

Experiments in Mapping Character Animation to Computational State

Richard Lachman
Interactive Cinema Group
MIT Media Lab
20 Ames St
Cambridge, MA 02139
richlach@media.mit.edu

Abstract

This paper describes a series of experiments mapping the state of a computer system to emotional behaviour in an animated character. Rather than focussing on intelligence in the interface agent, the projects seek ways to use the agent as an avatar for the computer. Facts, recommendations and statistics about the host-system affect behaviour through the subtle cues of gaze-direction, body language, and effects on the environment around the character. The domains of email-filtering, web-browsing and system-performance are the subjects of current explorations, using modified Unix web- and email-servers and a Windows NT client.

Introduction

“Animacy, then, is more properly understood as a framework or way of thinking...Animate thinking stems from a basic need to explain happenings and tell simple stories about them, and a need to fit things into roles in the stories as actors and objects of action”

-Keller, “Reflections on Gender and Science”

Character and setting are important parts of the stories we tell to ourselves, the ways we struggle to understand events around us. The power of an embodied character as computer-interface does not lie in its parsing of natural-language, or the accuracy of its user-models -- all of these features of an intelligent interface can be expressed through other, safer, more conventional means. An animated character,

however, brings into play factors about human nature previous interface-metaphors have rarely exploited: the subtle cues about emotional state and intention expressed through body language in our day-to-day life.

Many researchers have examined ways of creating believable animation and coherent characters, which can display emotion, motivation or internal state. [Bates94][Perlin96][Blumberg95][Laurel91] The goal of our current experiments, however, is to map meaningful stimuli to those same emotions and reactions.

This paper describes a set of experiments in using the emotional and behavioural state of a character to convey useful information using subtle, continuous and non-intrusive means. Rather than traditional dialog-boxes, full-text messages, or generic “alert” tones, we are creating emotionally charged mappings between events and behaviours to act as the medium of communication.

“...the characters can then be used much like an actor is used in a production. They can...both take direction from an outside force and perform reasonable actions autonomously when not directed” [Galyean95]

The goal is not to build entirely intelligent and self-contained systems. Rather, the intelligence, modeling capabilities and expert knowledge of future AI systems is simulated, and filtered through an interface designed to communicate in new and more effective ways. For example, a system-warning for closing an unsaved document called “Untitled” and, say, erasing a hard disk, should “feel” emotionally distinct to the user. Regardless of the intention of the user, the system should realize one is a dramatically more meaningful event, and react accordingly.

A character and the environment can deliver this kind of message, imply additional information in a manner requiring less time and attention than traditional interface elements.

Three working examples of these ideas will be discussed in the remainder of this paper, as well as directions for continued work in this area.

System Specifics

Current experiments use the beta-released Microsoft Agent™ Active-X control for animation and character appearance. The Active-X control is a windowless animated Genie, which can be instructed to perform distinct animations through a C++, Visual Basic, VBScript or Javascript interface. Instructions include such commands as “GlanceLeft”, “MoveTo(x,y)”, “Approval”, and “Sad”. In this project, the genie is controlled through dynamically generated Javascript (via Perl 5.0 CGI and shell scripts, served by a Netscape Communications server on an HP9000/750). Further control is provided through local Visual Basic programs. The character runs on an Intel Pentium/155 under Windows NT 4.0

Email filtering

The first project created filters to react to incoming email-messages according to a given priority. A flat-text file was written, containing email addresses and associated “priority” ratings. High-volume lists, for example, were given low-priority, while close friends or the author’s professor received higher rankings. The list simulated the work of an intelligent user-modeling agent, which might be able to glean such rankings through observing user-actions over a period of time, or by communicating with other user-agents in the same workgroups or net-communities.

When email is received at the user’s Post-Office server, it is stored for delivery as usual. A Perl script scans for the highest-ranked incoming message in the queue, determines behaviours according to hard-coded reactions, and generates Javascript instructions. These instructions are sent via HTTP to the Active-X genie on the user’s machine, which reacts accordingly.

The reactions include such behaviours as simply glancing at the icon of the user’s email-reader for low priority messages, or pointing more vigorously for higher rated email. Top priority messages cause the character to wave its arms to attract attention, move across the desktop to a position next to the email-reader application, and stare at the icon.




Figure 1 Genie as Email Agent

At any time, clicking on the agent itself reveals a list of the waiting email it is reacting to.

This example places the character “in the world” of the computer screen, allowing it to interact with the same icons, applications and desktop tools as the user.

Unfortunately, difficulties may arise as the agent gets more sophisticated responses to a larger number of stimuli – gaze-direction is a fairly arbitrary reaction to the specific domain of email handling. However, as the perceived importance of the message increases, the ambiguity of the character’s response decreases. Thus we can convey useful information, continuously, without necessarily interrupting the user’s original tasks.

Maitre-D

Maitre-D is a web-site based browsing agent being developed by the author. Through off-line textual analysis of the pages associated with a site, Maitre-D creates relations between groups of pages. As a visitor browses the site, Maitre-D notices which page the user is on and produces rated recommendations of related URLs, regardless of whether they are linked from the current page or not.

An animated character can work in concert with Maitre-D, as the interface to the intelligence. The Active-X character is spawned by a second browser-window, which also contains embedded Javascript functions to track events in the first, user-controlled window. Whenever the user moves to a new page, the new URL is reported back to Maitre-D, which runs server-side. Browser history, keywords which describe the new page, and fields simulating the input of a personal user-modeling agent are then used to produce a group of recommended “next pages” for the user to browse. The strength of Maitre-D’s highest recommendation is used to create behaviours in the client-side character.

For example, low-rated recommendations produce no visible reaction in the character. If the user wishes to view the agent’s suggested pages, they can click on the character directly for a plain-text, ordered listing. Better matches can cause the character to look around, nod approval, or wave arms to attract attention. Again, clicking produces a list of recommendations, and selecting from the list sends the browser to the new page.

Staging and Setting

The look and feel of the desktop environment provides a rich domain for subtly crafting informative emotional states. By treating the computer as a “setting”, we can leverage a huge body of experience in stagecraft, lighting-design and cinema to communicate with




Figure 2: Genie and Maitre-D

the user. The goal is to inspire a general feeling indicative of the state of the system, rather than conveying specific information.

The character can also react to the environment, allowing behavioural modes that don’t directly address the user, yet still serve a purpose. A-life techniques can thus be used to create a reactive character that responds to the stimuli of the computer system. The agent’s goals and desires can express valuable information implicitly, simply through its actions and behaviours.

This area is still being examined, with one attempt linking window/frame border sizes to web-download progress (or other duration-based gauges). A Visual Basic program polls an instance of a web-browser for progress information, and adjusts system-preferences accordingly.

A second program modifies desktop wallpaper according to system performance or network traffic, pushing up towards shades of red to show intense activity, and remaining in darker shades during normal conditions. Tinting the colour-scheme towards a sickly green could signify viral infection, fragmented disks or corrupted segments. Monitor keystoneing could signify low diskspace.

Again, the intention is not to replace discrete communications, or convey complete information. However, through continuous styles of interaction, general state can be implied,

and the user can take appropriate actions to determine specific problems and solutions.

Conclusions and Future Directions

The Microsoft Active-X control proved useful for rapid prototyping in the above experiments, but also revealed several shortcomings in the system's design. First and foremost is the focus on character-actions as discrete events: the API doesn't support the multiple simultaneous directives (eg: Glance Left AND look Sad) which are necessary for more complex behaviours. The ability to subtly vary actions – essentially, to convey information through the apparent emotional state of the character as well as through its discrete movements – is one of key benefits of an anthropomorphized interface. Parameters for controlling the execution of any given action, as well as influencing character-state over the course of many actions (persistent emotional state) might increase the versatility of the agent. Work in this area might allow more efficient communication while at the same time preventing prolonged interactions from growing extremely repetitive.

The Genie also has a lot of built in “idle” states. These animations, performed when no other instructions have been queued, apparently exist to preserve the illusion of life in the character. However, the movements and behaviours have no meaning whatsoever, despite the tendencies of users to attribute importance to them (eg “Why is he bored? Am I boring him?”). If the system is to be useful for conveying information through its actions, then those actions should be consistently meaningful.

On a related note, the animation-agent is, in general, too obtrusive. An agent should be able to carry out its activities with many levels of subtlety, attracting the user for explicit communication or going about its business in the periphery of attention, as needed. Adding transparency, or allowing software-control of the character's size and z-location (ie: as opposed to

being “always-on-top”) might help with this problem. It would also solve the problem of idle-states, allowing the agent to simply become less obtrusive (yet still present) when not otherwise needed.

Internally meaningful reactions (ie: those not specifically addressing the user, but part of the character's environment and role) should be explored further. If the animated agent is to be understood as an avatar of the computer's state, it should be continuously “displaying” information through general activity, rather than only animating to inform the user of a discrete event. For example, creatures might map network security levels to nervousness and agitation, or processor load as animation-speed or visual jitter.

The project will continue to explore these areas using new and more flexible animation system, as well generating more meaningful mappings, in the months to come.

Bibliography

- Bates, Joe “The Role of Emotion in Believable Agents”, *Communications of the ACM, Special Issue on Agents*, July 1994
- Blumberg, Bruce, and Galyean, Tinsley “Multi-level Direction of Autonomous Creatures” *Proceedings of ACM SIGGRAPH 1995*
- Galyean, Tinsley “Narrative Guidance of Interactivity” MIT PhD Thesis, 1995
- Keller, E.F. “Reflections on Gender and Science” Yale University Press (New Haven, CT 1985)
- Laurel, Brenda “Computers as Theatre” Addison-Wesley (New York, NY 1991)
- Perlin, Ken and Goldberg, Athomas “Improv: A System for Scripting Interactive Actors in Virtual Worlds”, *Proceedings of ACM SIGGRAPH 1996*
- Microsoft Agent homepage
<http://www.microsoft.com/intdev/agent/>