We are IntechOpen, the world's leading publisher of Open Access books Built by scientists, for scientists

6,300

Open access books available

172,000

International authors and editors

190M

Downloads

Our authors are among the

154

Countries delivered to

TOP 1%

most cited scientists

12.2%

Contributors from top 500 universities

WEB OF SCIENCE™

Selection of our books indexed in the Book Citation Index
in Web of Science™ Core Collection (BKCI)

Interested in publishing with us?
Contact book.department@intechopen.com

Numbers displayed above are based on latest data collected.
For more information visit www.intechopen.com

SIMPLEX – Simulation of Personal Emotion Experience

Henrik Kessler, Alexander Festini, Harald C. Traue,
Suzanne Filipic, Michael Weber and Holger Hoffmann
*University of Ulm, Medical Psychology, Institute for Media Computing
Germany*

1. Introduction

Emotion modelling is becoming increasingly important for emotion computing, computer games, interactive storytelling or life-like wizards and assistants because it is necessary to make human-computer interaction more natural. Reeves & Nass (1996) showed that humans like to communicate with computers as they do with people. Software applications which include models of emotional processes are needed to model the social and emotional aspects of human-machine interaction. Extending classic AI and logic by adding simulated emotions can be useful to improve the user's experience in many ways. This chapter will provide a brief overview of existing solutions and models used for artificial emotions (AE) and present a novel model of emotion simulation (SIMPLEX). Empirical data will be reported on its performance, especially the occurrence of emotions, in a game environment. This chapter concludes with a comment on the usefulness of separating AI and AE considering recent advances in cognitive neuroscience.

2. Models for artificial emotions

2.1 Historical roots

The 70s saw what might have been the first debate about emotions and artificial intelligence. The main and – as we know now – most important point was that purely cognitive systems lacked emotions, which strongly influence human thought processes. Two of the models that emerged at that time will be described here.

Simon's interrupt system

Herbert Simon was the first to propose that emotions should be part of a model of cognitive processes (Simon, 1967). His intention was to provide a theoretical foundation for a system incorporating emotions and multiple goals. Within this system, important processes could be interrupted so that more attention went into satisfying important needs (e.g. hunger, safety). Herbert Simon imagined two parallel systems, one designed to achieve goals (cognition, planning) and one observing the environment for events that require immediate attention (emotions). Indeed, the possibility of interrupting current cognitive processes is vital for survival, as it makes it possible to react to threats, but also to pay more attention to one's surroundings when a threat is expected.

Toda's Fungus Eater

Another step towards a theory for the computer modelling of emotions was made by the psychologist Masanao Toda (Toda, 1982) between 1961 and 1980, with a model called the Fungus Eater. This model resulted in the design of an autonomous robot system and partial implementations.

At first, Toda only wanted to create a scenario for a cognitive system that would require concentrating on multiple issues at the same time. In this scenario, the task was collecting as much ore as possible with the help of a mining robot. Operating this robot required energy that could only be gained by collecting a special fungus. Additionally, different Fungus Eaters were competing for the same resources, thus making the scenario more complicated. Toda came to the conclusion that in order to survive on their own, these Fungus Eaters would need to have emotions and to be partially controlled by them. However, Toda named them “urges” instead of emotions and on closer examination, it is apparent that some of these are actual emotions like joy or anger, while others are needs, goals or motives (e.g. hunger).

2.2 Theoretical approach and recent models

There are roughly three areas where emotion models are applied. Artificial emotions (AE) can be used to improve problem-solving in complex environments, as in the early approaches mentioned above. Emotion models can also be used to test psychological emotion theories in experiments using controlled scenarios. Finally, emotions are essential to make computer characters more believable. Emotion models which synthesize and express emotions are necessary to make AI characters more human-like. These models will be the focus of the next sections as they have inspired our own emotional model. The most influential theoretical approach, OCC, will be presented in detail, as it is the basis of many computational models of emotion. Then, three interesting recent models are briefly described.

OCC - a theoretical approach to simulate emotions

The OCC model by Ortony, Clore and Collins is an emotion theory based on appraisal which was explicitly developed to offer a foundation for artificial emotion systems (Ortony, Clore, & Collins, 1988). Its authors succeeded as it inspired many modern models and approaches to artificial emotions.

The basis of the model is that emotions are reactions to the attributes of objects, to events or to actions. Note that internal events (like bodily sensations or memories) which are a part of most modern emotion theories are neglected in the OCC approach. Objects, events and actions are evaluated in an appraisal process based on specific criteria, and result in multiple emotions of different intensities. Figure 1 gives an overview of the OCC approach.

Appraising the *aspects of objects* requires the agent to have *attitudes* (tastes or preferences) in order to decide whether the object is *appealing* or not. This appraisal process results in either love or hate.

Events, or rather *consequences of events*, are appraised by analyzing their impact on the agent's *goals*. This determines the *desirability* of events. The degree of desirability depends on how much closer to or further away from achieving the goal the agent will be after the event. The emotions of joy and distress are direct results of desirable and undesirable events, considering the consequences they have for the agent himself. Some emotions, like for example pity, are triggered when processing events that have consequences for other agents. An open issue is whether this appraisal should be based upon the agent's own goals

or rather the other agent’s goals. How much should an agent be empathic if another one looses something that is not important to the first agent? In an attempt to solve this issue, abstract goals were introduced (such as for example, not losing property). It eventually became clear that it is very important to keep the goals general and abstract, to avoid having to define too many specific goals. The emotions triggered by reacting to other agents’ good or bad fortune depend on how well-liked they are. Another agent’s bad fortune can trigger pity or gloating, while happy events can result in either feeling of happiness or of resentment, depending on the relationship between the agents.

Fig. 1. The OCC model

Appraising an event also means evaluating its *prospects* - hoping or fearing that something will or will not occur. Prospect-based emotions include disappointment and relief. The intensity of these emotions is usually based on the intensity of the preceding hope or fear. The criterion used to appraise the *actions of agents* is their *praiseworthiness*, which is based on the agent’s *standards*. Generally, praiseworthy actions cause pride and blameworthy actions cause shame, if the agent himself is the one acting. When the actions of other agents are appraised, the emotions triggered are admiration or reproach. Standards can be as complex as attitudes (aspects of objects) and goals (consequences of events), and are almost as subjective and individual. Again, the problem of listing them was solved by describing actions in an abstract way. An interesting phenomenon is the ability of feeling proud or ashamed of someone else's actions. Simply put, the closer an agent feels related to the acting agent(s), the more he will identify with him in appraising his actions. Examples of this phenomenon (called the strength of the cognitive unit) can range from parents being proud of their child to soccer fans being ashamed of their team’s performance.

One of the many practical implementations of OCC is the model by Staller & Petta (1999). They constructed a virtual agent which emotion architecture links discrete emotions categories to 14 action response categories, comprising a large range of individual actions. The OCC emotion model is also partly congruent with Nico Frijda's renowned theory of emotions (Frijda, 1986). For more details on emotion theory, see Traue & Kessler (2003).

Artificial Emotion Engine

The aim of the Emotion Engine (EE) is to control the behavior of an artificial agent in complex scenarios. It is made of three layers- emotions, mood and personality (Wilson, 2000). If an emotion is triggered, the actions will be based on this emotion. When emotions are not triggered, the engine bases its actions on the current mood; when no mood is activated, then personality serves as a basis for behavior. The emotion engine is based on the EFA model, which is a three-dimensional space, describing personality traits in terms of Extroversion, Fear and Aggression. Within this space, an area around the point representing an artificial agent's personality is determined and all traits located inside this area are considered to be available to the specific character. For Wilson, the EFA is congruent with the three central systems of the human brain which according to Gray (Gray & McNaughton, 1996) determine behavior: the Approach system, the Behavior Inhibition system and the Fight/Flight system. These three basic dimensions are intuitive, which makes programming easy.

Different personalities trigger some moods more frequently than others: extroversion is linked to good moods, and fear to negative moods. Aggression affects the speed of mood changes. Reward and punishment signals work as the main inputs, and this is comparable with the desirability of events in OCC. Inputs are adjusted based on personality, but also on how often this input occurred before. An agent can get used to a certain input, and this lowers the impact it will eventually have (habituation). On the contrary, a rare or unprecedented input will have more effect (novelty).

Needs are organized hierarchically. Physiological needs, such as hunger, thirst, and the need for warmth and energy are the most important. Each of these needs can become a priority, as when for example a very hungry agent will consider eating as his most important goal. Safety, affiliation and esteem needs are the remaining layers. While physiological needs are the most important, the order of the other layers can vary, depending on what is more important to the agent. Memory is very limited; an agent only remembers how much he likes the other agents. In the same way, in OCC, sympathy is used to cause different emotions for liked and disliked entities.

Only the six basic emotions of fear, anger, joy, sadness, disgust and surprise can be triggered. This might appear like a limited selection compared to the 24 emotions of OCC, but given the reactive nature of emotions in this model (working without inner events and triggers) and since some emotion theorists consider the broad spectrum of emotions as mixtures of these basic emotions, this is quite a sensible choice. Personality is used to adjust the intensity or the frequency of the occurrence of emotions, so that a character with personality that is "low in Fear" will simply not experience as much fear as others.

FLAME

The Fuzzy Logic Adaptive Model of Emotion (FLAME) is partially based on OCC, but what differentiates FLAME from other models is the use of fuzzy logic. This results in a relatively simple appraisal process.

FLAME can integrate multiple emotions at the same time (in a process called emotional filtering), as emotions at times inhibit one another. For example, imagine an agent feeling

joy and pride because he just obtained a new position, but who at the same time feels anger, because a relative of the boss of the company was given a higher position than himself. At this point, his anger may prevent him from feeling joy any longer. When opposite emotions occur, FLAME lets the stronger emotion inhibit the weaker one(s), giving a slightly stronger weight to negative emotions. Another way to handle conflicting emotions is through mood, which is determined by comparing the intensities of positive and negative emotions over the last few steps. If the summed up intensities of positive emotions are higher than that of the negative emotions, then the mood will be positive. If a positive and a negative emotion of comparable intensities occur at the same time, the mood determines which of these emotions will inhibit the other one.

As there is little research about the decay of emotions, FLAME uses a simple constant decay, though positive emotions decay faster than negative emotions. FLAME does not make it possible to implement an agent's personality; instead, differences in behavior are created through learning. For example, an agent may learn that reacting in an angry way will enable him to reach his goals, thus enticing him to be more choleric. FLAME implements multiple types of learning, such as classical conditioning (associating expectations with objects) which occurs in many situations, triggering fear or hope. Another type of learning is learning about consequences of actions or events. This is simple whenever an action directly causes a result. For example, learning that eating will result in feeling less hungry is rather trivial. In the case of more complex causal relations over time, FLAME is using Q-learning, a form of reinforcement learning.

Another form of learning, quite similar to model learning, is the ability to recognize patterns in the behavior of a user by observing sequences of actions. For this type of learning, FLAME simply counts the occurrences of sequences. The last type of learning in FLAME, but one of the most important, is learning about the value of actions. Remember that OCC relies on the praiseworthiness of actions, which is based on the agent's standards. In FLAME, these standards are not predefined knowledge, but they are learned from the interaction between users. Using learning instead of predefined knowledge seems like a very sensible way to avoid most of the troubling issues that come with using OCC. Additionally, learning allows agents to adjust, which makes them all the more believable.

ALMA

The intention in designing A Layered Model of Affect (ALMA) was to control agents in conversational scenarios. In interactive game or learning environments, the artificial characters display facial expressions of emotions and moods through their postures to appear more believable. Emotions, moods and personalities are implemented and interact with each other. Events and actions are described in terms of abstract tags which are then evaluated during the appraisal process and describe things like for example the expressed emotion or gesture accompanying an action or simply if something is a good or bad event. As ALMA is aimed at conversations, an action is often a statement. Hence, there are tags to describe the kind of statement, for example if it was an insult or a compliment. In addition, ALMA requires defining personality profiles for each agent. Essentially, these profiles already contain the desirability and praiseworthiness the agent assigns to certain tags.

Since our own emotion model shares some features with ALMA (see below) a key difference should be pointed out. In SIMPLEX we considered it impractical to explicitly specify this information, as this would have limited the model to a small number of agents. So instead of using tags, our model requires to specify goals and their priorities for an agent, where

generic goals can be used for all agents. Events still need to be described in a special way, but this is reduced to a relatively objective list of which agents goals are affected and in which way. All other information like praiseworthiness is automatically derived from this and the agent’s personality. Although this approach is providing less control over an agent’s appraisal process, it is better suited for a generic system meant to be used with minimal extra effort.

3. SIMPLEX – Simulation of Personal Emotion Experience

3.1 Overview

SIMPLEX is a context-independent module to create emotions as a result of primary application (environment) events. Goals, emotions, mood-states, personality, memory and relationships between agents have been modelled so they could interact as in real life. Figure 2 shows an overview of the model.

SIMPLEX is based on the OCC model by Ortony, Clore and Collins (1988) in that it creates discrete emotions by appraising events based on the desirability of their consequences and the praiseworthiness of the actions of agents. The appraisal process was modified by including the personality of virtual agents. The personality component is based on the Five Factor Model (FFM) introduced by psychologists McCrae & Costa (1987), which includes extroversion, conscientiousness, agreeableness, neuroticism and openness. The personality module influences the emotion module on multiple levels during appraisal processes and in the development of mood-states.

Fig. 2. The emotion module SIMPLEX

Other important aspects of the model are mood-states and relationships. Mood-states are represented in a three-dimensional space which dimensions are pleasure, arousal and dominance (Bradley & Lang, 1994), and they are based on active or recently experienced emotions (implemented by pull-functions). In the absence of emotions, a mood state will

slowly gravitate back to a default mood-state based on the agent's personality. A mood-state also functions as a threshold to determine whether an emotion is strong enough to become active at a given time.

Relationships are handled as if they were mood-states towards other agents (for instance a player in a game scenario): they are based on emotions caused by other agents and they can be considered as a simplified way to store memories of experiences with these agents. They are used as thresholds as well; for example, an agent will be more likely to become angry at another agent when their relationship is in the range of negative valence.

Personality (long-term), mood-state (mid-term) and emotions (short-term) thus represent three levels of the emotion module that interact with each other in order to create believable agents. Events from the scenario serve as the model's inputs. They are appraised according to the OCC algorithm (see figure 1). This appraisal is influenced by the agent's goals, his personality and his relationships with other agents. At the end of an appraisal one or several discrete emotions are generated. These emotions and the current mood-state are represented in the same three-dimensional PAD space: on the one hand, the emotion(s) serve(s) as an attractor for the recent mood-state position (pull function). On the other hand, the closer an emotion is located to the current mood-state, the more probable it will be that the emotion will be activated. The speed at which the mood-state changes, is influenced by the agent's neuroticism (a personality variable). Additionally, emotions that are caused by other agents will influence another mood-state representation (stored on another PAD space) representing the relationship with that agent. Thus, every agent has specific relationships with other agents, which influences his behavior towards others. Emotions, mood-states and relationships with other agents are the outputs of the model and can be used by the AI application.

Originally, the PAD space was designed to represent emotions in a dimensional rather than a discrete way (Russell, 1978). In our model, PAD is used as a common space where three different constructs (discrete OCC emotions, continuous mood-states and personality), are represented in order to be handled together by the SIMPLEX algorithm. An agent's current mood-state is thus the result of a mathematical function which takes into account the default mood (defined by personality), the pulling behaviour of OCC emotion(s) triggered by appraisals, and weighed factors influencing movement speed (see equation 1).

$$\text{Mood-state} = f(\text{PAD}_{\text{FFM}}, \text{PAD}_{\text{Emotions}}, \text{Filter}_{\text{FFM}}) \quad (1)$$

3.2 Basic components

Mood-state represented in the PAD-Space (Pleasure-Arousal-Dominance)

Beyond discrete emotions, which are typically short-term, mood-states are a powerful way to model emotional shifts and explain affective influences over longer periods of time. To implement mood-states in our model, we chose to use Russell's three-dimensional space to describe emotions (Russell, 1978) and Mehrabian's concept of how emotions are linked to personality traits (Mehrabian, 1996).

The dimension of Pleasure encompasses valence ranging from very positive to very negative. Arousal is an indicator of how intensely something is perceived, or of how much it affects the organism. Dominance is a measure of experienced control over the situation. For example, a different degree of dominance can make the difference between fear and anger. Both of these emotions are states of negative valence and high arousal, but not feeling in

control is what differentiates fear from anger. When an agent is angry, it is because he believes he can have a potential influence.

Although emotions are triggered by OCC appraisals and are therefore discrete, they are handled in a continuous three-dimensional space by SIMPLEX. The advantage of treating emotions in this way and not just as a fixed set of possible emotions is that it makes it possible to represent emotions that do not even have a name. It also creates the possibility to combine emotions, mood-state and personality in one space. First, a coordinate in PAD space can obviously represent an agent's mood-state. But emotions and personalities can also be described in terms of Pleasure, Arousal and Dominance values. For example, the value of arousal can be not only the degree of arousal associated with a specific emotion, but also the arousability of a person.

Mehrabian (1996) gives specific names to the resulting different octants in PAD-space and describes the diagonally opposite octants as Exuberant/Bored, Dependent/Disdainful, Relaxed/Anxious, Docile/Hostile. Thus mood-states are not points but octants of the PAD-space. However, positioning a personality (based on FFM) within a PAD-space could have been a rather difficult task, since there is no mathematically-correct way to make the conversion. Luckily, this transformation can be based upon empirical data. Mehrabian provided such a conversion table from FFM to PAD after correlational analyses of questionnaires measuring both constructs in healthy subjects (Mehrabian, 1996).

Five Factor Model of Personality (FFM)

The implementation of personality is a key factor when creating believable agents that differ from each other. OCC already offers a few possibilities: different goals, standards and attitudes automatically result in differences during the appraisal process. However, since personality goes beyond preferences, it was necessary to find a model of personality that made it possible to adjust the appraisal process, to shift the agent's perception and to influence mood-states.

The model chosen for SIMPLEX was the Five Factor Model (McCrae & Costa, 1987). After years of research, an agreement emerged that five groups of traits are sufficient to describe a personality. Using self-report questionnaires with multiple items, a personality profile can be provided for each individual scoring high or low in each of the five factors (this approach is called "dimensional"). In the case of our model, the value for each factor can be typed in when defining the artificial agent.

Agreeableness refers to a tendency to cooperate and to compromise, in order to interact with others in an agreeable way. High agreeableness often means having a positive outlook on human nature, assuming people to be good rather than bad. Low agreeableness is essentially selfishness, putting your own needs above the needs of others and not caring about the consequences your actions might have for others.

Conscientiousness is usually high in people who plan a lot, who think everything through, and who are very tidy or achievers. Extreme cases can appear to be compulsive or pedantic. The opposite personality trait includes sloppiness or ignoring one's duties.

Extroversion can be a measure of how much people experience positive emotions. An enthusiastic and active person that enjoys company and attention is extraverted, while a quiet individual who needs to spend more time alone is introverted.

Neuroticism is partly an opposite of Extroversion in being a tendency to experience negative emotions. However, being neurotic also means being more sensitive in general, and reacting emotionally to unimportant events that wouldn't usually trigger a response. Neurotics can

be prone to mood swings and tend to be more negative in their interpretation of situations. Low neuroticism means high emotional stability and describes calm people who are not easily upset.

Finally, those scoring high on *Openness to Experience* are creative and curious individuals, interested in art and more in touch with their own emotions than others. Those scoring low on that dimension are conservative persons with few interests, they prefer straight and simple things rather than fancy ones, and they do not care about art or science. It is suspected that Openness can be influenced by education.

3.3 Technical implementation

The appraisal process and the generation of emotions

There are three categories of inputs to the appraisal process of the emotion model: consequences of events, actions of agents and objects (see the OCC model in figure 1). The following section will describe the respective mechanisms applied when mapping each type of input to emotions.

Each event handled by a character is first adjusted according to the agent's personality.

First, the consequences are adjusted based on the agent's *neuroticism*. As neurotic people tend to see things more negatively, consequences are rated worse than what they actually are. The factor by which neuroticism can reduce the desirability of events is adjustable. Note that all personality traits are in the range of $[-1; 1]$, so that negative neuroticism actually makes consequences more positive. In real life, positive people could think "it could have been worse".

The desirability of events is determined by (predefined) goals during the event appraisal. A goal consists of two aspects: relevance $[0; 1]$ and state of realization $[0; 1]$, which means to which percentage the goal is already achieved.

Afterwards, the *praiseworthiness* of actions is determined. Basically, the more positive consequences an action has, the more praiseworthy it is considered to be. Sympathy plays a role in this process, as it is added to positive values and subtracted from negative ones. Consequences for self are considered to be more important than consequences for others, which are currently factored in at 50% of their value.

After the adjusted values for all consequences have been summed up, *conscientiousness* is used to obtain the final result, by being scaled and subtracted. Thus the more conscientious an agent is, the harder it will be to commit an action positive enough to be deemed praiseworthy. This applies to both actions of other agents and actions of the agent himself.

Agreeableness works the opposite way, but only for the actions of others. This is based on the psychological notion that agreeable people tend to be more forgiving in order to get along with others. Apart from having a different weight, factoring in agreeableness has the same results as negative conscientiousness.

The remaining factors serving as parameters for the action (*responsibility*, *unexpectedness*, *publicness*) are averaged and used to scale the result of the above calculations. Finally, as cost is attempted to be derived from consequences for self, it is subtracted, before the calculated praiseworthiness is averaged over the number of consequences or rather the number of affected agents. The resulting value of praiseworthiness is used as the intensity for admiration or reproach, depending on whether it is positive or negative. If the agent is appraising his own actions, the emotions are pride or shame instead of admiration and reproach.

Once the praiseworthiness has been calculated, a search is conducted through the list of prospects for all the ones that are active and that match the name of the event. For each, the *prospect appraisal* function is called, which determines the net desirability by multiplying it with the affected goal's relevance. This value will be compared to the expected desirability for this event. The simplest situation is when a positive consequence was expected but a negative one occurs. This would obviously cause disappointment. However, this is also the case if a very high desirability was hoped for and the actual consequences are less positive, but still not negative. Having a hope fulfilled results in satisfaction. If an event has exactly the expected consequences, it results in the full intensity for the emotion.

The intensity of emotions is the product of the determined *quality* of the event and of the intensity of the prospects. For example, if there was very little hope, there cannot be strong satisfaction. Which emotion is created depends on the kind of prospect and on the sign of the quality value. Hope and positive quality result in satisfaction, hope and negative quality in disappointment, fear and positive quality in fears-confirmed and fear and negative quality in relief. After the prospect appraisal is done, short term or one-shot prospects (only valid for one round) are removed.

Appraisal concerning joy and distress is done for each consequence affecting the agent himself, while appraisal for pity/gloating and happy-for/resentment is done for the remaining consequences. The former is straightforward, weighs the desirability with the goal's relevance and directly uses the absolute value as intensity. The intensity of joy and distress is obtained by multiplying the relevance of a goal with its desirability. To determine the intensity of emotions that are reactions to consequences for others, this value is additionally multiplied with the sympathy to this entity (see below). When the agent is indifferent to the entity, the emotions will have very low intensities. High desirability and sympathy for another agent leads to the emotion "happy for", high desirability and negative sympathy to resentment, low desirability and sympathy to pity, and low desirability and negative sympathy to gloating.

The remaining emotions are referred to as *compound emotions*, as they are the result of combining the consequences for the self and the praiseworthiness. If the agent himself was the cause, the emotion is determined based on the desirability of the event. Positive consequences cause gratification, negative ones result in remorse. Appraisal for compound emotions regarding other entities' actions are handled in a similar way and cause either gratitude or anger. An overview of emotions and their criteria is given in Table 1.

The interplay of emotions, mood-state and personality

Emotions (short-term), mood-states (mid-term) and personality (long-term) interact in multiple ways, which will be described in this section. First of all, the personality of an agent is stored under the form of user-defined values for the five personality dimensions (Five Factor Model, FFM): the values for extroversion, conscientiousness, agreeableness, neuroticism and openness are defined at the beginning and remain the same throughout the scenario. Empirical research has shown that with healthy subjects, FFM dimensions correlate with trait dimensions of pleasure, arousability and dominance (Mehrabian, 1996). Mehrabian also provided equations (2) to convert FFM into PAD. This is used in our emotion model to set the default mood-state according to an agent's personality.

$$P=0.21 \cdot \text{Extroversion}+0.59 \cdot \text{Agreeableness}+0.19 \cdot \text{Neuroticism}$$

$$A=0.15 \cdot \text{Openness}+0.30 \cdot \text{Agreeableness}-0.57 \cdot \text{Neuroticism}$$

$$D=0.25 \cdot \text{Openness}+0.17 \cdot \text{Conscientiousness}+0.60 \cdot \text{Extroversion}-0.32 \cdot \text{Agreeableness} \quad (2)$$

	Emotion	Cause
Positive Emotions	Admiration	Praiseworthy deed by other
	Gloating	Bad consequence for disliked other
	Gratification	Good consequence through own deed
	Gratitude	Good consequence through others deed
	Happy for	Good consequence for liked other
	Hope	Potential good consequence expected
	Joy	Good consequence
	Love	Attractive entity
	Pride	Praiseworthy deed
	Relief	Expected bad consequence not confirmed
	Satisfaction	Expected good consequence confirmed
Negative Emotions	Anger	Bad consequence through others deed
	Disappointment	Expected good consequence not confirmed
	Distress	Bad consequence
	Fear	Potential bad consequence expected
	Fears confirmed	Expected bad consequence confirmed
	Hate	Repelling entity
	Pity	Bad consequence for liked other
	Remorse	Bad consequence through own deed
	Reproach	Blameworthy deed by other
	Resentment	Good consequence for disliked other
	Shame	Blameworthy deed

Table 1. OCC emotions and their respective causes

When no other emotions are active, the current mood-state is slowly changing back to the default mood-state corresponding with the agent’s personality traits. Each of the five traits has a range of [-1; 1]. The current range is based on 0 typically being used as the average, and values above or below are in relation to an average persons rating in this factor. Neuroticism additionally influences mood-states as it is positively correlated with the speed of mood change (see below).

Each appraisal results in one or multiple emotions. On the one hand, these emotions are projected into PAD space and attract the current mood-state (pull function). On the other hand, an emotion is active or inactive depending on its proximity to the current mood-state. All of these changes are implemented in an *update* function with the time that has passed since the last emotional trigger. This update function will make the agent update its internal state, which results for example in the intensity of emotions being reduced depending on their decay values and the current mood-state. Active emotions dropping below their threshold will become inactive, emotions dropping to zero or less are removed from the update function. Changes to the mood-state may also cause inactive emotions to become active. All remaining active emotions are used to calculate the emotion centre, which then attracts the current mood-state. This is done with a pull function: if the mood-state is located between the origin and the emotion centre, the mood-state is moved closer to the emotion centre. The emotions should not all have the same influence for the emotion centre, so instead, the intensities of all emotions are added up and the ratio is used as a weight. The average intensity of all emotions is used to determine how much the mood-state is attracted

to the (calculated) emotion centre. Another factor is the agent's neuroticism. Higher neuroticism means a tendency to be moody and experience mood swings, which is simulated by simply allowing the mood-state to change faster than for other non-neurotic agents.

To decay an emotion, the distance of the emotion to the mood-state is used. The decay time is given as the time needed for the emotion to fully decay from 100% to 0.

Relations between agents are stored as simple PAD values linked to the name of every other agent that has caused emotions in the agent so far. Essentially, the relations are just an emotion average and can, for example, be considered as an associated mood-state towards the entity.

4. Empirical data in a game scenario

To test our emotion module a simplified game scenario was created, where players have to take over squares in a battlefield. The goal was to create an environment where the virtual agents can develop emotions and mood-states while the game would be kept simple enough, so the observed effects could be understood. Virtual players were made to react to game events by experiencing emotions and mood-state shifts, and these factors were also made to have an influence back on their decision making. The results of pure AI and AE players were compared. Apart from the fact that many different and realistic emotions appeared, AE agents have shown somewhat of an emergent behavior, resembling cooperation.

The simplest possible scenario that could still trigger all possible emotions was created. This scenario was reduced to one abstract resource and one possible action. The resource was squares on a board; the action was stealing a square adjacent to one of your own. An addition had to be made to allow for explicit cooperation: before an attempt to take over a square is made, the other players can be asked for help to improve chances of success. However, if the other players disagree, the chances to succeed are reduced, so the other player that is chosen to help should be chosen carefully.

Before describing the behavior of the emotional agents, a short introduction will be given on the way pure AI worked in this scenario. The AI module was set to always pick the player with the least number of squares to eliminate him from the game as quickly as possible in an attempt to reduce the potential attackers. Players with more than 80% of all squares are to be considered close to winning and be attacked as well. The decision to ask other players for help is purely based on their previous responses. If they agreed in 50% of all cases, then they are asked again. Before they have been asked at least five times, this statistic is ignored and they are always asked. The response of the agent(s) asked for help depends on the number of squares of the affected players. Typically, an AI agent would always agree, as long as the attacker had less squares than the agent himself.

Since this scenario was only used to test whether the emotion model was adequate to produce emotions as reactions to events, the emotion module did not affect AI decisions. Because of the simple nature of this scenario, the emotion module alone could play the game, which allows a direct comparison between AI and AE. Events were still evaluated by the module, but all decisions were based on sympathy. The least-liked player would be attacked and only players which 'inspired' positive sympathy would be asked for help or supported.

In order to test the emerging emotions, data was collected from thousands of test runs with virtual agents playing against each other. The first batch consisted of one thousand games

with four agents using the emotion module. Another thousand games mixed two emotional agents with two pure AI players, and the final batch included thousands of games of four AI players. To collect the data, a call-back function counted every time an emotion was changed from inactive to active. However, as the length of the game varied greatly, especially between the set using emotion and the set using AI, the results were divided by the number of turns. Multiple test runs were made and compared to make sure the results within the different groups were consistent and stable. While the different test groups showed significant differences between each other, the groups themselves did not.

The emotional agents always behaved in the same way, though in their case this happened naturally and was exactly the kind of emergent behaviour that the module was implemented for. Each player would end up with one liked ally and two disliked players, one of them usually disliked about twice as much as the other. These relationships were completely symmetrical. Since decisions were purely based on sympathy, it meant two players would always attack each other with one of the other two supporting him. As a result, games would last a very long time as squares would constantly go back and forth between these two players. Figure 3 demonstrates the longer duration of pure AE games in turns compared to mixed and pure AI groups. Once the first player was out, the remaining player would lose very quickly. Considering the very simple decision making rules of the application itself, the fact that agents automatically formed teams after a few moves is a good sign that complex behaviors could emerge in more interesting scenarios.

It turned out that games with only AI players showed the least cooperation, with only about 40% of the attacks being joint attacks and only 75% of the requests being answered positively. This is not surprising. The leading player would always be denied and would soon stop asking. This is also the reason why only 25% of requests were refused. Only the weakest players would regularly receive help and kept asking for it. Looking at the emotional players, over 90% of their attacks were joint attacks and almost 100% of them were granted. These 90% must be taken with caution, as an attack is counted as joint attack, if at least one other player is invited to join. As mentioned above, the emotional players would always have two enemies and one ally, so only one request was made each time. However, it is worth pointing out that the close to 100% success rate means that sympathy was always mutual. This is an interesting result, as no extra code was written to predict the other player’s answer and guarantee a positive response.

Fig. 3. Game duration comparing AE and AI players

Figure 4 (left side) shows examples of emotions “experienced” by agents in regard to their position in the game, i.e. whether they were winning or losing. Though almost all of the OCC emotions were triggered in the game we present only the most relevant ones for this analysis. Interestingly, the way the AI players immediately focused on the weakest player gave some insight into the different emotions of an agent that won without being attacked and one that was constantly attacked by everybody else. Surprisingly, the latter is not experiencing only negative emotions but shows a more interesting pattern. With a few exceptions, *all* emotions are experienced a lot more often by the losing player, including positive ones. An explanation for this is the fact that there are more emotional triggers for this player in the game since he is attacked (and hence “interacts” with other players) more often than winning agents. The losing player (fourth place in the figure) “experiences” distress because of being attacked all the time and disappointment because of eventually losing; these emotions are logical and show the realism of the emotion engine. Resentment is also natural since the losing player obviously does not appreciate that the other players win square after square, ever worsening his position. The fact that pity is equally distributed among players also makes sense since there is no reason why an agent’s rank in the game should influence his empathy and willingness to feel pity for other players. Why is the losing player experiencing more joy and relief, though? Although he is constantly under attack, a lot of those attacks fail because of mixed support the other players grant the winning player. Hence, there is high potential for joy (an attack successfully deflected, or the disliked other player not gaining support). Because the last player in rank is in constant fear of losing squares, relief will occur every time an attack fails.

Fig. 4. Emotions by ranking (left side) and emotions by personality (right side)

Finally, figure 4 (right side) demonstrates the influence of extreme personalities on the emotions experienced. This was tested in a different run where predefined personality variables were modified systematically. The first personality was completely average as in previous games. The second one was an absolutely open, extroverted and non-neurotic person, while the third one was his exact opposite, a highly neurotic and introverted conservative. Though negative emotions were usually experienced more often in the other test runs with average personalities, the extroverted player forms an exception: his experience consisted almost only of positive emotions. “Happy-for” is triggered extremely often and is the result of this player “being the friend of everyone”. Since there are multiple chances within a round to feel happy for other players’ achievements, “happy-for” has a high value. “Satisfaction” is also occurring naturally as the open and stable player is

optimistic and sees his hope often confirmed, partially due to good support from other players. In the opposite way, the neurotic person experienced nothing but negative emotions, again with the exception of relief, as his pessimism would not always turn out to be appropriate. This player seems to be always gloating and resentful. His lack in extroversion might also explain why only the most intense negative emotions came to the surface.

5. Conclusion

In this paper we presented SIMPLEX, a new model to simulate emotions. This model operates with three layers: personality, mood-states and emotions, which are interconnected in a believable way. Personality influences decisions on multiple levels. Mood-states are represented in a PAD space, they influence emotions and act as a means to form relationships to other agents. Emotions are generated according to OCC appraisal rules.

AI agents and AE agents whose emotions were generated by SIMPLEX were made to play a simple game. The data collected showed interesting emergent behavior. First, cooperation emerged between players. Second, the player losing the game is the one who experienced the most emotions, including positive ones.

It is questionable whether strictly separating AI from AE is realistic. Though it has been shown that in some cases, emotions may be triggered automatically, as when one freezes in front of a dangerous animal, in most cases emotions and cognition interact (Davidson, 2003; Ledoux, 1996). Cognitive evaluations of a situation, which are traditionally within the realm of AI, are also part of the appraisal processes necessary to trigger emotions. Frijda (1986) gives the example of how one can feel progressively overwhelmed by anger after having heard someone criticize a friend. On the contrary, Damasio (1996) has shown that emotions were necessary to cognitive processes.

Further research should thus investigate the behaviors of “Artificial-Emotionally-Intelligent” agents, which would be more realistic than AI or AE ones. This next generation of agents, which may be just as intelligent as AI ones, may then grow to be a little “emotionally creative” (Averill, 2004).

6. References

- Averill, J. R. (2004). A tale of two snarks: Emotional intelligence and emotional creativity compared. *Psychological Inquiry*, 15, 228-233.
- Bradley, M. M., & Lang, P. J. (1994). Measuring emotion: the Self-Assessment Manikin and the Semantic Differential. *J Behav Ther Exp Psychiatry*, 25(1), 49-59.
- Damasio, A. R. (1996). The somatic marker hypothesis and the possible functions of the prefrontal cortex. *Philos Trans R Soc Lond B Biol Sci*, 351(1346), 1413-1420.
- Davidson, R. J. (2003). Seven sins in the study of emotion: Correctives from affective neuroscience. *Brain and Cognition Affective Neuroscience*, 52(1), 129-132.
- Frijda, N. (1986). *The Emotions*. New York: Cambridge University Press.
- Gray, J. A., & McNaughton, N. (1996). The neuropsychology of anxiety: reprise. *Nebr Symp Motiv*, 43, 61-134.
- LeDoux, J. E. (1996). *The emotional brain: the mysterious underpinnings of emotional life*. New York: Simon & Schuster.

- McCrae, R. R., & Costa, P. T., Jr. (1987). Validation of the five-factor model of personality across instruments and observers. *J Pers Soc Psychol*, 52(1), 81-90.
- Mehrabian, A. (1996). Analysis of the big-five personality factors in terms of the PAD temperament model. *Australian Journal of Psychology*, 48(2), 86-92.
- Ortony, A., Clore, G. L., & Collins, A. (1988). *The Cognitive Structure of Emotions*. Cambridge: Cambridge University Press.
- Reeves, B., & Nass, C. (1996). *The Media Equation. How People Treat Computers, Television, and New Media like real People and Places*. New York: Cambridge University Press.
- Russell, J. A. (1978). Evidence of convergent validity on the dimensions of affect. *Journal of Personality and Social Psychology*, 36, 1152-1168.
- Simon, H. A. (1967). Motivational and emotional controls of cognition. *Psychological Review*, 74, 29-39.
- Staller, A., & Petta, P. (1999). Towards a Tractable Appraisal-Based Architecture for Situated Cognizers. In D. Canamero (Ed.), *Grounding Emotions in Adaptive Systems*. Zurich, Switzerland: Fifth International Conference of the Society for Adaptive Behaviour (SAB98).
- Toda, M. (1982). *Man, robot, and society*. The Hague, Nijhoff.
- Traue, H. C., & Kessler, H. (2003). Psychologische Emotionskonzepte. In A. Stephan & H. Walter (Eds.), *Natur und Theorie der Emotion* (pp. 20-33). Paderborn: Mentis.
- Wilson, I. (2000). *The Artificial Emotion Engine, Driving Emotional Behavior*. Paper presented at the AAAI Spring Symposium.

IntechOpen

Affective Computing

Edited by Jimmy Or

ISBN 978-3-902613-23-3

Hard cover, 284 pages

Publisher I-Tech Education and Publishing

Published online 01, May, 2008

Published in print edition May, 2008

This book provides an overview of state of the art research in Affective Computing. It presents new ideas, original results and practical experiences in this increasingly important research field. The book consists of 23 chapters categorized into four sections. Since one of the most important means of human communication is facial expression, the first section of this book (Chapters 1 to 7) presents a research on synthesis and recognition of facial expressions. Given that we not only use the face but also body movements to express ourselves, in the second section (Chapters 8 to 11) we present a research on perception and generation of emotional expressions by using full-body motions. The third section of the book (Chapters 12 to 16) presents computational models on emotion, as well as findings from neuroscience research. In the last section of the book (Chapters 17 to 22) we present applications related to affective computing.

How to reference

In order to correctly reference this scholarly work, feel free to copy and paste the following:

Henrik Kessler, Alexander Festini, Harald C. Traue, Suzanne Filipic, Michael Weber and Holger Hoffmann (2008). SIMPLEX — Simulation of Personal Emotion Experience, Affective Computing, Jimmy Or (Ed.), ISBN: 978-3-902613-23-3, InTech, Available from: http://www.intechopen.com/books/affective_computing/simplex_-_simulation_of_personal_emotion_experience

INTECH
open science | open minds

InTech Europe

University Campus STeP Ri
Slavka Krautzeka 83/A
51000 Rijeka, Croatia
Phone: +385 (51) 770 447
Fax: +385 (51) 686 166
www.intechopen.com

InTech China

Unit 405, Office Block, Hotel Equatorial Shanghai
No.65, Yan An Road (West), Shanghai, 200040, China
中国上海市延安西路65号上海国际贵都大饭店办公楼405单元
Phone: +86-21-62489820
Fax: +86-21-62489821

© 2008 The Author(s). Licensee IntechOpen. This chapter is distributed under the terms of the [Creative Commons Attribution-NonCommercial-ShareAlike-3.0 License](https://creativecommons.org/licenses/by-nc-sa/3.0/), which permits use, distribution and reproduction for non-commercial purposes, provided the original is properly cited and derivative works building on this content are distributed under the same license.

IntechOpen

IntechOpen