Title:

ADAPTIVE SECURITY SYSTEMS - COMBINING EXPERT
SYSTEMS WITH ADAPTIVE TECHNOLOGIES

Author(s):

Paul Argo
Rohan Loveland
Kirsten Anderson
Brian Kelley
Larry Ticknor
John Elling
James Howse
Kim Linder
Constance Buenafe
Bob Berglin
Kristin Adair
Chris Johnson
Susan Hruska

Submitted to:

1997 IEEE International Conference on Systems, Man, and
Cybernetics
October 12-15, 1997
Orlando, Florida

RECEIVED

AUG 27 1997

OSTI

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

MASTER

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

Los Alamos
NATIONAL LABORATORY

Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by the University of California for the U.S. Department of Energy under contract W-7405-ENG-36. By acceptance of this article, the publisher recognizes that the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. The Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy.

DISCLAIMER

**Portions of this document may be illegible
in electronic image products. Images are
produced from the best available original
document.**

Adaptive Security Systems - Combining Expert Systems with Adaptive Technologies

Paul Argo, Rohan Loveland, Kirsten Anderson, Brian Kelley,
Larry Ticknor, John Elling, and James Howse
Los Alamos National Laboratory
Los Alamos, New Mexico 87545
{pargo, loveland, kirsten, bkelley, lot, elling}@lanl.gov

Kim Linder, Constance Buenafe, and Bob Berglin
Allied Signal Federal Manufacturing and Technology
Albuquerque, New Mexico
{klinder, cbuenafe, bberglin}@kcp.com

Kristin Adair, Chris Johnson, and Susan Hruska
Department of Computer Science
Florida State University
Tallahassee, Florida 32305
{adair, cjohnson, hruska}@cs.fsu.edu

Abstract

The Adaptive Multisensor Integrated Security System (AMISS) uses a variety of computational intelligence techniques to reason from raw sensor data through an array of processing layers to arrive at an assessment for alarm/alert conditions based on human behavior within a secure facility. In this paper, we give an overview of the system and briefly describe some of the major components of the system. This system is currently under development and testing in a realistic facility setting.

1. Introduction

The Adaptive Multisensor Integrated Security System (AMISS) [3] is a system designed to provide security for nuclear facilities and other facilities for which safety and security are concerns and for which normal day-to-day activities change over time. Our system is different from the usual concept of a security system which is built from systems that provide access control and intrusion alarms. Rather we hope to operate in facilities that have significant ongoing activity, within which we must identify and place an alarm for

activities that might be “threatening” in that facility. In many cases this threatening activity will not have been specified beforehand.

The system combines expert systems, statistical evaluations, and adaptive processes (neural and fuzzy) at almost all levels of operation. Outputs of systems of similar sensors (e.g., Passive Infra-Red Motion detectors, or PIRs) are combined through a feed-forward neural network [6] that has been trained to localize the motion in a room.

From this straightforward case of sensor fusion we move to the process of combining similar information from different sensors. For example, our video processing system localizes the humans in the room, and this determination is combined with the PIR network to better define the locations and activities in the room.

In assessing the state of the room, an easy and obvious first evaluation is to compare the presence of people, either in the room or in specific parts of the room, with an authorization database. This job may be made easier if the system has also provided positive identification of the persons in the room.

At the level of most interest to this project, we compare the actual pattern of activity in the room both to a more complex authorization database, as well as a “learned” historical set of “usual” activi-

ties that the specific individual might undergo in the room. In this sense the system is used as an anomaly detector, as it looks for sufficient deviation from the usual activity pattern. Methods are also being constructed for characterizing the "interestingness" (or level of threat) of these activities.

The Integrated Monitoring system consists of three subsystems: data collection and fusion, analysis, and alarm/alert. An overview of the complete system is shown in Figure 1. This system is currently being tested in a realistic setting. The many vagaries of a room including cluttering equipment (which is frequently moved about), tracking more than one person, and human movements which are difficult to predict beforehand are providing an excellent training testbed for our system.

2. Data Collection and Fusion

The data collection subsystem consists of the various sensors and their communications with the central database. The primary parts of this subsystem are sensors that can track and identify personnel, and sensors that can track (or detect) radioactive sources. The fusion of information from a variety of sensors is a key feature of this subsystem.

Personnel tracking is done by a range of sensors, potentially including video camera, motion detection arrays, active radio-location systems, and active badge systems. Personnel recognition/verification involves a variety of biometric methods: face recognition, fingerprint identification, gait, voice, and/or iris recognition. It is our belief that having a variety of such systems will make the job of spoofing the system much harder. In addition, the use of active badges allows the system to do the much easier task of verifying an identity (rather than recognition).

Radioactive source tracking/monitoring is being performed by radiation sensor arrays, bar codes, and source characterization sensors. Part of the lower-level reasoning that takes place at the fusion level of this module involves attaching an identified person object to the source object, and tracking their combined movement around the room.

2.1. Video Tracking

Video cameras form one of the primary sensor systems of the AMISS project, due to the large amount of information that can be gained from them [5]. Cameras with wide fields of view can be used to survey large portions of the room simultaneously, allowing objects to be tracked as they move through the room.

Overlapping fields of view gives us the ability to perform fusion at the video identification level. The capability to zoom in on people objects once they are segmented out allows for face recognition. With sufficient processing, some activity classification is possible (e.g., Fred is kneeling in front of the vault).

The main thrust in the AMISS video processing to date has been in the isolation of people objects from the background, and in tracking them. The cameras provide one of the few reliable ways to obtain object position information. The goal of tracking is to determine the positions of all the objects moving in the room, in real-time, and to provide temporally consistent labeling, so that "Fred" remains recognized as "Fred" while he moves around the room. The problem is complicated by having a number of non-human objects in the room, many of which are large enough to occlude a person as he/she walks behind them, and some of which can be (and often are) moved around.

Our approach is to first build single camera tracking systems, optimizing their performance on large areas of the room, and then to fuse together the results from several camera systems in order to obtain more reliable labeling and position information.

The advantages of fusing several systems together are readily apparent. What is partially occluded by a file cabinet in one camera's view might be fully visible in the view of another camera. The apparent resolution due to perspective differences may be greater in one view than another, allowing more precise position estimation using more than one view. A person's face might be turned towards one camera, allowing face recognition, while only the back of the head may be seen in the other.

Using this approach, we are able to track multiple objects as they move around the room, obtaining fairly reliable labeling results in the face of occlusions and overlapping objects, and to accurately estimate the three-dimensional real world positions of those objects.

2.2. Other Sensor Systems

In addition to video, we use an array of PIR sensors to localize motion in the room. Individually, these PIR sensors give relatively low-level information on possible locations of moving person-sized objects. Taken as an array, in which the PIRs are installed throughout the room with a high degree of overlap in their fields of view, the PIR sensors can give fairly reliable readings localizing the sections of the room in which there is movement. These readings are used to confirm the position estimates given by the video system, and can also be used to focus the attention of

Figure 1. Overview of the AMISS system

the assessment system. For example, if movement is reported in a section of the room for which we have no other readings of identified objects, cameras can be focused on that area for further scrutiny, or the health of the other monitoring systems may be called into question for further checking.

Performing personnel identification, that is, attaching a name to the identified person-sized blob object, is done using a variety of off-the-shelf systems, including both active and passive badge readers, face recognition systems, and voice identification systems. These systems give independent readings as they are used, and the information is fused together and attached to the appropriate person object with an associated reading of reliability or certainty. The reliability reading is derived from fusing the individual reliabilities of the independent systems.

Tracking of radioactive sources is of interest in many of the types of facilities that we are targeting with this research. In our testbed system, we use an array of slab detectors which are used to triangulate the source in the room. Movement of the source in the room is tracked using position locations over time.

The inclusion of sensor subsystems in the AMISS system is by no means a static or completed process. As new sensor systems become available, or as we assess the need for better, different, or additional information, sensor systems can be added to the current sensor suite. The AMISS system design is modular, and is not tied to a particular sensor identifying and tracking a person or a radioactive source. Rather, these sensor system components can be switched in and out of the larger assessment system as needed. The underlying framework of taking in information from a variety of sensor sources, fusing it, and reasoning about it, is independent of any particular sensor.

2.3. Communication

In order to enhance the ability of the various sensor subsystems and to perform the fusion and reasoning, information must be readily passed between all sensor collection systems. In some situations, initial assessment of the safety and security of the facility can be made with an authorization database. In order to make this assessment, the sensor systems must be reporting to a central location in the facility management system.

If available for the facility and tests being performed, the authorization database can provide information on expected activities, e.g., how many people are expected to enter the nuclear material area, what the expected activity is, where the people will be, what nuclear material is to be handled, etc.

The AMISS data communications system is designed to ensure that all the needed information is stored in a central database, and to allow any particular sensor system to access any other information that it might need. Thus information can be passed across the system, as well as upward to the central database.

The database also serves as a repository for archived sensor data. These archives are called upon during some parts of the threat assessment process, when further information is needed by the reasoning system. For example, a falling level of reliability of personnel identification may trigger the need to have a guard look again at the last 10 seconds of video at the door entry-way, or to query the last 20 seconds of active badge readings from a particular location in the room.

The data collection design makes possible the implementation of the virtual sensor concept, which vastly simplifies sensor fusion.

2.4. Primitive and Virtual Sensors

One of the strengths of the AMISS project is that it is designed to allow the painless inclusion of *virtual* sensors into the system. Also called *smart sensors*, these systems are relatively independent of the physical sensor hardware, and rather emerge as a combination or post-processed version of potentially several sensor systems.

Many of the commercially-available sensors that we are using send out raw information that can be used for more than one purpose. For example, with radiation sensors it is possible to locate a radiation source in the room by processing the information one way, but it is also possible to identify various radioactive sources and distinguish them from one another, using different processing. The raw data from the "primitive" radiation sensors may well end up in the database, but one or more virtual sensors, providing information gained from processing that same raw data, will also appear giving readings in parallel with the primitive sensors.

Again, this concept lends itself to making data fusion much more modular than it would be otherwise, since a virtual sensor can take in any variety of raw data by asking for it. Virtual sensors also allow the easy inclusion of complex processing building blocks into the system, while making it seem as if the data is coming from just another sensor.

3. Analysis

3.1. Authorization Checking

The first line of defense against illegal or mistaken activity is to check ongoing activities against an authorization database, if one exists. Security personnel can enter information on the identities of individuals allowed to enter the secure area, the time they are allowed to enter, and the sources they are allowed to access. Then, as activities occur, such as the entrance of a new individual or the removal of a source, the activity can be checked against the authorized activities in the database.

Further security can be added by learning individuals' normal behaviors. This line of defense would be most appropriate for detecting when authorized individuals are behaving in a manner inconsistent with their usual behavior, perhaps indicating an insider threat. For example, this system might be expected to note that a specific individual had never handled a specific source before, as well as detecting that a specific source had been moved into a location in which it had never been before.

3.2. Reasoning

In the AMISS system, we learn patterns of behavior by clustering the patterns using an adaptive neural network [4]. The neural network is trained to learn the sequences of actions an individual makes as they move about the room, using sensor data stored in the database capturing the individual's normal behavior [1]. The sequence of actions being learned may also include information on the state of the room while the activity was occurring, for example, merging data on the location of the person with indications of the number of sources in the room, the number of other people in the room, and which doors are open.

After training on available data, the network can identify whether the current ongoing activity matches the system's learned pattern of behavior for that person or not. When the neural network triggers a false alarm, wrongly identifying an activity as abnormal, the network can add that pattern to its training set, and quickly learn that the misidentified behavior is normal.

The analysis subsystem uses both rule-based expert systems and adaptive technologies. Some rules for assessing the safety and security of a facility are known, and can be used directly in a rule-based system. Much of the expert knowledge on facilities monitoring, especially when it involves normal day-to-day human behaviors, is not immediately amenable to

rule-based processing. Our plans are to implement knowledge extraction methods which discover rules from data over time, and to enhance the existing rule-bases with that knowledge.

In areas where standard threat analysis provides specification of threatening activity patterns, the system can be directed to explicitly look for and detect those patterns. However, the real strength of the system lies in its own ability to learn those patterns over time, to detect in real time activity patterns that it has not seen before, and to bring these to the attention of the security analyst.

3.3. Materials Accounting and Source Tracking

Our plans are to implement additional security measures by using material accounting at the analysis level. This established process helps ensure that all nuclear material is in authorized locations. Analysis of longer term activities and sensor readings will help guard against more subtle threats such as repeated removals of small amounts of nuclear material.

3.4. Data Mining

Beyond the immediate assessment of behavior patterns developing in real time, our system plan includes a component for scrutinizing longer term patterns. We are interested in the ability of the system to capitalize on its rich repository of information in the central database, perhaps in detecting subtle changes in the facility operating patterns over time. Interactions between individuals, authorization patterns, and materials movement could also be gleaned from the data, even though not readily observable in real time.

Operational *data mining* will be used to find and categorize interesting patterns in the activity data. These patterns will then be presented to security analysts for further interpretation. Discovery of association rules [2] in the data relating personnel, sources, and the state of the room and its sensors is a later goal of this part of the project.

4. Alarm/Alert

The potential end users of our system include the facilities managers, guards, and security analysts. The interfaces for different users will provide information that is appropriate to their use of the system, but all will contain a basic component of the alarm/alert subsystem which operates in real time.

The alarm/alert subsystem provides for the interruption of unauthorized activities within the nuclear

facility. It may be used to enforce security directly, for example, by not opening a vault door for an unauthorized access attempt. At another level, it may be implemented to go to the extreme of calling automatically for actions by an armed force when the threat level is high enough. In most cases, the system will be providing real-time information to the security forces, with video images and annotations about detected activities. The intention of this type of use of the system is to provide automated assessment of the current state of the facility, including the people within it, and to provide the information on which that assessment was based.

One of our design criteria is that the system must degrade gracefully, as opposed to a system which would crash when overloaded with conflicting information. This means our system must recognize when the situation has become so complex (or ambiguous) that the system can no longer assess the appropriateness of the ongoing activity. At this "failure" point it must inform the security forces of both its failure and its best assessment of the situation.

In many cases, we expect that our system will be used to augment the observational power of the human guards/analysts, providing additional information for them to use in their assessment, and backup sensor data for that assessment. Our system may be used to focus attention on a particular section of the room during periods of assessed heightened threat potential. In this case, simply highlighting a portion of the video image to call the guard's attention to it would be sufficient. Information on sensor health (e.g., suspicion that a camera's lens has been obstructed) could also lead to alert or alarm conditions within the facility, again, highlighted for further human attention.

5. Discussion

This brief overview of the AMISS system is intended to show the relationships of the overall concepts of sensor fusion, analysis, and reasoning in our security system and how the various subsystems interact. One focus of this project is to bring to bear a variety of reasoning techniques, whether rule based, statistical, fuzzy, or neural network based, on this problem. The wide variety in the types of sensor data we can expect to reason with in a realistic facility with normal human activities forces us to consider a range of computational intelligence techniques. Melding them into a single coherent system which sounds appropriate alarms for a facility under real-time operating conditions is an ongoing challenge which we are meeting with the AMISS system.

References

- [1] K. L. Adair and P. Argo. Classification of behavior using unsupervised temporal neural networks. In *1997 IEEE International Conference on Systems, Man, and Cybernetics*, Orlando, Florida, October 1997. to be published.
- [2] R. Agrawal, T. Imielinski, and A. Swami. Mining association rules between sets of items in large databases. In *Proceedings of the ACM SIGMOD Conference on Management of Data*, pages 207–216, Washington D.C., May 1993.
- [3] P. Argo. Integrating materials protection, control, and accounting and facility monitoring. *Arms Control and Nonproliferation Technologies*, pages 25–27, First/Second Quarter 1996.
- [4] G. A. Carpenter, S. Grossberg, and D. B. Rosen. Fuzzy art: Fast stable learning and categorization of analog patterns by an adaptive resonance system. *Neural Networks*, 4(759-771), 1991.
- [5] H. Graf, E. Cosatto, D. Gibbon, M. Kocheisen, and E. Petajan. Multi-modal system for locating heads and faces. In *Proceedings of the Second International Conference on Automatic Face and Gesture Recognition*, pages 88–93, 1996. sponsored by the MIT Media Laboratory.
- [6] S. Haykin. *Neural Networks: A Comprehensive Foundation*, chapter 10: Self-Organizing Systems II: Competitive Learning, pages 397–434. Macmillan College Publishing Company, Englewood Cliffs, NJ, 1994.