Consistent Dynamic-Group Emotions for Virtual Agents

Joost Broekens, Niels Netten, Doug DeGroot

Leiden Institute of Advanced Computer Science, Leiden University, Leiden, The Netherlands, {broekens, cnetten, degroot}@liacs.nl

Abstract

The use of computational models of emotion in virtual agents enhances the realism of these agents in a variety of domains, including virtual reality training and entertainment computing. We consider these two domains as prototypical for Multi-emotional-Agent-Systems (MeASs), which are the focus of this paper. MeASs typically include groups of agents organised into clusters, for example a special-force unit. While each agent in such a group has its own emotional model, resulting in realistic individual emotional behaviour, the group as a whole can show unrealistic emotional behaviour. Currently there is no method to enforce emotional consistency of a cluster of agents while allowing agents to have individual emotions. Our approach introduces an emotional-state component that is a separate step in the computational model of emotion used by individual agents. The introduction of this emotional-state component enables multiple architectures for group emotions. We evaluate these architectures and conclude that several enable consistent integration of individual emotions and group emotions. We believe that our research enables agent- and scenario designers to benefit from the individual realism a computational model of emotion brings to virtual agents, without losing group consistency. Furthermore, by choosing one architecture versus another, designers can trade-off quality of the group emotion for computational performance.

1. Introduction

In cognitive psychology emotion is often defined as a psychological state or process that functions in the management of goals. This state consists of physiological changes, feelings, expressive behaviour and inclinations to act. Emotion is elicited by the evaluation of an event as positive or negative for the accomplishment of the agent's goals. According to this definition, emotion is a heuristic that relates events to the agent's goals [5]. Additionally, emotions are used in non-verbal communication. Computational models of emotion are embedded in virtual agents in a variety of domains, including:

- Entertainment computing; emotions are embedded in the non-player-characters for entertainment and realism purposes. The communicative aspect of emotional expression is used to create a sense of realism [2].
- Virtual-reality training environments; emotions are embedded in agents primarily to enhance a trainee's sense of realism during the training, using emotional expression and the interplay between emotions and plans [3].

We consider these two examples as prototypical for Multi-emotional-Agent-Systems (MeASs). MeASs typically include dynamic groups of agents organised into clusters, for example a special-forces unit. While each agent in such a group has its own emotional model, potentially resulting in realistic individual emotions, the group as a whole can show unrealistic emotional behaviour. For example, all members of a special-forces unity are afraid of a coming attack except one, who stands calmly with a big smile on his face. Currently there is no method to enforce emotional consistency of a cluster of agents while allowing agents to have individual emotions. In this paper, we propose a novel architectural approach to do this.

Our approach, like the majority of computational models of emotion that are embedded into virtual agents (many are based on the work by [6]), shares the appraisal theoretic assumption that the subjective evaluation of the environment, in relation to the agent's goals, is responsible for emotion [7]. It is this goal-related evaluation in terms of variables (e.g. novelty, pleasure and expectancy) that makes appraisal theories popular as basis for computational models of emotions in virtual agents.

In our approach, an Emotional Maintenance System that maintains the emotional-state of an individual agent integrates appraisal-results. This system is a separate step in the computational model of emotion. We evaluate, on paper, eight different architectures for group emotions that are made possible by the introduction of this system. Our evaluation is based on the criteria: computational performance, quality of the group emotion and design effort. In this paper we show that architectures based on our approach enables consistent integration of individual emotions and group emotions.

2. Computational emotions enhance realism.

In MeASs emotions are embedded in agents primarily to enhance a trainee's sense of realism, using emotional expression and the interplay between emotions and plans [3]. Reasons to use computational emotions include the following [4]. First, the ability to model virtual agents with emotional expressions provides richer and more variable scenarios. This would not be possible without the use of emotions. Second, emotional expressions are a key factor in achieving believable behaviour of virtual agents. Third, emotions provide more engaging VR training experiences, which might contribute to better memory retention and learning experience of a trainee as a result of VR training.

MeASs typically include dynamic groups of virtual agents organised into clusters, for example a special-forces unit. While each agent in such a group has its own emotional model, potentially resulting in realistic individual emotions, the group as a whole can show inconsistent emotional behaviour. This poses a problem for - for example - scenario designers who need groups of emotional agents to be consistent when they design a scenario to achieve a specific training experience. Problems related to the consistency of group emotions are discussed next.

3. Problems with Individual Emotions and groups.

The 'Titanic problem'. In a group of emotional agents that - for e.g. a trainee - should have individuals sharing the same emotion, one agent with an emotion that drastically

diverges from the this emotion may look unrealistic. For example, one agent is laughing while all others are crying and trying to get through a closed gate while the boat is slowly flooding. This laughing agent poses a serious thread for the trainee's suspension of disbelief. We call this the 'Titanic problem'.

The leader-selection problem. In an interactive simulation environment, groups of agents can be formed dynamically using an arbitrary grouping mechanism. Consequently, designers may lack prior knowledge of the amount and the type of agents in a group as well as the group's size. This is especially important for emotions based on appraisal of the environment. If the emotions of all individuals in a group are based upon the appraisal of a group-leader, and the group-leader is positioned on one side of the group, then agents on the other side of the group will react emotionally to the environment of their leader instead of to their own direct environment. This may result in unrealistic individual emotions, typically non-response to an emotional event. Grouping mechanisms as well as the possible constellations of the group should not be important in the consistent integration of individual emotions and group emotions.

The temporary group member problem. How to emotionally integrate temporary group members into a group while keeping a high degree of individual emotional realism? This problem is clarified best by an example. A panicked agent runs to the exit of a building. It passes by a group of happy agents. The agent gets integrated in the group by a clustering criterion. Two scenarios are now possible. First, the agent takes over the emotion of the group; an abrupt switch to a different emotion takes place, resulting in unrealistic individual emotional behaviour. Second, the agent's individual emotion slowly converges to the group's emotion, resulting in realistic individual emotional behaviour. Assumed that the agent's emotions are unrelated to its decisions (emotion is an 'agent add-on' to enhance the believability of the simulation), the agent continues to run to the exit. When exiting the group, this results in one of the following two scenarios. First, the agent keeps the emotion of the group, resulting in inconsistency between the agent's behaviour (running to the exit) and its emotion (happy). Second, the agent's emotion is reset to the emotion it had prior to entering the group, resulting again in an abrupt emotional change. Thus, overwriting the agent's emotion with the group's emotion may result in a problem if the agent is just briefly a member.

The group configuration problem. Scenario designers need the ability to define a group emotion for a whole group, with minimal loss of individual emotional expression of agents. In safety training, for example, a scenario may contain a panicked group as an important element of that scenario's specific training experience. So, apart from emotional consistency of a group, emotional configuration of a group is needed too.

4. Emotional-state based Architectures

To enforce a consistent group emotion, without losing individuality of the agents, we need to solve the problems mentioned above. According to appraisal theory the subjective evaluation of the environment in relation to the agent's concerns is responsible for emotion [7]. This evaluation is called appraisal, and results in values on

a set of appraisal dimensions[1] [8]. These values are responsible for emotion. To enforce consistent group emotions, we have defined a common emotion currency based on appraisal dimensions. This currency can be used for integration purposes if the computational model of emotion of every individual agent in a MeAS is based on the following three components (for a more detailed description of this approach and a discussion of some of the other benefits it has, see for example [1]).

- Appraisal System (AS). The agent's AS continuously evaluates the environment, resulting in a stream of n-dimensional appraisal vectors signifying the appraisal-results, with n equal to the number of appraisal-dimensions. These vectors are sent to the Emotional Maintenance System. The amount and type of appraisal-dimensions is configurable and need not be defined here. The AS can be a special component, but its functionality can also be integrated into the agent's controller.
- Emotional Maintenance System (EMS), The agent's EMS integrates the appraisal vectors and maintains the emotional-state. The emotional-state is also an n-dimensional vector. Appraisal-results thus induce changes to the emotional-state. Since the EMS both integrates appraisal-results from the AS and drives the Behaviour Modification System (described next), from a definitional point of view the EMS maintains both the appraisal-state and the emotional-state. In this paper we use *emotional-state* when we refer to the vector that is maintained by the EMS.[2]
- Behaviour Modification System (BMS). The agent's BMS selects, controls, and expresses emotional behaviour. The behavioural choices are based on the agent's emotional-state and additional knowledge the agent has. This allows the BMS to express the emotional-state in an agent-specific manner (e.g. "squad-leader" emotional expression versus "soldier" expression). In this paper we assume that the emotion represented by the emotional-state is the emotion the BMS expresses.

Separating the agent's appraisal (AS) from its emotional behaviour (BMS) via the introduction of an Emotional Maintenance System (EMS) enables us to use the agent's emotional-state as emotion-currency. This currency enables three types of architectures for the integration of individual emotions with group emotions. First, since the currency is a numerical vector, the emotional-states of agents can be integrated using arrhythmic operations like averaging, summing and interpolation. This type of architectures is based on emotional-state merging. Second, since all agents share the same emotion currency, the AS of one agent can send appraisal results to the EMS of a second agent which influences the emotional-state of the second. This type of architectures is based on appraisal grouping. Third, a shared emotion currency enables straightforward simulation of emotional communication via the creation of events with the emotional-state of one agent as data. Agent's who receive this event are able to 'understand' the communicated emotion directly, since they share the same emotion currency. This type of architectures is based on emotional communication.

Since all architectures are able to enforce consistent group emotions, we have investigated the advantages and disadvantages of these architectures. We have defined three main criteria to express the differences. *Computational performance*: MeASs are real-time environments and might need computationally light, but realistic emotions for

---

[1] An appraisal dimension is a variable - e.g. valence or novelty -, used to express the result of the appraisal of a perceived object - e.g. a friend - that influences emotion.
[2] Signal-based integration of appraisal-results is highly compatible with several appraisal-theoretic concepts including appraisal integration as proposed in [9].

all their agents; *Group-emotion quality*: some architectures are better than others in providing consistent group emotions while supporting individual emotions; *Implementation and design effort*: some architectures involve more design and developmental effort or considerations than others.

Criterion	Measure	Description
Perfor-mance	Appraisal scalability (++)	All agents in a group can share the appraisal system of one (or several) of its group members. Thus, the computation needed is expressed in the number of groups not the number of agents.
	Emotional state maintentance scalability (+)	All agents in a group can share the EMS of one (or several) of its group members. EMS scalability bonus is half the AS scalability bonus since the EMS performs a simpler function.
Quality	Group contributes to emotion (+)	All agents contribute via their appraisal to the emotion of the group, instead of just one (or several) group-leaders. The emotion of the group better reflects events in the environment.
	Agent-Group discrepancy (-)	If the agent's control architecture is such that the non-emotional behaviour (e.g. actions) of an agent is unrelated to the its emotion, discrepancies may arise between the emotional meaning of the individual agent's behaviour, and the individual agent's emotional expression based on the emotional state of the group (see also the 'temporary group member problem').
Design effort	Architectural changes to agent (-)	The number of architectural changes to the computational model of emotion that are needed to implement the integration.
	Group leader selection (-)	A leader is selected in the process of grouping agents. This increases design effort since a mechanism must be provided.
	Group emotional state merge (-)	The emotional state (or EMS) of multiple agents must be merged in the process of grouping agents.
	Clustering (-)	A group clustering method is needed.
	Group emotional state split (-)	In the process of group deformation, the emotional state of agents that leave the group must be recalculated.
	Appraisal System restart after deformation (-)	If the AS of an individual agent is stopped when the agent is in a group, difficulties restarting the AS may arise when the agent is detached from the group, especially if the AS is a complex cognitive appraisal system.
	Integration of emotion and control without additional assumptions (+)	The architecture allows integration of emotions and control (e.g. emotional influence on agent actions), without further assumptions about the agent's controller and AS. For example, the agent has a combined AS and controller that produces the appraisal vectors, or the agent's reasoning is influenced by its emotional state.

Table 1. Criteria used to score the architectures.

4.1. Architectures based on Emotional State Merging

The first class of architectures integrates individual emotions and group emotions by merging the emotional-state of the individuals into one emotional-state for the group, which is then managed by a special group EMS (Figure 1). This could be called a "centralised approach"; the group emotion is managed centrally, however individual agents still execute the other parts of their computational emotional system. Since every

agent has its own BMS, the agents are still able to express their emotions in a way that is characteristic to the agent (e.g. a "soldier" versus an "officer"). Positive characteristics of this class are the forced consistency of the group emotion due to a single EMS per group, and its performance due to EMS sharing, except in architecture 1 where the EMS of the group is not shared but used to influence the individual members. This approach can also be used in architectures 4ab and 5ab to solve the problem of group configuration. Architecture 3a specifically has high performance since the AS is also shared. However, a group EMS (arch. 2, 3ab) is either sensitive to agent-group discrepancy (the temporary member problem) or is in need of sophisticated merging and splitting mechanisms for the emotional-state and EMS.

Within this class of architectures, there are two main categories: individual-based appraisal (arch. 1 and 2) and leader-based appraisal (arch. 3ab). The main difference is that leader-based appraisal scales better due to the selection of a group-leader who is responsible for the evaluation of the environment, while individual-based appraisal may results in better group emotions since all individuals contribute to the emotional-state of the group. Also, leader-selection can become a difficult problem, since large groups can have very different environments at the extremities, resulting in an emotional reaction that is biased by the events on the leader's side of the group. Finally, to overcome stopping and restarting the agent's appraisal system (arch 3a) one can chose to leave the AS running (3b) without using its appraisal vectors. This downgrades performance but reduces design considerations. All comparison results are shown in table 2.


Figure 1. Architectures 1 (top-left), 2 (bottom-left), 3a (top-right) and 3b (bottom-right). Subscripts are used to denote which system belongs to which agent.

4.2. Architectures based on Appraisal Grouping

The second class of architectures integrates individual emotions and group emotions by selecting a group-leader agent who is responsible for the evaluation of the environment and who sends its appraisal vectors to the EMS of all group members (Figure 2). There is no central system to enforce group consistency. A first positive characteristic is the reduced agent-group discrepancy (the temporary member problem), which can be explained by the following. Instead of overwriting the agent's individual emotional-state, the group-leader takes over the appraisal vector stream to the EMS of all group members. This results in the individual integration of shared appraisal vectors. Newly arrived agents thus slowly converge to the emotional-state of the group, while keeping their individual emotional-state. An agent that merely passes by is influenced only if the

leader of the group actually appraises an event during the period of inclusion in the group. Once outside the group again, the agent still has an emotional-state that is close to the one it had before entering the group.

A second positive characteristic is that relatively little design effort is needed. There is no need for merging and splitting the EMS or emotional-state. And the AS of individual agents can be kept running (architecture 4b). Drawbacks of this class of architectures include scalability (especially arch. 4b) and the group-leader selection problem, as mentioned earlier.

Figure 2. Architectures 4a (left) and 4b (right).

4.3. Architectures based on emotional communication

The last class of architectures we have investigated is based on emotional communication (Figure 3). There is no central system to enforce group consistency. Communication is straightforward since both the emotional-state as well as the appraisal vectors share the same currency. Two forms of communication are possible. In appraisal-based communication (arch. 5a) appraisal vectors are send directly to the EMS of neighbouring agents, and in emotion-based communication (arch. 5b) events are created based on the emotional-state and broadcast for evaluation by neighbouring agents. Both architectures share the same characteristics. No EMS merging or splitting and special agent-grouping mechanisms are needed. Quality is good. All members of a group influence the emotional state of the agents in their direct environment by communication, which is psychologically plausible and supports emergent group emotions. The temporary member problem is solved in a natural way since agents are influenced either by the appraisal of neighbouring agents or by their own subjective evaluation of the emotional-state of neighbouring agents. However, performance is bad, since the agents in a group can neither share the AS nor the EMS.

Figure 3. Architectures 5a (left) and 5b (right).

5. Results-summary and conclusion

We conclude that the use of an emotional-state as emotion currency, made possible by separating the computational-model of emotion in three steps: appraisal, emotional-state maintenance and emotional behaviour, enables consistent emotions for dynamic groups of virtual agents. Architectures that are made possible by our approach range from high-performing to high-quality and trade-offs are possible between the two. Communication-based architectures (table 2, column 5a and 5b) that use our approach have high quality and little extra design considerations for group-emotions. Architectures based on group sharing of the emotional-state or appraisal system (table 2, column 3a,b and 4a) scale

better. Future work includes the extension of the simulation environment in [4] to test the different architectures. We believe that our research enables agent- and scenario designers to benefit from the realism a computational model of emotion brings to individual virtual agents, without losing group consistency. Also, designers can trade-off quality of the group emotion for computational performance.

Table 2. Overall comparison results

Criteria		Measure	1	2	3a	3b	4a	4b	5a	5b
Performance	p1	Appraisal scalability	0	0	1	0	1	0	0	0
	p2	Emotional state scalability	0	0	1	1	0	0	0	0
	P=	(p1*2+p2)/3	0.00	0.00	1.00	0.33	0.67	0.00	0.00	0.00
Quality	q1	Group contributes to emotion	0	1	0	0	0	0	1	1
	q2	Agent-Group behavioral discrepancy	0	1	1	1	0	0	0	0
	Q=	(q1+(1-q2))/2	0.50	0.50	0.00	0.00	0.50	0.50	1.00	1.00
Design effort	e1	# Architectural changes to agent	1	3	4	3	2	1	1	1
	e2	Group leader selection	0	0	1	1	1	0	0	0
	e3	Group emotional state merge	1	1	1	1	0	0	0	0
	e4	Group agent clustering	1	1	1	1	1	1	0	0
	e5	Group emotional state split	1	1	1	1	0	0	0	0
	e6	Appraisal System restart after deformation	0	0	1	0	1	0	0	0
	e7	Integration of emotion and control without further assumptions	1	0	0	0	0	0	1	1
	E=	(e7+(4-e1)/4+5-(e2+e3+e4+e5+e6))/7	0.54	0.32	0.00	0.18	0.36	0.68	0.96	0.96
Total measures		(P*2+Q*2+E*7)/11	0.43	0.30	0.18	0.17	0.44	0.52	0.80	0.80
Total criteria		(P+Q+E)/3	0.35	0.27	0.33	0.17	0.51	0.39	0.65	0.65
Quality versus Performance		IF(2*P<Q) then"q" elseif(2*Q<P) then "p" else "qp"	q	q	p	p	qp	q	q	q

References

[1] D. DeGroot and J. Broekens. Using Negative Emotions to Impair Gameplay. *BNAIC*, 2003.

[2] B. Mac Namee and P. Cunningham. Creating Socially Interactive Non Player Characters: The m-SIC System. *IJIGS* 2(1), 2003.

[3] S. Marsella and J. Gratch. Modeling the Interplay of Emotions and Plans in Multi-Agent Simulations. *Proceedings of the 23rd Annual Conference of the Cognitive Science Society*, Edinburgh, Scotland, 2001.

[4] N. Netten. Towards Believable Virtual Characters Using A Computational Model Of Emotion. *Master's Thesis*, LIACS, Leiden University, 2004.

[5] K. Oatley. Emotions. *The MIT Encyclopedia of the Cognitive Sciences (MITECS)*. Edited by Robert A. Wilson and Frank Keil, ISBN 0-262-23200-6, 1999.

[6] A. Ortony, G.L. Clore and A. Collins. *The Cognitive Structure of Emotions*. Cambridge University Press, 1988.

[7] I.J. Roseman and C.A. Smith. Appraisal Theories: Overview, Assumptions, Varieties, Controversies. In K. R. Scherer, A. Schorr, & T. Johnstone (Eds.), *Appraisal processes in emotion: Theory, Methods, Research*, Oxford University Press, 2001.

[8] K.R. Scherer. Studying the Emotion-Antecedent Appraisal Process: An Expert System Approach. *Cognition and Emotion* 7(3/4), pp325-355,1993.

[9] Smith, Craig A., Kirby, Leslie D. Consequences require antecedents: Toward a process model of emotion elicitation. In*: Forgas, Joseph P. (Ed), Feeling and Thinking: The role of Affect in Social Cognition*. Cambridge University Press 2000