Emotion-Based Framework for Multi-Agent Coordination and Individual Performance in a Goal-Directed Environment

Giovanni Vincenti

Computer and Information Sciences
Towson University
Towson, MD
giovanni.vincenti@gmail.com

James Braman

Computer and Information Sciences
Towson University
Towson, MD
Jbraman@towson.edu

Goran Trajkovski

Information Technology Programs
South University
Savannah, GA 31406
gtrajkovski@southuniversity.edu

Abstract

Through the use of a multi-agent system composed of emotionally enhanced sets of agents, we investigate how emotion plays a role in inter-agent communication, cooperation, goal achievement and perception. The behaviors of each agent are derived from its perception of the environment, its goals, knowledge gained from other agents and based on its own current emotional state, which is predicated on fuzzy sets. The state of each interacting agent is determined by its level of frustration combined with its interaction with other agents. The set of actions an agent may perform including that of its own ability to sense and understand the environment is limited by its current level of emotional context. In this work we focus on the analysis of the interaction of agents and review the grouping effect that arises from the inter-agent communications combined with the intra-agent emotional status as they perform the task required by the environment. The environment is based on previous work performed on navigational map learning by context chaining and abstraction.

Introduction

The extent of a person's current emotional state can affect one's behavior in unpredictable and varying degrees. This unpredictability in actions can lead to various levels instability when interacting with other humans or in general, other agents. Making decisions biased by an emotional condition can produce erratic, impulsive or risky decision making behaviors within a given context (Loewenstein, Weber, Hsee, & Welch, 2001). Various factors can contribute in eliciting such states and can be influenced by events in the environment, mental defect or disease, genetic disposition, traumatic events, social interactions or based from ones own perceptions. (Selyse & Fortier 1950; Loewenstein et al., 2001; Ohman & Wiens, 2004). Our aim is not to simulate complex emotions in an agent but to investigate how simple emotional bias lays the foundation for affecting goal-seeking agents within a simple environment. We envision such phenomena within a homogenous multi-agent system composed of emotionally enhanced sets of agents given both a finite set of options and emotional states. Following previous experimentation with limited perceptual context for a given agent and its combined effect on understanding and formation of personal goals, we now apply emotion in limiting an agent's perception and motivational attributes (Trajkovski, Collins, Braman & Goldberg, 2006). Emotions for our agents directly affect their ability to sense and interact within the world they inhabit. These states are derived from their success or failure in finding goal locations. The longer the span of time where no goals have been found, compounded with an agent running into many obstacles along its path, contributes to eliciting a stressful or angry response. An Agent may also be influenced by other agents within its perceptual field. We see emotions as a dynamic influence over response mechanisms in a given agent. It can be said that emotions can "cloud" our ability to make rational decisions which implies that they have a tendency to interfere with rational thinking and our ability to interpret perceptual information (Artz, 2000). In other situations they can however be extremely useful in making certain decisions "by rapidly reducing the options that one can consider" (Greenberg, 2002). It is these changes in perception and action that we are exploring by attributing certain basic emotions to goal seeking agents and studying its effect. Often an emotional response is induced by an event or an "object" that has been given meaning which is part of a stimulus interpreted by an agent. With various stimuli are attributed meanings which are a result of an appraisal process which derives significance to such stimuli or events (Planalp, 1999). Objects, events and interactions are interpreted by each agent which contributes to their particular emotional state. Following the distinct emotional conditions established by Elkman & Friesen (1975) which identify six emotional states that are innate across cultures, which are based on facial expressions (Anger, Fear, Sadness, Disgust, Surprise and Joy) we have chosen two focal emotions. From these basic states, anger and joy (or happiness) has been selected for our research. These two states can be attributed to individual factors concerning the achievement of goals which is a key component of this research (Planalp, 1999).

In general people who are in a less stressed or in a happier state of emotion are less inclined to make riskier judgments. People in a happier state would not wish to

take actions with risky or potentially negative outcomes so as not to disrupt their current positive state (Isen, Nygren & Ashby, 1988). An angry person or someone in a "bad" mood state is more likely to make poor judgments (Loewenstein et al, 2001). Anger often can influence us to act in ways that are not in our best interest (Borcherdt, 1993). Fear and anxiety often too play a role in behavior patterns as one avoids the object or causal of such stress. Depression as related to stressors can also distort judgment and the interpretation of perceived information as obtained from the environment (Gotlib, 1983). We envision the application of these emotional conditions to our artificial agents and apply these concepts as a basis of our experimentation.

Agents in our experiments will be limited to emotions ranging between two finite states (anger and happiness) which are based on fuzzy logic. An agent can be in any state between full happiness and complete anger where a numerical value is assigned to denote such state. We propose that each agent's range of perception and actions be limited based of their current emotional condition such that angry agents have the least options available compared to happy agents. In a broader sense we are limiting an angry agent's ability to make poor judgment while giving a happier (thus less influenced) agent more ability and control within the environment. With the simulation of these properties for each agent we investigate the possibility of emergent properties regarding inter-agent coordination in a goal directed system.

Existing Frameworks

The idea of adding an emotional component to agents is one that has been given much discussion. Bates (1994) first brings up the concept that it is important for an agent to have an emotional component because humans tend to relate more to what is like them.

The work that worked as the spark that initiated this branch of research comes with Damasio (1994). Since then, many research groups worked on creating emotional engines that resemble, more or less, real-life behaviors. It is important to realize though that there are two schools of thought when it comes to emotional engines for agents. The first incorporates the findings of Damasio (1994) into the models and tries to recreate an environment that is inspired directly by his work. The second instead creates frameworks that may take some concepts from Damasio's work, but are based on independent thinking.

Damasio-based frameworks

Although many researchers have followed the steps first left by Damasio (1994), we find in the work by Sloman (1998) the most interesting interpretation. Groups such as Ventura and Pinto-Ferreira (1998), Velásquez (1997), Gadanho and Hallam (1998) and others focus solely on the Damasio approach, Sloman (1998) focuses on a model that is created

through the evolution of the capabilities of life forms through history and pre-history.

Sloman (1998) affirms that there is a need to create control systems that do not rely on a fixed architecture and changing values, but more dynamic frameworks. This idea is motivated by the fact that humans have control structures that are dynamic in nature, both at the conscious as well as unconscious level. Sloman (1998) also states that, within our minds, we have other modules, some of which deal with inputs, others with outputs, and yet others with processing information. This does not necessarily mean that each module has a different architecture and behavior, or that a single module cannot take care of multiple functions. It does mean though that the overall functionality is achieved through the interaction of multiple components.

The model created by Sloman (1998) is the summary of the analysis of several fields unrelated to robotics and agency. He bases his main notions on evolution and the adaptation of the human mind to ever-changing natural conditions. The fields of biology, philosophy, psychology and many more all contributed to the refinement of our understanding of human control modules.

The first model that Sloman (1998) analyzes is the reactive model. In this particular model, which characterized many organisms through their evolution, deals with relatively simple reactions based on inputs. This model is characterized by outputs, or actions, that are generated by processes that do not take into consideration ramifications of behaviors and foresight. This means that this particular model lacks a planning module. New behaviors can be learned through positive or negative reinforcement, but they cannot be generated as an internal response to a set of conditions that recall past experiences. The motivations that drive such control mechanisms are instincts, such as hunger, fear or mating. Reactive models are also designed to work in conjunction with other organisms; great examples of such collaborative beings are insects.

The evolution of a reactive model is the addition to this system of deliberative models (Sloman, 1998). This component elaborates long-term memories and creates new plans and drives, to add to the ones built into the reactive system through evolution and conditioning. Besides a higher processing layer, this particular model requires the presence of reusable memory that can be accessed to store and retrieve information as the controller performs its functions. Sloman (1998) also affirms the necessity for a quick-response module, as careful evaluation of a situation is not always the ideal response to the environment. In dangerous situations, for example, an alarm module would take over and perform the necessary operations to place the organism back into safety, and then higher functions can process the information just received.

Damasio (1994) states that there are two kinds of emotions. Primary emotions are generated by external or internal stimulations of sense organs, and secondary emotions instead are generated by the cognitive system. When compared to the deliberative model (Sloman, 1998), we can see some similarities between the emotion system described

by Damasio (1994) and Sloman's alarm module. Damasio (1994) also states that secondary emotions always trigger higher cognitive reactions that, in turn, generate psychological changes. This analysis requires a slight modification of the alarm system, which should be divided into a section that deals with emergency situations as the deliberative model, and a second section that deals with similar situations, but interacts only with the deliberative, or cognitive, module (Sloman, 1998). The author compares this new component of the alarm system to the process of growing up and the acquisition of maturity, as also supported by Goleman (1994).

Sloman's proposed architecture involves the controllers discussed so far with the addition of a meta-managegement mechanism (Sloman, 1998). The meta-management process performs functions that work at a higher level compared to the other modules. Such component performs reflective kinds of elaboration, both on events and actions linked to the environment and on the internal state of the agent. Sloman (1998) affirms that, should a robot be given such a module, it may then start reflecting on the concept of "self" and "others". Moreover, this particular system may be overridden by the alarm module in certain cases, thus it may shift the focus of the operations on other tasks or problems. Sloman (1998) also introduces the idea of tertiary emotions. These emotions are initiated purely cognitively and may or may not trigger other physical changes. For example, a state of infatuation may lead to a decreased level of attention and also sweating and tension.

Non-Damasio-based frameworks

The application of Camurri and Coglio (1998), based on the model created by Camurri et al. (1997), works in a setting of the performing arts, introduces a very interesting framework that is not based directly on the work by Damasio (1994). The agent created by the researchers works by observing and being emotionally influenced by a dancer, creating as a consequence outputs of music and rhythm based on its internal emotional state.

This engine is explained in greater detail in Camurri et al. (1997). A macroscopic analysis reveals several components of control. There are five active components: Input, Output, Reactive, Rational and Emotional.

The Input component is responsible for either gathering information from the environment or receiving communications from other agents or humans. In turn, it will analyze and format it as appropriate to pass on to the Reactive, Rational or Emotional components. Given the responsibility of the Input module to forward the inputs where appropriate, this component is required to access the rational and emotional states. Moreover, the Input component can also receive feedback directly from Output, thus taking also this information into consideration when relaying messages.

The Output component is responsible for creating outputs based on the agent's internal state. The outputs of this module are generated by processing the outputs of the Reactive and Rational components. Although the Emotional component does not feed the Output one directly, it

will influence the calculations of the overall output by signaling the agent's emotional state to this last component. The Output component is also responsible for an internal feedback mechanism that affects the Input controller, as reported earlier.

The Reactive module is responsible for the real-time behavior of the agent, which is necessary given its application in the world of music and dance. This component collaborates closely with the Input and Output modules, and its processing is modulated by both the rational as well as the emotional state of the agent.

The Rational component maintains a view of the external world as well as one of the agent, consisting of its goal. This component has no real-time type of operations, leaving it the possibility to perform rather complex operations. This component of the agent interacts with all the other modules in several ways. Perhaps what is most important for this discussion is the role that self-awareness in relation to the goal plays in the overall emotional state of the agent. As the Rational module detects that goals are being accomplished, the emotional state of the agent increments towards a positive attitude, thus creating a better "mood" that will, in turn, affect the rest of the operations.

Finally, the Emotional component is governed by emotional stimuli generated from Input, Reaction and Rational. This module contains an emotional space. Such space is divided into several sectors that identify different emotive states. As the emotional state of the agent changes, it can be traced within this emotional space, which reveals the state of the agent.

Our Approach

The first choice that we need to make when it comes to creating a new emotion-based framework is whether or not we will try to mimic a Damasio-based architecture. For this work we want to focus on the emotional side, thus leaving out meta-cognitive states that may influence the behavior of our agents. We decided to design a framework that would resemble more closely the one described by Camurri et al. (1997), pictured in Figure 1.
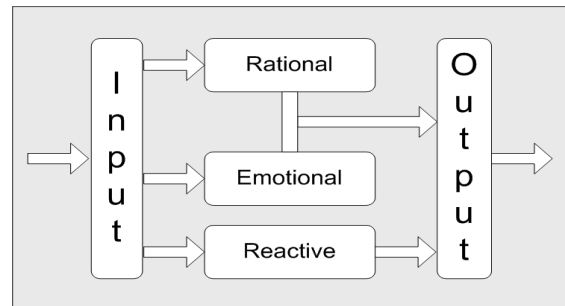

Figure 1: Architecture of the framework introduced in this work

We decided to keep all the components that Camurri et al. (1997) described. As we can see, our architecture contains a direct communication between the emotional mod-

ule and the output one. This is because we believe that the emotional state of an agent should modify the range of rational actions it can perform.

Camurri and Coglio (1998) state that, in order to implement an emotional component that evolves over time, it needs to be influenced by external events. We believe that it also needs an internal component, especially when operating in a multi-agent environment, as a corollary to Camurri and Coglio (1998). For this reason we created an elaborate Emotional module, composed of four components: Genetic, Internal, Interagent Verbal and Interagent Non-Verbal. The interaction among these components is shown in Figure 2.
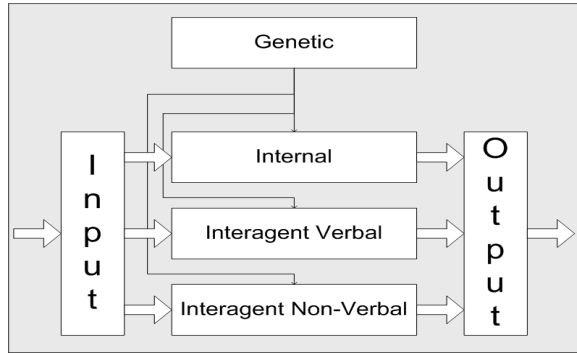

Figure 2: Detailed view of the "Emotional" module

The "Genetic" component does not have a direct influence on the calculation of the overall output based on the input sent to the module. Instead, it will define the weights to assign to the other components. It is our experience that some people seem to be influenced more by what others say, while some other people may be extremely emotionally unstable from causes internal to their person. We felt the need to analyze this relationship among factors using this architecture.

The "Internal" component analyzes information about the agent and elaborates the emotional state. An example of such analysis may be the time elapsed in the simulation, with the agent unable to either find clues or reach the overall goal. As time goes on, the morale of the agent may lower. The rate at which the morale of the agent is affected by internal events is dictated by its genetic predisposition.

The "Interagent Verbal" component instead relies on communications with other agents in order to influence the mood. Communications between agents will consist of clues that they will pass along. Each clue will be dictated by where the agent "believes" the goal is. Along with the belief, there is a weight assigned to the communication. The weight indicates the emotional state of the agent that is communicating the information. The agent will then internalize both the belief as well as the emotional state of the other agent. The emotive component of the communication will affect the agent's own emotional state at a rate dictated by the genetic component.

Finally, the "Interagent Non-Verbal" component relies on the concept that, in society, it is often easy to be able to

guess what mood a person is in by simply looking at them. In our framework, we believe that our agents should not only display their emotional state through communications, but they should also display it visually. When an agent senses the presence of other agents and then interacts with them through communication or simple proximity, then this component will analyze and read the apparent emotional state of its peers. Also in this case, the rate at which the agent will internalize emotions is set by the genetic component.

Preliminary Investigation

To investigate the emotional affects on coordination and interactions of agents in a simple system, we are developing an environment based on a simple two-dimensional world. This space will allow for agents to interact, communicate and react under similar conditions within the same environmental context. Agents will then be placed randomly within the environment and initiated with similar values for all internal variables and with the same level of understanding of the world. They will have limited knowledge of the space and must build their own intrinsic representation of the environment based on their own perceptions and emotional condition which changes during the course of interaction. The environment itself will not necessarily be static, as a dynamic space can be used to simulate a more natural environment. Obstacles within the environment will also be randomly created and distributed with the areas randomly marked as goals. The agents will only know that they have a certain level of perception and possible actions and that they are seeking a specific goal state and that they are able to communicate with other agents of their own type (if working in teams for instance). In a heterogeneous environment, other agents of different types will only be viewed as a dynamic obstacle within the space (Trajkovski, 2007). When an agent comes into perceptual range of another agent (within their particular scope allowed at their current emotional state) an agent may communicate his/her current state/condition, number of goals previously encountered and any relative information that may provide to be useful in the other agents quest to find a goal (i.e. general direction or location of the last encountered goal). An agent however does not keep track of all data for all goals that it encounters but is limited amount as feasible by memory size of an agent's associative memory. An agent also only communicates certain elements to other agents depending on the emotion of the communicating agent. Following experiments by Dabirsiaghi, we are integrating certain elements of context chaining and abstraction as a method for agents communicating certain directional data in search of previously encountered goals (2006).

The agent's ability to traverse the space is based on their current level of emotion. The happier the agent the more directions it can take in exploration and goal seeking (both possible modes of traversal). As shown in Figure 3, a happy agent has sixteen possible movements. An agent

that is in a transitional state has at least eight (N,S,E,W,NE,NW,SE,SW) and depending on the transitional status into another emotional state may have one or more or the directional capabilities as denoted by the dashed lines. The amount of these extra directions is based on how far the transition has progressed to the next state (see Figure 4); the particular directions that are chosen from the given set are chosen at random. An agent in an angry state shown at the bottom of Figure 3 is limited to only movements only in the four cardinal directions. We see this limitation on ability tied to that of emotion and motivation where an agent that is happier is more motivated thus having more options available. Anger can be said to be closely related to that of depression (Pelusi, 2006). We also see this in a wider scope and relate depression with that to a lessening of motivation in agents. It has also been observed in human subjects with depressive disorders that their cognitive flexibility is reduced due to emotional instability and negative thought patterns leading to reduced solutions to given problems. (Deveney & Deldin, 2006). Applying this to agent interactions we have limited the actions and perception of angry agents.
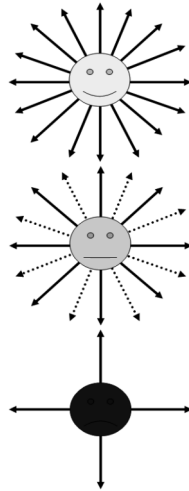

Figure 3: Three levels of Agent Directionality (In order: Happy State, Stressed State and Angry State)
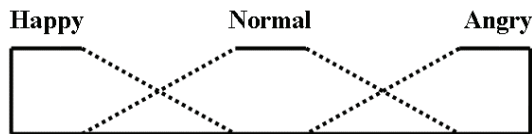

Figure 4: Transitions between States

During an agent's exploration of the environment it will traverse the environment in either two modes: 1. Traveling 2. Exploration. In traveling mode agents simply traverse the space within the limitations of their emotional state as it applies to their ability to move. An agent may switch to exploration mode if an object, goal or other agent falls within the sensing layer of their perception (see Figure 5). If an object is detected within this sphere an agent may switch from simply traversal, but to an exploration state so

that the agent may investigate the type of sensed object. If an obstacle is detected the agent may simply move to avoid it or follow a path around it to continue in a particular direction. If a goal is detected then it may record this event. If another agent is detected it will attempt to communicate with the detected agent. Limitations in the agents mobility based on their emotional state may make such activities more difficult.
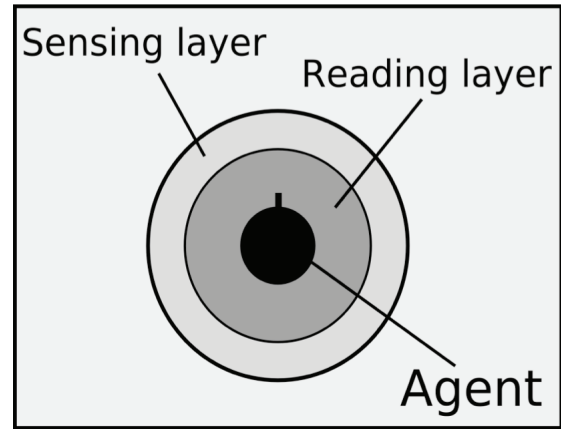

Figure 5: Perceptual Field

With this research we are trying to find answers to fundamental questions:

1. How can simulated emotions affect group dynamics and inter-agent cooperation?
2. What effects will these emotional states have on agent performance?
3. How can we successfully simulate simple emotional states in multi-agent systems?
4. Can the use of emotions benefit inter-agent communication and cooperation?
5. Will emotions benefit an agent's performance in satisfying their drive to find a goal?

Future Work

We will be conducting several sets of experiments changing various aspects of the agents and their environment. One experimental set up will investigate the performance of goal-seeking once a goal has been found by an agent where it can record such data into its associative memory and continue. The environment would then remove that found goal and may randomly place it in another location. This process would continue until the conclusion of the experiment. In another experiment, once a goal is found the agent is removed from the environment and another newly initialize agent is randomly placed back into the environment to take its place.

Communication differences are another aspect we are looking into further. Can we incorporate an agent's ability or propensity to communicate with other agents as a function of their current emotional state? An agent may be

more likely to communicate (and communicate truthfully) to other agents if they are in a happier state. Will an angry agent be less likely to communicate and more apt to "lie" to agents on opposing teams? As an agent traverses the space it keeps track of very generalized directional data in its associative memory to both help itself in the future perhaps but to also commutate helpful information to its own team of agents. Once an agent has found a goal it may communicate the knowledge to other agents in attempt to find other goals within the location. We envision two ways in which agents may communicate information once they are within perceptual range. One method is based on context chaining and abstraction where agent stores informational cues in the space where other agents including itself can use. These cues will help guide individual agent towards a goal. Another method is allowing each agent a certain finite memory space to serve as an associative memory where perceptual information can be stored. After a period of time and amount of movements within the environment data will be replaced with new information. This new stored data can be used as a guide for the agent as well as transmitted to other agents to assist in finding goals.

Future implementations of this project will account for team work and group dynamics within the multi-agent system. We plan to demonstrate the combined affects of emotion when the agents are divided into various teams and are in competition with other agents. Will opposing agents frustrate others by conveying misleading and false information to opposing agents? Will frustrations affect the entire team? Will they affect both teams if frustrations propagate throughout the whole system? Competing teams may have a sub-goal to purposely mislead opposing agents with false information or coordinate activities that hinder the efforts of the other team in their own effort to seek a goal. How will these types of activities affect the social interactions and group dynamics of the entire multi-agent system? After a certain amount of time has passed will the agent society stabilize in the sense of its emotional equilibrium and interaction? We believe that observing such automated behaviors that emerge from these systems can be beneficial in understating human group dynamics and interactions.

We are also looking into extending the research into several related components which are part of an agent's internal structure in developing emotional conditions. Some agents within the system may have predisposed mental or genetic defects in which certain interactions are hindered or exacerbated. We envision such defects as a way to simulate certain disorders in living organisms into artificial agents with which we may study the effects of interactions between defective and non defective agents. Can we also then simulate agents that produce offspring in which such defects are transmitted to their child agents. Certain mental diseases such as schizophrenia for instance have been shown to be related to genetics; some studies have found that some are predisposed to develop schizophrenia if parents and or siblings also develop the disorder even if they are raise in separate environments (Davison et al, 2004). In future projects we may allow for agents to reproduce and

possibly transmit information or "genetic" materials from parent to child agent. From these types of simulations we can also attempt to reproduce several symptoms of certain disorders in a way in which they can be controlled. For example we can have several symptoms i.e. hallucinations, paranoia, mania or nervousness that we can switch on, all at once or individually for various agents and then compare results depending on which disorders were active. Several complex models will need to be developed in order to accurately simulate such conditions within an agent. Hallucinations for example could be simply simulated by the introduction of random incorrect perceptual data or context chaining cues. This "hallucinated" data could affect individual agent's goal achievement but could also disrupt the team of agents if such data were communicated through the collective. We also plan to investigate the affects of environmental change, inter-agent interaction issues and various internal frustrations as elements that can adversely affect, change or disrupt agent emotional states and how the individual elements affect the multi-agent system itself.

Possible Applications

We see the application of this research to several areas of interest. In particular to the restriction of access or menu function based on the emotional state of the human user for a particular interface. Information is often delivered to users rapidly in ways that cause information overload and confusion; users typically want such information restructured and filtered with a minimum of effort (Shneiderman, 2005). As previously discussed users in certain emotional conditions may not be best suited to make critical unbiased decisions. If the interface automatically adjusts in limiting key components it would restrict the human user from taking actions that may have a negative impact or one in which he/she would later regret. We see an interface that adjusts to the user's emotional level to provide automatically the best displays and choices as to maximize their efforts. The human element of emotion in interface design is often overlooked. Consider an angry employee with a particularly high level of access or administration rights or an individual that may have access to key financial functions. An emotional driven interface could automatically account for these emotional contexts and shift the user's access to a lower and safer level at which they would be unable to take risky actions without consent from another user. It is known that many employees purposely put the goals of a company in jeopardy if they themselves are angry and feel vengeful towards their employer or even may involve themselves in a computer based crime in an attempt to "get even" (Quigley, 2002).

Another possible application of the concepts discussed in this paper consists of the analysis of traffic patterns. In a situation where the highway may be considered an environment of interactions among agents, and each car and driver is related to the concept of an agent, this framework could allow each agent to interact with others to be able to foresee emotion-based situations. An example would be

road rage, which this framework may be able to prevent by limiting the possible actions the driver can take, such as lower the maximum speed the car can reach, in certain emotional states. Research in this area by the authors of this work has already commenced.

Conclusions

This paper summarizes the basic framework for our investigation of the affects of emotional states of agents in a goal directed multi-agent system. We have also discussed the foundations of the theoretical background that supports the rational of the project as well as future applications that can be derived from experimentations.

We believe that this approach explores concepts that were so far left unattended. First of all, we take into consideration a genetic component that directs the inner workings of the emotional state. As a second focal point of this research, we believe that it is important to keep in mind that emotional states dictate the range of options of action that we have available.

Referemces

Artz, J. (2000) The Role of Emotion in Reason and its Implications for Computer Ethics. Computers and Society.

Bates, J. (1994). The Role of Emotion in Believable Agents. Communications of the ACM, 37(7): 122-125.

Borcherdt, B. (1993) You Can Control Your Feelings. Professional Resource Press. Sarasota, Florida.

Camurri, A. & Coglio, A. (1998) An Architecture for E-motional Agents. Multimedia, 5(4): 24-33.

Camurri, A., Coglio, A., Coletta, P. & Massucco, C. (1997). An Architecture for Multimodal Environment Agents. Proceedings of the International Workshop Kansei: Technology of Emotion, 48-53.

Dabirsiaghi, A. (2006) Navigational Map Learning by Context Chaining and Abstraction. Master's Thesis. Towson University, Towson Maryland.

Damasio, A. (1994). Descartes' Error – Emotion, Reason, and the Human Brain. New York, NY: Putnam Book.

Davison, G. Neale, J. Kring, A. (2004) Abnormal Psychology 9[th] Edition. John Wiley & Sons Inc. USA

Elkman, P, Friesen, W. (1975) Unmasking the face. Englewood Cliffs, NJ: Prentice Hall.

Gotlib, I. (1983). Perception and Recall of Interpersonal Feedback: Negative Bias in Depression. Cognitive Therapy and Research, Vol 7, No. 5. pp. 399-412.

Gadanho, S. & Hallam, J. (1998). Exploring the role of emotions in autonomous robot learning. In Cañamero, D., ed., Emotional and Intelligent: The Tangled Knot of Cognition, 84-89.

Goleman, D. (1994). Emotional Intelligence. New York, NY: Bantam Dell.

Greenberg, L. (2002) Emotion-focused Therapy: Coaching Clients to Work through Their Feelings. American Psychological Association. Washington, DC.

Isen, A. M., Nygren, T. E., & Ashby, F. G. (1988). Influence of positive affect on the subjective utility of gains and losses: It is just not worth the risk. Journal of Personality and Social Psychology, 55, 710-717.

Loewenstein, G. Weber, E. Hsee, C, Welch, N. (2001) Risk as Feelings. The American Psychological Association. Psychological Bulletin Vol. 127, No 2, 267-286.

Nermani, S. Allan, V. (2003) Agents and the Algebra of Emotion. AAMAS '03 14-18. Melbourne, Austria. ACM

Ohman, A. Wiens, S. (2004) The Concept of an Evolved Fear Module and Cognitive Theories of Anxiety. In Manstead A, Frijda, N, Fischer,A (Eds.), Feelings and Emotions, The Amsterdam Symposium. (pp. 58-80) Cambridge, UK.

Pelusi, N. (2006). Anger, Pain and Depression. Psychology Today.
http://psychologytoday.com/articles/index.php?term=pto3132.html&fromMod=popular_depression

Phanalp, S. (1999) Communicating Emotion. Social, Moral and Cultural Processes. Cambridge University Press. UK.

Quigley, A. Inside Job. (2002) Net Worker Volume 6. Issue 1 ACM Press New York, NY USA

Selye H, Fortier C. (1950) Adaptive reaction to stress. Psychosomatic Medicine; 12: 149–57. American Psychosomatic Society

Shneiderman, Ben. Plaisant, Catherine. Designing the user interface. Strategies for effective human-computer Interaction 4[th] edition. University of Maryland, College Park. Addison Wesley. 2005. 1

Sloman, A. (1998). Damasio, Descartes, Alarms and Meta-management. Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics, 2652-2657.

Trajkovski, G. Collins, S. Braman J. Goldberg, M. (2006) Coupling Human and Non-Human Agents. The AAAI Fall Symposium: Interaction and Emergent Phenomena in Societies of Agents. Arlington, VA.

Trajkovski, G. (2007) An Imitation-Based Approach to Modeling Homogenous Agents Societies. IGI Publishing. USA

Velásquez, J. (1997). Modeling emotions and other motivations in synthetic agents. Proceedings of AAAI-97.

Ventura, R. & Pinto-Ferreira, C. (1999). Emotion-based agents. Workshop of the Third International Conference on Autonomous Agents.