

# Evolving Agents to Recognize Plans and Emotions

H. Van Dyke Parunak, Sven Brueckner, Robert Matthews, John Sauter, Steve Brophy

Altarum Institute  
3520 Green Court, Suite 300  
Ann Arbor, MI 48105

{van.parunak, sven.brueckner, robert.matthews, john.sauter, steve.brophy}@altarum.org

## Abstract

Reasoning about agents that we observe in the world is challenging. Our available information is often limited to observations of the agent's external behavior. To understand these actions, we need to deduce the agent's internal state, which includes not only rational elements (such as intentions and plans), but also emotive ones (such as fear). In addition, we often want to predict the agent's future actions, which are constrained not only by these inward characteristics, but also by the dynamics of the agent's interaction with its environment. BEE (Behavior Evolution and Extrapolation) uses a faster-than-real-time agent-based model of the environment to characterize agents' internal state by evolution against observed behavior, and then predict their future behavior, taking into account the dynamics of their interaction with the environment.

## 1. Introduction

A central challenge in many applications is reasoning from external observations of agent behavior to an estimate of their internal state motivated by a desire to predict the agent's future behavior. Our observations are limited to the agent's external behavior, which can frequently be summarized numerically as a trajectory in space-time (perhaps punctuated by actions from a limited vocabulary). However, this behavior is driven by the agent's internal state, which may involve high-level psychological and cognitive concepts such as intentions and emotions.

This problem is traditionally known as "plan recognition" or "plan inference." Most work focuses on recognizing the rational state (as opposed to the emotional state) of a single agent (as opposed to an interacting community), and frequently takes advantage of explicit communications between agents (as in managing conversational protocols). Many realistic problems, including military combat, competitive business, and multi-player computer games, deviate from these conditions.

- Increasing the number of agents leads to a combinatorial explosion that can swamp conventional analysis.
- The dynamics of the environment can frustrate the intentions of an agent.

- The agents often are trying to hide their intentions.
- An agent's emotional state may be as important as its rational state in determining its behavior.

BEE (Behavioral Evolution and Extrapolation) is a novel approach to recognizing the rational *and emotional* state of *multiple* interacting agents based solely on their behavior, *without recourse* to intentional communications from them, and incorporating *environmental dynamics*. It is inspired by techniques used to predict nonlinear dynamical systems. The current version characterizes and predicts agents representing soldiers in urban combat [7].

Section 2 reviews relevant previous work, including plan recognition and prediction in nonlinear dynamical systems. Section 3 describes the architecture of BEE. Section 4 reports experimental results. Section 5 concludes. Details that cannot be included here for the sake of space are available in an on-line technical report [15].

## 2. Previous Work

BEE is related to previous research in AI (plan recognition) and nonlinear dynamical systems (trajectory prediction).

### 2.1 Plan Recognition in AI

An agent's goals guide its actions. Thus one ought to be able to learn something about an agent's goals by observing its past actions, and knowledge of the agent's goals facilitates predictions about the agent's future actions.

Reasoning from an agent's actions to its goals is known as "plan recognition" or "plan inference." This work (surveyed recently at [2]) covers both single-agent and multi-agent (e.g., robot soccer team) plans, intentional vs. non-intentional actions, speech vs. non-speech behavior, adversarial vs. cooperative intent, complete vs. incomplete world knowledge, and correct vs. faulty plans, among other dimensions. When repeated observations of the agent are available, one can fit a statistical model such as a hidden Markov model (HMM) or a dynamic belief network (DBN) [20], avoiding the expense needed to construct the plan libraries used by other approaches.

Plan recognition usually supports a higher-level function such as predicting the entity's future. An agent's plan is a useful input to prediction, but hardly sufficient. BEE in-

corporates two other influences, one internal and one external, that are not handled in traditional plan recognition.

The *external influence* is the dynamics of the environment, which may include other agents. The environment may do things on its own that interfere with the desires of the agent [3, 9]. Most interactions among agents, and between agents and the world, are nonlinear. When iterated, these can generate rapid divergence of trajectories with similar initial conditions. Actual simulation is one way (the only one we know now) to assess these effects.

The *internal influence* is the agent’s emotional state, which can modulate its decision process and its focus of attention (and thus its perception of the environment). In extreme cases, emotion can lead an agent to choose actions that appear irrational.

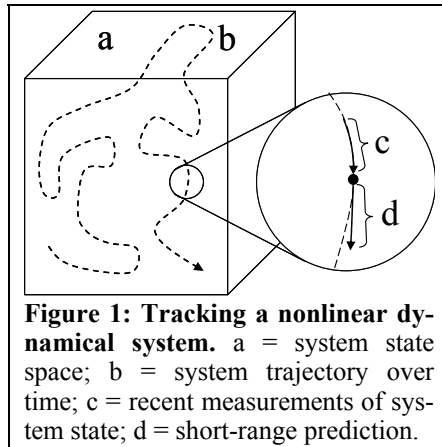
Current work on plan recognition for prediction focuses on the rational plan, and does not take into account either external environmental influences or internal emotional biases. BEE integrates all three elements.

BEE is statistical, but distinct from classical statistical plan recognizers. They reason from many observed trajectories to a single internal representation. BEE requires only one observed trajectory, and compares it with multiple candidate representations. When multiple trajectories are available, a HMM or DBN recognizer can advise BEE as described in Section 3.3.

## 2.2 Nonlinear Systems Analysis

Many systems can be described by a vector of real numbers that changes with time. The dimensions of the vector define the system’s state space. One typically analyzes such systems as vector differential equations, e.g.,  $\dot{\vec{x}} = f(\vec{x})$ . When  $f$  is nonlinear, the system can be formally chaotic, and starting points arbitrarily close to one another can lead to trajectories that diverge exponentially rapidly, becoming uncorrelated. Long-range prediction of such a system is impossible. However, one can anticipate the system’s behavior a short distance into the future, by fitting  $f$  to the system’s trajectory in the recent past, and then extrapolating this fit into the future (Figure 1, [6]). Iterating this process provides the user with a limited look-ahead into the future.

While robust and popular, this approach requires systems that can efficiently be described by mathematical equations that can be fit using optimization methods such as least squares. BEE takes its inspiration from this approach, but applies it to agent behaviors, which it fits to observed behavior using a genetic algorithm.



**Figure 1: Tracking a nonlinear dynamical system.** a = system state space; b = system trajectory over time; c = recent measurements of system state; d = short-range prediction.

## 3. Architecture

The BEE architecture include the model of an individual agent, the pheromone infrastructure through which agents interact, the information sources that guide them, and the evolutionary cycle that they execute.

### 3.1 Agent Model

BEE’s agents are inspired by two bodies of work. The first is our own experience with fine-grained agents that coordinate their actions through digital pheromones in a shared environment [1, 12, 16-18]. The second is previous agent-based combat modeling.

Digital pheromones are scalar variables that agents deposit and sense at their current location. Agents respond tropistically to local gradients of these variables. Their movements change the deposit patterns. This feedback, together with evaporation and propagation in the environment, support complex patterns of interaction and coordination among the agents [14]. For example, a living member of the adversary emits a RED-ALIVE pheromone.

Our soldier agents are inspired by EINSTEIN and MANA. EINSTEIN [5] represents an agent as a set of six weights, each in  $[-1, 1]$ , describing the agent’s response to six kinds of information. Four of these describe the number of alive friendly, alive enemy, injured friendly, and injured enemy troops within the agent’s sensor range. The other two relate to the agent’s distance to its own flag and that of the adversary, representing objectives that it seeks to protect and attack, respectively. A positive weight indicates that the agent is attracted to the corresponding entity, while a negative weight indicates that it is repelled.

MANA [8] extends the concepts in EINSTEIN. Friendly and enemy flags are replaced by waypoints pursued by each side. Enemies can pose different levels of threat. In addition, a set of triggers (e.g., reaching a waypoint, being shot at, making contact with the enemy, being injured) can shift the agent from one personality vector to another.

Personality vectors in MANA and EINSTEIN reflect both rational and emotive aspects of decision-making. Attraction to or repulsion from friendly or adversarial forces in various states of health is an important component of emotion (e.g., fear, compassion, aggression). Waypoints reflect goal-oriented rationality.

BEE embodies an integrated rational-emotive personality model.

A BEE agent’s rationality is a vector of seven desires (values in  $[-1, +1]$ ): ProtectRed (the adversary), ProtectBlue (friendly forces), ProtectGreen (civilians), ProtectKeySites, AvoidCombat, AvoidDetection, and Survive. Negative values reverse the sense suggested by the label. For example, a negative value of ProtectRed indicates a desire to harm Red. Different pheromones in the environ-

ment attract or repel an agent with a given desire. For example, an agent with a high positive desire to ProtectRed will be attracted to RED-ALIVE, RED-CASUALTY, and MOBILITY pheromone, and will move at maximum speed.

The emotive component of a BEE personality, described elsewhere [11], is based on the Ortony-Clore-Collins (OCC) framework [10]. OCC define emotions as “valenced reactions to agents, states, or events in the environment.” This notion is captured in MANA’s trigger states. BEE enhances this model with varying sensitivities of agents to triggers. For example, threatening situations tend to stimulate the emotion of fear, but a given level of threat will produce more fear in a new recruit than in a seasoned combat veteran. Thus our model includes Dispositions as well as Emotions. Each Emotion has a corresponding Disposition. Dispositions are relatively stable, and considered constant over the time horizon of a run of the BEE, while Emotions vary based on the agent’s Disposition and the stimuli to which it is exposed.

Interviews with military domain experts identified the two most crucial emotions for combat as Anger (with the corresponding disposition Irritability) and Fear (whose disposition is Cowardice). Each emotion has a set of triggering pheromones (Table 3 in [15]). For example, RED-CASUALTY pheromone stimulates both Anger and Fear in a Red agent, but not in a Blue agent.

Emotions grow in the presence of the triggering environmental condition and evaporate over time. A non-zero emotion modifies the agent’s actions. Elevated Anger increases movement likelihood, weapon firing likelihood, and tendency toward an exposed posture. Elevated Fear decreases these likelihoods.

Figure 2 summarizes the BEE’s personality model. The left side is a straight-forward Belief-Desire-Intention/Goal model [4]. The right side is the emotive component.

### 3.2 The BEE Cycle

BEE’s major innovation extends the nonlinear systems technique of Section 2.2 to agent behaviors.

Figure 3 is an overview of Behavior Evolution and Extrapolation. Each active entity has a persistent avatar that continuously generates a stream of ghost agents representing itself.

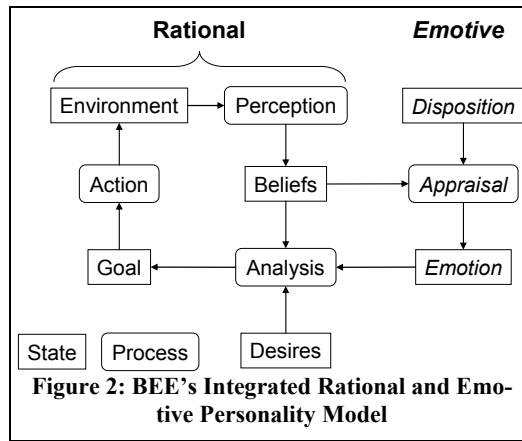


Figure 2: BEE’s Integrated Rational and Emotive Personality Model

ghost’s behavioral parameters (desires and dispositions) are sampled from distributions to explore alternative personalities of the entity it represents.

Each page between the insertion horizon and  $\tau = t$  (“now”) records the pheromone state of the world at the point in the past to which it corresponds. As ghosts move from page to page, they interact with this state, based on their behavioral parameters. Thus their fitness depends not just on their own actions, but also on the behaviors of the rest of the population, which is also evolving. Because  $\tau$  advances faster than real time, eventually  $\tau = t$  (actual time). At this point, we evaluate each ghost based on how close it is to its corresponding real-world entity.

The fittest ghosts have three functions.

1. The personality of each entity’s fittest ghost is reported to the rest of the system as the likely personality of that entity. This information enables us to characterize individual warriors as unusually cowardly or brave.
2. The fittest ghosts breed and their offspring return to the insertion horizon to continue the fitting process.
3. The fittest ghosts for each entity form the basis for a population of ghosts that run past the avatar’s present into the future. Each such ghost explores a different

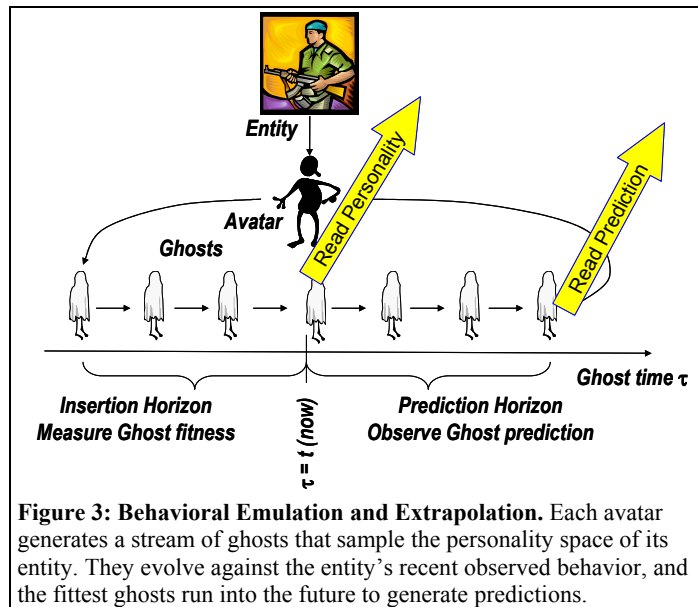


Figure 3: Behavioral Emulation and Extrapolation. Each avatar generates a stream of ghosts that sample the personality space of its entity. They evolve against the entity’s recent observed behavior, and the fittest ghosts run into the future to generate predictions.

The modeling entity consisting of avatar and its ghosts is a *polyagent* [13].

Ghosts live on a timeline indexed by  $\tau$  that begins in the past and runs into the future relative to the current time  $t$ . The timeline is divided into discrete “pages,” each for a successive value of  $\tau$ . The avatar inserts the ghosts at the insertion horizon. We currently insert at  $\tau - t = -30$ , meaning that ghosts are begin on a page representing the state of the world 30 minutes ago. At the insertion horizon, each

possible future, analogous to how people plan by mentally simulating different ways that a situation might unfold. Analysis of these possible futures yields predictions and recommended actions.

BEE must keep pace with the battle. Thus we use simple agents coordinated using pheromone mechanisms [1]. This infrastructure runs on the nodes of a rectangular lattice. Each node maintains a scalar value for each flavor of pheromone, and provides three functions: *aggregation* of deposits from

individual agents to fuse information across multiple agents and through time, *evaporation* over time to provide truth maintenance by discarding obsolete information, and *diffusion* to nearby places to share information with other agents.

Each page of the timeline is a complete pheromone field for the world at the BEE time  $\tau$  represented by that page. Ghost movement based on pheromone gradients is a simple process, so this system can support realistic agent populations without excessive computer load. In our current system, each avatar generates eight ghosts per shift. Since there are about 50 entities in the battlespace (about 20 units each of Red and Blue and about 5 of Green), we must support about 400 ghosts per page, or about 24000 over the entire book. Two features of the environment permit it to be scaled using multiple processors. First, agents (both ghosts and avatars) interact only with their local neighborhood. Second pheromones below a certain threshold are deleted, so they propagate only to a certain radius.

### 3.3 Information Sources

The BEE's pheromone infrastructure permits the integration of numerous information sources. Our current system draws on three, and others are easily added.

**Real-world observations.**—Observations from the real world are encoded into the pheromone field each increment of BEE time, as a new “current page” is generated.

**Statistical estimates of threat regions.**—Statistical techniques<sup>1</sup> estimate the level of threat to each force (Red or Blue), based on the topology of the battlefield and the known disposition of forces. For example, a broad open area with no cover is threatening, especially if the opposite force occupies its margins. The results of this process are posted to the pheromone pages as “RedThreat” pheromone (representing a threat to red) and “BlueThreat” pheromone (representing a threat to Blue).

**AI-based plan recognition.**—While plan recognition is not sufficient for effective prediction, it is a valuable input. We dynamically configure a Bayes net based on heuristics to identify the likely goals that each entity may hold.<sup>2</sup> The destinations of these goals function as “virtual pheromones.” Ghosts include their distance to such points in their action decisions, achieving the result of gradient following without the computational expense of maintaining a pheromone field. DBN techniques could be used to tune this network across multiple battles.

The additional contribution of the BEE beyond the reasoning provided by these components is three-fold. Its digital pheromone field provides a common representation for integrating their inputs. Its simulation brings to bear the dynamics of the environment in a computationally efficient way. Its evolutionary process selects among possible alter-

natives, both in characterizing an agent's internal state and in predicting its future behavior.

## 4. Experimental Results

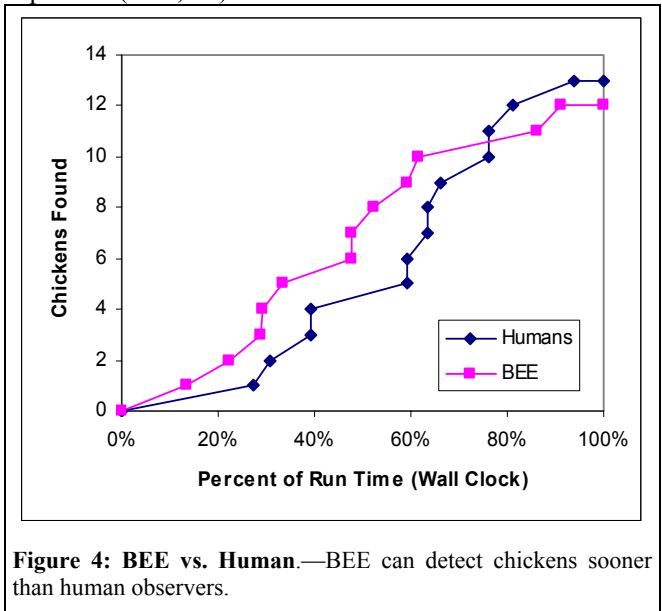
We tested BEE in a series of experiments in which human wargamers make decisions that are played out in a battlefield simulator (a different simulator than the one against which BEE evolves its agents). The commander for each side (Red and Blue) has a team of human operators who set waypoints for individual units in the simulator. Each operator is responsible for four to six units. The simulator moves the units, determines firing actions, and resolves the outcome of conflicts.

### 4.1 Fitting Dispositions

To test our ability to recognize emotions from behavior, one Red operator secretly selects two of his units to be cowardly (“chickens”) and two to be irritable (“Rambos”). He moves each unit according to the commander's orders until the unit encounters circumstances that would trigger the emotion associated with the unit's disposition. Then he manipulates chickens as though they are fearful (avoiding combat, avoiding Blue), and moves Rambos into combat as quickly as possible.

The difference between the two disposition values (Irritability - Cowardice) of the fittest ghosts proves an excellent indicator of their entity's emotional state. We maintain a 800-second exponentially weighted moving average of this Delta Disposition, and declare the unit to be a chicken or Rambo if this value passes a negative or positive threshold, respectively (currently,  $\pm 0.25$ ).

Figure 4 compares our emotion detector with the judgments of human officers. Two chickens were played in each of 11 experiments. The plot shows how many had been detected at a given point in the runs. For example, the square at (60%, 10) means that the total number of chick-



**Figure 4: BEE vs. Human.**—BEE can detect chickens sooner than human observers.

<sup>1</sup> This process was developed by R. Alonso, H. Li, and J. Asmuth at Sarnoff Corporation.

<sup>2</sup> This process was developed by P. Nielsen, J. Crossman, and R. Frederiksen at Soar Technology.

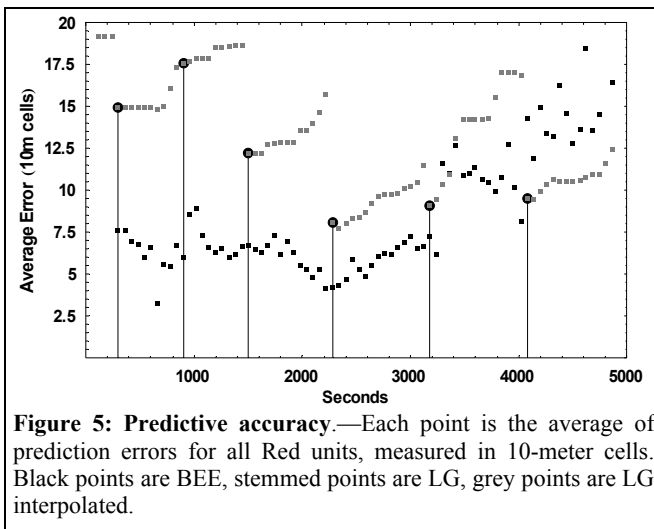
ens detected in all runs at the 60% point is 10. BEE was able to detect chickens earlier than humans, while missing only one chicken that the humans detected.

In addition to units intentionally played as emotional, BEE sometimes detects other units as cowardly or brave. Analysis shows that these characterizations are appropriate: units that flee in the face of enemy forces or weapons fire are detected as chickens, while those that stand their ground or rush the adversary are denominated as Rambos.

## 4.2 Integrated Predictions

Each ghost that runs into the future generates a possible future path that its unit might follow. The set of such paths for all ghosts embodies a number of distinct predictions, including the most or least likely future, the future that poses the greatest or least risk to the opposite side, the future that poses the greatest or least risk to one's own side, and so forth. In the experiments reported here, we select the future whose ghost receives the most guidance from pheromones in the environment at each step along the way. In this sense, it is the most likely future.

We evaluate predictions spatially, comparing an entity's actual location with the location predicted for it 15 minutes earlier. Figure 5 compares BEE's prediction errors (averaged over all 20 Red units) with those from a game-theoretic predictor based on linguistic geometry (LG) [19] on a typical run. The LG predictor can produce only one prediction (a 15 minute trajectory for each unit) about every 15 minutes (the stemmed points in the Figure), so the plotted errors for times in between these computations are based on a world state that becomes increasingly obsolete as time passes until a new prediction is issued. The BEE prediction is updated at each time step, by reading off from the any-time extrapolation process. BEE outperforms the LG predictor not only between successive LG predictions, but also when the LG prediction is current (stemmed LG points), except in the last 1000 seconds. The high BEE errors here are due entirely to a single Red unit, and if it is removed, the plot is essentially unchanged except that 11/13 of the BEE predictions in this interval become supe-



rior to the LG predictions.

## 5. Conclusions

Many domains require reasoning from an entity's observed behavior to an estimate of its internal state, and extrapolating that estimate to predict the entity's future behavior. BEE performs this task using a faster-than-real-time simulation of swarming agents, coordinated through digital pheromones. This simulation integrates knowledge of threat regions, a cognitive analysis of the agent's beliefs, desires, and intentions, a model of the agent's emotional disposition and state, and the dynamics of interactions with the environment. By evolving agents in this rich environment, we can fit their internal state to their observed behavior. In realistic wargames, the system successfully detects deliberately played emotions and makes reasonable predictions about entities' future movements.

BEE can only model internal state variables that impact the agent's external behavior. It cannot fit variables that the agent does not manifest externally, since the evolutionary cycle compares the outward behavior of the simulated agent with that of the real entity. This limitation is serious if our purpose is to understand the entity's internal state for its own sake. If our purpose of fitting agents is to predict their subsequent behavior, the limitation is much less serious. State variables that do not impact behavior, while invisible to a behavior-based analysis, are irrelevant to a behavioral prediction.

BEE lends itself to extension in several directions.

- In the experiments reported here, BEE has full knowledge of the state of the battlefield. We are currently experimenting with scenarios in which BEE has only partial knowledge.
- BEE's current inputs are only an example of the kinds of information that can be integrated. The basic principle of using a dynamical simulation to integrate a wide range of influences can support other inputs as well, requiring much less additional engineering than other more traditional ways of reasoning about how different knowledge sources come together in impacting an agent's behavior.
- Our initial emotional repertoire is a small subset of those that have been distinguished by psychologists, and that might be useful for understanding and projecting behavior. We expect to extend the set of emotions and supporting dispositions that BEE can detect.
- The mapping between an agent's cognitive-emotional state and its outward behavior is not one-to-one. Several different internal states might be consistent with a given observed behavior under one set of environmental conditions, but might yield distinct behaviors under other conditions. If the environment in the recent past confounds such distinct internal states, we will be unable to distinguish them. As long as the environment stays in this state, our predictions will be accurate, whichever of the internal states we assign to the agent. If the environment then shifts to one under which the different internal states lead to different behaviors, using the previously

chosen internal state will yield inaccurate predictions. One way to address these concerns is to probe the real world, perturbing it in ways that would stimulate distinct behaviors from entities whose psychological state is otherwise indistinguishable. Such probing is an important intelligence technique. We will explore how BEE's faster-than-real-time simulation may enable us to identify appropriate probing actions, greatly increasing the effectiveness of intelligence efforts.

- BEE has been developed for adversarial reasoning in urban warfare. It could be applied more widely, including computer games, business strategy, and sensor fusion.

## 6. 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).

## References

- [1] S. Brueckner. *Return from the Ant: Synthetic Ecosystems for Manufacturing Control*. Thesis at Humboldt University Berlin, Department of Computer Science, 2000.
- [2] S. Carberry. Techniques for Plan Recognition. *User Modeling and User-Adapted Interaction*, 11(1-2):31-48, 2001.
- [3] J. Ferber and J.-P. Müller. Influences and Reactions: a Model of Situated Multiagent Systems. In *Proceedings of Second International Conference on Multi-Agent Systems (ICMAS-96)*, pages 72-79, 1996.
- [4] A. Haddadi and K. Sundermeyer. Belief-Desire-Intention Agent Architectures. In G. M. P. O'Hare and N. R. Jennings, Editors, *Foundations of Distributed Artificial Intelligence*, pages 169-185. John Wiley, New York, NY, 1996.
- [5] A. Ilachinski. *Artificial War: Multiagent-based Simulation of Combat*. Singapore, World Scientific, 2004.
- [6] H. Kantz and T. Schreiber. *Nonlinear Time Series Analysis*. Cambridge, UK, Cambridge University Press, 1997.
- [7] A. Kott. Real-Time Adversarial Intelligence & Decision Making (RAID). 2004. <http://dtsn.darpa.mil/ixo/programdetail.asp?progid=57>.
- [8] 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.
- [9] F. Michel. *Formalisme, méthodologie et outils pour la modélisation et la simulation de systèmes multi-agents*. Thesis at Université des Sciences et Techniques du Languedoc, Department of Informatique, 2004.
- [10] A. Ortony, G. L. Clore, and A. Collins. *The cognitive structure of emotions*. Cambridge, UK, Cambridge University Press, 1988.
- [11] H. V. D. Parunak, R. Bisson, S. Brueckner, R. Matthews, and J. Sauter. Representing Dispositions and Emotions in Simulated Combat. In *Proceedings of Workshop on Defence Applications of Multi-Agent Systems (DAMAS05, at AAMAS05)*, pages 51-65, Springer, 2005.
- [12] H. V. D. Parunak and S. Brueckner. Ant-Like Missionaries and Cannibals: Synthetic Pheromones for Distributed Motion Control. In *Proceedings of Fourth International Conference on Autonomous Agents (Agents 2000)*, pages 467-474, 2000.
- [13] H. V. D. Parunak and S. Brueckner. Modeling Uncertain Domains with Polyagents. In *Proceedings of International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS'06)*, pages (submitted), 2006.
- [14] H. V. D. Parunak, S. Brueckner, M. Fleischer, and J. Odell. A Design Taxonomy of Multi-Agent Interactions. In *Proceedings of Agent-Oriented Software Engineering IV*, pages 123-137, Springer, 2003.
- [15] H. V. D. Parunak, S. Brueckner, R. Matthews, J. Sauter, and S. Brophy. Characterizing and Predicting Agents via Multi-Agent Evolution. Altarum Institute, Ann Arbor, MI, 2005. <http://www.altarum.net/~vparunak/BEE.pdf>.
- [16] 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 246-263, Springer, 2004.
- [17] H. V. D. Parunak, S. A. Brueckner, and J. Sauter. Digital Pheromone Mechanisms for Coordination of Unmanned Vehicles. In *Proceedings of First International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2002)*, pages 449-450, 2002.
- [18] J. A. Sauter, R. Matthews, H. V. D. Parunak, and S. Brueckner. Evolving Adaptive Pheromone Path Planning Mechanisms. In *Proceedings of Autonomous Agents and Multi-Agent Systems (AAMAS02)*, pages 434-440, 2002.
- [19] B. Stilman. *Linguistic Geometry: From Search to Construction*. Boston, Kluwer, 2000.
- [20] I. Zukerman and D. Albrecht. Predictive statistical models for user modeling. *User Modeling and User-Adapted Interaction*, 11:5-18, 2001.