On detecting and classifying aberrant behavior in unmanned autonomous systems under test and on mission

Rick Dove
Kennen Technologies and Stevens Institute of Technology

***Abstract*—Aberrant behavior in mobile unmanned autonomous weapons is likely. Regardless of the degree of autonomous control, aberrant behavior can be caused by design flaws, system malfunction, malevolent control penetration, and human error. In fully autonomous operation, unanticipated emergent behaviors are both likely and desirable in dealing with the infinite possibilities of situational reality. Simulation and test of individual units with these autonomous capabilities have their own sets of challenges, and cannot predict how these units will behave in group operations. Individual behavior cannot be ignored as simulation or testing advances to group behavior, and poses an explosive centralized monitoring and evaluation task with large groups. This paper explores four concepts that when combined indicate a promising ability for fast, accurate, sophisticated aberrant behavior detection and evaluation: 1) a pattern from organic life for social behavior monitoring, 2) trajectory recognition as augmentation of social monitoring, 3) massive pattern-based recognition exhibited in human domain expertise, and 4) a new VLSI processor architecture that can provide unbounded parallel pattern detection at constant speed.**

Social animal life exhibits built in systemic mechanisms for detecting aberrant behavior among its members, and mitigating that behavior if it is evaluated as intolerable. Studied examples include humans, elephants, penguins, ants and bees. This work identifies a foundation for employing that pattern in both pre-deployment device testing and in perpetual evaluations after deployment. The suggested approach was instigated by studies of Self-Organizing Systems-of-Systems in a graduate systems-engineering course at Stevens Institute of Technology, and deemed possible by capabilities of a new pattern detection-engine technology. The foundation explored in this research was shaped by this targeted technology. The paper lays a platform for subsequent investigation, and concludes with suggestions for that work, how an Unmanned Autonomous Test System might incorporate this capability, and why Test and Evaluation should be untethered from its static life-cycle position and become an integral part of complex systems.

Index Terms—anomalous behavior, multi-agent systems, pattern recognition, social behavior, testing.

Manuscript submitted November 30, 2008. This work was supported in part by the U.S. Department of Homeland Security award NBCHC070016.

Rick Dove is with Kennen Technologies and Stevens Institute of Technology, Box 289, Questa, NM 87556, USA (phone: 575-586-1536; e-mail: rick.dove@kennentech.com, rick.dove@stevens.edu).

I. Introduction

In *The Principles of Ethics (Spencer* 1893) Herbert Spencer[1] reaches into the animal kingdom to support his theories on the origins and enforcements of natural laws within social groups:

> "There arises such general consciousness of the need for maintaining the limits, that punishments are inflicted on transgressors — not only by aggrieved members of the group, but by the group as a whole. A 'rogue' elephant (always distinguished as unusually malicious) is one which has been expelled from the herd: doubtless because of conduct obnoxious to the rest — probably aggressive. It is said that from a colony of beavers an idler is banished, and thus prevented from profiting by labours in which he does not join: a statement made credible by the fact that drones, when no longer needed, are killed by worker-bees. The testimonies of observers in different countries show that a flock of crows, after prolonged noise of consultation, will summarily execute an offending member. And an eye-witness affirms that among rooks, a pair which steals the sticks from neighbouring nests has its own nest pulled to pieces by the rest."

Though stories of beaver and rook justice, and anecdotal witness to crow judgment and execution exist, scientific evidence is illusive; nevertheless the values and varieties of peer judgment constraining and enforcing societal behavioral are well known among humans and are studied and observed in some animal (Flack et. al. 2006) and insect societies (Monnin 2002, Heinze 2003, Ratnieks 2006).

This paper suggests that peer evaluation of behavior is necessary and valuable in the domain of autonomous unmanned systems (UAS) when they are working together as a team on a warfighting mission, and perhaps even more so when these systems are being tested, as they are less likely to be well behaved. The suggestion is prompted by the

[1] Wikipedia (01Nov2008): "British philosopher and sociologist, Herbert Spencer was a major figure in the intellectual life of the Victorian era. He was one of the principal proponents of evolutionary theory in the mid nineteenth century, and his reputation at the time rivaled that of Charles Darwin. Spencer was initially best known for developing and applying evolutionary theory to philosophy, psychology and the study of society."

positioning and planning for an Unmanned Autonomous Test System described in a 2008 Broad Area Announcement (Office of the Secretary of Defense 2008):

> "Due to the mobility inherent in all UAS, their close proximity to humans (e.g., soldiers, testers, population centers, etc.) and their capability for unpredictable behavior; a reliable fail-safe system is needed. This effort seeks technologies for all aspects of system safeties as they pertain to UAS, Systems of Systems, and Complex Systems. This includes safe test conduct, testing for top level mishaps, safety fail-safes, truth data assessment for safety, and safeties associated with large numbers of collaborating UAS.[2]"

It is also recognized that testing outcomes can have:

> "an almost infinite number of possibilities, depending on UAS cognitive information processing, external stimuli, operational environment, and even possible random events (hardware/software failures, false stimuli, emergent behavior, etc.).[3]"

Emergent behavior is later recognized as something less than random:

> "UAS formation control, swarming, and aggregate intelligent agent behavior are an emergent characteristic of this technology arena. ... System behavior, in a multi-agent system, can be difficult to predict and often unexpected system behaviors occur which lead to poor system performance. These unexpected system behaviors result from unforeseen group actions of agent groups and agent-group behavior that is not directly coded by the agent designers.[4]"

Such unexpected system behaviors can be good as well as bad. In fact, the goal of fully autonomous intelligent behavior is creative problem solving in situations without precedence. It is unlikely that unleashing a swarm of UAS that are only capable of dealing with well defined cataloged situations will be effective.

Isaac Asimov recognized this need when he first proposed the three laws of robotics in his original 1942 story of "Runaround" in the *I, Robot* anthology (Asimov 1942). If you are going to unleash intelligent autonomous entities to mingle with humanity you must have some rules that protect the humans. Even the awesome positronic brain Asimov conceived was incapable of consulting a policy and procedures manual prescribing required behavior for all possible human-machine interactions. He wrote many books exploring the ways in which three simple behavior-governing-principles would resolve sticky man-machine confrontations and interactions favorably.

The lesson from Asimov is that we cannot know the situations that will arise, nor can we directly control how things should play out. Instead, we must recognize and embrace uncertainty within a framework of governance

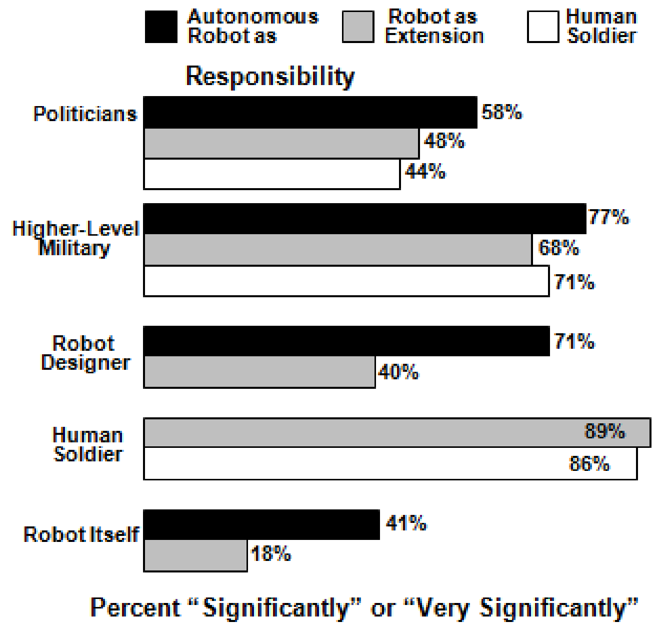

Fig 1: From the survey reported in (Moshkina and Arkin 2007) – Responsibility for Lethal Errors by Responsible Party. The soldier was found to be the most responsible party, and robots the least.

principles that will bound the outcomes within an acceptable space. The principle built on in this paper classifies behavior as unacceptable based on absolute boundary-infraction recognition, rather than attempts at imperfect reasoning or restrictions to specifically approved behaviors.

Asimov's positronic brain was invincibly obedient and not subject to psychoses, illness, error, or possession by another entity. UAS as we will know them are not likely to be so endowed. Nor will they have the real-time reasoning powers of the positronic brain. Which brings us to the fail-safe objective that unmanned autonomous system testing (UAST) seeks, and the challenge of "Confinement of UAS to safe 'play areas'[5]".

UAS will necessarily be tested and fielded in situations that have no precedence in cataloged responses. How will we constrain the outcomes to those we can live with? More to the point of this paper, how will we detect and classify unacceptable behavior in time to intervene if unacceptable consequences are the likely outcome?

From studies of agile systems maturing since 1991 (Dove 2001) I characterize the UAST domain as a class 1 (reconfigurable) agile system that must test class 2 (reconfiguring) agile systems.

This agile-system class distinction is relatively recent, becoming evident when formulating a collaborative research project as a graduate course in the School of Systems and Enterprises at Stevens Institute of Technology (Dove 2007). The course reviews the research literature in various bodies of knowledge relevant to self organizing systems of all kinds, and challenges collaborative analysis to identify recurring and necessary patterns across bodies of knowledge. Four cycles through this investigation to date is beginning to yield some

[2] From (Office of the Secretary of Defense 2008) p 48.
[3] ibid p. 23.
[4] ibid p. 54.
[5] ibid p. 23.

promising fundamental patterns. One in particular is the genesis of this paper's focus: successful social systems often exhibit a pattern of peer behavior-enforcement arising when the stability of the social system is put at risk. In certain animal societies, this enforcement ranges from one extreme of death to the other extreme of disciplining youth in proper behavior.

Though I've been unable to find proper research literature on the classic example widely seen and understood in the 2006 academy award winning documentary March of the Penguins, the lesson is Indelible. I refer to the scene where one female lost her egg to the Arctic cold and tried to replace it by stealing another's egg – other females in the immediate vicinity banded together and prevented this from happening.

A related body of work led by Ronald Arkin at Georgia Institute of Technology is concerned with ethical behavior of UAS used in military operations (Arkin 2007). The recent survey (Moshkina and Arkin 2007) investigated opinions about the use of, and responsibilities for, lethal autonomous systems among four demographic groups. A total of 430 respondents were distributed demographically as 54% robotics researchers, 30% military, 16% policymakers, and 27% general public.

Figure 1 depicts who the respondents felt responsible when behavior is unacceptable. Our interest here is in the autonomous devices, not the "robot as extension" case, where a human directs the unmanned system. Interesting to note: lethal mistakes made by a UAS are blamed on higher-level military, UAS designers, and politicians, in that order.

The survey showed that all four demographic groups want ethical standards for UAS to be considerably higher than that for soldiers. Figure 2 shows how they felt about specific constraints that should be enforced.

Monitoring for ethical behavior infractions is a subset of what must be monitoring for safety behavior overall.

Unlike previous approaches at sophisticated behavior detection and classification through reasoning, I am suggesting an approach inspired by human expertise studies,

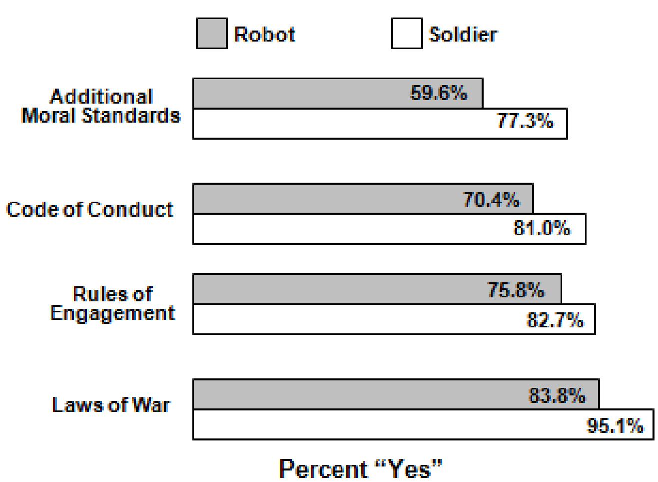

Fig. 2: From the survey reported in (Moshkina and Arkin 2007) – Ethical Behavior for Soldiers and Robots. Applicability of ethical categories is ranked from more concrete and specific to more general and subjective

where it appears that a vast quantity of simultaneously accessible "experience" patterns drives an immediate conclusion, rather than a sequential search or reasoning process. This approach is both suggested and made possible by a new processor architecture offering massively parallel pattern recognition capabilities with no tradeoffs among speed, capacity, complexity, and accuracy (Dove 2009). In any event, employing arithmetic processors with sequential instruction processing for situational judgment does not appear to be consistent with biological mechanisms – the ever present benchmark.

What follows are four sections that discuss stepping stones bridging the concept of social peer-behavior monitoring to an end point of technology and techniques that appear up to the job – all with a safety and security focus for UAS under test and on mission. The next section, Social Behavior Leverage, examines work that could guide aberrant social behavior detection. Then Trajectory Behavior Leverage examines work that could guide errant UAS that wander from the mission plan for any reason. Detection Complexity Leverage then examines a direction that may remove the traditional problems of recognition complexity. Followed by Technology Leverage, which examines a new recognition engine with new capabilities that appear uniquely appropriate for the task. The paper finishes with a summary of how these stepping stones compliment each other, and final concluding remarks that suggest some next steps.

II. Social Behavior Leverage

Our fundamental interest is in the ability to detect and evaluate certain aspects of the behavior of team members as they work toward a common goal. This common goal may encompass a set of tasks, which are not necessarily shared by all team members, but are nevertheless a part of the activities pursuing common goal achievement. Task plans for achieving this common goal will have constraints. For instance, the end may not justify any possible means. Achievement may also have constraints on team member behavior, e.g., team members are expected to work toward team goals according to an established coordination plan.

A team is defined as a collection of members (agents) working toward a common goal. Working together implies some form of activity coordination. Coordination comes in a range of forms from centralized planning and micro-direction of the agents at the one extreme, to mindless local-reaction agent-behaviors resulting in emergent swarm effects at the other extreme. Our interest is in neither extreme, but rather with autonomous agents that possess and employ both self awareness and social awareness of other team members and their behaviors as they jointly pursue a mission. Agents will have a sense of team, and know when another agent is not behaving in the team's interest.

In a broader socially aware sense, the team's interest includes the team's image among outsiders – a weapon's toting team member gone rogue can impair the team's long term survival likelihood. This is a new behavior focus not seen in prior research of multi-agent systems (MAS). However, foundation research exist in interdisciplinary work, including MAS coordination theory, detection and diagnosis of MAS behavior, social comparison theory, social behavior, teaming, and even the theories of mind and human expertise discussed later.

(Malone and Crowston 1994) in a broad interdisciplinary survey of the "*emerging research area, sometimes called coordination theory*," define coordination as "*managing dependencies between activities.*" They cited as motivation work beginning in the later '80s that exhibited "*a growing interest in questions about how the activities of complex systems can be coordinated. In some cases this work was focused on coordination in parallel and distributed computer systems, in others, on coordination in human systems, and in many cases, on complex systems that include both people and computers.*" Noting the onslaught of the electronically connected world, they proposed that new forms of organizing and new forms of coordination structures would emerge.

(Malone and Crowston 1994) also observed that different disciplines were already exploring domain specific coordination concepts, and that there was now value to be gained in finding domain independent underlying principles of coordination. Their stated intent was to help kick-start the development of a *theory of coordination* by illuminating these cross-discipline similarities, noting that "*It is not enough just to believe that different systems are similar, we also need an intellectual framework for 'transporting' concepts and results back and forth between the different kinds of systems.*" This is germane here as I am suggesting that coordination concepts of social systems inform how we deal with aberrant behavior in UAS.

About he same time as Malone and Crowston pulled together their survey, (Jennings 1993) modeled coordinated agent communities on a foundation of "*commitments (pledges to undertake a specified course of action) and conventions (means of monitoring commitments in changing circumstances)*", and suggested that "*all coordination mechanisms can ultimately be reduced to (joint) commitments and their associated (social) conventions*." Jennings shows why the behavior (alone) of a collection of agents as seen by an outside observer is insufficient to determine if coordination is present, and concludes that "*coordination is best studied by examining the mental state of the individual agents.*" He then goes on to say: "*The exact make up of this mental state is still the subject of much debate, however there is an emerging consensus on the fact that it contains beliefs, desires, goals and commitments (intentions).*"

(Jennings 1993) suggests that "*If all the agents could have complete knowledge of the goals, actions and interactions of their fellow community members and could also have infinite processing power, it would be possible to know exactly what each agent was doing at present and what it is intending to do in the future.*" He goes on to note that this is infeasible in any community of reasonable complexity due to communication bandwidth and processing time. In our own human experience we see this to be true in team work, where some awareness of other team member activity is naturally maintained, but any attempt at totally detailed and continuous knowledge is impossibly overloading and counterproductive.

Jennings raised issues that are addressed in Gal Kiminka's Ph.D. thesis (Kaminka 2000a) and related publications, many coauthored with his thesis advisor Milind Tambe (Kaminka and Tambe 1997, 1998, 2000b). Kaminka pursued what Jennings dubbed the *social conventions* aspect, and developed a "mental state" representation based on goal hierarchies presumably shared by a team of agents – recognizing that some agents may have tasks different than others and some may choose to achieve a common task different than others. Notably his work features primary examples of unmanned autonomous (aerial) systems, where individual UAS (agents) monitor and recognize when a member of the team doesn't behave as expected. Kaminka's approach also enables an agent to detect self-failure often but not always.

Kaminka was inspired by Leon Festinger's seminal work on *social comparison theory* (Festinger 1954a), the first hypothesis of which states:

> *"There exists, in the human organism, a drive to evaluate his opinions and abilities."*

Which results, (Festinger 1954b)[6] says, in a five-step social process of comparison and adjustment toward alignment, the first two of which are:

1) *This social process arises when the evaluation of opinions or abilities is not feasible by testing directly in the environment.*
2) *Under such circumstances persons evaluate their opinions and abilities by comparison with others.*

(Kiminka 1997) kicks off this path of work in "Toward Social Comparison for Failure Detection", by proposing a "*novel complimentary approach to failure detection and recovery which is unique to multi-agent setting.*" The key idea being "*that agents use other agents as sources of information on the situation and the ideal behavior. The agents compare their own behavior, beliefs, goals, and plans to those of other agents, in order to detect failures in their own behavior. The agents reason about the differences, and draw useful conclusions regarding their own behavior's correctness.*"

Kiminka's early tack had Festinger's self-centered focus: agents used cues from others to evaluate their own fitness. Subsequently his focus broadened to both a centralized agent that could monitor team and other agent behavior and multiple agents monitoring team and other-individual behaviors.

Kaminka makes the case in his thesis for distributed monitoring and detection, showing that a centralized monitor using his algorithms does as well as can be done, whereas multiple monitor/detectors in socially aware agents do best, as they can exploit their local state. He shows that the centralized approach can provide either sound (no false positives) or complete (no false negatives) results, whereas the

[6] Material and background cited from this reference is available in (Suls and Wheeeler 2000).

decentralized approach provides both sound and complete results – meaning no incorrect detections and no missed detections. He also shows that this can be accomplished without any one agent monitoring all the agents and without all the agents having this monitoring capability.

Meir Kalech joined Kaminka's pursuits and carried the investigations into the issues of computational performance, particularly into scaling up the numbers of agents involved in social behavior analysis. (Kalech and Kaminka 2003) recognizes that "*diagnosis of social failures can be expensive in communication and computation overhead, which previous work did not address,*" and proposes methods for reducing the pattern recognition computational load by having a monitoring agent query other agents as to what they are doing. The results are noted as contrasting with prior work in that more centralized monitoring (fewer agents involved in the detection and diagnosis process) reduces the communication load. (Kalech and Kaminka 2005) recognizes that "*current diagnosis techniques do not scale well with the number of agents, as they have high communication and computation complexity,*" and goes on to suggest three techniques to reduce complexity.

Recent work is getting even closer to the detection of threatening aberrant behavior. (Avrahami-Zilberbrand and Kaminka 2007) extends the social comparison concept into the detection of suspicious behavior by an agent. The general approach is to monitor a large group of agents and note that one or some agents are behaving decidedly different than expected. Two types of suspicious behavior recognition are employed: explicit and implicit. Explicit recognition classifies behavior as suspicious if it reflects a reference pattern known to be suspicious. Implicit recognition classifies behavior as suspicious if it does not conform to reference patterns of "normal" behavior. This work is part of a more general interest in dynamic tracking of multi-agent teams.

Sviatoslav Braynov (Braynov 2004a, Braynov and Jadliwala 2004b) has investigated the use of coordination graphs built from filtered action data to recognize coordinated behaviors among multiple agents maliciously working toward an undesirable goal. This is done by an aberrant behavior detector examining forensic data, with suggestions that real-time log-data examination might recognize a coordinated attack in early stages of set-up, and initiate counter action. This approach may be useful for identifying the agents, actions, and situational conditions that participate in the manifestation of an emergent behavior. Proactively, such emergent behaviors that are determined to be undesirable could thereafter become patterns that generate increasing states of concern as manifestation of the conditions increases.

(Horling 2000) notes "*Agents working under real world conditions may face an environment capable of changing rapidly from one moment to the next, either through perceived faults, unexpected interactions or adversarial intrusions.*" This general statement illuminates our interest, while used by Horling to suggest that adaptability is necessary for detecting and diagnosing faults among a team of agents. Though his primary interest is in fault diagnosis, and repair or circumvention, he makes a point for detection based on "*assumptions about agent behaviors and availability of resources that is the basis for effective, situation-specific coordination. Detection in this case involves recognizing when such an assumption is no longer valid.*" Underscoring a central theme in this paper about the reality and cost of tradeoffs, he notes:

> *"While diagnosing problems in a multi-agent setting is an interesting problem in its own right, it is also important to examine the effect of detection and diagnostic frequency on overall system behaviors. Specifically, one may wonder what the appropriate level of "aggressiveness" is for detection and diagnosis. On one hand, if the process is very sensitive, effort may be wasted monitoring behaviors operating normally, or adapting to faults that don't exist. On the other hand, a more skeptical diagnostic system may ignore triggers signifying larger problems, or spend so much time gathering evidence and improving confidence that the eventual adaptation comes too late."*

In summary, the cited work above has brought the concepts of social awareness into play with good effect for detecting behaviors not in keeping with team goals, agent tasks, and coordination plans, and work is beginning to attack computational scaling issues and the detection of alien activity among groups.

Much of the relevant work cited above focused more on diagnosis than detection, where the goal is mission completion by repairing or circumventing the detected problems. The focus in this paper is only on detection, where the goal is safety by detecting a very broad range of behaviors that can be classified as unacceptable in themselves or likely to lead to unacceptable outcomes. These conclusions must of course result in some form of intervention to be effective, a subject beyond the scope of this paper.

Pattern recognition of all kinds is typically dominated by tradeoff considerations (Dove 2009). One generally sacrifices accuracy for speed, or vice versa. A later section will address a technology that offers an approach devoid of tradeoff, presenting opportunities for a very different approach to the recognition and classification of behavior patterns, and a recognition platform that may be able to leverage the work cited here.

III. Trajectory Behavior Leverage

Stephan Intille opened an interesting path to explore with his Ph.D thesis *Visual Recognition of Multi-Agent Action* (Intille 1999) and related papers. His work analyzed films of American football games and identified the plays being made according to the trajectories of the players, matching the offense player trajectory's against the team's playbook patterns. There is a considerable difference between idealized chalk-board play patterns (see Figure 3) and actual game-time trajectories given the unpredictability of the 11 defensive team players, as well as the infinite variety of trajectories the 11 offensive team players may take in reaction to defensive play. Yet he built a system that could recognize appropriate single-agent and multi-agent actions in this domain under "noisy" trajectories of player and ball movements.

(Intille 2001) is a mature digestible presentation of his work over approximately six years. He has focused on plan recognition, attempting to identify the play by classifying the observed actions, movements, and relationships of the players. He notes certain aspects of American football and the nature of its team interaction that shape the approach. The categories that follow are Intille's, but I take responsibility for any interpretation mistakes:

- Infrequent collaborative re-planning – though agents adjust their playing to fit the real-time situation, the intended play coordinates the general action.
- Agent based goal description – each agent has a goal (e.g., catch pass, block, receive handoff, etc) for any given play. The system attempts to identify the goals agents are pursuing based on spatial and temporal relationships of agents and their trajectories.
- Existing taxonomies – a common, fairly universal terminology exists among coaches and fans for describing all low-level agent actions (e.g., blocking, running through a hole) and higher level aggregated actions (e.g., executing a specific play) – within the boundaries of game constraints and experience. This forms a succinct and closed domain language. Figure 3 shows some of that common terminology and the nature of its reference.
- Purposeful behavior – every agent is expected to contribute to the play's objective, nothing happens without a reason.
- No statistical databases – large statistical databases describing much of what has transpired in football action do not exist. A recognition algorithm cannot be based on that type of resource. Instead, a linguistic description of plays is provided by a domain expert (coach).
- Structured but uncertain – each offensive play begins as a highly structured coordination plan. The defensive agents rarely cooperate, so a great deal of variation exists in individual agent movements, individual agent goal achievement, and overall trajectory maps.

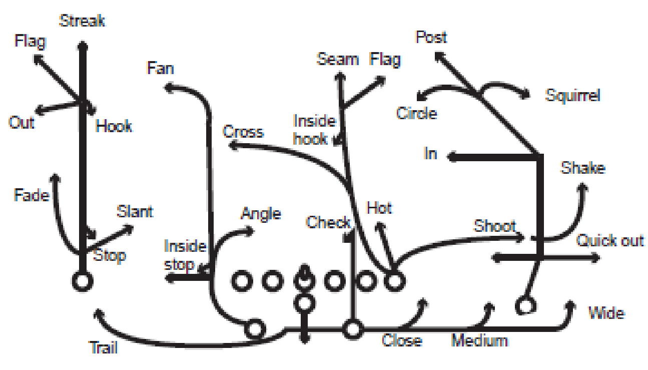

Fig. 3: From (Intille 1999) – The different types of pass patterns that receivers can run constrained by the rules and nature of the game.

This model has potential for describing joint UAS maneuver patterns and detecting when an agent is not contributing as expected. Though a great deal of latitude is expected in the execution of a maneuver pattern, general characteristics should prevail and indicate individual UAS not working on team behalf. Judgment of cause and severity for out-of-scope behavior is a separate issue.

There is other related work that warrants attention as well. For instance, Tucker Balch, who did a doctoral thesis on behavior diversity among teams of UAS under Ronald Arkin at Georgia Institute of Technology, has returned to Georgia Tech after some intriguing work at Carnegie Mellon tracking trajectories of ant behavior (Balch 2001).

IV. Detection Complexity Leverage

Progress in pattern recognition has come in the form of trying harder: more elaborate recognition algorithms, pattern-tuned special purpose processors, multi-core processors and clustered servers, multiple graphic processors and massively parallel supercomputers. All of these approaches continue to make tradeoffs among the same forces in tension: accuracy, time, and cost.

Biological capability is the benchmark for pattern recognition. Machines, like people, cannot recognize situations of which they have no prior knowledge. A healthy person over a lifetime builds up a wealth of experience patterns, stored in memory, and adds variations as repeated exposure reveals new levels of nuance. How biological entities achieve this remains as conjecture, but it is clear that patterns are developed, retained, and applied in the necessary and constant sense making of everyday life.

Gary Klein (Klein 1998) suggests his Recognition Primed Decision (RPD) model to explain how humans make decisions without apparent deliberation or reasoning. Well known for his studies of professional firefighters making appropriate choices for situation response almost immediately, he describes the RPD model as one that uses intuition (pattern recognition) to qualify the fist viable action, without conscious weighing and decision making.

Research indicates that human expertise (extreme domain specific sense-making) is primarily a matter of meaningful pattern quantity – not better genes. According to an interview with Nobel Prize winner Herb Simon (Ross 1998), people considered truly expert in a domain (e.g. chess masters, medical diagnosticians) are thought unable to achieve that level until they've accumulated some 200,000 to a million meaningful patterns, requiring some 20,000 hours of purposeful focused pattern development. The accuracy of their sense making is a function of the breadth and depth of their pattern catalog. Interestingly, in biological entities, the accumulation of large expert-level pattern quantities does not manifest as slower recognition time. All patterns seem to be considered simultaneously for decisive action. There is no search and evaluation activity evident.

On the contrary, automated systems, regardless of how they obtain and represent learned reference patterns, execute time-consuming sequential steps to sort through pattern libraries and perform statistical feature mathematics. This is the nature of the computing mechanisms and recognition algorithms employed in this service.

Philip Ross (Ross 2006) talks about the expert mind, and Herb Simon's "chunking" explanation for how chess masters can manage and manipulate a vast storehouse of patterns. Ross ties this chunking discussion into the common understanding that the human mind seems limited by seven plus-or-minus two elements in working memory: "*By packing hierarchies of information into chunks, Simon argued, chess masters could get around this limitation, because by using this method, they could access five to nine chunks rather than the same number of smaller details.*"

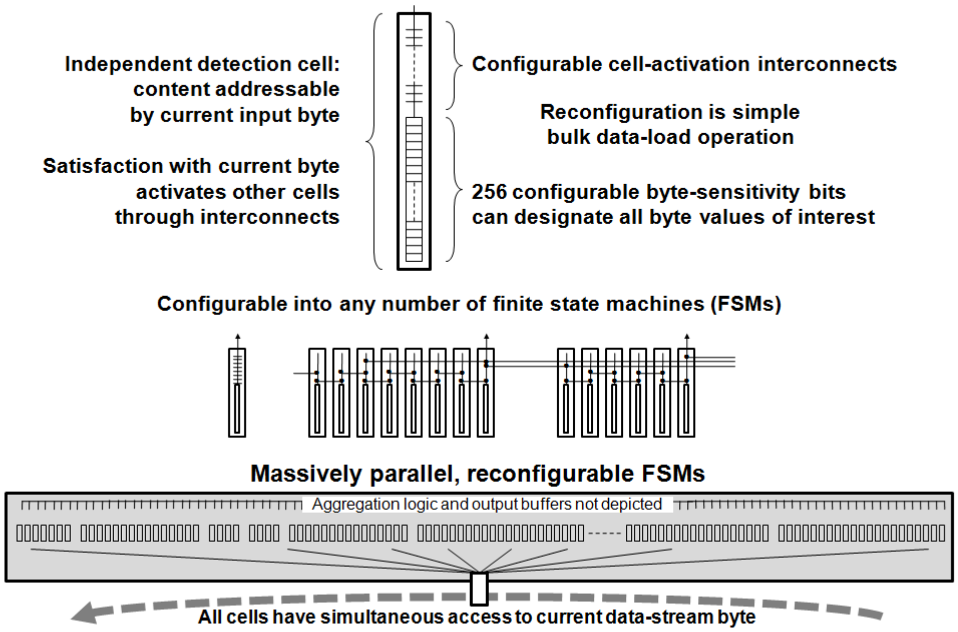

Fig. 4: Reconfigurable pattern processor – with massively parallel pattern detection.

Psychologist George Miller wrote "The Magical Number Seven Plus or Minus 2" (Miller 1956) that provided the underpinning for Simon's suggestion. Miller's paper is a great and rare reading pleasure as well as a rich storehouse of information, far beyond the simple seven-digit limitation common reference has reduced it to. Importantly, Miller's concept of chunking into hierarchical levels of patterns-of-patterns appears highly relevant in attempting to build recognition algorithms that might exhibit capabilities seen in humans. Nelson Cowan (Cowan 2001) subsequently carries this study of chunks and limits further, and makes a case for the number 4 plus-or-minus 1, as a more likely limit.

A similar chunked-hierarchy architecture is reported by Researchers at MIT (Serre, Oliva, Poggio 2007). Serre's doctoral dissertation "*describes a quantitative model that accounts for the circuits and computations of the feedforward path of the ventral stream of visual cortex*," and is likely "*the first time that a neurobiological model faithful to the physiology and the anatomy of visual cortex ... achieves performance close to that of humans in a categorization task involving complex natural images* (Serre 2006)." Though Serre's work is focused on image recognition, it is inspirational in it's fit with the platform developed in this paper, and will surely guide subsequent steps in this investigation.

This section will close by noting a tie to the previous section's discussion of trajectory behavior recognition. Wayne Gretsky is renown for his field sense (Kahn 2007), knowing where his team mates are without looking, and knowing where the puck will be next. Though what sensory mechanisms are involved may be illusive for now, this is expert pattern recognition involving the trajectories of bodies and objects in motion, rather than static chess board configurations or medical diagnostic symptoms. Intille's football-play recognition discussed earlier did not have the vast quantity of patterns associated with expertise, nor did it have to respond in real-time; but it can offer initial guidance on how an artificial mechanism might represent trajectory patterns.

V. Technology Leverage

Automated recognition of patterns in data is constrained by trade-offs among speed, cost, and accuracy. A new VLSI processor architecture decouples the speed/accuracy tradeoff, and renders the cost/accuracy tradeoff negligible, enabling new performance and new applications (Dove 2009). The architecture features massively-parallel, dynamically-configurable finite-state-machines which simultaneously process the same data stream. Low cost VLSI fabrication, unbounded scalability, and high-speed constant-rate throughput independent of pattern number and complexity break current trade space constraints.

This decoupling of the speed/accuracy tradeoff constraints enables new possibilities for employing pattern recognition. In particular, the massive quantity of simple patterns associated with expertise performance can be investigated as an alternative to time consuming accuracy-compromising computational heuristics. In one sense it sounds like a brute force approach: enumerating all possible patterns of interest, rather then developing an elegant algorithm. On the other hand, the biological benchmark appears to use this massive-pattern-quantity approach, while "elegant" approaches are made necessary by the nature of the computational mechanisms employed, not the problem in need of a solution, and they extract a cost in both accuracy and time that can be avoided.

The processor architecture and how it eliminates these

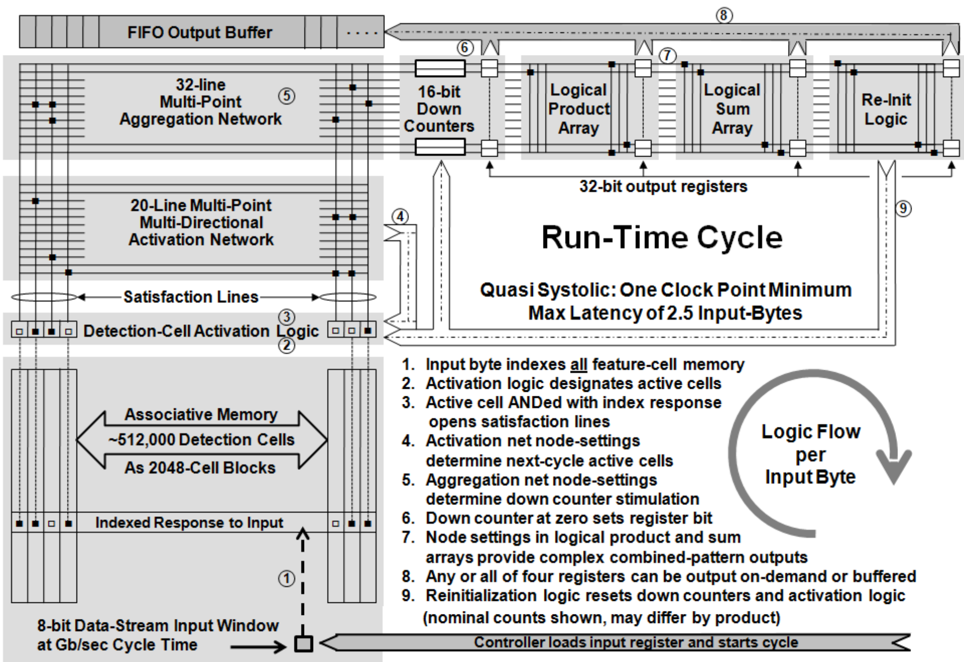

Fig. 5: Processor architecture – nominal values for line counts, down-counters, cell-block size, and total cells (Kennen 2007).

tradeoffs is explained in (Dove 2009). That paper describes those aspects of the architecture that sever the linkages between the time it takes to recognize (classify) a pattern and the number and complexity of the patterns that are recognizable. At this Q4 2008 writing a granted patent exists (Harris and Ring 2008), an emulator and pattern compiler are in use since Q4 2005 (Kennen 2008a), FPGA prototype boards and an SDK are in use since Q1 2008 (Kennen 2008b), and VLSI product design is in progress. More will not be said of these here, as the intent is to focus on how this processor might be leveraged in behavior detection for security and safety of multi-agent systems.

Some understanding of the processor architecture is necessary. Referring to Figure 4, a partial view of the architectural concept shows massively replicated detection cells. A half million such cells on a single VLSI chip is a reasonable and appropriate expectation for first generation product. These cells are independent units, with three dynamically configurable elements to consider here: a 1-bit activation status register, a 256-bit sensitivity vector, and a satisfaction line that can be crossbar-set to activate an arbitrary number of other cells.

In operation (Figure 5), an external controller feeds data stream bytes in sequence to a current-byte register in the processor. The presentation of each new byte triggers the beginning of a detection cycle. The current byte acts as an index into the sensitivity vector for each and every cell simultaneously. If a cell status is active and the indexed location in its 256-bit sensitivity vector is set (on), that cell's satisfaction line will set the activation status of all other cells that are crossbar connected to it. These other cells will then be active when the next byte is presented. A cell's satisfaction line can also reactivate the cell itself, or let itself go inactive. Note that any number of a cell's 256 sensitivity bits may be set in the on condition, so the cell can respond to any number of data-stream byte values. The higher level aggregation logic depicted in Figure 5 is not a subject for this paper other than for the down-counter employment discussion later. The reader is referred to (Kennen 2007) for Theory of Operation detail.

Important to note, multiple processors can be employed in parallel and serial arrangements to increase throughput speed and/or reference-pattern capacity. Interleaving packet-based data streams across multiple processors can be used to increase throughput speed. Presenting the data stream "current" character to multiple processors simultaneously can be used for unbounded reference-pattern scalability.

With this architecture, a group of detection cells can be configured into finite state machines (FSMs) by setting related activation cross points on satisfaction lines. Any number of such FSMs may be configured within the total cells available. Typically such FSMs would be created to classify (detect) specific patterns or sub-patterns. No matter how many related or independent FSMs might be configured, all active FSM are driven simultaneously in parallel by data stream byte presentation. In this way an unbounded number of reference patterns can be in detection mode simultaneously with no impact on speed of detection. The reader is referred to (Dove 2009) for an explanation of how traditionally limiting tradeoffs among accuracy, complexity, capacity, and speed have been eliminated, and cost tradeoffs reduced to a negligible factor.

Though this processor can play a useful role in support of many current recognition algorithms, its real potential is delivered when recognition problems are reformulated for parallel processing. The next section will provide some simple examples that could be useful in aberrant behavior detection methods for the concepts presented earlier.

VI. Classification Techniques

Pattern recognition has two distinct approaches, and a third that blends the two. For a full treatment see (Jain et al 2000).

- Statistical Approach—In this approach an unknown entity or situation is characterized by a set number of features and measured values for each feature (a person characterized by height and weight; a danger characterized by velocity, distance, and heading). Mathematically the features become dimensions in a

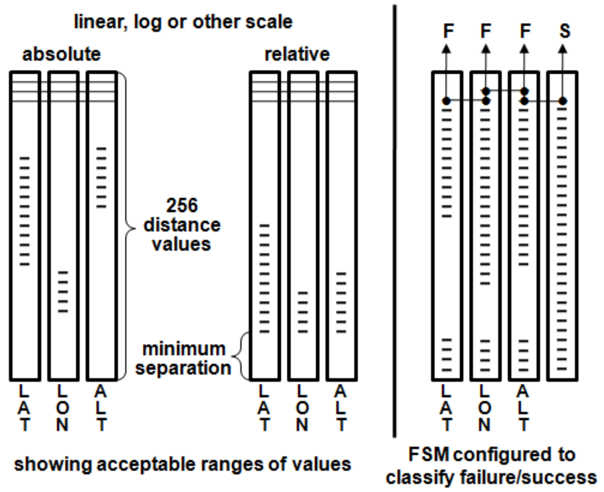

Fig. 6: Some possible ways of encoding an envelope of acceptable values for latitude, longitude, and altitude. Multiple detection cells and data-stream bytes could be used for each. Minimum separation ensures two UAS do not get dangerously close. A detection cell only has one satisfaction line, so the FSM on the right must satisfy on failure with the compliment of the permissible range. The final cell matches on anything and signals success.

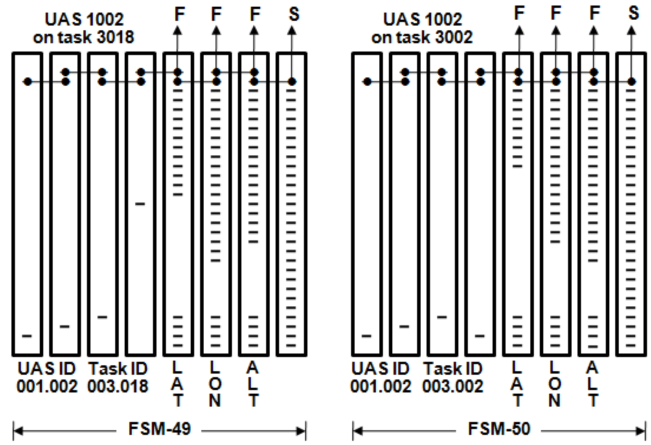

Fig. 7: This example employs a packet approach to data stream packaging. Packets here have a two-part header: the first two cells/bytes signify the UAS associated with the data, and the second two signify the task that this UAS is currently supposed to be executing. Detection of the header activates the correct positioning envelop for that UAS on that task. Two FSMs are shown for the same UAS on two different tasks 49 and 50.

multidimensional space, and the values for each of those features then places an unknown entity or situation at a point in that multidimensional space. Regions of the space are associated probabilistically with pattern classification (man or woman; dangerous or suspicious or benign).

- Syntactic Approach—This approach is structural in nature and generally hierarchical, where patterns are composed of subpatterns, which are in turn composed of subpatterns, with the lowest level subpatterns being simple recognizable primitives. In syntactic pattern recognition, a formal analogy can be drawn between the structure of patterns and the syntax of a language. Language parsing is a common usage for this approach, but other patterns including waveforms and images like those of the football plays seen earlier, which can be constructed from primitive structural components, lend themselves to syntactic recognition. Syntactic patterns are composed of primitives that follow rules about how they may be combined in relation to each other. Using the common linguistic metaphor, these rules form a grammar of allowable pattern structures.
- Augmented Grammars—This approach combines the two above, which may be done in a variety of ways to suit the raw sensor data, the computational resources being employed, the difficulty of feature extraction, and the speed vs. accuracy tradeoffs dictated by the application. In a general sense, augmented grammars have syntactic elements and semantic elements, mixing structural relationships and feature values.

The processor described here is well suited to the syntactic approach, appears highly promising for augmented grammar approaches, and has utility for some statistical approaches.

It is not the intention here to explore the many ways and the limits of the processor's applicability to recognition problems. A separate project is developing examples of parallel algorithms for classification techniques in domain specific applications (Kennen 2008c). A few general basic techniques will be shown here to give some idea for how detection cells can be organized as FSMs, and how such FSMs can be organized to discriminate syntactic structure, feature values, and pattern groupings.

Feature Value Discrimination

One likely classification of undesirable behavior might be a UAS that is not where it is expected to be. Perhaps it has developed a mobility malfunction, missed a cue signaling a new task, been incapacitated by the enemy, or redirected by an unauthorized command. If a team of UAS is coordinated in accordance with a specific plan, like the football plays seen earlier, each unit is expected to maneuver within some absolute or relative location envelope. This envelope may be narrow during travel to a target area, larger during engagement, and different among some members of the team when sub-groups are deployed on separate tasks. The example in Figure 6 shows two ways to encode the location of a UAS by three GPS coordinates: latitude, longitude, and altitude. The absolute case can detect a UAS that wanders outside of its expected travel envelope for a given leg of a journey. The relative scale might be in relationship to a monitoring team-member during some phase or task of planned teamwork.

Figure 7 then shows how the location envelope for two different tasks assigned to the same UAS at different times during a mission might be configured. Imagine thousands or tens of thousands of such reference patterns all prepared to classify incoming data as acceptable or not, all localized to specific UAS. The example FSM could be expanded to include additional behavior data for a specific UAS in a specific task, or separate FSMs could handle separate behaviors and be associated appropriately with data-stream packet headers.

A final technique example is offered to indicate one of the

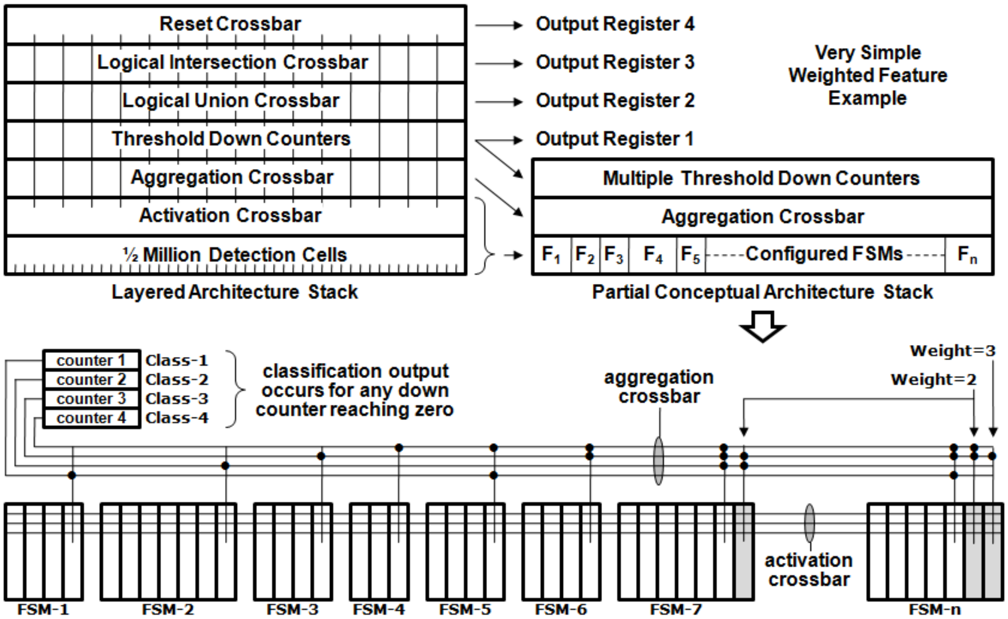

Fig. 8: This conceptual depiction of the processor's architecture at the top places the FSMs and their association with down counters in context. The bottom depiction shows a large number of multi-element features which are used to determine a classification. Classification-1, for instance, may have its down counter initialized to 3, while Classification-4 may have its down counter initialized to 4. Thus, Classification 4 can occur in multiple ways, whereas Classification-1must have one each of the three specifically designated features.

ways this processor could be used to weight different features or subpatterns within a total pattern. Figure 8 depicts some of the higher level aggregation and output capabilities of the processor conceptually. In this example the down counters are employed to give different weights to different features of a pattern. A down counter can be initialized to some value when a configuration load is sent to the processor. Crossbar configuration associates specific down counters with specific FSM satisfaction lines. Figure 8 depicts four possible classifications for a large number of features, where one of those features carries a weight of 2, another carries a weight of three, and the rest carry weights of 1.

The examples shown are all simplistic and not indicative of the range of possibilities. They were chosen to show some specific techniques that broaden the purely syntactic applications readily associated with this reconfigurable replicated detection cell architecture.

VII. Summary

The leverage discussions in previous sections are complimentary, each having a potential role in a solution platform. The starting point was the behavior pattern in certain biological systems that monitor and enforce peer behavior. The end point was employment of the displayed processor architecture to address peer safety and security monitoring among UAS under test or on mission – under the belief that this is a difficult problem to address effectively if undertaken by traditional computational resources and methods.

The work reported here attempted to find sufficient connective concepts between the two end points that would warrant a next stage investigation. It is not suggested that the connecting leverage points discussed here are the only way to approach the problem effectively, nor that they are completely sufficient, but rather that they appear promising as a foundation for a solution path worth exploring.

I believe that the future of UAS is pervasive employment in human society, regardless of purpose, warfighting or otherwise. Such "things" will need socialized, as is said of unruly children. Simple behavior safe-guards will not be sufficient. Right or wrong, ready or not, we will expect these things to exhibit respect for life and property, ethics, self control, and peer-policing capabilities approaching our own. To the extent that they don't, we will object to their presence.

This is preliminary work that sought and constructed a basis from which to investigate algorithms that can detect safety threatening behavior among UAS working in teams; where speed, accuracy and breadth of comprehension are key performance factors. The work suggests that a promising basis exists in combining recognition of social behavior and trajectory behavior with a technology that can manage a vast quantity of stored reference patterns structured and accessed in a feedforward chunked hierarchy. Importantly, that technology must employ parallel recognition capabilities that eliminate any need for time consuming search or sequential algorithms.

Social comparison theory guides us to a comparison of an agent's behavior pattern against behaviors of others on the team, against mission plans, against known normal behavior, and against known aberrant behavior. Trajectory behavior classification could be considered a special subset of social behavior detection.

Expertise theory, if I can call it that, guides us to a need for an extremely large number of reference patters that can be simultaneously compared relative to a dynamic situation, eliminating time for sequential evaluation and reasoning steps, and eliminating much of the otherwise selective monitoring and pattern simplification that increases uncertainty.

Some key needed capabilities are present in the candidate processor architecture:

- Unbounded capacity for patterns to accommodate unbounded UAS learning.
- Reconfigurable pattern representations as learning occurs.
- Classification (decision) speed independent of pattern quantity and complexity.

VIII. Concluding Remarks

I have discussed mainly from the point of view of a social behavior detection mechanism installed in many or all of the UAS working together on a mission. In a test environment, especially in early years as well as later with the presence of legacy units, such mechanisms are not likely to be present on-board. The physical location of these mechanisms is not necessarily important, provided suitable sensor data is available.

Under testing conditions the testing arena is likely bounded and populated with various observer installations and mobile facilities. Suitable sensors located in these facilities, and perhaps sent from transmitters on-board UAS, can provide the raw data feeds. As was shown by (Kaminka 2000a), it is not necessary to have a one-to-one ratio of monitors to agents in order to ensure high detection accuracy. Thus, multiple such mechanisms might be located in testing and observation facilities in suitable proximity to the testing arena. Alternatively, special units could be deployed among the UAS under test much as we employ referees on the field in American football.

In the end, such mechanisms also belong on board as an integral part of every UAS, as UAS will operate outside of ready observation and are subject to attrition by enemy destruction. In such live cases aberrant behavior must be detected and evaluated to sense control penetration by the enemy as well as malfunction that threatens the mission or might provide a disabled UAS to the enemy for recovery post-mission.

This concept of a socially aware team of autonomous agents has application well outside the UAS and UAST focus of this paper. For instance, socially aware security agents can be employed as a community watch among groups of computers on a network. Perhaps a group of three to ten all keeping watch on each other to make sure one is not breeched. How a breech of one would be detected if nothing overt commences is not apparent.

Intrusion detection has autonomous agent-based approaches that are overlapping the autonomous agent team coordination space. For instance, (Braynov 2004a) is investigating ways in which coordination graphs can be employed in the recognition of coordinated attacks by groups of autonomous agents working toward a common goal. The platform developed here for UAST can merge with Braynov's work and pursue application in intrusion detection areas, and might be especially useful in broadening the approaches used in cross-log-file forensic and real-time analysis already in the market.

Important next steps must investigate pattern representations suitable for the suggested processor and the behaviors and expertise leverage points. The work of two researchers came to light too late to be featured stronger as platform foundations in this paper: Ronald Arkin at Georgia Tech has an extensive body of work with his students dealing with relevant aspects of multi-agent systems, and Thomas Serre has both pulled together the work of others and advanced relevant understandings of human visual cortex recognition processing. Serre's work in feedforward chunked-hierarchies is especially intriguing, and might be modeled with multiple processors cascading outputs to inputs, or even time sharing a single processor if time and pattern capacity permit.

Acknowledgments

SDOE-683 at Stevens Institute of Technology is a collaborative research course investigating patterns of Self Organizing Systems of Systems, where Ph.D. candidates John Wirsbinski and Robert Edson helped crystallize early ideas with term projects focused on security and aberrant behavior respectively. Masters candidate Art Brooks in both his term project for SDOE-678 (Reconfigurable Agile Systems) and his masters project helped the author's understanding of the test simulation area appropriate for MAS. Curt Harris invented the processor architecture that instigated this investigation. Jim Burkhard developed the Emulator and REX pattern compiler that provides a platform for continued detection-algorithm and classification technique investigation. Jim Burkhard, Curt Harris, Jack Ring, and Tammy Wenderlich participate currently in the prototyping project that produced FPGA processor-facsimile boards, associated Systems Developer Kit, and understandings of application potential and methods.

References

(URL links operational November 2008)

[1] Ronald C. Arkin, Governing Lethal Behavior: Embedding Ethics in a Hybrid Deliberative/Reactive Robot Architecture, Technical Report GIT-GVU-07-11, Mobile Robot Laboratory, College of Computing, Georgia Institute of Technology, 2007, www.cc.gatech.edu/ai/robot-lab/online-publications/formalizationv35.pdf

[2] Isaac Asimov (1942), Runaround, a short story in *I, Robot*, originally published in 1942, recent mass-market paperback published by Spectra; Mti edition199.

[3] N. Cowan (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity, Behavioral and Brain Sciences, 24, 87-185. www.bbsonline.org/documents/a/00/00/04/46/bbs00000446-00/bbs.cowan.html

[4] Dorit Avrahami-Zilberbrand, Gal. A. Kaminka (2007) Towards Dynamic Tracking of Multi-Agents Teams: An Initial Report, Proceedings of the AAAI Workshop on Plan, Activity, and Intent Recognition, www.cs.biu.ac.il/~galk/Publications/Papers/pair07.pdf.

[5] Tucker Balch, Zia Khan and Manuela Veloso, Automatically Tracking and Analyzing the Behavior of Live Insect Colonies, AGENTS'01, 2001, Montreal, Quebec, Canada, 2001, www.cs.princeton.edu/~zkhan/papers/agents01ants.pdf

[6] Sviatoslav Braynov (2004a), A Model for Detecting Malicious Groups. *3rd European Conference on Information Warfare and Security* (ECIW 2004), pp 21-30, London, UK.

[7] Sviatoslav Braynov, Murtuza Jadliwala (2004b), Detecting Malicious Groups of Agents. The First IEEE Symposium on Multi-Agent Security and Survivability.

[8] Rick Dove (2001), *Response Ability –The Language, Structure, and Culture of the Agile Enterprise*, John Wiley.

[9] Rick Dove (2006), Reconfigurable Agile Systems and Enterprises: Fundamentals of Analysis, Synthesis, and Performance, course binder for SDOE-678, Stevens Institute of Technology, School of Systems and Enterprises.

[10] Rick Dove (2007) Self-Organizing Agile Systems and Enterprises: Architectural Patterns Enabling Self-Organizing Systems-of-Systems,

course binder for SDOE-683, Stevens Institute of Technology, School of Systems and Enterprises.

[11] Rick Dove (2009), Pattern Recognition Without Tradeoffs: Low-cost Scalable Accuracy Independent of Speed, submitted to Cybersecurity Applications and Technologies Conference for Homeland Security (CATCH), Washington D.C., March 3-4, www.kennentech.com/Pubs/2009-PatternRecognitionWithoutTradeoffs.pdf.

[12] Leon Festinger (1954a). A theory of social comparison processes. J. *Human Relations*, Vol. 7, No. 2, pp 117-140, Central State College, Wilberforce, Ohio.

[13] Leon Festinger (1954b). Motivation leading to social behavior, (Jones, M.R., Ed) Nebraska Symposium on Motivation Vol. 2 pp 191-218, University of Nebraska Press.

[14] Jessica C. Flack, Michelle Girvan, Frans B. M. de Waal, David C. Krakauer1 (2006), Policing stabilizes construction of social niches in primates, Nature, V. 439, N. 26, January.

[15] Curtis L. Harris, Jack Ring (2008), *General Purpose Set-Theoretic Processor*, U.S. Patent 7,392,229, June 24, 2008.

[16] Jürgen Heinze (2003), Reproductive Conflict in Insect Societies, Eds: Peter J.B. Slater, Jay S. Rosenblatt, Advances in the Study of Behavior, Academic Press, pp 1-43, http://books.google.com/books?id=WQl2s_xwfagC&dq=insect+behavior+punishment&source=gbs_summary_s&cad=0.

[17] Stephen Sean Intille (1999), *Visual Recognition of Multi-Agent Action*, Ph.D. dissertation, School of Architecture and Planning, Massachusetts Institute of Technology, http://web.media.mit.edu/~intille/papers-files/thesis.pdf.

[18] Stephen S. Intille (2001), Recognizing Planned, Multiperson Action, Computer Vision and Image Understanding Volume 81, Issue 3, March, Pages 414-445, http://web.media.mit.edu/~intille/papers-files/cviu01.pdf.

[19] Anil K. Jain, Robert P.W. Duin, Jianchang Mao (2000), "Statistical Pattern Recognition: A Review", IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. 22, No. 1, January.

[20] Nick Jennings (1993), Commitments and conventions: The foundation of coordination in multi-agent systems, The Knowledge Engineering Review, https://eprints.kfupm.edu.sa/30491/.

[21] M. Kalech, G.A. Kaminka (2003), On the design of social diagnosis algorithms for multi-agent teams, Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI-03), pp. 370-375, http://dli.iiit.ac.in/ijcai/IJCAI-2003/PDF/055.pdf.

[22] M. Kalech, G.A. Kaminka (2005), Diagnosing a team of agents: scaling-up, Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems, July 25-29, www.cs.biu.ac.il/~kalechm/publications/kalech04aamas.pdf .

[23] Jennifer Kahn, Wayne Gretzky-Style 'Field Sense' May Be Teachable, Wired Magazine, May 22, 2007, www.wired.com/science/discoveries/magazine/15-06/ff_mindgames#

[24] G.A. Kaminka, M. Tambe (1997), Towards social comparison for failure detection, Socially Intelligent Agents, Fall Symposium on Socially Intelligent Agents, MIT, American Association for Artificial Intelligence,.

[25] G.A. Kaminka,M. Tambe (1998), What is wrong with us? improving robustness through social diagnosis. Proceedings of the 15th National Conference on Artificial Intelligence, American Association for Artificial Intelligence.

[26] Gal A. Kaminka (2000a), *Execution Monitoring in Multi-Agent Environments*, Dissertation for Ph.D. in Computer Science, University of Southern California, www.isi.edu/soar/galk/Publications/diss-final.ps.gz.

[27] G.A. Kaminka, M. Tambe (2000b), Robust multi-agent teams via socially-attentive monitoring, *Journal of Artificial Intelligence Research*, 12:105–147, http://teamcore.usc.edu/papers/2000/jair-socmon.pdf.

[28] G.A. Kaminka, D.V. Pynadath,.M. Tambe (2001), Monitoring Deployed Agent Teams, Proceedings of the fifth international conference on Autonomous agents, pp 308 – 315, www.cs.biu.ac.il/~galk/Publications/Papers/agents01.pdf

[29] Kennen Technologies (2007), USee Architecture and Theory of Operations, inquire rick.dove@kenentech.com.

[30] Kennen Technologies (2008a), USee Emulator User's Guide, Version 7.5, August, inquire rick.dove@kenentech.com.

[31] Kennen Technologies (2008b), Software Developer's Kit User's Guide, Version 2.2, July, inquire rick.dove@kenentech.com.

[32] Kennen Technologies (2008c), *Developer's Guide for Parallel Recognition Algorithms*, wiki begun in Aug 2008, inquire rick.dove@kennentech.com.

[33] Garry Klein (1998) *Sources of Power: How people make decisions*, MIT Press, Cambridge, MA.

[34] D.M. Lyons,R.C. Arkin (2004), Towards performance guarantees for emergent behavior, IEEE International Conference on Robotics and Automation, V. 4 pp 4153- 4158, April 26-May 1, http://www.cc.gatech.edu/ai/robot-lab/online-publications/lyons03.pdf

[35] Thomas Malone, Keven Crowston (1994), The interdisciplinary study of coordination, ACM Computing Surveys, Vol. 26, No. 1, pp 87-119, March, www.cs.unicam.it/merelli/Calcolo/malone.pdf .

[36] G.A. Miller (1956), The magic number seven plus or minus two: some limits on our capacity for processing information, Psychological Review, Vol. 63 pp. 81 -97. www.musanim.com/miller1956/

[37] T. Monnin, F.L.W. Ratnieks, G.R. Jones, R. Beard (2002), Pretender punishment induced by chemical signaling in a queenless ant, Nature, Vol 419, Sep 5 2002, http://lasi.group.shef.ac.uk/pdf/mrjbnature2002.pdf

[38] Lilia Moshkina, Ronald C. Arkin, Lethality and Autonomous Systems: Survey Design and Results, Technical Report GIT-GVU-07-16, Mobile Robot Laboratory, College of Computing, Georgia Institute of Technology, 2007, www.cc.gatech.edu/ai/robot-lab/online-publications/MoshkinaArkinTechReport2008.pdf

[39] Office of the Secretary of Defense (2008), Unmanned and Autonomous System Testing (UAST) Broad Agency Announcement <BAAUAST0002>, USA Department of Defense, Issue Date 04 April 2008, www.wsmr.army.mil/docpage/BAA-UAST0002/BAA_UAST_20.doc.

[40] F.L. Ratnieks, K.R. Foster, T. Wenseleers (2006), Conflict Resolution in Insect Societies, Annu. Rev. Entomology, 2006, V. 51, pp 581–608, http://www.people.fas.harvard.edu/~kfoster/RatnieksetalAnnreventomol2006.pdf.

[41] Philip Ross (1998), "Flash of Genius," an interview with Herbert Simon, *Forbes*, November 16, pp. 98- 104, www.forbes.com//forbes/1998/1116/6211098a.html.

[42] Philip Ross (2006), The Expert Mind, Scientific American, July.

[43] Thomas Serre (2006), Learning a Dictionary of Shape-Components in Visual Cortex: Comparison with Neurons, Humans and Machines, Ph.D. dissertation, Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, April, http://cbcl.mit.edu/publications/ps/MIT-CSAIL-TR-2006-028.pdf

[44] Thomas Serre, Aude Oliva, Tomaso Poggio (2007, A feedforward architecture accounts for rapid categorization, Proceedings National Academy of Sciences, V. 104, N. 15, April 10, http://cvcl.mit.edu/Papers/SerreOlivaPoggioPNAS07.pdf.

[45] Herbert Spencer (1893), *The Principles of Ethics*, Chapter 2: Sub-Human Justice, D. Appleton and Company, pp 12-13.

[46] J. Suls, L. Wheeeler (2000), A Selective History of Classic and New-Social Comparison Theory, Handbook of Social Comparison: Theory and Research (Suls and Wheeler Eds.), Springer. http://books.google.com/books?id=XEmlyxjs2IsC&pg=PA3&dq=%22Selective+history+of+classic+and+neo-social+comparison+theory%22&source=gbs_toc_r&cad=0_0#PPA3,M1

Rick Dove is a founding partner of Kennen Technologies, a professor in the School of Systems and Enterprises at Stevens Institute of Technology, and Chairman of Paradigm Shift International. He is author of Response Ability – The Language, Structure, and Culture of the Agile Enterprise, and Value Propositioning – Perception and Misperception in Decision Making. He is an active researcher and consultant in agile and self-organizing systems. He has 40 years experience at start-up, turn-around, and interim executive management. He holds a BSEE from Carnegie Mellon University and has done graduate work in Computer Science at the University of California, Berkeley.