

Communicating Affect in a Multi-Agent System

Yasmine Arafa and Abe Mamdani

Intelligence Communication Systems
Department of Electronic and Electrical Engineering
Imperial College of Science, Technology and Medicine, London SW7 2BT, UK
{y.arafa/e.mamdani}@ee.ic.ac.uk

Abstract

We describe an operational approach to enabling the computational perception required for automated generation of affective behaviour through inter-agent communication in multi-agent real-time environments. The research is investigating the potential of extending current agent communication languages so as they not only convey the content of knowledge exchange but also so they can communicate affective attitudes about the shared knowledge. Providing a necessary component of the framework required for autonomous agent, with affective behaviour, development with which we may bridge the gap between current research in psychological theory and practical implementation of social multi-agent systems.

Introduction

Communicative Agents (CA) are autonomous Interface Agents that employ intelligence and adaptive reasoning methods to provide active, autonomous and collaborative services to the user. The CA metaphor aims towards providing effective highly personalised services. Personifying the CA with a context generated lifelike character is a visual dimension to providing personalised services. The motivation for this type of personalisation is that an animated figure, eliciting quasi-human capabilities, may add an expressive dimension to the agent's communicative features, which can add to the effectiveness and personalisation of the interface and the interactive experience on the whole.

Particularly important capabilities of an agent that must interact with the human user is Believability, in terms of how the agents behaves and expresses itself and that these expressions are appropriate to the context of the interaction [Bates, '94]. Affective computing is "computing that relates to, arises from or deliberately influences emotions" [Picard, '97]. For the agent to effectively achieve believable lifelike behaviour it must have the appropriate knowledge to handle and reason about *affect* so as to produce the believable response.

Since the agent works in a multi-agent system (MAS) environment, it's knowledge about users and the surrounding environment is maintained and exchanged through agent communication. In a real-time multi-agent environment the agent inhabits a world which is dynamic and unpredictable. To be autonomous, it must be able to

perceive its environment and decide its actions to reach the goals defined by its behavioural models. To visually represent the behaviour, the relevant actions and behaviour must be transformed into visual motional actions. Therefore the design of an animated communicative agent system requires components to endow them with perception, behaviour processing and generation, action selection, and behaviour interpretation into believable graphical representation. In this paper we discuss issues related to perception and how it is communicated.

Perception through Agent Communication

In order for the embodied communicative agent to select the appropriate actions, the behavioural system needs to be aware and able to perceive the state of the surrounding environment. This environment includes the user, other service agents, processes and events. Perception of events can be decomposed into different classes: expected or desired events, events occurring in the surrounding world, and potential anticipated events which may or may not occur depending on the progress of a situation. We consider how the communication language can provide a structure that conveys this perception.

Communication is the intentional exchange of information brought about by the production and perception of clues drawn from a shared system of conversational primitives [Russell, '89]. Agent communication languages like KQML [Finin & Fritzson, '94], and FIPA ACL [ACL'99] define a set of general primitives corresponding to speech acts (or communicative acts) [Maybury, '93], so known as performatives. These performatives define the protocols of dialogue between agents. They define *how* communication is conducted, rather than *what* is actually being communicated. In other words, *how* defines the class of actions considered as the basic building blocks of dialogue, and *What* is the core message which is included in the content part of a communicative act.

Inter-agent communication is the means by which conversation is mediated between an agent and the agent society wherein it is situated. We use this communication to acquire the information required for an agent's affective perception on both the *how* and *what* dimensions. We consider the development of CA perception as a process of

two well-defined, separate stages:

- inter-agent interaction between the various entities within a MAS society. We further consider three levels of inter-agent communication at which affect may be conveyed:
 - **content level:** referring to the actual raw message or object intended to be communicated among the entities;
 - **intentional level:** expressing the intentions of agents' communicative acts, usually as performatives of an agent communication language; and
 - **conversational level:** protocols that govern the conversations shared between agents when exchanging dialogue,
- CA affect model: dealing with the agents' inner behaviour (knowledge representation, reasoning, learning, etc.), the agents social and affective behaviour, and the generation of appropriate behaviour states that are transformed into scripts for visual embodiment in the interface.

Accordingly we make the distinction between the affect indicators that may be conveyed at the intentional level, which expresses semantics of communicative act performatives, and at the content level, which expresses facts about the domain. Such layering facilitates the successful integration of the language to applications while providing a conceptual framework for its understanding. Although current primitives could be extended to distinctively convey an affective message, the existing primitives capture many of our intuitions about what constitutes affect from the communicative act irrespective of application. We consider that semantic description provide a model of affect that is useful for modelling the overall behaviour.

Agent Communicative Acts

We use the performatives of agent communication language in terms of their communicative functions rather than being just a communication medium, and we provide a representation of the meaning in terms of how they can be used for affect modelling. This representation forms the base upon which we define a set of inference rules describing the affect process ongoing in agent conversation, and also has implications on how a receiver agent will interpret a conveyed message. These rules take into account the conversation instigator's conception of the context at hand, including the experience, type and degree of relation between both communicants and a model of the receiver's personality.

We consider affect in agent communication at the Intentional level to define the intentional meaning behind a performative coded as a set of goals and beliefs that sender agent A_s has about receiver agent A_r , when in conversation. This intentional meaning is represented in terms of propositions or affect descriptors that are declarative representations of semantic primitives conveying cues

about communication intention, attitude and/or emotions. The meaning of a communicative act is further considered through two dimensions of interpretation [Russell, '89]: Pragmatic Interpretation as the general objective expected in terms of a physical action, and Semantic Interpretation as propositional content. To illustrate, we take an example of a statement like: "Search for item X" made by agent A_s to agent A_r . The Pragmatic Interpretation is the type of action or expected goal task that A_s expects from A_r with regard to the propositional content "item X", where in our case here is the goal that A_r performs service search. Whereas, the Semantic Interpretation in terms of propositional content describes what sender A_s is pertaining, where in our case is the addressee A_r and item X.

We must further narrow the semantic interpretation to induce some affect. For example: the *agree* primitive infers a behaviour of positive nature since both parties are agreeing. Similarly, the *refuse* primitive infers a behaviour of negative nature since something is being rejected. However, in real-time agent conversation communicative acts are not standalone and can not alone determine behavioural inferences because they are uttered within particular context. What may have a positive impact on one individual may have neutral or yet negative impact on another. Moreover, a primitive like inform does not convey any affect. Nevertheless, these interpretations serve as the innate knowledge an agent may start with which can be further focused to the context at hand through the communicative act content. We use the propositional content to provide a more context sensitive interpretation. In our system we present information that specifies context by using meta-data description of a content object using the Content Object Description Language.

Content Language

The content language described provides the specific more context-oriented information. We use meta-level representation to convey more specific semantics of the data being exchanged. For the implemented components to be fully converged and effective in a real-time MAS, the agents must have some knowledge about the surrounding environment which includes affective knowledge about internal and external states. The agents must have a semantic and contextual understanding of the information being exchanged. This requires a theoretical framework for representing knowledge and belief of agents interacting within an agent society. This includes frameworks for representing uncertain knowledge about the surrounding environment evolving with experience and time, awareness of the implication of time constraints, and context-based behaviour. For this purpose we use the notion of meta-level knowledge representation objects which are annotations of data [McGuigan *et al.*, '98], being manipulated between the agents in a MAS. This provides an understanding of the content being handled and hence provides better awareness of the environment.

The abstractions hold physical and conceptual meaning as well as affective states. Offering a structured way of coordinating content and domain data for agents without the agents having to understand the raw content; as well as an automated process for returning information, and feedback regarding the state of the environment to agents in a language the agents can understand.

By including meta-representation of the data in the content part of the communication act, we provide an added dimension to agent comprehension of the actual meaning within the content exchanged. The content objects not only serve as data that agents manipulate but also serve as a basis for defining a social framework. Since any information can contain vast amounts of embedded knowledge, content object descriptions can therefore have potential to convey associated and inferred perception. Adding social attributes that can collectively serve as choices that influence agents' perception about the underlying semantics of the content and attributes to its behavioural and personality traits, provides an improvisational basis that serves as a key aspect in determining a CA's character [Arafa & Mamdani, '99].

Content object meta-descriptions essentially ground five important aspects of information-sharing between agents and entities [McGuigan *et al.*, '98]. The aspects defined to allow the sharing of content across a distributed system and facilitate scalable service provision are as follows:

- Object Description, providing an annotated view of raw data;
- Visual Description, aiding a presentation service in describing the temporal and spatial relationships among different content objects for renderable objects;
- Content Semantics, giving a semantic view of the stored objects;
- Service Descriptions, instantiated to an application domain;
- Initial Affect Indicators that a content object will convey in a non-context sensitive environment. These may change drastically according to the past and immediate states of the agent within an MAS and the state of the environment.

The affect model includes indicators to emotions expected when a defined content object is being handled in a normal or non-context sensitive way. These indicators are associated, using a set of selection rules, with a predefined intended behaviour. This planned behaviour is then modified to incorporate context-sensitive variables to select the appropriate affect-based behaviour. The resulting behaviour is then mapped to a visual (animated) representation. Variations of animated expression (intonation) can be realised by altering and adjusting the intensity of the behaviour according to the immediate situation.

Acknowledgements

The research is supported by a grant from the European Commission ESPRIT projects funding, under grant number EP28831.

Conclusion

The paper discusses work in progress for enabling, building and evaluating CAs with affective behaviour in a real-time multi-agent environment. We argued the need for an operational approach to enabling the computational perception required for the automated generation of affective behaviour through inter-agent communication in multi-agent real-time environments. The research investigates the potential of extending current agent communication content languages so as they not only convey the semantic content of knowledge exchange but also they can communicate affective attitudes about the shared knowledge, by including new communicative acts that can explicitly communicate emotion. We introduced the notion of meta-level knowledge representation, of affective relations, which are annotations of objects being manipulated between agents in a multi-agent system that can convey the current state. The work presented represents a promising first step towards an enabling technology for the creation of real-time, affect-based communicative agents.

References

- Arafa, Y. and Mamdani, E. 1999. Real-time Characters with Artificial Hearts. In *Proceedings of the 5th IEEE Conference on Systems, Man and Cybernetics*, 456-564, Japan: IEEE/SMC Society.
- Bates, J. 1994. The Role of Emotion in Believable Agents. *Communications of the ACM* 37(7), 122-125.
- Finin, T. and Fritzson, R.. 1994. KQML: A Language & Protocol for Knowledge & Information Exchange. In *Proceedings of 19th Intl. DAI Workshop*, 127-136.
- FIPA ACL 1999. <http://www.fipa.org/spec/fipa99spec.htm>
- Maybury, M. 1993. Communicative acts for generating natural language arguments. In *Proceedings of the 11th Conference of AAAI*, 357-364. AAAI Press.
- Mc Guigan, R., Delorme, P., Grimson, J., Charlton, P., Arafa, Y. and Mamdani, E. 1998. The Reuse of Multimedia Objects by Software Agents in the Kimsac System. In *Proceedings of Object Oriented Information Systems Conference*, 156-162.
- Picard, R. 1997. *Affective Computing*. The MIT Press.
- Russell, R. ed. 1989. *AI A Modern Approach*. Addison-Wesley.