Towards a Motivation-Based Approach for Evaluating Goals

Steve Munroe, Michael Luck
Electronics and Computer Science
University of Southampton
Southampton, UK
{sjm01r, mml}@ecs.soton.ac.uk

Mark d'Inverno
Computer Science
University of Westminster
London, UK
dinverm@wmin.ac.uk

ABSTRACT

Traditional goal-oriented approaches to building intelligent agents only consider absolute satisfaction of goals. However, in continuous domains there may be many instances in which a goal state can only be partially satisfied. In these situations the traditional symbolic goal representation needs modifying in order that an agent can determine a *worth value* of a goal state and also of any state approximating the goal. In our work we use the concept of *worth* in two ways. First, we propose a mechanism by which the worth of a goal is dynamically set as a function of the intensity of an underlying *motivation*. Second, we determine the worth of any state in relation to a goal through the use of a *metric* by which we can measure the *proximity* of an environmental state to a goal. In this way, it is possible to make judgements about the *relative* satisfaction an environmental state offers in regard to a goal.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence-Multiagent systems

General Terms

Algorithms, Design

Keywords

Decision making, Goal Generation, Motivation

1. INTRODUCTION

Overcoming the limitations of symbolically based representations as used in intelligent agents, to cope with more realistic domains, is an area growing in size. Work from robotic control [5], design-to-criteria scheduling [8] and cognitive appraisal theory [7] all relate to extending the abilities of intelligent agents into continuous, worth-oriented domains.

In this paper, we use the concept of *motivation* to combine the notion of goals as symbolic state descriptions with the notion of goals having dynamically set worth evaluations. This allows us to assign different values to a goal state depending on the current context of the agent.

2. MOTIVATED AGENTS

The notion of motivation is increasingly being used as the basis for control of autonomous agents. Indeed, motivation has already been investigated in terms of goal generation [1], proactive behaviour [4] and information processing [2].

AAMAS'03, July 14–18, 2003, Melbourne, Australia.
ACM 1-58113-683-8/03/0007.

Like the traditional notion of utility, motivation places value on actions and world states, but is a more wide-ranging concept than utility. In the traditional view, an agent examines its options and chooses one with the highest utility. Motivation also enables this to be done, but it is also more intimately involved with the agent's decision-making process than utility. A motivated agent has a dynamically changing internal environment, provided by motivations, which can influence its decisions. For example, in the presence of food, an agent may or may not choose to eat depending on the state of its internal environment (specifically its hunger motivation).

Motivation thus helps to influence and direct an agent's decision making. The influence of a given motivation over an agent's decision-making increases or decreases in response to changes in the environment or changes in the state of the agent's other motivations. We use the notion of *intensity* to capture this dynamic property. Intensity is here represented as a real-valued number in the range [0,1] where 0 represents no intensity and 1 represents maximum intensity. The more intense a motivation, the more influence that motivation will exert over decision-making. Thus, at any given time, each of an agent's motivations is characterised by an intensity that provides some indication of the motivation's suitability in the current environment. That is to say that when a motivation has high intensity, any goals generated should be highly *relevant* for the agent in the current environment. In order to achieve this, however, there must be some way of assessing the current environment in terms of *relevance* to motivations. The simplest way to achieve this is by attaching a set of *cues* to a motivation that determine when, and by how much, the motivation's intensity should be updated. Cues represent those salient aspects of the environment that an agent has some interest in registering or tracking. Their representation can be considered to take the form of simple attributes, which are simply those things that are potentially perceivable by an agent.

We introduce the given set $attribute$ to represent all such perceivable things[1].

$$[Attribute]$$

An environment is then a non empty set of attributes.

$$Environment == \mathbb{P}_1 \, Attribute$$

As an example of a cue, I may be interested in the tidiness of my room. One potential cue that I could use may be the number

[1]In this work we adopt the Z notation [6] which is based on set-theory and first order logic. Though we assume some familiarity with Z, the meaning should be clear.

of books lying around. When the number of books reaches some threshold, I may decide to take some action. Alternatively, I may be interested in whether or not some person is near, when they appear I may decide to take some action, such as smiling at them or running away. The appearance of a cue in an agent's view of its environment, or a value take from the cue, results in the agent's motivational intensity being updated by some amount. This amount can either be fixed, or depend on some *measurement*. Each of these update methods calls for a different type of cue, *categorical* and *continuous*.

Categorical cues update motivations by discrete amounts, whereas continuous cues alter the intensity of a motivation in proportion to the value of the cue, as with the *number* of books in the untidy room example above. Both types of cue are represented by sets of attributes.

$$CategoricalCue == \mathbb{P}_1\ Attribute$$
$$ContinuousCue == \mathbb{P}_1\ Attribute$$

The set of cues attached to a motivation can be a combination of both types of cues and as such we define cues to be either categorical or continuous.

$$\begin{aligned} Cue ::= &\ catcue\langle\langle CategoricalCue\rangle\rangle \\ |&\ \ concue\langle\langle ContinuousCue\rangle\rangle \end{aligned}$$

Motivation is considered to have seven basic components: a unique identifier; a current and maximum intensity value; a set of goals that can be used to mitigate the motivation (taking the form of state descriptions); a set of cues that lead to updating motivational intensity; and two functions that compute the effects that each type of cue has on intensity; a discrete effect function for categorical cues and a continuous effect function for continuous cues.

2.1 Updating Motivations

In order to update motivation an agent needs a way to perceive its environment so that it can register the cues it is interested in and activate the appropriate motivations. To represent this formally, we define the function *selectActive*, which takes an environment and a set of motivations, and returns a set of motivations whose cues appear in the current environment[2]

$$\begin{aligned} selectActive : Environment \to \mathbb{P}\ Motivation \\ \to \mathbb{P}\ Motivation \end{aligned}$$

Now that an agent can identify which motivations are relevant in the current environment, we need to update the motivations by the appropriate amounts defined by the cue types. Below we present a discrete update function $dUpdate$. It takes a categorical cue, a motivation and returns a motivation with the new intensity value that is the minimum of either the maximum intensity value or the new intensity determined by the cue.

$$dUpdate : Cue \to Motivation \to Motivation$$

We define the continuous update function similarly.

$$\begin{aligned} cUpdate : Cue \to View \to Motivation \\ \to Motivation \end{aligned}$$

2.2 Goals and Motivation

In this section we describe the two ways in which we use the concept of worth. Worth is related both to goals and environmental states. Our motivated agent is able to dynamically assign a worth to a goal depending on the state of its motivations. Once a goal is evaluated in this way, the agent can then place worth values on those environmental states that approximate the goal state. Below we show the signature of a function that calculates a goal's worth. It takes a motivation, a goal and the current environment and gives a value to the goal determined by both the motivational intensity and the current environment. This value represents the worth of the goal.

$$\begin{aligned} Goalworth : Motivation \to Goal \\ \to Environment \to RAT_0^1 \end{aligned}$$

The effect of any given state on a motivation is given by a *mitigation function*, which takes a goal, the associated motivation and the current environment, and returns a value that is used to *mitigate* the intensity of the motivation. The value placed on the current state represents the value gained from that state as a function of the *distance* of that state from the goal state and the current intensity value of the motivation. States that completely match those defined by a goal should maximally mitigate the associated motivation. Below we show the signature of the mitigation function.

$$\begin{aligned} mitigate : Goal \to Environment \\ \to Motivation \to RAT_0^1 \end{aligned}$$

By calculating the proximity of one state to another we can calculate the worth gained from the current environmental state in relation to a goal state.

3. CONCLUSION

We have presented a motivational mechanism that enables an agent to flexibly assign worth to a symbolically represented goal as a function of both the agent's current environment and the agent's motivational state. We have also presented a mechanisms to calculate the worth of a given environmental state to a goal state. This enables an agent to gain some value from environmental states that fall short of a goal state. Future work will involve expanding the model to cope with multiple goals, constraints and the social context of an agent.

4. REFERENCES

[1] M. d'Inverno and M. Luck. *Understanding Agent Systems*. Springer, 2001.

[2] D. Moffat and N. Frijda. Where there's a will there's an agent. In M. Wooldridge and N. R. Jennings, editors, *Intelligent Agents (LNAI Volume 890)*, pages 245–260. 1995.

[3] S. Munroe, M. Luck, and M. d'Inverno. Towards motivation-based decisions for worth goals. In *The Proc. of the 3rd International/Central and Eastern European Conference on Multi-Agent Systems*, To appear 2003.

[4] T. J. Norman and D. Long. Goal creation in motivated agents. In M. Wooldridge and N. R. Jennings, editors, *Intelligent Agents (LNAI Volume 890)*, pages 277–290. 1995.

[5] E. Spier and D. McFarland. In *A Finer-Grained Motivational Model of Behaviour Sequencing*, 1996.

[6] M. Spivey. *The Z Notiation, 2nd ed.* Prentice Hall, Hemel Hempstead, 1992.

[7] A. Staller and P. Petta. Towards a tractable appraisal-based architecture. In *From Animals to Animats: Proc. of SAB'98.*, 1998.

[8] T. Wagner and V. Lesser. Design-to-criteria scheduling: Real-time agent control. In *Proc of AAAI 2000 Spring Symp on Real-Time Autonomous Systems*, pages 89–96, Stanford, CA, March 2000.

[2] For this and all other functions we provide only their signatures. Interested readers should see [3] for complete specifications of these functions and more detail in general