# SCALABLE AND FLEXIBLE APPRAISAL MODELS FOR VIRTUAL AGENTS


Joost Broekens and Doug DeGroot
Leiden Institute of Advanced Computer Science (LIACS)
Leiden University
2333 CA Leiden
Netherlands
E-mail: {broekens; degroot}@liacs.nl



**KEYWORDS**
Computational Models of Emotion, Appraisal Theory, Scalability, Dynamic Models, Games and Virtual Agents



**ABSTRACT**

Computational models of emotion are useful in a variety of domains, including games, virtual realty training and HCI to name a few. Many of these models are inspired by appraisal theory. Most appraisal theories share with virtual agents the assumption that beliefs, desires and intentions are the basis of reasoning and thus of the emotional evaluation of the agent's situation. Consequently most computational models of emotion are deeply embedded into the agent model. In this paper we address the problem of how to emotionally instrument a system in a modular and extensible way, so that emotional sophistication can be added incrementally to a system. We propose a solution based on a modular, signal-based approach to computational emotions that allows us to develop scalable appraisal models that are easily added to non-emotional systems. Our approach allows runtime trade-off between emotional quality and performance, which makes it particularly useful in domains in which available computation time is unknown, like the gaming domain. We present experimental results that back-up our approach.


**INTRODUCTION**

In psychology emotion is often defined as a psychological state or process that functions in the management of goals and needs of an individual. 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. Thus, according to this view an emotion is a heuristic that relates the events from the environment to the agent's goals and needs (Oatley 1999). Additionally, emotion is a communication medium.

Computational models of emotion are embedded in agents in a variety of domains including HCI and electronic tutors, non-player characters (NPCs) in games, virtual-reality safety training environments and decision-making and planing. Emotions are embedded in virtual agents primarily to create an enhanced sense of realism, using emotional expression and the interplay between emotions and plans (Marsella and Gratch 2001). It has been argued that emotions and emotion-like phenomena are a good way of enhancing realism and thereby entertainment value of NPCs in games (Baillie-de Byl 2003, Mac Namee and Cunningham 2003).

The majority of computational models of emotion embedded in virtual agents is inspired by appraisal theories, cognitive theories of emotion that attempt to explain why a certain event results in one emotional response rather than another and why a certain emotion can be elicited by different events. The key concept of most appraisal theories is that the subjective cognitive evaluation of events in relation to the agent's goals and needs is responsible for emotion (Roseman and Smith 2001). More generically one can say that events have to be evaluated as having personal meaning (van Reekum 2000). This evaluation is called appraisal. Most appraisal theories assume that appraisal is a necessary and sufficient condition for emotion (Roseman and Smith 2001).

Agents often use a belief-desire-intention (BDI) based architecture (Jennings et al. 1998). If cognitive evaluation of events in relation to the agent's goals and needs is sufficient for emotion then the addition of a subjective evaluation of events related to the beliefs, desires and intentions of an agent is sufficient for computational emotions. This explains the popularity of appraisal theories in emotional agents.

Computational models of emotion must often be deeply integrated with the agent's non-emotional components because they depend on the BDI architecture of the virtual agent, as mentioned above. This deep integration has two problems. First, it takes quite some effort to add emotions to a non-emotional agent, because the computational model of emotion needs to be embedded into the agent's architecture. Second, computational models of emotion are difficult to adapt and upgrade in an incremental fashion for the same reason. Both problems are important for game development. From a marketing and sales point of view one wants to be able to incrementally add emotional sophistication to NPCs to sell upgrades of a game. From a technical point of view one wants to be able to evaluate which version of a computational emotional model to use based on e.g. performance, quality and stability.

We have investigated these problems and propose the FeelMe framework for computational emotions that is inspired by appraisal theory *and* allows the development of scalable appraisal models. The ability of the FeelMe framework to dynamically integrate the results of different emotional instrumentations that run simultaneously enables the development of scalable appraisal models. A scalable appraisal model can be used to emotionally instrument an agent (virtual agents, NPCs) in an incremental manner. A scalable appraisal model also allows runtime trade-off between emotional quality and performance. This trade-off ability makes these models particularly useful in domains in which computation power is an unknown factor, like the gaming domain.

In the next section we explain scalability of computational models of emotion and incremental instrumentation in more detail. Then we describe the FeelMe framework. We continue with a detailed description of how the FeelMe framework can be used to integrate different emotional instrumentations and what constraints exist for these instrumentations. Finally we present a proof-of-concept experiment with a game agent that uses our approach. The experiment shows that scalable appraisal models are possible, pointing out that the FeelMe framework enables game-character and virtual agents designers and developers to incrementally add more and more sophisticated emotions to their virtual agents.

## SCALABILITY AND GAMES

Any system for use in games must be efficient in terms of computation required (Mac Namee and Cunningham 2003). Additionally, games must run on different platforms so computation power is not a known factor. Users with a high-end PC want to have high-end effects so the minimum system requirements to run a game cannot be taken as development standard. A good solution to this dilemma is to use scalable systems that are able to trade-off quality versus performance. This trade-off is often seen in 3D-graphic engines and chess-engines, in which level of respectively graphical detail and intelligence can be dynamically traded-off for respectively frame-rate and total game-time.

The systems that make up a game engine are usually triggered at regular intervals. Flexibility regarding the frequency of this triggering enables a different form of scaling, namely scaling based on a trade-off between temporal quality and performance. For example a 3D-graphic-engine can be triggered 10 times per second, in which case the frame-rate is low and the frames are staggering. However, individual rendered frames are still consistent and the sequence of frames still consistently shows the motion of objects and agents in the game, although less detailed. A computational emotional model able to do the same, that is, scalable quality and flexible triggering, thus has a practical advantage compared to one that can't. We refer to these two kinds of scalability as *runtime-scalability*.

## INCREMENTAL INSTRUMENTATION

We define incremental emotional instrumentation of systems as a development process based on the step-by-step addition of complexity to a computational model of emotion resulting in meaningful and more sophisticated emotions of the agent that is consistent with the emotions of the simpler versions of the model. Consistent in our case means that the more complex version behaves equally meaningful as or more meaningful than the simpler version. By meaningful we mean that a human observer (e.g. the gamer) can - potentially in retrospect - understand why the agent exhibits a certain emotion or chooses to act in a certain way. In other words, more complex models should add sophistication to the emotions of the agent as well as add human understanding of the agent's emotion and related actions.

A possible first step to instrument a non-emotional system is by using event encoding. Based on common-sense, an event is given a specific emotional property, just like emotionally laden words in language already have. When an event is encountered by an agent, the agent's emotion is changed accordingly. For example, a Quake bot seeing his team-mate die could be configured to experience sadness by defining a high "sadness" property for the 'team-mate-died' event. It could also be configured to experience anger. Actually, many different emotions would make sense and are not depending on the event but more so on the evaluation of that event. Although efficient and sufficient in some situations, this way of directly encoding emotions into properties of events does not work in general and is not psychologically plausible[1]. According to appraisal theory, events are interpreted by the agent, after which the emotion is influenced. This interpretation includes reasoning about what the event means to the goals and needs of the agent, which is depending not only on the event but also on the current BDI state of the agent. So, an event does not *directly* influence the emotion - at least not in general - but the interpretation of the event does.

A BDI based approach to computational emotion evaluates events in the context of the current goal hierarchy of the agent, and determines the resulting emotion based on this evaluation. Switching from an event-encoding approach to a BDI-based approach is necessary for more meaningful computational emotions. However, some situations might be much easier to give emotional meaning using event encoding instead of using BDI-based appraisal. Also, when computing time becomes a bottleneck, an agent might need to switch to a simpler emotional instrumentation to save computation time for other sub-systems of the game.

Event-encoding and BDI based appraisal are two possible ways to emotionally instrument – or extend an emotional instrumentation of - a system. The relevant question for incremental emotional instrumentation of virtual agents is thus how to integrate the results of different concurrent emotional instrumentations? The ability to emotionally instrument a system in an incremental manner is referred to as *model-scalability*.

## FEELME: A DYNAMIC APPROACH TO COMPUTATIONAL EMOTIONS

The FeelMe framework (DeGroot 2004) is a modular approach to computational emotions and is based on a strict separation of the computational emotional process in five main steps (see Figure 1). These steps are described in more detail in this section. This framework for computational emotions has been developed to study the effects of emotion on decision-making by using emotions as first-order objects in reasoning (DeGroot and Broekens 2003).

- The Decision Support System (DSS) provides mediated access to the existing system (e.g. an existing arcade game). Since some information in the environment of an agent or in its own internal state is not directly suitable for appraisal, the DSS translates this information before sending it to the Appraisal System. The DSS constructs to-be-appraised objects, based on the events occurring in environment of the agent and sends these objects to the Appraisal System.

[1] In some domains psychological plausibility of emotions is of high importance, like virtual reality safety training.

- The Appraisal System (AS) continuously emotionally evaluates the constructed objects and interprets these in terms of values on a set of subjective measures, called appraisal dimensions. An appraisal dimension is a variable - e.g. arousal or valence - used to express the result of the emotional evaluation of a perceived object, for example a friend. The evaluation of the AS results in a continuous stream of *n*-dimensional vectors representing the *appraisal-results*, with *n* equal to the number of appraisal-dimensions. These vectors are sent to the Emotion Maintenance System. The number and type of appraisal-dimensions is configurable and need not be defined here. Just for the purpose of consistent terminology, in this paper we call any mechanism that produces appraisal-results an *appraisal mechanism*. Event-encoding can thus be called an appraisal mechanism, provided that it produces appraisal-results as defined above.
- The Appraisal Signal Modulator (ASM) can perform signal pre-processing on the incoming appraisal-results - like amplification of, dampening of and correlating certain appraisal dimension values - before these are sent to the EMS.
- The Emotion Maintenance System (EMS) continuously integrates the appraisal-results and maintains the agent's emotional-state. The emotional-state is also an *n*-dimensional vector. Appraisal-results induce changes to the emotional-state, thus for the EMS an appraisal-result is an *n*-dimensional vector of deltas of appraisal dimensions. This integration of deltas is what we refer to as a *signal-based* approach. The emotional-state of an agent can thus be understood as a continuously moving point in an *n*-dimensional space of appraisal dimensions. In this paper we use *emotional-state* when we refer to the vector that is maintained by the EMS. An emotional-state actually is not just a computer science approach to emotions. Many emotion theorists use this concept of a state to define emotion (Mehrabian 1980, Russell 2003, Reisenzein 2001, Scherer 2001).
- The Behaviour Modification System (BMS) selects, controls, and expresses the agent's emotional behaviour. The behavioural choices are based on the agent's emotional-state and additional knowledge the agent has.

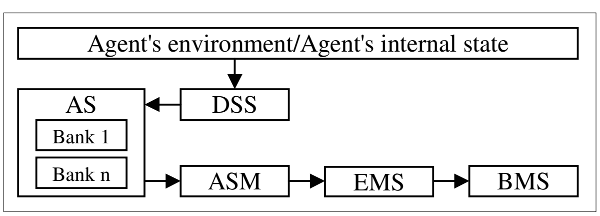


Figure 1. Overview of the components of the FeelMe framework relevant to this paper. Banks are explained later

Now, how does this help us integrate the appraisal-results of different concurrent appraisal mechanisms of an agent? The introduction of the emotional-state and it's interaction with the Appraisal System are key. The AS outputs vectors that represent changes to the emotional-state. Any mechanism that produces these vectors can be used. The Appraisal System can thus consist of multiple subsystems, provided that these produce the same kind of vectors. This opens up the possibility of adding more and more subsystems to add more and more detail to the virtual agent's emotions. The EMS integrates the appraisal-results. The most simple version of the EMS continuously adds-up all vectors sent to it by the AS (see formula 1 on page 6). Even such a simple paradigm allows integration of the results of different appraisal mechanisms into a meaningful representation of the emotional-state, provided that several guidelines for modular appraisal are used, as we will show. Also, this signal-based integration of appraisal-results from different concurrent mechanisms is highly compatible with the concept of appraisal integration by appraisal detectors as proposed by the appraisal theorists Smith and Kirby (2000).

## MODULAR AND SCALABLE APPRAISAL

In order to build scalable appraisal models we have defined *appraisal banks*. In this section we explain what an appraisal bank is, why it facilitates the development of scalable appraisal models and what kind of appraisal-results are needed for effective integration of these result.

### Context Sensitive Appraisal Banks

The Appraisal System (AS) is the complete appraisal system of an agent and an appraisal bank is a sub-systems of the AS (see Figure 1). An appraisal bank is an object (in the OO sense) that contains a set of functions that emotionally evaluate specific aspects of the agent's environment and internal state, for example all events related to survival. An appraisal bank is context sensitive, that is, the contribution of the bank's appraisal-result - as determined by the bank's evaluation functions - to the emotional-state of the agent depends on the situation of the agent. An appraisal bank can influence the contribution of another appraisal bank's appraisal-results through dependencies. Such dependencies allow the definition of causal connections which enable modelling of levels of appraisal and evaluation sequence (Scherer 2001, van Reekum 2000). To facilitate development of scalable appraisal models we enforce strict modularity of the AS by assuming that appraisal banks evaluate independent of each other.

We now explain why context sensitivity and evaluation sequence of appraisal banks facilitate the development of both model- and runtime-scalable appraisal models. First, context sensitivity facilitates the development of new - more elaborate - banks on top of older - more generic - banks. These new banks - for example based on BDI-based appraisal - can be sensitive to contexts where more meaningful emotions are needed but not achieved with the older banks - based on for example event encoding. The older banks can be sensitive to those contexts in which they work well. Context sensitivity of a bank can be configured by a game-character designer. Context sensitivity can also be implied based on the activity of appraisal banks (e.g. larger appraisal-results are more important then smaller ones). In this case the contribution to the emotional-state of one bank inhibits the contribution of another bank. Dependencies determine which bank influences which. These dependencies can be used to build an interconnected set of appraisal banks that influence each other. Context sensitivity facilitates model-scalability because banks evaluate situations they recognise while other banks are silent.

Second, simple - computationally cheap - appraisal banks, and elaborate - computationally expensive - appraisal banks can be active simultaneously. An agent can dynamically adapt its appraisal effort (and thus computation time needed) by switching between simple and elaborate banks, depending on the context and the maximum amount of computing time available for its computational emotion system. Alternatively the user can adapt the emotional detail of the NPCs by configuring his game. A third way of adapting appraisal effort can be based on the distance between the NPC and the user's viewpoint. All three adaptation examples show the runtime-scalability potential of appraisal banks. Interestingly, the dynamic adaptation of appraisal effort depending on the situation and available resources is consistent with certain appraisal theoretic approaches towards emotion (Scherer 2001).

The set of functions in an appraisal bank can be designed to produce meaningful appraisal-results. A second bank can be completely separated from the first and also produce meaningful appraisal-results. If both banks work well together at the same time, they can be active at the same time. If they don't then one bank has to inhibit the other, or both banks have to be configured to be sensitive to mutually exclusive contexts (see below). Grouping appraisal in context-sensitive banks facilitates debugging of the appraisal model, since designers can focus on specific banks and assume other banks are deactivated. For example, in a typical RPG scenario this allows the development and debugging of an NPC's appraisal banks for battle, travel/quest and village/city situations.

In this paper we analyse the results of an experiment testing the difference between one bank, and two dependent banks in which the first inhibits the second.

### Constraints for appraisal banks

If we assume that the Appraisal Signal Modulator (ASM) does not pre-process appraisal-results and the Emotion Maintenance System (EMS) only integrates appraisal-results by addition, what kind of appraisal-results do appraisal banks need to output for successful integration? Multiple constraints exist, of which we describe three.

First, appraisal-results need to be defined on the interval scale, addition of appraisal-results must be meaningful. To be more concrete, a 0.5 on the Pleasure dimension produced by bank 1 must emotionally mean the same as a 0.5 on the Pleasure dimension produced by bank 2. Furthermore, in order for the EMS to meaningfully integrate the values by adding up, a 0.5 increase or decrease on a certain dimension must always mean the same for that dimension. These two criteria do *not* have to hold *between* appraisal dimensions, a 0.5 Pleasure increase does not have to mean the same in terms of intensity-change as a 0.5 Arousal increase, but we will not go into this issue here.

Second, the set of appraisal banks (the Appraisal System) together must be able to produce non-zero appraisal-values for all appraisal dimensions. These values must be both positive and negative. This allows the emotional-state to potentially be driven in all directions. Note that this does not need to hold for one bank in particular, since it is perfectly fine if one bank produces mostly positive values for a certain appraisal dimension while another produces mostly negative ones. This would still potentially drive the emotional-state in both directions. Being able to drive the emotional-state in all directions is needed to maximise emotional coverage. Emotional coverage is the ability of the computational model of emotion to attain all possible emotional-states, as defined by the appraisal-dimensions used in the computational model of emotion. Emotional coverage is important for several reasons, of which we mention only one. An agent that is designed to express a set of emotions must also be able to attain these emotions. Not being able to do so presents a huge loss of development effort (i.e. facial expression rendering, emotional behaviours, etc).

Third, appraisal banks need to respond to mutually exclusive contexts. This can be explained by the following. If we assume that $r_1$ and $r_2$ are the absolute values of two appraisal-results produced by respectively bank $B_1$ and $B_2$ at a certain time, and $r_2 \neq 0$, and the newer version of an appraisal model contains both $B_1$ and $B_2$ while the simpler version contains only $B_1$, then the simpler model produces $r_1$ while the newer version produces $r_1+r_2 \neq r_1$. Nothing can be said about how meaningful $r_1+r_2$ is, even though $r_1$ and $r_2$ may be meaningful by themselves. At least two ways to ensure model-scalability - i.e. incremental emotional instrumentation - exist: first, appraisal banks are never active together in which case $r_1+r_2$ never happens; second, $B_1$ knows about $B_2$ or vice versa so that they can adapt $r_1$ and $r_2$. This introduces a dependency between two versions of the appraisal-model, and such a dependency limits model-scalability. There are several other issues that relate model-scalability, choice of appraisal dimensions and emotional coverage to each other, but these would diverge us too much from the main point.

Mutual exclusiveness is a rather restricting constraint. Fortunately another option is available. If $r_2 \approx 0$ then $r_1+r_2 \approx r_1$. This means that, if $B_1$ and $B_2$ are active at the same time and $B_2$ is an appraisal bank that "fine-tunes appraisal" while $B_1$ "looks at the big-picture", then both banks can be active at the same time. Now $B_2$ incrementally adds more meaning to the appraisal model while staying consistent with the model only containing $B_1$.

If we assume the Appraisal Signal Modulator (ASM) is pre-processing appraisal-results before the EMS integrates these, do the guidelines stay the same? Yes and no, appraisal-results still need to be defined on the interval scale, because the EMS still has to integrate them and the set of banks still must be able to produce non-zero positive and negative values for all appraisal dimensions in order to maximise emotional coverage. However, other scenarios are possible for the interplay between two or more banks. We explain one of these scenarios. If the ASM constructs a weighted average[2] of $r_1$ and $r_2$ where the weighing function is based on the intensity of the appraisal-result - intensity can be calculate using for example the length of the appraisal-result

[2] A similar weighted influence of appraisal-results – using attention as weighing function - has been proposed to explain the effects of concurrent appraisals on human emotion (Schimmack et al 2001).

vector, as shown in equation 2 page 6 - then $min(r_1,r_2)<r_1+r_2<max(r_1,r_2)$. This means that if both $r_1$ and $r_2$ are meaningful when used separately, the appraisal-result as integrated by the EMS is between $r_1$ and $r_2$ and has a high chance of also being meaningful. The problem with this approach is that appraisal-results from appraisal banks that should "fine-tune" the appraisal model - like the above example of $B_2$ - should be added to the appraisal-results of "big-picture" appraisal banks instead of integrated with these results in an average. In equations this means: if $r_2 \approx 0$ then $(r_1+r_2)/2 \approx r_1/2$, while $B_2$ was designed to achieve $r_1+r_2 \approx r_1$. To conclude, without further assumptions the ASM cannot solve the mutual exclusive contexts constraint, but it can soften it. Appraisal banks that "fine-tune" appraisal can be configured to be left untouched by the ASM, and all other appraisal banks can be either averaged by the ASM or mutually exclusive. When needed, the ASM enables a range of different mechanisms to pre-process the appraisal-results of appraisal banks making these results suitable for integration by the EMS.

## EXPERIMENTAL ASSUMPTIONS

We have instrumented a Java version of the arcade game of PacMan (Chow 2003). Since we want to test if our signal-based, modular approach facilitates incremental emotional-instrumentation, and that this incremental instrumentation is feasible even for existing non emotional systems, programming a game ourselves would have seriously diminished the convincing power of our results.

The game of PacMan consists of an "eater" in a rectangular maze, filled with dots, power-pills, fruit and several ghosts. A human player controls the "eater". The goal it is to collect as many points as possible by eating the objects in the maze. When a ghost touches the "eater", it loses a life. When no lives are left, the game is over. However, if the "eater" eats a power-pill, it is temporarily able to eat the ghosts, thus reversing roles. When all dots are eaten, the game advances to the next - more difficult - level.

### PacMan as Experimental Platform

We have chosen PacMan for the following reasons. First, PacMan has easy to define goals, like survival and collecting points. This facilitated development of an appraisal model with one bank related to survival (e.g. avoiding ghosts) and then extend this model with a bank related to the goal of collecting points (e.g. eating dots). Second, the "eater" in PacMan potentially has many different emotions that make sense. Eating ghosts, eating dots, being chased, chasing, etc. are all different situations relating to different emotions. This allows us to test to what extend emotional coverage changed depending on the appraisal-model. Third, PacMan is an 'action-packed' environment, which allows us to test our signal-based approach, under continuous-time constraints.

### Pleasure Arousal Dominance Dimensions

Our approach does not prescribe a specific set of appraisal dimensions. We have chosen the Pleasure, Arousal, Dominance (PAD) personality-trait and emotional-state scales by Albert Mehrabian (1980) for the following reasons. First, even though Mehrabian is not an appraisal theorist and his dimensions are generally not considered to be appraisal dimensions, he argues that any emotion can be expressed in terms of values on these three dimensions, and provides extensive evidence for this claim (Mehrabian 1980). This makes his three dimensions suitable for a computational approach[3]. Second, since the PAD scales are validated for both emotional-states and traits, they provide a useful basis for a computational framework that consistently integrates states and traits (even though we don't use traits in the experiments, this is very valuable for further instrumentation experiments). Last, Mehrabian (1980) provides an extensive list of emotional labels for points in the PAD space. Figure 2 gives an impression of the emotional meaning of combinations of Pleasure, Arousal and Dominance.

The following sample ratings illustrate definitions of various emotion terms when scores on each PAD scale range from -1 to +1:

angry (-.51, .59, .25)
bored (-.65, -.62, -.33)
curious (.22, .62, -.01)
dignified (.55, .22, .61)
elated (.50, .42, .23)
hungry (-.44, .14, -.21)
inhibited (-.54, -.04, -.41),
loved (.87, .54, -.18)
puzzled (-.41, .48, -.33)
sleepy (.20, -.70, -.44)
unconcerned (-.13, -.41, .08)
violent (-.50, .62, .38).

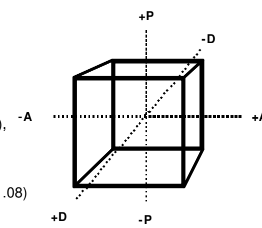


The emotional state "angry" is a highly unpleasant, highly aroused, and moderately dominant emotional state. The "bored" state implies a highly unpleasant, highly unaroused, and moderately submissive state.

From: Albert Mehrabian's (1980) PAD Scales.

Figure 2: The Mehrabian P-A-D Temperament Scale

In an attempt to instrument PacMan in such a manner that the appraisal-results are defined on the interval scale - i.e. adding results from different banks means something -, we need guidelines to think about events in terms of appraisal-dimensions instead of emotions. What does a certain event mean in terms of P, A and D?

The Pleasure dimension is highly related to the impact an event has on the probability of fulfilment of an agent's desires (e.g. Mehrabian 1980, Reisenzein 2001, Scherer 1993). If the event increases the probability of the outcome, Pleasure is positive, else it's negative. The "eater" in the PacMan game has two desires: survival and collecting points. According to Mehrabian (1980), arousal is highly correlated with activity and alertness. This relates to the expectancy and novelty of an event. The appraisal dimensions expectancy and novelty are related to the amount of attention a certain environmental change gets (van Reekum 2000). Based on these observations we assume that Arousal is the amount of attention needed to address a certain event. A ghost needs a lot of attention, while eating a dot needs only a little.

Dominance is a measure for the influence the situation has on the agent's freedom of choice to act in different ways (Mehrabian 1980). High Dominance implies a large freedom of choice, while low Dominance implies little choice. This maps well to the PacMan game. For example, seeing a ghost decreases Dominance, while seeing an edible ghost increases

[3] It might even be very interesting to study in detail how these non-appraisal dimensions behave in a signal-based appraisal setting.

Dominance. We have used these guidelines to think in terms of P, A and D.

**Instrumentation of Appraisal Banks: Survival and Collecting Points**

To test if context sensitive appraisal banks facilitate the development of scalable appraisal models, PacMan is instrumented in two ways. First, a simple instrumentation based one appraisal bank that emotionally evaluates events related to survival. Second, a more complex instrumentation based on two appraisal banks, one related to survival the other related to collecting points. In both banks we have used event-encoding to simulate emotional meaning of events. The DSS constructs the actual events. We now describe how events are interpreted by the two appraisal banks..

*"Survival" bank*
This bank appraises only survival related events (Table 1). The rationale for the P, A and D values is based on the guidelines described above. Pleasure depends on the level of obstruction versus conductance of an event related to a goal. For example, seeing a ghost is moderately obstructing for survival, while being eaten is highly obstructing. Arousal is related to the amount of attention an event needs. For example, seeing a ghost needs a moderate amount of attention while losing a ghost needs no attention (because the ghost poses no thread anymore). Dominance is related to the amount of freedom the "eater" has. For example seeing a ghost decreases the amount of freedom, while losing a ghost increases the amount of freedom.

Table 1: "Survival" bank

| Event | Pleasure | Attention | Dominance |
|---|---|---|---|
| See_ghost | -.5 | 0.5 | -.5 |
| Lost_ghost | 1.0 | 0.0 | 0.5 |
| Eaten_by_ghost | -1.0 | 1.0 | -1.0 |

*"Points" bank*
This bank appraises only events related to the goal of collecting points. Table 2 shows the events for the "points" bank. Again, the rationale for the P, A and D values is based on the guidelines described above.

Table 2: "Points" bank

| Event | Pleasure | Attention | Dominance |
|---|---|---|---|
| eaten_ghost, , | 1.0 | 1.0 | 0.0 |
| see_edible_ghost | 0.5 | 0.5 | 1.0 |
| eaten_fruit, | 0.5 | 0.2 | 0.0 |
| eaten_dot, | 0.2 | 0.2 | 0.0 |
| eaten_power | 0.2 | 0.2 | 0.0 |

**Appraisal-results**

Appraisal-results are produced by both appraisal banks and are based on situational change. This means that whenever an event is interpreted by an appraisal bank at time *t*, it compares if this event has already been encountered at time *t-1*. If this is not the case, the appraisal dimension values associated with the event are sent as appraisal-result. If it is the case, nothing is sent. If an event is no longer encountered at time *t* while it was at time *t-1*, a relaxation function kicks in. This function is responsible for sending enough small values over a short time period - say until *t+x* - so that these values - when summed - are the exact opposite of the appraisal dimension values associated with the event encountered at time *t-1*. The mechanism has been adapted to work for multiple events, but we will not go into this here.

One of the reasons for implementing appraisal banks in this way is that we are now sure that an appraisal bank outputs both positive as well as negative appraisal values for all appraisal dimensions that are used by the events the bank interprets. As mentioned above, this is an important criterion for emotional coverage. Another reason is that continuous exposure to, for example, eating dots would permanently drive the emotional-state to an extreme value (remember that an appraisal-result is a delta - a change - and that these are just added up by the EMS). Not going into the discussion of whether this is or isn't plausible, it is a problem we would have had to solve in one way or the other for the current experiment. We have chosen for a simple but effective appraisal mechanisms using both situational habituation -i.e. measuring situational change - and subsequent relaxation. We would like to stress, however, that this is just one of many ways an appraisal mechanism could be implemented in order to produce appraisal-results that maximise emotional coverage as well as protect the emotional-state from "walking to extremes".

**Context sensitivity**

In the simple instrumentation - using only the "survival" bank - context sensitivity is irrelevant. There is just one bank active at all times. In the complex instrumentation context sensitivity is of importance and implemented in the following way. Since survival is more important than points, the "points" bank is inhibited by the "survival" bank. This is implemented by weighing the contribution to the emotional-state of the appraisal-result of the "points" bank relative to the amount of emotional activation (appraisal-intensity) in the "survival" bank. Formula (1) implements the weighing function, where $\Delta_{goal}'$ is the weighted appraisal-result vector as to send to the EMS by the "points" bank, $\Delta_{goal}$ is the non-weighted vector, $|\Delta_{survival}|$ is the length of the appraisal-result vector of the "survival" bank and the cubic root of 3 is the maximum length of an appraisal-result vector[4].

$$\Delta_{goal}' = \Delta_{goal} * \left(1 - \frac{|\Delta_{survival}|}{\sqrt[3]{3}}\right) \qquad (1)$$

If the "survival" bank is highly active, the appraisal-results from the "points" bank have are almost no influence on the final appraisal-result sent to the EMS and vice-versa. This mechanism should result - and the experiment shows it does - in emotions produced by the complex model that are consistent with the emotions produced by the simple model in survival-related situations. This mechanism exemplifies context-sensitivity of appraisal banks. Note that appraisal

[4] Calculating intensity in a Pleasure-Arousal theory of emotion based on the length of the Pleasure-Arousal vector is psychologically plausible (Reisenzein 1994).

banks allow abstraction from the actual event interpretation, facilitating a modular approach to the appraisal model.

### Integration of Appraisal-results

Appraisal-results are integrated by the EMS using equation (2), where $E_t$ is the emotional-state at time $t$, $E_{t+1}$ is the new emotional-state, $n$ is the number of appraisal banks and $\Delta PAD_{ti}$ the appraisal-result vector of bank $i$ at time $t$.

$$E_{t+1} = E_t + \sum_{i=0}^{n} \Delta PAD_{ti} \qquad (2)$$

The EMS adds up appraisal-results produced by the banks.

## EXPERIMENTAL RESULTS

The experiment itself consists of a human that controls PacMan, and who plays the first level of the PacMan game (by eating all dots), while losing a life two times during the process, and eating at least one Ghost. To be able to compare the two different instrumentations and the effect of triggering the appraisal banks of the instrumentations at different frequencies - i.e. at 5 times per second and 10 times per second - , we have configured PacMan in such a way that we were able to test all four instrumentations -i.e. both instrumentations at 5 and 10 times per second - in just one test-run. We instantiated four different versions of the emotion system and events were delivered to the appraisal banks of all four instantiations. Plots of the emotional-state changing over time have been generated.

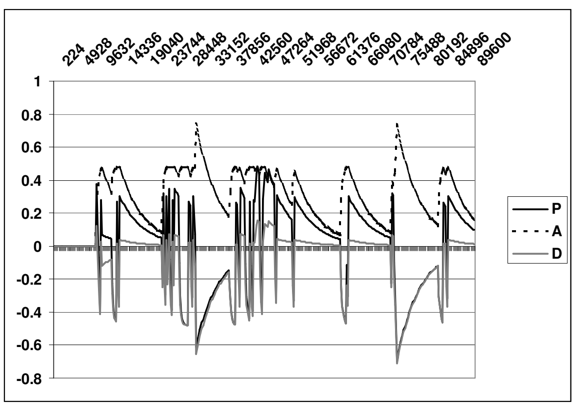


Figure 3: "Survival" PacMan, 200ms instrumentation

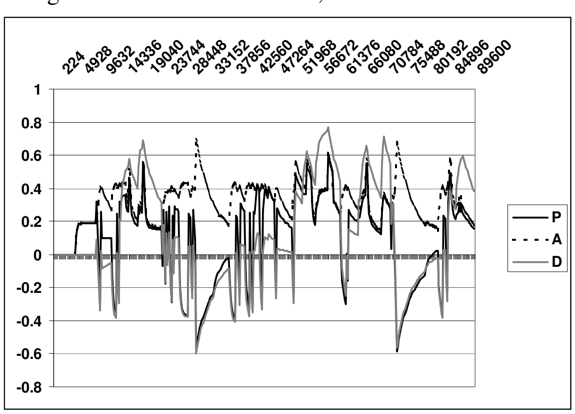


Figure 4: PacMan using both banks, 200ms instrumentation

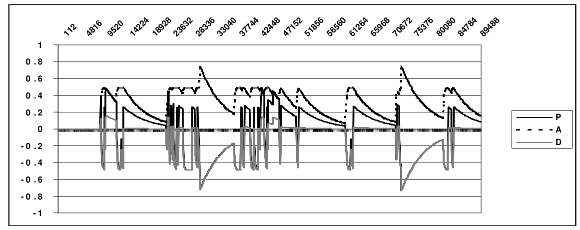


Figure 5: "Survival" PacMan, 100ms instrumentation

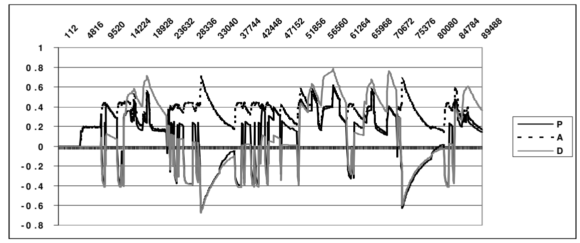


Figure 6: PacMan using both banks, 100ms instrumentation

### Model scalability

All graphs clearly show broad emotional coverage (the 2-bank instrumentation shows broader coverage, however), irrespective of appraisal rate. Even with a minimal instrumentation based on 3 events, the emotional-state varies substantially from (-P, +A, -D) to (-P, -A, -D) to (+P, -A, 0D). Furthermore, the ability to define context-sensitive appraisal-banks clearly allows us to first define this minimal instrumentation and subsequently scale-up the model by adding a new bank. The context sensitivity of the banks - using inhibition of the "points" bank - results in consistent behaviour of the emotional-state. In the "both banks" instrumentation the emotional-state is more meaningful due to the second bank, shown by the overall difference between Figure 3 and 4 and in particular the effect the "eaten_ghost" event has on P and D around $t=56672$. The emotional-state is at least as meaningful in those situations where the "survival" instrumentation already produced a meaningful emotional-state, shown by the effect the "eaten_by_ghost" event has on P, A and D around $t=28448$ and $t=70784$ in Figure 3 and 4. This shows that context-sensitive appraisal banks - enabled by our signal-based approach - facilitate model-scalability.

### Runtime-scalability

Comparison of the results between the 100ms and 200ms instrumentations shows that our signal-based, context-sensitive appraisal banks are insensitive to a 100ms difference in triggering frequency. We can see that Figure 3 and Figure 5 as well as Figure 4 and 6 are pair-wise identical. This insensitivity is mainly the result of appraising situational *change* instead of the situation itself. Appraisal-results are identical assumed that the difference between the frequency of both instrumentations is not so large that the slower-frequency-instrumentation completely skips both the delivery and the retraction of an event from its current "blackboard". A large difference is thus a risk for appraising situational change, but many ways exist to solve this problem using more sophisticated "event delivery". This shows that our approach supports flexible triggering.

The potential of our approach for runtime-scalability related to quality/performance trade-off is indicated by the fact that the "points" bank actually fills-in the non-emotional episodes of the "survival" bank. During run-time the "points" bank can be switched off, still resulting in meaningful but less detailed emotions, as shown by Figure 3. Of course in our case both banks consume virtually no resources, but in a situation where two different appraisal mechanisms are used, this runtime-scalability becomes useful.

## CONCLUSION

In this paper we have addressed the problem of incremental emotional instrumentation of systems. That is, how to develop computational models of emotion based on a step-by-step addition of sophistication to a such a model resulting in meaningful and more sophisticated emotions of the agent that is consistent with the emotions resulting from the simpler models. We have proposed the FeelMe framework (DeGroot 2004) as a solution to this problem. The FeelMe framework is a modular, signal-based approach to computational emotions. In this paper we have focussed on the Appraisal System in the FeelMe framework. Context-sensitive appraisal banks are introduced to facilitate the development of scalable appraisal models. The results of an experiment we have conducted with a game agent show the following. An appraisal model using two appraisal banks - the first being sensitive to all events related to survival and the second being sensitive to all events related to collecting points - results in more sophisticated emotions than an appraisal model with just the "survival" bank. The "survival" appraisal bank in our experiment inhibits the "points" appraisal bank. This inhibition provides consistency between the two instrumentations. Consistency between appraisal models and incremental emotional sophistication are two of the requirements for model-scalability and runtime-scalability, indicating that context sensitive appraisal banks - enabled by our signal based approach - facilitate the development of scalable appraisal models.

Runtime-scaling of appraisal models is useful in domains in which computation time is an unknown factor, because it enables trading-off emotional quality with available computation time, just like 3D-graphic-engines and chess engines. The experiment also indicates that our signal-based approach is flexible regarding the frequency of appraisal. This flexibility enables a different form of runtime-scaling, namely scaling based on a trade-off between temporal quality and performance.

Even though the number of appraisal banks was small in order to test scalability, we think that these results show that our modular, signal-based approach to computational models of emotion has many benefits for the gaming and virtual agent arena.

## FURTHER WORK

Possible extensions of our dynamic approach to computational emotions include modelling the mood of an NPC, modelling the effect mood can have on the emotional state, and the use of our approach in multi-agent environments.

## REFERENCES


Baillie-de Byl, Penny. 2003 "A Six-Dimensional Paradigm For Generating Emotions in Virtual Characters." *IJIGS 2(2)*: 72-79.

Chow, Benny. PacMan. 2003 http://www.bennychow.com/pacman_redirect.shtml.

DeGroot, Doug and Broekens, Joost. 2003. "Using Negative Emotions to Impair Gameplay." *Proc. of the BNAIC:* 99-106.

DeGroot, Doug. 2004 "Integrating Emotional Traits and States in Virtual Humans", in preparation.

Jennings, N. and Sycara, K. and Wooldridge, M. 1998. "A Roadmap of Agent Research and Development," *Autonomous Agents and Multi-Agent Systems 1(1)*: 7-38.

Mac Namee, Brian and Cunningham, Padraig. 2003. "Creating Socially Interactive Non Player Characters: The m-SIC System." *IJIGS 2(1)*: 28-35.

Marsella S. and Gratch J..2001 "Modeling the Interplay of Emotions and Plans in Multi-Agent Simulations." *Proc. of the 23rd Annual Conference of the Cognitive Science Society*.

Mehrabian, Albert. 1980. *Basic Dimensions for a General Psychological Theory*. OG&H Publishers. Cambridge Massachusetts.

Ortony, A. Clore G.L., Collins A. 1998. *The Cognitive Structure of Emotions*. Cambridge University Press.

Reekum, van C. 2000. *Levels of Processing in Appraisal: Evidence from Computer Game Generated Emotions*. Phd Thesis nr. 289, University de Geneve, Section de Psychology

Reisenzein, R. 1994. "Pleasure-Arousal Theory and the Intensity of Emotions." *Journal of Personality and Social Psychology 67(3)*: 525-539.

Reisenzein, R. 2001. "Appraisal Processes Conceptualized from a Schema-Theoretic Perspective: Contributions to a Process Analysis of Emotions." In *K.R. Scherer, A. Schorr and T. Johnstone (eds), Appraisal processes in emotion: Theory, Methods, Research*. Oxford University Press.

Russell, J.A. 2003. "Core Affect and the Psychological Construction of Emotion." *Psychological Review 110(1)*: 145-172.

Schimmack, U., Colcombe, S., & Crites, S. Pleasure and displeasure in reaction to conflicting picture pairs: Examining appealingness and appallingness appraisals. *Unpublished Manuscript*. 2001.

Smith, Craig A., Kirby, Leslie D. 2000. "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.

Scherer, K. R. 1993. "Studying the Emotion-Antecedent Appraisal Process: An Expert System Approach." *Cognition and Emotion 7(3/4)*: 325-355.

Scherer, Klaus R. 2001. "Appraisal Considered as a Process of Multilevel Sequential Checking." In*: K.R. Scherer, A. Schorr and T. Johnstone (eds), Appraisal processes in emotion: Theory, Methods, Research*. Oxford University Press