WORKING TOWARDS AN EMOTION-BASED AGENT ARCHITECTURE

Christopher S. Tingley

MEng Computer Science and Cybernetics, siu01cst@rdg.ac.uk

ABSTRACT

It is known that emotions provide useful functions within animals and humans. It has been proposed that agent systems can also benefit from biologically inspired emotions, but to do this an adequate emotional-based architecture must be developed. This project has developed such an architecture using a test platform consisting of three physical robots and a number of virtual agents. From observing the emergent behaviors of the agents equipped with the developed emotion-based architecture, several interesting and beneficial traits were identified. 1) A reduced complexity in behavioral programming and action selection due to the abstraction of inputs in which turn leads to simpler creation of emotional responses. 2) Reduced bandwidth between the agents and controllers and 3) more realistic behaviors.

1. AN INTRODUCTION TO EMOTIONS

Emotions are known to be linked into how humans make decisions and help guide the way the species lives in both work and play. They are also important in controlling and regulating ways that humans communicate and how humans express themselves verbally and non-verbally. It is important to note that there are several different definitions for emotions as they are often used in many different contexts including neuro-endocrine, physiological, behavorial, cognitive, communicative and social. When referring to biological systems, this paper assumes that "an emotion" may comprise of any one or more of these elements.

A conclusion may be drawn that emotions are simply a result of neuro-mechanisms linked to actions and consequences, which might explain the physical processes of emotions and why humans have them. Rolls [3] defines emotions as "the states elicited by reward and punishers, including changes in those rewards and punishments", by which "a reward is anything for which an animal will work, i.e. a Positive Reinforcer" and "a punishment is anything that an animal will work to escape or avoid, i.e. a Negative Reinforcer". This description providex a straight-forward explanation of the role of emotions within animals. It is considered that emotions are a result of the actions humans take and are a sum of the rewards and punishments obtained along the way - much like the actions and consequences that result from chemical processes in the brain.

2. WHY USE EMOTION-BASED REASONING IN AUTONOMOUS AGENTS?

Any agent (humans included) operating in a dynamic and unpredictable environment suffers from the problem of selecting between conflicting short and long term goals. In the biological world, emotions are one of many mechanisms employed to deal with such situations and their associated problems. To attack similar problems in autonomous agents, biological emotions can be looked towards for inspiration. From this point on, the word emotion (in the context of autonomous agents) is used purely as a analogy providing a simple description of biologically inspired control mechanisms for autonomous agents.

Sloman [6] has put forward a straightforward "Triple Layer Perspective" (hybrid) architecture produced from a low-level reactive layer, a middle deliberative layer and a higher "Meta-management" layer. This architecture allows the classical reactive and deliberative layers to operate as ever, but with the addition of a meta-management layer designed to allow other operations such as emotions to be placed within the system. By adding a meta-management level, Solman hypothesises that the following abilities can be realised within autonomous agents:

- The ability to monitor, categorise, evaluate and (to some extent) control other internal processes.
- Control attention and thought processes, e.g. a deliberative layer can be made up from a series of "what ifs?" generated by the inputs into the system. Some of these "what ifs?" may not be relevant to the current task and so can be ignored, thus directing attention to the most important considerations.

- Have varying levels of sophistication for example, having to depend on social learning to complete a complex or unfamiliar task.
- Self evaluation - e.g. How do I feel?
- Ability to characterise mental states - e.g. What would I call how I feel (happy, frustrated, etc.)?

In essence, adding a meta-management layer is very similar to the hybrid architecture. Assuming that this is possible, Scheutz [5] gives 12 roles of emotion within emotional agents: *Action selection* - what to do next based on current emotional state(s) or mood. *Adaption* - short or long-term changes in behaviour due to emotional state(s). *Social regulation* - communicating or exchanging information with other agents via emotional expressions. *Sensory integration* - emotional filtering of data or blocking of specific sensory integration. *Alarm mechanisms* - fast, reflex-like reactions in critical situations that interrupt other processes. *Motivation* - creating motives as part of an emotional coping mechanism. *Goal management* - creation of new goals or reprioritisation of existing ones. *Learning* - emotional evaluations as Q-values in reinforcement learning. *Attentional focus* - selection of data to be processed based on emotional evaluation *Strategic processing* - selection of different search strategies based on overall emotional state. *Self model* - emotions as representations of "what a situation is like for the agent", e.g. self evaluation and or categorisation.

3. ARCHITECTURE DESIGN

3.1. Traditional architectures

Following on from the introduction above, agent architectures can be defined by a collection of components or modules (layers) with clearly defined interfaces and connections, structured in a hierarchial manner. Components can be descried on a number of levels:

Low-level components such as actuators, sensors or circuit boards (on a physical agent i.e. a robot) and the connections between them. *Mid-level* components may refer to the programmatic behaviours of the agent through observation of the agents' interaction with the world via its actuators and sensors (whether they are simulated or real-life).*High-level* encompasses the emergent behaviors of multi-agent systems or meta-management layers of single agent systems.

Agent architectures must also have some degree of *modularity* allowing components to be easily changed, removed, appended or swapped without having to reprogram or redesigning the remaining or underlying architecture.

3.2. An emotional architecture

Figure 1 depicts a diagram of the authors' emotional-based architecture designed to facilitate the traits of emotions within humans. It has also been designed to work towards some of the roles of emotions proposed by Scheutz, whilst adhering to issues raised in [1] and [4].

The architecture consists of three layers, each varying in their level of complexity. The lowest layer allows for reactive (reflex like) behaviors, the middle layer controls goal management whilst the highest (and most abstract) layer governs the overall emotional state of the agent. Lower levels are able to subsume emotions generated by higher levels and connection to a multi-agent system controller is possible by communicating non-observable emotions.

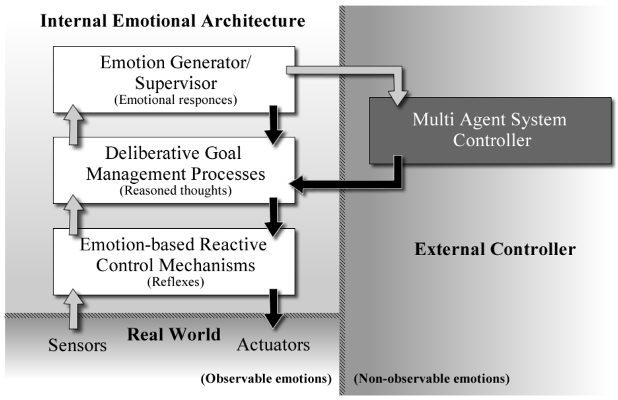

Figure 1 . The developed emotion-based architecture

This architecture encompasses all of the important features of the hypothesised uses of emotions within autonomous agents whilst fitting them into a simple, three layered architecture. The multi-layered approach means that different levels (complexities) of emotional responses can be computed at different levels of autonomy within the system. It is important to note that the authors' implementation did not include any deliberative goal management as the main focus was on development of low level reactive control mechanisms (more commonly know as action selection).

4. THE IMPLEMENTATION PLATFORM

The research platform used to visualise the effectiveness of the developed emotional-based architecture are three Miabots developed by Merlin Systems [2]. The Miabots are equipped with an ultrasonic ring comprising of eight transmitter/reciever pairs and one ultrasonic beacon developed as part of the localisation system. Each of the Miabots has a bluetooth link to the server PC, in turn facilitates the relay of messages to and from the robot and the controller.

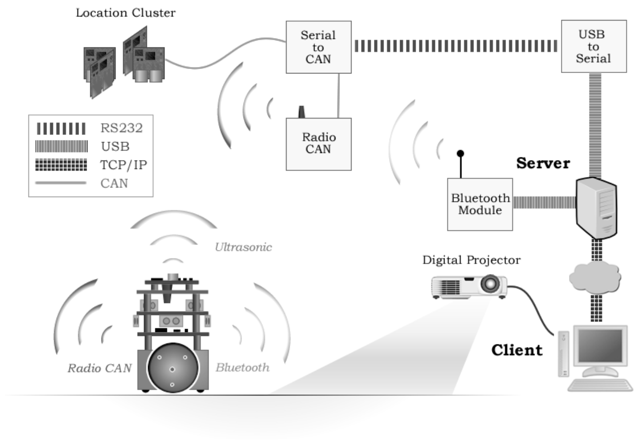

Figure 2 . The implementation platform

4.1. Localisation system

A novel system to uniquely identify and locate each of the Miabots was adapted from a system already in use by the Madlab. The original system uses three receivers boards (each consisting of two ultrasonic receivers) to determine the angle and elevation of a transmitted ultrasonic pulse. This system was modified to include another receiver board so as to locate the source of the pulse in three axes (i.e. three dimensions).

The new system employs the new configuration of four receiver boards (a.k.a. locator boards), a radio CAN bridge and a CAN to RS232 (serial) converter. This complete system is known as the location cluster and is hung from the ceiling directly above the environment. Each robot is equipped with one beacon programmed with a unique identification number (ID) and has the ability to send and receive CAN messages via a radio link to the location cluster on the ceiling.

4.2. The server-client configuration

As the hardware for the project was being shared, it was necessary to to develop a method of distributing control of the Miabots to more than one machine. For this reason it was decided that one PC should be dedicated to the control and collation of the Miabot data. This machine is equipped with a bluetooth link to communicating with the Miabots and a serial link to communicate with the location cluster. This machine is known as the Miabot Server. Software was written (in C#) that allows simultaneous connection to one or more of the Miabots, providing directional control and collation of the ultrasonic ranges from each of the Miabots.

As the range finders and the location system both operate in the ultrasonic spectrum, they were found to interfere with each other producing incorrect results. These undesired results were avoided by giving the server the ability to synchronise the ranging and location pings, ensuring operations did not overlap. As the ranging data is less important than the location, the ranging for each robot is done simultaneously and the location pulses occur independently, one Miabot at a time.

The server software runs continuously gathering data and has the ability to relay this information to any other PC on the computer network via a TCP/IP connection. The client software can connect to the server and periodically requests ranging and location data. Processing then occurs on this data and each connected Miabot has its speed and direction recomputed appropriately. The updated wheel speeds are then sent to the server which in turn relays this information over the bluetooth link to the Miabot.

This arrangement of a server-client allows control of the Miabots to be completely off-board and to operate at a sufficiently high refresh rate for the control system to be event-driven in real-time. The processing overheads on the client are also lower as the data pre-processing occurs on the server machine.

4.3. Virtual Bots

To experiment with the interaction between more than the three physical Miabots owned by the department, a computational model (using forward kinematics) of the miabot was created which uses the same architecture and rules as the physical Miabots. The virtual and real robots are able to "see" each other by computing fake ultrasonic readings within the client software, this means every robot/agent in the environment is able to interact with each other.

5. RESULTS

The proposed architecture has been implemented into scalable, completely object-orrientated multi-agent system controller applied to three physical Miabots and several virtual agents. The resultant system can be seen to exhibit three important features, all of which can be seen in some form in the list of important roles proposed by Scheutz (described in 2). During development of the system, many areas of further research have become apparent as the current development platform sacrifices obvious optimisations in favour of extensibility.

5.1. Consolidation of inputs

This refers to the ability to easily consolidate a high number of external inputs into fewer "abstracted emotional inputs". The developed architecture allows this to be done from a more "natural", pragmatic and ultimately real-life approach than current architectures, allowing much simpler control mechanisms to be developed from a more natural viewpoint, i.e. when curious, move faster or when scared, move away etc.

5.2. Lower communication bandwidth

After emotional abstraction has taken place, fewer inputs mean that the communication channels between individual agents and the multi agent controller benefit from requiring less bandwidth. The current implementation of the platform has the physical robot controlled completely off-board. Future developments could reduce bandwidth requirements even further by implementing low-level parts of the emotional architecture on-board the robots.

5.3. Realistic responses

The emergent emotional behaviors exhibit more natural and realistic responses, making the agents seem more life-like and "likable". This can be very useful for any agent system in the general consumer market and may be even more desirable for virtual agents, where it is known to be difficult for people to relate or interact with virtual agents.

5.4. The mixed reality system

Figure 3 shows the accuracy of the location system to be under a centimeter for each location computation (note that over ∼25 readings there were no obvious anomalies). This accurate location ensures that the mixed reality system is able to provide a very accurate overlay of the real Miabots within the virtual environment.

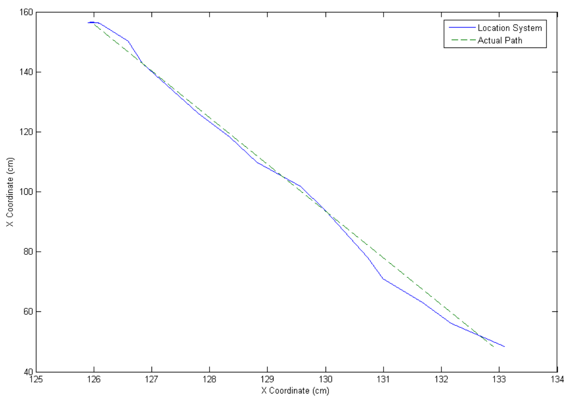

Figure 3 . Actual stright-line path of one of the Miabots plotted against the output of the locator system, illustrating the deviation between the measured and actual paths.

6. CONCLUSION

This project has identified several benefits for a emotional-based agent architecture and has implemented the proposed architecture in a robust, physical and virtual robotic environment. An accurate three dimensional localisation system has been constructed and integrated into a powerful server providing a scalable interface for communicating with the departments' Miabots. The server software has been written to allow control of the system through any network connection within the University, allowing distributed and remote control of the Miabots.

Additionally, a mixed reality system has been developed as part of the client control software allowing physical and virtual robots to coexist in the same environment. This allows multiple observers to view the mixed reality system in real-time, from any angle and without obfuscation of the virtual robots due to physical objects within the environment. With respect to the developed architecture, it can been seen to be effective by providing a simple platform on which to build complex behaviors.

Acknowledgments: Many thanks go to Dr. Will Browne for his supervision throughout the project and both Dr. Ben Hutt and Iain Goodhew for their invaluable help and discussions during the project. Fellow undergraduate, Chris Forster should also be mentioned for his collaboration in several aspects of the project, most predominantly the location system.

7. REFERENCES

[1] Canamero, D. Issues in the design of emotional agents. In Canamero, D., editor, Emotional and Intelligent: The Tangled Knot of Cognition: Papers from the 1998 Fall Symposium, pages 23–27, 1998.

[2] Merlin, Systems Corp. Corporate website, March 2006. http://www.merlinsystemscorp.com.

[3] Rolls, E. T. The brain and emotion. Oxford University Press, 1999.

[4] Scherer, K. R. Criteria for emotion-antecedent appraisal: A review. In Hamilton, V., Bower, G., and Frijda, N., editors, Cognitive perspective on Emotion and Motivation, pages 89–126. Kluwer Academic, 1988.

[5] Scheutz, M. Useful roles of emotions in artificial agents: A case study from artificial life. In AAAI, pages 42–48, 2004.

[6] Sloman, A. When will real robots be as clever as in the movies? In ASE Conference, Birmingham, UK, January 2003. Department of Applied Mathematics and Theoretical Physics, University of Birmingham.