The Evolution of Affective States and Social Control

Matthias Scheutz
Department of Computer Science and Engineering
University of Notre Dame, USA
mscheutz@cse.nd.edu
http://www.nd.edu/˜mscheutz/

We present a research strategy to study the effects of various components in an agent architecture (e.g., feedback circuits that can implement various control states) on the dynamics of species interactions in a simulated environment. In particular, we discuss the consequences of behaviors elicited by affective states (like "hunger" or "fear") for whole agent societies. Construing affective phenomena as processes implemented in part by the agent, but partly also resulting from the agent-environment interaction (including interactions with other agents), we report simulation results that suggest that certain affective states are very likely to evolve in competitive multi-species environments because they lead to behavior that is beneficial at the society level. From this we conclude that further studies along the lines we suggest might reveal evolutionary trajectories from simple individual reflex-like agents to much more complex multi-agent societies, possibly also hinting at why relatively few species have developed a complex deliberative control system.

Affective phenomena are wide-spread in nature. Especially in social interactions, we can observe many different kinds of affect at work, from expressive prosody in speech, to facial expressions of emotions, to bodily gestures expressing attitudes. Humans, and to a lesser degree many other animals, seemed to be locked into an intricate network of affective control that involves expressing one's own affect as well as recognizing and reacting to the affect of others. While it seems obvious from their mere presence that affective phenomena must play some role in social situations–otherwise they would not have survived evolutionary pressures–their precise function is not that clear. Even less clear is how affective phenomena should be categorized and how the so-obtained categories could be mapped onto possible organizations of agent control architectures that are able to bring about or contribute to bringing about these phenomena.

We believe that part of the problem with a satisfactory characterization of affect has to do with the limiting view of affect as "state" (of an agent), which can only provide a "snap shot picture" of what may be better understood as an enduring process of *behavior coordination*: actions and reactions, adjustments and modifications, anticipation and compensation of behavior in various social situations. Often it is not a single inner state of an agent that determines whether an agent has, experiences or displays some form of affect, but rather a whole sequence of such states in combination with environmental states (e.g., “fear” does not refer to the make-up of an agent at a particular moment in time, but to the unfolding of a sequence of events, starting from the perception of a potentially threatening environmental condition, to a reaction of the agent’s control system, to a reaction of the agent’s body, to a change in perception, etc.). Consequently, we construe affective phenomena–different from most research on affect–as being intrinsically process-like rather than state-like phenomena, and use the term “affective state” only where common language usage favors it over other (better) ways of referring to affective phenomena. We believe that the process-view of affect can help to elucidate the functional role of affective phenomena as coordination mechanisms that allow many animals (and also humans) to regulate group behavior.

Methods

In accordance with the above mentioned view of affect as means of behavior coordination, we then construe affective processes as instances of “control processes”. As such they serve the primary purpose of initiating, ending, interrupting, modulating, regulating, predicting, and adjusting agent behavior. From simple homeostatic control states like hunger, to primary emotions like fear, to complex emotional states like guilt, affective control underlies most (if not all) interactions of animals. Understanding what exactly the functional roles of the different kinds of affect are in the context of a social group or species and how these functions can be implemented in agent architectures may contribute to theories of (1) how individual behavior (as guided by affective control) can influence group behavior, (2) how changes in individual behaviors can impact groups and societies, and (3) why evolution seemed to have produced so many “simple affective, social creatures” (e.g., such as colony-forming insects), but only relatively few complex, deliberative, social species.

Answers to these questions are hard to find and will involve an understanding of the logical space of possible affective phenomena. In particular, there seem to be three questions that are relevant in the context of our understanding of affective states. The first class concerns affective phenomena as they occur in nature, asking (1) what affect is and what different kinds of affect there are, (2) how and why affective mechanisms came about, and (3) what the functions of affective states are (assuming they have some function) in information processing architectures.

We believe that an answer to any of these question will likely not come forth from independent inquiries, but from the interplay of conceptual analyses, empirical findings and concrete experiments with agent architectures. The proposed research strategy then is to start with a notion of affective state, which is applicable to natural systems, determine/define its function in a particular agent architecture and subsequently try to explore the properties of this state for concrete agents in different environments with the goal of extending the notion to more complex cases. This includes investigating ways in which slight changes in environments can change the trade-offs between design options for the architecture and hence for the functional role of the affective state. Such explorations of “neighborhoods in design and niche space” (e.g., Sloman 2000b) will help us understand what the competitive advantage of a particular change in architecture or mechanism might be in a particular environment, and how the benefits change in slightly different environments.

To be able to study the origins of affective phenomena from an evolutionary perspective and effectively experiment with different kinds of agent architectures, a genuine *artificial life simulation environment* is required, within which different species of agents (with different architectures and possibly different bodies) can *coexist* and *procreate*. Both requirements are crucial; the first, because affective control mechanisms in natural systems did not evolve in isolation, but rather in competing multi-species societies. Hence, to fully appreciate the benefits of affective control, we need to study the trade-offs between different control architectures in competition with one another. A model employing affective mechanisms in the control of a particular isolated agent or a group of agents *with identical architectures* is necessarily silent about the evolutionary advantage of affective control over other ways of controlling and regulating behavior (e.g., by virtue of various kinds of non-affective reactive or deliberative processes) in a multi-species environment. The fact that agents of one kind perform better than agents of another kind if tested independently does not shed any light on their performance in mixed groups.

The second requirement is equally important, because classic genetic algorithms (GAs) assess the fitness of agents based on a static, predetermined fitness function and can hardly (if at all) do justice to the dynamics of the local interactions of agents with their (changing) environments, which in the end determines reproductive success (e.g., see Kauffman 1995). There are several problems with specifying fitness explicitly besides evolutionary plausibility. For one, it is not clear what architectural features to select for if the task at hand is to evaluate the role and potential of affective states in different agent architectures from an evolutionary perspective. Furthermore, as agents and their architectures change over time together with the environment, adaptive fitness changes as well, which would have to be somehow reflected in the fitness function (for a more detailed description of the differences between *exogenous* and *endogenous* fitness and some reasons why endogenous fitness is to be preferred in such a simulation setup, see Menczer and Belew 1996).

Other desiderata include spatial continuity (to eliminate any potential influence of grid structures), temporal sensitivity (to be able to study temporal trade-offs of actions and processing mechanisms), at least two resources that agents need to obtain (to make the decision problem interesting, e.g., Tyrrell 1993, or Spier and McFarland 1998), and Lamarckian mutation mechanisms (to be able to control modifications and extensions of certain components of an architecture).[1]

We have developed the SimWorld Artificial Life Simulation Environment based on the above requirements in order to be able to study the origins and roles of affective phenomena in agent societies with possibly many different kinds of agents. SimWorld consists of a potentially unlimited continuous surface populated with various spatially extended objects such as various kinds of agents, static obstacles of varying size, and food and water sources, which pop

---

[1] Note that this is for methodological reasons only. As long as these mutation operations are feasible using Darwinian mutation, we can justify performing operations directly on the architecture instead of performing them on genetic representations. As an aside, it is always possible to regard architectures as representations of themselves, although it is doubtful that organisms would use such an uncompressed code.

up within a particular area (800 by 800 units for the experiments reported in this paper) and disappear after a pre-determined period of time, if not consumed by agents earlier. Agents are in constant need of food and water as moving consumes energy and water proportional to their speed–even if they do not move, they will still consume a certain amount of both. When the energy/water level of an agent drops below a certain threshold ω, agents “die” and are removed from the simulation. They also die and are removed, if they run into other agents or obstacles.

All agents are equipped with exteroceptive “sonar”, “smell”, and “touch” sensors. Sonar is used to detect obstacles and other agents, smell to detect food and water, and touch to detect impending collisions with agents or obstacles as well as consumable food and water sources. In addition, the touch sensor is connected to a global alarm system, which triggers a reflex beyond the agent’s control to move the agent away from other agents and obstacles. These movements are somewhat erratic and will slightly reorient the agent (thus helping it to get out of “local minima”). Furthermore, agents have two proprioceptive sensors to measure their energy and water levels, respectively, which are also connected to global alarm systems. These alarms will be triggered if the energy/water level drops below a critical threshold and make the agent ignore any other need (e.g., if the water level is critically low, food will be ignored).

On the effector side, they have motors for locomotion (forward and backward), motors for turning (left and right in degrees) and a mechanism for consuming food and water (which can only be active, when the agent is not moving). When agents come to a halt on top of a food or water source, their ingestion mechanism suppresses the motors for locomotion until the item is consumed, which will take a time proportional to the amount of energy/water stored in the food/water source depending the maximum amount of food/water an agent can take in at any given time.

After a certain age α (measured in terms of simulation cycles), agents reach maturity and can procreate asexually. Since the energy for creating offspring is subtracted from the parent, agents will have a variable number of offspring depending on their current energy level (from 0 to 4), which pop up in the vicinity of the agent one at a time. Since a mutation mechanism modifies with a certain probability μ some of the agent’s architectural parameters (e.g., such as connection weights in a neural network), some offspring will start out with the modified parameters instead of being exact copies of the parent. Note that both parameters, α and ω, can be used to specify, whether the simulation is used as an exogenous or as an endogenous fitness model.

While different agents may have different (implicit) short-term goals at any given time (e.g., getting around obstacles, consuming food, reaching a water source faster than another agent, or having offspring), common to all of them are two (implicit) long-term goals: (1) *survival* (to get enough food/water and avoid running into obstacles or other agents), and (2) *procreation* (to live long enough to have offspring).

In the following experiments, we study different kinds of related agents, which all possess the same architectural components (but not all the same links among them). All agents process sensory information and produce behavioral responses using a schema-based approach (Arkin 1989). Let $Ent = \{f, w, o, a\}$ be an index set of the four types of objects *food*, *water*, *obstacle*, and *agent*–all subscript variables will range over this set unless stated otherwise. For each object type in Ent, a force vector F_i is computed, which is the sum, scaled by $1/|v|^2$, of all vectors

v from the agent to the objects of type i within the respective sensory range, where '$|v|$' is the length of vector v. These four *perceptual schemas* are then mapped into motor space by the transformation function $T(x) = g_f \cdot F_f + g_w \cdot F_w + g_o \cdot F_o + g_a \cdot F_a$ for $i \in Ent$, where each g_i is the respective gain value. These gain values are provided by the output layer of a three-layer *interactive activation and competition* (IAC) neural network with four input units in, four hidden units hid, and four output units out (Rumelhart and McClelland, 1986) via individual scaling functions $f_i(x) = x \cdot c_i + b_i$ (where b_i is the *base gain value* and c_i the scaling factor for the activation of out_i).

The input layer is connected (again via similar scaling functions) to the internal water (in_w) and energy level sensors (in_f) as well as the global alarm mechanism (which sends an impulse to in_o or in_a units depending on whether the alarm was triggered by an impending collision with an agent or an obstacle). Note that neural networks employed in other simulations to control the behavior of agents (Menczer and Belew 1996, Seth 2000, et al.) usually compute the mapping from sensors to effectors, while the neural network here is intended to implement the affective system, thus adding another layer on top of the input-output mapping (which is accomplished in a schema-based manner; of course, this mapping, in turn, could have been implemented as neural network as well).

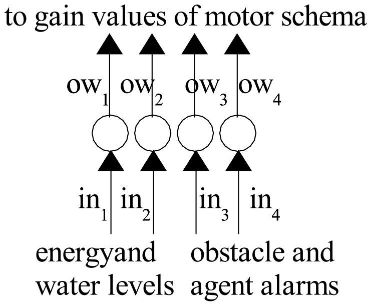

Figure 1: A simplified sketch of the neural network architecture that is used to implement the affective control (only hidden units and their incoming and outgoing connections are depicted).

The activation value $act_i(t)$ of an IAC unit i at time t is defined by

$$act_i(t) = \begin{cases} (max - act_i(t-1)) \cdot net_i(t) - decay, \\ \hfill net_i(t) \geq 0 \\ (act_i(t-1) - min) \cdot net_i(t) - decay, \\ \hfill net_i(t) < 0 \end{cases}$$

where min and max are the minimum and maximum activation level, respectively, $decay$ is a decay factor defined by $d \cdot (act_i(t) - rest)$ (where d is a constant), $rest$ the rest level, and $net_i(t)$ the weighted sum of all inputs to unit i at time t.

Although fully connected IAC networks are possible, we will focus on a subset of networks at this point to avoid complexity, where weights between in_i and hid_i are always non-zero and weights among the hidden units hid_i, call them "matrix weights" denoted by $hw_{i,j}$, and between

hid_i and out_i, call them "gain weights" denoted by ow_i, may be non-zero. In *basic affective agents*, then, each ow_i and $hw_{i,j}$ is zero and as a result the corresponding gain value $g_i = b_i$, i.e., constant. Consequently, the behavior of such agents is completely determined by their inputs: inner states, as possibly implemented by the hidden units, *do not contribute* to their behavior, which is entirely reactive. Basic agents are contrasted with *extended agents*, where some ow_i (and possibly some $hw_{i,j}$ as well) are non-zero and gain values in T can consequently vary depending on the state of the neural network.

As one might expect, the differences in behavior between the various kinds of agents can be very subtle as the influence of the hidden units on the gain values can be very gradual, and hence very difficult to detect. It is therefore crucial to look at a time-frame larger than the life-time of a single agent to be able to evaluate the advantages and disadvantages of different weight values, in particular, in competitive multi-agent environments. In fact, most trade-offs are only visible in simulations of many generations of agents in different combinations under different environmental conditions. Nevertheless, it is possible to sketch a few general behavior tendencies. The basic agents, for example, always behave in the same way given that their gain values are constant: with positive $g_f = g_w$ they behave like the "consume nearest" strategy in environments without obstacles (Spier and McFarland 1998). Negative $g_o = g_a$ values will make them avoid obstacles and other agents. In extended agents (with the same gain values) the degree to which they engage in the respective behaviors will in addition to the sign and strength of the weights depend on the activation of the respective hidden units and hence vary from time to time (e.g., they tend to avoid food, if they are not "hungry").

Results

Various previous experiments have shown that affective control mechanisms can and will evolve in various different environmental settings. For example, in (Scheutz and Sloman 2001) we show (although in a slightly different setup) that agents with primitive motivational states (i.e., "hunger" and "thirst" drives) are likely to evolve from basic agents if the ow_f and ow_w weights are mutated. These *motivational agents* are likely to evolve from basic agents independent of many environmental conditions such as the frequency of appearance of new food and water sources, or the numbers and initial distributions of food and water sources, obstacles and agents. Furthermore, we found that starting with *motivational agents*, different kinds of agents with different combinations of primitive emotional state like "fear" or "anger" will evolve, if the ow_o and ow_a weights are mutated (Scheutz 2001). In both cases, we argued in detail that the evolved mechanisms (i.e., positive or negative gain weights) indeed implement affective processes, based on (1) the functional characterizations of the involved affective processes, (2) the observable behavior of the agents in the environment, and (3) the evolved architectural components (i.e., connection weights).

What we did not answer, are the following two important questions: will emotional agents evolve directly from basic agents (without first evolving motivational agents)? Are mutations of the matrix weights $hw_{i,j}$, which were kept at zero in the previous experiments (largely to reduce complexity), going to result in more adapt, more complex agents? And finally, will emotional agents that can “perceive” emotions of other emotional agents and use these perceptions to influence their behavior have an evolutionary advantage over agents without such perceptual mechanisms? This last question addresses in a new way the role of expressing emotional states as a means of social regulation.

We will answer all three questions in the following for a variety of different environmental and simulation parameters.

Table 1: The result of placing at random 20 basic agents and 40 static obstacles in the environment using a food rate of 0.25 and water rate of 0.16 averaged over 20 runs of 100000 simulation cycles with mutation of either the gain weights or both gain and matrix weights at a mutation factor $\tau = 0.05$.

	gain weights only			matrix and gain weights		
Cat.	μ	σ	Con	μ	σ	Con
alive	9.4	7.88	(7.96,10.84)	13.05	8.22	(11.55,14.55)
age	414.74	11.05	(412.72,416.75)	416.74	9.65	(414.97,418.5)
energy	1366.19	47.98	(1357.43,1374.95)	1389.79	10.74	(1387.83,1391.75)
water	829.96	83.07	(814.79,845.12)	872.34	65.93	(860.31,884.38)
distance	918.58	17.02	(915.47,921.68)	922.71	15.34	(919.91,925.51)
EEF	1.04	0.05	(1.03,1.05)	1.07	0.03	(1.06,1.07)
WEF	0.46	0.08	(0.44,0.47)	0.5	0.06	(0.49,0.51)

Table 2: The result of placing at random 20 basic agents and 40 static obstacles in the environment using a food rate of 0.25 and water rate of 0.16 averaged over 20 runs of 100000 simulation cycles with mutation of either the gain weights or both gain and matrix weights at a mutation factor $\tau = 0.5$.

	gain weights only			matrix and gain weights		
Cat.	μ	σ	Con	μ	σ	Con
alive	10.95	7.86	(9.52,12.38)	6.2	8.9	(4.58,7.82)
age	375.41	21.76	(371.44,379.38)	400.61	12.83	(398.26,402.95)
energy	1381.15	15.65	(1378.29,1384.01)	1368.39	36.57	(1361.71,1375.06)
water	784.93	83.49	(769.69,800.17)	786.5	67.53	(774.17,798.82)
distance	831.97	43.12	(824.1,839.84)	878.7	30.33	(873.17,884.24)
EEF	1.22	0.09	(1.21,1.24)	1.11	0.07	(1.09,1.12)
WEF	0.5	0.07	(0.49,0.51)	0.44	0.06	(0.43,0.45)

We performed four major sets of experiments, where we populated various environments (from 0 static obstacles to 70 static obstacles) with basic agents and allowed *either* for mutation of the gain weights only (“g-mutation”), or for mutation of both matrix and gain weights ("gm-mutation") by a fixed mutation factor of *either* $\tau = 0.05$ or $\tau = 0.5$. Whenever an agent has offspring, the probability μ for modification of any of the two weights is 1/3 (i.e., 1/6 for increase or decrease by τ, respectively). Results of these experiments for environments with 40 static obstacles and a "food frequency" (i.e., frequency of appearance of new food source) of 0.25 are shown in (Table 1) and (Table 2). All statistics are calculated for 20 parallel runs of the simulation for 100000 update cycles each, where 20 basic agents were randomly placed in an environment. For each set of experiments, these tables show mean, standard deviation and confidence intervals at $\alpha = 0.05$ for the average number of agents alive after 100000 update cycles, as well as additional parameters of agents (averaged over averages) intended to measure the effectiveness of the evolved architectures: average age, average energy and water consumption, average distance traveled, and the two efficiency parameters E(energy)EF(ficency) and W(ater)EF(ficency). These two parameters capture the average energy and water consumption per distance unit traveled: the lower the parameter, the better the agent's ability to find and get food or water.

The answer to the first question, then, is "yes": emotional agents will evolve from basic agents, as evolution makes use of the architectural features (i.e., the connection weights) in ways that give rise to affective states, including the primitive motivational and emotional states we found previously. In the case of g-mutation, the most frequently evolved state is typically "hunger" (realized by a positive weight ow_f). Much less so we find positive ow_w weights, indicating that a "thirst" state has evolved. This has partly to do with the fact that water is typically readily available in any of the experiments and not a scarce source like food. Hence, ignoring/going after water is not as critical for survival as ignoring/going after food. Even less uniformity we find with the remaining two gain weights ow_o and ow_a. In about a third of the cases, both weights will be negative and implement "fear-like" states (i.e., "fear of obstacles" and "fear of other agents", see Scheutz 2001 for a discussion about how to assess what affective state is implemented). These are simulations with a typically high number of surviving agents, since avoiding obstacles *and* other agents makes agents spead out better in the environment, thus providing better support for coexisting agents (Seth 2000 calls the degree to which agents are spread out "clumpiness of the environment).

In the remaining cases, we find a mixture of negative and positive weights (which, in turn, implement "anger-like" states). While there are examples of surviving agents for any possible combination of the above states, agents within each simulation are very homogeneous, i.e., have all very similar weights. This indicates that while there is a great variety of different viable agent configurations with respect to affective control, the diversity among agents in each group is limited, which is partly due to the fact that different affective states result in different behaviors that are not mutually compatible (another factor, for example, is the spatial extension of the simulated world). Hence, in the competition of different affective control sets, one set of affective states will eventually prevail and agents that implement the other will become extinct.

The picture is more complex with gm-mutation, because the matrix weights allow for intricate interactions of affective states. While the overall results are very similar, the kinds of affective states implemented by gm-mutation agents are more varied and nuanced, and require careful analysis of the flow of activation in the whole IAC network. Various sorts of self-excitatory and inhibitory connections augmented by complex temporal influences usually complicate the picture, and sometimes it is simply not possible to label the kinds of interactions that can take place (not the least for lack of appropriate terms to describe the particular kind of affect). Space limitations do not allow us to include a detailed analysis here.

Answering the second question, we do not find large differences in performance as measured in term of EEF and WEF between agents with g-mutation and agents with gm-mutation (for either value for τ). However, it is more likely in simulations with higher rates of τ that no gm-mutation agent will be alive after 100000 update cycles than that no g-mutation agent will be alive, as higher rates of τ can lead to greater variations, which increases the search space for appropriate weights that needs to be explored by evolution. For g-mutation this is beneficial as it increase the otherwise low search space. For gm-mutation, however, this is of disadvantage as it already has to deal with a large search space, which is only increased further. In the experiment shown in (Table 1) with low τ, for example, agents with g-mutation were alive in 13 runs, whereas agents with gm-mutation were alive in 17 out of 20; in the experiment reported in (Table 2) with high τ, agents with g-mutation were alive in 15 runs, agents with gm-mutation only in 8.

We ran the above four sets of experiment for different food frequencies (from 0.25 to 0.125) keeping the water frequency constant at 0.16. It turns out that while at at food frequency of 0.2 the picture is still similar, at food frequencies of 0.16 and below, agents typically only survive in low obstacle environments (less than 20 obstacles). This is true for both the g-mutation and the gm-mutation agents. In none of the above simulations did agents survive in environments with 70 obstacle (or more).

Finally, to answer the third question, we first allowed emotional agents to express their emotional states in a way that can be perceived by other agents, and then defined a new kind of emotional agent, called "s-emotional agent" ('s' for "social"), which can use this perception to influence its behavior. Specifically, an s-emotional agent A compares its own fear level to the fear level of another emotional agent B, which A can perceive when B's proximity triggers A's alarm, and either initiates its reflex (as before) if $fear(A) - fear(B) \geq \Delta$, or otherwise suppresses the reflex and remains in its current position. This additional mechanisms can be beneficial if two agents compete for a close-by food source, for example, as knowing whether the other agent will stick around or move away can be an important factor in deciding what to do next (when the agent alarm is triggered).

We ran several experiments with different values for Δ, and typically the s-emotional agents performed as well as emotional agents if placed in homogeneous groups in the same kinds of environments for Δ values up to 0.2.[2] Then, we start to see a degradation in performance (again as measured in terms of the average number of surviving agents), which becomes significantly worse for values of 0.5 and greater. A large Δ value effectively means that the decision of the agent to suppress the alarm is largely independent of the difference in fear levels, hence it is very probable that this decision is not beneficial. For sticking around only pays off in the long run, if the other agent eventually retreats, which it will not do if its fear level is smaller that the one of the agent. Hence, in most encounters, both agents will sit in a tie situation, while their fear levels are increasing, and keep losing valuable energy and water without getting any closer

---

[2] The detailed results of these experiments are not reported here for lack of space.

to food or water. The tie is only broken when either (1) the difference in fear levels exceeds Δ, (2) the source of attraction (i.e., food or water) that attracted the agents in the first place disappears, or (3) at least one of the agents dies.

In sum, perceiving another agent's emotions (as specified above) does not lead to worse performance among s-emotional agents. But is it beneficial in mixed groups? To answer this question, we placed emotional and s-emotional agents in various environments and ran simulations for 10000 update cycles. It turns out that with Δ values up to 0.2 s-emotional agents outperform emotional agents significantly. This advantages vanishes as Δ increases and for values of 0.5 and 0.6 s-emotional and emotional agents show the same performance. The superior performance of s-emotional agents for low Δ, say $\Delta = 0$, is not surprising: whenever they predict that the other agent will not stick around, they are right. For emotional agents will always back up, and s-emotional agents will also back up if their fear level is greater than that of their opponent. This does not mean that s-emotional agents get all of their predictions right. In fact, they will wrongly predict that another emotional agent is going to stick around if its fear level is lower than theirs, as they do not have any means to distinguish emotional from s-emotional agents. This error in prediction, however, is not problematic as emotional agents will also back up (and, hence, not be able to reap the benefits of the s-emotional agent's wrong prediction). Hence, being right 3 out of 4 times on average and being able to stick around as opposed to being forced to back up all the time (as in the case of emotional agents), is beneficial enough to make s-emotional agent superior to emotional agents. Hence, there is an evolutionary advantage to expressing one's emotions, which can then be perceived by others and used to make better decisions. The emotions in the above setting, therefore, have a social, regulatory role that increases fitness of the whole species.

Discussion

The above experiments demonstrate the research strategy suggested earlier, which we believe will help us understand the role and origins of affective states as well as the potential uses affective states can be put to in the control of agents, especially with respect to their potential as social regulators of behavior. Furthermore, the experiments confirm that if there are architectural components that can implement them, affective states like "hunger", "thirst", "fear", and "anger" are likely to evolve in different combinations, even in very competitive multi-agent environments, as they are not only beneficial to the individual agent, but also lead to behavior, which benefits the whole species (for more a more detailed description of the evolved behaviors, see Scheutz 2001). Furthermore, they show that expressing emotional states can serve a regulatory role in decision making, which can benefit not only individual agents, but aforteriori the whole group.

These results are 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.

The investigations reported in this paper are part of an ongoing, long-term project, in which we investigate the logical space of affective states, their potential as coordination mechanisms for individual agents and groups of agents, and their requirements for and possible implementations in agent architectures. Many more experiments using different kinds of agent architectures with different capacities to implement affective control mechanisms are needed to explore the space of possible uses of affective control and the space of affect itself.

References

Arkin, R. C. 1989. Motor schema-based mobile robot navigation. *Intern. Journal of Robotic Research* 8:92–112.

Cañamero, D. 1997. Modeling Motivations and Emotions as a Basis for Intelligent Behavior. In *Proceedings of the First International Symposium on Autonomous Agents*. Marina del Rey, CA: The ACM Press.

Kauffman, S. 1995. *At home in the universe: The search for laws of complexity*. London: Penguin Books.

Menczer, F. and Belew, R. K. 1996. From Complex Environments to Complex Behaviors. *Adaptive Behavior* 4 (3/4):317–363.

Rumelhart, D., and McClelland, J. 1986. *Parallel distributed processing: Explorations in the microstructure of cognition*. Cambridge, Mass: MIT Press.

Scheutz, M. 2001. The Evolution of Simple Affective States in Multi-Agent Environments. In *Proceedings of AAAI Fall Symposium 01*, AAAI Press.

Scheutz, M. and Sloman, A. 2001. Affect and Agent Control: Experiments with Simple Affective States. In Proceedings of IAT 2001, World Scientific Publishers.

Seth, A. 2000. On the Relations between Behaviour, Mechanism, and Environment: Explorations in Artificial Evolution. Ph.D. diss., Univ. of Sussex.

Sloman, A. 2000. Interacting trajectories in design space and niche space. In M. Schoenauer et al. eds. *PPSN VI LNSC* 1917, 3–16. Berlin: Springer.

Spier, E., and McFarland, D. 1998. Possibly optimal decision making under self-sufficiency and autonomy. *Journal of Theoretical Biology* 189:317–331.

Tyrrell, T. 1993. Computational Mechanisms for Action Selection. Ph.D. diss., Univ. of Edinburgh.