Using Negative Emotions to Impair Game Play

Doug DeGroot and Joost Broekens

Leiden Institute of Advanced Computer Science

Leiden University, Leiden, The Netherlands

Abstract

Emotions make significant contributions to rational thought and behavior at multiple levels. Most research with computational models of emotion focuses on reactive behaviors caused by base or learned emotions; yet emotions can play a role in purposeful behaviors as well, leading to emotion-related actions rather than reactions. Higher-level cognitive processes can reflect on the emotional state, treating emotions as data rather than forces, and incorporate them into consciously selected behaviors, plans, and strategies. We have developed a computational framework for exploring cognitive decision processes that reason about emotional actions as an integral component of strategy and control over emotions. To validate this framework, we are implementing a prototype gaming system that exhibits conscious use of the emotional state and negative emotional behaviors during a game of chess in an attempt to impair its opponent's game play.

1. Introduction

There is increasing interest in the integration of models of computational emotions with models of cognition and behavior — especially within the arenas of cognitive neuroscience and intelligent virtual humans and agents — as it is becoming increasingly accepted that emotions comprise a significant component of rational thinking and human behavior. Most of the research with intelligent virtual humans, however, has focused on emotionally reactive, behavioral response mechanisms in an attempt to produce consistent, believable animations of virtual humans, including facial movements, gestures, voice, etc. In these models, the goal of computational emotions is to achieve increased "believability" of the agents, increased enjoyment from interacting with the agents, or to achieve a more meaningful, value-added communication channel between the human and the agent [1]. This is *communication-driven* emotion modeling [2]. Other research efforts have explored the use of emotional stimulus evaluation (appraisal) to maintain representations of an agent's emotional state. The state changes in relation to the agent's evaluation of external events compared to its internal desires and goals. The resulting emotional states are used to influence the action decisions of the agent, to modify its beliefs, or to influence its reasoning and planning activities. This is *appraisal-driven* emotional modeling [2].

In essentially all these efforts, the main goal of incorporating emotions is achieving higher performance and/or higher satisfaction levels on the part of the human through increased believability and emotional intelligence of the agent. To make agents more predictable and understandable by the humans interacting with them, the focus is on the direct, positive correlation of emotions with behaviors, wherein positive (negative) emotions result in positive (negative) behaviors.

Things are a bit different in the computer gaming arena — but not much. Emotional modeling in games is also used primarily to make computer characters behave more "believably". Unlike intelligent virtual agent applications, though, computer games are often designed specifically to beat their human opponents. They attempt to reduce the human's performance level rather than increase it. However, even here, the focus has remained on positively correlated, believable behaviors.

2. A Shift to Negative Emotional Behaviors

We are exploring the utility of negatively correlated emotional behaviors on the part of a game-playing software robot in order to determine whether and how they might prove effective in emotionally impairing a human opponent's game play. We use computational emotions neither to improve the human's performance nor his satisfaction with the game; rather, our goal is achieving better game play on the part of the computer at the expense of the human.

Our research involves two fundamental shifts in focus from these other efforts. First, we integrate appraisal-driven with communication-driven emotion modeling. The effect is that the agent's own emotion-driven actions and behaviors loop back to influence its own desires, plans, and beliefs. Additionally, its conscious plans and desires can elicit actions on its part that are intended specifically to influence both its own emotional state and that of its human opponents. Our model thus embodies an "end-to-end" approach. This is similar to the approach of Franklin's Conscious-Mattie [6]. Second, unlike C-Mattie, we support negatively coupled emotions and behaviors, conscious, purposeful control of emotions, and dissimulation.

We have developed a general-purpose framework – the NTIM (for "intimidation") framework – that incorporates a computational model of emotions into a self-aware reasoning engine. The reasoning engine is capable of emotional-response control and purposeful selection of behavioral responses. We have implemented a prototype emotional, chess-playing agent based on this framework in order to explore the cognitive aspects of conscious control of emotional behavior. We refer to the prototype as the NTIM game-playing robot below.

Rather than employing emotions in a direct, positive feedback loop to achieve behavioral homeostatis [3], our use of emotions is in expressing negative emotions and behaviors in an attempt to defeat the human more resoundingly. We strive to *disrupt* the emotional channel. This requires cognitive processing of emotions as data rather than behavioral forces. Specifically, our use of negative emotional behavior is intended to imbue within the human opponent a state of negative emotions that will impair the human's play, thereby reducing both his performance and his satisfaction levels. We believe this inverse use of computational emotions is novel.

3. Chess and Emotions

We selected Chess as the initial application within which to study our framework. There are several reasons for this choice. Foremost among these is that chess has stood as the quintessential game for showcasing artificial intelligence over the past 50 years. Numerous chess playing systems have been explored during this time, but all with a complete reliance on analytical capabilities for playing and winning the game. Emotions have not yet played a role in the battle — at least not on the part of the computer.

However, when Kasparov lost to Deep Blue in 1997, emotions did play a role — they even proved to be paramount. While Kasparov is reputedly adept at imposing emotional pressure on his opponents, he is naturally (currently) unable to intimidate a computer; but in an ironic twist of fate during the 1997 match, several unexpected behaviors by Deep Blue caused Kasparov to unwittingly "turn on himself". Following the match, Kasparov admitted to suffering significant emotional distress – even paranoia – as a result of Deep Blue's unexpected behavior [4]. Kasparov's resulting self-imposed distress was partially responsible for his losing the match, and in essence he partially defeated himself. Clearly emotions played a key role in the game, albeit in a very surprising and perhaps totally unexpected manner.

Our second motivation for selecting Chess, then, was to explore the possibility of using negative emotional behaviors (e.g., intimidation) on the part of the computer to achieve similar self-impairing reactions on the part of human opponents. Even if the computer were capable of beating the human without resorting to negative behaviors, doing so might make winning take less time or prove more spectacular. Finally, because chess is a board game rather than an immersive, virtual reality game, we would be limited to low bandwidth in the emotional communication channels [1], resulting in stricter success criteria. In our initial prototype, we have limited our model to vocal behaviors only.

4. Combining Emotional Traits and States

We chose Mehrabian's PAD Temperament Model for the basis of our computational emotion system [5]. This model decomposes emotions into three distinct components — pleasure, arousal, and dominance — each of which is represented numerically in a normalized scale from -1 to +1. Because the three components are (nearly) orthogonal, they can be used as the basis for a 3-dimensional vector space into which various human emotions can be mapped. Figure 1 illustrates this space and lists numeric P-A-D values for several of the more than 240 emotions defined for Mehrabian's model.

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).

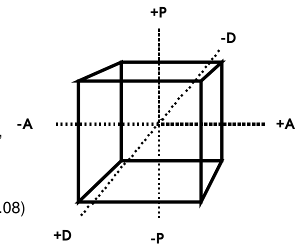

The emotional state "angry" is a highly unpleasant, highly aroused, and moderately dominant emotional state. The "bored" state implies a highly unpleasant, highly unaroused, and moderately submissive state.

Figure 1: The Mehrabian P-A-D Temperament Scale
From: Albert Mehrabian's PAD Scales (see [5]).

Although Mehrabian's PAD system is a largely static system and is used for characterization and classification purposes, we developed a dynamic PAD-inspired model that continually computes the robot's changing emotional state and biases it by the robot's set of long-term, emotional traits. Emotional states are transitory; they can vary substantially and rapidly throughout a game; an emotional trait is a long-term predisposition to be in a particular emotional state.

5. The NTIM Framework

The NTIM framework integrates a traditional chess-playing program with an overarching, emotionally aware behavioral model. It incorporates parameterizable, personality-specific belief models, dynamic emotional states, scalable affective behaviors, and model-based reasoning (see

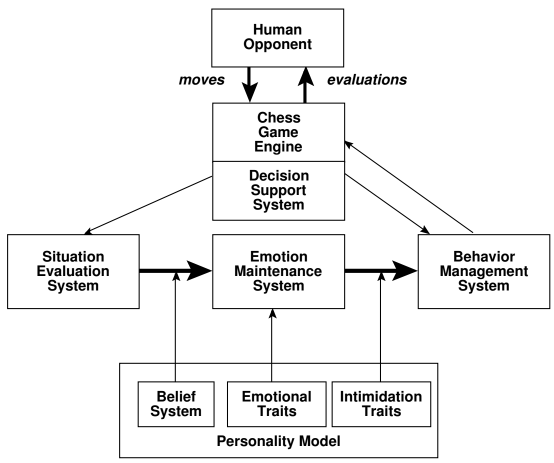

Figure 2: High-Level View of the NTIM Architecture

Figure 2). The NTIM framework contains the following key components:

For the **Chess Playing Engine**, we chose to adopt an “off the shelf” chess program and to tinker with it as little as possible. This greatly simplifies our situational appraisal tasks since the chess engine already computes numerous performance variables and statistics. The **Decision Support System** simply provides mediated access to these variables and game statistics.

The **Situation-Appraisal Manager** performs stimulus evaluation and situational appraisal. Given the computed performance variables and statistics, it “thinks” about the current game situation (board, control, score, etc.) and decides what it “means” to the robot and how the robot “feels” about it -- good or bad. These feelings update the emotional state. The **Emotional-State Manager** maintains the robot’s emotional state. Game play and situation-appraisal induce changes in the emotional state, but the ESM uses the robot’s temperament to ensure that the robot’s temporary states are kept consistent with its personality.

The **Behavior Management System** selects, controls, and expresses the robot’s negative, intimidating behaviors. The behavioral choices are purposeful; they are based on the robot’s emotional state and beliefs that represents the robot’s thoughts on how intimidation best works.

A **Personality Specification Module** can be used to tune the NTIM cognitive, emotional, and behavioral systems. It consists of two separate personality components consisting of “thought amplifiers” — sets of numeric weights used to modify the base signals — and the emotional temperament model — specifications of the robot’s emotional traits and their associated strengths. [7]

6. Situational Influence on Emotions

The Situation-Appraisal Manager (SAM) performs low-level, unconscious stimulus evaluations and situational appraisals. It symbolically “thinks about” the current situation and decides what it “means” to the robot. This includes evaluating the current board position, the way the game is proceeding, and the like. These and other performance variables calculated by the chess engine are treated as stimuli by the SAM and converted to situational variables for use in a situational calculus. The situational calculus is used to define potential low-level “thoughts” and “impressions” on the part of the robot; these thoughts affect the robot’s emotional state.

Situational appraisal may also involve complex beliefs and memories of game positions and play. We limit these to learned emotional beliefs, and thus we treat them as additional stimuli as well. In this way, though, higher-level cognitive inputs can influence the situational appraisals.

In an attempt to maintain cognitive and neurological plausibility, we model the result of each situational assessment function as a set of signal strengths, or here, as a triple of PAD deltas — potential changes to the system's current emotional state. These objective PAD deltas are weighted by the robot's personality-specific "thought amplifiers" to convert them into subjective PAD deltas. These are then averaged using a weighted sum and passed on to the Emotional State Manager as a single three-component vector for use in updating the robot's current emotional state.

7. The Emotional-State Manager

The Emotional-State Manager (ESM) maintains internal PAD-scale representations of both the robot's static temperament and its dynamic, emotional state. The emotional-state PAD model is biased by the temperament model. We use the following mechanisms to implement this:

- a personality-specific emotional temperament, specified as a set of points and densities in the 3-dimensional PAD space [5]
- an update function that gradually 'walks' through the 3D space, dampened by the temperament model; thus dynamic emotional swings stay relatively consistent with the personality profile
- an emotional state fading mechanism that slowly moves toward the [0,0,0] neutral state when the SAM produces no emotional input

During game play, the SAM reaches various conclusions about how it "feels" about the game; these feelings are represented as PAD deltas and sent to the ESM. The ESM uses these deltas and the current emotional state to compute a new dampened, emotional state. If the SAM does not know what to think about a situation, the emotional state fading mechanism kicks in. All this results in dampened emotional swings, emotional states that fade over time, and emotional states that remain reasonably consistent in terms of the personality's emotional traits.

8. Strategic Use of Negative Emotional Behaviors

We are initially exploring two simple cases in which negative, intimidating behaviors can be exhibited by the robot. First, losing positions may cause it to become angry and choose to rant and rave at the human — to express displeasure, to hopefully upset the human's future play, or whatever. Second, winning situations might make it gleeful and choose to gloat or ridicule the human. The human opponent would likely perceive both behaviors as negative, and either or both might serve to impair his playing ability. However, the first type of behavior might inadvertently please the human, since in this case the human would already likely perceive himself to be in the superior situation. Observing the robot's pathetic exhortations at losing might at that point serve only to encourage the human and to improve his play. We chose to focus on the second case. Because the human will likely have already perceived that he is in a losing situation, his emotional state is likely to already be negatively oriented at the moment. Thus if the robot chooses to then gloat or ridicule the human, it is in essence "piling on", and the human is more likely to suffer further negative shifts in his emotional state — or so we hope.

Thus we adopt the fundamental assumption that grousing when the human is ahead is more likely to reinforce the human's already positive emotional state and make his play stronger, whereas gloating when the human is behind is more likely to reinforce the human's already negative emotional state and make his play weaker

The NTIM framework contains a model-based method of purposeful intimidation of the human opponent. It controls the type and frequency of exhibited behaviors in an attempt to protect against inappropriately exhibiting negative behavior or exhibiting too much or too little negative behavior. Were it not to do so, then the human may become inured to the negative behavior and recover from its impact. Because we do not yet have empirical results, the NTIM framework's belief model is parameterizable through the personality "thought amplifiers".

9. The Behavior Management System

Once the robot's emotional state has been updated, the Behavior Management System (BMS) decides whether and how to act out in its attempt to intimidate the user. The determination is a function of the robot's emotional state, a belief model of how intimidation best works, and the history of recent intimidation actions. Even though an emotional state may predispose the robot to certain types of behaviors, these can be controlled, over-ridden, or modified; further, the robot can dissimulate.

Our current behavioral repertoire consists only of negative vocalizations (comments, sighs, finger tapping, etc.) recorded as sound files. They are played on cue when the BMS determines that a specific type of negative behavior should be invoked, such as contempt, impatience, disbelief, etc. Each negative behavior type has associated with it one or more negative vocalization files; this helps minimize potential repetitious behavior. We can also record and select from multiple levels of emotional intensity for each negative behavior.

The BMS in the NTIM framework is not a simple "reactive" translator of emotional states to emotional behaviors. It is a self-reflective, cognitive part of the computational emotional process, able to consciously access the underlying primary emotional evaluation (i.e. the ESM), and other data from the system to actively choose an appropriate behavior within the context of a goal (e.g. intimidating). Therefore, the emotional state from the ESM predisposes the possible set of behaviors, it certainly does not prescribe it, as would be the case in systems implementing a more reactive emotional expression.

10. Summary

We designed the NTIM framework to study the integrated use of emotions in both reactive and conscious, purposeful behaviors. The core of this framework is a computational model of emotions that dynamically integrates emotional states with emotional traits. The framework also incorporates separate appraisal and behavioral modules that are loosely coupled to the emotional maintenance system. To explore the NTIM framework, we are developing a conscious chess-playing robot whose goal is to reason about its own internal emotional state and select appropriate negative behavioral responses that it believes will impair the game playing skills of its human opponent. Our initial prototype is limited to emotional vocal behaviors only; these are expressed by the robot as an integral part of and during the actual playing of the game of chess.

REFERENCES

1. Johnson, Colin and Jones, G.: Design precepts for incorporating affect into intelligent virtual agents, May 2000. (http://www.cs.kent.ac.uk/people/staff/cgj/papers/vr.ps)

2. Marsella, Stacy and Gratch, Jonathan: A step toward irrationality: using emotion to change belief, First Int'l Conf. on Autonomous Agents and Multiagent Systems, Bologna, ACM, July 2002.

3. Plutchik, R.: *The Emotions*, Univ. Press of America, 1991.

4. Chelminski, Rudy: This Time It's Personal: Humankind battles to reclaim the chess-playing championship of the world, *Wired*, Issue 9.10, Oct. 2001.

5. Mehrabian, Albert: Pleasure-Arousal-Dominance: A General Framework for Describing and Measuring Individual Differences in Temperament, *Current Psychology*, Winter, 1996, V14, N4, pp. 261-292.

6. McCauley, Thomas L. and Franklin, Stan: An Architecture for Emotion, AAAI 1998 Fall Symposium, *Emotional and Intelligent: The Tangled Knot of Cognition*, AAAI, 1998.

7. DeGroot, Doug and Broekens, Joost: Integrating Emotional Traits and States in Virtual Humans, in preparation.