Agents with Attitude

Penny Ray
Dept. Maths and Computer Science
University of Southern Qld.
Toowoomba 4351 AUSTRALIA
617 4631 1511
ray@usq.edu.au

Dickson Lukose
Brightware Inc.
90 Park Avenue, Suite 1600,
New York, NY, 10016, USA
1 215 981 5678
dickson.lukose@brightware.com

Mark Toleman
Dept. Maths and Computer Science
University of Southern Qld.
Toowoomba 4351 AUSTRALIA
617 4631 2533
markt@usq.edu.au

ABSTRACT

Incorporating emotional factors into logical machines is a relatively new and exiting area of research. The motivation behind much of this work has been to improve our knowledge about emotional affects in humans and animals, exploration of affective theories and to better understand human-computer interaction with respect to usability and acceptance. In this paper, we propose a method for affective reasoning and decision making. Our domain-independent, simulation environment, GOMASE, allows its agents to assess their beliefs and satisfy their goals based on their *feelings* about themselves, 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. To examine our agents, we have presented a solution to a challenging problem-solving scenario posed by Picard [6] endeavoring to examine the affective reasoning capabilities of our agents.

Keywords

Affective Computing, Emotions, Intelligent Agents, Attitude.

1. INTRODUCTION

Ajzen and Fishbein as noted by Petty and Cacioppo in [5] 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, for example, where it is being performed; and,
4. the temporal alignment of the action, for example, the time of day or month.

These four elements can be used to define a behavioural event and we will refer to them collectively as the *event space.*

One way of assessing the mechanisms for a behavioural event is to calculate the person's attitudes towards the elements in the event space. In general, attitudes serve as convenient summaries of our beliefs [5], for example, "I liked the movie" rather than explicitly stating all the reasons why. They also help others to predict our behaviour. For example, if you were to say, "I do not like horror movies" your friend may predict that you will not be going to see "Scream II".

In the psychological literature, there is no clear definition of attitude. Attitude has been equated with other terms such as attraction, attribution of disposition, opinions, morale and behavioural intentions [2]. There is however 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 [5].

One such model to incorporate these positive and negative *feelings* is the cognitive appraisal theory developed by Ortony, Clore and Collins (OCC) [4]. It examines the valenced reactions toward three elements: the consequences of an event; the actions of an agent; and, the aspects of an object. The model assesses the human reaction to one or more of these elements as being either, pleasing, displeasing, approving, disproving, liking or disliking. The OCC cognitive appraisal model, and subsets of it, has been successfully implemented in a number of emotional agent architectures such as Elliott's Affective Reasoner [1] and Reilly's Believable and Emotional Agents [9].

What we are proposing focuses on the reactions to the three elements mentioned above. Our approach differs from other implementations of the cognitive appraisal theory in that it takes a micro perspective of an event and evaluates it based on beliefs and the measurements of attitude that we apply to the elements in the event's *event space.*

To this end, this paper is organized in the following manner: in Section 2, we will describe how attitudes can be used to determine the agents feelings about decision choices; in Section 3, we discuss the relationships between event appraisal and emotion generation; in Section 4, we define the event space for a situation; and in Section 5, we address the appraisal of an event using attitude as a mechanism for biasing the intentions of an agent.

Following this, we look at our test bed environment GOMASE. In Section 6, we briefly overview the architecture of our system; in Section 7, we define an affective reasoning problem domain; in Section 8, we give examples of our agents' affective problem solving abilities; and in Section 9, we conclude with a summary of achievements outlined in this paper and the future directions and work that are yet to be done to realize the full potential of this research.

2. Attitudes as Biasing Mechanisms

According to Velasquez [10] most theories about human performance in reasoning and decision-making can be classified into two different positions. Firstly, humans reduce all decisions to a list of possible outcomes and use logic to conduct a feasibility study that will give a list of best possible choices. Secondly, new decision situations are compared to past experiences and we tend to make choices that follow consistent patterns of previous choices. In many situations, it is impractical to analyse all possible courses of action and make a decision based on the measured plausibility of each. It is an ideal course of action to take the second position.

In her affective decision making scenario, Picard [6] suggests a model of decision-making using intuition as a guide to reasoning. High-level, general decisions are made using measures of *bad* or *good* 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, these measurements of the choices, mentioned above, are the attitudes.

This method of affective reasoning will be examined further in Section 7.

3. Attitudes and Emotions

The method we have already described for determining overall attitude can be implemented in emotional agents to enhance the modelling of emotion generation theories.

The OCC model has its foundations in six distinct emotion categories (well being, fortunes-of-others, prospect-based, confirmation, attribution and attraction) and two hybrid categories (well being/attribution and attraction/attribution). Many of the categories describe emotions arising from the appraisal of an event (past, present or future) and whether it is pleasing or displeasing.

This also holds true for Frijda's model of emotion as described in [11]. In this theory, Frijda suggests that emotions are generated through a procedure of appraising a situation for signals of pleasure and displeasure.

In keeping with both of these theories, it is possible for us to appraise any situation as pleasing or displeasing using attitudes to evaluate the elements that exist within the event space for that situation.

4. Defining the Event Space

As mentioned in the introduction to this paper, rational human behaviour can be predicted best by assessing the attitudes toward the elements that would influence the behaviour. The set of these elements defines, what we have termed, the *event space*.

The event space E, can be defined as:

$$E = \{a, o, c, t\} \tag{1}$$

where a is the set of actions that relate to the event, o is the set of objects involved or effected by a, c is the context or conditions in which a is taking place or being performed and t is the temporal component of a.

Each of these elements may vary along a dimension of explicitness. At the most exact level of defining an event space, a person will intend to perform a specific action, with or towards a certain object, in a particular context or situation at an exact point in time. For example, Floyd may intend to *buy a bunch of flowers* (action) for *Georgina* (object) from *Barry's House of Posies* (context) at *12:30 p.m. on Tuesday* (time). At the other end of the spectrum, Floyd may intend to be affectionate without referral to any exact elements of the event space. For example, Floyd may intend to *buy a gift* for a *friend*.

For our purposes, how each component of the intended behaviour is categorized is inconsequential. The *bunch of flowers* could also be defined as an object of Floyd's behaviour. He could also just as easily intend to buy Georgina a *box of chocolates*. As there may be multiple elements of the same type in the event space (and even event spaces within event spaces) and each element is evaluated in the same manner, it is unimportant how we precisely define each component.

Our use of the event space is to set motivational levels in our agents with respect to performing actions that will satisfy a goal. Given a goal and a number of actions/tasks that will satisfy the goal, our agents calculate their preferred choice of action from values, or measures of attitude, that they hold about the elements in the corresponding event spaces.

Once evaluated, the actions can be prioritized in order of intention. This is then used by the agent in determining its behaviour.

5. Calculating Behaviour from Attitudes

There are a number of theories that have been developed for the purpose of describing attitude. For an analysis of these approaches see [2]. The theories describe how attitudes are formed[1] and how attitudes can predict behaviour. For our

[1] At this point, we would like to mention the importance of attitude formation. Rosenberg, sited in [Fishbein, 1975 #63] defined attitude as a "relatively stable effective response to an object". He also formulated a theory of attitude formation that stated that the attitude towards an object was proportional to that objects level of participation in obtaining or blocking the attainment of a goal. As this is not the primary focus of this paper we will not explore this theory here, except to say that it is important to belief revision and we will consider it in future research.

purposes, we would like to calculate attitude in order to generate a set of intentions for our agents from which to extrapolate their behaviour. To these ends, we are applying principles of human behaviour prediction, grounded in attitude theory, in order to mimic this behaviour in our agents.

In a BDI agent, intentions are the plans an agent has formalized in response to an activated goal. These intentions are scheduled for performance. The behaviour of the agent is determined by an intention being acted on. Not all intentions become agent behaviours, and when one intention/behaviour satisfies an agent goal, any other intention/behaviours that would also satisfy that goal are no longer necessary and may be discarded [8]. What we are interested in is developing a motivational mechanism for prioritizing these sets of intentions.

Our chosen model of intention prediction is Fishbein and Ajzen's *theory of reasoned action*. Their theory is twofold. Firstly, the attitude A, towards performing a behaviour B, can be determined by the n number of beliefs b, that performing the behaviour will lead to consequences i, with respect to the person's evaluation e of those consequences. Thus:

$$A_B = \sum_{i=1}^{n} b_i e_i \qquad (2)$$

Although this equation was developed to reason about actions, according to Petty and Cacioppo [5], it can also be used to assess attitudes towards people, objects and issues.

Secondly, the predictability that an intention will become a behaviour must also be evaluated with the respect to the person's motivation to comply with society (also called the subjective norm) or:

$$SN_B = \sum_{j=1}^{p} b_j m_j \qquad (3)$$

where b is the set of p number of beliefs about a behaviour B, that a group of individual j think are acceptable or unacceptable and m is the motivation to comply with j.

We now have a formula for the prediction of our agent's intentions and thus behaviour, which is determined by its current intentions I. Thus:

$$I = w_1 \sum_{i=1}^{n} b_i e_i + w_2 \sum_{j=1}^{p} b_j m_j \qquad (4)$$

where w_1 and w_2 are weightings added to represent the fact that attitudes and subjective norms are not always evaluated equally in the formation of behavioural intentions. The belief components, b_i and b_j, are the collection of the agents beliefs in the *event space* for the intention or:

$$E_I = b_i \, Y \, b_j \qquad (5)$$

where:

$$b_i \subseteq \{a_I, o_I, c_I, t_I\} \quad and \quad b_j \subseteq \{a_I, o_I, c_I, t_I\} \qquad (6)$$

Due to the complexity of event space, given a goal and corresponding number of relevant actions that may satisfy the goal, a number of intentions may be generated. Whilst Fishbein stated that a person's intention to perform a given behaviour is the best single predictor of whether or not the person will perform the behaviour, he also noted that predictions may be improved by measuring all of a person's intentions and alternative courses of action.

Petty and Cacioppo[5] suggest using the following formulae to calculate difference scores between intentions toward the target behaviour and intentions toward alternative behaviours. Thus:

$$I = I_{target} - [\sum_{i=1}^{p} I_{alternative\_i} / p] \qquad (7)$$

where alternative intentions are numbered from 1 to p.

6. Goal Setting, Motivation and the Theory of Reasoned Action

In the struggle to define human behaviour, Koestler [3] makes the observation that the human organism is not merely a mechanical device, but reacts to an ever-changing world and how humans react to the world is based on temporally dynamic goals.

In our test bed environment GOMASE (Goal-Orientated, Multi-Agent Simulation Environment), agents are driven by a goal hierarchy [8]. In this, we assume that a goal can be either abstract or primitive. An abstract goal can be broken down 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 [8]. 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.

Before the implementation of attitude theory, GOMASE agents simply used an opportunistic approach to carrying out tasks in order to satisfy an active goal. By now establishing an algorithm by which the agent can apply the theory of reasoned action in decision-making, we are creating a BDI agent that can make affective decisions and order intentions with human reasoning-like ability.

7. Applying Attitude Theory to Make Affective Decisions

We will now look at the example of affective decision making suggested by Picard [6]:

> *"Albert, a very busy scientist, has a beloved eight-week-old boy, and is trying to decide how to provide for his son while he works during the day. He does not know any family members or friends who could help. He acquires lists for three kinds of day care providers: a list of ten nanny referral services, a list of 145 licensed family car providers, and a list of 24 day care centers located nearby. He contemplates posting notices in newspapers and on bulletin boards. Albert loves his son, and want's to choose the best care for him. He needs a care-provider within a month. Albert is a highly rational man; how does he decide what to do?"*

In Picard's proposed solution each of the options available to Albert are weighted against each other by his attitude towards them. Albert's attitude towards the options is based in associated concepts that Albert already has attitudes about. From our point of view, each option presents a new event space to Albert. To evaluate his options, Albert forms an attitude about the new event space based on other overlapping event spaces about which he has already formed attitudes. For example, Albert believes that advertising attracts weirdoes. From this, we may assume that Albert has had a *bad* experience with advertising and he is applying that belief to this new situation. We will now examine one of Albert's options and how he formulated an attitude towards it.

One of Albert's options is to choose from a list of nannies. In order to calculate Albert's intention toward this option we need to calculate two parts; his attitudes and subjective norms toward the behaviour of choosing this option. Table 1 displays Albert's beliefs (taken from Picard's narrative solution) about using a nanny for day care and his evaluation of each of these beliefs. The beliefs are extracted from the elements that Albert considers to be in the event space for this behaviour. We will apply a similar 7 point scale as used in [5] for assessing his evaluations and beliefs such as +3 for 'good' or 'likely' and -3 for 'bad' or 'unlikely'.

Given that these are all of Albert's beliefs that are relevant to the hiring of a nanny and if this is how he assessed them, we can calculate Albert's attitude towards this option to be 6. To fully assess Albert's intention to hire a nanny, we must also examine his belief regarding others opinions about this action and his motivation to comply with them. Picard's narrative is not clear on the influence that other's opinions have on Albert's decision making, so we have collated an example to calculate his subjective norm in Table 2. Belief is again rated on a 7 point scale where 3 infers that the person is likely to approve of hiring a nanny and -3 infers that they do not. Motivation to comply is rated in a similar manner where 3 implies that Albert generally does what this person wants him to do and -3 implies that Albert hardly ever complies with this person's wishes.

Consequences of Hiring a Nanny for Day Care of Son	Belief (b) 3 = likely -3 = unlikely		Evaluation (e) 3 = good -3 = bad		$(b_i)(e_i)$
1. nanny would come to the house during the day.	3	×	3	=	9
2. A nanny would be expensive.	3	×	-1	=	-3
3. nanny-referral services want fees up front	3	×	-2	=	-6
4. nanny could be with us for many years	2	×	3		6
5. nanny may not be available to start within four weeks - same newspaper ads for nanny positions repeated over several weeks*	2	×	-3		-6
6. nannies can be abusive - a documentary he saw on TV*	-2	×	-3		6
$A_B = \sum_{i=1}^{n} b_i e_i$				=	6

* how these beliefs were formed.

Table 1. Determining Attitude (A) from b_1 and e_I

Important Referents	Belief (b) 3 = likely -3 = unlikely		Motivation to Comply (e) 3 = always -3 = never		$(b_i)(m_i)$
1. his mother Anna	3	×	2	=	6
2. his friend Joe	-1	×	-2	=	2
3. his brother Phillip	-3	×	0	=	0
$SN_B = \sum_{j=1}^{p} b_j m_j$				=	8

Table 2. Determining Subjective Norm (SN) from b_j and m_j

From the examples given, we calculate Albert's subjective norm about hiring a nanny to be a value of 8. If we assume that

Albert weights his attitude to be twice as important as his subjective norm i.e. $w_1 = 2$ and $w_2 = 1$, then we can assess his intention to hire a nanny to give a value of 20. Although this would suggest a high motivation in Albert to perform this intention, we must also take into consideration the other options that Albert has available to him and how he assesses these.

A similar method used to calculate Albert's intention to hire a nanny could be applied to his other options. Ranking these options in order of intention would give us further insight into and an improved prediction of what Albert's choice of behaviour regarding the search for a day care giver would be.

8. Simulating Albert

In this section, we will present the results from our simulation of Albert's reasoning and decision-making process. Our artificial agent, Albert, has been created in our test bed simulation environment GOMASE.

Initially Albert is given the task of finding day care for his son. Since our simulated Albert has never done this task before he asks other agents in the environment (simulated friends and relatives) for advice. As Albert gathers the information it is sorted and stored in his goal hierarchy (see [8] for a detailed description of this). Albert gathers together six options for day care. These include: ask a family member to baby-sit; ask a friend to baby-sit; put son in day care center; advertise for a day carer; place son in family home care; and, hire a nanny. A partial goal hierarchy of these options can be seen in Figure 1.

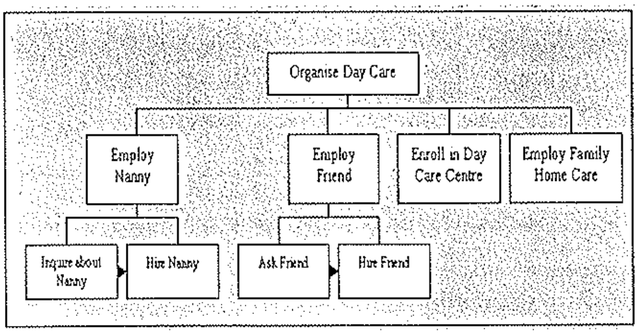

Figure 1. A Partial view of Albert's Goal Hierarchy.

The goal hierarchy structures Albert's goals from most general terms and decomposes them into sub-goals. At the very atomic level these goals translate into tasks which are arranged as activity digraphs. The digraphs for this example have been kept simple. Each sub-goal has two associated tasks. For example, for Albert to *Employ Nanny* he has a plan to first *Inquire about Nanny*, where he will gather more information about nannies, and if successful he can then do the second task, *Hire Nanny*.

Having developed a goal hierarchy and identified a number of plans that could be executed in order to satisfy his initial goal (to get day care for his son), Albert begins to evaluate the tasks using subsets of his beliefs that are in each task's event space. These are coupled with a subjective norm that he calculates from his friends and relatives opinions about the options.

From these assessments, Albert's intentions as to each option are formalized and the value of the intention is used as a prioritizing mechanism that determines his behaviour. Once each option has been assessed Albert's highest priority intention becomes his behaviour and the agent begins work on that intention.

Albert assesses his intentions
Agent Albert begins Inquire About Nanny
Albert asks Penny for information on Inquire About Nanny
Albert gets new information from Penny for Inquire About Nanny
Penny tells Albert that Inquire About Nanny is , Work at His Home, Expensive, Well Trained, Abusive, Permanent
Albert has a good feeling about Inquire About Nanny
Fortunately Penny can help Albert at this time for Inquire About Nanny
Albert assesses his intentions
Agent Albert begins Inquire About Family Care Givers
Albert asks Fiona for information on Inquire About Family Care Givers
Albert gets new information from Fiona for Inquire About Family Care Givers
Fiona tells Albert that Inquire About Family Care Givers is , Licensed, Stable, Have Big Dog
Albert has a bad feeling about Inquire About Family Care Givers and puts it off until later.
Albert assesses his intentions
Agent Albert begins Inquire About Friends
Albert asks Joe for information on Inquire About Friends
Albert gets new information from Joe for Inquire About Friends
Joe tells Albert that Inquire About Friends is , Imposing, Inexpensive
Albert has a bad feeling about Inquire About Friends and puts it off until later.
Albert assesses his intentions
Agent Albert begins Inquire About Day Care Centres
Albert asks Daniel for information on Inquire About Day Care Centres
Fortunately Daniel can help Albert at this time for Inquire About Day Care Centres
Albert assesses his intentions
Agent Albert begins Hire Nanny
Albert asks Penny for information on Hire Nanny
Fortunately Penny can help Albert at this time for Hire Nanny

Table 3. Narrative Output from GOMASE Simulation of Albert.

The narrative output from one run of GOMASE can be seen in Table 3. At the beginning of this simulation, Albert evaluated *Inquire About Nanny* as his highest priority intention. The agent then began simulating this task. During the simulation, the agent identifies and locates another agent in the environment that can assist him. Albert has preprogrammed knowledge that the agent called Penny is a nanny and he asks for her[2] assistance. He receives new information from her about hiring a nanny. After Albert integrates this new information into his knowledge base, he re-evaluates his attitude towards this option. As can be seen from the narrative, Albert has formulated a *good* feeling about the task.

However, Albert does not continue with this plan. Although he feels good after inquiring about a nanny, it does not necessarily mean hiring a nanny is the option that he feels the most positive about. As can be seen in Table 3. Albert inquires about a number of options before he decides to go ahead and hire the nanny. As he gathers more information, the priorities on his

[2] Although our simulated agents do not have a specific gender, by referring to an agent as her or him (rather than it) we can better represent the simulated scenario in the mind's eye of the reader and better relate the situation with real-life, human-social, settings.

intentions change. A negative evaluation of the information gives Albert a bad feeling and a positive evaluation, a good feeling. Intentions are prioritized from the most positive option to the most negative.

9. Conclusion

In this paper we have begun to explore the development of an affective reasoning agent. The mechanism on which we base our affective calculations is that of attitude. By using attitude as a measurement, we are able to program an agent with a means of appraising an event as either pleasurable or displeasurable. This pleasure rating on the event acts as a device for the prioritizing of an agent's intentions with respect to the agent's current set of goals.

Fishbein and Ajzen's theory of reasoned action, fits well with BDI agent architecture as the belief and intention structure is already in place. The subjective norm aspect adds to the agent's social responsibilities, concern for others and his own moral beliefs. This acts as the agent's conscience and gives the agent a means of interacting socially with other agents using an aspect of human behaviour.

In Picard's suggested solution to Albert's problem, she explains that in making his decision, Albert does not evaluate all possible solutions but works along a path of *feel good* options until a negative feeling arises at which point he changes tack and begins evaluating other options. In this example we are allowing Albert to evaluate all of his options and to choose from the best one based on his feelings towards it. However, in an environment where the options are innumerable, Picard's suggestion would be the most appropriate approach in finding a solution by affective reasoning.

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 BDI agent architecture.

10. References

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

[2] Fishbein, M. and I. Ajzen, Belief, Attitude, Intention and Behaviour. 1975, London: Addison-Wesley.

[3] Koestler, A., The Ghost in the Machine. 1967, London: Penguin Books Ltd.

[4] Ortony, A., G.L. Clore, and A. Collins, The Cognitive Structure of Emotions. 1988, Cambridge, UK.: Cambridge University Press.

[5] Petty, R.E. and J.T. Cacioppo, Attitudes and Persuasion: Classic and Contemporary Approaches. 1996, Boulder, Colorado: Westview Press.

[6] Picard, R., Affective Computing. 1998, London: The MIT Press.

[7] Rao, A.S. and M.P. Georgeff. An abstract architecture for rational agents., In: Proceedings of Third International Conference on Principals of Knowledge Representation and Reasoning. 1992. San Mateo, C.A.: Morgan Kauffmann Publishers.

[8] Ray, P. and D. Lukose. To Believe, to Desire, to Intent: Motivating Autonomous Intelligent Agents with a Mind of their Own. In: Proceedings of Managing Enterprises 99. 1999. Newcastle.

[9] Reilly, W.S.N., Believable Social and Emotional Agents, P.h.D. Dissertation in School of Computer Science . 1996, Pittsburgh: Carnegie Mellon University.

[10] Velásquez, J. 1998 "Modeling Emotion-Based Decision-Making." In: Proceedings of the 1998 AAAI Fall Symposium Emotional and Intelligent: The Tangled Knot of Cognition (Technical Report FS-98-03). Orlando, FL: AAAI Press.

[11] Wright, I.P., Emotional Agents, P.h.D Dissertation in School of Computer Science. 1997, University of Birmingham: Birmingham. p. 267.