FAtiMA Modular: Towards an Agent Architecture with a Generic Appraisal Framework

João Dias, Samuel Mascarenhas, Ana Paiva

INESC-ID and Instituto Superior Técnico, Technical University of Lisbon
Tagus Park, Av. Prof. Cavaco Silva, 2780-990 Porto Salvo, Portugal
Email: joao.dias@gaips.inesc-id.pt, samuel.mascarenhas@gaips.inesc-id.pt, ana.paiva@inesc-id.pt

Abstract. This paper presents a generic and flexible architecture for emotional agents, with what we consider to be the minimum set of functionalities that allows us to implement and compare different appraisal theories in a given scenario. FAtiMA Modular, the architecture proposed is composed of a core algorithm and by a set of components that add particular functionality (either in terms of appraisal or behaviour) to the architecture, which makes the architecture more flexible and easier to extend.

1 Introduction

FAtiMA (**F**earnot **A**ffec**TI**ve **M**ind **A**rchitecture) is an Agent Architecture with planning capabilities designed to use emotions and personality to influence the agent's behaviour[3]. During the last years, the architecture was used in several scenarios (such as FearNot![10], ORIENT[14], and a process Model of Empathy[12]) and by different research institutions, which led to the architecture being extended with several new features and functionalities (e.g. Cultural Behaviour[8] and Drives[5]). The resulting architecture has become perhaps difficult to use since the complexity of understanding the architecture escalated with the number of existing features. For this reason, a major effort was put in creating a modular version of the architecture, where functionalities and processes are divided into modular independent components. This enables us to use lighter and simpler versions of FAtiMA with just some of the components (for instance only the reactive and emotional components) for the simpler scenarios. During the refactorization, the most relevant processes were generalized making the architecture easier to extend by allowing us to work independently on different components whose functionality can be easily added to the architecture.

One of the most relevant processes generalized in the architecture was the appraisal process. The rationale followed was that the architecture should be able to later incorporate several distinct appraisal mechanisms and even appraisal theories. There are currently several different appraisal theories that model the process of emotion generation (OCC[9], Roseman's[13], Scherer[15]). Due to the difficulty of implementing them all in a same scenario, there has not been an effort (at least to our knowledge) in directly comparing and evaluating the different theories together. We believe that creating scenarios where the several emotion theories can be integrated and evaluated can help us to learn more about them. For instance, it would be interesting to determine if Scherer's appraisal theory, where facial expressions derive directly from the stimulus evaluation checks instead of the resulting emotions, can indeed generate more natural and believable emotional expressions.

The aim of this paper is thus to present a generic core architecture for emotional agents, and to describe how the generic appraisal mechanism could be used to implement different appraisal theories. This is a first step towards a contribution to the standards in emotion modeling.

2 FAtiMA Core

FAtiMA Modular is composed of a core layer (named FAtiMA Core) on which components are added in order to add functionality. FAtiMA Core is a template that generally defines how the Agent Architecture works. Added components can provide specific implementations for the generic functions defined in the Core. Figure 1 shows a diagram of FAtiMA Core with the basic functionalities for an emotional agent architecture. An agent is able to receive perceptions from the environment (events) which are used to update the agent's memory (or internal state) and to trigger the appraisal process. The result of the appraisal process is stored in the affective state[1], and later used to influence the action selection processes which will make the agent act upon the environment.

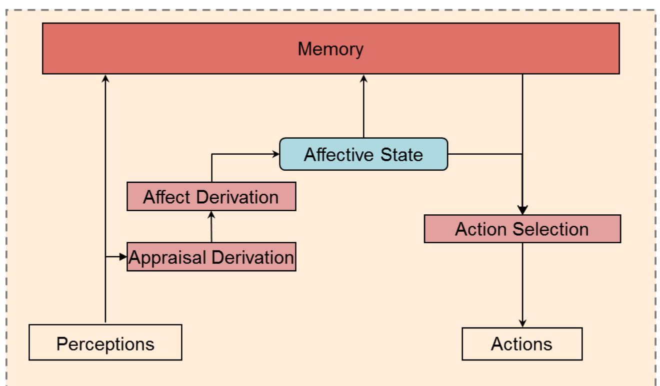

Fig. 1. FAtiMA Core Architecture

The Core architecture is strongly based in [7] and thus it divides appraisal into two separate processes. The first one, appraisal derivation, is responsible for evaluating the relevance of the event to the agent and determine a set of appraisal variables (e.g. desirability and desirability for others in the case of OCC). The second process, affect derivation, takes the appraisal variables as input and generates the resulting affective states (emotions or mood) according to an Appraisal Theory. This division into distinct

[1] We call it affective state since it is used to store affective states such as emotions and moods.

processes is also in accordance with the formal characterization of Structural Appraisal Theories proposed by Reisenzein[11] and with the Set based Formalism proposed by Broekens[2], which identify three main processes in an Appraisal Theory: perception, Appraisal and Mediation (Appraisal Derivation and Affect Derivation in our terminology).

It is important to point out that the Core architecture does not commit itself with the particular methods used. In fact a FAtiMA agent that only has a Core will not do anything. Behaviour is added by adding components that implement the mentioned functionality. However, a component is not required to implement all functionality in the Core, it can implement just one of the processes. In order to differentiate components when adding them to the Core, they are categorized according to the implemented functionality. For instance, an AffectDerivation Component will have to implement an Affect Derivation process, and a Behaviour Component will have to implement the action selection function. All components are designed following two main properties: they must be interchangeable - i.e, being able to be replaced, added or removed with a minimum effort; and they must be loosely coupled - dependencies between components should be avoided unless strictly needed.

2.1 Appraisal Process

We will now look with more detail to the Appraisal Process. One of the main goals when designing the appraisal mechanism was that it would have to be powerful and flexible enough to represent most Appraisal Theories. Given that Scherer's [15] theory is one of the most complex Appraisal Theories, an effort has been made to make the appraisal mechanism compliant with it. Other appraisal theories can be easily implemented by modeling simpler single-component processes of appraisal. Scherer defines appraisal as a process of multilevel sequential checking: a set of evaluation checks are performed in sequence in order to assess the relevance of an event, the implications for self of the same event, coping potential and finally normative significance. Appraisal is sequential because some of the appraisal evaluation checks require others to complete. According to Scherer, appraisal is also done at several levels, a sensory-motor level, a schematic level and conceptual level. The lower levels are usually faster to determine a stimulus evaluation check, but often may offer less than perfect appraisal information. In situations where low-level appraisals cannot properly evaluate a situation, higher level (but more complex and heavier) appraisals take place.

Scherer's notion of evaluation check can easily be aggregated with other theories concept of appraisal variables, so we will use the term appraisal variable henceforth. Furthermore, in order to make the appraisal process as flexible as possible, FAtiMA appraisal is grounded on a set of design principles, which try to generalize Scherer's requirements for its appraisal theory:

1. Appraisal is incremental, i.e. appraisal is not a one-shot process and different appraisal variables and the corresponding emotions may be generated at different points in time. For instance, in the case of OCC, desirability (associated with Joy and Distress) can be determined first while desirability for other could be determined later (associated with Fortune of Others emotions such as Pitty). Conse-

quently, the appraisal process must be continuously executed and is up to each component to decide if and when to return an appraisal variable.

2. An appraisal component may depend on an appraisal variable determined by another component and thus need to access all the appraisal variables generated in the appraisal process. This implies that we need a way to store the result of an appraisal component (a structure akin to SEC registers) and use it as an input to other appraisal components. Furthermore, these first two principles allow us to model sequential evaluation checks (or sequential appraisals). By making a component start the appraisal only after an appraisal variable is defined (or defined with a particular value), we create an implicit dependency and ordering between the components. In order to favor modularity, only the dependent component needs to know the dependency, if we remove it from the architecture the initial component will continue working as normal.
3. Lazarus pointed out that appraisal is often followed by reappraisal, with the purpose to correct the evaluation based on new information or more thorough processing[4]. Therefore, the value of an appraisal variable may change over time, and the affective state must reflect this change. For example, an event may be initially appraised as undesirable by one fast reactive appraisal component, but later considered desirable by a more elaborate cognitive appraisal component. Thus considering an OCC affect derivation component as example, the initial Distress emotion that was triggered by the event should change into a Joy emotion. This means that the Core architecture must enforce that the affect derivation processes are executed whenever there is a new result returned from the appraisal components.
4. Any component can contribute to any appraisal variable. For this reason, we must store the contribution of all components to a given appraisal variable. In order to determine the final value of an appraisal variable, an explicit policy is used. The policy can be choosing the highest value, performing a weighted sum, using the value determined by the latest component to set the variable, or using a priority mechanism. These last two principles enumerated enable us to model Scherer's multi-level appraisal by creating one component for each level. A reactive component can be used to implement Scherer's sensory-motor level (or schematic level) for a given appraisal variable, while a deliberative component can implement the conceptual level for the same variable.

In order to store the intermediate results of the appraisal variables generated, and to allow any appraisal derivation component to access that information, an Appraisal Frame structure is used. An Appraisal Frame is usually associated to an event (internal or external[2]) and stores a list of appraisal variables associated with the event. As mentioned, the Appraisal Frame must store the contribution of all components to an appraisal variable and use a policy to determine its final value. Figure 2 presents a diagram of the appraisal process. An event generated will trigger the start of an appraisal process creating an appraisal frame. The arrow from the appraisal frame to the appraisal components represents that appraisal can take several cycles and depend on the appraisal of

[2] External events correspond to events that happen in the environment, such as John pushing Luke, while internal events correspond to events that are triggered by architecture's internal processes, for instance the activation, success or failure of a given goal.

other components. In the example provided, one of the components initially defines the appraisal variable with a positive value (4), but another component later defines it as being slightly negative (-1).

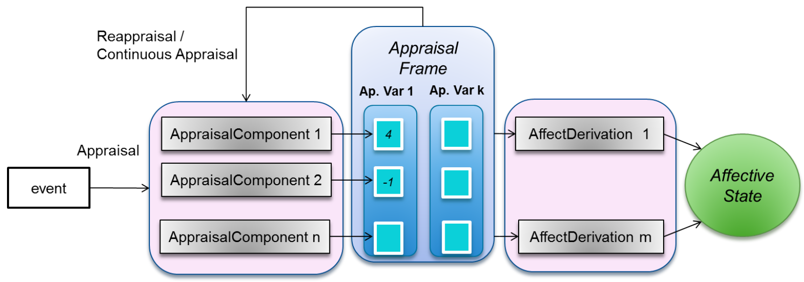

Fig. 2. Appraisal Process

Whenever there is a change in an Appraisal Frame, the emotional state of the agent needs to reflect that change. This is handled by the Affect Derivation components, whose sole responsibility is to generate and update the agent's emotions. Emotions are defined as having a type, valence and an intensity which decays with time. Each affect derivation component is informed of changes in the appraisal variables stored in the appraisal frame. This information is then used to decide which emotions to create and their corresponding intensity. Additional modulation factors, such as a personality or cultural bias, might also be used by the component to determine the emotion intensity. Note that affect derivation components are independent from each other and their number is not restricted. As such, several emotions may be added simultaneously to the agent's emotional state, which then results in the agent experiencing mixed emotions. Besides emotions, the agent's affective state also integrates the notion of mood. Mood represents an overall affective state which is influenced by the emotions experienced by the agent (and also decays with time): positive emotions increase mood, while negative emotions decrease it.

One particular affect derivation component that is already implemented in the architecture is the OCC Affect derivation component that, as the name implies, is based on the OCC theory. Whenever this component is informed of a change in any OCC appraisal variables such as desirability, desirabilityForOther, or praiseworthiness, it will generate the corresponding emotion (e.g. given a positive desirability value, a Joy emotion is generated). This component also looks for the agent's predefined emotional thresholds to dampen the final emotion intensity. These emotional thresholds are a mechanism to model the agent's personality by making it harder/easier to experience certain emotions.

One last important mechanism is the way that emotions are stored and organized in the affective state. The third principle requires us to properly update emotions when the value of an appraisal variable changes. This is particularly important because a change in an appraisal variable may even change the type of an emotion generated. To do so, emotions are indexed in the affective state by the set of appraisal variables used to generate them as well as the corresponding event. When a new emotion is added to the emotional state, the emotional state checks if any emotion caused by the same event and triggered by exactly the same appraisal variables already exists. If it does, then it will replace the existing emotion by the new one. If not, it will simply add the new emotion. Imagine an event that is initially appraised with a negative desirability, according to OCC this will generate a distress emotion. However, if the event is later appraised (by another component or by reappraisal) with positive desirability, the distress emotion will be removed and replaced by a new joy emotion. Note that with this mechanism, it is possible to generate several distinct emotions associated with the same event as long as they are triggered by a different set of variables.

Figure 3 presents the resulting pseudocode for FAtiMA Core. The first update function is used to update the components every cycle (used for instance to simulate processes of decay). Then the perceptual process checks if there is a new perception taking place in the environment. If a new event is perceived, it is used to update the agent's memory and all existing components and a new Appraisal Frame will be created associated to the event. The Appraisal Frame will then be used to start the appraisal process.

```
while(shutdown != true)
  for each Component c
    c.update();

  e <- perceiveEvent();

  if(a new event is perceived)
    memory.update(e);
    for each Component c
      c.update(e);

    aF <- newAppraisalFrame(e);

    for each AppraisalComponent aC
      aC.startAppraisal(e,aF);
      updateEmotions(aF);

  for each AppraisalComponent aC
    aF <- aC.continuousAppraisal();
    updateEmotions(aF);

  for each BehaviorComponent bC
    bc.actionSelection();

  a <- selectAction();
  executeAction(a);
```

Fig. 3. FAtiMA's Core Pseudocode

Once the Appraisal Frame is determined, the function updateEmotions will be used to initiate the AffectDerivation process. According to figure 4, whenever there is a change detected in the Appraisal Frame (meaning that one of the components has set at least one appraisal variables), all existing appraisal derivation components are used to generate the corresponding emotions, which are added to the emotional state. Finally, at the end of an agent's cycle, the Core will ask the behaviour components if they want to perform any action. This can be a simple reactive mechanism or a complex deliberative one. As example, FAtiMA reactive component is a behaviour component that will trigger simple action tendencies based on the current emotional state, while the deliberative component uses goals and planning to determine the best course of action to follow. If more than one action is triggered by different components, a priority mechanism will be used to select the most relevant one. The selected action will then be executed in the environment.

```
updateEmotions(aF)

  if(aF.hasChanged())
    for each AffectDerivationComponent aD
      emotion <- aD.affectDerivation(aF)
      AffectiveState.add(emotion);
```

Fig. 4. Update Emotions method

3 FAtiMA Modular

As previously mentioned, a FAtiMA agent that only has a Core will not do anything. FAtiMA Modular architecture is created by adding a set of components to the core. For instance, in the scenario described in [8], the architecture was defined with the following components:

- **Reactive Component -** this component uses predefined emotional reaction rules to determine the value of the following OCC appraisal variables: Desirability, DesirabilityForOthers, Praiseworthiness and Like. When an event is perceived, the reactive appraisal matches the event against a set of emotional rules. A rule may define a particular value for each of the appraisal variables and can then target a specific event (e.g. the agent finds it desirable whenever it receives a compliment from agent B) or it can be more general (e.g. the agent finds it undesirable whenever the action cry is performed).
- **Deliberative Component -** handles goal-based behavior and adds planning capabilities to the agent. It uses the state of plans in memory to generate appraisal variables for OCC Prospect Based Emotions. These appraisal variables are GoalStatus, GoalConduciveness, and GoalSuccessProbability (see [3]) for more details on the deliberative and reactive components).
- **OCCAffectDerivation Component -** generates emotions from the appraisal variables according to the OCC Theory of Emotions (see Table 1). For instance an event with a positive desirability value for the agent will generate a Joy emotion

if it surpasses the agent's predefined threshold for Joy. On the other hand, if the event's desirability is negative, then a Distress emotion is generated instead.

– **Motivational Component -** component that models basic human drives, such as energy and integrity and uses them to help select between competing goals in the deliberative component. The more a certain need is low/high, the more higher/lower the utility of a goal that contributes positively for that need is. Additionally it is also used to determine an event's desirability according to the effects it had on the agent's drives (e.g. an *eat* action lowers the energy need). When an event lowers/raises the agent's needs, it is evaluated as desirable/undesirable for that agent (for more details please refer to [5]).
– **Theory of Mind Component -** creates a model of the internal states of other agents. This component determines the desirability of an event for others by simulating their own appraisal processes.
– **Cultural Component -** implements cultural-dependent behaviour of agents through the use of rituals, symbols and cultural dimensions. It it also used to automatically determine the Praiseworthiness appraisal variable based on cultural values and also on the impact actions have on the motivational states of the agents. For instance, the more collectivistic the agent's culture is, the more praiseworthy is an action that positively affects the need of other agents in detriment of the agent's own needs.

Table 1. Association between implemented OCC appraisal variables and OCC emotion types

Appraisal Variables	Associated Emotion Types
Desirability	Joy, Distress
Desirability, DesirabilityForOthers	HappyFor, Gloating, Pitty, Resentment
Praiseworthiness	Pride, Admiration, Shame, Reproach
Praiseworthiness, Desirability	Gratification, Gratitude, Remorse, Anger
Like	Love, Hate
GoalStatus, GoalConduciveness, GoalSucessProbability	Hope, Fear, Relief, Satisfaction, FearsConfirmed, Disappointment

To illustrate how these several components are combined in the scenario to create a complex appraisal process consider the following example. During a dinner party enacted by five different agents, one of the agents informs the others that he is feeling sick. When the agent's culture is defined as highly collectivistic, there is an agent that decides to offer his own medicine. After receiving the medicine, the agent that is sick elicits a Joy and Admiration emotions.

The appraisal process that takes place for the receiving agent occurs in the following manner. First, the motivational state component will quickly determine that receiving the medicine will have a positive effect on the agent's integrity drive (which is low at the moment) generating a strong positive desirability, and store that value in the Appraisal

Frame. FAtiMA core will detect a change in the Appraisal Frame and initiate the AffectDerivation process by calling the OCCAffectDerivation Component. This leads to the initial creation of a Joy emotion.

A few cycles later, the Theory of Mind component will also determine that the same event will have a negative effect on the giver agent's integrity drive (since it lost its medicine), determining the same event as undesirable for the giver. AffectDerivation will be called again, but this time no emotion is generated. Since the cultural component requires the two others appraisal variables to generate the appraisal, it will not trigger until both values are determined by the other components. Once this happens, the cultural component determines the praiseworthiness variable according to the formula presented in [8]. In this example, given the collectivistic culture of the appraising agent, he will appraise the event as highly praiseworthy[3]. The OCCAffectDerivation component will then generate an Admiration emotion based on praiseworthiness.

4 Related Work

There is some related work relevant to the generic appraisal model presented here. The FeelME framework proposed by Broekens[1], similarly to the work presented here, addresses the problem of creating a computational emotional system in an modular and extensible way. In the framework, appraisal processes are separated into five main steps: the Decision Support System that corresponds to a perceptual system; the Appraisal System (AS) that evaluates perceived objects and returns a vector of appraisal dimensions; the Appraisal Signal Modulator that can change (for instance amplifying or dampening) the appraisal result vectors; an Emotion Maintenance System responsible for integrating the appraisal-results and maintaining the emotional state, which is also represented as a n-th dimensional vector; and finally the Behaviour Modification System that controls the agent's emotional behaviour. Modularity and scalability is attained by making the AS to be composed of a set of appraisal banks that evaluate specific aspects of the agent's environment. Although the banks are assumed to evaluate independently, an appraisal bank can influence the contribution of another appraisal bank to the EMS through dependencies, thus allowing the framework for instance to model Scherer's levels of appraisal and evaluation sequence.

Comparing the FeelME framework with the architecture proposed shows us that there are some similarities. For instance, the way that the appraisal system if composed of independent appraisal banks is similar to our approach of using a set of differentiated appraisal derivation components. However, although computationally flexible, modular and efficient, the framework does not model affect derivation process and uses a continuous representation of emotions directly obtained from the sum of the resulting appraisals. Although a continuous model of emotion per se is not necessarily a disadvantage, the fact that the appraisals of distinct events are combined together may originate a loss of information. As example, if two events e1 and e2 evaluated with opposing appraisal values are perceived close together, the resulting emotional state will be neutral instead of an emotional state with equally strong but opposing emotions.

[3] A collectivistic culture values self-sacrifice for the well being of others.

With the goal of facilitating the comparison and integration of different Cognitive Appraisal Theories and systematically analyzing computational models of emotion, Broekens et al.[2] have put forward a theory-independent formalization that can be used to describe the structure of appraisal. The formalism proposed is based on set theory and models appraisal with the three types of processes identified by Reisenzein[11]: perception processes (P), appraisal processes (A) and mediating processes (M). Perception processes map the external world to mental objects, appraisal processes evaluate perceived mental objects to appraisal dimensions and mediation processes relate appraisal dimensions values to emotion components that can be used to represent emotion categories and intensities. Broekens et al.'s formalization models explicit inhibitory and excitatory dependencies between processes (of perception, appraisal and mediation) through guards and links. It is important to point out that dependencies can be created between processes of the same type (e.g. the appraisal of novelty with high value activates the appraisal of relevance) or processes of different types (the presence of an object returned by a perception process activates an appraisal process).

The formalization proposed by Broekens et al. is more complete than the FeelME framework. It can represent explicit Mediation processes and Emotion Components or categories similarly to our model. As such, there are even more resemblances between this formalization and the model proposed by us. Nonetheless, there are some relevant differences. While in Broekens formalization dependencies are created between processes, in FAtiMA-Core appraisal components depend on appraisal variables and not on other components. Ideally, an appraisal component does not need to know what other appraisal components exist. We argue that this is important in order to make the architecture more flexible and modular. Since several components may generate the same appraisal variable, a third component that depends on that variable will still work even if some of the original components are removed. An additional relevant difference is that the proposed semi-formal model was designed with the goal of being easily mapped to a computational model, while the set-based formalization aims at a systematic formal analysis of particular computational models.

The last model analysed, EMA [6], is a very complete computational model of appraisal processes. It focus on explaining both the rapid dynamics of emotional reactions and slower deliberative responses. However, in contrast to multiple-level approaches that use different levels for the reactive and deliberative appraisals, authors argue that a single and automatic appraisal process can be used to generate similar dynamics that result from deliberative and reactive processes operating on a person's relationship with the environment. Moreover, authors argue that appraisal checks occur in parallel and any sequential relation between appraisal is due to the requirements of the cognitive processes involved in the construction of the representation of the appraised event. EMA consists of a set of processes that interpret a representation of the person environment relationship in terms of a set of appraisal variables and a set of coping processes that manipulate this representation in response to the appraisals. To do so, EMA uses a representation built on the causal representations created by decision-theoretic planning, which are capable of easily capturing concepts such as utility and probability that can be directly translated to appraisal variables such as desirability and likelihood. These appraisal variables, which are continuously updated, are stored in a data structure called

Appraisal Frame that is created for each proposition inside the causal interpretation. In EMA, emotions correspond to the existing appraisal frames. Although there is mapping between the appraisal variables and emotion labels, it is mainly used to facilitate the facial expression of emotions. The model features a interesting cyclic relationship between appraisal, coping and reappraisal. The initial appraisal of a situation triggers several cognitive and behavioral responses that changes the person's relationship with the environment. These changes will lead to a reappraisal of the initial situation, which may eventually result in additional responses.

While the focus of EMA is on creating a single-level process model of appraisal that can explain the dynamics of emotions, the goal of the structural architecture proposed here is to be able to model several distinct theories of emotions, be it single-level or multi-level sequential theories. In fact, while most of the paper focused on being able to model a multi-level sequential appraisal theory, the presented generic model is currently easily integrated with a single-level appraisal theory. As example, the Deliberative Component, which was not mentioned in detail in this paper has an appraisal mechanism where events are used to update planning structures, which will generate appraisals of goal conduciveness and goal probability, albeit this is done at a much simpler level than EMA. We do not argue which approach is the most correct. Instead we agree with Broekens in that these generic structural models provide an interesting tool to analyse, compare and evaluate distinct appraisal theories.

5 Conclusions and Future Work

This paper proposes a model for an Emotional Agent Architecture that can be easily extended by adding new components that define a set of functionalities used by a core generic algorithm. One of the processes that can be extended is the appraisal process. An effort has been made to make the appraisal mechanism as flexible and dynamic as possible in order to support the implementation of distinct appraisal theories. With the proposed architecture it is possible to create scenarios where several emotion theories can be properly compared. This is in our opinion a first step towards an important contribution to the creation of standards in emotion modeling.

In the future we pretend to implement Scherer's Appraisal Theory and create a scenario where one can directly compare the emotional behaviour of an agent with OCC Appraisal Theory to the one of an agent with Scherer's Theory. We are particularly interested in determining whether the more dynamic appraisal process of Scherer's Theory, which allows an appraisal to change from an initial negative towards a positive one, can in fact generate emotional expressions and behaviour that is perceived as more natural and believable by users.

6 Acknowledgments

This work was partially supported by European Community (EC) and by two scholarships (SFRH BD/19481/2004 and SFRH BD/62174/2009) granted by the Fundação para a Ciência e a Tecnologia (FCT). It is currently funded by the EU FP7 ICT-215554 project LIREC (LIving with Robots and IntEractive Companions). The authors are solely responsible for the content of this publication. It does not represent the opinion of the EC or the FCT, which are not responsible for any use that might be made of data appearing therein.

References

1. J. Broekens and D. DeGroot. Scalable and flexible appraisal models for virtual agents. In *5th Game-On International Conference: Computer Games: Artificial Intelligence, Design and Education.*, 2004.
2. J. Broekens, D. DeGroot, and W. Kosters. Formal models of appraisal: Theory, specification, and computational model. *Cognitive Systems Research*, 9(3):173–197, 2008.
3. J. Dias and A. Paiva. Feeling and reasoning: a computational model for emotional agents. In *Proceedings of 12th Portuguese Conference on Artificial Intelligence, EPIA 2005*, pages 127–140. Springer, 2005.
4. R. S. Lazarus. *Psychological stress and the coping processes*. McGraw Hill, New York, 1966.
5. M. Lim, J. Dias, R. Aylett, and A. Paiva. Creating adaptive affective autonomous npcs. *Special Issue Journal of Autonomous Agents and Multi-Agent Systems (to appear)*, 2011.
6. S. Marsella and J. Gratch. Ema: A process model of appraisal dynamics. *Cognitive Systems Research*, 10(1):70–90, 2009.
7. S. Marsella, J. Gratch, and P. Petta. *A blueprint for an affectively competent agent: Cross-fertilization between Emotion Psychology, Affective Neuroscience, and Affective Computing.*, chapter Computational Models of Emotion. Oxford University Press, 2010.
8. S. Mascarenhas, J. Dias, R. Prada, and A. Paiva. A dimensional model for cultural behaviour in virtual agents. *Applied Artificial Intelligence*, 24(6):552–574, 2010.
9. A. Ortony, G. Clore, and A. Collins. *The Cognitive Structure of Emotions*. Cambridge University Press, UK, 1998.
10. A. Paiva, J. Dias, D. Sobral, R. Aylett, S. Woods, L. Hall, and C. Zoll. Learning by feeling: Evoking empathy with synthetic characters. *Applied Artificial Intelligence*, 19(3):235–266, 2005.
11. R. Reisenzein. *Appraisal Processes in Emotion. Theory, Methods, Research*, chapter Appraisal Processes Conceptualized from a Schema-Theoretic Perspective: Contributions to a Process Analysis of Emotions. Oxford University Press, 2001.
12. S. Rodrigues, S. Mascarenhas, J. Dias, and A. Paiva. I can feel it too! emergent empathic reactions between synthetic characters. In *Affective Computing and Intelligent Interaction and Workshops, 2009. ACII 2009. 3rd International Conference on*, pages 1 –7. IEEE, 2009.
13. I. Roseman, P. Jose, and M. Spindel. Appraisals of emotion-eliciting events: testing a theory of discrete emotions. *Personality and Social Psychology*, 59(5):899–915, 1990.
14. A. Ruth, V. Natalie, A. Elisabeth, A. Paiva, S. Enz, and L. Hall. But that was in another country: agents and intercultural empathy. In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems - Volume 1*, AAMAS '09, pages 329–336. International Foundation for Autonomous Agents and Multiagent Systems, 2009.
15. K. Scherer, A. Schoor, and T. Johnstone. *Appraisal Processes in Emotion. Theory, Methods, Research*. Oxford University Press, 2001.