

An Action Selection Architecture for an Emotional Agent

Gertjan J. Burghouts, Dirk Heylen, Mannes Poel, Rieks op den Akker and Anton Nijholt

University of Twente, P.O. Box 217, 7500 AE Enschede, The Netherlands
burghout@science.uva.nl, {heylen, mpoel, infrieks, anijholt}@cs.utwente.nl

Abstract

An architecture for action selection is presented linking emotion, cognition and behavior. It defines the information and emotion processes of an agent. The architecture has been implemented and used in a prototype environment.

Introduction

We present an architecture and a formal model for a reconstruction of what we think are essential aspects of the relations between human *emotions*, *cognition* and *behavior*.

We distinguish three types of motivation for looking at emotions with respect to software systems or for developing 'emotional agents'. These types correlate each with different aspects of emotions or with different perspectives on emotions.

1. The Control-engineering motive. The emotional process can be viewed as an efficient, low-cost control mechanism especially useful in resource-constrained competitive environments in which deliberative control is not appropriate. This view focusses on one typical property of some emotional behavior, namely that it leads to an immediate response to environmental events.
2. The Experimental-theoretical motive. Suppose one's primary interest lies in human emotions and human emotional behavior as such. In this context of emotion research a computer system with emotional agents is viewed as an implementation of a model based on a theory of emotions and emotional behavior. The system is built and used as an experimental environment to verify or falsify hypotheses based on the theoretical insights expressed in the emotion theory.
3. The Believable-agent-motive. Embodied conversational agents that show emotions in the way they act or behave in environments where they interact with humans (such as computer games or tutoring environments) are more believable and engaging than similar agents that do not show emotions. One of the underlying ideas here is that the computer should, in particular situations, behave like humans to be convincing to the user.

Notice that none of the three types of motivation necessarily imply or exclude an opinion on the question whether these systems 'really' have emotions.

We are interested in emotional agents primarily because of the third motive. Our main interest is to research and develop intelligent systems for human computer interaction, including embodied conversational agents. For this we need models that relate emotional and cognitive processes with behavior, so that the behavior shown by the software implementing this model is felt to be natural and expressing recognizable emotions and personality traits.

In "The Emotions", Frijda (1988) presents a functional theory of emotions, that forms the main theoretical background for the construction of our model and architecture. It is complemented by insights from many other psychologists.

This view on emotions shows a huge interest in the change in action *readiness* or 'action tendency' that emotions give rise to. It shows how emotions for an individual are often *functional* and motivated by one or more *needs* of the individual. In this paper we are particularly concerned with how emotional and cognitive processes influence the selection of actions. This paper can be seen as a sequel to (Poel *et al.* 2002) in which we presented a system that was concerned with learning to appraise events emotionally. Here we present an architecture for the next stages in the emotion process that lead to the selection of the actions. We present our experimental environment in which these processes are enacted by agents trying to survive in a grid world.

This paper is organized as follows. First we present the architecture for emotional action selection. Next, we look at the implementation of the model in our grid world, followed by some observations about test runs of this system. More details about this architecture can be found in (Burghouts 2002).

An Action Selection Architecture

We have turned to the literature on the emotional process to identify the various components that have their role to play and to establish the kinds of function they serve. In this section we present the result of the integration of numerous sources.

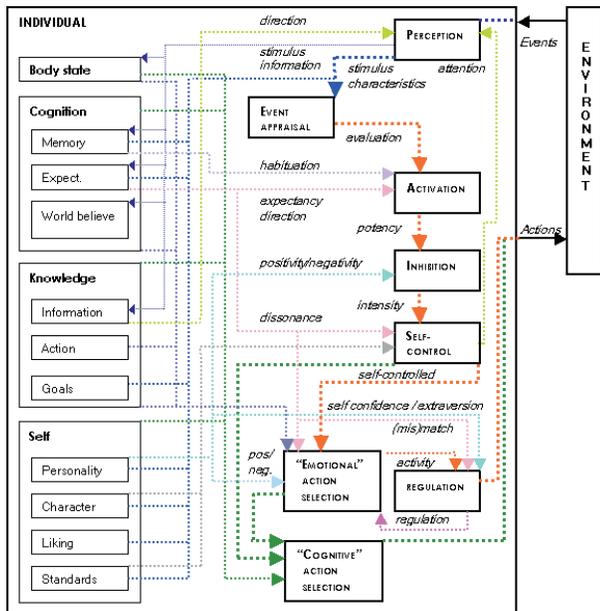


Figure 1: Action selection architecture

Action Selection and Emotion

The emotion system is said to be one of the major factors in the processes that motivate agents to act in a certain way, if not its *raison d'être* (Scherer 1987). Starting from the appraisal of *events* that lead to *motivational action tendencies*, (Frijda 1988), the agent is led by the emotional and other cognitive processes to an actual *response*: a change in attitude or a specific behavior.

Broadbent (1958) organizes the stages in the emotional process from events to responses. First events are detected by the individual. One or more events are selected and processed further. This can lead the agent to change his beliefs about the probabilities of these and related events and thereby to change his or her *expectancies*. Information about the event, the individual's *knowledge* and prior expectancies are all integrated. The information, for instance about the conditions in which the types of events experienced occur, is *stored* in memory. A *response* is selected and executed. Through environmental *feedback* responses may lead to new events from the environment. We specified Broadbent's general description of this process.

Figure 1 shows our architecture that models the *emotion process*, *cognition* and *behavior*. The boxes to the left under the label *Individual* list various characteristics of an individual (agent) that play a role in the emotion process. We have included aspects of an individual's body state, memory, expectations, beliefs about the world, knowledge about actions and their consequences, the different goals an agent is pursuing, its likings, standards, character and personality traits. We view personality as the attitude w.r.t. 'the world' such as optimism, confidence and extraversion (Lisetti & Gmytrasiewicz 2001). We think of character as a set of typical behavioral strategies one applies in certain situations.

The boxes to the right decompose the emotion process.

The Emotion Process

The emotion process is defined by a number of modules that each change the intensity of the emotions in a rule-based fashion. Each module will typically involve different contributing factors, i.e. characteristics of the individual or of the environment. We distinguish the following sequential (Frijda 1988) subprocesses that modify the intensity of the emotions as experienced by an agent: appraisal, activation, inhibition and self-control.

Event Appraisal The agent continually evaluates its environment. Emotions are triggered or elicited depending on certain conditions and decay over time. In (Poel *et al.* 2002) we described our implementation of the appraisal mechanism using neural nets. We based ourselves on the OCC model (Ortony, Clore, & Collins 1988) to associate intensity values with different types of emotions such as *joy/distress*, *hope/fear*, *relief/disappointment*, *pride/shame*, *admiration/reproach*, *happy-for/resentment* and *glloating/pity*. The emotional state of an agent was encoded as a vector of values representing the intensities of emotion types.

$$eventAppraisal :: Event \rightarrow [(Emotion, Intensity)]$$

Activation Event appraisal is the first stage in the emotion process. During this process emotions are elicited. The next stage we call activation. It involves a number cognitive factors such as memory and expectations. For example, the degree of discrepancy between a prior expectancy and a stimulus plays an important role in determining whether an individual assimilates or contrasts his reaction to an event (Frijda 1988). Habituation also influences the activation of emotional behavior: responses to unconditioned stimuli decrease when these stimuli occur repeatedly.

In the activation stage, the emotional state is thus transformed to what we call *potency* values for activated emotions, depending on *habituation*, *expectancy* and *discrepancy*:

$$emotionActivation :: [(Emotion, Intensity)] \rightarrow (ElapsedTime, PreviousIntensity) \rightarrow Expectancy \rightarrow Discrepancy \rightarrow [(Emotion, Potency)]$$

Inhibition Suppression of emotions may occur due to *inhibition*. Inhibition is especially evident in the case of negative emotions: an individual will then try to reduce these emotions by emphasizing positive emotions (Ben-Ze'ev 1997). This reducing factor depends on the individual's personality such as his inclination towards positive or negative thoughts and feelings.

$$emotionsInhibition :: [(Emotion, Potency)] \rightarrow [(Emotion, EmotionImpact)] \rightarrow [(Emotion, Intensity)].$$

To account for inhibition, we derived inference rules based on the relations between potencies and actual activations of emotions as laid down in (Osgood, Suci, & Tannenbaum 1957). When potency and actual activation of two emotions are positively correlated, an emotion does not inhibit another emotion. When they are negatively or not significantly correlated, the emotion with the higher activation/potency-ratio inhibits the other emotion.

Self-control of Emotional Behavior *Self-control* is needed to resolve the conflicts between present and future gratification (Gifford 2000). Bandura (2001) makes the point that self-control of emotions and the potentially corresponding actions includes “self-monitoring and self-guidance via personal standards, and corrective self-reactions”.

Emotions differ with respect to the ease with which they can be controlled. It is, for instance, often claimed that some ‘hot’ or ‘old’ emotions, such as fear, are difficult to control (Frijda 1988).

Self-control is manifested in the reduction or removal of *dissonances* between *expected behavior* (based on the individual’s *character*) and the *standards* one upholds (i.e. the way he judges the praiseworthiness of actions).

$$\begin{aligned} selfcontrolExpectedBehavior :: \\ [(Emotion, Intensity)] \rightarrow \\ Expectancy \rightarrow Strategies \rightarrow \\ [(Action, PraiseworthinessOfAction)] \rightarrow \\ [(Emotion, Intensity)] \end{aligned}$$

Action Selection

Various factors play a role in the selection of an action. For instance, how one perceives the success of behavior depends heavily on one’s confidence (Marakas, Johnson, & Palmer 2000), the level of extraversion (Luminet *et al.* 2000) and the costs in energy and resources associated with the behavior (Gifford 2000). The action selection components are rule-based and determine a strategy to actually select an action.

Emotions do not just influence action selection but also information processes such as perception for instance (e.g. ‘tunnel vision’ when experiencing fear). Emotions also are the key factor in expressions and expressive behavior, but these are not taken into account here.

Previous Work

The work reported on in (Poel *et al.* 2002) was solely concerned with generating emotions by event appraisal: agents living in an environment had an emotional life. However, the emotional experience would not feed back into their behavior depending on the action tendencies that are linked to the emotions. The architecture presented above served to refine the emotional process on the one hand and to let the emotions influence the selection of actions. Next, we consider the implementation of this architecture and the effects on the lives of our agents.

A Prototype Implementation

A prototype environment was set up to see how the components in the architecture could be made operational for certain agents and how this would influence the behavior of these agents living in the environment. We will first describe the environment and then the effects of the architecture’s support for emotional action selection.

The Emotional Agents and their Environment

The prototype is a predator/prey environment which consists of a grid world inhabited by four emotional agents, the preys, and four predators. The world contains trees, pools, and rocks. Apples grow on the trees for food. The pools contain water (though they may also dry out) and herbs grow on rock (herbs provide medicine, i.e. the health of agents eating herbs improves). The pools, trees and stones produce their resources in a nondeterministic fashion. In time, the amount of food and water of the agents decreases and each time one of our agents is attacked by a predator, the health decreases. When resources become below the critical level, an agent dies. Agents walk around in the environment trying to survive. They wander around, looking for resources, fighting predators, teaming up together in social groups to withstand predators better or to distribute resources among the group. The leader of a group is selected via a social ordering between the agents.

Implementation of the Emotion Process Module

The architecture presented in figure 1 was derived from observations, experiments and theories scattered throughout the psychological literature. In the process of implementing the architecture we had to decide on more specific values for the various parameters that enter into the model. For instance, observations from the literature such as “very negatively evaluated disappointment results in moderate activation” contains vague terms such as “very negative” that have to be made precise. Rules are also often incomplete. A statement such as “habituation involves memory to determine a habituation signal” leads to the question how elapsed time and the occurrence of stimuli are related exactly and what the impact could be of a certain habituation signal on intensity?

To convert the vague linguistic descriptions into more precise numerical values we have used fuzzy rules.

The Emotion Process Module

The grid world was implemented to be able to observe the various effects that result from changing parameters of the architecture. Emotional behavior should differ according to the *personalities* of the agents. In our world we introduced two agents, *Hero* and *Grumph*, with different personalities. This is reflected, for instance, in the way they appraise events. Hero is fairly optimistic which means that he will be try to keep his hope up. He is more confident and will for instance be more inclined to attack predators. He is also self-centered, extraverted and idealistic. Grumph is inclined to blame others for his misfortunes. He is easily disappointed, reproaches others more often and is introverted.

We wanted to see what effects these differences had for various components of the emotion process. For this we ran our world a number of times. Each run starts with an initial setting of the agents and the world. It ends when both the agents have died. We recorded for each run what happens to the different values of the emotion modules for the agents Hero and Grumph and how frequently some output was generated. We have no test results of what happens when a component is left out of the architecture.

From the experiments it followed that the main effect of the difference between the two personalities was, not surprisingly, noticeable in the activation of the emotions. Grumph activated more negative emotions and Hero activated all emotions more gradually, while both activated very common emotions such as distress and fear in a similar way.

Self-control is an important process in determining the action selection given activated emotions. Emotions that are not self-controlled tell us most about how the action selection manifests itself. The more gradually activated and very basic emotions of Hero sometimes resulted in emotional action selection, as some of these were not inhibited or self-controlled. The dominant and old emotions fear and distress were difficult to self-control, and therefore these emotions influenced the action selection often. For Grumph however, also other negative emotions (such as shame, reproach, and resentment) influenced the action selection, even though they were self-controlled occasionally. This influence on the selection of actions was more or less as we would expect from the definitions of the personalities Hero and Grumph.

On another level, we can observe the differences in the kind of behaviors that are displayed by the two personalities. For instance, by counting the number of times they either attack or flee from a predator we can measure *aggression*. Counting the number of times they either take or refuse resources we can measure *altruism* or *generosity* and *cooperativity* can be evaluated by looking at the times the characters either join or leave groups.

The different figures below (figures 2, 3, 4) show the values for aggression, generosity and cooperativity of behaviors during the agents' lives. The figures are based on a case study. The horizontal axis in the figures depicts the agents' age t . Above the age t the percentages of decisions that the agent makes concerning these behaviors until t is given. For instance, 40 for attacking predators (i.e. aggression) reflects that 40% of the agent's decisions taken until the corresponding time t involve the attacking of predators. We have recorded the different behaviors of the two agents *Hero* and *Grumph* for two situations: in which Hero is the *leader* of a group and Grumph is just a *member* and vice versa. The red lines are used for the leader, the gray lines for the member.

With Hero as a leader, Grumph is less likely to attack a predator if Hero already attacks the predator. Since Hero is more likely to attack a predator than Grumph, Grumph only attacks predators when Hero earlier attacked a predator (see figure 2 (a), 1).

When Grumph is a leader of the group, Hero and Grumph attack predators together most of the time. However, Hero is more inclined to attack (see figure 2 (a), 1). In addition, Hero seems more persistent in attacking predators (2).

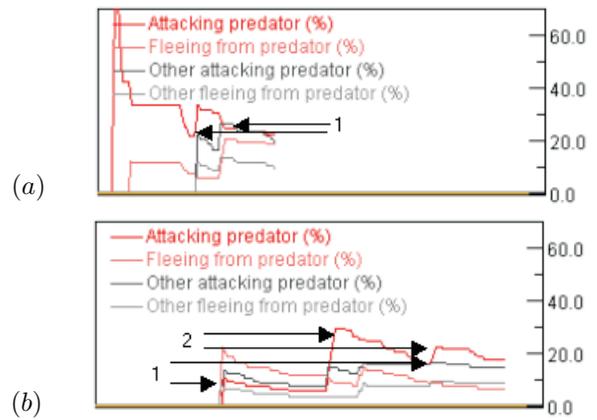


Figure 2: Aggression of Hero and Grumph. (In figure (a) Hero is the leader and Grumph a member of the group. In figure (b) the roles are interchanged.)

When Hero is leader of a group, the first resources are taken by Hero (see figure 3 (a), 1). Later on, Grumph also gets resources. Eventually, the resources seem to be equally divided.

With Grumph as a leader, Hero and Grumph get resources equally from the moment they formed a group. See figure 3 (b). During a lifetime, Grumph joins the group of Hero nor-

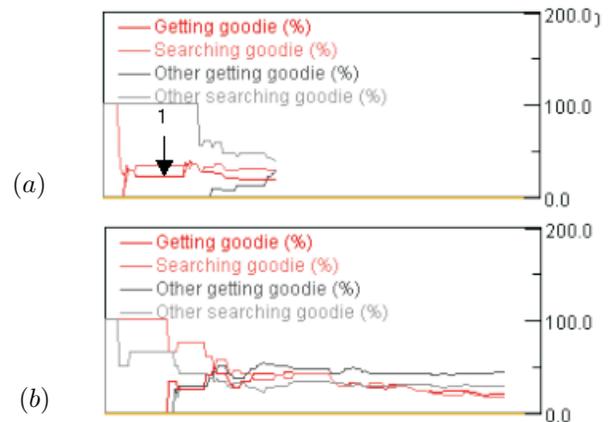


Figure 3: Generosity of Hero and Grumph. (In figure (a) Hero is the leader and Grumph a member of the group. In figure (b) the roles are interchanged.)

mally once. Only if Hero's health gets really worse, and Hero therefore lets the group search for a resource Grumph does not want, Grumph leaves the group. He then immediately joins another group (see figure 4 (a), 1).

Hero joins a group one or two times in a lifetime (see figure 4 (b), 1), which is also the case when Grumph is the leader. Hero seems to leave and join the same group again sometimes (2).

All in all: Hero attacks more predators than Grumph, and flees less; Grumph is more generous than Hero; Hero

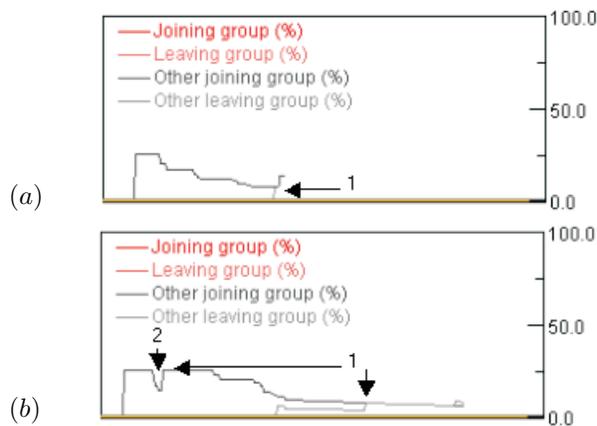


Figure 4: Cooperation of Hero and Grumph. (In figure (a) Hero is the leader and Grumph a member of the group. In figure (b) the roles are interchanged.)

is slightly more flexible in joining and leaving groups than Grumph.

Fitness and Emotional Behavior

What are the consequences of allowing emotions co-determine action selection? One way to say something about this is to look at the *fitness* of an agent measured in terms of hunger, thirst and health.

Hero's impulsive behavior concerning resources does not improve his fitness sufficiently to compensate for the intense attacking of predators. However, Hero's impulsive resource behavior does not seem to affect the group's fitness significantly, as the group's fitness only gets slightly worse. While at the same time Hero requests more resources, the group finds more resources as a consequence of Hero's impulsive behavior. However, one should take into consideration that Hero takes more resources and may 'die young', leaving a group with few resources, while there still may be predators around.

Also Grumph's fitness is only slightly improved, while he distributes resources equally among his group members and occasionally attempts persuasively to attack a predator.

To sum up, it seems that Grumph's more resolute behavior is beneficial for the group, because more resources are found and these are equally distributed.

Discussion and Future Work

Can we consider our agents' behavior to be *believable*? All in all, the actions of the agent do not allow for real complex emotional behavior. Our main achievement is that we can consider the architecture's generated shallow emotional behavior to be *believable* from observation: the apparent differences between Hero's and Grumph's behavior relate to our intuition of an idealistic, extraverted, self-confident but self-centered personality, and an introvert and negative personality, respectively. In addition, we found arguments from Chou's 'anneagram' of human personalities (Chou 2001)

that the behavior of Hero and Grumph generalizes for humans in non-specific terms: our agents' behaviors compare well to the basic behavior characterization of the 'aggressive power-seeker' and 'withdrawn power-seeker', respectively.

A next step would be to do more experiments concerning founding and testing the believability of our agents, taking into account the work of Ortony (2003) and Picard (2003).

Further research could also focus on including emotions in *communication* and *relations* between individuals. In addition, the architecture should be endowed with *social learning* capabilities that involve emotion (e.g. conformation).

References

- Bandura, A. 2001. Social cognitive theory: an agentic perspective. *An Reviews of Psychology* 52:1–26.
- Ben-Ze'ev, A. 1997. The affective realm. *New Ideas in Psychology* 15:247–259.
- Broadbent, D. 1958. *Perception and communication*. Oxford: Pergamon Press.
- Burghouts, G. 2002. Developing an action selection mechanism for emotional agents. Master's thesis, University of Twente.
- Chou, T. 2001. A directional theory of the enneagram or Karen Horney^{squared}. *Enneagram Monthly* 7.
- Frijda, N. 1988. *The emotions*. Cambridge: University Press.
- Gifford, A. 2000. Emotion and self-control. *J of Economic Behavior and Organization* 49:113–130.
- Lisetti, C., and Gmytrasiewicz, P. 2001. Emotions and personality in agent design and modeling.
- Luminet, O.; Zech, E.; Rimé, B.; and Wagner, H. 2000. Predicting cognitive and social consequences of emotional episodes: The contribution of emotional intensity, the 5 factor model, and alexithymia. *J of Research in Personality* 34:471–497.
- Marakas, G.; Johnson, R.; and Palmer, J. 2000. A theoretical model of differential social attributions toward computing technology: when the metaphor becomes the model. *Int J of Human-Computer Studies* 52:719–750.
- Ortony, A.; Clore, G.; and Collins, A. 1988. *The cognitive structure of emotions*. Cambridge: University Press.
- Ortony, A. 2003. On making believable emotional agents believable. In Trapp, R.; Petta, P.; and Payr, S., eds., *Emotions in Humans and Artifacts*. MIT Press.
- Osgood, C.; Suci, G.; and Tannenbaum, P. 1957. *The Measurement of Meaning*. University of Illinois Press.
- Picard, R. 2003. What does it mean for a computer to "have" emotions? In Trapp, R.; Petta, P.; and Payr, S., eds., *Emotions in Humans and Artifacts*. MIT Press.
- Poel, M.; op den Akker, R.; Nijholt, A.; and van Kesteren, A. 2002. Learning emotions in virtual environments. In Trapp, R., ed., *Proc of the 16th European Meeting on Cybern and System Res*, volume 2, 751–756.
- Scherer, K. 1987. Toward a dynamic theory of emotion: the component process model of affective states.