AFFECT AND AGENT CONTROL: EXPERIMENTS WITH SIMPLE AFFECTIVE STATES

MATTHIAS SCHEUTZ
Department of Computer Science and Engineering
University of Notre Dame, Notre Dame, IN 46556, USA
E-mail: mscheutz@cse.nd.edu

AARON SLOMAN
School of Computer Science
The University of Birmingham, Birmingham, B15 2TT, UK
E-mail: axs@cs.bham.ac.uk

We analyse control functions of affective states in relatively simple agents in a variety of environments and test the analysis in various simulation experiments in competitive multi-agent environments. The results show that simple affective states (like "hunger") can be effective in agent control and are likely to evolve in certain competitive environments. This illustrates the methodology of exploring neighbourhoods in "design space" in order to understand tradeoffs in the development of different kinds of agent architectures, whether natural or artificial.

1 Introduction

Affective states (such as emotions, motivations, desires, pleasures, pains, attitudes, preferences, moods, values, etc.) and their relations to agent architectures have been receiving increasing attention in AI and Cognitive Science.[1,2,3] Detailed analyses of these subspecies of affect should include descriptions of their functional roles in contributing to useful capabilities within agent architectures[4], complemented by empirical research on affect in biological organisms and concrete experiments with synthetic agent architectures, to confirm that the proposed architectures have the claimed properties. Our approach contrasts with most evolutionary AI research, which attempts to discover what can evolve from given inital states. Instead, we explore "neighbourhoods" and "mini-trajectories" in design space, by starting with examples of agent architectures, then explicitly provide possible extensions with evolutionary operators that can select them, and run simulations to investigate which of the extensions have evolutionary advantages in various environments. This can show how slight changes in environments alter tradeoffs between design options.

To illustrate this methodology, we next analyse functional roles of affective states and then describe our simulation experiments which show how certain simple affective control mechanisms can be useful in a range of environments and are therefore likely to evolve in those environments.

2 What Affective States are and aren't

If we attempt to define "affective" simply in terms of familiar examples, such as "desiring', "having emotions", "enjoying", etc. we risk implicitly restricting the notion to organisms with architectures sufficiently like ours. That could rule out varieties of fear, hunger, or aggression found in insects, for example. We need an architecture-neutral characterisation, which is hard to define if it is to be applicable across a wide range of architectures (such as insect-like reactive architectures or deliberative architectures with mechanisms able to represent and reason about non-existent and possible future states). Our best hope is to define "affective" in terms of a functional role which can be specified independent of the specific features of an architecture.

The intuitive notion of "affect" already has two aspects that are relevant to a variety of architectures, namely *direction* and *evaluation*. On the one hand there is direction of internal or external behaviour, for instance, wanting something or trying to avoid something. On the other hand there is positive or negative evaluation of what is happening, for instance, enjoying something or finding it unpleasant. However, even evaluation is linked to direction insofar as enjoying involves being disposed to preserve or repeat and finding painful involves being disposed to terminate or avoid. Either way affective states are examples of *control states*[5].

Yet, not all states in control systems are affective states, even if they have some effect on internal or external behaviour. For instance, perceiving, knowing, reasoning, and self-monitoring can influence behaviour but are not regarded as affective. Suppose an agent can use structures as representations of states of affairs (never mind how). Anything that represents must be capable of failing to represent. There are various kinds of mismatch, and in some cases the mismatch can be detected, for instance perceiving that some desired state has or has not been achieved, or that a goal is being approached but very slowly. If detection of a mismatch has a disposition to cause some behaviour to reduce the mismatch there are (to a first approximation) two main cases: (1) the behaviour changes the representation to fit the reality, or (2) the behaviour changes reality to fit the representation. In (1) the system has a "belief-like" state, and in (2) a "desire-like" state. In other words, belief-like states tend to be changed to make them fit reality, whereas attempts are made to change reality to make it fit desire-like states. It is this distinction between belief-like and desire-like control states that can give us a handle on how to construe affective states, namely as "desire-like" control states whose role is initiating, evaluating and regulating, internal or external behaviour, as opposed to merely acquiring, interpreting, manipulating, or storing information (that might or might not be used in connection with affective states to initiate or control behaviour).

A state representing the current position of an effector, or the location of food in the environment, or the agent's current energy level is, therefore, not an affective state. However, states derived from these which are used to initiate, select, prioritise, or modulate behaviour, either directly or indirectly via other such states would be affective states. An example might be using a measurement of the discrepancy between current energy level and a "target" level (a "hunger" representation), to modulate the tendency of the system to react to perceived food by going for it. This might use either a "hunger threshold" to switch on food-seeking or a continuous gain control.

In complex cases, the "reference states" used to determine whether corrective action is required may be parametrised by dynamically changing measures or descriptions of the sensed state to be maintained or prevented, and the type of corrective action required, internally or externally. For instance, an organism that somehow can record how frequently food sources are encountered might use a lower hunger threshold to trigger searching for food. If sensitive to current terrain it might trigger different kinds of searches in different terrains. Thus while the records of food frequency and terrain features are acquired they function as components of perceptual or belief-like states, whereas when they are used to modulate decision making they function as components of affective states.

Affective states can vary in cognitive sophistication. Simple affective mechanisms can be implemented within a purely reactive architecture, like the "hunger" example. More sophisticated affective states which include construction, evaluation and comparison of alternatives, or which require high-level perceptual categorisations, would require the representational resources of a deliberative architecture. However, recorded measurements or labels directly produced by sensors in reactive architectures can have desire-like functions, and for that reason can be regarded as affective states that use a primitive "limiting case" class of representations[6].

The remainder of this paper describes simulation experiments where agents with slightly different architectures compete for resources in order to survive in a carefully controlled simulated environment. Proportions surviving in different conditions help to show the usefulness of different architectural features in different contexts. It turns out that simple affective states can be surprisingly effective.

3 The Simulation Environment

The simulated environment consists of a rectangular surface of fixed size (usually around 800 by 800 units) populated with various kinds of agents and other objects such as "lethal" entities of various sizes, some static and some moving at different speeds in different directions, and "food items" (i.e., energy sources which pop up at random locations and disappear after a pre-determined period of time unless consumed by agents). Agents use up energy at a fixed rate, when stationary, and require additional energy proportional to their speed, when moving. Hence, they are in permanent need of food, which they can consume sitting on top of a food source in a time proportional to the energy stored in the food source depending on the maximum amount of energy an agent can take in at any given time. Agents die and are removed from the simulation if they run out of energy, or if they come into contact with lethal entities or other agents.

All agents are equipped with a "sonar" sensor to detect lethal entities, a "smell" sensor to detect food, a "touch" sensor to detect impending collisions and an internal sensor to measure their energy-level. For both sonar and smell sensors, gradient vectors are computed and mapped onto the effector space (see below), yielding the direction in which the agent will move. The touch sensor is connected to a global alarm system, which triggers a reflex to move away from anything touched, unless it is food. These movements are initiated automatically and cannot be controlled by the agent. They are somewhat erratic and will slightly reorient the agent (thus helping it to get out of "local minima").

On the effector side, agents have motors for locomotion (forward and backward), motors for turning (left and right in degrees) and a mechanism for consuming food.

After a certain number of simulation cycles, agents reach maturity and can procreate asexually, in which case depending on their current energy level they will have a variable number of offspring which pop up in the vicinity of the agent one at a time (the energy for creating a new agent is subtracted from the parent, occasionally causing the parent to starve). While different agents may have different short term goals at any given time (e.g., getting around lethal entities or consuming food), common to all of them are the two implicit goals of *survival* (i.e., to get enough food and avoid running into/getting run over by lethal entities or other agents) and *procreation* (i.e., to live long enough to have offspring).

For evolutionary studies, a simple mutation mechanism modifies with a certain probability some of the agent's architectural parameters (e.g., the parameters responsible for integrating smell and sonar information). Some offspring will then start out with the modified parameters instead of being exact copies of the parent. This mutation rate as well as various other parameters need to be fixed before each run of the simulation (a more detailed description of the simulation and its various control parameters is provided elsewhere)[7].

In is worth pointing out that our setup differs in at least two ways from other simulated environments that have been used to study affective states.[8,9,10,11,12] First, by allowing agents to procreate (i.e., have exact copies of themselves as offspring) we can study trajectories of agent populations and can thus identify properties of architectures that are related to and possibly influence the interaction of agent populations. And second, by adding mutation, we can examine the potential of architectures to be modified and extended over generations of agents. In particular, by controlling which components of an architecture can change while allowing for randomness in the way they can change, we are able to study evolutionary tradeoffs of such extensions/modifications. From these explorations of "design space" and "niche space"[13] we cannot only derive advantages and disadvantages of architectural components, but also the likelihood that such components would have evolved in natural systems using natural selection.

4 The Agents and their Architectures

In the following we consider two kinds of agents: *reactive agents* (R-agents) and *simple affective agents* (A-agents) (other studies have compared different kinds[7]). R-agents process sensor information and produce behavioural responses using a schema-based approach, which obviates the need for a special action selection mechanism: both smell and sonar sensors provide the agent with directional and intensity information of the objects surrounding the agent within sensor reach, where $intensity = 1/distance^2$ (i.e., the distance of the object from the current position of the agent). The sum of these vectors (call them S and F for sonar and food, respectively) is then computed as a measure of the distribution of the respective objects in the environment and passed on to the motor schema, which maps perceptual space into motor space yielding the direction, in which to go: $\delta S + \gamma F$ (where δ and γ are the respective gain values).[a]

A-agents are *extensions* of R-agents. They have an additional component, which can influence the way sensory vector fields are combined by altering the gain value γ based on the level of energy. In accordance with our earlier analysis of affective states as modulators of behaviours and/or processes, this component implements an affective state, which we call "hunger".

The difference in the architecture gives rise to different behaviour: R-agents are always "interested" in food and go for whichever food source they can get to, while A-agents are only interested in food when their energy levels are low. Otherwise they tend to avoid food and thus competition for it, reducing the likelihood of getting killed because of colliding with other competing agents or lethal entities.

5 The Behavioural Potential of a Simple Affective State

We start our series of experiments by checking whether each agent kind can survive in various kinds of environments on its own. Five agents of the same kind are placed in various environments (from environments with no lethal entities to very "dangerous" environments with both static and moving lethal entities) at random locations to "average out" possible advantages due to their initial location over a large number

[a] Note that this formula leaves out the details for the touch sensor for ease of presentation.

Table 1. Surviving agents in an n-environment when started with 5 agents of only one kind.

	R-agents			A-agents		
Env	μ	σ	Con	μ	σ	Con
0	14.60	2.80	1.73	19.20	2.74	1.70
5	13.20	4.78	2.96	17.20	3.05	1.89
10	11.90	3.81	2.36	17.20	3.77	2.33
20	11.60	3.47	2.15	15.40	3.95	2.45
30	7.50	4.43	2.75	13.00	3.56	2.21
40	2.90	3.57	2.21	10.40	3.57	2.21
50	0.20	0.63	0.39	8.00	3.56	2.21

Table 2. Surviving agents in an n-environment when started with 5 agents each of boths kinds.

	R-agents			A-agents		
Env	μ	σ	Con	μ	σ	Con
0	0.00	0.00	0.00	17.20	3.61	2.24
5	0.00	0.00	0.00	16.30	2.91	1.80
10	1.60	5.06	3.14	14.50	6.54	4.05
20	0.10	0.32	0.20	14.50	4.22	2.62
30	0.00	0.00	0.00	15.10	3.35	2.08
40	0.00	0.00	0.00	12.80	2.49	1.54
50	0.00	0.00	0.00	10.00	3.16	1.96

of trials. The "food rate" is fixed at 0.25 and the procreation age at 250 update cycles. Table 1 shows for each agent kind the average number (μ) of surviving agents as well as standard deviation (σ) and confidence interval (Con) for $\alpha = 0.05$ taken over 10 different runs of the simulation, each for 10000 environmental updates for a given environment (where "n-environment" is intended to indicate that n static and n moving lethal entities were placed at random in the environment). Note that A-agents do significantly better than R-agents if measured in terms of the average number of agents in each environment at any given time.

Given that each agent kind can survive on its own, we now compare the performance of mixed groups of R- and A-agents. It turns out that A-agents reliably outperform R-agents in all considered environments (see Table 2).

The results depend neither on the initial number of agents nor on the distribution of moving and static lethal entities: experiments with different numbers of initial agents of each kind as well as experiments with different numbers of moving and static lethal entities (that added up to the same total) yield very similar results. Higher food rates (e.g., of 0.5) do not change the picture either, rather they show even more clearly the ability of affective agents to coexist in large groups. With lower food rates the advantage of A-agents over R-agents slowly decreases as waiting for hunger to grow before moving towards food is not a good strategy. Eventually, at food rates of 0.125 and below, survival in crowded environments becomes impossible for any agent kind–there are simply too many lethal entities obstructing the paths to food.

The superior performance of A-agents might not seem very surprising, since the additional information about the current energy level, ignored by R-agents, but utilized by A-agents, allows for a more complex mapping between sensory input and behavioural output. However, using more information does not automatically lead to better performance, as can be seen from the fact that A-agents may lose out against R-agents if the "rules of the game" are slightly modified: in a simulation without procreation, where either the numbers of surviving agents of each kind are counted after a predetermined number of cycles or the average lifespan of an agent is used as a measure of fitness, R-agents almost always perform (slightly) better than A-agents (in all of the above environments). Only in combination with procreation is the tendency of A-agents to distribute themselves better over the whole environment (in Seth's terminology: their lower degree of "clumpiness"[12]) by virtue of being at times less attracted to food beneficial, as their offspring will benefit from not having to compete immediately with many other agents in their vicinity. In this light, the answer to the following question whether A-agents can be produced by some evolutionary process is not obvious at all.

6 Simple Affective States Can Be Evolved

To study the degree to which simple affective states like "hunger" can be evolved in a competitive environment, we allowed for mutation of the link between the component connected to the energy sensor (which is supposed to assume the role of the affective "hunger" state) and the component encoding the food gain value γ in the mapping from perceptual to motor space. This link, expressed as a multiplicative factor and called "foodweight", is initialised at random in the interval $(-0.2, 0.2)$. Whenever an agent has offspring, the probability for "genetic modification" of the foodweight is 1/3 and the probability for weight increase/decrease (by the given factor $\tau = 0.05$) is 1/6, respectively. Everything else remains the same.

Of all seven environments, A-agents did not survive in the 40- and 50-environments, which are very tough in that wrong moves are punished right away: there is simply no room for genetic trial and error.[b]

In the other five environments, A-agents evolved using the state in the expected way, although to varying degrees: the less crowded an environment, the better the use

[b]The only agents that survived on 2 out of 10 runs were the R-agents in 40-environments.

Table 3. Average weight values, standard deviation and confidence level at $\alpha = 0.05$ for the "foodweight" of the surviving affective agents in one run in an n-environment.

	Foodweight				Foodweight		
Env	μ	σ	Con	Env	μ	σ	Con
0	0.26	0.09	0.05	10	0.19	0.07	0.05
	0.27	0.05	0.03		0.30	0.04	0.03
	0.23	0.08	0.05		0.29	0.09	0.05
	0.07	0.05	0.03	20	0.24	0.12	0.00
	0.17	0.06	0.04	30	0.17	0.11	0.07
5	0.33	0.12	0.07		0.13	0.07	0.04
	0.19	0.06	0.03	40	0.00	0.00	0.00
	0.10	0.09	0.06	50	0.00	0.00	0.00
	0.18	0.04	0.03				

Table 4. Surviving affective agents in an n-environment when started with 5 R-agents, which can have randomly initialised A-agents as part of their offspring with a probability of 0.25. No R-agent survived a single run.

	A-agents			Foodweight		
Env	μ	σ	Con	μ	σ	Con
0	17.90	3.60	2.23	0.18	0.10	0.06
5	14.90	1.91	1.19	0.19	0.11	0.07
10	15.10	3.03	1.81	0.19	0.11	0.07
20	11.53	8.39	5.20	0.18	0.09	0.05
30	3.80	5.45	3.38	0.17	0.09	0.06
40	4.10	5.57	3.45	0.21	0.10	0.06
50	1.00	2.31	1.43	0.21	0.09	0.06

of the state can be evolved, the reason being that agents with initial random weights are very likely to be inefficient in navigating through the environment, if able at all. In such cases it is helpful if food is not obstructed by too many lethal entities. Table 3 shows for each environment mean, standard deviation and confidence interval (again for $\alpha = 0.05$) for weights for all those runs on which affective agents survived. The above experiment also works for different mutation rates as well as different values of τ. Note that while in 5 out of 10 runs the "affective use" of the state was evolved in 0-environments, only in 2 out of 10 runs was the use evolved in 30-environments.

The positive value of the foodweight indicates that the hunger state deserves its name. Yet, the magnitudes of the weight seem small given the procreation age of 250 and the increment/decrement factor $\tau = 0.05$. On closer inspection, however, it turns out that evolution was quite fast, since assuming that there are only about 40 generations of agents in each run, and given that the probability of a positive increase of the weight by τ is 1/6, then starting at a slightly positive hunger weight, say, the maximum we should expect is about 1/3.

We have not dealt with issues of genetic coding and how genetic codes relate to the "added machinery" in the cognitive architecture of affective agents. Rather, we assume that adding a realizer of such a state (e.g., a neuron) is an evolutionarily feasible operation (e.g., which could result from some sort of duplication operation of segments of genetic information[14]) and that mutation on genetically coded weight information can lead to an increase or decrease of weight values.

We have, however, considered an evolutionarily more plausible variant of the experiment. Starting with R-agents, let some of their offspring have additional architectural capacities with a certain probability (in our case, the capacities of A-agents). The probability, with which R-agents have such randomly initialised A-agents as offspring is 0.25 (the results are also valid for much lower rates such as 0.05). It turns out that environments with only R-agents in the beginning will eventually be also populated by A-agents (most of the time exclusively, see Table 4).

It is worth mentioning that the results of this section also hold for extended simulations, where agents need a second resource (e.g., water) for survival. Multiple affective control states (e.g., "hunger" and "thirst") are even more beneficial when agents have multiple needs, which can be seen from the fact that R-agents can hardly survive on their own in such a setting (to "always go for the nearest resource" is not simply not a good strategy, e.g., see [11]). They even lose against A-agents if fitness is determined without procreation (see the end of the last section).

7 Discussion

The above experiments help us understand some of the conditions in which affective states like hunger have survivial value, and indicate that in certain competitive environments, if there is an option to develop new architectural resources that implement such affective states, then these resources will likely evolve. Especially the last result is not obvious, for a reason that makes the question why higher species with more complex and sophisticated control architectures evolved in the first place so fascinating: every species along an evolutionary trajectory has to have a viable control architecture, which allows its individuals to survive and procreate, otherwise it will die out. This is a very severe constraint imposed on trajectories in design and niche space, which we are only slowly beginning to understand.

Our investigations are, of course, just a start. Many more experiments using different kinds of affective states are needed to explore the space of possible uses of affective states and the space of possible affective states itself. We have begun to explore a slightly different neighbourhood in design space by allowing some agents to have deliberative capabilities, and comparing them with A-agents[7]. In a surprising variety of environments the deliberative agents do less well, though a great deal more investigation is needed. Further work on the capacities of affective states as control mechanisms and the likelihood of their evolution in certain environments should thus also help to explain why evolutionary developments that increase intelligence by adding a deliberative layer were favoured by so few species!

Acknowledgments

The work was conducted while the first author was on leave at the School of Computer Science at the University of Birmingham and funded by the Leverhulme Trust.

References

1. K. Oatley and J.M. Jenkins. *Understanding Emotions* (Blackwell, Oxford, 1996).
2. R. Picard. *Affective Computing* (MIT Press, Cambridge, Mass, 1997).
3. G. Hatano, N. Okada, and H. Tanabe, eds. *Affective Minds* (Elsevier, Amsterdam, 2000).
4. A. Sloman, In *Cognitive Processing*, 1 (2001).
5. A. Sloman. In *Philosophy and the Cognitive Sciences*, eds. C. Hookway and D. Peterson (Cambridge University Press, 1993).
6. A. Sloman. In *Forms of representation: an interdisciplinary theme for cognitive science*, ed. D.M.Peterson (Intellect Books, Exeter, U.K., 1996).
7. M. Scheutz and B. Logan. In *Proceedings of AISB 01* (AISB Press, 2001).
8. P. Maes. In *From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior*, eds. J.A. Meyer and S.W. Wilson (MIT Press, Cambridge, MA, 1991).
9. T. Tyrrell. *Computational Mechanisms for Action Selection* (Ph.D. Thesis, University of Edinburgh, 1993).
10. D. Cañamero. In *Proceedings of the First International Symposium on Autonomous Agents* (AA'97, Marina del Rey, CA, The ACM Press, 1997).
11. E. Spier *From reactive behaviour to adaptive behaviour* (Ph.D. Thesis, University of Oxford, 1997).
12. A. Seth. *On the Relations between Behaviour, Mechanism, and Environment: Explorations in Artificial Evolution* (Ph.D. Thesis, University of Sussex, 2000).
13. A. Sloman. In *Parallel Problem Solving from Nature – PPSN VI, LNCS 1917*, eds. M. Schoenauer, *et al.* (Berlin, Springer-Verlag, 2000).
14. J. Maynard-Smith and E. Szathmáry. *The Origins of Life: From the Birth of Life to the Origin of Language* (Oxford University Press, Oxford, 1999).