SAFIRA- Supporting AFfective Interactions in Real-time Applications

A.Paiva*, E. André**, Y. Arafa±, M. Costa*, P. Figueiredo+, P. Gebhard**, K. Höök §§, A. Mamdani±, C. Martinho*, D. Mourão*, P. Petta §, P. Sengers++, M. Vala*

* INESC/IST, Rua Alves redol, 9, 1000 Lisboa, Portugal
+ ADDETI- Associação para o Desenvolvimento das Telecomunicações Técnicas de Informática, Portugal
** Deutsches Forschungszentrum für Künstliche Intelligenz GmbH – DFKI, Germany
++ GMD – Forschungszentrum Informationstechnik GmbH, Germany
± Imperial College of Science, Technology and Medicine, UK
§ Austrian Research Institute for Artificial Intelligence - ÖFAI
§§ Swedish Institute of Computer Science - SICS

safira@gaiva.inesc.pt

Abstract:
This paper provides an overview of the SAFIRA project, an IST-KA4 project in the area of Affective Interactions. We explain the rationale for providing computers with the ability to understand, model, and express human-like emotions. We present the toolkit for affective computing that we are building, and describe 3 demonstrators that illustrate different approaches to affective interaction that are using our toolkit.

Keywords:
Affect in Human/Machine Interactions.

Project URL:
http://gaiva.inesc.pt/safira/

Introduction

Although emotions have been, for a long time, considered undesirable for rational behaviour, there is now evidence in neuroscience and psychology that place emotions as an important factor in problem solving capabilities and intelligence in general [1] [2]. As a result, organizations now use **emotional intelligence** [3] tests as an evaluation procedure, and a strong new field is emerging in computer science: **affective computing**, *i.e.* "computing that relates to, arises from or deliberately influences emotions" [4].
As humans interact directly with computers, it is critical for this interaction to be empowered with affective components that enrich it and make it more adequate for each individual user, while also providing "machine intelligence" with otherwise impossible capabilities. In the words of Minsky: "the question is not whether intelligent machines can have any emotions, but whether machines can be intelligent without emotions." [5].
Based on these recent findings, SAFIRA addresses the resulting need to enrich the interaction between humans and computers with an affective dimension. We propose that such enrichment can be achieved through the provision of a framework and a toolkit that capitalize on recent research done on the area of affective computing and address the problem of building believable characters.
The main objective of the SAFIRA project is to bring to the software community an enabling technology to support affective interactions. To achieve that, we are investigating a framework to enrich interactions and applications with an affective dimension and implementing a toolkit for affective computing. This toolkit combines a set of components addressing affective knowledge acquisition, representation, reasoning, planning, communication and expression.
We are also creating a set of basic demonstrators of the potential of such systems: 2D and 3D virtual environments shared by synthetic characters and users, and "anthropomorphised" personality-rich personal service assistant based applications. Those demonstrators explore the concept of a virtual environment improved by an emotional channel seamlessly integrated with the audio-visual representation of the environment. Multi-sensory interaction between participants' projected cyber-selves, other participant avatars and the autonomous agents inhabiting the environment will be the means to evaluate the models developed during the project.

In this paper we will describe some of the design decisions concerning the construction of the SAFIRA toolkit and we will provide a brief description of the demonstrators that we are building using that toolkit.

The SAFIRA TOOLKIT

One of the SAFIRA project objectives is to provide domain-independent components that can be put together to build agents/machines capable of exhibiting both emotionally controlled and emotionally expressive behaviour grounded in the internal architecture. To attain this objective we have classified such components as follows:

1. Affective Acquisition, which deals with the mechanisms required in order to obtain information about the user, in particular, concerning affective elements.
2. Emotion Processing, which maintains and handles the mechanisms for embedding emotion and cognition in a software agent.
3. Affective Expression, encapsulates the mechanisms, which are required in order to generate real-time emotional expression.
4. Inter-component Messaging, which provides a rich communication language to share information and service access at a flexible level, thus permitting both sub-components and components to be integrated in a loosely coupled manner.

These components are derived from the background research conducted in highly multi-disciplinary approaches which involved psychology, agent technology, knowledge representation and manipulation, visualisation and moreover, the integration of all these.

The challenges of building affective synthetic agents are not only in the specifications of engines and architectures that can understand, manipulate and generate affective behaviour and expression, but also the integration of both in a common standard API that will allow for the various components to seamlessly interact and exchange common knowledge spaces and data.

To attain that, we have specified an API to bridge the gap between the available underlying emotion engines and the various agent tools. The defined API provides a map between these tools by automating the movement of information from XML Schema structure definitions into the appropriate relational parameters required to generate the appropriate behaviour.

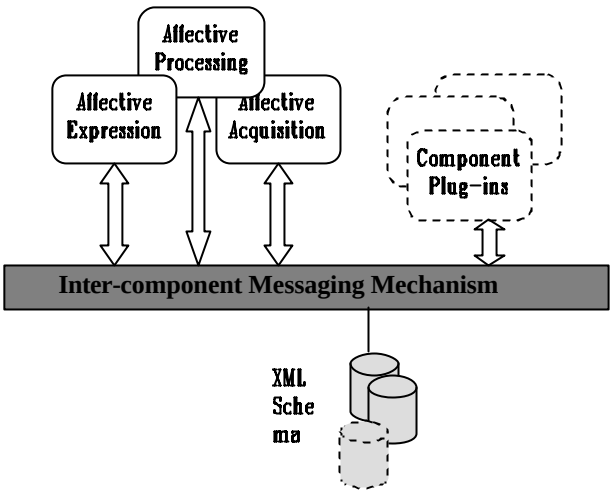

Figure 1 Toolkit architecture

Given this architecture, the toolkit will allow application developers to use API as a glue-like mechanism to tie the various visual and underlying affective behaviour generation tools together, seamlessly, regardless of the platform that they run on and the language they are developed with. To test this, we are developing 3 demonstrators, which will be described next.

The SAFIRA Demonstrators

In order to test all the relevant properties in the toolkit, such as integration, communication between modules, affective processing, etc. we are developing three demonstrators. There are three goals for these demonstrators: to provide a platform for showing the range of possible contexts in affective interaction that can be covered by the toolkit, to provide a testbed for evaluating the correctness and usefulness of the components developed, and to further in different ways our understanding of affect in interaction. Below is a description of these demonstrators.

James the Butler

The James the Butler demonstrator builds on the well-recognised metaphor of a natural, human-like computer character interacting with its user. James the Butler tries to design a *believable* character: one with which users may interact over a longer time span, to which users may build a relationship, and which will react in a comprehensible and relevant manner in dialogue with the user. James the Butler will have affective facial (integrated with the body behaviour rather than the facial affective output module in the toolkit), body and speech behaviours (though bound to the limitations of these modules in the toolkit), and will attempt to model users' emotions and react accordingly. It will use a combination of synthetic emotions and personality

traits. The role affect plays in the system will be to support functional goals for e-commerce, and this provides a constrained field in which to understand affect in interaction.

Figure 2 James the Butler

The believable character is a common metaphor for interaction and AI research in the field, and one that the SAFIRA toolkit must be able to support. Although the field is well-defined, there are still numerous unresolved research challenges, making this an interesting and important field for research for SAFIRA.

FantasyA

James the Butler is an application for e-commerce which builds on computer science research which is relatively mainstream. In recent years, however, researchers in the areas of AI, graphics, and HCI have become aware of the importance and interest of applying advanced computing research to computer games and entertainment. The second demonstrator, FantasyA, explores the possibilities for bringing affective computing to this newly fashionable area. FantasyA is a role-playing game with 3D characters, where affect is the key to advancing in the game. Only when the user is able to make his/her avatar portray the "right" kinds of affective expressions will s/he be allowed to move to the next "island" in the game.

FantasyA presents several additional challenges, which the toolkit addresses. The first is how to enable the user to alter the affective expression of his/her avatar in a natural way; to answer this, we are experimenting with a plush toy as an input device for affect, something that has not been done before in the field of affective interactions. Also, the characters in the game will need more complex affective reasoning components. Finally, FantasyA provides an additional testbed for bodily expression.

Figure 3 FantasyA

FantasyA also builds on the complexity and richness of emotion that we support. The notion of emotion in FantasyA is more complex than in James the Butler, because emotional modelling is connected to a larger narrative context in which the emotions are given meaning and can be understood, thereby relating FantasyA to the Narrative Intelligence research field. In FantasyA, we also have the added complication of several characters interacting, which moves us away from the single, user-facing agents in the James the Butler scenario. The kind of affective interaction supported in FantasyA is more open-ended than in James the Butler, because users can explore the game in a less guided way and interact with richer characters.

Influencing Machine

The two demonstrators described so far use affective interaction for believable characters, one in the established realm of computer science, the other in the newly emerging area applying computer science to games. But if SAFIRA is to address the issues of affective computing more broadly, it cannot limit itself strictly to computer science alone. The goal of affective computing is to allow a humanization of the interface, to make interfaces address a wider range of human experience than those traditionally addressed by computer science. In order to accomplish this, affective computing must address a broader, culturally appropriate notion of emotion, not only as understood by the cognitive science models of emotion which motivate much computer science research in the area, but also in such areas as depth psychology, the humanities, and art. In these areas, emotions are not binary or well-defined; they are soft, fuzzy, ambiguous, and open to interpretation. The goal of the Influencing Machine demonstrator is to show that the SAFIRA toolkit can also address these broader aspects of emotion, and to explore the nature of affective interaction in these more human, more analog, and therefore less clearly defined areas.

The Influencing Machine is a museum installation which, in the manner of an art piece, engages users in an intimate relationship with an affective, intelligent system by letting the user influence its emotional and developmental state. The relationship between the affective input from the user and the affective output of the machine is intentionally made fuzzy, complex, and rich. The user can input emotionally evocative art postcards in a "mailbox". The machine interprets these postcards as if they portray certain emotions which influence both the emotional and the developmental state of the machine. Depending upon which postcards the machine is exposed to, and when they arrive, the development will go in different directions.

Many affective systems with complex agents relate to the user either by trying to extract and respond to the user's emotions, or by allowing the user to directly control an agent in the system. The Influencing Machine explores a different relationship, in which the user neither directly controls nor is modelled by a character, but indirectly influences the system's emotions and thereby its behaviours.

The Influencing Machine expresses itself through child-like drawings, inspired by how a child develops between 9 months and 3 years old, going from scribbling to primitive human figures. The drawings express emotions both through their contents, and through *how* they are drawn, the dynamic properties. The drawings are underscored by emotion-driven sound, developed by a sound designer. The drawings are intended to be intriguing, being given their meaning through users' interpretation and imagination. The Influencing Machine thus allows users to reflect on the notion of synthetic emotions, on what it means to interact with an affective, intelligent system, and on their own emotions and reactions.

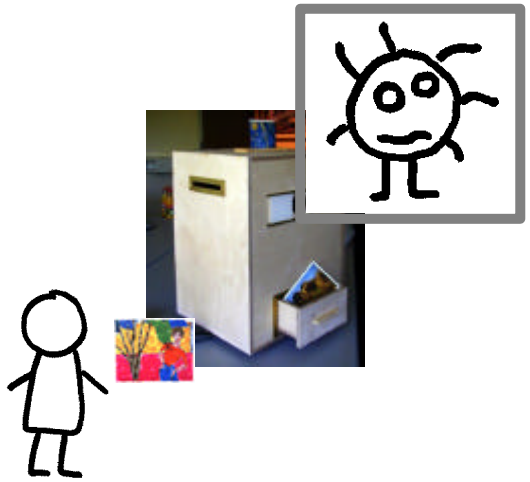

Figure 4 The Influencing Machine

From a user interaction perspective, the Influencing Machine allows SAFIRA to explore affect in an open-ended interaction over time, in which the user's reactions and interpretations are central to the experience.

Summary

These three demonstrators, still under development, provide different dimensions along which to expand our understanding of how to design for affect in interaction with end users. They also address an expanding range of contexts in which affective computing is meaningful, from within computer science, to the boundaries where computer science is expanding into the neighbouring field of computer games, to a broad context in which technology is designed with an eye to including cultural and humanist perspectives.

References

[1] António Damásio: Descartes' Error, Avon Books, 1995.

[2] R. de Sousa: The Rationality of Emotion, MIT Press, 1987.

[3] Daniel Goleman: Emotional Intelligence, Bantam Books, 1997.

[4]Rosalind Picard: Affective Computing, MIT Press, 1997.

[5] Marvin Minsky: The society of Mind, Simon & Schuster, 1985.