Emotion Eliciting in Salt & Pepper

Luis Botelho*
Luis.Botelho@iscte.pt

Pedro Ramos*
Pedro.Ramos@iscte.pt

Pedro Figueiredo*
Pedro.Figueiredo@we-b-mind.org

ADETTI/ISCTE
We, the Body and the Mind Research Lab

Abstract

This paper describes the emotion elicitation process of the Salt & Pepper architecture for autonomous agents. Its main contributions are a general framework upon which it is possible to define different theories of emotion generation, the view of artificial emotion as an adaptive mechanism that is designed at the level of the agent architecture, and use and the discussion of several emotion eliciting mechanisms with special emphasis on non-cognitive emotion generation processes. We propose to add *emotion eliciting inhibition times* to emotion eliciting rules in order to avoid the repeated generation of the same emotion due to the same circumstances. Finally, we propose the *conditioned emotion eliciting chunk*, which sequentially generates conditioned emotions. This new emotion eliciting structure is useful to initiate emotion processes in stereotypical situations in which the initiated emotion depends on the emotion that has previously been generated. Our experiments were done during the SAFIRA European Project.

1 Introduction

The goal of our work regarding emotions in artificial autonomous agents is to build agent architectures with general adaptive mechanisms capable of improving agent performance. Some of those mechanisms, which we call artificial emotions, have been inspired by roles played by natural emotions for living organisms. The result is a component based agent architecture called Salt & Pepper. Salt & Pepper agents are made of several components, each of which may be composed of several modules, including interaction management modules [Gonçalves, Jesus and Botelho 2003].

In the Salt & Pepper architecture, emotion is a sequential possibly iterative process comprising three stages:

Evaluation of the agent global state (i.e., external environment plus internal state)

Emotion generation (i.e., generation of emotion-signals)

Emotion-response, which is the way the agent responds to generated emotion signals.

In the evaluation stage, the agent global state (internal state and external environment) is partially evaluated so that it may be determined if the conditions for emotion elicitation are met. In the second stage, an emotion is initiated in case some emotion eliciting conditions are met, giving rise to an emotion signal. The generated emotion signal is sent to other components of the agent architecture possibly causing modification in the agent subsequent behaviour. One of the possible effects of emotion signals is to interrupt the current process being carried out by the agent working memory. This is exactly what was put forth by Simon more than 30 years ago [Simon 1967].

The paper describes experiments performed during the SAFIRA IST[1] project to create agents using the Salt & Pepper architecture. The main concern of the paper is the description of the used emotion generation mechanisms. The remaining of the paper is organized as follows. Section 2 describes the agent component architecture. Section 3 describes our experiments regarding mechanisms responsible for emotion generation (stages 1 and 2 of the emotion process described above). Section 4 describes related work. Finally, section 5 presents conclusions and it points directions for future developments.

2 Agent Architecture

In the Salt & Pepper architecture, emotional and cognitive processes interact in somehow unpredictable ways to control the agent behaviour.

Using the Salt & Pepper component architecture, an agent is composed of a set of independent and replaceable components possibly running on different processors, which interact with each other via XML messages. An

* This work was partly supported by UNIDE/ISCTE and partly by the FP5 ITS Project IST-1999-11683

[1] IST – Information Society Technology, 5th Framework Programme

agent can have more or less components (for example, it can have several appraisal components).The agent must have components with the following functions: sensors, emotion eliciting, action control, and effectors. It is not mandatory to have one component for each of the required functions. A single component may provide more than one of these functions, and there may be more than one component for the same function.

In our conceptual and computational framework for component-based agents, each component registers its services with a special component called the Central Registry Service (CRS). This service provides the dynamic discovery of services allowing any component to subscribe or request a service without having prior knowledge about the provider identity, and the run-time replacement and/or deletion of components [Gonçalves, Jesus and Botelho 2003].

In one of our experiments (Figure 1), the implemented agent has four components, additionally to the CRS: an appraisal component, a control component (Em-PSys), an effector and a sensor.

The appraisal component receives information from the sensor and from the Em-Psys and generates adequate emotion signals.

The Em-PSys (Emotional Production System) is an action control component for deliberatively controlling the agent behaviour. In the absence of other influences, the Em-PSys works as a regular production system: in each cycle of its operation, it selects all satisfied rules to the conflict set, and then it picks the action part of one of those rules. The selected action is sent to the effector for being executed. When an emotion-signal is generated by one of the emotion eliciting mechanisms it is sent to the agent components that have subscribed it. If an emotion-signal is received in the action control component (Em-PSys), its regular functioning may be overridden by the emotion response. Emotion responses are specified by special purpose control structures called *pattern-action control structures*. If the emotion-signal matches the pattern of a pattern-action control structure, it is possible that the specified action is sent to the effector to be executed overriding the deliberative control of the Em-PSys.

The Salt & Pepper memory is not an independent component. It is a software resource used by the Em-PSys component to store and manipulate the required knowledge structures. The Salt & Pepper memory comprises two distinct memories: working memory and long-term memory. Long-term memory is an associative memory with spreading activation that permanently stores the agent's facts about the world, production rules, pattern-action structures and other kinds of knowledge, including procedural knowledge and descriptions of action effects. Working memory is a limited memory space and executor that only holds information that is being processed at each moment in time. Working memory represents the agent focus of attention.

Since the Em-PSys production rules are stored in the nodes of the Salt & Pepper long-term memory, it uses the dynamics of long-term memory nodes accessibility as a conflict resolution policy. When the conflict set contains more than one rule (i.e., when more than one rule is satisfied), the Em-PSys picks the action of the satisfied rule that is contained in the most accessible long-term memory node of those containing rules in the conflict set.

In summary, emotion conditions agent behaviour through two mechanisms: it may influence the selection of the rule of the conflict set whose action should be executed, and it may also directly select an action to be executed through the use of pattern-action control structures. Figure 1 shows how emotional responses can change agent behaviour.

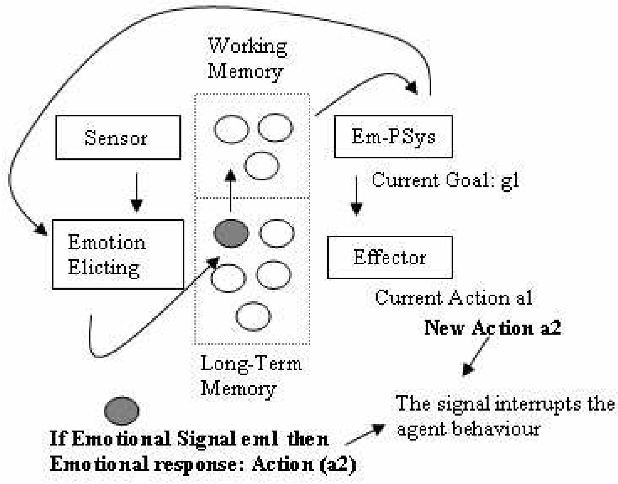

Figure 1 - Agent Architecture and its Interruption Mechanism

We have developed a software tool called the Appraisal Compiler to assist the creation of emotion eliciting mechanisms. This tool generates appraisal modules or full appraisal components that communicate with other components using the defined interaction mechanisms. Besides helping agent designers creating emotion eliciting mechanisms, the Appraisal Compiler offers two built-in domain-independent modules – Expectations Evaluator Module and Unpredicted Actions Module. The offered modules may be used, if desired, without any further modification. However, they may also be modified at will by agent designers to fulfil their specific needs.

3 Emotion eliciting mechanisms

In the SAFIRA IST Project, we have created and implemented the Monga World, a game designed to evaluate the Salt & Pepper architecture for artificial autonomous agents. We have implemented several versions of a Salt & Pepper agent to play the Monga World and abstracted some design principles allowing both better software engineering practice and better agent performance.

We have used the Monga World as a challenging unpredictable and unreachable environment in which we can assess if emotion can help the agent to better face and solve its problems and to test different agent designs. This paper focuses on the created emotion eliciting mechanisms.

3.1 Distributed Emotion Eliciting

Using the component-based Salt & Pepper architecture it is possible to create distributed emotion eliciting mechanisms. It is possible to implement different dedicated emotion eliciting components each one focused on particular kinds of emotion eliciting conditions, or on specific kinds of emotion signals. It is also possible to include emotion eliciting modules in components whose main purpose is not emotion generation. For instance, it is possible to include emotion eliciting modules in sensing components.

From an engineering point of view this is an advantageous possibility for both design and efficiency reasons. According to good practice software engineering design principles, emotion eliciting modules handling emotion eliciting conditions defined on the information naturally available in a given component should be built into that component.

Efficiency can be improved by reducing the amount of information exchanged among the different components of the agent. Therefore the integration of an emotion eliciting module in a component may be a sensible design choice if it contributes to reduce the amount of information that must be exchanged among the components. However, if the emotion eliciting mechanism is to receive information from several components, it is better to create a separate emotion eliciting component.

From a scientific point of view there is also evidence that there are different mechanisms of emotion generation which are distributed in different locations in the brain. Neurological findings suggest that different although cooperating brain systems are responsible for different groups of emotions [Damasio 2003][LeDoux 1996]. Cognitive and artificial intelligence scientists have also put forth the idea of multi-layered emotion generation processes [Leventhal and Scherer 1987][Staller and Petta 1998]. This lends some support to the distributed emotion eliciting hypothesis.

If desired by agent designers, the Salt & Pepper component-based architecture also allows the design of a single centralized appraisal component. In our experiments in the SAFIRA project we have created an agent with a centralized appraisal component, a second emotion eliciting mechanism in the action control component (the Em-PSys), and yet a third emotion eliciting mechanism integrated in the agent effector.

The centralized appraisal component is used mainly to generate performance evaluation emotion signals whenever the expected results of action execution are not met by its actual results. The appraisal component receives the action expected results from the action control component (Em-PSys) and the actual results from the sensor component.

The emotion eliciting mechanism integrated in the agent action control component generates emotion signals that detect when the actions that have been sent to the effector for execution were not deliberately planned. This happens when the normal operation of the action control component is overridden by unplanned emotion responses.

Finally, the emotion eliciting module that has been built within the agent effector generates performance evaluation emotion signals indicating the failure or the success of action execution.

3.2 Cognitive Processes in Emotion Generation

Cognitive appraisal theorists of emotion generation [Ortony, Clore and Collins 1988] have put a strong emphasis on cognitive appraisal processes. One of the most mentioned appraisal process is imminently cognitive – the evaluation of the expected results of a planned action against its actual effects. When an agent, especially as the result of a planning algorithm, chooses an action to be executed with the purpose of achieving certain goals, it may compare the actual results of its execution with the expected results which have lead to its choice. The explicit comparison of expected results with actual results involves the representation of world states (those expected to be achieved and those actually achieved) which, by definition, is a cognitive process.

In one experiment in the SAFIRA project, the agent action control module computes the desired expected results of the actions it plans for being executed. In the same process, the agent sensor is responsible to acquire information regarding the actual results of executing the planned action. These two pieces of information are sent to the emotion eliciting component that compares the expected with the actual results possibly generating the corresponding emotion-signal. The intensity of the emotional signals depends on the importance of the goal that should have been fulfilled by the action that failed to achieve it, and the effort the agent has already spent to achieve the goal (the effort variable was pointed out in [Ortony, Clore and Collins 1988], goal importance is indicated in [Gratch 2000]). A heuristic measure of the importance of the goal in the Salt & Pepper architecture is its accessibility in the agent long-term memory.

3.3 Non-Cognitive Emotion Generation

[Botelho and Coelho 2001] stresses the possibility of non-cognitive emotion eliciting processes. There, the main criterion for distinguishing a cognitive appraisal process from a non-cognitive appraisal process was grounded in the agent architecture. Appraisal processes carried out by the cognitive engine would be dimmed cognitive. Appraisal processes carried out by non-cognitive components of the architecture would be considered non-cognitive (*affective*, in the original).

In our experiments in the SAFIRA project, one agent was created with an emotion eliciting module integrated in the agent effector. This emotion eliciting mechanism takes action execution statuses and the recent history of action execution and generates performance evaluation emotion signals.

If action execution fails, a negative emotion signal of performance evaluation is generated. The signal is more intense if the execution of the same action has recently failed more often.

If action execution succeeds, a positive emotion signal of performance evaluation is generated. The intensity of the signal decreases as the number of times the successful execution of the same action increases. When the action execution status changes (from success to failure or vice versa) the intensity of the emotion signal is proportional to the difference between the number of succeeded and failed actions.

This is obviously not a cognitive emotion eliciting process. Firstly, it is performed by the effector which will certainly not be considered a cognitive component of the architecture. Secondly, there are no representations involved. The process just takes action execution statuses and the history of action execution and generates an emotion.

Our design decision is consistent with other authors [Leventhal and Scherer 1987][Staller and Petta 1998], which have proposed architectures including sensory-motor level emotion eliciting mechanisms.

3.4 Emotion Eliciting Structures

In the implementation we have created in the SAFIRA project of the Salt & Pepper architecture, emotion eliciting was originally done by a prioritised set of JESS production rules of the form *IF <condition> THEN generate-emotion-signal(<emotion-signal>).* Soon we came to the conclusion that this simple structure, while providing enough means for several applications, was not enough.

One difficulty we found with this structure was that, while the specified condition holds, the system would keep on generating the same emotion signals. This caused several problems: the communication between components was flooded with emotion signals; the receiving components had no time to process all received emotion signals on time; and there was no easy way of distinguishing a relevant emotion signal from its mere repetition.

The above problem led us to the decision of attaching emotion eliciting inhibition times to emotion eliciting rules. Emotion eliciting rules became 3-tuple structures containing a condition, an emotion signal and the emotion eliciting inhibition time. When the rule condition is satisfied the specified emotion signal is generated only if at least a minimum time interval has passed since the last time the rule fired. When the condition is satisfied but not enough time has passed since the last time the rule was used, the specified emotion signal is not generated and the rule is not used. This minimum elapsed time is specified by the *emotion eliciting inhibition time* parameter. For the time being, the value of the *emotion eliciting inhibition time* parameter associated with each emotion eliciting rule is predetermined quite arbitrarily by agent designers.

In the Monga World game several agents and some dangerous living beings, like the Grosserontes, inhabit in their world. One appraisal rule stands that when an agent sees a grosseronte an emotional signal (Fear) is generated. In order to avoid a permanently generation we used the emotion eliciting inhibition time parameter to ensure that the gap between two generations is not less than 8 seconds.

Another shortage of the usual rule-based concept for emotion generation is the inefficiency it creates and the difficulty to consider dependencies between previously generated emotion signals and current emotion generation. Sometimes, it is necessary to capture the sequential generation of emotion signals in which each signal is conditionally generated. The process should be terminated as soon as one of the conditions fails.

We propose a new emotion eliciting structure capable of capturing the described concept – the so-called *conditioned emotion eliciting chunks*. Conditioned emotion eliciting chunks are sequences of pairs *<condition*$_i$*, emotion-signal*$_i$*>*. When a conditioned emotion eliciting chunk is executed, the condition of the first pair *<condition*$_1$*, emotion-signal*$_1$*>* is evaluated. If it is satisfied, the associated emotion signal is generated and the execution process continues to the next pair until the sequence is exhausted. If the condition is not satisfied, the execution process halts.

Conditioned emotion eliciting chunks resemble conditioned plans but the similarity does not resist a deep analysis. First of all, conditioned plans are aimed at controlling the agent actions while conditioned emotion eliciting chunks only control the generation of emotion signal which do not directly lead to immediate action. Secondly, while conditioned plans are dynamically generated by the planning algorithm in face of a new problem, conditioned emotion eliciting chunks are programmed by the agent designer or learnt through a learning algorithm. In the third place, the execution of conditioned plans continues even if a certain condition is found that is not satisfied, while the sequential generation of emotion signals by the conditioned emotion eliciting chunk halts as soon as a non-satisfied condition is found.

Conditioned emotion eliciting chunks are very useful progressive alarm systems that warn the agent not to be involved in its current course of action because it will likely lead to negative results. [Figueiredo and Botelho 2003] presents an algorithm for emotion learning capable of learning conditioned emotion eliciting chunks.

3.5 Evaluation

The evaluation of our proposal regarding emotion eliciting cannot be a quantitative process. In fact, it is not at stake whether or not, agents with emotions are better than agents without emotions. Although that kind of evaluation is necessary, it was not the purpose of this paper. This paper presents a set of mechanisms that may be used in the Salt & Pepper architecture for emotion eliciting.

The evaluation of the proposed mechanisms consist of determining if they comply with good design principles, if they provide flexible and powerful ways of designing emotional agents covering a wide range of possibilities, if it is easy to use them when creating agents with emotions, and maybe if they provide the means for experimenting diverse psychological and neurological mechanisms.

Most of these criteria have already been discussed in this section or will be discussed in section 4. The only evaluation criterion that has not yet been analysed is determining if agent designers and programmers find it easy and intuitive to use the proposed emotion eliciting mechanisms when creating agents with emotions. Unfortunately, we have not assessed this criterion because our architecture is not yet mature enough to be used by a large community of agent designers.

4 Related Work

The artificial agents communities have used emotion models with a variety of purposes, namely to improve agent believability [Ortony 2003], to improve the interaction between software or robotic agents and people [Breazeal and Velásquez 1998][Peters II et al 2001][Picard 2003], and to improve agent performance at specific tasks [Elliott, Rickel, and Lester 1997][Numao et al 1997]. Emotion models have also been implemented in artefacts as a way of understanding or simulating natural emotions [Adams et al 2000][Cañamero 1997].

Our work regarding emotions in artificial agents has followed a different path. We are trying to implement artificial mechanisms that play the same beneficial roles for artificial agents as natural emotions do for people and animals [Damasio 1994][Damasio 2003][LeDoux 1996].

Sloman and the Cognition and Affect project [Sloman 2003], and Petta and his co-workers [Staller and Petta 1998] propose a three layered architecture in which emotion is generated at the three layers of the architecture, corresponding to different types of information processing. In the experiments described in this paper, we have implemented three different types of emotion generation processes – the cognitive level (comparison of the expected and actual results of action execution), the schematic level (detection of non-planned action execution) and the sensory-motor level (status of action execution by the effector). We stress the fact that the emotion eliciting mechanism we have embedded within the agent effector cannot be considered a cognitive process. [Peters II et al 2001] also propose a similar emotion eliciting mechanism.

Other researchers have suggested other non-cognitive processes of emotion generation, namely the generation of emotion responses due to learnt associations between stimulus and emotion responses and emotion generation through imitation [Breazeal and Velásquez 1998]. Both of these mechanisms are consistent with neurological theories and findings [Damasio 2003]. Due to the workings of its memory, the Salt & Pepper architecture is capable of automatically learning associations between stimuli and emotion responses. Although we recognize the importance of imitation both as an emotion generation process and as a general learning mechanism, we have not yet tried to create it in the Salt & Pepper architecture.

If we need fast emotion processes, emotion eliciting should rely upon specific and short emotion eliciting conditions and should avoid lengthy explicit evaluations of the satisfaction of the agent motives [Botelho and Coelho 2001]. However, for practical reasons, one cannot specify the emotion that should be generated for each specific significant event or stimulus. Some abstraction should be used [Bartneck 2002]. Since these two requirements are somehow conflicting, we have only implemented the most basic, general and necessary set of domain independent emotion eliciting conditions and rely upon learning mechanisms to acquire the remaining necessary conditions. This approach is described in [Figueiredo and Botelho 2003].

Ortony and colleagues have developed a domain-independent theory of emotion elicitation [Ortony, Clore and Collins 1988]. This model has been used in artificial agents to generate and to explain emotions from a given stimulus [Elliot 1992, 1993][O'Rorke and Ortony 1994]. Although the emotion eliciting theory is domain-independent, its practical application requires domain knowledge necessary to interpret situations. Although we have implemented general domain-independent emotion eliciting mechanisms, we acknowledge the role played by domain specific knowledge. This however should be acquired through emotion learning mechanisms. Instead of trying to increasingly use domain independent emotion eliciting conditions, we have been more worried with finding general architectural principles and general mechanisms that may be used for emotion generation.

5 Conclusions and Future Work

We have created a component-based agent architecture called Salt & Pepper, which can be used by agent designers to implement their agents according to different agency theories, including different emotion theories, in particular, different theories of emotion generation.

Emotion eliciting was designed as a distributed process carried out by a centralized cognitive appraisal component that generates performance evaluation emotion-signals reflecting the comparison of the expected results of action execution with their actual results, a schema-based emotion eliciting module integrated within the action control component that generates appropriate emotion signals indicating the execution of non planned actions, and a non-cognitive sensory motor level emotion eliciting module integrated within the agent effector that generates performance evaluation emotion-signals depending on the status of action execution.

Finally, we have associated an *emotion eliciting inhibition time* to emotion generation rules, and we have created a new more efficient and flexible emotion eliciting structure called *conditioned emotion eliciting chunk*. Conditioned emotion eliciting chunks allow the generation of emotion-signals for stereotypical situations in which it makes sense to generate emotion signals even before all emotion eliciting conditions have been evaluated.

Some work remains yet to be done in the Salt & Pepper architecture regarding emotion and, in particular, the emotion generation stage. Although we have used domain independent emotion eliciting mechanisms, some domain dependent knowledge is still necessary. We must investigate more domain-independent emotion eliciting mechanisms, paying more attention to the relation between emotions. It is necessary to create emotion learning algorithms, including emotion learning by imitation, capable of acquiring domain specific emotion eliciting structures.

Currently, emotion eliciting inhibition times have been arbitrarily predefined and encoded in emotion eliciting rules. The use of arbitrary numeric parameters has often revealed difficult to tune in implemented systems. The alternative would be to create the mechanisms to automatically determine such numeric parameters.

We need to evaluate our architectural principles in concrete applications so that they may be more thoroughly validated in more professional domains such as eCommerce and process control.

6 References

[Adams et al 2000] Adams, B; Breazeal, C.; Brooks, R.; and Scassellati, B. 2000. "Hummanoid Robots: A New Kind of Tool". *IEEE Intelligent Systems*, August 2000

[Bartneck 2002] Bartneck, C. 2002. "Integrating the OCC Model in Embodied Characters". In *Workshop on Virtual Conversational Characters: Applications, Methods and Research Challenges*

[Botelho and Coelho 2001] Botelho, L.M. and Coelho, H. 2001. "Machinery for artificial emotions", *Cybernetics and Systems* 32(5):465-506

[Breazeal and Velásquez 1998] Breazeal, C.; and Velásquez, J. 1998. "Toward teaching a robot 'infant' using emotive communication acts". In *Proc. of the SAB98 Workshop on Socially Situated Intelligence*

[Cañamero 1997] Cañamero, D. 1997. "Modeling Motivations and Emotions as a Basis for Intelligent Behavior". In *Proc. of the First International Symposium on Autonomous Agents (AA'97)*

[Damasio 1994] Damásio, A.R. 1994. "Descartes' Error: Emotion, Reason and Human Brain", Picador, London, UK

[Damasio 2003] Damasio, A.R. 2003. Looking for Spinoza: Joy, Sorrow, and the Feeling Brain. Harcourt

[Elliott, Rickel, and Lester 1997] C. Elliott, J. Rickel, and J.C. Lester. 1997. "Integrating affective computing into animated tutoring agents". In *Proceedings of the IJCAI Workshop on Animated Interface Agents: Making Them Intelligent*, pages 113-121

[Figueiredo and Botelho 2003] Figueiredo, P.; and Botelho, L.M. 2003. "Salt & Pepper Emotion Learning: the magic SPEL". *Working Paper*

[Gonçalves, Jesus and Botelho 2003] Gonçalves, B.; Jesus, N.; and Botelho, L.M. 2003. "Salt & Pepper Architecture and Toolkit". In *Proc. of the Joint International Conference of Cognitive Science*

[Gratch 2000] Gratch, J. 2000. "Emile: Marshalling Passions in Training and Education". In *Proc. of the 4th International Conference on Autonomous Agents*

[LeDoux 1996] LeDoux, J. 1996. "The Emotional Brain. The mysterious underpinnings of emotional life". Simon and Schuster. New York, N.Y. USA

[Leventhal and Scherer 1987] Leventhal, H.; and Scherer, K.R. 1987. "The relationship of emotion to cognition: a functional approach to a semantic controversy". In *Cognition and Emotion*, 1(1):3-28

[Numao et al 1997] Numao, M., Kobayashi, M. and Sakaniawa, K. 1997. "Acquisition of human feelings in music arrangement", *Proceedings of the 15th International Conference on Artificial Intelligence (IJCAI'97)*, p268-273

[O'Rorke and Ortony 1994] O'Rorke, P.; and Ortony, A. 1994. "Explaining Emotions". *Cognitive Science*, 18:283-323

[Ortony 2003] Ortony, A. 2003. "On Making Believable Emotional Agents Believeble". In Trappl, R.; Petta, P.; and Payr, S. (eds) *Emotions in Human and Artefacts*

[Ortony, Clore and Collins 1988] Ortony, A; Clore, G; and Collins, A. 1988. "The Cognitive Structure of Emotions". Cambridge University Press. Cambridge

[Peters II et al 2001] Peters II, R.A.; Kawamura, K.; Wilkes, D.M.; Hambuchen, K.A.; Rogers, T.E.; and Alford, A. 2001. "ISAC Humanoid: An Architecture for Learning and Emotion". In *Proc. of the IEEE-RAS International Conference on Humanoid Robots*

[Picard 2003] Picard, R.W. 2003 "What does it mean for a computer to have emotions?". In Trappl, R.; Petta, P.; and Payr, S. (eds) *Emotions in Human and Artefacts*

[Simon 1967] Simon, H. 1967. "Motivational and emotional controls of cognition". *Psychological Review*, 74(1):29-39

[Sloman 2003] Sloman, A. 2003 "How many separately evolved emotional beasties live within us?". In Trappl, R.; Petta, P.; and Payr, S. (eds) *Emotions in Human and Artefacts*

[Staller and Petta 1998] Staller, A.; and Petta, P. 1998. "Towards a Tractable Appraisal Based Architecture for Situated Cognizers". *Proc. of SAB'98 Workshop Grounding Emotions in Adaptive Systems*, http://www.oefai.at/oefai/agents/#Publications