

Emotions as Durative Dynamic State for Action Selection

Emmanuel Tanguy, Philip Willis and Joanna J. Bryson

University of Bath

Computer Science Department

e.tanguy@cs.bath.ac.uk, j.bryson@cs.bath.ac.uk, p.willis@cs.bath.ac.uk

Abstract

This paper presents a representation system for maintaining interacting durative states to replicate realistic emotional control. Our model, the Dynamic Emotion Representation (DER) integrates emotional responses and keeps track of emotion intensities changing over time. The developer can specify an interacting network of emotional states with appropriate onsets, sustains and decays. The levels of these states can be used as input for action selection, including emotional expression. We present both a general representational framework and a specific instance of a DER network constructed for a virtual character. The character's DER uses three types of emotional state as classified by duration timescales, in keeping with current emotional theory. The system is demonstrated with a virtual actor.

1 Introduction

Emotion is a popular topic in AI research, but most of existing work focuses on the appraisal of emotions or mimicking their expression for HCI (see review below). Our research is concerned with their role in evolved action selection mechanisms. In nature, emotions provide decision state which serves as a context for limiting the scope of search for action selection [LeDoux, 1996]. This state is sustained more briefly than traditional (life-long) learning, but longer than simple reactive response. To study the role emotions play in intelligence, and to improve the realism of virtual reality characters, we have developed a mechanism for modelling both the temporal course of emotional state and the interactions between such states, which can be both positive and negative. Our work is based on current emotion theory.

Our model, the Dynamic Emotion Representation (DER), assumes other mechanisms for eliciting emotional responses from an agent's environment, but keeps track of emotion intensities changing and interacting over time. Modellers can specify the number and descriptions of fundamental emotions and express how they interact. Each emotion has characteristic intervals of onset, sustain and decay, and each emotion may either excite or inhibit any other. In this respect, the DER can be used as spreading activation system [Maes, 1991, e.g.]

which might be taken as the root of the agent's goal structure. The future state of a DER depends on its current state as well as its environment. Note that in this fully modular system, additional 'higher order' emotions may either be interpreted as emerging from the interactions of fundamental emotions, or they can be introduced with explicit representations — the choice is left to the developer or experimenter.

Although this system could be used to represent any similar durative action selection state (such as drives like hunger), for the purpose of development and evaluation we have focused on emotional state. After presenting the DER framework, we present also here an instance of a DER model. This model explores theories which postulate three types of emotional state changing over three different timescales — emotional *behaviours activations*, conventional emotions, and moods. For evaluation, we have integrated this system into an Emotionally Expressive Facial Animation System (EE-FAS) used to drive a 'talking head' virtual actor. This paper shows that the system provides more diverse visual speech deliveries according to the original state of the DER.

2 Representing Emotional State

Computational emotion models should include two parts:

- mechanisms eliciting emotions from external and internal stimuli, including potentially the agent's own goals, beliefs and standards;
- emotion representations keeping track of the emotional states and their changes over time.

In the design of emotion models the distinction between mechanisms eliciting emotions and emotion representations is useful; the assessment of an emotional event can be the same but its impact on the actions and future emotional state of the virtual actor can vary according to its *current* emotional state. For instance the event of knocking over a cup of tea might make somebody already angry lose their temper, whereas if this person was happy in the first place this negative event might have little impact, just a slight reduction of happiness. An appropriate emotion representation can enable programmers to reduce the complexity of mechanisms eliciting emotions by assessing an identical event in the same way.

Most existing emotion theories are concerned with mechanisms eliciting emotions [Plutchik, 1980; Lazarus, 1991; Frijda, 1986; Izard, 1993; Sloman, 2003]. In contrast, the

duration of emotions and their interactions are not the focus of much research. The same imbalance is found in computational models of emotions — the focus is on the mechanisms eliciting emotions and on their expression, but their representation is typically trivial. The Dynamic Emotion Representation (DER) model described in this paper is novel because:

- it can represent any number of persisting states, such as moods, emotions, drives;
- any emotion impulse can affect any state positively or negatively;
- state responses to emotional impulses are influenced by the other states represented in the DER;

In addition, the DER focuses on usability — it is totally customisable through an XML file.

Besides simplifying the emotion elicitation process, the DER makes it easier to generate variety in the behaviour of real-time interactive character-based AI systems, since the same stimuli can result in significantly different (but locally coherent) responses, depending on the agent’s emotional state. The DER also greatly simplifies scripting for virtual actors by decomposing the problems of specifying the emotionally salient events from describing the agent’s individual reaction to or expression of characteristic emotional states. For instance, a script could specify the action of grabbing an object but characteristics of this action, such as its space and effort [Badler *et al.*, 2002], would be modified per the current DER state producing different animations and adding different interpretations. Finally, the DER presents a powerful mechanism for specifying virtual agent personalities, by allowing developers to for example specify moods or other emotion-related attributes such as tension level.

3 Three Responses to one Series of Events

Before explaining the details of the DER we will first clarify its utility with a concrete example. This example involves a specific instance of a DER system (described below) where there are representations corresponding to three different time courses:

- *behaviour activations*: happy, angry. These correspond most directly to emotion-related behaviours such as smiling. Behaviour activations *trigger* pre-organised behaviours with different intensity levels which follow their own time course after their activations.
- *emotions*: happiness, anger. These build more slowly in response to events and provide a context influencing the current actions.
- *mood variables*: tension, energy. These change very little during the period of the demonstration, but differentiate the various conditions. They are the most persistent of the three state types.

Figure 1 shows a DER in three different mood contexts responding to the same series of six emotional impulses, three happiness impulses followed by three anger impulses. Graph *d* shows these impulses, while graph *a* shows the response in a negative mood (high tension, low energy), *b* a neutral mood, and *c* a positive mood (high energy, low tension). The




Figure 1: Changes of DER states due to six emotional impulses in three contexts. Graphs *a*, *b* and *c* show state changes in the contexts of *Negative Mood*, *Neutral Mood* and *Positive Mood*, respectively. Graph *d* shows the emotional impulses sent over time to the DER.

Happy and Angry graphs have + symbols on them when they result in signal being sent to the EE-FAS system to generate a facial expression (see Figure 2). The EE-FAS also contains a muscle model which results in smoothing between expressed states.

Figure 2 shows the EE-FAS output. The intensities of the behaviour activations determine the strength of the displayed facial expressions. However the durations of facial expressions are innate characteristics and do not relate to the duration of behaviour activations, which are triggers influenced by other DER states. Between each screen shot shown Figure 2 expressions decay slowly to return to its neutral position, or to be replaced by a new expression. Typically a happy expression is shown by raising the lip corners and low-eyelids;

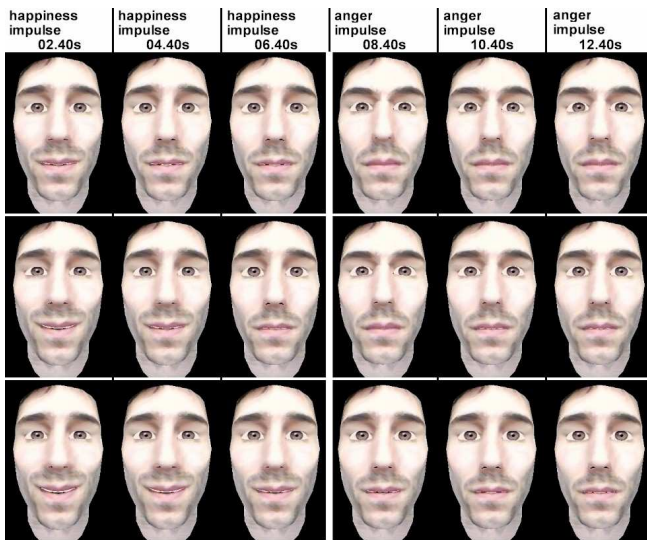


Figure 2: The top through bottom rows show EE-FAS screenshots of the *Negative*, *Neutral* and *Positive Mood*, corresponding to charts *a*, *b* and *c* of Fig. 1. Columns correspond to the time when an emotional impulse is sent to the DER.

and an angry expression is represented by eyebrow frowns and tight lips.

In this example, emotional impulses could be due to the appraisal of six declarations from the employer of the character: “You are our best employee (first happiness stimulus), you have helped us tremendously (second stimulus) and I should increase your salary (third); however I can’t do this (first anger stimulus) — the company is not doing well (second anger), so we will need to work harder (third)”. This sequence of events produces various reactions (such as a smile) depending on the agent’s mood, but they also change the emotional state of the character. This could in turn influence the character’s next action.

In row 1, the first screen-shot shows some response to the happiness impulse. However this response is reduced due to the influence of negative mood. The next two impulses show no response because the happiness impulse intensities are too low. In the neutral mood context, row 2, the first happiness impulse produces a stronger response than in the context of negative mood. Here the second impulse also produces a response but the intensity of the third impulse is still too low to generate a response. The positive mood, row 3, amplifies the happiness impulses, therefore they all produce facial expressions, and these expressions have stronger intensities than in the previous contexts.

The effects of anger impulses are influenced by mood states but in addition, they are also influenced by the level of the emotion *happy* generated by the previous happiness impulses. In row 1, all three anger impulses produce responses, showing an expression of anger represented by an eyebrow frown and tight lips. The different impulse intensities, 100%, 50% and 50%, produce different strengths of expression, but the differences are not marked since the effect of negative mood and the building anger amplifies the lower intensity im-

pulses.

In row 2, the neutral mood, the response of the first anger impulse is more stronger than the responses of the two other anger impulses. In fact, Figure 1 shows that no behaviour activation has been triggered by the second anger impulse, therefore no facial expression should be displayed. The expression shown at the time of the second anger impulse is due to the visual momentum of the expression produced by the first anger impulse which is still fading. The third anger impulse does produce an emotional expression, since happiness has decreased. The difference between these is the emotional momentum produced by the previous happiness impulses. Happiness impulses increase the level of happiness, which takes time to disappear and inhibits the effects of anger impulses. In contrast, the negative agent (row 1) never became very happy in the first place. While for the positive agent, (row 3) no anger impulses produce any behavioural response. This is due to two reasons. First, the positive mood reduces the effects of anger impulses. Second, the happiness momentum produced by the previous happiness impulses also reduce the effects of anger impulses. The largest impact of the bad news is the reduction of happiness which had previously soared.

4 The DER Basic Representation

The basic unit of the DER model is a modular representation based on the Picard description of emotion intensity and emotion filters [Picard, 1997]. A DER network consists of a system of modules connected by filters.

We assume an emotion appraisal mechanism, such as those based on the OCC model [Ortony *et al.*, 1988], classifies events, actions and objects, outputting emotion types and intensities. We call this output *emotional impulses*. Emotional impulses are defined by the name of an emotion, an intensity value and a valence, which can be 1, 0 or -1. The valence specifies whether an emotional impulse is positive, neutral or negative.

Emotional impulses are transformed by the DER into emotion stimuli. The stimuli include timing information typical of an emotion type and represent intensity changes over time. The activation curve representing emotional stimuli, shown at the bottom of Figure 3, was chosen to represent the slow decay of emotion intensity [Picard, 1997]. The effects of small emotional events are cumulative. Therefore emotional stimuli are summed to compute the intensity of an emotion, as shown by the top curve in Figure 3.

As shown in the example (Section 3), the effect of an emotional event depends on an agent’s current emotional state. This is implemented in the DER by connecting modules through a system of *dynamic filters*. Sigmoid functions have been chosen for this role because they describe “a large variety of natural phenomena” [Picard, 1997]. The use of sigmoid functions is inspired by the description of the influence of mood on these functions given by Picard [Picard, 1997]. The sigmoid function parameters can be modified depending on the DER model’s state, resulting, for example, in a shifted sigmoid curve as show in Figure 4. This effect simulates the change of sensitivity, e.g. emotion threshold activation, to

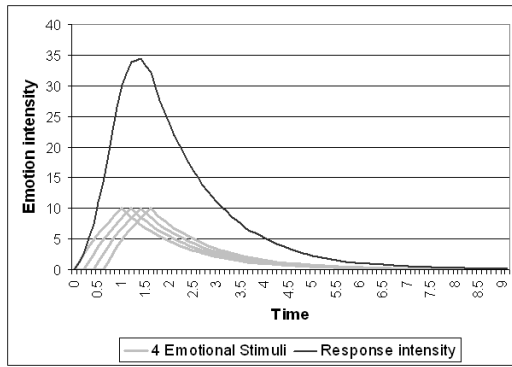


Figure 3: Emotional stimuli are summed to compute the intensity of an emotion.




Figure 4: Sigmoid functions are used as dynamic filters by changing their parameters.

particular emotional stimulus in relation to the current emotional state of a person.

In the DER model, we refer to each module as a *dimension*. A dimension may represent for instance a particular emotion or a particular behaviour activation. Dimensions consists of a list of dynamic filters. One dynamic filter modifies one parameter of a particular type of emotional stimulus, for example its peak intensity. Some characteristics, such as the decay duration, can be modified. In Figure 4, the horizontal axis represents the input value, such as the peak intensity of an emotional stimulus, and the vertical axis represents the output value, such as a new peak intensity value.

Modules influence each other by passing their output to a bus which modifies the parameters of other module's dynamic filters. This bus is the same one conducting the emotional stimuli; the influence on other modules occurs through their 'input' filter. Figure 5 shows states affected by two types of emotional stimulus, *Anger* and *Happiness*, and whether the influence is positive or negative for each type of emotional stimuli. For instance, happy stimuli affect positively the emotional state of *Happiness* and negatively the state of *Anger*. This figure also shows that the emotional state *Happiness* amplifies the effects of happy stimuli and reduces the effects of angry stimuli. In practise, the higher the level of the emotional state *Happiness*, the more the sigmoid func-

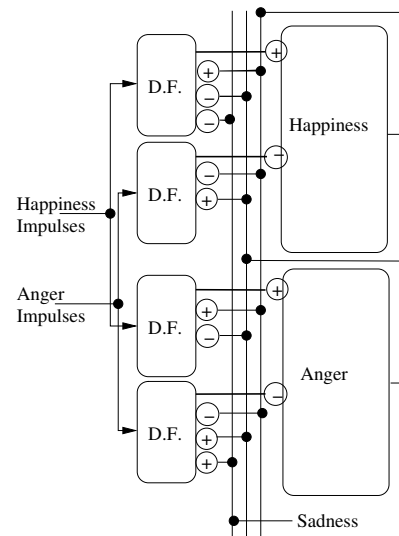


Figure 5: Example of a network of influences between emotional states and emotional impulses in the DER model. D.F. stands for Dynamic Filters.

tions controlling the effects of emotional stimuli anger are shifted to the right. This mechanism decreases the positive effect of anger stimuli on the emotional state anger. This simulates the effects of good and bad news on different moods as demonstrated above.

The system can be configured such that any type of emotional impulse can affect any emotional state and any emotional state can influence the effects of emotional impulses on any emotional state. For the instance of the DER model integrated in the EE-FAS (described below), the tuning has been carried out using heuristic values and visualisation software plotting the resulting sigmoid functions as well as the modified emotional stimuli.

5 An Embedded Instance of a DER

Using the representations just described, we created an instance of a DER and integrated it into the EE-FAS. As mentioned earlier, this DER is composed of three types of modules: *behaviour activations*, *emotions* and *moods*. Each represents persisting states changing on different timescales. Figure 6 shows a graphical representation of this DER.

The behaviour activations are generated by the DER due to emotional impulses, which currently can be either scripted or provided in real time through button presses on a GUI interface. Each behaviour activation is displayed by the EE-FAS as one of the Ekman's emotional facial expressions. However, the duration of an emotional facial expression is different to the duration of the corresponding behaviour activation, they have their own innate time courses. The graphs *a*, *b* and *c* in Figure 1 show the behaviour activations *Happy* and *Angry*. Emotions, such as *Happiness* and *Anger*, also produced by emotional impulses, last longer than behaviour activations. In the EE-FAS, emotions are also used to select facial signals corresponding to communicative functions, e.g. semi-

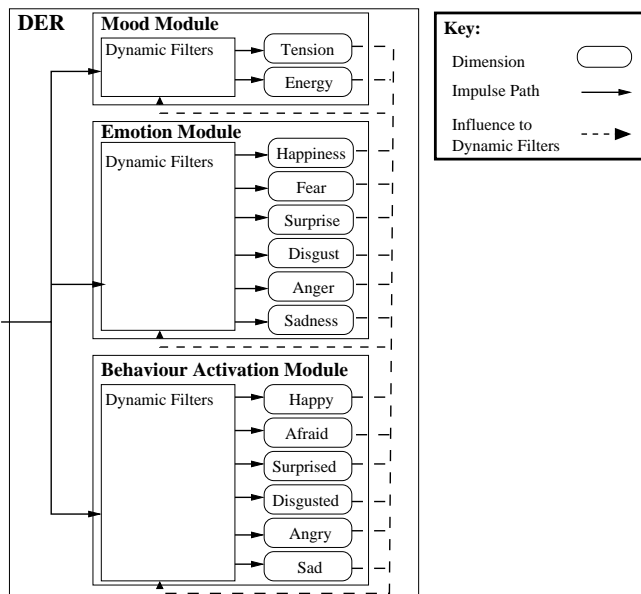


Figure 6: A DER composed of three types of state changing on different timescales: *behaviour activations*, *emotions* and *moods*.

deliberate facial expressions synchronised with the speech, such as emphasis or deliberate smiles. Mood changes on a slower timescale than emotions and it influences and is influenced by the effects of emotional impulses on the DER state. More detail on the design of this DER can be found elsewhere [Tanguy, 2006]. These include details of experiments on the human perception of emotion from components of facial expressions which were conducted using this tool.

6 Discussion and Related Work

A fundamental concept of the DER model is the division of emotion models into mechanisms eliciting emotions and emotion representations. This greatly simplifies scripting or direction for characters, since once the DER is set and an initial mood described, the script or direction can describe communicative actions abstractly rather than describing precise facial expressions. For agents situated in long-term real-time domains such as web pages or interactive environments, this also provides a mechanism for greater variety of behaviour per scripted response, since the mood can depend on the recent interaction history.

Emotion models are heavily researched now; of these a number of researchers have made systems to which the DER is more or less similar. The main advantage of the DER is its modularity, flexibility, ease of configuration and relative autonomy (lack of required direction) once completed. Its main disadvantage is that it takes some time “out of the box” to configure, but this is partially addressed by the fact we distribute it with the EE-FAS instance described earlier. The DER can *a)* represent any number of emotions, *b)* represent emotions or other forms of durative state with different time scale, *c)* define interactions between emotions, and *d)* customise the influences of emotional impulses on each emotion.

The DER is inspired by the various appraisal theories [Lazarus, 1991; Ortony *et al.*, 1988; Plutchik, 1980; Frijda, 1986; Izard, 1993], and the pioneering work of several influential researchers [Picard, 1997; Sloman, 2003]. Many systems use appraisal to influence the decision processes and behaviours of virtual actors [Cañamero, 2003; Gratch and Marsella, 2004; André *et al.*, 1999; Delgado-Mata and Aylett, 2004].

Relatively few systems provide dynamic emotion representations on multiple time scales like the DER, and none of these are as configurable as our system. Paiva *et al.* [Paiva *et al.*, 2004] present an emotion model which assigns a decay function to each emotion elicited with a value higher than a personality threshold. In contrast to the DER, Paiva *et al.*'s work does not implement any interaction between emotions in its emotion representation. Egges *et al.* [Egges *et al.*, 2004] describe a generic emotion and personality representation composed of two types of affective states, moods and emotions. Any number of moods and emotions can be represented. In the implementation of their model the only influences between states is the influence of mood on emotions. Egges *et al.* compute the intensity of affective states by linear functions through the use of matrix operations. In the DER, sigmoid functions are used to control and change the influences of emotional stimuli on emotions and the influences between emotions. This mechanism produce non-linear behaviours closer to natural phenomena.

The Duy Bui does use a decay function to represent the durations of emotions, but the effects of new emotional impulses on an emotion are influenced by the intensity of the other emotions. Their decay functions are also influenced by personality parameters [The Duy Bui, 2004]. Velásquez's representation, the computation of the intensity changes of an emotion takes into consideration the intensity of other emotions, the decay in intensity and the previous intensity of the emotion itself [Velásquez, 1997]. The influences of emotions on others are of the types inhibitory or excitatory. The DER model differs less from these, however it is customisable; any durative state, such as mood, and any number of emotions can be represented; influences of emotional stimuli on emotions can also be defined by the researcher. The main advantage of the DER model is that its representation can be adapted to different emotion theories and to different mechanisms eliciting emotions from the agent's environment. It is a tool that can help the community to model different emotion theories.

Compared to some systems though we have simplified the system slightly to make it more generic. Our definition of emotional impulses carries less information than the emotional structures described by Reilly, which also contain the cause or referent of the emotions [Reilly, 1996]. Reilly and colleagues' interest was primarily in creating ‘believable’ agents, like classic animated cartoon characters which exist to entertain and communicate. We have focussed instead on *realistic* models, which are more useful for experiments and long-term plotless semi-autonomous applications. The DER model focuses on the duration and interaction of emotions, cognitive referents can be tracked in other parts of an agent's intelligence, or the definition of impulse can be expanded.

In our system every emotions decays over time. An emo-

tion such as hope could persist as long as the situation is the same. In fact, we take the position that hope will decay even if the situation stays the same. However new appraisal of the same situation produces new emotional stimuli increasing the level of hope. Similarly, our system could be used to represent drives like hunger. The only requirement is inverting the levels and decay function. A drive slowly increases over time, but then consumatory actions (rather than emotional impulses) such as eating can abruptly reduce its level.

7 Conclusion

This paper has presented a representational system for the relatively persistent action selection state often associated with moods, emotions and drives. The Dynamic Emotion Representation enables a programmer to represent any number of persisting states and the interactions between them, whether excitatory or inhibitory, linear or not. The system integrates these with ordinary temporal decay. We also presented a working system using these representations. In this paper we have emphasised the use of the DER to simplify the programming of realistic real-time emotional agents. The integration of a DER within a virtual character architecture simplifies the requirements for emotion elicitation and the production of varied behaviour. In future work, we hope to extend the system to integrate emotions and drives into the goal structure of artificial life agents. We hope with this to achieve more realistic and perhaps more optimal high-level goal arbitration, while leaving detailed action selection to more simple representations.¹

References

- [André *et al.*, 1999] Elisabeth André, Martin Klesen, Patrick Gebhard, Steve Allen, and Thomas Rist. Integrating models of personality and emotions into lifelike character. In *Proc. of the Workshop on Affect in Interaction.*, pages 136–149, Siena, Italy, October 1999.
- [Badler *et al.*, 2002] Norman Badler, Jan Allbeck, Liwei Zhao, and Meeran Byun. Representing and parameterizing agent behaviors. In *Proc. Computer Animation*, pages 133–143, Geneva, Switzerland, 2002. IEEE Computer Society June 2002.
- [Cañamero, 2003] Dolores Cañamero. Designing emotions for activity selection in autonomous agents. In Robert Trappl, Paolo Petta, and Sabine Payr, editors, *Emotions in Humans and Artifacts*, pages 115–148. MIT Press, 2003.
- [Delgado-Mata and Aylett, 2004] Carlos Delgado-Mata and Ruth S. Aylett. Emotion and action selection: regulating the collective behaviour of agents in virtual environments. In *Proc. 3rd International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 1304–1305, Washington, DC, USA, 2004. IEEE Computer Society.
- [Egges *et al.*, 2004] Arjan Egges, Sumedha Kshirsagar, and Nadia Magnenat-Thalmann. Generic personality and emotion simulation for conversational agents. *Computer Animation and Virtual Worlds*, 15:1–13, January 2004.
- [Frijda, 1986] Nico H. Frijda. *The Emotions*. Cambridge University Press, 1986.
- [Gratch and Marsella, 2004] Jonathan Gratch and Stacy Marsella. Evaluating the modeling and use of emotion in virtual humans. In *Proc. 3rd International Joint Conference on Autonomous Agents and Multiagent Systems*, vol. 1, pages 320–327. IEEE Computer Society, 2004.
- [Izard, 1993] Carroll E Izard. Four systems for emotion activation: Cognitive and noncognitive processes. *Psychological Review*, 100(1):68–90, 1993.
- [Lazarus, 1991] Richard S Lazarus. *Emotion and adaptation*. Oxford University Press, 1991.
- [LeDoux, 1996] Joseph LeDoux. *The Emotional Brain: The Mysterious Underpinnings of Emotional Life*. Simon and Schuster, New York, 1996.
- [Maes, 1991] Pattie Maes. The agent network architecture (ANA). *SIGART Bulletin*, 2(4):115–120, 1991.
- [Ortony *et al.*, 1988] Andrew Ortony, Gerald L. Clore, and Allan Collins. *The cognitive structure of emotions*. Cambridge University Press, 1990/1988.
- [Paiva *et al.*, 2004] Ana Paiva, João Dias, Daniel Sobral, Ruth Aylett, Polly Sobreperez, Sarah Woods, Carsten Zoll, and Lynne Hall. Caring for agents and agents that care: Building empathic relations with synthetic agents. In *Proc. of the 3rd International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 194–201, Washington, DC, USA, 2004. IEEE Computer Society.
- [Picard, 1997] Rosaline W. Picard. *Affective Computing*. MIT Press, Cambridge, Massachusetts, 1997.
- [Plutchik, 1980] Robert Plutchik. A general psychoevolutionary theory of emotion. In Robert Plutchik and Henry Kellerman, editors, *Emotion: Theory, research, and experience*, pages 3–33. Academic press, inc, 1980.
- [Reilly, 1996] W. Scott Neal Reilly. *Believable social and emotional agents*. PhD thesis, School of Computer Science, Carnegie Mellon University, Pittsburgh, 1996.
- [Sloman, 2003] Aaron Sloman. How many separately evolved emotional beasts live within us? In *Emotions in Humans and Artifacts*, pages 35–114. MIT Press, 2003.
- [Tanguy, 2006] Emmanuel Adrien Raymond Tanguy. *Emotions: the Art of Communication Applied to Virtual Actors*. PhD thesis, University of Bath, June 2006. University of Bath, Technical Report CSBU-2006-06.
- [The Duy Bui, 2004] The Duy Bui. *Creating emotions and facial expressions for embodied agents*. PhD thesis, University of Twente, 2004.
- [Velásquez, 1997] Juan D. Velásquez. Modeling emotions and other motivations in synthetic agents. In *Proc. of the Fourteenth National Conference on Artificial Intelligence*, pages 10–15, 1997.

¹This work was supported by a studentship from the Department of Computer Science, University of Bath, and the UK ESPRC AIBACS initiative, grant GR/S/79299/01.