Detecting Suspicious Behavior from Positional Information

Neil C. Rowe
MOVES Institute, U.S. Naval Postgraduate School
Code CS/Rp, 833 Dyer Road, Monterey CA 93943 USA
ncrowe@nps.edu

Abstract

Suspiciousness is not the same an anomalousness. Suspicion requires evidence of deception in observed attempts at concealment. We propose metrics for measuring suspiciousness of agents moving in a sensor field based on only periodic knowledge of their positions (as with large numbers of "small and cheap" sensors). This has applications to electronic sentries and counterterrorism. This theory requires assessment of the behavior, visibility, and noticeability of the average agent as well as the anomalousness of the position, velocity, and acceleration vectors of a particular agent. We conclude with a report on experiments with an implementation of our theory on a simulated sensor network.

Content areas

Sensor networks, deception, detection, kinematics, visibility, anomalies.

1 Introduction

Wireless sensor networks [Callaway, 2004] are increasingly used for surveillance for "suspicious" behavior, as with the currently-popular topic in the United States of "homeland security". For instance, we would like to detect thieves wandering about looking for theft opportunities or terrorists planting bombs [Hackwood and Potter, 1999]. Suspicious behavior is rare and monitoring for it is tedious; witness the difficulties of detecting it in casinos [Powell *et al*, 2003], one of the most active applications of surveillance. So it would be desirable to automatically detect suspicious behavior.

We will assume a persistent wireless sensor network with large numbers of small, inexpensive, limited-capability sensors in just one sensing modality (" microsensors"). We assume sensors without cameras that can only estimate distances and/or directions to moving agents; cameras are more expensive, and image processing to automatically find suspicious behavior is difficult and problematic in public areas. Microsensors can monitor sound levels, metal objects passing nearby, infrared body heat, or electromagnetic emissions. For processing, we will assume a smaller number of "collector" sensors with strong antennas and more power that aggregate the data from the microsensors, a common design approach today [Horton *et al*, 2002]. Such large numbers of microsensors must be dispersed by semi-random methods such as dropping them from the air [Hynes and Rowe, 2004]. An initialization phase can localize them by, for instance, sending a noisy device on a fixed path through the sensor field and analyzing the pattern of reports with least-squares fitting.

Collector sensor can localize moving agents using triangulation or other types of fitting. From this they can estimate velocity and acceleration, and (unambiguously except in very busy environments) track the agent. With sufficient redundancy, tracking can be done in the presence of occlusions of microsensors using best-fit methods [Shin *et al*, 2003].

A question is how much higher-level understanding of agent behavior we can obtain from their positions, velocities, and accelerations. We believe that much inference is possible, since gross motions reflect plans and deceptive agents execute two or more inconsistent plans. Then more sophisticated sensors or people could be invoked to study the phenomena more closely.

2 Previous work

Previous work on detecting suspicious behavior has focused heavily on the detection of anomalous behavior. For instance, [Wu *et al,* 2003] looked for suspicious behavior in a parking lot such as circling, zigzagging, and back-and-forth motions. However, there are many reasons that behavior could be anomalous without being suspicious, like surveys, repair work, and waiting to meet someone. If we label such activities as suspicious and take measures against them, we may be engaging in discriminatory or illegal behavior.

We postulate that suspicious behavior shows deliberate deception with concealment. Deception is an important social phenomenon with many applications in law, business, military operations, psychology, and entertainment, and a number of nonverbal clues can be used to detect it [Decaire, 2000; Qin *et al*, 2004], including:

- 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 have analogies. 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 and acceleration of the agent. So a sensor network can look for such clues.

Deceptive behavior also has higher-level clues in the form of "discrepancies" from normal behavior [Heuer, 1982]. So a sensor network could keep baseline statistics and note when behavior of an agent exceeds them. For instance, an agent with high acceleration magnitudes around other agents may be a thief engaged in theft.

Deceptive activities are common in warfare. A classic problem for the recognition of suspicious behavior with primarily positional data is air defense, the process of analyzing aircraft motions, obtained from radar, to discover aircraft which pose potential threats to a site [Liebhaber and Smith, 2000]. An enemy aircraft that wants to threaten will try to conceal their intentions, and will not often provide electronic identification or other obvious clues. But it is difficult to conceal locations and we can track them over time.

Detection of suspicious behavior is done routinely with intrusion-detection systems for computer networks [Proctor, 2001] where the issue of false alarms is also important. Most methods use strong clues such as attack signatures, but some measure discrepancies from baselines.

3 Assessing suspicious behavior

We will address detecting suspicious behavior in a two-dimensional sensor field, especially for urban or indoor terrain. Video imagery [Gibbins, Newsam, and Brooks, 1996; Wu *et al*, 2003] provides a variety of clues, but we will confine ourselves here to the positional state vectors only: positions, velocities, and accelerations. Such information also often is all that is available from military intelligence gathering from encrypted electronic transmissions [van Meter, 2002]. Atypical values for positions, velocities, and accelerations are suspicious, and are worthy of additional study or "attention" as in [Chu *et al*, 2004]. [Shao *et al,* 2000] defined suspicious behavior as that having high velocities, but unusual locations and unusual accelerations are equally suspicious when velocities are normal. Accelerations are important since F=ma relates them to forces which usually reflect volition on the part of agents.

3.1 A general formula

We will follow Occam's Razor and propose the simplest possible theory that covers the major observed phenomena. If we choose our key factors carefully to be independent, Naive Bayes inference should be a reasonable assumption [Korb and Nicholson, 2004]. We will use the odds form:

$$o(S \mid (A \,\&\, D)) = \frac{o(S \mid A)}{o(S)} \frac{o(S \mid D)}{o(S)} o(S)$$

where o represents odds (so o(X) = p(X)/(1-p(X)) where p(X) is the probability of X), S is the condition of suspiciousness of an agent, A is the condition of anomalousness, and D is the condition of deceptiveness. The odds of A and D will depend on additional factors. Statistics can be collected for a sensor field to estimate these odds.

We could use this formula at every time interval, but this will magnify short-term perturbations if the time intervals are closely spaced. So we will compensate with:

$$o(S \mid (A \,\&\, D)) = [\frac{o(S \mid A)}{o(S)} \frac{o(S \mid D)}{o(S)}]^{1/M} o(S)$$

for some number of timesteps M, the "learning time". This slows down the rate by which o(S) is changed for M>1 since a fractional power of a positive number is always closer to 1 than the original number.

Note that odds can differ for different times of day and times of the week. For instance, if the sensor field is an commercial area, we would expect to see more activity during the daytime on weekdays.

3.2 Measuring anomalies in state vectors

Figure 1 shows an example sensor field with unidirectional arrows representing the mean of the acceleration vectors in that vicinity, bidirectional arrows representing the mean of the velocity vectors, and circles representing visit frequency to that area of the terrain. These statistics can be obtained by dividing the area into a uniform rectangular grid and counting observations for each grid cell.

Then we can estimate the anomalousness of the state vector of an agent at time t that is inside grid cell (i,j) as a weighted average of the atypicality of location, velocity, and acceleration as r(t):

$\frac{c_1 g}{p(i,j)} + c_2 \mid v(i,j) - v(t) \mid + c_3 \mid a(t) \mid$ where t is the time a particular agent is observed, c_1, c_2, and c_3 are constants set by experiments, g is the "traffic rate" or average frequency of an agent in a random grid cell, p(i,j) is the average frequency of an agent in grid cell (i,j), v(i,t) is the average velocity of agents, v(t) is the observed velocity of the particular agent, a(i,t) is the average acceleration of agents, a(t) is the observed acceleration of the particular agent, and the bars mean the norm. Velocity and acceleration are handled differently because the least

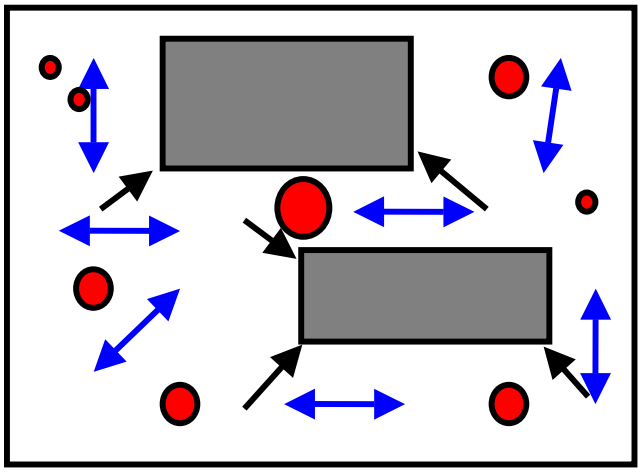

Figure 1: Example mean acceleration vectors (unidirectional arrows), velocity vectors (bidirectional arrows), and positional frequency (circles) for a simple sensor field.

anomalous velocity is that of previous visitors to that grid cell, whereas the least anomalous acceleration is zero.

The idea is that the sum of three factors will be more likely to be a normally-distributed random variable, by the Central Limit Theorem of probability theory, and be a more reliable indicator of an anomaly. Note that this formula will flag a variety of suspicious behavior: hiding (suggesting attempts to surprise), too-fast motion (suggesting attacks), too-slow motion (suggesting loitering), speed changes (suggesting avoidance), and sudden lateral accelerations (suggesting surprise).The directions of velocity vectors are as important as their magnitudes, in order to recognize turns. However, many observed velocity-direction distributions will be bimodal with peaks 180 degrees apart, indicating bidirectional corridors of travel; we eliminate this problem by doubling the directional angles of our observed velocity vectors before averaging them, then halving the average.

3.3 Excuses for anomalies

Not all anomalous state vectors are equally suspicious. One issue is that there may be excuses for anomalies, nearby events that could have influenced them. Such events should set r(t) to zero. An excuse can be a large grid-wide event like a loud noise. It can also be collision avoidance, defined as situations in which the projected distance between two agents, assuming their current velocity vectors, is less than a minimum safety radius. The safety radius will vary with the danger of the agents, so it will be larger for a vehicle agent than a human agent.

An important class of excuses relate to social activities such as meeting someone (which can happen with vehicle agents as well as people). Conversation events can be detected when two or agents converge, decelerate, and stop at a typical conversation distance from one another with visibility between participants. The typical distance is closer under more crowded conditions. The duration of a conversation event is important, and is usually at least 10 seconds; shorter periods suggest theft instead. Note that stopping is generally important, although a more sophisticated theory could allow for conversations between agents moving side by side, but not where one is following the other.

Apparent conversations can either decrease or increase suspicion, depending on how suspicious we are of the participants individually. Generally speaking, the difference in the suspiciousness of the participants should decrease after a conversation. So if an unsuspicious agent meets a suspicious agent, that should increase the former's suspiciousness. Relaxation methods can be used to iteratively improve estimates when many social interactions are observed. An important part of antiterrorism intelligence analysis is finding connections from known suspicious people to others [Coffman *et al,* 2004].

3.4 Visibility

On the other hand, events associated with visibility changes should be additionally suspicious. For instance, when an agent doubles back on its path after another agent disappears from view, that is suspicious because it suggests the agent has things to do that it does not want seen. Visibility estimation must take into account occlusion by obstacles, terrain, and other agents. It is often difficult for people to do this because they have eyes only on one side of their heads, and it is hard at a distance to tell whether someone is watching you. Nonetheless, we assume that deceptive agents can estimate their visibility based on current known locations of other agents and a general knowledge of the degree to which they are likely to be visible in an area.

Visibility decreases monotonically with distance between agents. Suspiciousness assessment depends mostly on linear resolution which decreases inversely with the distance. However, there is a minimum distance (since an observer cannot observe well if they are too close) and a maximum distance (since faint signals fall below the threshold of noticeability). We need to estimate the average unoccluded distance within these limits in every direction from a point, as illustrated in Figure 2. With gridded data, we can estimate average visibility of an agent at location (x(t),y(t)) to every other point as w(t,i,j):

$$\frac{\sum_i \sum_j p(i,j)b(t,i,j)/\sqrt{(x(t)-i)^2+(y(t)-j)^2}}{\sum_i \sum_j p(i,j)/\sqrt{(x(t)-i)^2+(y(t)-j)^2}}$$

where t is the time, p(i,j) is the probability of an agent being in cell (i,j), and b(t,i,j) is the probability that the line between the agent at (x(t),y(t)) and grid cell (i,j) is unoc-

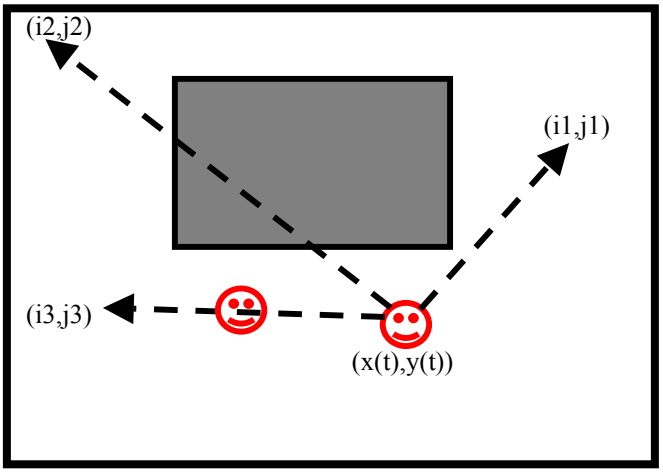

Figure 2: Average visibility calculation: b(t,i1,j1)=1, b(t,i2,j2)=0, and b(t,i3,j3)=0.

cluded by an obstacle or another agent. So w(t,i,j)=1 in an unobstructed area.

It is valuable to distinguish visibility from noticeability n(t) or lack of "cover". Given two people at an equal distance, the one in the larger crowd will appear less noticeable. We propose $n(t) = \frac{c_4}{c_5 + p(x(t), y(t))}$.

3.5 Mapping anomalousness and visibility to odds

Given r(t), w(t,i,j), and n(t), we must calculate the conditional odds of suspicious behavior. Both of these mappings are subjective and will vary with context. But both are monotonic since greater anomalousness, lesser visibility, and lesser noticeability all mean greater suspicion. Following Occam's Razor, we can use

$$\frac{o(S \mid A)}{o(S)} = c_6 r(t) \text{ and } \frac{o(S \mid D)}{o(S)} = \frac{c_7}{v(t)n(t)}.$$

This gives overall:

$$o(S \mid (A \,\&\, D)) = [\frac{c * r(t)}{w(t,i,j)n(t)}]^{1/M} o(S)$$

This gives a way to update the suspiciousness estimate at each timestep.

This formula assumes that our sensor network is not known to the agents, at least in all of its detail. Thus the formula is not the degree of suspicion that agents moving in the sensor field have. For them, invisibility and unnoticeability have the opposite effect because a normal agent will not expect suspicious behavior and will only notice it when it is quite visible and noticeable. We thus propose for reasoning by nonsuspicious agents in the sensor field:

$$o(S_n \mid (A \,\&\, D)) = [c_f r(t) w(t,i,j) n(t)]^{1/M} o(S_n)$$

4 Experiments

To test our ideas, we created a program to estimate suspiciousness of behavior. We use a rectangular grid to represent terrain. Figure 3 shows an example of an indoors room like an office; shaded areas represent untraversable and view-blocking obstacles, and thickened lines represent entry doors.

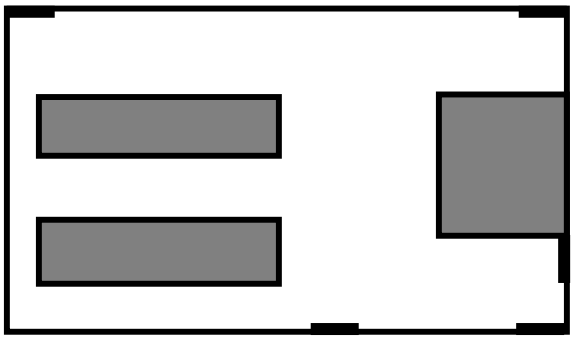

Figure 3: Example indoor terrain (a room).

Lacking statistics, we first estimated frequency of agents within each grid cell by finding all possible straight unblocked paths between doors and/or corners of obstacles. Corners were used because optimal paths in regions of uniform cost with obstacles are found by searching the graph of the vertices of the obstacles and possible start points. The resulting relative frequencies of visits are shown in Figure 4 for a 10 by 12 approximation of the room of Figure 3, where "00" indicates an obstacle. This incorporates a baseline visit rate of 5 (representing nontraversal behaviors such as accessing files in an office, assumed to be evenly apportioned over the room) and some averaging of frequencies into their neighbors (modeling the tendency of agents to wander a bit from the optimal paths). These numbers are on a scale of 0 (never visited) to 1000 (visited all the time); values were normalized so they would average to a specified traffic rate of g=0.01 (10 on the 0-1000 scale), the probability that some agent would be observed in a randomly chosen grid cell at a random time.

```
19 13 06 06 06 05 05 05 06 07 10 11
20 08 08 09 09 08 13 07 13 07 06 06
23 14 09 08 08 15 24 21 20 10 03 02
18 00 00 00 00 00 16 10 17 00 00 00
19 00 00 00 00 00 16 09 13 00 00 00
28 15 08 07 07 15 28 16 13 00 00 00
24 13 09 10 10 15 29 18 19 00 00 00
14 00 00 00 00 00 17 19 29 17 14 17
12 00 00 00 00 00 17 19 29 17 14 17
14 11 08 08 08 12 18 23 30 21 17 18
```

Figure 4: Estimated visit frequencies for the room.

Then we compute the average visibilities of all points in the room using our above formula. This calculation includes a penalty representing that the line of sight is more likely to be occluded in high-traffic areas. Results for Figure 3 are shown in Figure 5, with a scale of 0 (invisible) to 100 (perfect visibility).

```
62 58 59 62 69 77 80 81 80 64 56 52
64 58 56 57 61 74 81 83 81 61 56 52
60 51 46 47 48 66 82 82 79 54 45 41
44  0  0  0  0  0 74 74 73  0  0  0
47  0  0  0  0  0 77 80 72  0  0  0
63 53 54 63 64 74 84 85 76  0  0  0
59 52 52 50 58 78 84 87 84  0  0  0
39  0  0  0  0  0 77 80 81 65 67 65
36  0  0  0  0  0 78 81 81 72 70 68
53 38 39 45 44 57 79 82 81 74 69 65
```

Figure 5: Calculated visibility for each point.

Finally, suppose we have four agents visiting this room as indicated by the numbers in Figure 6. Each number represents an observation of an agent at that time at the location. Agent locations were matched between timesteps by projecting observed velocities and finding the closest match, and the approximate inferred paths are shown as thin lines; newly appearing agents were assumed to come from the nearest door.

We then applied the suspicion formula to track the cumulative suspiciousness of each agent's path. Noticeability was not used because there were so few agents. We added to the a priori visibility of the agent w(t,i,j) a weighting of the specific visibility by other agents u(t,i,j), to suggest that deceptive agents only know half the time who can see them; w(t,i,j) is the average visibility of an agent at cell (i,j) at time t, while u(t,i,j) is the number of agents that can see (i,j) at time t. The "learning time" M was set to 2, and we used only the difference in the angle of twice the velocity vectors because the velocity magnitude was not meaningful with our method of estimating visit frequencies discussed above. Reasonable weights for r(t) were obtained by experiment. In total, we used the updating formula:

$$o(S \mid (A \,\&\, D)) = \left[\frac{3 * g * r(t)}{w(t,i,j) + g * u(t,i,j)}\right]^{0.5} o(S),$$

$$r(t) = \frac{g}{p(i,j)} + 0.3(1 - (2v(t) \circ 2v(i,j))) + 2 \mid a(t) \mid$$

where "$\circ$" denotes the inner product of vectors.

Our program assessed an average suspicion of 0.34 for the agent going south to northwest, 0.91 for the agent going east to west to northeast, 0.25 for the agent moving in the northwest corner, and 0.25 for the agent moving only in the northeast corner. These are reasonable because the first agent is unsuspicious, the others make significant direction changes, the second agent spends time in the "alley" with lower visibility, and the third and fourth agents apparently meet at times 6 and 7 (thus excusing their nonzero accelerations during the meeting). We would need to conduct experiments with human subjects to compute a threshold value to decide at what level of suspicion we should take action; for instance, we could observe a public area where subjects would be directed to simulate tasks such as planting a bomb.

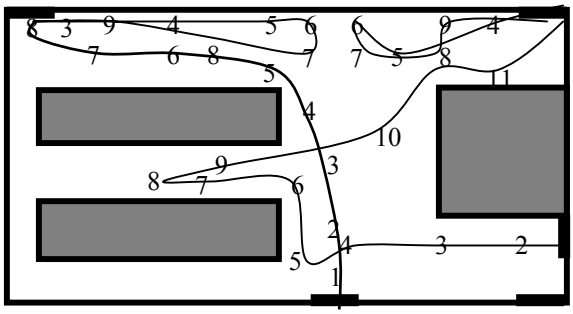

Figure 6: Some agent paths in the room.

5 Strategies for the deceptive agent

The theory we have developed also predicts heuristic strategies for a deceptive agent to minimize o(S) in some task such as stealing a wallet or planting a bomb:

- Suspiciousness-minimizing strategy: Reduce r(t) by following well-traveled routes with minimal changes of direction. However, deceptive agents must necessarily try to achieve at least some goals incompatible with this.
- Visibility-minimizing strategy: Reduce w(t,i,j) by spending much time in low-visibility parts of the field (e.g., loiter in alleys). However, it is hard to anticipate all viewing angles, and an agent observed trying to be invisible is additionally suspicious.
- Fast-execution strategy: Do suspicious actions quickly so the timesteps for which r(t) is a nonnegligible are few (for instance, run in and throw a bomb, then run out). But fast activities are intrinsically anomalous and suspicious.
- Loitering-timing strategy: Delay suspicious activities to periods of low visibility from other agents (for instance, pace back and forth until there aren't people around, then plant a bomb). This makes r(t) only large when w(t,i,j) is small, thereby keeping their product low. But loitering itself is suspicious since it requires acceleration changes and velocity reversals.
- 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 r(t) is a concave function of the number of suspicious clues offered. This requires difficult coordination, and a concave function may be impossible. For instance, both an agent that brings a bomb and an agent that installs it must visit the same suspicious low-visibility location.

- Diversion strategy: Create a diversion during suspicious activities, reducing n(t) then. This also requires coordination by the deceptive agents, and will not work repeatedly as the non-deceptive agents will eventually realize they are being manipulated.

6 Conclusions

We have proposed a new approach to detecting suspicious behavior in areas under low-level surveillance. If one can afford many cameras and real-time image-processing software, much such surveillance can be automated. But this is an expensive solution, only appropriate for high-security areas. For monitoring for theft and terrorist activity in broad areas, simple footfall or motion sensors may be all that are feasible. We have showed that a good deal can be learned from even such simple sensors, by noticing anomalous location, velocity, and acceleration data and correlating it with visibility and noticeability by other agents. Our next step will be to test our ideas with real-world experiments.

References

[Callaway, 2004] E. Callaway. *Wireless sensor networks: architectures and protocols*. Boca Raton, FL: Auerbach Publications, 2004.

[Chu *et al*, 2004] M. Chu, P. Cheung, and J. Reich. Distributed attention. Proc. 2nd Intl. Conf. on Embedded Networked Sensor Systems, Baltimore, MD, 313, 2004.

[Coffman *et al*, 2004] T. Coffman, S. Greenblatt, and S. Markus. Graph-based technologies for intelligence analysis. *Communications of the ACM, 47* (3), 45-47, March 2004.

[Decaire, 2000] M. Decaire. The detection of deception via non-verbal deception clues. www.uplink.com.au/lawlibrary/Documents/Docs/Doc64.html, 2000.

[Gibbins *et al*, 1996] D. Gibbins, G. Newsam, and M. Brooks, M.. Detecting suspicious background changes in video surveillance of busy scenes. Proc. 3rd IEEE Workshop on Applications of Computer Vision, 22-26, December 1996.

[Hackwood and Potter, 1999] S. Hackwood and P. A. Potter, Signal and image processing for crime control and crime prevention. Proc. Intl. Conf. on Image Processing, Kobe, Japan, October 1999, Vol. 3, 513-517.

[Heuer, 1982] R. J. Heuer. Cognitive factors in deception and counterdeception. In *Strategic Military Deception*, ed. Daniel, D. C., and Herbig, K. L. (New York: Pergamon), 31-69, 1982.

[Horton *et al*, 2002] M. Horton, A. Broad, M. Grimmer, K. Pisler, S. Sastry, J. Rosenberg, and N. Whitaker. 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, 290-295, 2002.

[Hynes and Rowe, 2004] Hynes, S., and Rowe, N., Multi-agent simulation for assessing massive sensor deployment. *Journal of Battlefield Technology*, 7(2), 23-36, July 2004.

[Korb and Nicholson, 2004] K. Korb and A. Nicholson, *Bayesian artificial intelligence*. Boca Raton, FL: Chapman and Hall/CRC, 2004.

[Liebhaber and Smith, 2000] Liebhaber, M. J., and Smith, P., (2000, June). Naval air defense threat assessment: cognitive factors and model. Command and Control Research and Technology Symposium, Monterey, CA.

[van Meter, 2002] K. van Meter, Terrorists/liberators: Researching and dealing with adversary social networks. *Connections*, vol. 24, no. 3, 66-78, 2002.

[Powell *et al*, 2003] G. L. Powell, L. A. Tyska, and L. J. Fennelly, *Casino surveillance and security: 150 things you should know*. New York: Asis International, 2003.

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

[Qin *et al*, 2004] T. Qin, J. Burgoon, and J. Nunamaker. An exploratory study of promising cues in deception detection and application of decision tree. Proc. 37th Hawaii Intl. Conf. On Systems Sciences, 2004.

[Shao *et al*, 2000] H. Shao, L. Li, P. Ziao, and M Leung, ELETVIEW: an active elevator video surveillance system. Proc. Workshop on Human Motion, Los Alamitos, CA, 67-72, December 2000.

[Shin *et al*, 2003] J. Shin, L. Guibas, and F. Zhao. A distributed algorithm for managing multi-target identities in wireless ad-hoc sensor networks. *Lecture Notes in Computer Science, vol. 2634* (Heidelberg, Germany: Springer-Verlag), 223-238, 2003.

[Wu *et al*, 2003] G. Wu, Y. Wu, L. Jiao, Y.-F. Wang, and E. Chang. Multi-camera spatio-temporal fusion and biased sequence-data learning for security surveillance. Proc. 11th ACM Intl. Conf. on Multimedia, Berkeley, CA, 528-538, 2003.

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