, d) \wedge \mathit{champagne}(d))$

Assumption about context

b. $[\textsc{cr}](\mathit{man}(M) \wedge \mathit{with}(M, D) \wedge \mathit{champagne}(D)) \supset$
$A\,\mathit{utters}(2) \supset$
$[\textsc{n}][\textsc{cr}]\mathit{happy}(M)$

Dynamic semantics

c. $A\,\mathit{utters}(2)$

Hypothesized course of action

d. $[\textsc{n}][\textsc{cr}]\mathit{happy}(m)$

Desired effect, by modus ponens and instantiation from (3a–3c).

We could also record such inferences as terms in a suitable type theory, as in [33], or even as the trace of a suitable logic-programming interpreter, as in [24]. One advantage of the inferential form of (3) is thus that it supports the treatment of local pragmatics from earlier inferential accounts. For example, the inference from (3a) to the antecedent of (3b) leaves room for bridging in the resolution of presuppositions. And the inference from the consequent of (3b) to (3d) leaves room in the interpretation for certain kinds of relevance-guided implicature.

Pragmatic interpretations such as (3) are linguistic representations that support Gricean inference. They not only describe an utterance in linguistic terms but map out a reason to use the utterance; they can figure not just in understanding but in general models of deliberation and cooperation. As we develop such structures, we are no longer just applying theoretical ideas from AI, linguistics, and philosophy. We are synthesizing them into something new.

For example, pragmatic interpretations mark a departure from the tradition of formal specification of agents within the knowledge representation community. This tradition originates in Levesque's view of KR as a tool whose main use is to characterize, validate and interact with agents [29], and is epitomized by the intricate attitude-definitions for cooperative agents that occupy Cohen and Levesque [10, 11, 12], and even to some extent the related work of Pollack and Grosz and colleagues [34, 23, 22]. Of course, the attitudes of agents in using utterances should conform to normative specifications of cooperation, but for us that is secondary. Fundamentally, pragmatic interpretations formalize the content of intentions as data structures for agent implementations.

Pragmatic interpretations also steer clear of traditional oppositions from linguistics, such as that between unstructured representationalist views of interpretation, as in DRT, and the more information-theoretic approach of dynamic semantics. Consider the treatment of presupposition in particular. Presupposition figures in (3) in the content of axiom (3b), which describes the context change potential of the utterance. According to (3b), the contribution of the utterance is contingent on finding a man M and a drink of champagne D as part of the conversational record. The inference from (3a) derives the specific instance m for M and d for D. This link between the utterance and its context is part of the interpretation and has to be recognized to recover the speaker's intention. This is very much an anaphoric and representationalist treatment of presupposition, in the spirit of [46, 26].

However, there is not, and cannot be, in structures such as (3) the kind of surgical process of accommodation that van der Sandt has proposed. Presuppositions don't disappear from pragmatic interpretations once they are resolved; on the contrary, the resolution is itself recorded as part of the inference. Moreover, in any proof, you can use a rule only by deriving its antecedent from the available assumptions—a proof records the logical consequences of assumptions about the world. An interpretation in particular must spell out sufficient assumptions about the context to derive the presuppositions required by any causal axioms, such as (3b), that figure in it. This is in keeping with Beaver's observation that the interpretation of presupposition always involves a "top-level" assumption about the conversation [4]. Indeed, this argument shows that a representationalist view of presupposition may rule out local accommodation for principled reasons, just like Stalnaker's information-based view of pragmatic presupposition does [39].

Pragmatic interpretations even mark a departure from Grice's characterization of communication! Grice was particularly concerned with the process of communication, and observed that understanding a meaningful utterance such as (2) depends on recognizing the intention behind it. So with (2) for example, a hearer B will recognize i_1. From this, B knows the speaker intends the conversation to evolve in a certain way. In accepting and grounding A's contribution [9, 7], B cooperatively sees to it that these effects are taken up. In the normal case, A anticipates and indeed INTENDS this uptake. In other words, the speaker A asserts (2) WITH THE INTENTION that B will recognize i_1, adopt it spontaneously and cooperatively, and make it clear that B is doing so. Let's call this intention i_2. It represents an EXPECTATION ABOUT CONVERSATIONAL DYNAMICS, an uncertain conclusion about the unfolding of a collaborative exchange within a group of language users, which allows a speaker to select a good utterance in context.

Both the linguistic reasoning encapsulated in pragmatic interpretations, such as i_1, and the psychological reasoning encapsulated in expectations about conversational dynamics, such as i_2, are essential to the process of meaningful communication that is characteristic of language use. Grice crystallizes this in his famous analysis of non-natural meaning. Recasting this analysis in present terms, we might offer (4).

(4) To mean p is to deliberately attempt to use CONVERSATIONAL DYNAMICS to contribute p to the conversational record, by manifesting a PRAGMATIC INTERPRETATION which represents an utterance as contributing p to the conversational record.

This doesn't look much like Grice's definition. For Grice, a communicative intention spells out ITS OWN causal role in directing the hearer's cognitive processes; Grice assumed that $i_1 = i_2$ and understood this intention as self-referential. These complex speaker meanings must inherently conflate linguistic conventions with psychological generalizations. If that's right, research on interpretation that sidesteps psychology, as undertaken by computational semanticists for example, can only be regarded as short-circuiting Gricean reasoning, not interfacing with it; see Asher and Lascarides [2]. Grice's suggestion has a geeky appeal, but it is not inevitable. Grice's position in fact was formulated with the anticipation that it would help philosophers REDUCE speakers' grammatical knowledge of linguistic meaning to speaker's general knowledge of human psychology—a tendentious outgrowth of other aspects of Grice's philosophy of language. Thomason [45] argues that the key to the Gricean account is just that the speaker's overriding intention plays out through a suitable process of intention-recognition. Thomason introduces an account of meaning which allows for a distinction between the overriding intention, such as i_2, and a recognized intention, such as i_1.

The dual status of utterances as intentional actions allows us to acknowledge the richness of the processes that unfold in real conversation and to articulate a principled role in these processes for pragmatic interpretations (and the semantic representations they contain). The characterization of these processes is now central to formal mdoels of dialogue; for example, Ginzburg and Cooper [17, 13] keep track of the contextual links of an utterance as well as its intended contribution, and define dialogue transitions that update the conversation in different ways by drawing flexibly on these representations. In dialogue, hearers sometimes take up the new presuppositions they find a speaker making. But other times they reject them, or ask about them—or simply delight with the speaker in their absurdity. These four scenarios, for example, can all start from B's successful recognition of the same representation (3) as manifested by A in uttering (2).

—It really is mutual knowledge that m is the man with the champagne. B knows that A is sincere and cooperative and so intends the conversation to evolve as mapped out in (3). In accepting A's contribution, and in grounding it, B cooperatively sees to it that the effects envisaged in (3) are taken up, and thus that the conversational record does come to provide that m is happy.

—B does not not already know what kind of drink m has. But from recognizing (3), B knows that A presumes m has champagne. Judging A sincere and cooperative, B takes up the assumption that m has champagne, and goes on to accept A's contribution. A and B proceed with a conversational record where m has champagne and m is happy.

—B knows that m does not have champagne; m has water. From recognizing (3), however, B knows A presumes otherwise. Judging A sincere and cooperative, B concludes that A is in error, and follows up with a correction: *it's not a champagne, it's water.* From this point, the interlocutors still have to reach agreement on what drink m has, and what m's emotional state is.

—B looks to m, sees a sour face, and laughs. That is, since m is obviously not happy, A could not have been sincere in offering this interpretation. (3)

represents a pretense, and may not contribute explicit information to the conversational record. But perhaps it has its other conversational effects. It might, for example, cement A and B's relationship, by reinforcing some implicit understanding they have (e.g., the party is too far gone to be saved).

3 Working with Pragmatic Interpretation

I think pragmatic interpretations embody an appealing theory of language use with close ties to our existing practice. But what I like most about them is their ability to inject syntactic and semantic knowledge into pragmatic processes based on Gricean reasoning. In this section I'll sketch two examples of this, both of which focus on reconstructing a speaker's assumptions in using an utterance.

3.1 Interpreting Vagueness

Utterances using vague words can achieve specific effects in context. Imagine describing an arrangement of three two-cm squares (call them s_1, s_2 and s_3) and one four-cm square (call it s_4) as in (5).

(5) The large square is nice.

With this utterance, the speaker can be understood to refer to s_4. In this section, I'll sketch how the intention behind (5) might be recognized in its context, even if we assume that vague adjectives presuppose a standard of comparison, and the context here does not inherently supply one. What's more, even though (5) is associated with a flawed communicative intention in this context, simply by being recognized, (5) can achieve all the effects we would normally associate with it; and it can, in addition, update the context to include a standard of comparison for large squares—by accommodation. This explanation shows how we can capture the pragmatic perspective on vagueness common to many recent proposals [18, 28, 3] within general computational models of utterance understanding.

I begin by specifying the presupposition of (5), as given in (6).

(6) $\mathit{square}(X) \wedge \mathit{size}(X, S) \wedge \mathit{large\text{-}std}(\langle D, \infty\rangle) \wedge \mathit{in}(S, \langle D, \infty\rangle)$

In words, X is a square, the size of X is S, the interval $\langle D, \infty\rangle$ (lower bounded by some scalar value D and without an upper bound) provides the standard for large size in the context, and S lies within this interval.

The presupposition arises as part of a communicative intention such as that represented in (7).

(7) a. $[\textsc{cr}](\mathit{square}(s_4) \wedge \mathit{size}(s_4, 4cm) \wedge \mathit{large\text{-}std}(\langle d, \infty\rangle) \wedge 2cm < d < 4cm)$
Assumption about context

b. $[\textsc{cr}](\mathit{square}(X) \wedge \mathit{size}(X, S) \wedge \mathit{large\text{-}std}(\langle D, \infty\rangle) \wedge \mathit{in}(S, \langle D, \infty\rangle)$
$A\ \mathit{utters}(5) \supset$
$[\textsc{n}][\textsc{cr}]\mathit{nice}(X)$
Dynamic semantics

c. $A\, utters(5)$

Hypothesized course of action

d. $[\textsc{n}][\textsc{cr}]nice(s_4)$

Desired effect, by logic from (7a–7c).

In particular, the presupposition originates in the antecedent of the dynamic semantic clause (7b). As before, the intention derives a specific instance of this presupposition by inference from a hypothesis (7a) about the conversational record. But now in this case, the standard is represented as an arbitrary or underspecified term d that must lie somewhere between the size of squares s_1, s_2 and s_3 and the size of square s_4. The vagueness of the interpretation consists in the underspecification; the speaker is not committed to a specific value for d and the hearer's cannot identify one.

Let's consider the hearer's inference in recognizing the plan in (7). The hearer is tracking the speaker's deliberation, and knows:

(8) a. The speaker is acting as though a certain context obtains. This pretend context is different from the actual context only in certain potentially predictable ways; in particular, the pretend context may supply standards for vague predicates that the actual context does not.

b. The pretend context supplies some intended instance of (6).

c. In this pretend context, this instance can be recognized as intended.

By (8a) and (8b), the hearer can infer that X is one of the four squares in the context, and that S is the size of X. By (8a) and (8b), the hearer can also infer that the pretend context specifies an interval of size that includes the size S of X. There now remain two qualitatively different standards (less than 2 cm; or between 2 cm and 4 cm). But (8c) eliminates the smaller standard, since it provides no way to recognize which of the four squares is X: the presupposition of the plan can be satisfied in the pretend context with any of the four possible squares. On the other hand, (8c) confirms the larger standard, since the presupposition now has only the resolution where X is square s_4. This inference combines abductive reasoning to reconstruct the speaker's context [24, 44] with constraint satisfaction to resolve references.

Once the hearer recognizes the plan, the hearer is free to respond to it in any reasonable way. The most cooperative strategy would be to update the representation of the actual context, to provide the standard of size the speaker appealed to, in a step of accommodation, and then to respond naturally to the utterance in the revised context. For example, if the hearer would have handed the large square to the speaker at this juncture if both had already agreed that this was the large square, the hearer could also hand the large square to the speaker here. Note, however, that this kind of cooperative strategy probably diverges from actual practice in face-to-face conversation, where more conservative and interactive coordination would be expected [9, 7].

3.2 Constructing Models for Generation

The discussion so far has centered on reasoning processes that a dialogue system might use in interacting with a user. But perhaps the most flexible tools we need for reasoning about meaning are those that will make it easier to develop and extend such systems in the first place—tools that help us construct semantic resources, for example. Such problems may also benefit from the application of Gricean reasoning to linguistic representations.

Consider generation of referring expressions (GRE), a well-studied and important subtask of NLG (see, e.g., [14]). The input to GRE is an entity, as characterized by a database of information about it that is presumed to be part of the common ground. The output is the semantics and perhaps also the syntax of a linguistic description that will allow the hearer to distinguish the entity uniquely from its salient alternatives. To build a GRE module requires identifying the context set of individuals that are explicit or implicit in application discourse, formalizing the relevant properties of all these individuals, and specifying how to compose these properties into linguistic descriptions. The GRE module thus epitomizes the resource-intensive character of many natural language processing tasks.

The simplest way to build a GRE module would be to supply it with examples of desired behavior, pairing entities with descriptions of them that would be satisfactory for a system to use in context. Automatic methods would then construct a suitable context set, and knowledge base for a satisfactory GRE module for the system. These NLG resources would account for the sample descriptions the designer has supplied, but could also generalize to other possible forms of referring expressions and to other possible contexts.

This sets up a problem of Gricean reasoning. The problem is to frame GRE tasks for the system in such a way that it can make its choices in a predictable way to match the specified examples. A solution involves reconstructing the choices a system has in GRE and reconstructing the reasons it must have to make those choices one way or another. Consider an example: we want the system to describe d_1 as *the furry black dog*. Then we must at least have the information in (9).

(9) a. $dog(d_1)$

b. $black(d_1)$

c. $furry(d_1)$

This information is required to SUPPORT the system's choices. But that can't be all the conversational record contains. Otherwise we would expect *it* or *the dog*. The generator must have alternatives to d_1 in mind. For example, maybe there is another dog d_2 that is furry but not black, and a third dog d_3 that is black but not furry, as in (10).

(10) a. $dog(d_2)$ $\quad$ $dog(d_3)$

b. $black(d_3)$

c. $furry(d_2)$

((10) assumes a convention of negation-by-failure; we specify that d_2 is not big and that d_3 is not black by simply omitting the relevant formulas.) Now the generator has a MOTIVATION for every element in its description of d_1: without *black*, the referent might be d_1 or d_2; without *furry*, the referent might be d_1 or d_3; and without *dog*, the generator wouldn't have an English noun phrase expressing the right properties. By supporting and motivating the generator's choices, we can ensure that the generator realistically should and would be able to use *the big black dog* to identify d_1 in this context.

To pursue this idea in a general way, we can use a model of linguistic choice after [43], where a grammatical derivation adds lexical elements one-by-one. We reconstruct a rationale for such choices. Each of these choices must be supported. And each must be motivated in the context of the generator's other commitments, and the goals of reference; it must fulfill a syntactic function that is required in a complete derivation, or else it must rule out some distractor.

Applying Gricean inference to such system-building problems allows us to draw on flexible, general modules. For example, the general approach makes predictions about how the same individuals could be described in new contexts, perhaps using more concise, context-sensitive descriptions. Krahmer and Theune [27] investigate the inherent context-sensitivity of general approaches to GRE; see also [14]. The general approach also offers an inexpensive diagnostic that specified linguistic behavior portrays a world of individuals in a way that is consistent and that interlocutors can be expected to recognize. This check would fail for some sets of examples. For instance, a specification that said d_1 could be described as *the black dog* and d_2 could be described as *the dog* could not be supported and motivated. In this specification *the dog* must be ambiguous. Failure in such cases flags a genuine defect in the specification for NLG, for which the appropriate response is to revise the NLG examples [36].

4 Conclusion

An essential ingredient of language use is our ability to reason about utterances as intentional actions. Linguistic representations are the natural substrate for such reasoning. Intentions are resources for deliberation; they therefore abstract away from considerations that don't bear on planning and choice. In the case of communication, abstracting away from the cooperative processes of conversation distills a communicative intention that records the grammatical description of the utterance, its links with the context, and its contributions to the conversational record.

From this perspective, models from computational semantics can often be seen as providing an infrastructure to carry out such inferences from rich and accurate grammatical descriptions. Exploring such inferences offers a productive pragmatic perspective on problems of interpretation, and promises to leverage semantic representations as part of more flexible and more general tools that compute with meaning.

References

[1] J. Allen, G. Ferguson, and A. Stent. An architecture for more realistic conversational systems. In *Intelligent User Interfaces*, 2001.

[2] N. Asher and A. Lascarides. *Logics of Conversation*. Cambridge, 2003.

[3] C. Barker. The dynamics of vagueness. *Linguistics and Philosophy*, 25(1):1–36, 2002.

[4] D. Beaver. *Presupposition and Assertion in Dynamic Semantics*. CSLI, 2001.

[5] N. Blaylock, J. Allen, and G. Ferguson. Managing communicative intentions with collaborative problem solving. In R. Smith and J. van Kuppevelt, editors, *Current and New Directions in Dialogue*. Kluwer, 2002.

[6] M. E. Bratman. *Intention, Plans, and Practical Reason*. Harvard, 1987.

[7] S. E. Brennan. *Seeking and Providing Evidence for Mutual Understanding*. PhD thesis, Stanford University, 1990.

[8] P. Brown and S. C. Levinson. *Politeness: Some Universals in Language Use*. Cambridge, 1987.

[9] H. H. Clark and E. F. Schaefer. Contributing to discourse. *Cognitive Science*, 13:259–294, 1989.

[10] P. R. Cohen and H. J. Levesque. Intention is choice with commitment. *Artificial Intelligence*, 42:213–261, 1990.

[11] P. R. Cohen and H. J. Levesque. Rational interaction as the basis for communication. In P. R. Cohen, J. Morgan, and M. E. Pollack, editors, *Intentions in Communication*, pages 221–256. MIT, 1990.

[12] P. R. Cohen and H. J. Levesque. Teamwork. *Nous*, 24(4):487–512, 1991.

[13] R. Cooper and J. Ginzburg. Using dependent record types in clarification ellipsis. In J. Bos, M. E. Foster, and C. Matheson, editors, *EDILOG 2002*, pages 45–52, 2002.

[14] R. Dale and E. Reiter. Computational interpretations of the Gricean maxims in the generation of referring expressions. *Cognitive Science*, 18:233–263, 1995.

[15] K. Donellan. Reference and definite description. *Philosophical Review*, 75:281–304, 1966.

[16] G. Ferguson and J. F. Allen. TRIPS: An intelligent integrated problem-solving assistant. In *AAAI*, pages 567–573, 1998.

[17] J. Ginzburg and R. Cooper. Clarification, ellipsis and the nature of contextual updates in dialogue. King's College London and Göteborg University, 2001.

[18] D. Graff. Shifting sands: An interest-relative theory of vagueness. *Philosophical Topics*, 28(1):45–81, 2000.

[19] H. P. Grice. Meaning. *Philosophical Review*, 66(3):377–388, 1957.

[20] H. P. Grice. Utterer's meaning and intention. *Philosophical Review*, 78(2):147–177, 1969.

[21] H. P. Grice. Logic and conversation. In P. Cole and J. Morgan, editors, *Syntax and Semantics III: Speech Acts*, pages 41–58. Academic Press, 1975.

[22] B. Grosz and S. Kraus. Collaborative plans for complex group action. *Artificial Intelligence*, 86(2):269–357, 1996.

[23] B. J. Grosz and C. L. Sidner. Plans for discourse. In P. R. Cohen, J. Morgan, and M. E. Pollack, editors, *Intentions in Communication*, pages 417–444. MIT, 1990.

[24] J. Hobbs, M. Stickel, D. Appelt, and P. Martin. Interpretation as abduction. *Artificial Intelligence*, 63:69–142, 1993.

[25] H. Kamp and U. Reyle. *From Discourse to Logic*. Kluwer, 1993.

[26] H. Kamp and A. Rossdeutscher. DRS-construction and lexically driven inference. *Theoretical Linguistics*, 20:97–164, 1994.

[27] E. Krahmer and M. Theune. Efficient context-sensitive generation of referring expressions. In K. van Deemter and R. Kibble, editors, *Information Sharing: Reference and Presupposition in Language Generation and Interpretation*, pages 223–265. CSLI, 2002.

[28] A. Kyburg and M. Morreau. Fitting words: Vague words in context. *Linguistics and Philosophy*, 23(6):577–597, 2000.

[29] Hector J. Levesque. Foundations of a functional approach to knowledge representation. *Artificial Intelligence*, 23(2):155–212, 1984.

[30] D. Lewis. Scorekeeping in a language game. In *Semantics from Different Points of View*, pages 172–187. Springer, 1979.

[31] S. Neale. *Descriptions*. MIT, 1990.

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

[33] P. Piwek. *Logic, Information and Conversation*. PhD thesis, Eindhoven University of Technology, 1998.

[34] M. E. Pollack. Plans as complex mental attitudes. In P. R. Cohen, J. Morgan, and M. E. Pollack, editors, *Intentions in Communication*, pages 77–103. MIT, 1990.

[35] M. E. Pollack. The uses of plans. *Artificial Intelligence*, 57:43–68, 1992.

[36] E. Reiter and S. Sripada. Should corpora texts be gold standards for NLG? In *INLG*, pages 97–104, 2002.

[37] J. R. Searle. Indirect speech acts. In P. Cole and J. Morgan, editors, *Syntax and Semantics III: Speech Acts*, pages 59–82. Academic Press, 1975.

[38] D. Sperber and D. Wilson. *Relevance: Communication and Congition*. Harvard, 1986.

[39] R. Stalnaker. Presuppositions. *Journal of Philosophical Logic*, 2(4):447–457, 1973.

[40] M. Stokhof and J. Groenendijk. Dynamic semantics. In R. Wilson and F. Keil, editors, *MIT Encyclopedia of Cognitive Science*. MIT, 1999.

[41] M. Stone. Representing communicative intentions in collaborative conversational agents. In B. Bell and E. Santos, editors, *AAAI Fall Symposium on Intent Inference for Collaborative Tasks*, 2001.

[42] M. Stone. Communicative intentions and conversational processes in human-human and human-computer dialogue. In J. Trueswell and M. Tanenhaus, editors, *World-Situated Language Use*. MIT, 2003.

[43] M. Stone and C. Doran. Sentence planning as description using tree-adjoining grammar. In *ACL*, pages 198–205, 1997.

[44] M. Stone and R. H. Thomason. Context in abductive interpretation. In J. Bos, M. E. Foster, and C. Matheson, editors, *EDILOG 2002*, pages 169–176, 2002.

[45] R. H. Thomason. Accommodation, meaning and implicature. In P. R. Cohen, J. Morgan, and M. E. Pollack, editors, *Intentions in Communication*, pages 325–363. MIT, 1990.

[46] R. van der Sandt. Presupposition projection as anaphora resolution. *Journal of Semantics*, 9(2):333–377, 1992.