"Agent *A*, Can You Pass the Salt?" The Role of Pragmatics in Agent Communication

Heather Holmback, Mark Greaves, Jeffrey M. Bradshaw*
Applied Research and Technology
The Boeing Company
P.O. Box 3707 MC 7L-43
Seattle, WA 98124-2207
{heather.holmback,mark.t.greaves,jeffrey.m.bradshaw}@boeing.com

October 3, 1998

Abstract

Developers of current agent communication languages (ACLs) have exploited the parallels between agent communication and human communication in developing the syntax and semantics of their languages. In this paper we extend the comparison of agent communication language to natural language by proposing an ACL pragmatics. Just as in natural language, the pragmatics of an ACL is the relationship between the language and its communicative context. Agent communication via ACLs is different from human communication with natural language, and so not all aspects of natural language pragmatics apply. But there are specific areas where it makes sense to talk about ACL pragmatics when we consider an ACL in its communicative context. We discuss two such areas: 1) the cooperative use of speech acts in

*The authors gratefully acknowledge the support of the DARPA CoABS Program (Contract F30602-98-C-0170) and the NASA Ames Research Center (Contract NCA2-2005) while preparing this paper.

ACLs and 2) ACLs in the conversational context. We argue that these pragmatic issues are important considerations in advancing the state of agent communication.

1 Introduction

Agent communication languages (ACLs) traditionally have been defined via a collection of syntactic and semantic specifications for each of the types of allowable messages. This style of analysis has its roots in logic and the formal study of natural languages, and results in a convenient theoretical framework in which to understand the various properties of ACLs. Modern ACLs, such as KQML [5], FIPA [6], and KAoS [3], have also borrowed ideas from the pragmatic theory of natural languages. Specifically, ACLs such as these explicitly incorporate certain concepts from speech act theory. This paper will discuss the framework of natural language pragmatics from which these ideas were borrowed, and investigate how additional ideas from pragmatics can inform aspects of ACL design.

Our paper is in three parts. Section 2 reviews the relevant theory from linguistics, and discusses the basic application of these concepts to ACL analysis. In section 3, we describe two specific areas of ACL application where the analytic framework provided by natural language pragmatics can be helpful. We also propose a set of principles for ACL design and generation. Section 4 concludes with some observations about the benefits of analyzing ACLs in this manner.

2 Natural Language and ACLs

2.1 The Study of Natural Language

The theoretical study of natural languages begins with the division of the domain into three areas: syntax, semantics, and pragmatics. Syntax is the study of the surface expressions of the language, its words, and the rules which dictate how these words combine into legal or well-formed expressions of the language. Semantics and pragmatics combine to account for what the expressions of the language mean. Semantics is often viewed as an account of the core truth conditions of a sentence: the conditions under which the sentence or the proposition it expresses is true. The semantics of an expression therefore defines its literal meaning – that aspect of a sentence's meaning that is common across every context of usage. Pragmatics, on the other hand, is concerned with that aspect of meaning which arises from *specific contexts of use*, and how that context contributes to both the total meaning and the effects of of an utterance. Here, "context" refers to the relevant circumstances in which a specific communicative act occurs, such as previous discourse, physical setting, social setting, properties of the speaker and hearer, shared assumptions, overall communicative and rhetorical goals of the participants, and so forth. Thus, pragmatics addresses additional factors beyond the static truth conditions specified by a given semantics – factors which determine those aspects of natural language interpretation which are dependent on the specific communicative setting.

One major subfield of pragmatics is speech act theory (See [10, 11]). Briefly, an important part of the pragmatic interpretation of a natural language utterance is determining the speech act (or *illocutionary* act) that the expression is used to convey. For example, a sentence like "It is cold in this room" has a syntactic analysis and a literal, semantic meaning which is constant across all of its possible uses (see [12]) – namely, that the temperature in the room is cold relative to the speaker. However, the speech act that a speaker intends to perform by using this sentence depends on the context of its utterance. The sentence could be used to state a fact, request that the listener close a window, warn the listener not to enter the room,

or for other kinds of purposes. In fact, natural language utterances are frequently and consciously used for several purposes at once. Each type of goal that a speaker intends with a particular utterance roughly corresponds to a given type of speech act (declaring, requesting, warning) that the utterance is designed to perform.

2.2 The Theoretical Framework of ACLs

In developing ACLs, agent researchers have generally relied on the sort of theoretical framework described above. So, for example, ACLs are commonly specified with a syntax and a semantics. The ACL syntax is the set of elements that comprise the ACL and the rules that govern how the elements combine to form well-formed messages. Further, in speech-act based ACLs, the syntax of the language includes, at least, a set of specifiable ways to designate a particular speech act which is to be associated with an instance of an ACL message.

ACL designers have frequently adopted the framework of speech acts for their languages because they recognized an important similarity between the way that human languages are used and the goals of an ACL. Human languages support far more than the pure transmission of facts – they function as *actions* which achieve larger goals for the speaker, such as placing a bet or enlisting a helper. Agents, as goal-driven autonomous actors, also require this capability in their language. In classic speech act theory for natural language, certain verbs have been recognized to explicitly name what speech act the speaker is performing when they occur in a special syntactic construction and are being used literally. Such utterances were referred to as *explicit performatives* in Austin [1], since merely by uttering the verb in the right syntactic expression, the speaker is performing the act which the verb names. For example, the utterance "I promise to be there" generally counts as a promissory act (unless context dictates otherwise).[1] It wears its

[1]Not every use of the verb "promise" constitutes a promise. The use of "promise" in "John promised to be there" is not an explicit performative and the utterance of the sentence does not count as a promise. Rather, it is a report of a promise. Further, there are other ways to promise in English that do not involve any use of the verb "promise."

illocutionary force on its sleeve. Current ACLs reflect the consensus that agent communication is best analyzed by viewing the messages which agents exchange as designed to achieve certain ends, on par with all the other actions which an agent might do to achieve its goals. That is, ACL designers chose to analyze agent communication in an *intentional* way. Speech act theory, which was developed to analyze the particular intentional components of individual natural language utterances, was therefore imported into ACL design.

One result of this decision was that the formal analysis of ACL expressions has become very complex, involving an entire taxonomy of relevant theory types. The failure to correctly specify the level of the taxonomy at which a particular theory resides has been the source of much confusion in ACL design. Modern ACL expressions are composed of sub-expressions drawn from different formal languages, and so their analysis involves the interaction of several different theories. For this paper, at least the following three targets of analysis are relevant:

1. **ACL content languages:** An ACL content language will have a syntax which specifies its legal expressions, and a semantics which specifies its truth conditions.[2] For example, many ACLs use KIF as a content language. KIF's syntax is a variant of first-order predicate calculus, and its semantics is derived from standard model theory. In general, ACL content languages are not an appropriate subject for independent pragmatic analysis, because they (similar to the propositional content of natural language expressions) are only used as a part of a larger communicative entity.

2. **ACL speech act designators:** ACLs such as KQML and FIPA also have a method to specify the speech act that was intended by an agent's use of that particular message. The syntax of these speech act designators is typically straightforward – just the choice of one out of a small

"I will be there" is not an instance of using the verb "promise" as a performative, but it can still be an act of promising in an appropriate context.

[2]It is typical to refer to this language as specifying the *propositional content* of an ACL expression, *i.e.* the proposition that is expressed.

set of expressions (called "performative names" in KQML, and "communicative act types" or "message types" in FIPA) denoting allowable speech acts.[3] Note, however, that this syntax typically forbids the construction of ACL expressions which would specify two or more speech acts for a single utterance. However, the semantics of these speech act designators – the meaning conditions which they impose on a message independently of the context of use – can be extremely complex. A formal semantic analysis of KAoS's `OFFER` speech act designator, of the type carried out in [ref smith, cohen] involves a complex account of agent intentions and plans, and is cast in the language of quantified multimodal logic (which itself has a syntax and semantics drawn from logical theory). Pragmatics plays an indirect role here in that the semantics of these designators is typically based on the pragmatic analysis of the speech acts associated with the eponymous natural language performatives. We note, however, that specifying the semantics of the speech act designators is *not* equivalent to specifying the semantics of the ACL expressions which use these designators.

3. **Complete ACL expressions:** Complete ACL expressions are formed by specifying their parts. There is both a syntax and a semantics for these expressions. For KQML and FIPA, the syntax is a LISPy frame-based scheme, and the semantics is an (in almost all cases informal) account of how the speech act designator, propositional content specification, ontology designator, language designator, and so forth, combine to yield the meaning of a particular message. Because complete ACL expressions are the only objects which can constitute a legal utterance

[3]The use of the word "performative" in the KQML community, while intuitive, is not strictly accurate given its standard use in natural language pragmatics (see also [4]). Performatives for Austin were specific English verbs with their own preexisting syntax and semantics. Sentences which contained these verbs in the right syntactic construction could be generally inferred to be used to perform the speech act associated with the performative verb, although there were many other ways to perform the same speech act. So, the use of a performative verb is one way of performing a given speech act; it is not the speech act. In contrast, the KQML community uses "performative" to refer to an entire parameterized KQML message. The performative name (such as `tell` or `ask-all`) is intended to refer to an actual speech act. Because we will be referring to both natural language pragmatics and ACL theory in this paper, we will use the neutral term "speech act designator" for the ACL message type, and reserve "performative" for its traditional use. This admittedly is a fine point.

in an ACL, only these can properly be the subject of pragmatic analysis.

Given this brief history and taxonomy, we can now state more precisely the goal of this paper. We have noted above that pragmatic analysis has already been used in ACL design to specify the *semantics* of speech act designators in an ACL. The purpose of this paper is to investigate how additional elements of the pragmatic analysis of natural language expressions can contribute to the *pragmatic* analysis of complete ACL expressions.

2.3 Applying Pragmatics to ACLs

Ordinary natural language pragmatic theory has to be modified in several ways in order to be applicable to the analysis of complete ACL expressions. Most obviously, the target data of the theory are of different form (natural language utterances *vs.* ACL messages) and, although ACLs have been designed to exploit similarities to natural language where appropriate, they rely on very different syntactic and semantic theories. For example, the pragmatic analysis of an utterance of "I am cold" would allow it to be linked to a set of speech acts, typically including informing and requesting (*i.e.,* if the statement were intended as a means of getting someone to close the window). Much of standard natural language pragmatic theory is designed to accomplish this complex mapping of natural language expressions onto a defined set of illocutionary acts. However, because speech act-based ACLs are deliberately constructed in a way which is supposed to parallel the use of explicit performatives in English utterances, this mapping of ACL messages to speech acts is presumed to be trivial.[4]

However, the pragmatic theory surrounding the human use of language is also designed to account for additional factors besides the pure relationship between speech acts and utterances. Natural language has important rhetorical functions, such as holding the floor or glorifying the speaker. And, more importantly, human language serves critical social functions such

[4]There are important subtleties here, though. See [13].

as politeness, saving face, not revealing too much information, avoiding confrontation or argument, continuing the conversation, ending the conversation, and having any number of emotional effects on the hearer. Consider three classic examples from pragmatic theory:

1. Please pass the salt.
2. Can you pass the salt?
3. Why is there never any salt on this table?.

These three utterances each have a different syntactic form, a different semantic description, and a different pragmatic analysis. But, in a given context, they can all be mapped to the same speech act and propositional content:

A REQUEST B : B give the salt to A.

The choice of the utterance used to convey this meaning is typically made according the politeness requirements of the context. A pragmatic theory for ACLs will not need to account for this type of politeness constraint, because (for now!) politeness is not a feature of communication between agents.

Ignoring for the present these aspects of natural language pragmatic analysis, this paper will focus on two areas of pragmatics which we believe are promising for the analysis of complete ACL expressions. First, an important pragmatic principle called the *Cooperative Principle* [9] has been suggested to explain the ability of speakers and hearers to figure out the intended illocutionary act of an utterance. We believe that an important analog of this principle exists in agent communication. Second, the aspects of natural language pragmatics which address conversation management between humans can be adapted to inform issues of conversation policy design in ACLs.

3 Steps Toward an ACL Pragmatics

In the previous section we discussed the scope of pragmatic analysis in natural language and in agent communication. In this section, we address two types of pragmatic conditions which are important for the design and use of the current generation of ACLs.

3.1 Selecting the Best Speech Act Designator

Our first proposal for an ACL pragmatics concerns the question of how to select the proper speech act designator in such a way that the resulting ACL message will best communicate a given intention. Assuming that the ACL contains an adequate set of speech acts, what principles govern the communicative appropriateness of an ACL expression containing a speech act designator and a particular propositional content? How should an agent programmer compose the correct ACL expression, so that the ACL message can be interpreted by the recipient as conveying the correct intention – no more and no less?

In human communication, the speaker generally intends the hearer to recognize the illocutionary act he is performing. If he is being cooperative, he chooses an utterance that will accomplish this, and simultaneously serve his social and rhetorical goals. Human communication provides an extensive set of linguistic and non-linguistic devices for these purposes. Further, the hearer typically can make use of a very rich situational context when deducing the primary illocutionary act that the speaker intends. So, a request to give the speaker an object X may be achieved in any number of ways, including:

1. I request that you give me X.
2. I want X.
3. X, please.

4. Are you able to give me X?

5. Why don't you give me X?

6. I'm waiting for X.

7. (Pointing at X with a pleading look.)

Any of these utterances might be more or less appropriate depending on the context. In some contexts not all of these examples would count as a request. In a rich context, where it is easy for the hearer to recognize the speaker's intent, an indirect manner of requesting is often more appropriate because it is more polite.

In basic agent communication, however, there are no social and rhetorical needs to be served. The illocutionary act being performed, along with the propositional content, is the sole thing that must be communicated. This is what makes speech acts so attractive as an explicitly-named component of an ACL message, and also what makes the proper selection of a speech act designator so critical. Let us continue our example of an agent requesting X. In most cases of agent communication, the best and most obvious way to request X is to use an ACL expression with something like `REQUEST` as the speech act designator and something like "give me X" as the propositional content. And, although there may be no reason why in principle the agent could not use `INFORM` as the speech act designator and something like "I want you to give me X" as the propositional content, it is less satisfactory for the agent to construct the message in this fashion.

This much seems intuitive. But notice that the choice of `REQUEST` over `INFORM` for an ACL message with this intention does *not* follow from the syntax or semantics of `REQUEST` or `INFORM`.[5] Presumably the semantics of both the ACL expression with `INFORM` and the expression with `REQUEST` would be satisfied in our example. So, what accounts for the intuition that using `REQUEST` is more appropriate than using `INFORM` in this case?

[5]For a slightly different example, see [13].

Our conjecture is that a something analogous to Grice's Cooperative Principle (CP) is at work here. This principle was proposed as part of an explanation of how natural languages achieve their communicative purpose:

> *Make your conversational contribution such as is required, at the stage at which it occurs, by the accepted purpose or direction of the talk exchange in which you are engaged.*

The CP is admittedly vague and not without its critics, but it is generally accepted that speakers and hearers implicitly follow some such principle when they use and interpret natural language. As mentioned above, the CP exploits the fact that there is generally a very rich context in human communication relative to which utterances are always interpreted and this allows very indirect speech acts to be understood.

Our claim is that a similar principle should govern agent communication, but that it should be narrowed to account for a more restricted context and set of goals. We propose the following Agent Cooperative Principle (ACP):

> *An ACL expression sent by an agent should contain that speech act designator which most directly expresses the illocutionary act intended at that point in the communication.*

Just as the CP applies both to the speaker and the hearer using natural language, so would the ACP apply to both the sending and receiving agents. Thus, the receiving agent would assume that the sending agent has used that speech act designator which most directly expresses the illocutionary act intended at that point in the communication, and this would guide understanding and further conversation. If an agent received a message with the speech act designator `INFORM`, the agent should interpret that utterance as intended purely to inform, and not as an indirect request or other speech act. If the receiving agent assumes that the sender is being cooperative, then it would violate the ACP to use `INFORM` if a request was intended.

It also follows from the ACP that the illocutionary act cannot be indicated in the ACL expression other than as the speech act designator, since otherwise it would be possible that the speech act designator denoted an illocutionary act other than the one intended. Thus, according to the ACP, it would be uncooperative to form a request using an ACL expression with `INFORM` as the speech act designator and an indication that a request is being made somewhere in the propositional content. This has important ramifications for the expressivity of the propositional content language, for it provides a reason to block the type of problematic `INFORM`/`OFFER` equivalence that was described in [13]. Basically, the expressive power of the propositional content language of an ACL would need to be constrained so as to forbid reference to speech acts, *i.e.* that it could not implicitly modify those ACL elements that serve as speech act designators for an ACL message.

It should be noted that lack of directness can be a serious problem in the use of natural languages. If the context does not allow the hearer to easily construct the intended illocutionary act, then miscommunication can occur. In low-context situations where politeness is less important than getting the message across, using the strongest, most direct form congruent with the intended speech act is often required. An example of such a context is air traffic control communication. In these cases, a construction with an explicit performative such as "Request permission to land" is used rather than a more indirect request such as "Would it be possible for me to land?" Because current agent communication is a very low-context communicative setting where politeness need not be considered and efficiency and comprehensibility are of prime concern, directness is desirable.

Three objections can be made to the ACP:

1. **The ACP is obvious and intuitive.** We agree that the ACP is a reasonable principle. However, the real interoperability problems faced by fielded agent systems show us that the agent developer community has no current consensus or guidance about how to determine the correct speech act in a given ACL. Further, the formal semantics of ACL expressions or speech act designators does not entail the ACP.
2. **The ACP needlessly limits the expressive power of ACLs.** One

of the purposes in making the speech act designator explicit was to lessen the inferences necessary by the message recipient (see [7]). Without the guarantees provided by the ACP, this burden may well increase, because the content language may be powerful enough to express indirect speech acts. [13] shows that this can happen in subtle ways. We acknowledge that the pragmatic restriction provided by the ACP definitely involves a tradeoff of this sort over the expressive power of the ACL content language, but see the implications of this issue as important for agent researchers interested in the practicalities of fielding agent systems for large-scale applications.

3. **The ACP is hopelessly vague.** The pivotal phrase in the ACP is "most directly expresses." What criteria should we use to determine whether one ACL expression is more direct than another? The remainder of this section addresses this question.

We can begin to address the vagueness objection by borrowing from the English speech act analysis of Vanderveken [14]. Vanderveken describes certain relations between illocutionary acts. One such relation is described as follows:

> "... many illocutionary acts have *stronger conditions of satisfaction* than others, so that whenever they are satisfied in a possible context of utterance, the other illocutionary acts are also satisfied. For example, if a promise to be nice is kept then the assertion that the speaker can be nice is *eo ipso* true. This is why the performative sentence 'I promise to be nice' truth conditionally entails the sentence 'I can be nice!'." (p. 33).

Applying this criterion to the directness of the speech acts performed by ACL messages, we can strengthen our ACP:

> *An agent should use a speech act which has the strongest conditions of satisfaction consistent with its intended communicative purpose at that stage of the communication.*

Call this the Strengthened ACP (SACP). The conditions of satisfaction of a speech act will be part of or derivable from the semantics of the speech act designator. The relative strength of conditions of satisfaction of speech acts will be derivable from the entailment relations between these conditions.

As an example, consider the communicative context where Agent A wishes to have Agent B send it four widgets. Both of the following abstract ACL expressions are consistent with Agent A's communicative intent:

1. `REQUEST`: Send me four widgets
2. `QUERY`: Do you have four widgets

The ACL expression with `REQUEST` has stronger conditions of satisfaction than the ACL expression with `QUERY`, because the conditions of satisfaction of the former include those of the latter (*i.e.,* part of satisfying a request to be sent four widgets is that the receiving agent be able to determine if it has four widgets). Thus, if the goal is to be sent the four widgets, then the `REQUEST` should be the ACL message chosen. However, as should be obvious, if A's communicative intent is to poll B just to determine if B has four widgets, then the `REQUEST` is too strong for the communicative intent and the `QUERY` should be sent.

So, as agent communication moves beyond the simple cases, principles such as SACP will be relevant in governing how agents communicate in a direct and cooperative manner. Detailed analyses of speech acts like that in Vanderveken (1990) and others can provide the conditions of success and satisfaction which should be included in the semantics of the speech acts. Relations like strength of conditions of satisfaction will provide a basis for stating pragmatic constraints on the usage of those speech acts in ACL expressions.

We further note that a broader goal of ACL theory is to determine an optimal inventory of possible speech acts and speech act designators for ACL expressions. This will certainly depend on the different communicative goals which are necessary for agents to accomplish their tasks. There is a

certain minimal subset of these acts which extant ACLs *de facto* agree upon. However, entailment relations between speech acts can also be helpful in deciding on an inventory of speech act designators for an ACL. An ACL must, of course, include designators for the speech acts with the weakest conditions of satisfaction (*e.g.* `INFORM`) because sometimes that is all the agent intends to communicate. But the ACL should also include speech acts with the strongest conditions of success and satisfaction that can be foreseen to apply to a communicative situation, so as to ensure that the communication required can be as direct as possible and conform to the SACP.

3.2 ACL Pragmatics for Conversations

One of the common criticisms of traditional natural language pragmatic analysis is that its practitioners have too often used as data sentences and descriptions of contextual features in isolation, rather than in the context of ongoing interaction between participants. This is unfortunate, since conversational *non sequiturs* "immediately strike us as much more inappropriate than merely syntactical or semantic errors in conversation." [15, p. 36] It is important that we learn from such lessons as we examine the role of pragmatic analysis in ACLs, and hence we turn our attention to a second area: the structure of agent conversations. A conversation policy is a goal-directed set of message sequencing conventions centered around a type of speech act reflecting the initiating agent's intention.[6] The policy prescriptively encodes regularities that characterize communication sequences between agents, constraining and defining the types and succession of permissible messages in a particular kind of conversation. In ACLs with conversation policies, a conversation is a specific sequence of messages among agents based on a mutually-agreed-to conversation policy. The development of conversation policies is an active area of ACL research and involves basic issues such as the number and types of relevant conversation policies; the way conversation policies are individuated, represented, and composed into other policies;

[6]The term *conversation policy* is used in the KAoS agent framework. It is roughly equivalent to FIPA's *interaction protocols* or Jackal's *conversation specifications.*

and the degree of reasoning power required of an agent which follows them. Answering these questions will help to refine a set of conversation policies with a certain structure. But even the most explicit, well-worked out, exhaustive set of conversation policies can not account for all those aspects of agent conversations that we can envision being needed for agent communication. In anything beyond very simple conversations, there will be properties of the conversation that depend on the situation in which the agents are communicating, *i.e.* the context.

As discussed in section 2.3, an ACL pragmatics will involve how actual ACL expressions can be used in the context of communication. ACL pragmatics therefore comes into play when we consider how the context will affect or determine actual agent conversations. Contextual factors might include things like time constraints, moment-to-moment reliability of the channel, resource availability, and the special capabilities of individual agent participants. These factors may not have a large role in very simple conversations, but as agents move beyond simple conversations or engage in dialogues with humans, contextual factors will be important in determining when conversations are initiated and what form they will take. Real-time context might affect the duration and sequencing of individual messages, the number of iterations in an iterative sequence (*e.g.,* offer, counter-offer, counter-counter-offer, etc.), and how long a conversation can and needs to last. In other words, contextual factors can have a large influence on how an actual conversation gets realized.

An obvious target for pragmatic analysis is in the selection or composition of the appropriate conversation policy to govern an intended conversation. But there are many other less obvious application. For example, consider the acknowledgment of an inform message. In some cases, the designer of an agent conversation might decide across the board to not have the receiving agent acknowledge the inform in a given application because it is not critical at that point in the conversation and it costs scarce time and resources. In another application, an acknowledgment of an inform message might always be important for the sending agent. Perhaps its continuing on a task requires that the receiving agent understand and confirm that it was informed. In still other cases, it could be that the agent designer does not know ahead of time if an acknowledgment should be required. Perhaps it depends

on dynamic aspects of the current conversation, such as whether or not some other agent has acknowledged the inform or whether or not there is time to wait for the acknowledgment. In this case, the specification of whether acknowledge is the appropriate reply to inform is partially dependent on the actual state of the conversation and how the messages are being used in a particular context. Such conditions are pragmatic conditions in that they are dependent on how a message type is being used to in the communicative context.

Other pragmatic conditions are not particular to a given conversation policy type, but rather apply to conversations in general. One such pragmatic condition might involve how agents interrupt each other. Certain applications or contexts may benefit from allowing one agent to interrupt another. A given agent communication system may not allow interruption at all, but if it does then there should be some conditions on when and how the interruption can occur to ensure successful communication. The conditions should state what type of messages can be interrupted, when in the conversation interruption is permissible, and how to recover from an interruption. Agents may also need to be able to negotiate if they are going to allow interruptions.

A final example of a pragmatic condition on ACLs involves the control of insertion sequences in a conversation. A given conversation policy may specify the speech act types in a conversation and their relative sequence, but allow for insertion sequences under certain conditions (e.g. clarification) but not under others (e.g. issuing a second request while one is pending). Insertion sequences allow for more flexibility in conversations.

Here we have exemplified a few conditions to include in our ACL pragmatics. The common pragmatic thread involves the use of an ACL in a communicative context, *i.e.* a conversation.

4 Conclusions

Though natural language pragmatic theory has informed the specification of the semantics of ACL speech act designators, we have described two additional areas of ACL design and usage that can benefit from this theoretical orientation: (1) selection of the most appropriate speech act designator, and (2) the selection and unfolding of agent conversations in context. We have motivated these by showing how agent communication and ACLs share certain goals and structure with human communication and natural languages and arguing that aspects of natural language pragmatics relevant to these goals and structure should apply to speech-act based ACLs as well. To illustrate this, we have proposed an Agent Cooperative Principle patterned after Grice's Cooperative Principle, and we have shown how Vanderveken's work can be used to clarify and strengthen the ACP. In this section we discuss further how a theory of ACL pragmatics can contribute to important issues in agent communication.

Interoperability is one of the main worries which surrounds the development and deployment of practical large-scale agent systems. Even working from the same ACL specification, agents developed by one developer group can rarely interoperate successfully with those developed by an independent group. The problems range from ontology and content language issues to disagreements about the meaning and commitments that different ACL messages entail. However, as we have argued, interoperability is not just a matter of using the same set of speech acts designators with the same meaning for those speech acts, it is also a matter of being as cooperative as possible so that the intended speech act of an ACL message can be correctly and efficiently inferred. This is where the SACP provides guidance.

ACL pragmatics is not exhausted by the SACP. There are several other aspects of natural language pragmatics which can help ACL pragmatics. For example, non-trivial agent communication may include complexities relating to the sincerity and reliability of agents, differences between sophisticated and simple agents, and variations in the reliability of transmissions. All these factors are generally assumed to be taken care of or unproblematic in most discussions of agent communication. The assumptions made are anal-

ogous to the assumption of the ideal speaker-hearer in discussions of natural language syntax and semantics. But these assumptions may not be warranted when considering agent communication in actual practice. Non-ideal speaker/hearer/channel/context issues can have a major impact on agent conversations. It is the realm of ACL pragmatics to deal with non-ideal contexts and conditions such as miscommunication and communication repair. These phenomena have been well studied in natural language pragmatics, and we believe that more principles can be found which can be used to guide ACL designers.

An additional poorly understood pragmatic issue that is vital to the creation of sophisticated agents is the usage conditions for particular ACL messages. For example, in human usage, offers always have an implicit or explicit expiration, based on time or other conditions. This issue arises when we consider how utterances intended as offers might be used by agents, or human and agent teams, in a conversational exchange. In those situations, expiry conditions may be very important. We have begun to inventory and analyze conditions like these as they arise in specific conversation policies in the context of a project to build a general agent conversation design tool (see [2, 8]). We believe that exposing agent communication designers to typical usage conditions for a given illocutionary act in context as part of the use of this tool will lead to more effective, efficient, and error free agent communication.

Finally, it is also useful to be aware of the pragmatic issues of natural language as we may want to incrementally bring in more of the social function of language into agent systems. The agents of the future will deal more and more with people, and may therefore need to interpret indirect speech acts, politeness constraints, rhetorical issues, and so on. In order to support this type of vision for agents, we will require a strong and supple ACL pragmatic theory. We believe this paper has given the first steps towards such a vision of agents.

References

[1] J. L. Austin. *How to Do Things With Words*. Clarendon Press, Oxford, 1962.

[2] J. Bradshaw, M. Greaves, and H. Holmback. Mediating representations for an agent design toolkit. In *Proceedings of the 11th Workshop on Knowledge Acquisition, Modelling, and Management (KAW-98)*, 1998.

[3] J. M. Bradshaw. KAoS: Toward an industrial-strength open agent architecture. In J. M. Bradshaw, editor, *Software Agents*. MIT Press/AAAI Press, Cambridge, MA, 1997.

[4] P. Cohen and H. Levesque. Communicative actions for artificial agents. In J. M. Bradshaw, editor, *Software Agents*. MIT Press/AAAI Press, Cambridge, MA, 1997.

[5] T. Finin, Y. Labrou, and J. Mayfield. KQML as an agent communication language. In J. M. Bradshaw, editor, *Software Agents*. MIT Press/AAAI Press, Cambridge, MA, 1997.

[6] Foundation for Intelligent Physical Agents. *FIPA 97 Agent Communication Langauge Specification*, 1.2 edition, 1997.

[7] M. Genesereth. An agent-based framework for interoperability. In J. M. Bradshaw, editor, *Software Agents*. MIT Press/AAAI Press, Cambridge, MA, 1997.

[8] M. Greaves, H. Holmback, and J. Bradshaw. Cdt: A tool for agent conversation design. In *Proceedings of the 1998 Workshop on Software Tools for Developing Agents*, 1998.

[9] H. P. Grice. Logic and coversation. In P. Cole and J. L. Morgan, editors, *Syntax and Semantics 3: Speech Acts*. Academic Press, New York, 1997.

[10] S. C. Levinson. *Pragmatics*. Cambridge University Press, Cambridge, 1983.

[11] J. Searle. *Speech Acts: An Essay in the Philosophy of Language*. Cambridge University Press, Cambridge, 1969.

[12] J. Searle. Indirect speech acts. In P. Cole and J. L. Morgan, editors, *Syntax and Semantics 3: Speech Acts*. Academic Press, New York, 1997.

[13] I. Smith, P. Cohen, J. Bradshaw, M. Greaves, and H. Holmback. Designing conversation policies using joint intention theory. In *Proceedings of the 1998 International Conference on Multiagent Systems (ICMAS-98)*, 1998.

[14] D. Vandervecken. *Meaning and Speech Acts*. Cambridge University Press, Cambridge, 1990.

[15] P. Watzlawick, J. H. Beavin, and D. D. Jackson. *Pragmatics of Human Communication*. W. W. Norton, New York, 1967.