HOW TO GIVE AN AGENT AN ATTITUDE

Penny Ray
Department of Mathematical and Computing,
University of Southern Qld,
Toowoomba, Qld., 4350,
AUSTRALIA.
Email: ray@usq.edu.au

Dickson Lukose
Brightware Inc.,
90 Park Avenue, Suite 1600,
New York, NY, 10016, USA.
Email:
dickson.lukose@brightware.com

Mark Toleman
Department of Mathematical and Computing,
University of Southern Qld,
Toowoomba, Qld., 4350,
AUSTRALIA.
Email: toleman@usq.edu.au

KEYWORDS

Agent, Attitude, Emotion, Reasoning and Decision Making

ABSTRACT

The need to improve our knowledge about the emotional influences in humans and animals and explore affective theories in order to help us understand human-computer interactions with respect to usability and acceptance, has in recent years seen the advent of affective computing. In this paper, we propose a method for generating affective reasoning and decision making based on psychological theories of attitude. Our domain-independent, simulation environment, GOMASE, allows its agents to assess their beliefs and satisfy their goals based on their *attitudes* towards other agents and objects in their environment. This mechanism allows the agents to deal with situations where problem solutions are innumerable or time doesn't permit for finding the optimal answer. The motivation behind this research is to develop a computational model that will allow us to give artificial agents the abilities to evaluate their thoughts and actions. These assessments will form the building blocks on which to develop intelligent emotion generation.

INTRODUCTION

Recently there has been much research into the design and development of artificial agents that implement affective models. The motivation behind this growth is threefold. Firstly, it aids to improve our knowledge about affects and their influence on behaviour in all animals (including human beings). Secondly, the exploration of affective theories and the role that emotions play in human heuristic problem solving could increase machine performance by incorporating these models into computer systems, and thirdly, we could improve human-computer interaction with respect to the acceptance and usability of a system for the user.

There is much scientific evidence that emotions play an important part in cognitive processes and that they impact significantly on rational thinking, perception, learning and other psychological functions that mold human behaviour (Picard 1998).

Several software systems have been developed in recent years that incorporate models of emotion into their architectures. These include Elliott's Affective Reasoner (Elliot 1992) and Reilly's Believable and Emotional Agents in (Reilly 1996). Both of these systems implement the Ortony, Clore and Collins (OCC) cognitive appraisal theory (Ortony, Clore et al. 1988). Using this model allows the agents in their system to synthesize emotions and to be behaviourally influenced by them.

In our research, we have taken a general overview of the OCC model and other psychological models of emotions such as Frijda's model of emotion as described in (Wright 1997) and Koestler's in (Koestler 1967). These authors agree (as do psychologists Freud, Thorndike and Hull (Koestler 1967)) that there is a general production of emotions related to pleasant and unpleasant experiences and these experiences influence future behaviour.

Our work focuses on the appraisal mechanisms by which the OCC model defines the categories of emotion. For example, the category known as *conformation* is the set of emotions that are generated when an individual is *displeased* or *pleased* about an event appraised as *desirable* or *undesirable*. The emotion of *disappointment* is generated when an individual is *displeased* about a *desirable* event failing.

In this paper we describe a mechanism for appraising an event and using this to drive the behaviour of an artificial agent. This behaviour is reinforced in the agent by performing tasks and rating them as either pleasant or unpleasant. Our aim is to implement a model by Fishbein and Ajzen (Fishbein and Ajzen 1975) used to predict human behaviour and apply it to intelligent agents designed to mimic human behaviour.

This paper is formatted in the following manner. Initially we will define attitude and it's importance in being used as a foundation for emotion generation. Next we will examine how attitudes lead to behaviour and how behaviours influence attitudes. Finally, we will give an outline of our initial experimentation in comparing our computational models of attitude with that of a human subject in a reasoning and decision-making scenario, a discussion of our results and thoughts on future works.

ATTITUDE

While there is no clear definition of attitude, there is popular consensus among social psychologists that the term refers to the general enduring disposition to feel positively or negatively towards an object, person or issue (Petty and Cacioppo 1996). By using a relevant collection of a person's attitudes, it is possible to closely predict the behaviour of that person where the same set of attitudes is applied (Fishbein and Ajzen 1975).

Ajzen and Fishbein as noted by Petty and Cacioppo in (Petty and Cacioppo 1996) concluded that behaviour can be viewed as consisting of four key elements. They are 1) the action being performed; 2) the target or targets that are the object of the action; 3) the context of the action e.g. where it is being performed; and, 4) the temporal alignment of the action e.g. the time of day or month.

In the experiment outlined in this paper, we have constrained these four key elements and isolated just one on which to perform our investigations. We have chosen to look at objects and how attitudes are formed towards them. Therefore, the elements of action, context and time have been set to static values.

According to Rosenberg, as cited in (Fishbein and Ajzen 1975), the more a given object is instrumental in the success or failure of an individual's goal the more the individual's attitude towards that object is favourable or unfavourable.

We can express this attitude towards the object, A as:

$$A = \sum_{i=1}^{n} W_i V_i \qquad 1$$

where W_i is the weighting of blame or praise associated with the object being instrumental in the failure or success of the goal i, V_i is the degree of satisfaction or dissatisfaction obtained from performing the goal and n is the number of times that the goal has been performed using the object.

TURNING ATTITUDE INTO ACTION

In many situations, it is impractical to analyse all possible courses of action and make a decision based on the measured plausibility of each thus, it is an ideal course of action to consider the second theory. Picard (Picard 1998) suggests a model of decision making using intuition as a guide to reasoning where measures of *good* or *bad* are used to assess choice options. While choices deemed to be bad or negative overall are not dismissed as alternatives, choices that are good or positive overall are explored further. By our definition, the weightings on the choices mentioned above are the attitudes towards the options.

In our agent, there exists a mechanism by which it can calculate it's overall attitude towards an option or event by assessing its attitude towards individual objects that are involved in the execution of the event. The overall attitude value for an option is used to prioritize it.

Given the set of options, the agent will assess its attitudes towards each option. The option with the highest calculated attitude value is used to represent the option towards which the agent *feels* most *favourable*. This ordered list of options becomes the agent's intentions towards satisfying the activated goal that began the attitude evaluation exercise. These intentions are scheduled for performance. The outward behaviour of the agent is an intention being acted upon. Not all intentions become agent behaviours, and when one intention satisfies the active goal, any other intentions that would also satisfy that goal are no longer necessary and may be discarded (Ray and Lukose 1999).

TURNING ACTION INTO ATTITUDE

As previously mentioned, the attitude of an object is related to the way in which an individual perceives that object as being involved in the success or failure of a goal. However, where the individual and the object in question have a past history of successful and unsuccessful usage, attitudes towards that object will already exist. This means that an unsuccessful attempt to use the object will not automatically cause a disliking of the object, especially if the object is already held in high regard (Davidson 1995).

Research has identified that it is not only the success or failure associated with an object that determines the attitude towards it. The most significant factor that contributes to an attitude value is the amount of information that the individual already knows about the object (Davidson 1995).

Each new piece of information that the individual receives about the object will have an effect on the attitude held about it. There are several information processing models that have been developed to explain what is referred to as *information integration* (Davidson 1995).

By applying an information integration model to our agent, we can update the attitude value that is held about each object by assessing a new attitude based on equation 1 and combining it with the previously stored attitude.

SCENARIO AND TESTING

In our test bed environment GOMASE (Goal-Orientated, Multi-Agent Simulation Environment), agents are driven by a goal hierarchy (Ray and Lukose 1999). In this we assume that a goal can be either abstract or primitive. An abstract goal can be decomposed into sub-goals (of which some will be abstract goals, while others may be primitive goals). Primitive or atomic goals correspond to an activity (or action) that needs to be carried out to achieve the goal. When a goal becomes the focus of an agent's belief and the agent wants to satisfy that goal, each sub-goal of that goal becomes active.

An agent may have any number of atomic goals for which it can perform tasks in order to satisfy the goal. In many cases where the agent has been given numerous task sets, not all of these need to be executed to satisfy the goal (Ray and Lukose 1999). Often when a subset of these tasks has been successfully completed the goal will be satisfied and the remaining active atomic goals and tasks can be deactivated.

For our attitude experiment we have chosen to give the agent one simple goal. The agent is preprogrammed with the knowledge that there are 10 different ways in which to satisfy this goal. The agent is then asked to satisfy this goal 100 times. Data is gathered about the choices that the agent makes.

This same goal was also been given to a human subject. At this point it may be simpler to insert the narrative given to the human subject at the outset of the experiment in order to explain it.

You are in a plane crash and find yourself stranded in the Jungle. It is 9am on a clear sunny morning. You have to find your way back to civilization. In the back of the plane you find 10 survival packs. Each survival pack contains a number of items to help you with your journey. Each pack weighs the same and is comfortable to carry. You have been given a list of the packs and their contents. Each item in the pack can be responsible for your success or failure. If just one item fails when you try to use it then your mission fails.

When asked, select a pack by entering the number. Each item in the pack will be examined to see if you can use it.

To complete a game you have to keep trying to escape the Jungle by selecting a mission survival pack.

Each survival pack represents one of ten ways to attempt a task that will satisfy the goal of escaping from the Jungle (alive). The list given to the subject inventories the ten packs and their contents. The packs randomly contain up to 5 items named a, b, c, d and e. For example, survival pack 1 may contain items a, c and e and survival pack 5 may contain items a and b. Each item's alphabetical designation refers to exactly the same item. For example, item a in pack 1 is the same as item a in pack 5. We have chosen not to associate the items with real world objects in an attempt to eliminate any preconceived attitudes that could influence the human decision making process.

The agents (both artificial and human) have to try and escape from the jungle 100 times. Each of these times we have called a *game*. To complete a game the agents must find a pack that helps them to escape successfully.

Before the start of the experiment, each item in the pack is given a static probability rating of failure. When either agent chooses a pack, the success or failure of that pack is calculated by assessing a success or failure on each item in the pack. If one item in the pack produces a failure rating, the whole task is rated as a failure. The agents are then made aware of which items are responsible.

As the experiment is run, we gather data about the number of packs chosen in a single game before a success is found, the order in which the packs are chosen and how the agents feel about the items. The aim is to determine if the attitude model implemented in the agent performs the task in a similar manner to a human.

RESULTS

Our initial results have been quite encouraging. Figure 1 is a box plot showing a 95% confidence interval for the mean number of survival packs chosen in a single game before a successful one was found. As can be seen there is basically no difference between the GOMASE agent and the human subject.

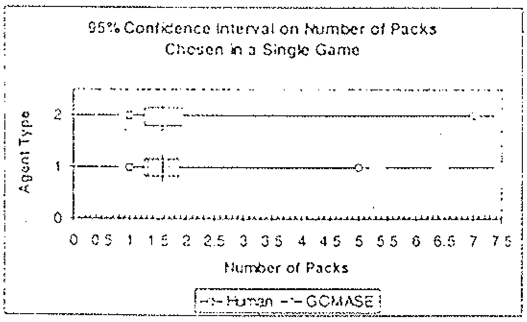

Figure 1. 95% Confidence Intervals on Mean Number of Pack Choices per Game

Further investigation showed different results for the choice patterns made by the two agents. Figures 2 and 3 display the sequence of choices that the agents made in selecting a survival pack. The dark points show each pack chosen and the white dots represent when the chosen pack was successful. Although the packs chosen are quite different, similar patterns can be seen in each graph, as explained below.

Firstly, plateaus are apparent where the agents selected a pack and the pack was successful, so it was chosen the next time. For the human this is mostly evident for packs 7, 2 and 1. For the GOMASE agent they are packs 9, 4 and 3.

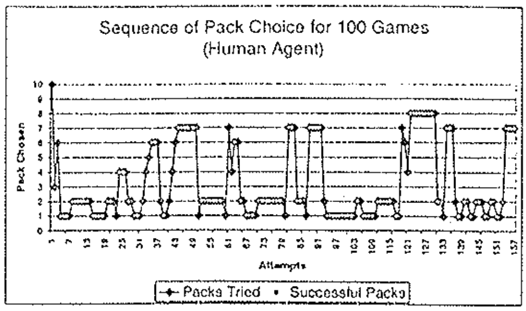

Figure 2. Human Agent Pack Choice Sequence

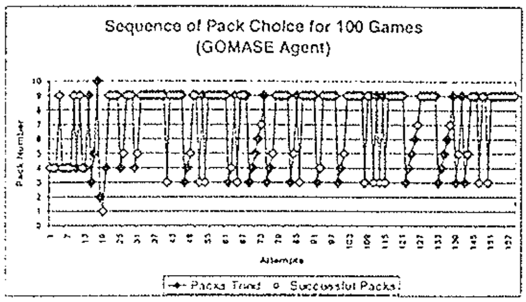

Figure 3. GOMASE Agent Pack Choice Sequence

For each of the agents, these are the packs that were successful most of the time and therefore the agents returned to use them again.

Secondly, there are set sequences of pack choice. Should a chosen pack fail, apparently each agent had a *backup or alternative* pack from which to select. For the human agent this sequence consisted of the packs identified in the graph plateaus. One emerging pattern was to select pack 1 should pack 2 fail and vice versa. Pack 7 was also chosen several times when packs 1 or 2 failed or both and a sequence of packs 2 then 1 were chosen should pack 7 fail.

In the case of the GOMASE agent the patterns were more distinct. This should be expected as the GOMASE agent is a logical machine basing it's reasoning on discrete values. This agent's *favourite* sequence was 9, 3, 4, 5, 6, 7.

Thirdly, in the beginning, each of the agents tried several tasks before establishing a pattern. These initial tries also had an effect on future choices. In the case of the human agent at the beginning of the graph it can be seen that several tasks are tried. The failure of these tasks seems to have an impact on other pack choices by the human subject. Packs 10 and 3 are never chosen again and pack 6 was chosen again only after a succession of other failures.

The GOMASE agent's graph was slightly different. In the beginning the agent was fortunate enough to select packs 4 and 9 that had a high success rate, so it continued using them. Once they both failed in succession around the 15th game, it can be seen that the agent tried a number of other packs before finding one that worked. During this time it tried three packs that it never chose again. These were packs 2, 10 and 1 (even though 1 was successful).

Finally, in each case there were packs that were not chosen: The human agent did not select pack 9 and the GOMASE agent didn't choose pack 8.

DISCUSSION

While computers can outperform humans in many logical problem-solving situations, there are often times when optimal solutions to problems cannot be enumerated with the available resources or within a finite time range. Humans have the remarkable ability to respond in situations with limited knowledge, limited memory and comparatively slower processing speeds (Picard 1998).

These initial experiments show promising results for using attitude prediction theories to produce humanistic reasoning patterns in artificial agents. In this experiment neither agent knew the success rates of the survival packs. Only by using the packs could the agents form an opinion about them. In both agents, the packs that succeeded the most had a higher attitude formed towards them. This caused these *favourite* packs to be chosen time after time.

In the human agent, a popular pack did not lose favour with the subject after several failures. This too can be seen in the artificial agent. The information integration formulae used to update the GOMASE agent's attitude towards the packs increases the attitude value for a pack each time it succeeds. After many successful implementations, the GOMASE agent can have a high attitude value for a pack. This attitude value may be much higher than any other pack. When such a pack fails, the GOMASE agent decreases its attitude value towards the pack. However, because the attitude value was so high, one failure of a favourite pack, will not cause it to become unpopular. Depending on its attitude value relative to other packs, the pack may remain the favourite.

The exercise in this experiment was to find a pack that helped the agent escape from the jungle. The agents did not have to identify the optimal pack. For commentary sake, the pack with the least probability of failure was pack 7. The human subject identified this pack although it was ranked as third most used, as can be seen in Figure 4. The GOMASE agent chose pack 7 very little. Each agent identified a small number of packs that could help it achieve its task in the least number of choices.

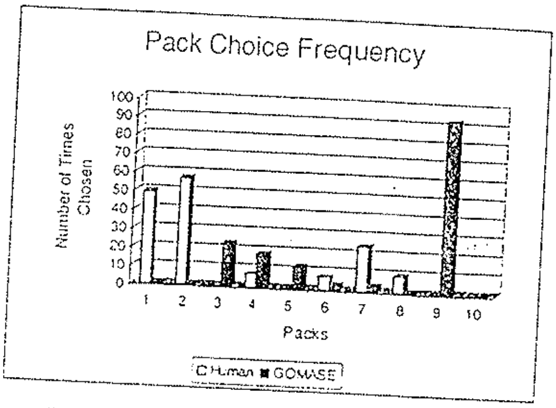

Figure 4 Pack Choice Frequency for Agents

In many real life situations, human's tend to select a solution that works within specified parameters and if it is successful, there is often no need to look for another solution unless the previously selected one fails. For example, a person may drive to work the same way everyday. The route may not be the optimal route (shortest, better road, less traffic etc.) but for one reason or another, this route has worked successfully for the person day after day. Imagine that one day the route is closed due to road works. The person will find another route that works for them. When the road works are finished they may return to the old route, or if they found the new route to be better[1], they may stick with it.

For this experiment initially it is not possible for either agent to access the success probabilities without using the packs. Only through past attempts can the agent form an attitude towards the packs. The more games the agent plays the more information they have about the packs success probabilities. This is an example of a

[1] The term *better*, as used here, refers to the person's individual preference. In the case of a *route to work* this might include things such as the scenery being more pleasant.

problem where it is not possible to calculate the optimal solution before applying one. In order to calculate the success probabilities of all the packs, by using them first, an agent would have to choose each pack at least 100 times to get an accurate picture of one pack's success rate (this being calculated as a failure probability out of 100). Multiply this by each pack and the attempts needed to determine the best survival pack far outweigh the 10 attempts (approximately) in the GOMASE and human agents for determining a short sequence of packs that work.

CONCLUSIONS AND FURTHER WORK

This paper has introduced the concept of modelling attitude in artificial agents in order to produce humanistic reasoning capabilities. Two domains in which this work is applicable are social simulation and interactive entertainment. Both of these areas require more accurate models of human behaviour at the micro level (Carley 1996) (Stern 1999) whether it be for simulating workers in an enterprise to monitor behavioural change versus environmental change or producing believable and unpredictable artificial enemy and companion characters for computer games.

The experiments in this paper are preliminary and are initial attempts to compare human reasoning and attitude with the models implemented in the GOMASE software. Further experimentation with human subjects in differing scenarios will assist us in improving our artificial agents. This will allow us to develop and test complex social scenarios with our agents in virtual environments with human interfaces and interaction for evaluation.

To these ends, our continuing research and the development of the GOMASE environment will endevour to expand our understanding of attitude theory and it's affect on affective reasoning in humans and the translation of these models into our affective agents.

REFERENCES

Carley, K. (1996). "A comparison of artificial and human organizations." *Journal of Economic Behavior and Organization* 896: 1 - 17.

Davidson, A. R. (1995). From Attitudes to Actions to Attitude Change: The Effects of Amount and Accuracy of Information. *Attitude Strength*. R. E. Petty and J. A. Krosnick. Mahwah, New Jersey, Lawrence Erlbaum Associates.

Elliot, C. (1992). *The Affective Reasoner: A process model of emotions in a multi-agent system.* Evanston, Illinois, Northwestern University.

Fishbein, M. and I. Ajzen (1975). *Belief, Attitude, Intention and Behaviour.* London, Addison-Wesley.

Koestler, A. (1967). *The Ghost in the Machine.* London, Penguin Books Ltd.

Ortony, A., G. L. Clore, et al. (1988). *The Cognitive Structure of Emotions.* Cambridge, UK., Cambridge University Press.

Petty, R. E. and J. T. Cacioppo (1996). *Attitudes and Persuasion: Classic and Contemporary Approaches* Boulder, Colorado, Westview Press.

Picard, R. (1998). *Affective Computing* London, The MIT Press

Ray, P and D Lukose (1999). *To Believe, to Desire to Intend Motivating Autonomous Intelligent Agents with a Mind of their Own.* Managing Enterprises 99, Newcastle

Reilly, W S N (1996) *Believable Social and Emotional Agents* School of Computer Science Carnegie Mellon University, Pittsburgh

Stern, A (1999). *AI Beyond Computer Games* 1999 AAAI Spring Symposium, Artificial Intelligence and Computer Games, Menlo Park, AAAI Press.

Velasquez, J D (1998). *Modeling Emotion-Based Decision-Making* Proceedings of the 1998 AAAI Fall Symposium Emotional and Intelligent: The Tangled Knot of Cognition (Technical Report FS-98-03), Orlando, FL, AAAI Press

Wright, I P (1997). *Emotional Agents* School of Computer Science, Birmingham, University of Birmingham: 267