

Towards Structured Libraries of Communicative Acts *

Juan Manuel Serrano

Artificial Intelligence Group

School of Engineering

University Rey Juan Carlos

C/Tulipán S/N, 28933 Móstoles

email: jserrano@escet.urjc.es

Sascha Ossowski

Artificial Intelligence Group

School of Engineering

University Rey Juan Carlos

C/Tulipán S/N, 28933 Móstoles

email: s.ossowski@escet.urjc.es

Abstract

The FIPA Agent Communication Language includes a Library of Communicative Acts, which agents use to interact socially. This library allows agents to exchange information, and to speak about action performing, two generic social functions that are needed in any multiagent system. However, when developing a specific application, the agent interaction should be described in less abstract terms, by means of more expressive communicative acts. In this paper, a structured approach to the design of communicative act libraries for multiagent applications is put forward, which takes into account both, the criteria of reusability and expressiveness. The approach is illustrated for the case of the design of an advisory subcatalogue of communicative acts, appropriate for a large number of multiagent domains such as Intelligent Tutoring, Decision Support or Personalized Agents.

1 Introduction

Agent Communication Languages (ACLs) are considered to be the centerpiece of today's multiagent systems (MAS). This is not surprising, as in most multiagent settings agents need to influence their acquaintances' behaviour through communication. Modern ACLs, such as FIPA-ACL [FIPA, 2000a], are grounded in Speech Act Theory [Austin, 1962][Searle, 1969]. In essence, they provide a catalogue of communicative acts (CA), that artificial agents are supposed to use when communicating with each other, as well as with their human users. A primary goal of a general purpose catalogue, such as FIPA ACL's *communicative act library (CAL)*, is to provide a *standardized* set of CAs with a precise semantics which, if shared by all agents, ensures a smooth "*interoperation*" among them. The decision which CAs to include in a CAL is usually based on a trade-off between generality and expressiveness: CAs of a general-purpose ACL are supposed to be generic enough to be used in a variety

of contexts. FIPA ACL, for instance, abstracts completely from potential domain ontologies or particular social or organizational roles that agents are bound to. By consequence, it supports conversations based on such generic concepts as information exchange and action performing.

Although this approach, and its specific focus on interoperation, seems reasonable at the first glance, it has given rise to several criticisms (e.g. [Pitt and Mamdani, 1999]). In particular, it ignores the fact that the variety of *social functions* that an ACL needs to support increases with the complexity of the domains as well as with the social and organizational structures that agents are involved in. For instance, modeling complex dialogues between control engineers and decision-support agents in the domain of traffic management results rather unnatural when only CAs from the FIPA CAL are used [Serrano and Ossowski, 2000]. In addition, several ad hoc CAs, different from those found in the FIPA CAL, have been proposed whose expressiveness is adapted to support social interactions that are particularly relevant to certain MAS settings, such as different types of negotiation [Sierra *et al.*, 1998][Amgoud *et al.*, 2000], intelligent tutoring [Porayska-Pomsta *et al.*, 2000] and decision support [Serrano and Ossowski, 2000]. Still, a principled analysis of the interplay between the adequacy of CAs for artificial agents and the physical, social and organizational contexts in which they are used is still to come.

In this paper, we aim at exploring the potential of context-specific *extensions* to FIPA's general-purpose CAL: agents in a specific MAS domain use certain well-designed "dialects" of FIPA ACL, while maintaining cross-domain inter-operability through a shared core CAL. The paper is organized as follows: In section 2 we analyse FIPA-ACL as a domain independent ACL, and provide a simple taxonomy that structures FIPA-CAs according to their social function. Section 3 discusses strategies for developing context-specific extensions to FIPA ACL. Section 4 gives an example of such an extension with respect to the design of an advisory subcatalogue of communicative acts. Finally, in Section 5, we present some conclusions respecting potential benefits and drawbacks of our approach.

*Research sponsored by MCyT, project TIC2000-1370-C04-01, and by CAM, project 07T/0011/2000

The catalogue proposed by FIPA contains 22 CAs. The meaning of these CAs is established in term of the FIPA SL language [FIPA, 2000b]. The semantic framework underlying the formalisation of CAs sets out from a first-order language, including modal operators of different mental attitudes (belief, uncertainty and choice) and actions (feasible, done)[Sadek, 1992]. The logical model for the belief and choice operators is a *KD45*-model (with choice and belief necessitation rules). Thus, beliefs and choices are consistent, and closed under logical consequence. Moreover, agents are positively and negatively introspective with respect to their beliefs and choices. Uncertainty is such that the following set of formulae are mutually exclusive [Sadek, 1991]: $\{B_i \neg \phi, U_i \neg \phi, U_i \phi, B_i \phi\}$.

Within this framework, the assumption of rationality is captured by stating different axioms (or properties), which account for the relationships between the agent's mental attitudes [Sadek, 1991]. For instance, the *realism constraint*: $\models B_i \phi \Rightarrow C_i \phi$, avoids agents to adopt preferences which are in conflict with their beliefs. In addition to basic rationality principles and constraints, such as the ones described in [FIPA, 2000a], further axioms may be added to account for a more *cooperative* behaviour between agents, such as the adoption (or transfer) of intentions and beliefs of another agent [Bretier and Sadek, 1997]. However, it should be noted that these cooperative axioms are not part of the current FIPA semantic framework.

With respect to the modeling of actions, their meaning is given in terms of a set of formulae, which is structured in two major parts. The first one establishes the *rational effect* (RE) of the action, i.e. the reasons to plan it. The second one states the *feasibility preconditions* (FP) which must be fulfilled if the action is to be planned. In the case of communicative actions, the RE corresponds to the perlocutionary effect of the CA, and the FPs can be further decomposed into *ability preconditions* and *context-relevance preconditions* [Sadek, 1991]. The following axiom establishes that whenever an agent observes the *consummation* of some action *a* (communicative or not) it believes any persistent feasibility preconditions or effects *p* [Bretier and Sadek, 1997].¹

Axiom 1 $\models B_i(Done(a) \Rightarrow p)$

An analysis of the FIPA CAL from a functional point of view reveals that two major classes of social activities are supported: those referring to the *beliefs* of some agent (*informational* functions), and those centered around the *behaviour* of some agent (*volitional* functions). In the first class, we distinguish between CAs attempting to modify the beliefs of the hearer *j*, such as *inform*, *confirm* and *disconfirm*, and those intended to be used in social interactions that target the beliefs of a speaker *i* (*query-if*, *query-ref* and *subscribe*). Volitional activities are related to commitments to perform some action. As before,

¹It basically correspond to property 5 of the FIPA CAL underlying semantic model [FIPA, 2000a]

within this class we can distinguish between those CAs used in interactions that refer to a potential action of the hearer *j* (*request*, *cfp*, *accept* etc.) and those targeting the speaker *i* (*propose agree*, *refuse*). Table 1 summarizes our functional analysis.

Obviously, informational and volitional social activities are ubiquitous, which accounts for the general applicability of CAs that support them and justifies their inclusion in a general-purpose CAL. However, it is not obvious that the CAs to be included are those specified by FIPA. For instance, to support an exchange of information in principle two performatives, say *tell* and a *query*, would be sufficient. We will just analyse the semantics of the first one, which is given by the following formulae: $FP(\langle i, tell(j, p) \rangle) \equiv B_i p$ and $RE(\langle i, tell(j, p) \rangle) \equiv B_j p$.

Due to the consummation axiom 1, whenever the speaker issued one of these CAs, the hearer could infer that the speaker believes its pre-conditions (conventionally specified by the semantics of the CAs) to be fulfilled. However, these pre-conditions would be rather weak in the case of *tell*, as there would not be any context-relevance precondition, and the only ability precondition would just require the speaker to believe the proposition being conveyed, which accounts for their lack of expressiveness. Still, on the other extreme, CAs which are over-specified with respect to a desired social function would have so many context-relevance preconditions that few situations actually satisfied them. So, they would no longer be useful. FIPA takes an intermediate position in the above trade-off between reusability and expressiveness, by defining three more specialised versions of the above *tell* performative (*confirm*, *disconfirm* and *inform*), which account for more explicit context-relevance preconditions.

3 The structure of a context-specific CAL

In this section we tackle the problem of designing the CAL of a specific Multiagent System. The different design alternatives for such a MAS CAL are outlined first, and criteria to choose among them are introduced. Finally, we propose a simple procedure for identifying CAs for such context-specific CALs leading to a particular structure for extended catalogues.

A MAS CAL can be related to the general-purpose CAL of FIPA in different ways. There are three major alternatives for the design process:

1. to define a new CAL from the scratch, that includes ad-hoc CAs which are specially relevant to the domain (" $MAS\ CAL \cap FIPA\ CAL = \{\}$ ");
2. to stick to the FIPA CAL, but to encode the additional information to be conveyed in the message content (" $MAS\ CAL = FIPA\ CAL$ ");
3. to use the current FIPA CAL plus new subcatalogues of communicative actions (" $MAS\ CAL \supset FIPA\ CAL$ ").

As argued in the last section, the presence and the characteristics of a specific CA in an ACL can be ex-

Informational functions		Volitional functions	
$B_i\phi$	$B_i\phi$	$Done(< j, act >, \phi)$	$Done(< i, act >, \phi)$
<i>inform</i> (ϕ)	<i>query</i> – <i>if</i> (ϕ)	<i>request</i> (a)	<i>propose</i> (a, ϕ)
<i>confirm</i> (ϕ)	<i>query</i> – <i>ref</i> ($\phi(x)$)	<i>cfp</i> ($a, \phi(x)$)	<i>agree</i> (a, ϕ)
<i>disconfirm</i> ($\neg\phi$)	<i>subscribe</i> ($\phi(x)$)	<i>accept</i> (a, ϕ)	<i>refuse</i> (a, ϕ, ψ)
		<i>reject</i> (a, ϕ, ψ)	
		<i>cancel</i> (a)	

Table 1: Social Activities Supported by the FIPA CAL

plained with respect to the relevance of its social function. As long as the social activities to be supported in the MAS are generic, the CAs of a general-purpose ACL such as the FIPA ACL are sufficient. Still, different domains may require different functions from the CAL and, in particular, the more complex the setting to which a MAS is applied, the more specific will be the resulting dialogues. For instance, to describe agent interaction in a complex MAS like, for instance, an intelligent tutoring system, as an exchange of information is too generic. Rather, what tutoring agents do is to *explain* concepts, *correct* the student, etc. The main point here is to keep the level of abstraction at which agents communicate closer to the intuitions of the designers and users of the application. This increased level of expressiveness has important advantages:

1. A more anthropomorphic view of the system to be developed is achieved (a higher-level programming paradigm).
2. More information about the attitudes of the speaker is conveyed in a single message (due to the more complex preparatory conditions, and the consummation axiom). As a result:
 - the task of ensuring the consistency of the beliefs of agents is facilitated;
 - less messages are usually needed so that the traffic load is reduced;
 - interaction protocols can become simpler.
3. From the standpoint of human agent interaction, more information is available to multimodal interaction planners.

If we attended just expressiveness criteria, the first design alternative would be preferred. However, a CAL designed from scratch would severely limit agent interoperability (especially in open systems), and no reuse would be granted. Although reuse and interoperability criteria are met by the second design choice (to use core FIPA CAL), this option suffers from a lack of expressiveness. This can be overcome by the third design alternative: to extend the FIPA CAL with a set of new communicative actions specially suitable to that multiagent domain. In this way, we seek a compromise between expressiveness, interoperability and reusability. The second criteria is guaranteed as the FIPA CAL is maintained as the core of the catalogue. We pursue the third criteria by grouping the new CAs in function-oriented subcatalogues so that

reuse in other multiagent domains can be more easily detected.²

By consequence, the question arises as to how the CAs of such a context-specific extensions to the FIPA CAL can be designed. We propose the following three step procedure:

1. **Identification of Social Activities** The first step is to identify the main social activities that are particularly relevant to the MAS domain. These can be obtained from a pragmatic analysis of several natural language dialogues, representing prototypical interaction patterns between the agents of the system. There are three possible types of social activities: (1) directly supported by the FIPA CAL; (2) not supported, but potentially reusable in other multiagent applications; and (3) particular to the MA application (reuse is not possible). Advisory interactions are an example of the second case, as the corresponding catalogue is applicable to personal assistant applications, decision support agents, intelligent tutoring agents, etc. This subcatalogue will be analyzed in the next section of this paper. Another candidate for reusability is a negotiation by argumentation subcatalogue [Sierra *et al.*, 1998], which might include persuasive communicative actions such as *threaten*, *appeal* and *reward*.
2. **Identification of Communicative Acts** The second task is to identify for each social interaction a subcatalogue of CAs that best support it. This requires a pragmatic analysis (based, for instance, on the work of philosophers of language [Searle, 1969], [Searle and Vanderveken, 1985] or linguists [Wierzbicka, 1987]).
3. **Formalisation of Communicative Acts** The formalisation of the subcatalogues is the last stage, which aims at describing the formal semantics of the new communicative acts. The FIPA-ACL semantic model allows for a well-founded extension of the catalogue.

Table 2 summarizes the proposed structure for the communicative act library of a multiagent application

²Note that the extension of the FIPA CAL is currently promoted by the standardization body FIPA (see [FIPA, 2000a]). However, the difference between our approach and FIPA is that we suggest a principled extension in terms of subcatalogues, structuring the new CAs with respect to social functions.

MAS CAL					
FIPA CAL		Reusable CALs		Particular CAL	
Informational	Volitional	Advisory	Negotiation	...	
<i>inform(ϕ)</i>	<i>request(a)</i>	<i>warn(...)</i>	<i>threaten(...)</i>
...	...	<i>recommend(...)</i>	<i>appeal(...)</i>		
		<i>suggest(...)</i>			

Table 2: Structure of a MAS CAL

as suggested by our design procedure.

4 An Advisory Subcatalogue of Communicative Acts

In this section we attempt to illustrate the design method described above, by providing a set of communicative acts to support advisory interactions. We assume that in some domain the need for this type of social function has been already identified (e.g. in a decision support or a tutoring domain), which allows us to focus on steps two and three of our design process in this section. Accordingly, a set of communicative acts for advisory interactions is identified first, taking as input relevant work from the field of linguistics [Wierzbicka, 1987]. Next, these performatives are formalised, giving rise to the desired subcatalogue of CAs, whose use is finally illustrated by a brief example in the domain of Intelligent Assistants.

4.1 Identification of communicative acts

Our aim in this subsection is to obtain the illocutionary acts for a catalogue to support advisory dialogues. More specifically, we will focus our attention on the acts necessary for the advisor, which is the role an Intelligent Assistant will play. In general, the illocutionary actions which are relevant in this case are those looking for the interest of the addressee. As we seek a high degree of expressiveness, we will look for these actions in speech act catalogues for *natural languages*. Particularly well-suited in this context is the work by Wierzbicka [Wierzbicka, 1987], who aims at an analysis of the English-categorisation of the universe of speech acts. She provides definitions for around two hundreds English speech act verbs, which are classified in groups according their similarity. Definitions of individual speech act verbs are given in a *Natural Semantic Metalanguage* (NSM) which consists of a set of nearly 150 English primitive terms.

In Wierzbicka’s analysis there are seven illocutionary verbs which, in principle, could be used by our artificial advisors: *advise*, *counsel*, *recommend*, *suggest*, *warn*, *propose*, and *offer*.³ Lack of space prevents a complete analysis, so that we will only analyze some of these natural language illocutions. Our goal is to obtain from them a set of artificial illocutionary actions, which will be defined using the NSM language. As we focus on software agents, the resulting definitions of the artificial performatives will be simpler than those of their natural language counterparts.

³Another verb which might be considered as well, is *to explain*.

Warnings

The Wierzbicka’s definition of the performative verb *warning* is not very explicit, as it attempts to capture all possible frames in which the verb *warn* may happen: *warn about*, *warn to*, *warn not to*, *warn against*, etc. From a conceptual point of view there are three components related to all these syntactic frames: the “bad thing” expected to happen to the addressee (“I warn you about the possibility of catching a cold”), its cause (as in “I warn you not to go out – or against going out, that way”) and the possible action to avoid those negative consequences (“I warn you to take the umbrella”).

Our proposal for a performative to be used by an artificial agent, similar to the English illocutionary verb *warn*, focuses on the “bad thing” expected to happen, and its possible cause. This leads to the following definition:

Definition 1 Warn

I think Y will be done
 I think of X as something that could be bad for you
 I think X could be caused by Y
 I say: Y could cause to happen something bad (X) to you
 I say this because I want to cause you to know that X could happen

Advice, Counsel and Recommend

The verbs *advise* and *counsel* are quite similar. Indeed, as Wierzbicka suggests, *counsel* could be defined as “professional advice”. Hence, they differ basically in the supposed mode of achievement (the appeal to knowledge or expertise in the case of counseling, the personal “touch” in the case of advise). *advise* may have moral connotations: the hearer *should* do it because of moral or societal concerns. However, we are only interested in what users *should* do with respect to the *instrumentality* of their actions for their choices. With these restrictions, advise and counsel are similar to *recommend*, as a recommendation is an attempt of the speaker to make the hearer know that some action is good for him in an instrumental sense, i.e. as a “means to an end”.

Moreover, we propose to make explicit the particular “aim” or goal that our agent is supporting, as this is what justifies the recommendation. Also, we want the agent to recommend something whenever it believes that it is currently feasible for us to perform that action. So, this component is included in the model (the first one, in the following definition):

Definition 2 Recommend

I think Y could be done
 I think of X as something that could be good for you

I think X could be caused by Y
 I say: I think Y could be done to achieve X
 I say this because I want to cause you to know that Y
 could be done

4.2 Formalization of the Catalogue

In this section we propose a formalisation of the above communicative actions, based on the *Semantic Language* (SL) used by FIPA[FIPA, 2000b]. The FIPA SL language is obviously more restricted than the *Natural Semantic Metalanguage* proposed by Wierzbicka (not only syntactically, but also semantically). However, as our artificial illocutionary actions are simpler than their natural language counterparts, a proper formalisation is still possible.

Also, note that we make use of some modal operators which are not part of the SL specification[FIPA, 2000b]. These include the choice operator *C*, and the temporal modalities *Henceforth* and *Possible*. Although the choice modality is described in the FIPA ACL semantic framework [FIPA, 2000a] as being part of the FIPA SL language, the FIPA SL specification does not capture it (maybe, because it is not needed in order to define the semantics of any FIPA CA). On the other hand, the temporal modalities are useful abbreviations described in [Sadek, 1992], which could be similarly included in the FIPA SL specification. They have the following meaning: $Possible(p) \equiv \exists e Feasible(e, \phi)$ and $Henceforth(p) \equiv \neg Possible(\neg \phi)$.

Warn

Definition 1 indicates that the propositional content of a warning consists of some state of affairs p and some action a . The speaker expresses his belief that it is possible that action a be done in the future, and that this action will cause p to happen. Moreover, the speaker expresses his belief that p is a bad thing to happen to the hearer. This is obviously related to preferences, and so to the choice operator. A possible formalisation is: $B_i C_j \neg q$:

$$\begin{aligned} <i, warn(j, p, a)> \\ FP: & B_i Possible(Done(a)) \wedge \\ & B_i Henceforth(Done(a) \Rightarrow p) \wedge \\ & B_i C_j \neg p \\ RE: & B_j Possible(p) \end{aligned}$$

In order for the agent to plan this communicative action, we think it convenient to establish the following cooperative axiom, which provides the reasons to perform a warning (similar axioms could be given for the other performatives).

Axiom 2 $\models B_i(Possible(q \wedge C_j \neg q) \Rightarrow I_i B_j Possible(q))$

Of course, if the agent is equipped with the FIPA ACL catalogue then there are other candidates to achieve the given intention: either *inform*, *confirm* or *disconfirm*. However, in the case that one of these CAs and also the *warn* CA, satisfied the preconditions, it seems intuitive that the agent should issue a warning, instead of a (dis)confirmation or inform action. This is because of the Gricean cooperative maxim of quantity and relevance, which favours the use of the CA

that conveys more information relevant to the present context. The additional information provided by the *warn* CA comes from the extra feasibility preconditions, and the observation axiom 1. The “expressiveness principle” can be formalised as follows:

Axiom 3 $\models I_i Done(a_1 | \dots | a_n) \Rightarrow I_i Done(a_j)$, where $a_k, k \in 1 \dots n$, are actions with the same RE (rational effect), and a_j is the more expressive action⁴.

Recommend

The dictum proposed in definition 2 suggests a compound propositional content, consisting of the action a recommended to the hearer, and some state of affairs p . The agent expresses its belief that the recommended action can be done: $B_i Possible(Done(a))$, and that p is a (persistent) rational effect of the action: $B_i Henceforth(Done(a) \Rightarrow p)$. The context-relevant feature, that can be assumed when the agent performs this action, is that p is (currently) good for the hearer, which could be simply modelled as $B_i C_j p$. The essential perlocutionary effect is simply to make the hearer believe that the action can actually be done.

$$\begin{aligned} <i, recommend(j, a, p)> \\ FP: & B_i Possible(Done(a)) \wedge \\ & B_i Henceforth(Done(a) \Rightarrow p) \wedge \\ & B_i C_j p \\ RE: & B_j Possible(Done(a)) \end{aligned}$$

4.3 An Example: a Personalized Trading Assistant

In this section, we illustrate the relevance of the proposed CAs, by an example of their use in the domain of personal assistant agents. We describe the interaction with personalized trading assistants (PTA): agents that play the double role of broker and advisor to non-expert users in electronic markets, whose main goal is to keep their users (the investors) “risk free”. Let us suppose a PTA (acting as a broker) has been told to transfer some order to the market. The PTA (acting as an assistant) analyses the order, and concludes that it implies too much risk, given the investor profile of the client. This context matches the following generic scenario: some agent i believes that some other agent j intends to perform some action a , which has negative consequences for it. Formally, the informational state of the agent includes the following first three formulae as premises; the following propositions can be derived from them (p stands for the PTA, c for the client, o for the order and i for the proposition that the order is inconsistent with the profile of the client):

1. $B_p I_c Done(o)$
2. $B_p Henceforth(Done(o) \Rightarrow i)$
3. $B_p Henceforth(C_c \neg i)$

⁴The more complex feasibility preconditions and rational effect some action has, the more expressive is.

4. $B_p Possible(Done(o))$; from (1)

5. $B_p Possible(i)$; from (4) and (2)

6. $B_p Possible(i \wedge C_c \neg i)$; from (5) and (3)

Then, in accordance with the cooperative axiom 2, the PTA is committed to make its client aware of this situation, i.e., $I_p B_c Possible(p)$. Given this intention, the rationality principles, and the extended FIPA CAL with the advisory subcatalogue, two actions can be planned: $\langle p, inform(c, Possible(i)) \rangle$, and $\langle p, warn(c, i, o) \rangle$. Still, given the greater expressiveness of the latter and an adequate operationalization of axiom 3, a warning act is chosen for execution.

5 Conclusions

This paper is a first step towards the design of structured and expressive CALs. Its first contribution is a review of FIPA ACL so as to identify and justify the structure of its CAL. Second, we have addressed the problem of the design of a CAL for a given multiagent application, providing a simple strategy for the identification and structuring of the CAs that constitute the library. This library is composed of the FIPA CAL, plus a number of different subcatalogues with a relevant social function in the MA domain. Some of these subcatalogues may be generic enough to be reused in other domains. In that way, the structuring of the reusable catalogue of communicative actions, and the expressiveness of the extended subcatalogues, will help multiagent system designers to choose those performatives which best fit their particular applications. These characteristics partially compensate the increase of the catalogue complexity due to the new CAs, with respect to that of the FIPA CAL. Finally, we have proposed a set of CAs to support advisory dialogues. Although these CAs are primitively defined, new CAs for other applications might be compositionally defined in terms of existing FIPA CAs.

We hope that the ideas presented in this paper can contribute to the way in which FIPA ACL will be extended in the future, an issue that has received special attention in the latest version of the definition of FIPA CAL. In particular, they pave the way for practitioners within the MAS community to participate actively in the future evolution of FIPA ACL, by contributing CAs to subcatalogues that tackle social functions necessary in their particular application domains. Finally, the Natural Semantic Metalanguage outlined in section 4 may be useful to facilitate an initial understanding of the semantics of the FIPA CAL, as a language that “mediates” between unrestricted natural language descriptions and the logical language SL.

⁵Note that the PTA acts as a broker as well, so that the possibility of having the order transferred to the market depends on itself. Due to its *competence* in the performance of this kind of actions, it can be inferred that the satisfaction of the client's intention is actually possible

REFERENCES

- [Amgoud *et al.*, 2000] L. Amgoud, N. Maudet, and S. Parsons. Modelling dialogues using argumentation. In *E. Durfee, ed., ICMAS-2000*, pages 31–38. IEEE Press, 2000.
- [Austin, 1962] J.L. Austin. *How to do Things with Words*. Clarendon Press, Oxford, 1962.
- [Bretier and Sadek, 1997] Philippe Bretier and David Sadek. A rational agent as the Kernel of a cooperative spoken dialogue system: Implementing a logical theory of interaction. In Jörg P. Müller, Michael J. Wooldridge, and Nicholas R. Jennings, editors, *ATAL III*, volume 1193 of *LNAI*, pages 189–204, Berlin, August 12–13 1997. Springer.
- [FIPA, 2000a] FIPA. *FIPA Communicative Act Library Specification*. <http://www.fipa.org/specs/fipa00037>, 2000.
- [FIPA, 2000b] FIPA. *FIPA SL Content Language Specification*. <http://www.fipa.org/specs/fipa00008>, 2000.
- [Pitt and Mamdani, 1999] J. Pitt and A. Mamdani. Some remarks on the semantics of fipa's agent communication languages. *Autonomous Agents and Multiagent Systems*, 4:333–356, 1999.
- [Porayska-Pomsta *et al.*, 2000] K. Porayska-Pomsta, C. Mellish, and H. Pain. Aspects of speech act categorisation: towards generating teachers' language. *International Journal of Artificial Intelligence in Education*, 11:254–272, 2000.
- [Sadek, 1991] M. D. Sadek. Dialogue acts are rational plans. In *Proceedings of the ESCA/ETRW Workshop on the structure of multimodal dialogue*, pages 1–29, Maratea, Italie, 1991.
- [Sadek, 1992] M. D. Sadek. A study in the logic of intention. In William Nebel, Bernhard Rich, Charles Swartout, editor, *Proceedings of the 3rd International Conference on Principles of Knowledge Representation and Reasoning*, pages 462–476, Cambridge, MA, October 1992. Morgan Kaufmann.
- [Searle and Vanderveken, 1985] J.R. Searle and D. Vanderveken. *Foundations of illocutionary logic*. Cambridge University Press, 1985.
- [Searle, 1969] J.R. Searle. *Speech Acts*. Cambridge University Press, 1969.
- [Serrano and Ossowski, 2000] J. M. Serrano and S. Ossowski. Domain extensions to the FIPA-ACL: an application to decision support systems. In *Proceedings of the 2nd Spanish Symposium on Distributed Computing (SEID-2000)*, pages 1–14. University of Vigo, 2000.
- [Sierra *et al.*, 1998] C. Sierra, N. R. Jennings, P. Noriega, and S. Parsons. A framework for argumentation-based negotiation. *Lecture Notes in Computer Science*, 1365:177–193, 1998.
- [Wierzbicka, 1987] A. Wierzbicka. *English speech act verbs. A semantic dictionary*. Academic Press, Australia, 1987.