2005

Detecting Suspicious Behavior from Only Positional Data with Distributed Sensor Networks

Rowe, Neil C.

Monterey, California. Naval Postgraduate School

Proceedings of DTEC '05, 2005 ASME International Design Engineering Technical Conferences and Computers and Information in Engineering Conference
http://hdl.handle.net/10945/36816

Proceedings of DTEC '05
2005 ASME International Design Engineering Technical Conferences and Computers and Information in Engineering Conference

DETC 2005-84420

DETECTING SUSPICIOUS BEHAVIOR FROM ONLY POSITIONAL DATA WITH DISTRIBUTED SENSOR NETWORKS

Neil C. Rowe
Modeling, Virtual Environments, and Simulation (MOVES) Institute
U.S. Naval Postgraduate School
Code CS/Rp, 833 Dyer Road, Monterey, CA 93943
ncrowe@nps.edu

ABSTRACT

Wireless sensor networks are increasingly popular, and are being used to measure simple properties of their environment. In many applications such as surveillance, we would like them to distinguish "suspicious" behavior automatically. We distinguish here between suspicious and anomalous behavior, and develop a mathematical model which we illustrate on some sample data. We show the model predicts six classic deception strategies. We conclude with analysis of more sophisticated deceptions that exploit system responses to simpler deceptions.

Keywords: sensors, networks, suspicious, behavior, locations

INTRODUCTION

Wireless sensor networks [1] are increasingly being implemented for surveillance applications. Such surveillance is often concerned with detection of behavior that is suspicious in some way, as for the growing field of “homeland security”. For instance, we would like to detect people planting bombs in a public urban area by monitoring by a sensor field. However, suspicious behavior is usually rare, and it can be extremely tedious for people to monitor for those rare occurrences. It would be desirable, therefore, to automate the detection of suspicious behavior.

The work reported here assumes a long-term (“persistent”) wireless sensor network with large numbers of “small and cheap” sensors (each with limited capabilities) in just one sensing modality. Such microsensors could monitor sound levels, metal objects passing nearby, or electromagnetic emissions. Such microsensors that monitor one-dimensional signals without the ability to directionalize them can be made cheaply in millions of units. Because of their simplicity, they also pose less a threat to privacy of individuals than more sophisticated sensors such as imaging ones. For processing, we will assume a much smaller number of “collector” sensors with strong antennas and significant power that aggregate the data from the microsensors, a common design approach today [2]. We assume that the primary purpose of processing is to localize animate agents moving within the sensor field as by tracking footfalls and body heat.

Large numbers of microsensors cannot be manually emplaced but must be dispersed with semi-random methods such as firing from a gun or dropping from an aircraft [3]. An initialization phase is then required where microsensors report to collector sensors to enable their localization. For instance, a sensation-creating source such as a noisy device can be sent on a fixed path through the sensor field, and the microsensors are designed to report when they sense it together with their identification number; the pattern of reports can be fitted to formulae to approximate the location of each reporting microsensor. Tracking then can be done in the presence of occlusions using best-fit methods [4].

CLUES FOR DETECTING DECEPTION

The question is how to define suspicious behavior of a source in a sensor field based on just its position over time and its derivatives. Our approach here is to postulate that suspicious behavior is behavior that shows evidence of deliberate deception. The problem of detecting deception is important with many applications in law, business, psychology, and military operations, and a number of nonverbal clues can be used to detect it [5, 6]. [7] outlines the main clues in human interactions:

- Visual: increased blinking, increased self-grooming, increased pupil dilation;
- Vocal: increased hesitation, shorter responses, increased speech errors, higher voice pitch;
- Verbal: increased overgenerality, increased irrelevance, more frequent negations, more frequent hyperbole

These are not directly observable from positional information, but there are analogies to many of them. Increased blinking and self-grooming have an analogy in uncertainty about path direction and speed; increased hesitation and increased errors have an analogy in unnecessary stops and starts; and shorter responses and higher voice pitch have an analogy in increased speed of the agent. So a sensor network can look for such clues.

Deceptive activities are common in warfare [8]. A classic instance of the recognition of suspicious behavior with primarily positional data is naval air defense, the process of analyzing sensor data, primarily from radar, to discover which aircraft pose potential threats to a military ship [9]. An enemy aircraft that wants to threaten or attack a ship will try to conceal their intentions. That means they will not often provide electronic identification or other obvious clues. But it is very difficult to conceal a radar profile and a course, and we should still be able to track their position over time.

Personnel in the Command Information Center of U.S. Navy ships are trained thoroughly on the skill of detecting suspicious behavior of “air contacts” or radar-visible aircraft from the few clues available. [10] did surveys with air-defense experts to identify the following factors in decreasing order of value: (1) electronic emissions, (2) altitude, (3) course, (4) speed, (5) whether the aircraft is in an airlane, (6) the airport of origin, (7) the distance from the ship, (8) evidence of a weapons system, (9) evidence of coordination with other aircraft, (10) IFF (Identify Friend or Foe) system information if available, (11) time to reach the ship, (12) whether the aircraft is over water, (13) whether the aircraft has launched missiles (important but rare), (14) time in the air, (15) number of apparent manuevers, (16) apparent fuel remaining, and (17) visibility conditions. [10] treats these factors as equal in import. But clearly some of the factors (1, 2, 5, 6, 9, 10, 12, 15) concern the suspiciousness of the contact and others the dangerousness of the contact. Suspiciousness functions like a probability and dangerousness functions like a cost, so we ought to multiply their cumulative assessments. Cumulative probability of suspiciousness can be obtained most simply by a Naïve Bayes model for the suspiciousness probabilities and cumulative cost by a weighted average of danger factors.

HEURISTICS FOR RECOGNIZING DISCREPANCIES FROM NORMAL BEHAVIOR

We would like to detect suspicious behavior in a sensor field. If we have video imagery as [11] and [12] did, we could analyze more subtle clues to suspicious behavior, but we would like to confine ourselves to the positional information only, obtained by simple sensors. So we propose that a sensor network try to observe the following clues. These represent "discrepancies" [13] that could be noted from normal behavior, a key factor in detecting deception [14]. Most have analogies in naval air defense.

- Loud noises and noticeable signals associated with the agent (these are not positional but are often used to infer position with the kind of sensors we envision, like footstep-sensing ones).
- High speed, since it contributes to surprise effects.
- Abrupt changes in path direction or speed, since this suggests real-time planning of something and is atypical of normal behavior [15].
- Communication with other agents that are themselves suspicious (a “contagion” effect). Much of intelligence analysis consists of finding enemy communication channels since this can reveal the full nature of the enemy [16].
- Avoidance of visibility by other agents, to enable secrecy.
- Isolation and lack of communication with other agents, since this suggests the agents are outsiders.
- Lack of goal-directed behavior (“lurking”), since it suggests looking for trouble – but this is not a strong clue since legitimate agents may lack immediate goals too.

Much as with anomaly-based “intrusion detection” of attackers on computer systems [17], we can define a distribution of values for each of several numeric metrics for the above characteristics. Then the sum of those random variables will approach a normal distribution by the Central Limit Theorem of probability theory, whose mean and variance will be defined by the means and variances of the individual distributions. If an agent's discrepancy total exceeds more than a threshold number of standard deviations from the mean of the joint distribution, we will consider it suspicious. The first four clues are positive (that is, larger values are suspicious) and the last three are negative.

- High speed and loud noises can be defined from statistics on agent speeds or sounds.
- Abrupt changes in path direction or speed can be defined from statistics on the derivatives of speed and heading with respect to time.
- Isolation of the agent can be observed by the history of rendezvous-type actions among other agents. A rendezvous is when two agents approach each other closely for a period of time.
- Lack of visibility to other agents can be measured by the average sum of the inverses of the distances of all other agents (since visibility is roughly inversely proportional to distance), ignoring occluded lines of visibility.
- Lack of goal-directed behavior can be observed by the consistency of the overall direction of movement. This is different from observing abrupt changes in path direction because it should ignore local changes in the path by smoothing techniques.
- To these we can add domain-specific clues, like the suspicious behaviors of vehicles in parking lots cited in [12] that included circling and back-and-forth motions.

GENERAL MODELING OF SUSPICIOUS BEHAVIOR

However, atypicality is not the same as suspiciousness. We will get too many false alarms if we presume so, since people engage in many different legitimate kinds of activities including rare ones. To do better we need to model the notion of deception in the activity.

Following [5], we postulate that deception activities have goals or methods that must be concealed. Thus visibility of the deceptive agent by other agents is key factor, since visible activities needing of deception, or especially deception that is itself visible, can often invoke a response that thwarts the goals of the deceptive agent either directly or indirectly. So a metric for the “exposure” of a deceptive agent D is

$$E = (1/T)\int_{t=0}^{T} s(t)\sum_{i=0}^{N} v(i,t)n(i,t)dt \qquad (1)$$

where T is the duration of the time interval of observation, s(t) is the intrinsic suspiciousness of D at time t on a scale of 0 to 1, N is the number of non-deceptive agents in the sensor field, v(i,t) means the visibility of D at time t from non-deceptive agent i on a scale of 0 to 1, and n(i,t) means the "noticeability" of D's suspiciousness at time t by agent i when D is visible to agent i on a scale of 0 to 1. If we receive information from the sensor network at discrete times T, we can approximate the integral by a summation to get

$$E = (1/T)\sum_{t=0}^{T-1} s(t)\sum_{i=0}^{N} v(i,t)n(i,t) \qquad (2)$$

But this is usually more of an approximation because it is difficult to include point events such as explosions occurring between discrete time points; for compatibility with the integral, the values of s, v, and n for each time interval must aggregate all events of that time interval.

The s(t) must reflect two kinds of factors: Fixed factors that do not change during D's appearance in the sensor field, such as that D is wearing a bulky raincoat on a sunny day, and dynamic factors that do, like D's pacing back and forth. Since we are limiting the sensor networks in this paper to those which sense only position of agents, most fixed factors are unavailable to us. The dynamic characteristics that can affect s(t) are the coordinates of the agent, its velocity vector, and its acceleration vector. Normal activities exhibit a population of state vectors; s(v) represents the degree to which an agent's state vector is atypical of this population, obtained from a distribution fitted from statistics.

Estimation of v(i,t) must take into account occlusion phenomena by obstacles and terrain. It is also difficult for the deceptive agent to know if they are visible to a particular agent because human agents have eyes only on one side of their heads and hearing is not especially helpful in this task. So we must postulate that deceptive agents can only approximately estimate their v functions. Nonetheless, it must clearly decrease monotonically with distance between agents. Sight, sound, vibration, chemical, and radiation clues generally decrease with the square of the distance since they can be modeled as being spread out on the surface of a sphere expanding about a point source. It is also helpful to postulate a "virtual observer" for i=0 which represents the notion that someone may be watching that the deceiver cannot see since it is hard to check every possibility.

The n(i,t) reflects the degree to which the deceptiveness of D, when clearly visible to an agent, will be noticed by the agent. This is the lack of "cover" as the term is used in surveillance. As with s(t), there are factors fixed over the course of the visit to the sensor field, such as the distraction of the agent by personal worries, and dynamic factors, such as the number of people present. We cannot recognize fixed factors except quite indirectly, but most of the dynamic factors can be estimated from the number of people nearby the agent i and D; in other words, D "blends in" when many people are present.

APPROXIMATION OF THE SUSPICIOUSNESS FORMULA

Useful approximations of the E function can be obtained given a distribution p(x,y), representing the probability of observing an ordinary agent at that location (x,y), and a distribution q(s,h), representing the probability of observing an ordinary agent with that speed and heading (we ignore acceleration because it is harder for people to observe than the velocity vector). We can obtain p and q from statistics from observations of the sensor field over a long period of time. Position and the velocity vector are mostly independent of one another, so the joint probability of observing a particular agent at a particular location with a particular velocity vector is p(x,y)q(s,h). Then a first-order approximation of intrinsic suspiciousness is the reciprocal of the probability of observation, or

$$s(t) = k_1 / (k_1 + p(x,y) q(s,h)) \qquad (3)$$

Noticeability is also related to the density of agents, since suspicious features are harder to distinguish in a crowd than in isolated individuals. So we can approximate noticeability by

$$n(i,t) = k_3 / (k_3 + p(x,y)) \qquad (4)$$

Visibility is closely related to angular resolution of details, which varies inversely with the square of the distance from the source (since at twice the distance, the area of a visible object is one-fourth as large). In areas in which there are fewer agents on the average, visibility of suspicious actions is decreased because there are fewer observers. But more importantly, the presence of visual obstacles will impede visibility. In general,

$$v(t) = k_2 \int_{\theta=0}^{\theta=2\pi} \int_{r=M}^{r=\max(b(\theta),M)} (p(x + r\cos\theta, y + r\sin\theta)/r)\, dr\, d\theta \qquad (5)$$

where $b(\theta)$ is the distance of the nearest obstacle from (x,y) in the direction of polar angle θ, m is the minimum radius (at worst, the radius of agent D), and M is the maximum radius (to prevent the integral from being unbounded). There is a single r in the denominator since the r in the numerator due to the polar integration is canceled by one of the two in the denominator due to the inverse square relationship.

In many areas of terrain we can assume p(x,y) is uniform, as in travel corridors where there is no reason to stop. If there are no obstacles in an area, we will assume v(t)=1 or its maximum possible. But if there are visual obstacles, they remove some of the space from visibility consideration. Consider Figure 1 where a deceptive agent is distance d from a rectangular obstacle that subtends a range of heading $\theta = \alpha$ to heading $\theta = \beta$ as viewed from the agent, where $\theta = 0$ is the ray perpendicular to the nearest surface of the obstacle. We assume there is a minimum distance m that other agents can approach the deceptive agent (which could be, at least, the radius of the deceptive agent) as well as a maximum distance M at which they cannot see anything significant. Then v is reduced by both the range of heading that the obstacle subtends and the portion of distance within that range beyond which an observer is blocked by the obstacle. Hence:

$$v(t) = 1 - \frac{\beta - \alpha}{2\pi} \int_{\theta=\alpha}^{\theta=\beta} \int_{r=d/\cos\theta}^{r=M} (1/r)\, dr\, d\theta \Big/ \int_{\theta=\alpha}^{\theta=\beta} \int_{r=m}^{r=M} (1/r)\, dr\, d\theta$$

$$= 1 - \frac{(\beta - \alpha)^2 \ln \frac{M}{m}}{2\pi} [(\beta - \alpha) \ln(\frac{M}{d}) + \int_{\theta=\alpha}^{\theta=\beta} \ln(\cos\theta)\, d\theta]$$

(6)

We can approximate this last integral numerically. The whole formula gives us a way of calculating visibility for a given position of a deceptive agent near a straight obstacle edge, and can be generalized to obstacles of many edges and multiple obstacles.

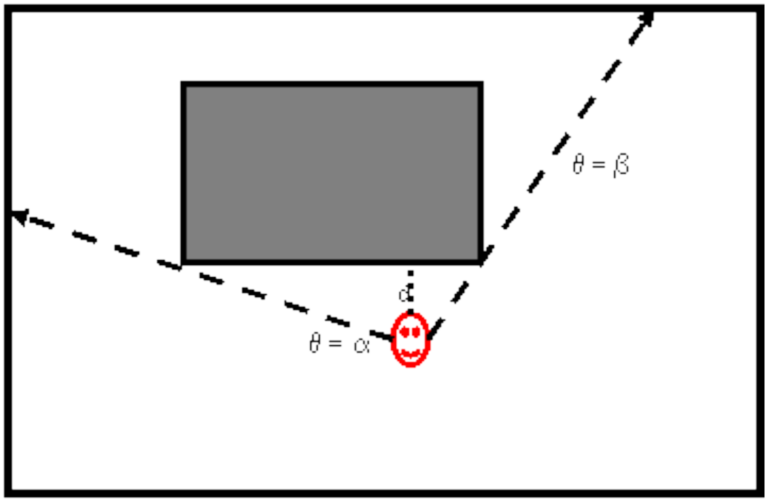

Figure 1: Visibility around a single rectangular obstacle.

STRATEGIES FOR THE DECEPTIVE AGENT

Deceptive agents D have several heuristic strategies to minimize E over the accomplishment of some required task such as planting a bomb:

- Suspiciousness-minimizing strategy: Try to reduce s(t) by camouflage of D's appearance and behavior (for instance for a terrorist, try to look, dress, and act like a normal citizen engaged in normal activities). However, deceptive agents must necessarily try to achieve goals they must conceal (like planting bombs) that are intrinsically incompatible with minimal suspiciousness.
- Visibility-minimizing strategy: Try to reduce v(i,t) by hiding as much as possible from other agents (for instance, have the deceptive agent spend time waiting in alleys). However, it is hard to make invisibility complete, and an agent observed trying to be invisible (like by hiding in a bush) is very suspicious, more so than an agent that is accidentally suspicious.
- Fast-execution strategy: Have D do suspicious actions quickly so v(i,t) is large over only a small time interval (for instance, run in and throw a bomb, then run out). This minimizes the period of time in which s(t)v(i,t) is nonnegligible, decreasing the value of the integral for E. The weakness of this is that fast activities are intrinsically suspicious, and increased s(t) function may counteract the benefits of the reduced time interval.
- Loitering strategy: Delay suspicious activities to time periods of low total visibility from other agents (for instance, wait or pace back and forth until there aren't many people nearby, then plant a bomb). This makes s(t) only large when v(i,t) is small, thereby keeping their product low. But loitering itself is suspicious to some degree: Suspiciousness is also a function of the motion vector, so an abnormally low speed (as in stopping and waiting) or abrupt turns (as in pacing) are suspicious too.
- Distributed-suspiciousness strategy: Distribute the suspiciousness over several deceptive agents (for instance, have one agent bring a bomb and another agent install it). This works when s(t) is a nonlinear concave function of the number of suspicious clues offered, something true of the popular "Naive Bayes" formula which multiplies factors from each clue. But suspiciousness distribution requires difficult coordination problems since communications cannot call attention to the deception. Also the degree of gain by distributing the suspiciousness may be insufficient; for instance, both the agent that brings the bomb and the agent that installs it must visit the same suspicious low-visibility location.
- Diversion strategy: Use one or more deceptive agents to create a diversion attracting the attention of the non-deceptive agents while a primary deceptive agent D engages in especially suspicious activity (for instance, have another deceptive agent create a loud noise while D plants a bomb). This has the effect of decreasing n(t) while s(t) is large, thereby decreasing E. This requires coordination by the deceptive agents, and also will not work repeatedly as the non-deceptive agents will realize they are being manipulated.

Figure 2 shows example two-dimensional terrain with three buildings (shaded) as obstacles. Assume an explosion occurs in the alley. Assume corridors of frequent travel indicated by the thick dotted lines, as determined by statistics. Solid-line path A exemplifies a suspiciousness-minimizing strategy (or fast-execution if done quickly); small-dotted-line path B exemplifies a loitering strategy; and C exemplifies a visibility-minimizing strategy. Both A and B could be used together in a distributed-suspiciousness strategy. Path D is an example nonsuspicious path.

As an example, Table 1 estimates E for each of the four paths in Figure 2. Each row gives the product of suspiciousness, visibility, and noticeability for that time interval

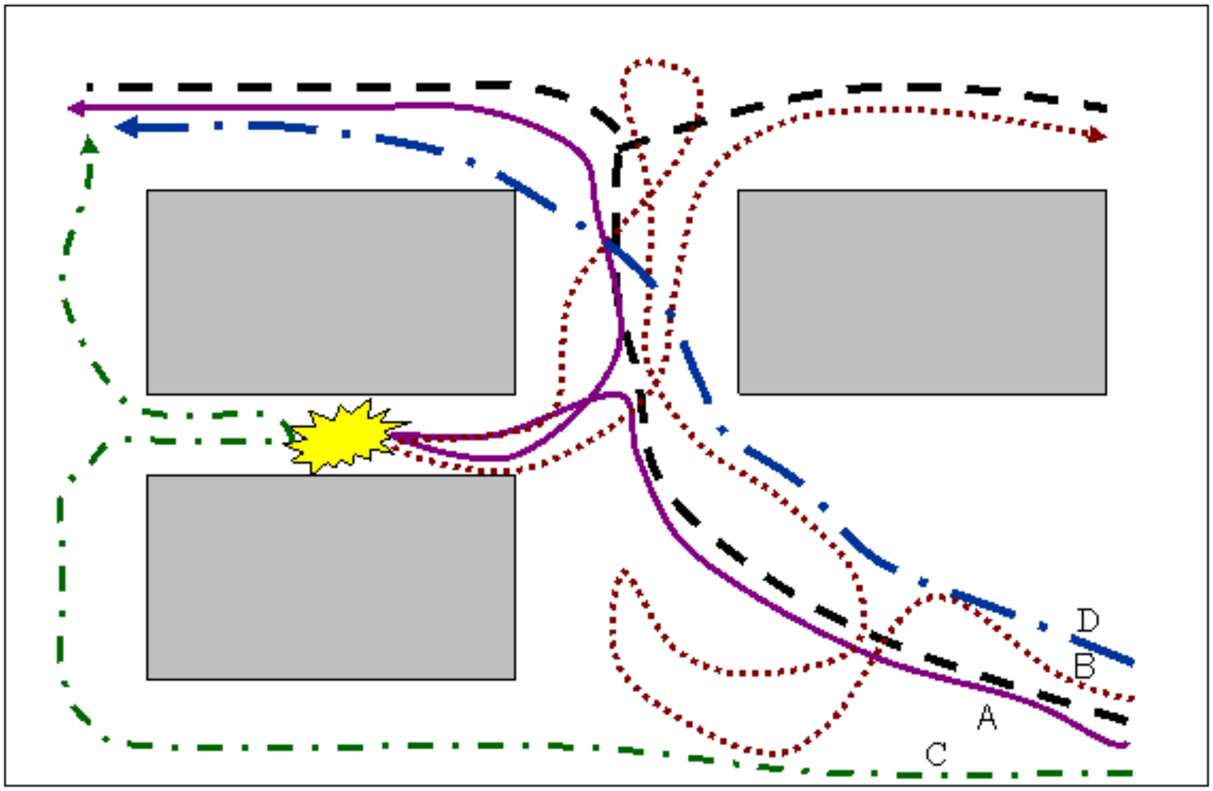

Figure 2: Example paths of deceptive (A, B, and C) and nondeceptive (D) agents in terrain with obstacles (shaded areas) and travel corridors (thick dotted line).

and the last row adds the columns to estimate E. Parameters were estimated by intuition. Time is given relative to the start of the path, so the paths are not expected to start simultaneously. Visibility by other agents was assumed constant. With these assumptions, the suspicious-minimizing strategy seems the best choice for a deceptive agent.

Table 1: Example calculations of the suspiciousness-visibility-noticeability product for Figure 2.

Time	Path A	Path B	Path C	Path D
0	.0	.0	.0	.0
1	.0	.2	.1	.0
2	.0	.2	.2	.0
3	.0	.0	.3	.0
4	.1	.0	.0	.0
5	.2	.2	.2	.0
6	.1	.0	.0	.0
7	.0	.2	.0	.0
8	.0	.1	.0	.0
9	.0	.0	.0	.0
10	.0	.0	.0	.0
Total	.4	.9	.8	.0

SECOND-ORDER DECEPTION

One problem with the clues so far discussed is their locality. A clever enemy may realize this and exploit it with "second-order" deceptions involving relationships between first-order deceptions, as with the distributed-suspiciousness and diversion strategies. For instance, they may create a disturbance in one part of the sensor field, then a similar disturbance in another part of the sensor field, and so on. This can have the effect of forcing the network to focus its resources in one location, then suddenly be required to shift resources to another, in a continual game of "catching up" which is hard to play adequately. Military sensor networks can be prone to this kind of manipulation since decoys and surprises are key elements in military tactics.

To represent this problem sufficiently for analysis, we need a model of processor resource allocation in sensor networks since the problem is due to processing-load redistribution strategies. Many algorithms have been proposed [18, 19]. For an architecture with many microsensors reporting to collector sensors, the collectors should do the work of tracking, signal analysis, and other forms of analysis. Then a common resource distribution algorithm is "offloading": Overloaded collectors broadcast requests for help to other neighboring collectors; those collectors that have excess processing capacity make offers to take some of the load; the load is then sent to them. This requires that the overloaded processor calculate an additional load for the effort of partitioning its load and sending it off. Common refinements include prenegotiating offloading paths in advance to save communications effort at critical times, and permitting offloaded loads to be further offloaded to other still more remote sites, and so on.

Deceptive adversaries can attack an offloading algorithm by creating high overloads or high rates of change of the overloads. With a load high over a sustained period of time, offloading cannot be done sufficiently quickly to permit all necessary processing. But there is a good solution, for processors to shift to processing aggregates instead of raw data. For instance, rather than reporting the time of every footstep, a microsensor can report the number of footsteps per minute for each minute. If this is done optimally, this results in an encoding that requires the logarithm of the number of original bits, but the optimum cannot be approached very closely in practice. Detection of deception can then look for abnormally large sensor events in abnormally large numbers.

Deceptive adversaries can also attack an offloading algorithm by varying the amount of sensation input considerably, say by creating disturbances at widely scattered places in the sensor field at closely spaced times. So sensor networks need some delay in adjusting to changed circumstances so they do not oscillate in the presence of varied input. How much to delay needs to be determined by experience.

CONCLUSIONS

Detection of deception has focused on the inadvertent physical clues that deceivers provide through lesser communication channels like body posture and gestures, or verbal clues when available. But even with only positional information about the center of mass of an agent, as with networks of primitive sensors, important clues are still available for detecting deception in the observed changes in goals. These clues will be important in such applications as detecting terrorism and criminal activity.

REFERENCES

[1] Callaway, E., 2004, *Wireless Sensor Networks: Architectures and Protocols*, Auerbach Publications, Boca Raton, Florida.

[2] Horton, M., Broad, A., Grimmer, M., Pisler, K., Sastry, S., Rosenberg, J., and Whitaker, N., 2002, "Deployment Ready Multinode Micropower Wireless Sensor Network for Identification, Classification, and Tracking," *SPIE Vol. 4708, Sensors and Command, Control, Communications, and Intelligence technologies for homeland defense and law enforcement*, pp. 290-295.

[3] Hynes, S., and Rowe, N., 2004, "Multi-Agent Simulation for Assessing Massive Sensor Deployment," Journal of Battlefield Technology, **7**, 2, pp. 23-26.

[4] Shin, J., Guibas, L., and Zhao, F., 2003, "A Distributed Algorithm for Managing Multi-Target Identities in Wireless Ad-Hoc Sensor Networks," *Lecture Notes in Computer Science, Vol. 2634*, Springer-Verlag, Heidelberg, Germany, pp. 223-238.

[5] Ford, C. V., 1996, *Lies! Lies!! Lies!!! The Psychology of Deceit,* American Psychiatric Press, Washington, DC.

[6] Decaire, M., 2000, "The Detection of Deception via Non-Verbal Deception Clues," retrieved January, 2005 from www.uplink.com.au/lawlibrary/Documents/Docs/Doc64.html.

[7] Rowe, N., 2005, "Detecting Online Deception and Responding to It," in the *Encyclopedia of Virtual Communities and Technologies*, S. Dasgupta, ed., The Idea Group, Hershey, Pennsylvania.

[8] Dunnigan, J. F., and Nofi, A. A., 2001, *Victory and Deceit, Second Edition: Deception and Trickery in War*, Writers Club Press, San Jose, California.

[9] Calfee, S., and Rowe, N., 2004, "Multi-Agent Simulation of Human Behavior in Naval Air Defense," Naval Engineers Journal, **116,** 4, pp. 53-64.

[10] Liebhaber, M. J., and Smith, P., 2000, "Naval Air Defense Threat Assessment: Cognitive Factors and Model," *Command and Control Research and Technology Symposium,* Monterey, California.

[11] Gibbins, D., Newsam, G., and Brooks, M., 1996, "Detecting Suspicious Background Changes in Video Surveillance of Busy Scenes," *Proc. 3rd IEEE Workshop on Applications of Computer Vision*, Sarasota, Florida, pp. 22-26.

[12] Wu, G., Wu, Y., Jiao, L., Wang, Y.-F., and Chang, E., 2003, "Multi-Camera Spatio-Temporal Fusion and Biased Sequence-Data Learning for Security Surveillance," *Proc. 11th ACM Intl. Conf. on Multimedia*, Berkeley, California, pp. 528-538.

[13] Heuer, R. J., 1982, "Cognitive Factors in Deception and Counterdeception," in *Strategic Military Deception*, D. Daniel and K. Herbig, eds., Pergamon, New York, pp. 31-69.

[14] Whaley, B., and Busby, J., 2002, "Detecting Deception: Practice, Practitioners, and Theory," in S*trategic Denial and Deception,* R. Godson and J. Wirtz, eds., Transaction Publishers, New Brunswick, New Jersey, pp. 181-221.

[15] Chu, M., Cheung, P., and Reich, J., 2004, "Distributed Attention," *Proc. 2nd Intl. Conf. on Embedded Networked Sensor Systems*, Baltimore, MD, p. 313.

[16] Coffman, T., Greenblatt, S., and Markus, S., 2004, "Graph-Based Technologies for Intelligence Analysis," Communications of the ACM, **47**, 3, pp. 45-47.

[17] Proctor, P. E., 2001, *Practical intrusion detection handbook*, Prentice-Hall PTR, Upper Saddle River, New Jersey.

[18] Gupta, G., and Younis, M., 2003, "Load-Balanced Clustering of Wireless Sensor Networks," *Proc. IEEE Intl. Conf. on Communications*, Baltimore, Maryland, Vol. 3, pp. 1848-1852.

[19] Haenggi, M., 2003, "Energy-Balancing Strategies for Wireless Sensor Networks," *Proc. Intl. Symposium on Circuits and Systems*, Bangkok, Thailand, Vol. 4, pp. 828-831.

ACKNOWLEDGEMENTS

This work was supported by the Chief of Naval Operations, U.S. Navy, N61F22.