ELSEVIER

Journal of Pragmatics 32 (2000) 1719–1741

journal of PRAGMATICS

www.elsevier.nl/locate/pragma

Perlocutions: The Achilles' heel of speech act theory☆

Daniel Marcu*

Information Sciences Institute and Department of Computer Science, University of Southern California, 4676 Admiralty Way, Suite 1001, Marina del Rey, CA 90292–6601, USA

Received 22 February 1999; revised version 4 October 1999

Abstract

This paper criticizes previous approaches to perlocutions and previous formalizations of perlocutionary effects of communicative actions by showing that some of their fundamental assumptions are inconsistent with data from communication studies, psychology, and social studies of persuasion. Consequently, it argues for a data-driven approach to pragmatics, one that permits pragmatic theories to be falsified and improved. The paper also offers an introductory account of a formal theory that can explain the difference in persuasiveness between messages that are characterized by the same set of locutionary and illocutionary acts; and the difference in persuasiveness of the same message with respect to different hearers. The formal theory is developed using the language of situation calculus.

Keywords: Perlocutions; Persuasion; Speech acts; Situation calculus

1. Introduction

Since Austin published his most influential book (1962), researchers in linguistics, computational linguistics, and artificial intelligence have analyzed, challenged, and modified to an impressive extent his initial conception of locutionary and illocutionary acts. In contrast, perlocutionary acts have received much less attention. To my knowledge, apart from a few studies (Cohen, 1973; Campbell, 1973; Gaines,

☆ I am grateful to Graeme Hirst for suggesting that I review the empirical literature on persuasion and for his comments on previous drafts of this work. Most of this research was carried out while I was at the University of Toronto and was supported by the Natural Sciences and Engineering Research Council of Canada and by Technology Ontario.

* E-mail: marcu@isi.edu

0378-2166/00/$ – see front matter
PII: S0378-2166(99)00121-6

1979; Davis, 1979; Gu, 1993; Kurzon, 1998), the linguistics literature provides only marginal discussions with respect to perlocutions (Black, 1969; Searle, 1969; Sadock, 1974; van Dijk, 1977; Bach and Harnish, 1979; Leech, 1983). The computational linguistics and artificial intelligence literature is not much richer either: in most cases, perlocutions are treated only peripherally (Cohen and Perrault, 1979; Allen and Perrault, 1980; Perrault and Allen, 1980; Cohen and Levesque, 1985; Appelt, 1985; Cohen and Levesque, 1990a; Perrault, 1990) or are studied from a predominantly argumentative facet (Zukerman and McConachy, 1995; Reed and Long, 1997; Pautler and Quilici, 1998).

The purpose of this paper is to show that there is a significant gap between empirical data that was derived by researchers in communication studies, psychology, and social studies and the assumptions that underlie current studies of perlocutions. The inconsistencies between empirical data and perlocution theories do not seem to stem only from the distinction between illocutions and perlocutions that worried Austin (1962: 109), but also, to a certain extent, from some assumptions that underlie speech act theory as a whole.

Despite the fact that the data that I present here falsify some of the assumptions that constitute the foundations of speech act theory, my intention is by no means to dismantle it. Rather, the message that I want to convey is that if we want pragmatics to have a substantial impact both on linguistics and on real-world applications, we need to avoid language idealizations and artificially constructed examples, and anchor our research and theories in real-world data. My intent in this paper is not only to offer an example of how empirical data can be used to falsify and improve pragmatic theories, but also to provide a sociolinguistic perspective on speech acts from which applied researchers in multi-agent communication, natural language generation, argumentation, and user modelling could benefit. If inter-agent communication languages (Finin et al., 1997; Shoham, 1993; Sidner, 1994; Cohen and Levesque, 1995), inter-agent negotiation protocols (Sycara, 1989; Cohen and Levesque, 1993; Chu-Carroll and Carberry, 1998; Carberry and Lambert, 1999), and natural language generation systems that argue persuasively (Hovy, 1988; Elhadad, 1992; Maybury, 1993; Marcu, 1996; Reed and Long, 1997; Pautler and Quilici, 1998) and adapt to different users (Hovy, 1988; Bateman and Paris, 1989; Paris, 1991; Moore, 1995; Zukerman and McConachy, 1995) are to become more flexible and proficient, they will eventually have to make use of more knowledge of pragmatics: this paper provides a corpus of such knowledge that is well grounded in empirical experiments.

My inquiry into a data-driven approach to the study of perlocutions, as defined by Austin (1962), starts (section 2) with a brief introduction to speech act theory, perlocutionary effects, and persuasion. Once the background information is laid out, I present and discuss a set of seven fallacies that are characteristic to current linguistic and formal models of perlocutions (see section 3). The fallacies are justified by empirical results in communication studies, psychology, and social studies of persuasion. On the basis of the empirical data that is presented in section 3, I then propose (section 4) an introductory account of a formal theory that can explain the difference in persuasiveness between messages that are characterized by the same set of

locutionary and illocutionary acts; and the difference in persuasiveness of the same message with respect to different hearers.

2. Speech act theory and persuasive communication: Background information

2.1. Perlocutionary acts in speech act theory

Central to speech act theory is the understanding that there are three facets to what can be said:

1. The *locution* is the act *of* saying something. This "includes the utterance of certain noises [the phonetic act], the utterance of certain words in a certain construction [the phatic act], and the utterance of them with certain 'meaning' in the favourite philosophical sense of that word, i.e., with a certain sense and with a certain reference [the rhetic act]" (Austin, 1962: 92).
2. The *illocution* is the act done *in* saying something. More precisely, an illocution explains in what way one is using a locution: "for asking or answering a question, giving some information or an assurance or a warning", etc. (Austin, 1962: 98).
3. The *perlocution* is the act done *by* saying something. "Saying something will often, or even normally, produce certain consequential effects upon the feelings, thoughts, or actions of an audience, or of the speaker, or of other persons" (Austin, 1962: 101).

For example, if you say to me 'Don't smoke!', you do not merely utter the words 'don't' and 'smoke', which subsume the locutionary act, but you also perform an illocutionary act, namely, that of urging, advising, or ordering me not to smoke. If, as a consequence of this utterance, I don't smoke, the perlocution is that you convinced me not to smoke.

In this paper, I make no philosophical claims with respect to the role of intention in interpreting perlocutions and speech acts (Campbell, 1973; Gaines, 1979; Gu, 1993); the difference between 'direct' and 'associated' perlocutions (Cohen, 1973); or the fallacies that follow from a straightforward adoption of the multiplicity, infinity, or causation thesis (Gu, 1993) – the multiplicity thesis asserts that speaker's saying something may produce multiple effects in different hearers; the infinity thesis asserts that almost anything could result from a speech act; and the causation thesis asserts that the hearer's being affected should be treated as a consequential effect of speaker's saying something. My critique of previous approaches to perlocutions and speech act theory relies instead on a set of inconsistencies that I prove to hold between the assumptions on which these theories have been built and data derived from empirical studies in persuasive communication.

2.2. A short introduction to persuasive communication

Researchers in communication studies, psychology, and social studies have shown that, given a speech act, one cannot predict whether that speech act will

achieve its intended perlocutionary effect. Nevertheless, one can use linguistic mechanisms to increase the chances of a speech act and its perlocutionary effect to be successful. Altering the chances of *uptake* for a speech act involves decisions that pertain to a very large number of factors, which range over features that characterize the speaker, hearer, message, communication channel, and destination of the message that is associated with that act. In presenting these factors, I use a concrete task: the creation of a persuasive message that is designed to warn teenagers about the temptation to smoke, which they should resist. This message should not be understood as a cooked-up laboratory example of the kind that I criticize, but rather as an opportunity for the reader to understand and visualize how different mechanisms and techniques modify a small text from one that is neutral to one that is persuasive. The incremental creation of the persuasive message is constructed in such a way that it emphasizes at each step weaknesses and false assumptions that are used in previous approaches to perlocutions and speech act theory.

Behavioral studies (Pfau et al., 1992) show that children are born and live through the first decade of their lives with a natural aversion for smoking. However, during the transition from primary to middle grades, this attitude deteriorates (Evans and Raines, 1982; Pfau et al., 1992) and produces in some adolescents a growing chance of experimental and regular smoking. One way to prevent this is by *inoculation*, i.e., "by exposing the person to a weak dose of the attacking material strong enough to stimulate his defenses but not strong enough to overwhelm him" (McGuire, 1970: 37). According to Miller (1980), who defines a persuasive act as a message that is intended to shape, reinforce, or change the response of another, or others, inoculation *is* a form of persuasive communication. Consequently, an inoculation act *is* an instance of a perlocutionary act, since its result is a message intended to reinforce a response with respect to a certain behavior, in this case, smoking.

Obviously, inoculation is just a particular case of persuasion; but as it will become obvious in section 3, designing a message that falls into the inoculation category raises as many problems as designing one that falls into the more general category, the persuasive one. Moreover, such an enterprise uncovers an important number of issues that pertain to the 'likelihood' of a given message to achieve certain perlocutionary effects.

To simplify further my enterprise, I will deal in this paper only with written texts; this means that this paper makes no claims about perlocutionary effects that are achieved through phonetic acts, or through different pitches, intonations, etc.

3. Persuasive communication: A case study in perlocutionary acts

3.1. The fallacy of considering perlocutionary acts to be consequences of simple locutionary acts

Assume that one is given the task of creating a message that will persuade teenagers to stay smoke free. From a philosophical perspective, analyzing the locutionary, illocutionary, and perlocutionary facets of 'Don't smoke!', 'Stay smoke

free!', or 'Smoking could kill you' is a fascinating exercise. However, if we use any of these messages in the real world, it is very unlikely that they will achieve the intended perlocutionary effect. If we want to be persuasive, we will have to do much more: we will have to give explanations, reinforce the beliefs of the teenagers that are consistent with our goal, refute the beliefs that are not, appeal to emotions, etc. (McGuire, 1968; Stiff, 1994). In other words, we will have to build much more sophisticated texts.

As soon as we accept that perlocutionary effects need not be direct consequences of elementary locutionary or illocutionary acts, we open a whole new realm that is beyond the reach of speech act theory, as it was initially proposed by Austin (1962).[1] Let us assume that in our enterprise of creating a message that will persuade teenagers to stay smoke free, we come up with the following text:[2]

(1) No matter how much one wants to stay a non-smoker, the truth is that the pressure to smoke in junior high is greater than it will be any other time of one's life: only 25% of the young adults will not pick up a cigarette and let curiosity take over. Only about 70% will not become experimental smokers. Of those who will start smoking, about 90% will end up with a pack and a lighter for the rest of their lives. We know that 3,000 teens start smoking each day, although it is a fact that 90% of them once thought that smoking was something that they'd never do.

Text (1) has a much more sophisticated structure than any of the simple, one-clause messages that I considered in the previous paragraph. It starts by contrasting the current attitude that is incompatible with smoking and the future attack from the peer group. The contrast is then made more salient by an enumeration of statistical data that support the argument presented in the introductory sentence.

Despite its richness, we will shortly see that text (1) is still unlikely to achieve persuasion, i.e., it is still unlikely to cause the perlocutionary effect of persuading a teenager to stay smoke free. However, let us assume for a moment that it does so. There are two classes of questions that we can then ask: the first class pertains to the organization of knowledge, beliefs, and attitudes of the hearer, which could explain how and why the perlocutionary effect succeeded. The second class of questions pertains more to the message itself and concerns the procedures that could be applied in order to identify the individual speech acts in a sequence of utterances and the precise determination of the locutionary or illocutionary elements that produced the intended perlocutionary effect. After all, assuming that text (1) is persuasive, what were the factors that produced the effect? Was it the warning about the pressure to start smoking that teenagers were going to be exposed? Was it some of the evidence that supported this warning? Or was it the text as a whole? Speech act theory provides no answers to these questions. In fact, providing answers to them amounts, among other things, to solving the 'many forms for one function' mapping problem

[1] This fallacy has been pointed out by Austin (1962) and Searle (1969) as well. I use it here only for the sake of making my argument complete.

[2] Example (1) is created from a text discussed by Pfau (1995).

(Schiffrin, 1994), which occurs when it is the complete full text or the discourse that has a function, rather than any single sentence. Unfortunately, as noted by Taylor and Cameron, "the question of how illocutionary acts are sequenced in actual episodes of connected speech is not one that looms large in the lives of philosophers ... For the analyst working with natural data, however, it is an issue on a par with that of classification/identification" (Taylor and Cameron, 1987: 58).

3.2. *The fallacy of considering that 'new' information is needed in order to produce perlocutionary effects*

A hidden assumption that is used by researchers in speech act theory (Austin, 1962; Davis, 1979; Gu, 1993) is that perlocutionary effects occur as a consequence of having the hearer understand an utterance that 'brings in' some *new* information, i.e., information that is unknown to the hearer. This assumption constitutes the basis of formal models of perlocutions in the artificial intelligence literature as well. For example, Cohen and Perrault (1979) associate with each speech act two plan operators: one plan operator corresponds to the illocutionary act and has its effects defined in terms of hearer's beliefs about the speaker's beliefs. For instance, the effect of an *inform*(*speaker*, *hearer*, *proposition*) speech act is that the *hearer* believes that the *speaker* believes *proposition* – *believes*(*hearer*, *believes*(*speaker*, *proposition*)). A second plan operator, called 'mediating act', is supposed to bridge the gap between speech acts and their perlocutionary effects. For instance, *convince*(*speaker*, *hearer*, *proposition*) is applicable when the *hearer* believes that the *speaker* believes that *proposition* holds, and has the effect that the *hearer* will believe that *proposition* holds. Hence the effect of *inform*(*speaker*, *hearer*, *proposition*) is *believes*(*hearer*, *proposition*). The same assumption is used by Appelt (1985), who states the effects of an action as knowledge of the hearer about knowledge of the speaker.

Given Cohen and Perrault's and Appelt's formal definitions, which reflect the pervasive interpretation of perlocutions, it is not at all clear how one could account for perlocutionary effects in cases in which a *proposition* is already believed by the hearer, or in which the hearer already believes that the speaker believes that a *proposition* holds. For example, one can smoke despite believing that smoking is unhealthy, or one could refuse to speak louder despite knowing that all other participants in the dialogue want one to speak louder.

In fact, new information by itself rarely changes attitudes (Karlins and Abelson, 1970). On the contrary, if a text is to be persuasive, it will have to remind the hearer of information that she already knows (McGuire, 1968; Holtgrave et al., 1995; Witte, 1995). For example, in our case, we can remind the reader that smoking is a detrimental habit and that it costs money. A key factor in persuasive communication is to go beyond simple reminders, continuously reinforce the beliefs that the receivers hold, and emphasize the cumulative risks that otherwise people tend to underestimate (Holtgrave et al., 1995). If we apply these techniques to increase the persuasiveness of text (1), we can reinforce the beliefs that smoking smells and is bad and that no matter how one feels about smoking now and no matter how bad one

thinks it is, or even if one promises they'll never start, 75% of the teenagers will try it at least once (see text 2 below).

If a speaker wants to be persuasive, he must not only include in the message information that is already known to the hearer, but must also make explicit the conclusions that are implicit (McGuire, 1968). In doing so, he will have to pay attention to the particular place in the text where these conclusions are embedded: explicit conclusions should go either first or last, but not in the middle. If the conclusion is shattering, as in the case of telling a patient that she has cancer, the conclusion should go last because after such news, the patient will not be able to follow the rest of the message (Witte, 1995; Bostrom, 1983). However, the introductory part should not be too long, since when one has bad news, it is inappropriate to 'beat around the bush'.

Formal approaches to perlocutions (Cohen and Perrault, 1979; Appelt, 1985; Sycara, 1989) assume that the hearer is a rational agent that knows everything that pertains to the deductive closure of his knowledge. Therefore, they cannot account for the difference in perlocutionary effect between two messages that differ only with respect to a conclusion that is implicit in one message, but explicit in the other one. To account for this, models that distinguish between the short- and long-term memory of conversational participants (Walker and Rambow, 1994) and models that operate with beliefs that can be assigned different strengths (Galliers, 1992) would need to be considered.

In my ongoing refinement of the inoculation message, I make explicit the conclusions of each part: I use a direct inference – *but something or someone convinced them otherwise* – as the final conclusion of the second paragraph; and a metaphoric comparison between smoking and baseball as the final conclusion. This corresponds to the cue of using a climax ordering (McGuire, 1968) and avoiding the saying of the best things, or the conclusions, in the middle of the text.

Text (2), i.e., the text that is obtained from text (1) when one applies the changes discussed in the last two sections, is still unlikely to be persuasive. However, I present it here in order to give the reader the opportunity to understand better one of the key ideas that I want to emphasize, namely that the persuasiveness of a text is not a feature that can be characterized appropriately by a binary value (yes or no), but rather by a continuum that spans over an interval that has at one extreme texts that are persuasive and at the other extreme texts that are dissuasive, i.e., texts that have the opposite of their intended effect.

(2) No matter how much one wants to stay a non-smoker, the truth is that the pressure to smoke in junior high is greater than it will be any other time of one's life: only 25% of the young adults will not pick up a cigarette and let curiosity take over.
No matter how one feels about smoking now and no matter how bad one thinks it is, or even if one promises one will never start, 75% of the teenagers will try it at least once. Only about 70% will not become experimental smokers. Of those who will start smoking, about 90% will end up with a pack and a lighter for the rest of their lives. We know (that 3,000 teens start smoking each day, although it

is a fact that 90% of them once thought that smoking was something that they'd never do. *But something or someone convinced them otherwise.*
And all this happens despite the fact that teenagers know that smoking is a detrimental habit: It smells, it's bad, and it costs money. As they say in baseball 'three strikes and you're out'.

3.3. *The fallacy of attaching little or no importance to the role of the hearer in the perlocutionary act*

A careful analysis of Austin's definition of perlocutions shows that he assigns to the hearer a passive role in the success of a perlocutionary act. However, as Gu (1993) notes in his critique of the causation and infinity theses, treating the hearer's response act as a simple consequence of speaker's speech act denies the status of agent for the hearer. Researchers in user modelling and communications have gone far beyond this point. *Demassification* (Chamberlain, 1994; Strecher et al., 1994), i.e., the degree to which specialized content can be delivered to different individuals, is nothing but a recognition of the importance of the hearer's features in persuasion. An illustration of the successful applicability of demassification in health communication is provided by the fact that over 400 different HIV/AIDS programs exist today in San Francisco alone, each aimed at a particular subaudience (e.g., gay Hispanic male prostitutes) (Rogers, 1994).

Moreover, recent research in the field of behavioral decision making (Yates, 1990; Holtgrave et al., 1995) shows that persuasion is not a constant function of beliefs and attitudes, but that it rather varies depending on the hearer's level of involvement, i.e., "the motivational state induced by an association between an activated attitute and some aspect of the self-concept" (Johnson and Eagly, 1989: 293), and the *stage of change* that the hearer (recipient) is in. In the stages-of-change model (DiClemente and Prochaska, 1985; Holtgrave et al., 1995), for example, there are five stages through which a recipient is thought to pass as they change their attitudes, beliefs, and behaviors:

1. *Pre-contemplative* – the recipient does not recognize the problem or the need for a change.
2. *Contemplative* – the recipient is seriously thinking about the problem and the need for a change.
3. *Preparation* – the recipient is making a commitment to change and is taking steps to prepare for a change.
4. *Action* – the recipient has already modified their behavior for a period of one day to six months.
5. *Maintenance* – the recipient has changed from six months to an indefinite period.

A recipient may be stalled at one stage and need a 'push' to the next. It is obvious that in order to be persuasive, the content of a message should differ radically according to the stage that the recipient is in: the information provided to a person who is not aware that smoking is unhealthy should be different from that provided to

a person who has been attempting to quit smoking for more than three months. The stages-of-change model, which has been validated empirically, also creates complications concerning the adequacy of a definition of perlocutions in terms of modification in hearer's beliefs or behavior because a good definition will have to explain why a recipient changes her behavior only when she moves from the preparation to the action stage.

The hearer's level of involvement affects the likelihood of messages being persuasive as well. If it is high, the arguments that support the claims in a message should be of high quality; if it is low, it is not the quality of the arguments that counts, but rather, their number: the more arguments, the better (McGuire, 1968; Bettinghaus and Cody, 1987). Another important factor is the moral profile of the hearer: if they are *achievement seekers*, it is better to include in a message information that explains how they can advance and achieve a goal; if they are anxious, it is better to avoid the use of fear appeals because this may lower the persuasive effect.

3.4. The fallacy of assuming the hearer to be a 'rational' agent

Most theories of perlocutions assume that perlocutionary effects can occur only when the hearer recognizes both the speaker's intention (Campbell, 1973; Gaines, 1979) and when the speaker fully understands the message (Austin, 1962; Cohen, 1973; Campbell, 1973; Gaines, 1979; Davis, 1979; Gu, 1993). However, empirical data (McGuire, 1968; Bettinghaus and Cody, 1987) show that this is not the case. In fact, the most popular cognitive models of persuasion, the *elaboration likelihood model* (Petty and Cacioppo, 1986) and the *heuristic model* (Chaiken, 1987), rely on a distinction between *central* and *peripheral* routes to persuasion, and *systematic* and *heuristic* processing, respectively. The central route is marked by a careful scrutiny of the message content: it posits that attitude change is a function of message content and elaboration. The peripheral route reflects an attitude change process that is marked by the association of message recommendations with positive or negative cues in the message environment, without effortful scrutiny of the message content. To my knowledge, no approach to perlocutions is able to account for these differences.

3.5. The fallacy of attaching little or no importance to the role of the speaker in the perlocutionary act

Speech act theory acknowledges only the causative role that a speaker has in producing a perlocutionary act; however, studies in persuasion have shown that the credibility, attractiveness, and power of the speaker and of the source to whom a message is attributed plays a much richer role in achieving persuasion than theories of perlocutions have acknowledged so far. As expected, the relevant features of the speaker are not absolute: they are rather dependent on the hearer (Stiff, 1994) and they are influenced by the groups to which the hearer belongs or wants to belong. It is normal for different hearers to assign to the same speaker a high, moderate, or low credibility.

In our ongoing example, we can increase the chances of achieving persuasion by assigning the message to a high credibility source. For example, instead of saying 'We know that 3000 teens start smoking every day', we can say '*Research tells us* that 3,000 teens start smoking each day'. Teenagers may be more receptive to messages whose sources are highly involved. In our example, the involvement of the speaker could be increased by replacing the formulation 'And all this happens despite the fact that teenagers know that smoking is a detrimental habit' with '*Stupid habit! I don't know why people waste their breath on these cigarettes. You* know that smoking is a detrimental habit'. These changes are shown in text (3) below.

(3) No matter how much one wants to stay a non-smoker, the truth is that the pressure to smoke in junior high is greater than it will be any other time of one's life: only 25% of the young adults will not pick up a cigarette and let curiosity take over.
No matter how one feels about smoking now and no matter how bad one thinks it is, or even if one promises one will never start, 75% of the teenagers will try it at least once. Only about 70% will not become experimental smokers. Of those who will start smoking, about 90% will end up with a pack and a lighter for the rest of their lives. *Research tells us* that 3,000 teens start smoking each day, although it is a fact that 90% of them once thought that smoking was something that they'd never do. But something or someone convinced them otherwise.
Stupid habit! I don't know why people waste their breath on these cigarettes. You know that smoking is a detrimental habit: It smells, it's bad, and it costs money. As they say in baseball 'three strikes and you're out'.

3.6. The dismissal of the role that the structure of locutions plays in the success of perlocutionary acts

The arguments that I have brought, so far, against speech act theory pertain more to what pragmaticists will call external factors. As Gu (1993) would probably say, the issues discussed so far pertain more to the way a transaction is carried out, i.e., pertain to a specialized domain of pragmatics, that of rhetoric. However, in the rest of this section I will show that empirical data that was gathered by researchers in persuasive communication challenge not just external aspects, but also the very basic understanding of speech acts as fundamental units of communication.

Structuring a message for persuasive purposes is very important, because although understanding decreases smoothly as messages become 'less structured', the chance for their perlocutionary effects to be successful decreases rapidly (Darnell, 1963; Bettinghaus and Cody, 1987; Reed and Long, 1997). Depending on hearer's awareness of the pros and cons of a problem, different orderings of the information to be presented should be chosen. For a hearer who is not aware of the cons, to increase the chances of a message being persuasive, it should be composed in such a way that supporting information is presented first, and refuting information

second. For a hearer who is aware both of the pros and cons, the message should first present the arguments against the goal that is intended to be achieved and support it afterwards (McGuire, 1968). In our example, when inoculation is performed, the teenagers are unaware of the pressure that they are going to be exposed to: therefore, it is better to move the third paragraph in text (3) to the beginning of the message (see text 4 below). In this way, we support better the aversion for smoking that teenagers have, and prepare them for the presentation of the arguments that refute this aversion. In the field of natural language generation, Reed and Long (1997) and Marcu (1997) have acknowledged the importance of ordering in the generation of persuasive texts; however, the mechanisms that they propose do not provide a full description of the relation between argumentative and rhetorical structures and perlocutions.

In other domains, other techniques may be useful. For example, if some arguments are disagreeable while others are agreeable, it is better to present the agreeable ones first (McGuire, 1968). Section 3.2 above discussed the variations in perlocutionary effects that are observable when different orderings are chosen with respect to the position of the conclusion in an argument.

The persuasive technique of *framing* seems to be one that could give plenty of trouble to speech act theorists. Despite controversies that exist with respect to psychological explanations of fear appeals and the relation between fear and persuasion (Hale and Dillard, 1995), researchers in communication agree that in some domains, in order to persuade, a message has to be framed in such a way that bad consequences are emphasized (Meyerowitz and Chaiken, 1987; McNeil et al., 1982). Hence, although *25% of the young adults will not start smoking* and *75% of the young adults will start smoking* have the same meaning, the second one is more likely to achieve persuasion. Given the equivalence of the locutionary acts in these examples, and given exactly the same speaker and hearer, speech act theory has no ability to distinguish between the different perlocutionary effects that are associated with them.

If we apply the techniques described above to our ongoing example, we obtain text (4):

(4) *Stupid habit*! I don't know why people waste their breath on these cigarettes. We know that smoking is a detrimental habit. It smells, it's bad, and it costs money. As they say in baseball 'three strikes and you're out'.
No matter how much one wants to stay a non-smoker, the truth is that the pressure to smoke in junior high is greater than it will be any other time of one's life: *75%* of the young adults *will pick up* a cigarette and let curiosity take over.
No matter how one feels about smoking now and no matter how bad one thinks it is, or even if one promises one will never start, 75% of the teenagers will try it at least once. About *30% will become* experimental smokers. Of those who will start smoking, about 90% will end up with a pack and a lighter for the rest of their lives. Research tells us that 3,000 teens start smoking each day, although it is a fact that 90% of them once thought that smoking was something that they'd never do. But something or someone convinced them otherwise.

3.7. The fallacy of considering speech acts to be the basic units of communication

The assumption that the speech act is the basic unit of communication (Searle, 1969) is strongly refuted by an important number of persuasive techniques that apply at levels that are much more finely grained than speech acts. If we take speech act theory literally, there is no possibility of explaining why given a common set of conditions, speech acts that are characterized by the same locutionary and illocutionary acts, yield different perlocutionary effects. Let us consider a few examples.

People have troubles reasoning with probabilities (Holtgrave et al., 1995) and there is great disagreement about how different probabilistic values are linked with qualitative adjectives (Mostellor and Youtz, 1990). Because understanding the message is an important prerequisite for the success of a persuasive communication act (McGuire, 1968; Stiff, 1994), in persuasive communication it is better if instead of quantitative discriminators, one uses the qualitative discriminators that are suggested by sociolinguistic research. For example, *virtually always* and *high probability* have average quantitative probability meanings relatively near 99%, while *unlikely*, *doubtful*, and *low probability* are consistent with a probability value of 20% for some audiences (Bryant and Norman, 1980). In changing a quantitative discriminator into a qualitative one, it seems that the locutionary and illocutionary parts of a speech act remain the same. However, empirical data supports the observation that perlocutionary effects differ.

Lexical choice in general plays an important role in persuasive communication. One way to increase persuasiveness is by eliminating all hedges, such as *perhaps*, *maybe*, and *possibly*. The use of metaphors (McGuire, 1968), 'stronger' language and 'pathetic' words (Bettinghaus and Cody, 1987), and specific terms instead of abstract terms (Toulmin et al., 1979: 142) increases the degree of persuasiveness. In our ongoing example, these recommendations will make us prefer *cancer sticks* over *cigarettes*, *messy* over *detrimental*, and *bucks* over *money*.

Increasing the denotative, spatial, and temporal specificity of a text is another way to make a text more persuasive (Parrott, 1995). The denotative specificity can be increased by stating the agent, object, and action explicitly, the spatial immediacy by using demonstratives that are 'close' (*this*, *these*, *here* rather than *that*, *those*, *there*), and temporal immediacy by using the present tense in reference to present events:

(5) Stupid habit! I don't know why people waste their breath on these *things – literally. We used to call them 'cancer sticks'. I* know that *you think that* smoking is a *messy* habit, *right*? It stinks, it's bad, and it costs *bucks*. As they say in baseball "three strikes and you're out'.
No matter how much *you* want to stay a non-smoker, the truth is that the pressure to smoke in junior high is greater than it will be any other time of *your* life. *Three out of four* young adults – *and you are one of them* – will pick up a cigarette and let curiosity take over.
Think about it! No matter how *you* feel about smoking now, no matter how bad *you* think it is, or even if *you swear you*'ll never start, *three of every four of you* will try it at least once. *Almost a third of you* will become experimental smokers,

and *many of you* will end up with a pack and a lighter *full time. Researchers tell* us that 3,000 teens start smoking each day, and *you can count on the fact* that *most* of them, *like you*, once thought that smoking was something that they'd never do. But something or someone convinced them otherwise.

In explaining some of the differences in persuasion between text (4) and text (5), I am sympathetic with the position expressed by Campbell (1973: 291) who argues that Austin's notions of locution and perlocution exclude both the speaker and hearer as determinants of meaning. According to Campbell, a word like *cat* has one meaning for a pet lover, and a different meaning for one who suffers from allergy. It is true that Campbell's position presupposes the existence of a level that is finer grained than locutions, but it still does not explain why for a given hearer, semantically equivalent terms yield different perlocutionary effects.

In discussing the fallacies that underlie the foundations of speech act theory, I have showed how one can take a text, text (1) in our example, and modify it iteratively using heuristics that have been empirically shown to be successful for increasing the likelihood of that text being persuasive. The differences between any two consecutive versions of the text are not prominent. But the difference between text (1) and (5) is. I have not carried out any experiments to assess empirically whether text (5) is more persuasive for a teenage audience than text (1). My intuition is that it is, i.e., it is more likely to produce the intended perlocutionary effect.

4. Towards a formal theory of persuasion

4.1. Theoretical background

Given our empirically demonstrated inability to draw a distinct line between perlocutions that are successful and perlocutions that are not, it seems that attempts such as those of Gaines (1979) and Davis (1979) to classify perlocutionary acts are inadequate. I believe that their two-faceted classifications, which are reminiscent of the Searlean mapping between locutionary and illocutionary acts and which propose sets of necessary and sufficient conditions that make the performance of a perlocutionary act successful, cannot capture the continuum that characterizes the likelihood of a message to be persuasive, a continuum exemplified by texts (1) to (5). Hence, instead of developing a theory in which I can discuss whether a perlocutionary act is successful, I will develop a theory in which I can estimate the likelihood of a perlocutionary effect being achieved. I believe that such a theory can more easily accommodate the empirical findings that I presented in section 3.

Explaining how the beliefs and attitudes of hearers are organized and revised as a consequence of being exposed to certain messages are knowledge representation issues that go beyond pragmatics and that I do not address in this paper. Instead, I prefer to give an introductory account of a formalism that accommodates the continuum of likelihoods of messages being persuasive and that can be used to compare the persuasiveness of different messages.

Central to the formalization of persuasiveness that I propose here is the acknowledgment of the roles that the speaker and hearer, and the structural, locutionary, and illocutionary facets of a message have on increasing or decreasing the likelihood of that message being persuasive. In order to be consistent with the empirical data discussed in section 3, the formalization will account for perlocutionary effects not only at what is traditionally called the speech act level, but also at both finer and coarser grained levels. The finer grained level will account for the effects that various lexical constructs have, while the coarser level will account for the effects that various rhetorical orderings have.

The formalization that I propose has its roots in previous work on action-based theories of language (Cohen and Perault, 1979; Allen and Perrault, 1980; Perrault and Allen, 1980; Cohen and Levesque, 1985; Appelt, 1985; Cohen and Levesque, 1990a,b), which is consistent with Austin's view in assuming that communication is a form of action and that coherent texts can be generated/understood in terms of the sequences of actions that subsume them.

In outlining the formalization, I rely on a variant of the language of situation calculus that was developed by Reiter (1991, 1999). The situation calculus is a first-order language for representing dynamically changing worlds in which all of the changes are the result of named *actions* performed by some agent. Throughout this paper, I use the convention that free variables are universally quantified, that variables start with lower-case letters and constants with upper-case letters. In the situation calculus, terms are used to represent states of the world, i.e., *situations*. If a is an action and s a situation, the result of performing a in s is also a situation, represented by $do\ (a, s)$. The constant S_0 is used to denote the initial situation. Relations whose truth values vary from situation to situation, called *fluents*, are denoted by predicate symbols taking a situation term as the last argument. For example, KNOWS (*Speaker*, *P*, *S*) means that in state *S*, the Speaker knows that *P*. Functions whose values vary from situation to situation, called *functional fluents*, are denoted by function symbols taking a situation term as the last argument. For example, PROCESSED-WORDS(*Hearer*, *S*) = 15 means that up to state *S* the *Hearer* has processed 15 words.

4.2. Outline of the theory

As I have shown in section 3, essential for the understanding and producing of perlocutionary effects is the ability to account for effects of actions that pertain to finer grained levels than the level of traditional illocutionary acts. A possible solution is one that is built on the basis of a set of simple *surface-based actions* that pertain to the lexico-grammatical, locutionary level of messages.[3] In this set we will find actions such as USE_NOUN(*noun*, *elementary_message*), USE_VERB(*verb*, *elementary_message*), USE_PRONOUN(*pronoun*, *elementary_message*), START(*elementary_message*), and CLOSE(*elementary_message*), where an *elementary_message* subsumes the locutionary facet of a traditional speech act. Hence, we assume that each

[3] As this section shows, although I use the notion of *surface-based action* in a way that is similar to Appelt's (1985) or Stone and Doran (1997), the effects that I associate to them are much richer.

locution corresponds to a sequence of elementary actions and that each such sequence starts with a START(*elementary_message*) action and ends with a CLOSE(*elementary_message*) action. At the surface level, a CLOSE(*elementary_message*) action could correspond, for example, to the insertion of a period, or question mark.

In addition to these surface-based actions, we also use a simple action called INTERPRET(*elementary_message*). The role of this action is to model the effects that fully realized elementary messages (locutionary acts) may have on the state of the world, the message, and the hearer. These effects can be of illocutionary nature, i.e., they can model, for example, the illocutionary effects of a locutionary act (Searle, 1969); can be of rhetorical nature, i.e., they can model how a locutionary act is integrated into a more complex structure that reflects the discourse organization of all elementary messages that make up a message (text) (Asher, 1993; Kamp and Reyle, 1993); or can be of perlocutionary nature, i.e., they can model how a locutionary act affects the likelihood of a message being persuasive. In this paper, I will focus only on perlocutionary effects.

According to the requirements of the language of situation calculus, each simple action is characterized by a precondition axiom that specifies the conditions in which that action can be performed. Given the nature of the surface actions that I propose here, their precondition axioms will need to make reference not only to the beliefs and knowledge of the speaker, as most locutionary actions do (Cohen and Perrault, 1979; Allen and Perrault, 1980; Perrault and Allen, 1980; Cohen and Levesque, 1985; Appelt, 1985; Cohen and Levesque, 1990a,b), but also to lexico-grammatical constraints. For example, a precondition axiom for USE_VERB(*verb*, *elementary_message*) will specify that in a state *s* the action is possible, POSS(USE_VERB(*verb*, *elementary_message*), *s*) $\leftrightarrow$ *number*(*verb*, *s*) = *number*(*noun*(*elementary_message*), *s*), if and only if *verb* has the same number as the corresponding *noun* in the *elementary_message*. Similarly, a CLOSE(*elementary_message*) action is possible in a state *s* only if there exists a previous state in which the *elementary_message* was started. And an INTERPRET(*elementary_message*) action is possible in a state *s* only if the action that comes immediately before it in the sequence of actions that subsumes the whole message is a CLOSE(*elementary_message*) action.

In order to capture the continuum that characterizes the likelihood of a *message* being persuasive (a *message* consists of a sequence of *elementary_messages*), I introduce a functional fluent, PERSUASIVENESS(*message*, *speaker*, *hearer*, *goal*, *s*), which takes values over the set of integer numbers – *speaker* and *hearer* are the agents involved in the message exchange, *goal* is a persuasion-related goal, such as 'convince hearer to stay smoke free' or 'persuade hearer to buy flowers', and *s* is the state of the world. A high value of PERSUASIVENESS with respect to a certain *message* does not guarantee any response from the *hearer*; it rather estimates the likelihood of eliciting a response. Hence, this approach does not fall prey to the causation fallacy discussed by Gu (1993), because it poses neither necessary nor sufficient conditions between a message and a response. The way the hearer interprets and responds to a message is uncoupled from the message and specific to each individual; as a consequence, the 'effect = act' fallacy (Gu, 1993: 243) that follows from the multiplicity and infinity theses and that leads to associating with

messages a triggering, perlocutionary role independent of the hearer is also avoided.

The functional fluent PERSUASIVENESS is constructed so that it also acknowledges that the same message can have different likelihoods of being persuasive with respect to different persuasive goals, even when the speaker and hearer are fixed. It is entirely possible, that a message has a high likelihood of being persuasive with respect to a speaker convincing a hearer to go to a party, for example, but a low likelihood with respect to convincing the same hearer to get there by bus. By incorporating the variable *goal* into the fluent, we acknowledge that goals are an important facet of persuasive acts as well.

In the process of generating a message, a speaker starts with a communicative goal and plans the sequence of actions that will lead from state S_0, in which the value of the functional fluent PERSUASIVENESS is set to 0, to a state in which the communicative goal is satisfied. In the process of interpreting a message, a hearer reconstructs the sequence of surface actions that generated that message and makes use of INTERPRET actions in order to understand how the elementary messages fit together and affect its beliefs and attitudes. In either case, a message is assumed to be a sequence of surface and interpret actions. From a perlocutionary perspective, each of the actions in such a sequence can have a positive or negative contribution with respect to the likelihood of achieving some perlocutionary effect. The bigger the value associated with the functional fluent PERSUASIVENESS is in the final state S, the more likely is a message to be persuasive.

Depending on the factors that affect the likelihood of a message being persuasive, simple actions can be partitioned into two classes: agent-independent and agent-dependent actions.

Agent-independent actions are actions whose perlocutionary effects are independent of the speaker, hearer, structure of the message, etc. Using specific nouns instead of general ones, and using denotative, spatial, or temporal specific nouns are strategies that can be formalized using actions from this class. For example, the Specific-noun effect axiom (6) states that if it is possible in a state s to perform an action USE_NOUN(*noun*, *elementary_message*) and the *noun* is 'specific', then, in the resulting state, *do*(USE_NOUN(*noun*, *elementary_message*), *s*), the likelihood of the overall *message* being persuasive will be increased. The General-noun effect axiom (7) specifies that the use of a general noun decreases the persuasiveness of the message independent of the speaker, hearer, goal, or structure of the message. In what follows, the function symbols *succ* and *pred* are used to denote the successor and predecessor functions defined over the set of integer numbers. The predicate POSS formalizes the preconditions that need to be satisfied in state s in order for an action to be possible. The symbol $\rightarrow$ stands for material implication. And the symbols *m*, *em*, *sp*, *h*, and *g* are used as abbreviations of *message*, *elementary_message*, *speaker*, *hearer*, and *goal* respectively.

(6) *Specific-noun axiom*

$$\text{POSS}(a, s) \wedge a = \text{USE\_NOUN}(noun, em) \wedge is\_specific(noun) \wedge \text{PERSUASIVENESS}(m, sp, h, g, s) = n \rightarrow \text{PERSUASIVENESS}(m, sp, h, g, do(a, s)) = succ(n)$$

(7) *General-noun axiom*
POSS(a, s) AND a = USE_NOUN(*noun*, *em*), AND *is_general*(*noun*) AND
PERSUASIVENESS(m, sp, h, g, s) = n →
PERSUASIVENESS(m, sp, h, g, *do*(a, s)) = *pred*(n)

Agent-dependent actions are actions whose perlocutionary effects are dependent on the agents that participate in a conversational exchange, on the agents to which locutionary acts are attributed to, or indirectly, on the structure of the message. To illustrate these three sources of dependence, I discuss now a set of effect axioms that are associated with the action INTERPRET(*em*).

Source-sensitive effect axiom (8) specifies that if an INTERPRET(*em*) action is performed in state s (assuming it is possible to perform it – POSS(INTERPRET(*em*), s) is true), and if the elementary message *em* is attributed to a source *source* who has a *High* credibility in the eyes of the hearer h, then in the subsequent state *do*(INTERPRET(em), s), the likelihood of the message being persuasive is increased.

(8) *Source-sensitive axiom*
POSS(a, s) ∧ a = INTERPRET(*em*) ∧
SOURCEOF(*em*) = source ∧
CREDIBILITY(h, *source*) = *High* ∧
PERSUASIVENESS(m, sp, h, g, s) = n →
PERSUASIVENESS(m, sp, h, g, *do*(a, s)) = *succ*(n)

Hearer-sensitive effect axiom (9) specifies that if an INTERPRET(*em*) action is performed in state s, the hearer h already believes in some y, and *em* reinforces the hearer's belief in y, then the likelihood of the message to be persuasive is increased.

(9) *Hearer-sensitive axiom*
POSS(a, s) ∧ a = INTERPRET(*em*) ∧
(∃y)(BELIEVES(h, y) ∧ REINFORCE(*em*, y)) ∧
PERSUASIVENESS(m, sp, h, g, s) = n →
PERSUASIVENESS(m, sp, h, g, *do*(a, s)) = *succ*(n)

Structure-sensitive effect axiom (10) specifies that if an INTERPRET(*em*) action is performed in state s, there is a component y of the overall message m that has been constructed up to state s, and *em* is a conclusion of y in a rhetorical sense, then the likelihood of the message being persuasive is increased.

(10) *Structure-sensitive axiom*
POSS(a, s) ∧ a = INTERPRET(*em*) ∧
(∃y)(PARTOF(y, *message*) ∧ ISCONCLUSION(*em*, y)) ∧
PERSUASIVENESS(m, sp, h, g, s) = n →
PERSUASIVENESS(m, sp, h, g, *do*(a, s)) = *succ*(n)

Obviously, the classification above has primarily a didactic role; it only provides a convenient way of organizing the space of factors that influence the likelihood of

a message being persuasive. A complete axiomatization will be characterized by effect axioms whose antecedents use predicates that characterize both the agents involved, the structure of the message, and its lexical realization. If one enumerates all precondition and effect axioms that pertain to a domain of discourse, then it is possible to automatically derive a set of 'successor state axioms', which completely characterize the way the values of the fluents and the values of functional fluents change as effect of the simple actions (Reiter, 1991, 1999).

4.3. The theory in action

Assuming that one has axiomatized completely a given domain, i.e., provided precondition and effect axioms for all possible actions that are relevant to that domain, one can then also characterize the likelihood of messages being persuasive. Consider for example, the following excerpts from text (4) and text (5) respectively:

(11) No matter how *one* feels about smoking now, no matter how bad *one* thinks it is, or even if *one promises* that *one* will never start, 75% *of teenagers* will try it at least once.
(12) No matter how *you* feel about smoking now, no matter how bad *you* think it is, or even if *you swear* that *you*'ll never start, *three of every four of you* will try it at least once.

If we analyze these texts, we notice that text (12) is identical with text (11) with the exceptions of using more specific denotative terms (*you* instead of *one*), stronger language (*swear* instead of *promise*), and avoiding percentages (*three of every four* instead of 75%). If we consider the two sequences of actions that produced these texts, they will be similar with the exception of the surface actions that characterize the differences that I have mentioned above. A parallel representation of the two sequences in shown below.

(13) a_1; ...; a_i; USE_PRONOUN(*one*); ...; a_k; USE_VERB(*promise*); ...
(14) a_1; ...; a_i; USE_PRONOUN(*you*); ...; a_k; USE_VERB(*swear*); ...

Since each of the specified surface actions that pertain to text (12) increases the likelihood of that text to be persuasive ('you' is a specific pronoun and 'swear' is a strong verb), while the ones associated with text (11) does not ('one' is a general pronoun and 'promise' is a weak verb), in the states that result from executing the two sequences of actions, the value associated with the fluent PERSUASIVENESS will be higher for the sequence that characterizes text (12), i.e., sequence (14), than the value associated with the sequence that characterizes text (11), i.e., sequence (13). Therefore, the proposed formalization enables one to determine among a set of messages, for a given speaker and hearer, which message is most likely to be persuasive.

If the formalization proposed here is used in a natural language setting similar to the one described by Appelt (1985), for example, and if the goal is to generate persuasive text, a planning system could choose among the different plans (sequences

of actions) that satisfy a given communicative goal, the one that is most likely to be persuasive, i.e., the one for which the value associated with the PERSUASIVENESS fluent in the final state is the highest.

In a similar way, on the grounds provided by the actions that belong to the agent-dependent class, one could explain why the same message has different likelihoods to achieve its persuasive effects for hearers who have different beliefs and attitudes. Consider, for example, that we have completely axiomatized two hearer models H_1 and H_2 and consider that they differ only in the way they perceive the credibility of some source S: H_1 believes that S has high credibility, while H_2 does not believe that. The interpretation of the sequence of actions that corresponds to a message (see, for example, the sequence in 15) below) will yield an increase in message persuasiveness when the action INTERPRET(*em*) that is attributed to source S is executed in the context of user model H_1, but no increase in persuasiveness when the same action is executed in the context of user model H_2.

(15) a_1; ...; a_i; INTERPRET(*em*); ...;

Everything else being equal, the sequence of action (15) will be considered more persuasive when interpreted from the perspective of H_1 than the perspective of H_2.

The distinction between surface- and interpret-like actions can also model the difference between peripheral and central routes to persuasion (Petty and Cacioppo, 1986; Chaiken, 1987). If the understanding of a message is formalized to be part of the truth conditions that make an INTERPRET action possible, then the persuasiveness of a message can still be high due to the persuasive effects of surface-based actions, which do not require understanding. However, when a hearer actually pays attention and understands a message, the likelihood of that message being persuasive will be higher because both surface and interpret actions can increase the value associated with the persuasiveness fluent.

5. Conclusion

In this paper, I studied the adequacy of previous approaches to perlocutions from a perspective that is firmly grounded in empirical data that have been obtained by researchers in communication, psychological, and social studies. I studied a particular case of perlocutionary acts, namely, persuasive acts, and showed that there were inconsistencies between the empirical data and the assumptions that perlocution theories have been built upon. The inconsistencies are due to the simple correlation in these theories between illocutionary acts and perlocutionary effects and the neglect or diminution of the importance that the characteristics of the hearer and speaker, the structure of the message, and the linguistic choices and cues have on the likelihood of a message being persuasive.

Given our inability to draw a distinct line between perlocutions that are successful and perlocutions that are not, it seems that attempts such as those of Gaines (1979) and Davis (1979) to classify perlocutionary acts are inadequate. The data pre-

sented in this paper shows that their 'all or nothing' classification of perlocutionary acts, which is reminiscent of the Searlean mapping between locutionary and illocutionary acts and which proposes sets of necessary and sufficient conditions that make the performance of a perlocutionary act successful, cannot capture the continuum that characterizes the likelihood of a message to be persuasive.

To circumvent the problems of current formalizations of perlocutions, I have proposed a framework that can explain the difference in persuasiveness between messages that are characterized by the same set of locutionary and illocutionary acts; and the difference in persuasiveness of the same message with respect to different hearers. In the proposed framework, we pose neither necessary nor sufficient conditions for a message being persuasive. By uncoupling the persuasive effects from the hearer, we hence avoid the 'effect = act' fallacy discussed by Gu (1993), which consists in associating with messages a triggering, perlocutionary effect.

References

Allen, James F. and C. Raymond Perrault, 1980. Analyzing intention in utterances. Artificial Intelligence 15: 143–178.

Appelt, Douglas E., 1985. Planning English sentences. Cambridge: Cambridge University Press.

Asher, Nicholas, 1993. Reference to abstract objects in discourse. Dordrecht: Kluwer.

Austin, John L., 1962. How to do things with words. Cambridge, MA: Harvard University Press.

Bach, Kent and Robert M. Harnish, 1979. Linguistic communication and speech acts. Cambridge, MA: MIT Press.

Bateman John A. and Cécile L. Paris, 1989. Phrasing a text in terms the user can understand. Proceedings of the 11th International Joint Conference on Artificial Intelligence, vol. 2: 1511–1517.

Bettinghaus, Erwin P. and Michael J. Cody, 1987. Persuasive communication (fourth edition). New York: Holt.

Black, M., 1969. Austin on performatives. In: K.T. Fann, ed., Symposium on J.L. Austin, 401–411. London: Routledge.

Bostrom, Robert N., 1983. Persuasion. Englewood Cliffs, NJ: Prentice-Hall.

Bryant G.D. and G.R. Norman, 1980. Expressions of probability: Words and numbers. The New England Journal of Medicine 302: 411.

Campbell, Paul N., 1973. A rhetorical view of locutionary, illocutionary, and perlocutionary acts. Quarterly Journal of Speech 59: 284–296.

Carberry, Sandra and Lynn Lambert, 1999. A process model for recognizing communicative acts and modeling negatiation subdialogues. Computational Linguistics 25(1): 1–54.

Chaiken, Shelly, 1987. The heuristic model of persuasion. In: M.P. Zanna, J.M. Olson and C.P. Herman, eds., Social influence: The Ontario symposium, vol. 5, 3–39. Hillsdale, NJ: Erlbaum.

Chamberlain, Michael A., 1994. New technologies in health communication: Progress or panacea. American Behavioral Scientist 38(2): 271–284.

Chu-Carroll, Jennifer and Sandra Carberry, 1998. Collaborative response generation in planning dialogues. Computational Linguistics 24(3): 355–400.

Cohen, Philip R. and Hector J. Levesque, 1985. Speech acts and rationality. Proceedings of the 23rd Annual Meeting of the Association for Computational Linguistics: 49–60.

Cohen, Philip R. and Hector J. Levesque, 1990a. Intention is choice with commitment. Artificial Intelligence 42: 213–261.

Cohen, Philip R. and Hector J. Levesque, 1990b. Persistence, intention, and commitment. In: P.R. Cohen, J. Morgan and M.E. Pollack, eds., Intentions in communication, 33–69. Cambridge, MA: MIT Press.

Cohen, Philip R. and Hector J. Levesque, 1993. Preliminaries to a collaborative model of dialog. International Symposium on Spoken Dialogue: New Directions in Human and Man-Machine Communication, 261–266.

Cohen, Philip R. and Hector J. Levesque, 1995. Communicative actions for artificial agents. Proceedings of the First International Conference on MultiAgent Systems ICMAS '95: 65–73.

Cohen, Philip R. and C. Raymond Perrault, 1979. Elements of a plan-based theory of speech acts. Cognitive Science 3(3): 177–212.

Cohen, Ted, 1973. Illocutions and perlocutions. Foundations of Language, 9: 492–503.

Darnell, D., 1963. The relation between sentence order and comprehension. Speech Monographs 30: 97–100.

Davis, Steven, 1979. Perlocutions. In: J.R. Searle, F. Kiefer and M. Bierwisch, eds., Speech act theory and pragmatics, 37–55. Dordrecht: Reidel.

DiClemente, C.C. and J.O. Prochaska, 1985. Processes and stages of self-change: Coping and competence in smoking behavioral change. In: S. Shiffman and T.A. Willis, eds., Coping and substance abuse, 319–343. San Diego, CA: Academic Press.

Elhadad, Michael, 1992. Using argumentation to control lexical choice: A functional unification implementation. Ph.D. thesis, Columbia University.

Evans, R.I. and B.E. Raines, 1982. Control and prevention of smoking in adolescents: A psychosocial perspective. In: T.J. Coates, A.C. Petersen, and C. Perry, eds., Promoting adolescent health: A dialogue on research and practice, 101–136. New York: Academic Press.

Finin, Tim, Yannis Labrou and James Mayfield, 1997. KQML as an agent communication language. In: Jeff Bradshaw, ed., Software Agents. Cambridge, MA: MIT Press.

Gaines, Robert N., 1979. Doing by saying: Toward a theory of perlocution. Quarterly Journal of Speech 65: 207–217.

Galliers, Julia Rose, 1992. Autonomous belief revision and communication. In: Peter Gärdenfors, ed., Belief revision, 220–246. Cambridge: Cambridge University Press.

Gu, Yueguo, 1993. The impasse of perlocution. Journal of Pragmatics 20: 405–432.

Hale, Jerold L. and James P. Dillard, 1995. Fear appeals in health promotion campaigns: Too much, too little, or just right? In: E. Maibach and R.L. Parrott, eds., Designing health messages: Approaches from communication theory and public health practice, 65–80. Beverly Hills, CA: Sage.

Holtgrave, David R., Barbara J. Tinsley and Linda S. Kay, 1995. Encouraging risk reduction: A decision-making approach to message design. In: E. Maibach and R.L. Parrott, eds., Designing health messages: Approaches from communication theory and public health practice, 24–40. Beverly Hills, CA: Sage.

Hovy, Eduard H., 1988. Generating natural language under pragmatic constraints. London: Erlbaum.

Johnson, B.T. and A.H. Eagly, 1989. Effects of involvement on persuasion: A meta-analysis. Psychological Bulletin 106: 290–314.

Kamp, Hans and Uwe Reyle, 1993. From discourse to logic: Introduction to modeltheoretic semantics of natural language, formal logic and discourse representation theory.(Studies in Linguistics and Philosophy, vol. 42) Dordrecht: Kluwer.

Karlins, Marvin and Herbert I. Abelson, 1970. Persuasion: How opinions and attitudes are changed. (Second edition.) Berlin: Springer.

Kurzon, Dennis, 1998. The speech act status of incitement: Perlocutionary acts revisited. Journal of Pragmatics 29: 571–596.

Leech, Geoffrey N., 1983. Principles of pragmatics. Cambridge: Cambridge University Press.

Marcu, Daniel, 1996. The conceptual and linguistic facets of persuasive arguments. Proceedings of the ECAI–96 Workshop, Gaps and Bridges: New Directions in Planning and Natural Language Generation: 43–46.

Marcu, Daniel, 1997. From local to global coherence: A bottom-up approach to text planning. Proceedings of the Fourteenth National Conference on Artificial Intelligence (AAAI–97): 629–635.

Maybury, Mark T., 1993. Communicative acts for generating natural language arguments. Proceedings of the Eleventh National Conference on Artificial Intelligence (AAAI–93): 357–364.

McGuire, William J., 1968. The nature of attitudes and attitude change. In: G. Lindzey and E. Aronson, eds., The handbook of social psychology, vol. 3: 136–314. Reading, MA: Addison-Wesley.

McGuire, William J., 1970. A vaccine for brainwash. Psychology Today 3: 36–39.
McNeil, B.J., S.G. Pauker, H.C. Sox and A. Tversky, 1982. On the elicitation for preferences for alternative therapies. The New England Journal of Medicine, 306: 1259–1262.
Meyerowitz, B.E. and S. Chaiken, 1987. The effect of message framing on breast self-examination attitudes, intentions, and behaviors. Journal of Personality and Social Psychology 52: 500–510.
Miller, Gerald R., 1980. On being persuaded: Some basic distinctions. In: M.E. Roloff and G.R. Miller, eds., Persuasion: New directions in theory and research, 11–28. Beverly Hills, CA: Sage.
Moore, Johanna D., 1995. Participating in explanatory dialogues: Interpreting and responding to questions in context. Cambridge, MA: MIT Press.
Mostellor, F. and C. Youtz, 1990. Quantifying probabilistic expressions. Statistical Science 5: 2–34.
Paris, Cécile L., 1991. Generation and explanation: Building an explanation facility for the explainable expert systems framework. In: Cécile L. Paris, William R. Swartout and William C. Mann, eds., Natural language generation in artificial intelligence and computational linguistics, 49–82. Dordrecht: Kluwer.
Parrott, Roxanne L., 1995. Motivation to attend health messages: Presentation of content and linguistic considerations. In: E. Maibach and R.L. Parrott, eds., Designing health messages: Approaches from communication theory and public health practice, 7–23. Beverly Hills, CA: Sage.
Pautler, David and Alex Quilici, 1998. A computational model of social perlocutions. Proceedings of the 36th Annual Meeting of the Association for Computational Linguistics and 17th International Conference on Computational Linguistics (ACL/COLING'98): 1020–1026.
Perrault, Raymond C., 1990. An application of default logic to speech act theory. In: P.R. Cohen, J. Morgan and M.E. Pollack, eds., Intentions in communication, 3, 161–185. Cambridge, MA: MIT Press.
Perrault, Raymond C. and James F. Allen, 1980. A plan-based analysis of indirect speech acts. American Journal of Computational Linguistics 6(3): 167–182.
Petty, Richard E. and John T. Cacioppo, 1986. Communication and persuasion: Central and peripheral routes to attitude change. New York: Springer.
Pfau, Michael, 1995. Designing messages for behavioral inoculation. In: E. Maibach and R.L. Parrott, eds., Designing health messages: Approaches from communication theory and public health practice, 99–113. Beverly Hills, CA: Sage.
Pfau, Michael, S. van Bockern and J.G. Kang, 1992. Use of inoculation to promote resistance to smoking initiation among adolescents. Communication Monographs 59: 213–230.
Reed, Chris and Derek Long, 1997. Content ordering in the generation of persuasive discourse. Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence (IJCAI–97): 1022–1027.
Reiter, Ray, 1991. The frame problem in the situation calculus: A simple solution (sometimes) and a completeness result for goal regression. In: Vladimir Lifschitz, ed., Artificial intelligence and mathematical theory of computation: Papers in honor of John McCarthy, 359–380. San Diego, CA: Academic Press.
Reiter, Ray, 1999. Knowledge in action: Logical foundations for describing and implementing dynamical systems. Unpublished manuscript. Available at http: //www.cs.toronto.edu/cogrobo/book.ps.Z.
Rogers, Everett M., 1994. The field of communication today. American Behavioral Scientist 38(2): 208–214.
Sadock, Jerrold M., 1974. Toward a linguistic theory of speech acts. New York: Academic Press.
Schiffrin, Deborah, 1994. Approaches to discourse. Oxford: Blackwell.
Searle, John R., 1969. Speech acts: An essay in the philosophy of language. Cambridge: Cambridge University Press.
Shoham, Yoav, 1993. Agent-oriented programming. Artificial Intelligence 60(1): 51–92.
Sidner, Candace L., 1994. An artificial discourse language for collaborative negotiation. Proceedings of the Twelfth National Conference on Artificial Intelligence: 814–819.
Stiff, James B., 1994. Persuasive communication. New York: Guilford.
Stone, Matthew and Christine Doran, 1997. Sentence planning as description using tree-adjoining grammar. Proceedings of the 35th Annual Meeting of the Association for Computational Linguistics, 198–205.

Strecher V.J., M. Kreuter, D.J. Den Boer, S. Kobrin, H.J. Hospers and C.S. Skinner, 1994. The effects of computer-tailored smoking cessation messages in family practice settings. The Journal of Family Practice 39: 262–270.

Sycara, Katia P., 1989. Argumentation: Planning other agents' plans. Proceedings of the 11th International Joint Conference on Artificial Intelligence 517–523.

Taylor, T. and D. Cameron, 1987. Analyzing conversation. New York: Pergamon.

Toulmin, S., Richard Rieke and Allan Janik, 1979. An introduction to reasoning. New York: Macmillan.

van Dijk, Teun A., 1977. Text and context. London: Longman.

Walker, Marilyn A. and Owen Rambow, 1994. The role of cognitive modeling in achieving communicative intentions. Proceedings of the Seventh International Workshop on Natural Language Generation: 171–180.

Witte, Kim, 1995. Fishing for success: Using the persuasive health message framework to generate effective campaign messages. In: E. Maibach and R.L. Parrott, eds., Designing health messages: Approaches from communication theory and public health practice, 145–166. Beverly Hills, CA: Sage.

Yates, J. Frank, 1990. Judgement and decision making. Englewood Cliffs, NJ: Prentice-Hall.

Zukerman, Ingrid and Richard McConachy, 1995. Generating explanations across several user models: Maximizing belief while avoiding boredom and overload. Proceedings of the Fifth European Workshop on Natural Language Generation (ENLG–95): 143–161.

Daniel Marcu is a Research Scientist at the Information Sciences Institute and a Research Assistant Professor in the Department of Computer Science, University of Southern California. His published work addresses topics in discouse and text theories, pragmatics, knowledge representation for natural language, and computational linguistics. His current focus is on building theoretical and algorithmic foundations for the discourse parsing and discourse-based summarization of free, unrestricted texts; and on applying discourse and pragmatic theories to real-world natural language applications.