Submitted to the Workshop on Defence Applicaitons of Multi-Agent Systems (DAMAS05) at AAMAS05, Utrecht, Netherlands, July 2005.

Representing Dispositions and Emotions in Simulated Combat

H. Van Dyke Parunak, Robert Bisson, Sven Brueckner, Robert Matthews, John Sauter
Altarum Institute
van.parunak@altarum.org

1 Abstract

Emotion is an essential element of human behavior. Particularly in stressful situations such as combat, it is at least as important as rational analysis in determining a participant's behavior. Yet combat models routinely ignore this factor. DETT (Disposition, Emotion, Trigger, Tendency) is an environmentally mediated model of emotion that captures the essential features of the widely-used OCC model in a computationally tractable framework that can support large numbers of combatants. We motivate and describe this architecture, and report preliminary experiments that use it in simulating combat scenarios.

2 Introduction

Simulation and modeling are extensions of the experimental method for studying systems. Direct experimentation is the approach of choice when the systems in question are common (so that one can find instances for study), malleable (open to manipulation by the experimenter, so that one can fix some variables while varying others), and observable (so that one can see the results of experimental manipulations). When these conditions are not met in the real world, a computer model can provide them in a simulated one.

Real-world warfare lacks malleability and observability. Making deliberate changes in the face of combat is extremely difficult, a phenomenon that the Prussian General Carl von Clausewitz termed the "friction of war," and their outcome is obscure not only to those participating in the conflict, but sometimes also to historians years after, a situation he called the "fog of war." Thus it is not surprising that the military is one of the leading users of simulation and modeling, in understanding how combat situations develop and how commanders and troops should respond.

Emotion is an essential element of human behavior. Particularly in stressful situations such as combat, it is at least as important as rational analysis in determining a participant's behavior. Yet combat models routinely ignore this factor. DETT (Disposition, Emotion, Trigger, Tendency) is an environmentally mediated model of emotion that captures the essential features of the widely-used OCC model in a computationally tractable framework. A unique feature of our approach is the definition of a Disposition parameter to distinguish different agents' susceptibility to various emotions.

In Section 3, we review previous work in both combat modeling and computational emotions. Section 4 describes the DETT model. Section 5 reports on some experiments with the model. Section 6 concludes.

3 Previous Work

Our work is related to two bodies of previous work: a long tradition of computer-assisted combat modeling, and more recent research on computational models of emotion.

3.1 Combat Modeling

The roots of combat modeling go back well before the computer era, and follow two distinct lines, one mathematical and the other behavioral.

3.1.1 Mathematical Models

Mathematical models of combat fall into two broad categories: Lanchester theory, and game theory.

Lanchester Theory.—In 1916, F.W. Lanchester published a set of differential equations that expressed how the change in strength of each side in a conflict varies with the current strength of the other side [9]. In their simplest form, his equations define the evolution through time of the strength of the two sides, $R(t)$ and $B(t)$, as a function of the effective firing rates α_R and α_B of the two sides,

$$\frac{dR}{dt} = -\alpha_B B(t)$$

$$\frac{dB}{dt} = -\alpha_R R(t)$$

His system is formally equivalent to the Lotka-Voltera equations for modeling predator-prey populations. When computers were first introduced, their most immediate application to military modeling was integrating the Lanchester equations, and many of the military's leading models today are still based on refinements of this model, for example, the Bonder-Farrell Attrition Algorithm equations [1].

Game Theory.—Game theory was originally developed in context of economic analysis [19, 20], but after WWII, it became a central tool for military planning at the DoD-sponsored RAND Institute and elsewhere. Game theory focuses on the rationality of the parties in conflict, and assumes that each seeks to maximize its own utility while recognizing that the other party is seeking to do the same.

Game theory and Lanchester theory differ in two important ways.

1. Lanchester theory models the combattants as physical forces with no reference to rationality. Game theory assumes that players are rational beings who seek to maximize a utility function.
2. Lanchester theory describes the evolution of combat through time. Game theory in its simplest form is concerned with the final outcome.

In spite of these differences, the two mathematical theories are similar in two ways.

1. They treat the opposing sides as aggregates, and do not consider the detailed interactions of the individual soldiers and weapons of which they are composed.
2. They lump the effect of a soldier's emotions on the outcome of combat in with other factors, such as firepower or positional advantage, and thus do not permit them to be studied in their own right.

3.1.2 Behavioral Models

On the behavioral side, armies from time immemorial have conducted wargames, either with real troops or on sand tables on which experimenters alternatively move playing pieces to explore possible tactics. With the advent of inexpensive computer cycles and the development of multi-agent systems techniques, there is increased interest in constructing models of combat in which each entity is represented by an individual computer agent.

Agent-based models of combat are superior to traditional mathematical models because they can capture the individual evolution of entities as they interact, rather than modeling them as averages over the population. Interactions in combat are strongly nonlinear, and population averages often miss important divergences in individual trajectories. As a result, entity-based models can often yield much more realistic results than Lanchester or game-theoretic models.

A disadvantage of agent-based models is that they can be much more expensive to run than classical mathematical models. One answer to this problem is the discovery that relatively simple entity models, embedded in an environment based on cellular automata, are often sufficient to capture much of the complexity of warfare [6]. One explanation for this fortunate outcome is the phenomenon of universality [15], which recognizes that the structure of a system's interactions may overwhelm differences in the processing carried out by individual agents.

Once one is characterizing individual combatants, it becomes natural to capture characteristics that are naturally thought of as emotional in nature.

Thus, for instance, Ilachinski's EINSTein model [6] represents an agent's personality as a set of six weights, each in [-1, 1], describing the agent's response to six kinds of information. Four of these depend on its sensor range, and describe the number of alive friendly, alive enemy, injured friendly, and injured enemy troops within the agent's sensor range. The other two weights relate to the model's use of a childhood game, "capture the flag," as a prototype of combat. In this outdoor game, each team has a flag, and seeks to protect it from being captured by the other team while simultaneously capturing the other team's flag. The fifth and sixth weights in the EINSTein personality vector describe how far the agent is from its own and its adversary's flag. A positive weight indicates that the agent is attracted to the entity described by the weight, while a negative weight indicates that it is repelled.

The MANA model [10] extends the concepts in EINSTein. The notion of friendly and enemy flags is replaced by the more generic notion of the next waypoint being pursued by each side. MANA includes four additional components in the personality vector: low, medium, and high threat enemies. In addition, MANA defines a set of triggers (e.g., reaching a waypoint, being shot at, making contact with the enemy, being injured) that shift the agent from one personality vector to another. A default state defines the personality vector when no trigger state is active.

The notion of being attracted or repelled by friendly or adversarial forces in various states of health is an important component of what we informally think of as emotion (e.g., fear, compassion, aggression), and the use of the term "personality" in both EINSTein and MANA suggests that the system designers are thinking anthropomorphically, though they do not use the term "emotion" to describe the effect they are trying to achieve.

3.2 Emotional Modeling

The study of emotion has a rich history in the psychological and physiological literature, reaching back well over a century [4]. This history has produced a wide range of theories, identifying emotions with outward expressions, physiological responses, distinct behaviors, or cognitive processes, among others.

Agent-based software has been growing in importance in two areas where realistic simulation of human behavior is important: agent-based modeling, and human interfaces (including gaming). This growth has led to a flurry of interest in computational models of emotion, surveyed accessibly in [17], each drawing on different segments of the psychological tradition.

Emotion clearly has facets related to an organism's outward expressions and physiological reactions, and these are important for applications of emotions in human interfaces and robotics

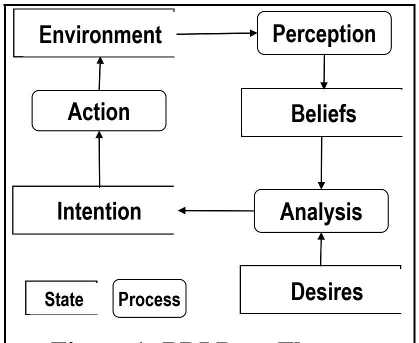

Figure 1: BDI Data Flow

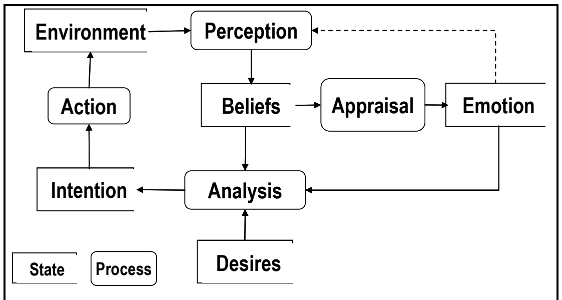

Figure 2: BDI enhanced with OCC

(e.g., [11, 12]). For our purposes, a cognitive perspective on emotions is more appropriate, and we draw on the OCC model, named after the authors of the book that sets it forth [13]. The fundamental insight of this model is that emotions are “valenced reactions to events, agents or objects, with their particular nature being determined by the way in which the eliciting situation is construed.” That is, the strength of a given emotion depends on the existence of events, agents, or objects in the environment of the agent exhibiting the emotion. This existence is mapped to a “valence,” a positive or negative score, by a process that is sometimes called “assessment” or “appraisal.”

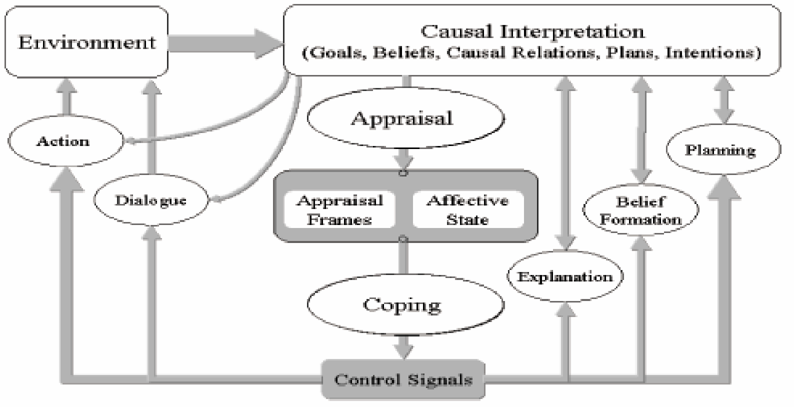

Figure 3: Gratch-Marsella Model of Cognitive-Motivational-Emotive System ([5], Fig. 2)

Once an emotion exists, it impacts the subject in several ways. It focuses attention, increases the prominence of an event in memory, affects cognitive style and performance, and influences judgments [2]. In particular, according to OCC, “behavior is a response to an emotional state in conjunction with a particular initiating event.” In our application, we focus on the impact of emotion on an agent’s analysis and judgment, the process by which it selects its intentions from its desires.

To put this system in a broader perspective, consider the basic Belief-Desire-Intention [7, 18] data flow summarized in Figure 1. Beliefs (derived from the environment by perception) and Desires feed an analysis process that produces Intentions, which in turn drive actions that change the environment. Figure 2 shows a simple enhancement of this model with the OCC model of emotion. Beliefs feed not only analysis, but also the appraisal process that generates emotions. These emotions in turn influence analysis and perception (the latter link shown dashed because we do not emphasize it in our current system).

Table 1: Comparison of Models

BDI + OCC (Figure 2)	Gratch-Marsella (Figure 3)
Environment	Environment
Perception	Causal Interpretation
Beliefs	Causal Interpretation
Appraisal	Appraisal
Emotion	Affective State
Analysis	Coping
Desires	???
Intention	Control Signals
Action	Action

Gratch and Marsella [5] offer one of the more mature current computational models of agent emotion. Figure 3 sketches their model, and Table 1 summarizes the correspondence between salient elements of the two models.

We can interpret the decision systems in EINSTein and MANA as subsets of Figure 2.

Figure 4 casts EINSTein in this framework. EINSTein's personality vector guides the agent's decisions, but is itself fixed, and does not change in response to the agent's beliefs about the events, objects, or agents in its environment. Thus it is not a "valenced reaction," and is best considered a representation of the agent's desires. In this sense, EINSTein does not capture combatant emotion.

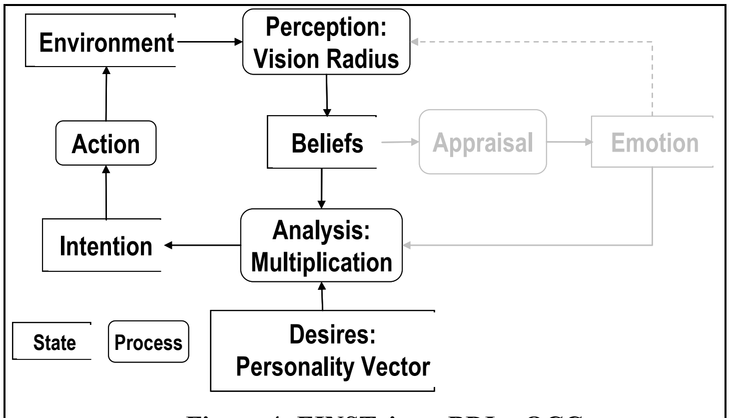

Figure 4: EINSTein as BDI + OCC

Figure 5 represents MANA in this framework. Because MANA's personality vectors are switched on the basis of a trigger state, they qualify as a valenced reaction. The default personality vector that applies when no trigger state is active continues to function as a representation of the agent's desires.

In both EINSTein and MANA, the analysis process consists of multiplying the environmental information available to the agent by the personality vector, directly yielding a movement vector to guide the agent's subsequent actions. In both models, the perception process is represented by a vision radius within which the agent has perfect knowledge of its environment. These processes are considerably simpler than the mechanisms of symbolic AI applied in [5] for the same functions. The differences reflect the differing objectives of the systems. Gratch and Marsella are supporting a training environment with relatively few agents, and regular interaction with humans slows the pace to the point that significant computation can occur. EINSTein and MANA manipulate dozens or even hundreds of agents in non-interactive simulations of combat, and need to minimize the overall execution time to permit the execution of many instances of a scenario.

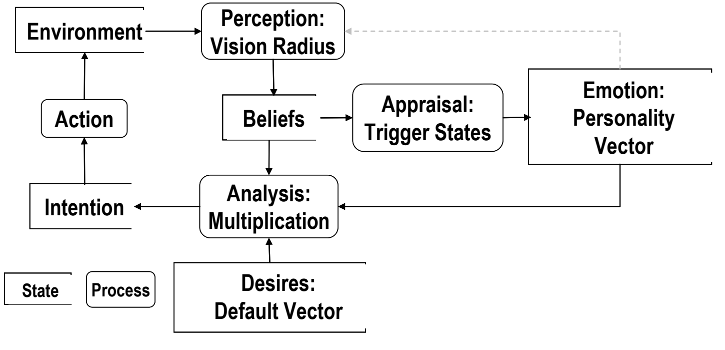

Figure 5: MANA as BDI + OCC

4 The DETT Emotion Model

Our work is being conducted in the context of two DoD projects that require the ability to simulate large numbers of combatants very rapidly. As a result, our mechanisms favor numerical computation.

4.1 Application Contexts

The DETT model was developed in the context of the DARPA RAID program, and is also being used to model noncombatants in another project.

The objective of the DARPA RAID program [8] is to anticipate enemy actions and deceptions, in order to provide real-time support to a tactical commander. We are constructing a module that reasons about the adversary's likely state and actions (thus, an Adversarial Reasoning Module, or ARM). Our particular ARM synergizes three distinct tactical reasoners: statistical reasoning for early detection of anomalous situations that might indicate risk, knowledge-based inference to reason about possible states of the world, and behavioral evolution and extrapolation, using swarms of fine-grained agents to explore possible futures of the battlespace. In this third reasoner, we evolve agents against observed reality to learn their characteristics and determine which ones are most likely to reflect future behavior. Because many of these agents must execute faster than real time, they cannot conduct complex symbolic reasoning, but use numerical computation. These agents use the DETT model described in this paper.

In the other project, MAROP (Multi-Agent Representation of the Operational Environment), we are developing methods to enrich a new military modeling system (Combat XXI) by automating the reactions of non-combatants with combatants. This capability requires us to recognize the possibility of different personality types among non-combatants and to incorporate these differences in their behavior.

4.2 Architecture

We need a computationally efficient way to take emotional tendencies into account in modeling combat. This reasoning takes place at two locations in Figure 2: Appraisal and Analysis. We have defined numerical methods for both of these.

4.2.1 Appraisal

MANA's use of triggered personality vectors that specify numerical weights for translating beliefs into intentions is a useful model for appraisal, but it is limited in two ways. First, MANA defines vectors and triggers at the level of the squad, and all members of the squad share the same values. In practice, individual combatants will differ widely in their susceptibility to different emotions. A firefight that might stimulate a high level of fear in a new soldier may have much less effect on a seasoned veteran. In order to use evolution to learn the characteristics of entities, we must parameterize this kind of difference. Second, MANA assumes that an agent in the presence of a trigger immediately adopts the associated emotion, and that when the trigger is removed, the emotion ceases immediately. Empirically, the rise of an emotion, while rapid, is not instantaneous, and the emotion is likely to persist after the trigger is removed.

We address the first concern adding a new component, which we call Dispositions, to the model (Figure 6). There is a one-to-one mapping between Emotions and Dispositions. Like Desires, Dispositions are persistent (that is, their values are constant over the time horizon of our simulations). A Disposition modulates the Appraisal process to determine the extent to which a given belief triggers the corresponding emotion. The emotion then modulates the analysis process to impose a Tendency on the resulting intention. The main elements of this model are thus the Disposition, Emotion, Trigger (the beliefs that lead to the emotion), and Tendency (the effect on intentions) (DETT). Table 2 illustrates two Dispositions that we are using, with their associated Emotions and illustrative Triggers and Tendencies.

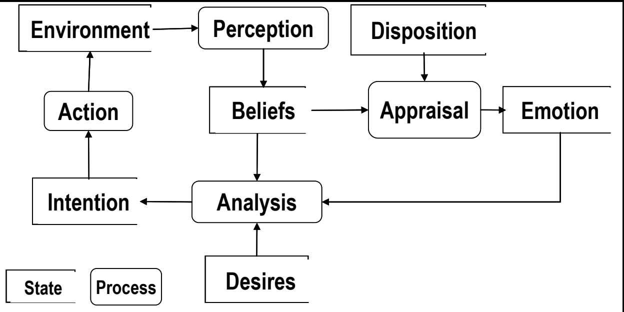

Figure 6: Incorporating Disposition in BDI + OCC

Our agents live in a digital pheromone infrastructure [3] in agents sense one another's presence through labeled scalars that they deposit in the environment and that diffuse spatially and evaporate over time. The dynamics of these pheromones models (very crudely) the Perception process that maps environmental reality into agent beliefs: an agent believes

Table 2: Sample DETT Semantics

Disposition	Emotion	Trigger	Tendency
Cowardice	Fear	Presence of armed enemy Incoming attack	Less attention to commander's orders Tendency to move away from threat
Irritability	Anger	Presence of enemy troops	Increased likelihood to engage in combat Tendency to move toward threat

what it senses in the form of pheromones in its environment. Table 3 summarizes our pheromone vocabulary in the case of RAID.

Let ***P*** be the vector of pheromone strengths at an agent's location. The agent's Disposition is a matrix ***D*** whose elements are in [0,1]. ***D*[*i*,*j*]** represents the relevance of *i*th pheromone flavor to the *j*th emotion. Then the agent's *j*th emotion depends (nonlinearly) on the *j*th element of the product $P^T D$.

To allow emotions to vary more realistically in time, agents have internal pheromones (compare [16]), one for each Emotion. The product $P^T D$ at a given time step determines the deposit to the vector ***E*** of emotion pheromones at that time step, so the longer an agent is exposed to a trigger pheromone, the higher the level of the associated emotion grows. Conversely, when the relevant trigger is removed, the corresponding emotion decays exponentially. Also, the higher the value for the disposition, the more quickly the associated emotion grows in the presence of a trigger. An agent with high irritability will grow angrier more quickly in the presence of a triggering pheromone than an agent with low irritability.

4.2.2 Analysis

Analysis draws on the same pheromone vector ***P*** of beliefs as does Appraisal, and takes as input the current state of the emotion vector ***E***. In addition, it considers the values of the agent's vector of Desires or Wants ***W***. The desires we are modeling are Protect Red, Protect Blue, Protect Green, Protect Key Sites, Avoid Combat, Avoid Detection, and Survive. Each has a value in the range [-1,1], where a negative value indicates that the agent wants the opposite state of affairs described by the desire. A movement matrix ***M*** indicates whether a given Desire tends to attract or repel the agent toward a given flavor of pheromone: ***M*[i,j]** is 1 if desire *j* is attracted to pheromone *i*, -1 if it is repelled, and 0 if the pheromone is irrelevant to the desire.

In the absence of emotions, the agent's behavior is a function (again nonlinear) of $P^T M W$. Emotions modulate these behaviors. An elevated level of Anger will increase movement likelihood, weapon firing likelihood, and tendency toward an exposed posture, while elevated Fear will decrease these likelihoods. Note that level of a particular emotion actually models the extent to which the emotion modulates the agent's behavior. Someone who experiences a high level of fear, but is able to continue to behave as if he were not experiencing that level of fear would be modeled as having a low level of fear in our approach. We are not trying to

Table 3: Pheromone Flavors in RAID

RedAlive	Emitted by a living or dead entity of the appropriate group (Red = enemy, Blue = friendly, Green = neutral)
RedCasualty	
BlueAlive	
BlueCasualty	
GreenAlive	
GreenCasualty	
WeaponsFire	Emitted by a firing weapon
KeySite	Emitted by a site of particular importance to Red
Cover	Emitted by locations that afford cover from fire
Mobility	Emitted by roads and other structures that enhance agent mobility

model the level of emotion being experienced by an agent, only the apparent level of emotion that can be perceived by its impact on the agent's behavior.

5 Experimental Results

We report here the initial experiments that we are conducting in the context of the two projects that are using the DETT model. At this point, our experiments are focused on demonstrating that the software mechanisms linking dispositions, pheromone-mediated beliefs, emotions, desires, and resulting intentions are working correctly.

5.1 MAROP

Our initial experiments measure the effect of the emotions of fear and anger on the spatial correlation of non-combatants with other entities of interest. All experiments are conducted on a road network modeling an urban area. In our experimental schema, Red and Blue forces follow scripted movements that carry them from opposite sides of the town to a central location where they engage in a firefight. Initially, Green agents are distributed randomly throughout the town.

MAROP uses a simplified version of DETT in which different dispositions and emotions are precompiled into an agent's attraction to or repulsion from each of four different pheromone flavors: RedAlive, BlueAlive, GreenAlive, and WeaponsFire. We define two emotions, Fear and Anger, as summarized in Table 4. (These encodings assume that the combat takes place on Red's "turf."). We can measure the resulting effect on the relative distribution of Green and each of the other classes (Red agents, Blue agents, and conflict events) by computing the spatial correlation of the associated pheromone fields. A correlation of 1 indicates that the two classes of agents tend to be in the same regions of the town, while a correlation of -1 indicates that they tend to avoid one another.

Table 4: Emotional Tendencies in MAROP

Pheromone	Fear	Anger
RedAlive	Repulsive	Neutral
BlueAlive	Repulsive	Attractive
GreenAlive	Attractive	Neutral
WeaponsFire	Repulsive	Attractive

Distinct scenarios are run with Green agents coded as fearful, angry, and unemotional. Thus each experimental scenario forms a point in a three-dimensional space. Figure 7 shows a projection of this space on the plane defined by Green-Red and Green-Blue correlations, illustrating how Green agents of different emotional configuration assume different relations to the other agents in the scenario.

5.2 RAID

RAID uses the avatar-ghost system we developed in earlier work [14]. Each real-world entity has exactly one agent representative, its avatar, but the avatar explores alternative possible futures by constantly sending out a swarm of ghost agents whose pheromone-based self-organization then guides the avatars.

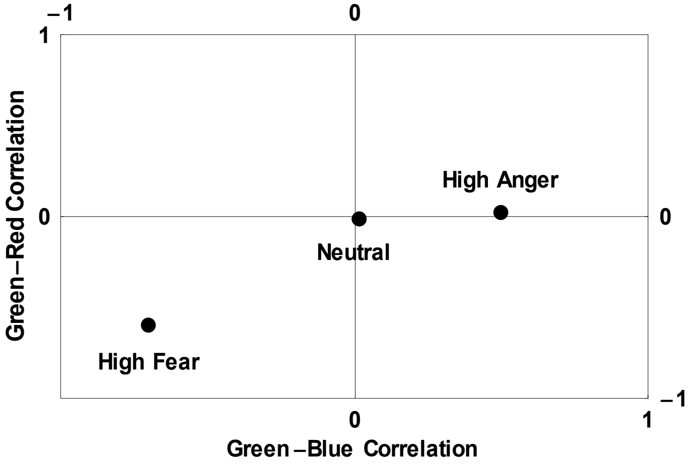

Figure 7: Effect of Emotions in MAROP (notional data)

To test the effect of emotions in RAID, we arrange ten units (one avatar per unit) of each color in files, and have the Red and Green march through the Blue in formation. When a file reaches one extreme of the arena, it reverses its direction. The units reach their original locations after 189 time steps, and the scenario repeats. Thus the units repeatedly pass through one another, depositing pheromones that indicate their presence and sensing the pheromones deposited by the other agents. Each unit emits eight ghosts per time step, and each ghost explores the future for five time steps before dying.

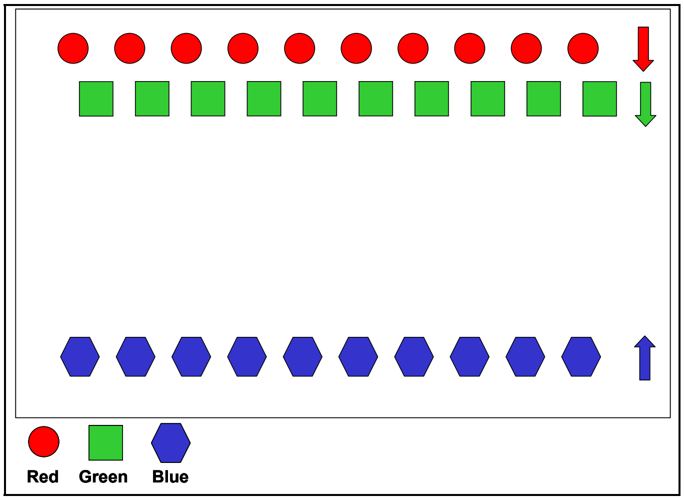

Figure 8: Experimental Configuration for RAID

The ghosts for each avatar are generated with range of values in [0,1] for the “cowardice” disposition. In our full system, an evolutionary process narrows these down on the basis of comparisons between the ghosts and actual history, but this process is not operating in this experiment. Thus, averaged over time, the ghosts for each avatar have a cowardice of 0.5.

According to Figure 6, the combination of a disposition with beliefs about the environment yields emotions. Figure 9 (top)shows the average value of the “fear” emotion across the ghosts for Red unit #1, as a function of time. This value peaks each time the unit crosses the line of Blue units, reflecting an interaction between the ghosts’ cowardice disposition and the BlueAlive pheromone that they sense in the environment.

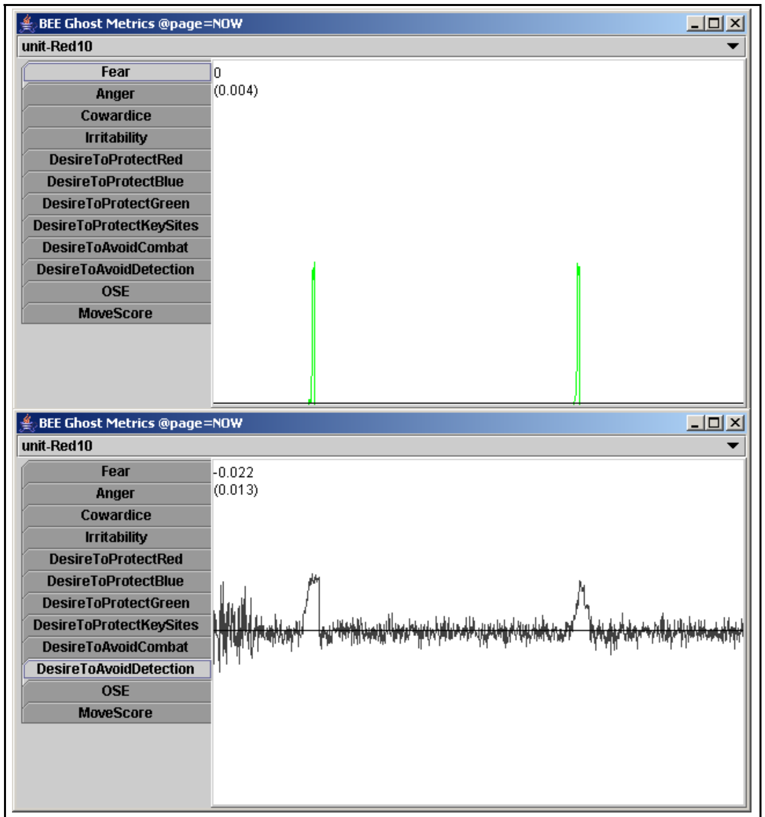

Figure 9: Fear emotion (top) and intent to avoid detection (bottom) resulting from Cowardice disposition in the presence of adversarial pheromone. X-axis is time.

Again following Figure 6, emotion affects the agent’s analysis to determine its intentions. Figure 9 (bottom) shows the average level of the “avoid detection” intention across this unit’s ghosts. (The use of the term “desire” in the legend is a typo that will be corrected in revisions to the software.) As required by the architecture outlined in Section 4, this intention increases when the fear emotion is active.

The net effect of the architecture is thus to modulate the agent’s intrinsic desire to avoid detection on the basis of its emotional state (specifically, fear), as determined by its disposition (cowardice) and its beliefs about the environment (the presence of adversaries).

6 Conclusion

An important next step in the development of combat models is the implementation of realistic models of combatant emotion. The realities of human combat make this refinement necessary, while the maturation of agent-based models of combat makes it feasible. The Gratch-Marsella model offers a sophisticated implementation of current psychological theories of emotion, but is computationally too expensive to apply to large populations of combatant agents. Some fine-grained agent-based models embed a notion of personality (EINSTein and MANA), but do not recognize the important distinctions between individual combatants.

The DETT model (Dispositions, Emotions, Triggers, Tendencies) combines the theoretical richness of the Gratch-Marsella model with the computational efficiency of EINSTein and MANA. We are using the model in two different contexts, and have demonstrated the basic computational cycle in implemented software.

Our ongoing research includes embedding this mechanism in an evolutionary loop whose fitness is determined by comparison of simulated behavior with real-world status to estimate the actual emotional state of observed combatants, and developing methods for verifying and validating the DETT architecture against actual human behavior.

7 Acknowledgements

This material is based in part upon work supported by the Defense Advanced Research Projects Agency (DARPA) under Contract No. NBCHC040153. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the DARPA or the Department of Interior-National Business Center (DOI-NBC). Distribution Statement "A" (Approved for Public Release, Distribution Unlimited).

This study was supported in part by the TRADOC Analysis Center, Naval Post Graduate School, Monterey under Contract No. GS-35F-4912H, Order No. GST0904BH6603. The views and conclusions in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the TRADOC Analysis Center the Naval Post Graduate School, the Department of Defense, or the US Government.

8 References

[1] S. Bonder and L. W. Farrell. Development of Models for Defense Systems Planning. SRL 2147 TR70-2, Systems Research Laboratory, University of Michigan, Ann Arbor, MI, 1970.

[2] S. Brave and C. Nass. Emotion in Human-Computer Interaction. In J. A. Jacko and A. Sears, Editors, *The human-computer interaction handbook: fundamentals, evolving technologies and emerging applications*, pages 81-96. Lawrence Erlbaum Associates, Inc., Mahwah, NJ, 2003.

[3] S. Brueckner. *Return from the Ant: Synthetic Ecosystems for Manufacturing Control*. Dr.rer.nat. Thesis at Humboldt University Berlin, Department of Computer Science, 2000. Available at http://dochost.rz.hu-berlin.de/dissertationen/brueckner-sven-2000-06-21/PDF/Brueckner.pdf.

[4] C. Darwin. *The expression of the emotions in man and animals*. London, John Murray, 1872.

[5] J. Gratch and S. Marsella. A Domain-independent Framework for Modeling Emotion. *Journal of Cognitive Systems Research*, 5(4):269-306, 2004.

[6] A. Ilachinski. *Artificial War: Multiagent-based Simulation of Combat*. Singapore, World Scientific, 2004.

[7] D. Kinny, M. Georgeff, and A. Rao. A Methodology and Modelling Technique for Systems of BDI Agents. In W. VandeVelde and J. W. Perram, Editors, *Agents Breaking Away. 7th European Workshop on Modelling Autonomous Agents in a Multi-Agent World (MAAMAW'96). Lecture Notes in Artificial Intelligence 1038*, pages 56-71. Springer, Berlin, 1996. Available at ftp://www.aaii.com.au/pub/aaii-technotes/technote58.ps.gz.

[8] A. Kott. Real-Time Adversarial Intelligence & Decision Making (RAID). 2004. http://dtsn.darpa.mil/ixo/programdetail.asp?progid=57.

[9] F. W. Lanchester. *Aircraft in Warfare: The Dawn of the Fourth Arm*. London, Constable and Co, Ltd., 1916.

[10] M. K. Lauren and R. T. Stephen. Map-Aware Non-uniform Automata (MANA)—A New Zealand Approach to Scenario Modelling. *Journal of Battlefield Technology*, 5(1 (March)):27ff, 2002. Available at http://www.argospress.com/jbt/Volume5/5-1-4.htm.

[11] R. Marinier and J. Laird. Toward a Comprehensive Computational Model of Emotions and Feelings. In *Proceedings of Sixth International Conference on Cognitive Modeling*, pages 172-177, Lawrence Earlbaum, 2004.

[12] R. R. Murphy, C. L. Lisetti, R. Tardif, L. Irish, and A. Gage. Emotion-Based Control of Cooperating Heterogeneous Mobile Robots. *IEEE Transactions on Robotics and Automation*, 18(5 (October)):744-757, 2002.

[13] A. Ortony, G. L. Clore, and A. Collins. *The cognitive structure of emotions*. Cambridge, UK, Cambridge University Press, 1988.

[14] H. V. D. Parunak, S. Brueckner, and J. Sauter. Digital Pheromones for Coordination of Unmanned Vehicles. In *Proceedings of Workshop on Environments for Multi-Agent Systems (E4MAS 2004)*, pages (forthcoming), Springer, 2004. Available at http://www.altarum.net/~vparunak/E4MAS04_UAVCoordination.pdf.

[15] H. V. D. Parunak, S. Brueckner, and R. Savit. Universality in Multi-Agent Systems. In *Proceedings of Third International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2004)*, pages 930-937, IEEE, 2004. Available at http://www.altarum.net/~vparunak/AAMAS04Universality.pdf.

[16] H. V. D. Parunak, S. A. Brueckner, R. Matthews, and J. Sauter. Pheromone Learning for Self-Organizing Agents. *IEEE SMC*, May, 2005.

[17] R. W. Picard. *Affective Computing*. Cambridge, MA, MIT Press, 2000.

[18] A. S. Rao and M. P. Georgeff. Modeling Rational Agents within a BDI Architecture. In *Proceedings of International Conference on Principles of Knowledge Representation and Reasoning (KR-91)*, pages 473-484, Morgan Kaufman, 1991.

[19] J. von Neumann. Zur Theorie der Gesellschaftsspiele. *Mathematische Annalen*, 100:295-320, 1928.

[20] J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton, Princeton University Press, 1944.