Intruder detection using a wireless sensor network with an intelligent mobile robot response

YuanYuan Li and Lynne E. Parker
Distributed Intelligence Laboratory, Department of Electrical Engineering and Computer Science
The University of Tennessee, Knoxville, Tennessee 37996-3450
Email: {yli, parker}@eecs.utk.edu

Abstract

In this paper, we present an intruder detection system that uses a wireless sensor network and mobile robots. The sensor network uses an unsupervised fuzzy Adaptive Resonance Theory (ART) neural network to learn and detect intruders in a previously unknown environment. Upon the detection of an intruder, a mobile robot travels to the position where the intruder is detected to investigate. The wireless sensor network uses a hierarchical communication/learning structure, where the mobile robot is the root node of the tree. Our fuzzy ART network is based on Kulakov and Davcev's implementation [6]. We enhanced the fuzzy ART neural network to learn a time-series and detect time-related changes using a Markov model. The proposed architecture is tested on physical hardware. Our results show that our enhanced detection system has a higher accuracy than the basic, original, fuzzy ART system.

1. Introduction

There are many advantages of using Wireless Sensor Networks (WSNs) to detect changes in the environment. Each individual node in the network can monitor its local region and communicate through a wireless channel with other nodes to collaboratively produce a high-level representation of the environment's states. By using such a network, large areas can be monitored to detect intruders with low cost. Furthermore, we enhance the detection process by adding autonomous mobile robots. Thus, the system becomes more flexible upon the detection of an intruder. The mobile robot can reach places and perform tasks that static sensors cannot.

The sensor nodes first learn an initial model of the environment using a fuzzy ART neural network; we refer to this as the *normal* model of the environment. After the training period, any changes compared to the learned normal model are treated as anomalies possibly caused by an intruder. Upon the detection of the anomaly, an intrusion alert is generated, and an autonomous mobile robot responds to the alert by traveling to the place where the sensor nodes have detected the anomaly. The robot uses its additional sensors (e.g., a camera) to verify if there is an intruder in the area. The mobile robot uses a camera to track moving objects. If there is a moving object, the robot declares that an intruder is detected.

We have incorporated a machine learning technique into the WSN so that the networks can automatically learn to recognize normal and abnormal modes of operation. Our approach makes use of a fuzzy Adaptive Resonance Theory (ART) neural network, which was first implemented in a WSN by Kulakov and Davcev [6]. The fuzzy ART neural network system is an unsupervised Artificial Neural Network (ANN) that can perform dimensionality reduction and pattern classification. The network can continually learn from new events without forgetting what has already been learned. No off-line training phase is required. The algorithm is simple enough to be implemented in the tiny platform of the Crossbow motes [1], yet still achieve good performance.

However, a shortcoming of the original fuzzy ART approach is that it does not detect time-related changes. We have, therefore, enhanced the basic fuzzy ART system to enable it to learn a time-series through the use of a Markov chain. The approach builds a state transition model on-line during the initial period of deployment, and considers the built model as the normal model. After the training phase is over, any events that occur in the environment that do not fit the existing transition model are considered as abnormal events.

In this paper, we first review related work in Section 2. Then, we present our approach in Section 3. In Section 4, we present the hardware platform that we have used to test our system. Our experimental results from the physical experiments are presented in Section 5. Finally, we summarize our findings in Section 6.

2. Related work

To our best knowledge, no previous work addresses intruder detection by using both a WSN and mobile robots. There are some works that make use of mobile robots together with a WSN in other applications. Laibowitz and Paradiso [7] explore a specific type of mobility in wireless sensor networks — parasitic mobility. They propose a solution to the problems of power usage, node size, and node complexity in the form of parasitically actuated nodes. LaMarca, et al. [8] used robots to increase the feasibility of WSNs because sensor networks can acquire data but lack actuation, while robots have actuation but limited coverage in sensing. Schaffert [13] adapts sensor network models for

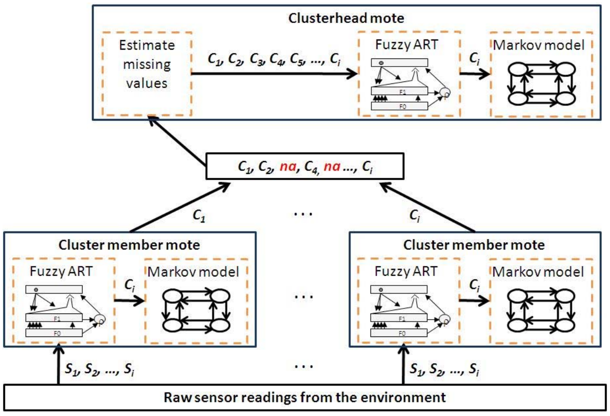

Fig. 1. Our fuzzy ART architecture, extended to estimate missing data and perform time-series analysis.

use with information maps and verifies the ability of such maps to improve robot localization. Ren, et al. [10] focus on a fire and intruder detection application by using sensors only on a mobile robot.

Many machine learning algorithms have been proposed in the WSN area for anomaly detection. Several models of neural networks have been implemented in the area of WSN, such as Multi-layer Perceptrons (MLPs) [11], Self-Organizing Maps (SOMs) [3], and Adaptive Resonance Theory (ART) [6]. We have chosen the fuzzy ART model proposed by Kulakov and Davcev in [6] for our WSN implementation, because of its unique ability to solve the stability-plasticity dilemma, to learn in a short period of time in the fast-learning mode, and to continually learn from new events without forgetting what has already been learned [2], [6], [12]. Nevertheless, Kulakov and Davcev's ART models cannot detect time-related changes. In this paper, we propose an algorithm to enhance the model so that it can detect time-related changes.

3. Approach

In this section, we first introduce our network architecture. Then, we describe the basic fuzzy ART network. Subsection 3.3 then discusses our approach to incorporate a Markov model for time-series analysis. Section 3.4 briefly mentions our approach to missing data estimation.

3.1. Architecture for the sensor networks

In our approach, sensors are arranged hierarchically, as shown in Figure 1. The hierarchical structure has been used in many sensor networks, such as [5]. In our WSN, sensor nodes are divided into clusters. Each cluster has a clusterhead and multiple cluster members. Each cluster covers a geometric region and is responsible for detecting the environmental changes in that region. Both cluster members and clusterheads run an identical detection system — a missing data estimator, a fuzzy ART network, and a Markov model. Cluster members read in raw sensor readings, s_i, (e.g., light and sound) from the environment as input, and then classify data into categories c_i. After the classification process, cluster members send their category labels to their clusterheads. The clusterheads first pre-process the collected category labels by identifying and estimating the missing values. Then, the processed categorizations are used as input to the fuzzy ART neural network and are fused together to reduce the dimensionality of the data. The output of the fuzzy ART network is a category label c_i. After the categorization process is finished, the system further checks if there are time-related changes. Clusterheads may have higher level clusterheads which classify their output class labels. Finally, the root mote obtains the final model of the environment.

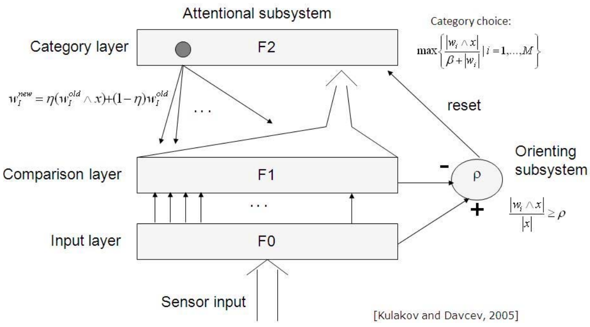

Fig. 2. A typical fuzzy ART architecture (see [6]).

The fuzzy ART model alone cannot detect time-related abnormal events. For example, if people turning on the lights during the day and turning off the lights when they leave work is considered as a normal event, then an intruder only turning on the lights briefly in the evening should trigger an alarm. We enhance the basic fuzzy ART network by using a Markov chain to detect these abnormal events. The Markov model takes the category/state and checks if the transition to the current state is probable based on the existing history. We believe that with this architecture, we can detect abnormal environmental changes as well as time-related changes. Our design is flexible because it allows the operator to turn off the Markov model if time series data is not of interest.

3.2. The fuzzy ART network

Adaptive Resonance Theory (ART) is a neural network architecture developed by Grossberg and Carpenter [2]. The basic fuzzy ART network is an unsupervised learning model, which is able to take analog input. Kulakov and Davcev proposed a fuzzy ART model for change detection in a WSN in [6]. Our basic fuzzy ART network is implemented in the same way. Figure 2 gives a representation of their fuzzy ART network. A typical fuzzy ART network has three layers: an input layer (F0), a comparison layer (F1) and a category layer (F2). For completeness, we present the learning process of the neural network in detail as follows.

Compare each input vector I ($I_j \in [0, 1]$; and $j = \{1, 2, ..., N\}$) with each category/prototype in F2 to classify it with its best match. The choice function A_j is defined:

$$A_j(I) = \frac{|I \wedge w_j|}{\epsilon + |w_j|} \quad (1)$$

where parameter w_j is the binary weight vector of category j, and parameter $\epsilon \in [0, 1]$. The operator $\wedge$ is defined by $x \wedge y \equiv min(x, y)$ and the operator $|\cdot|$ is defined by $|x| \equiv \sum_{i=1}^{M} x_i$. Then, the approach selects the F2 node J that has the highest match ($A_J = max\{A_j | j = 1, ..., N\}$). The weight vector of the winning node (w_J) is compared to the current input at the comparison layer. The training process starts if the match function of the chosen category meets the vigilance; that is:

$$\frac{|I \wedge w_J|}{|I|} \geq \rho \quad (2)$$

where parameter $\rho \in [0, 1]$ represents the vigilance level. The number of developed categories can be controlled by the vigilance level. The higher the value, the more sensitive the network is to changes in the environment, which can result in a larger number of finer categories. When $\rho = 1$, the network creates a new category for every unique input pattern.

The prototype node J in F2 captures the current input, and the network learns by modifying the weight vector w_J according to:

$$w_J^{new} = \gamma(I \wedge w_J^{(old)}) + (1 - \gamma) w_J^{(old)} \quad (3)$$

where parameter $\gamma \in [0, 1]$ is the learning rate. Fast learning occurs when $\gamma = 1$. If the stored prototype w_J does not match the input sufficiently (i.e., Equation (2) is not satisfied), the winning F2 node will be inhibited until a new input vector is applied. Then, another F2 node is selected with the highest A_j value, whose prototype will be matched against the input. This process is repeated until the network either finds a stored category whose prototype matches the input, or selects the uncommitted F2 node if all prototypes result in mismatches. In this case, learning a new category is initiated according to Equation (3).

The fuzzy ART system has a category proliferation problem. Complement coding can be used to overcome this problem [12]. The complement of input I can be achieved by preprocessing each incoming vector a by $a^c \equiv 1 - a$. After the complement coding process, input I becomes a $2M$-dimensional vector, $I = (a, a^c) \equiv (a_1, ..., a_M, a_1^c, ...a_M^c)$. Note that normalization of input I can be achieved by preprocessing each incoming vector a, $I = \frac{a}{|a|}$. Inputs preprocessed into complement coding form are automatically normalized.

3.3. Markov model extension

We enhance the existing fuzzy ART network by adding a Markov chain to detect time-related changes. By definition, a Markov chain is a discrete-time stochastic process with the Markov property, which states that, for a given process, knowledge of previous states is irrelevant in predicting the probability of subsequent states. At each time increment, the system may either stay in the same state, or transition to a new state. A Markov chain is formally defined as a sequence of random variables $X_1, X_2, ...$, which, given the current state, the previous and next states are independent. Formally, $Pr(X_{n+1} = x | X_n = x_n, ..., X_1 = x_1) = Pr(X_{n+1} = x | X_n = x_n)$.

Algorithm 1 Building the Markov model

```
for Each time step do
    if The current state is the same as the last time step then
        Record the time spent in this state.
    else
        Record the state transition.
    end if
end for
for Each state i do
    Find the mean μ_i and standard deviation σ_i of the time the system remains in state i.
    Find the transition probability p_ij for each possible state j.
end for
```

In a WSN setting, the Markov model is built during the training phase using the algorithm shown in Algorithm 1. Sensor motes periodically sense the environment at a fixed rate and feed the normalized sensor readings to the neural network to build categories of the environment. For each category/state (i), we keep an average time and the variance of the time the system remains in that particular state. Additionally, for each state we record the state transition probabilities, p_{ij}, to the next set of states. By doing so, an alarm will trigger if the amount of time in a state is either too short or too long. In a similar fashion, if a state transition is not probable, then this may also trigger an abnormal alarm. Thus, we can capture an anomaly from state transitions and from state occupancy time.

3.4. Missing data estimation

Based on our prior work, we have found that 40% of the sensor values are often missing due to the unstable wireless connections, synchronization issues, packet collisions or malfunctioning sensors. Our system deals with missing data by using a spatial-temporal imputation technique which is discussed in detail in our companion paper [9].

4. Hardware platforms

Our wireless sensor network consists of both static motes and mobile robots. The static sensors are Crossbow motes and the mobile robots are Pioneer 3 robots. Human operators can use a laptop to interact with the wireless sensor network. We describe the hardware in the following subsections.

A Crossbow [1] mote contains a processing unit, a sensor module, and a communication module. The processing board contains an 8-bit processor at $8MHz$, a $128KB$ programming memory and a $512KB$ additional data flash memory. The wireless transmission range is around 10

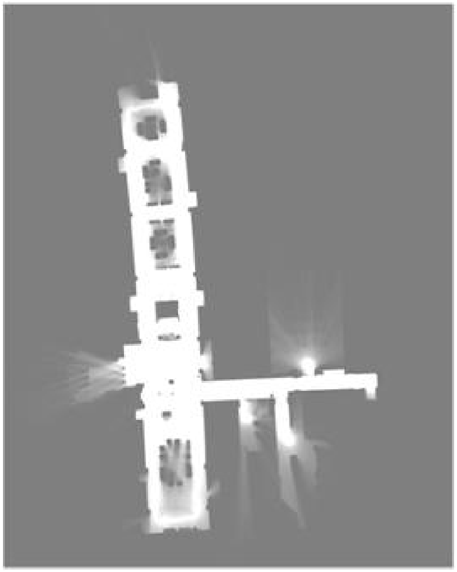

Fig. 3. Robot-generated laser maps of Oak Ridge National Laboratory's JICS building.

meters inside a building. The sensor board has a buzzer, a light sensor, a microphone, 2 magnetometers and 2 accelerometers. For the experiments reported in this paper, we used the light and sound sensing components.

The mobile robot used in these experiments is a Pioneer 3 robot. Pioneer 3 is a mobile robot with a two-wheel differential drive. The mobile robot uses the Linux operating system and runs the Player-client/server device driver [4]. The robot uses a SICK LMS-200 range-finding laser for localization.

The mobile robot can communicate with the sensor motes by having a mote attached to an MIB500 programming board through a serial connection. In our intrusion detection application, the mote on the robot runs the fuzzy ART program and acts as a clusterhead. Thus, the mobile robot is a mobile clusterhead with higher processing power and more sensing capabilities.

5. Experimental results

5.1. Intruder detection system

In order to detect abnormal events in a previously unknown environment, the sensor network first learns what is normal for the environment. Abnormal states of the environment are not kept in the sensor nodes due to space limitations. Therefore, any events that do not match the existing normal model will be treated as abnormal events by the sensor motes. When an intruder is detected, a mobile robot moves to the area to investigate. We assume the robot knows the location of each cluster in advance. If the higher level clusterhead detects an anomaly (i.e., a category change after stabilization), the robot moves to the location of the cluster that detected the change. The mobile robot is the root clusterhead of the hierarchical fuzzy ART system.

In order to navigate in the environment, the mobile robot first creates a laser map using Simultaneous Localization and Mapping (SLAM). An example of the learned map is shown in Figure 3. After an intruder has been detected by the sensor network, the mobile robot uses a wavefront path planning algorithm to plan a path from its current position to the goal position. During motion, it localizes itself using Monte Carlo localization.

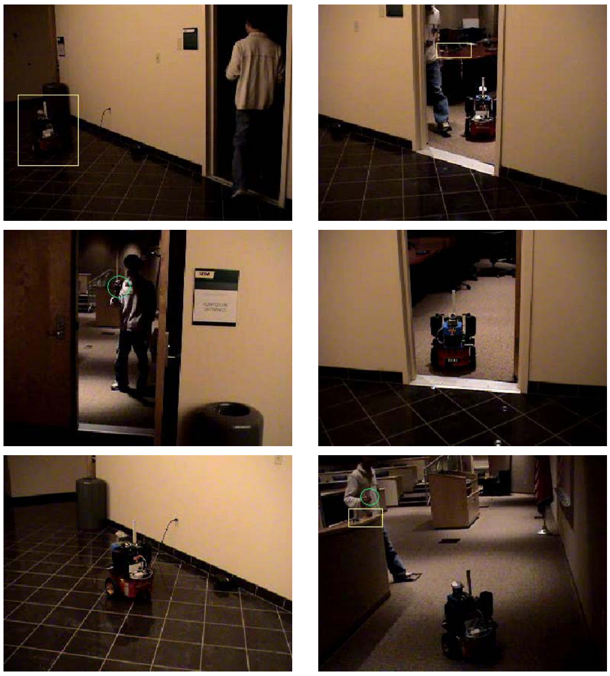

Fig. 4. Snapshots of the intruder detection system in operation at ORNL. Motes and the mobile robot are indicated by rectangles on the picture. The sound device carried by the intruder is indicated by a circle (read left to right, top to bottom).

We implemented and tested the intruder detection system on real motes along with a mobile robot and experimented with the system at both the University of Tennessee and Oak Ridge National Laboratory (ORNL). Figure 4 shows snapshots from the experiments at ORNL. We deployed 2 clusters of sensor motes in the environment. The first cluster was deployed into a conference room of ORNL's JICS building. The second cluster was deployed in an auditorium close by. The mobile robot was stationed in the hallway listening for abnormal changes. The robot was the root clusterhead. It detected abnormal changes by learning the combination of changes of the 2 clusterheads (sensor motes) deployed in the 2 rooms. The mobile robot ran the same learning algorithm as the sensor motes, namely, the fuzzy ART system. In the beginning, it was quiet and the lights were off in both rooms. The WSN learned that "quiet" and "dark" were normal in this environment. Then, an intruder entered the conference room and turned on the lights. The WSN detected the abnormal event and notified the robot. The robot planned a path using its wavefront path planner and moved to the conference room to check on the abnormal event — "light on". The intruder then moved to the auditorium. He turned on the lights and a buzzer to make noise in the auditorium. The robot detected the abnormal

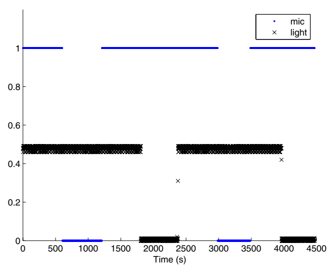

Fig. 5. An example: normalized light and microphone readings collected by a sensor mote. From time 0 to 510, light was on (0.5), microphone was on (1); From time 510 to 1200, light was on (0.5), microphone was on (0); From time 1200 to 1800, light was on (0.5), microphone was off (1); From time 1800 to 2300, light was off (0.1), microphone was off (1), etc.

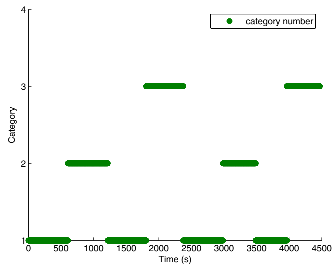

Fig. 6. An example of the categorizations showing changes detected by the fuzzy ART network during the training phase from sensory data shown in Figure 5.

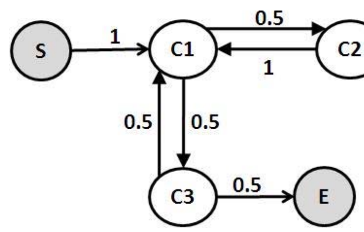

Fig. 7. An example of a learned Markov model for the training phase. The model was the normal model of the environment. States "S" and "E" denote the start and end states, respectively. They were manually added to the system. State C1 denotes lights were on and buzzer was off. State C2 denotes lights were on and buzzer was on. State C3 denotes lights were off and buzzer was off.

activities in the auditorium — "buzzer on" and "lights on". The robot then planned a path and moved to the auditorium to check on the abnormal event. Once the robot arrived at the auditorium, it used its camera to track the intruder.

5.2. Performance metrics

To evaluate our system, we collected statistics on the miss rate, false alarm rate, sensitivity and specificity. The miss rate is calculated as $\frac{FN}{(TP+FN)}$, where False Negative (FN) denotes the number of faults that the system failed to detect, and True Positive (TP) denotes the number of true faults that are detected by the system. The false alarm rate is defined as $\frac{FP}{(FP+TN)}$, where False Positive (FP) denotes the number of detected faults that were not true faults, and True Negative (TN) denotes number of "no faults" that were detected by the system. The sensitivity is defined as $\frac{TP}{(TP+FN)}$. The false alarm rate is defined as $\frac{FP}{(FP+TN)}$. Ideally, the values of sensitivity and specificity are at 100%, and the false alarm rate and miss rate are at 0%.

5.3. Temporal change detection

In this experiment, we began by having the system learn the normal model; then, the testing began. Both training and testing were performed online. All sensors sampled the environment at a rate of 1 sample per second. Six motes were used during this experiment. One mote acted as a clusterhead, and the rest as cluster members of that mote. The cluster member motes were uniformly deployed around the clusterhead and all cluster members were within the communication range of the clusterhead. The vigilance levels for cluster members were set to 0.90, while those for the clusterheads were set to 0.97.

The training process took approximately 1.5 hours. We treated this as a normal environment. Two sensors were used by cluster members — light and microphone. Raw light readings between 0 and 2000 indicated dark and light, respectively. Microphone readings came from a hardware detection system onboard. The values were binary — 1 indicates no noise is detected, and 0 indicates noise is detected. We used a buzzer as a sound source, which operates at $4Hz$. The sound sensor can detect the buzzer within a radius of 3 to 4 meters in our testing environment.

Figure 5 is an example of the typical sensor readings collected from the environment. The light sensory readings are normalized between 0 and 1. Figure 6 shows the categories learned by the fuzzy ART neural network from the data shown in Figure 5 during the training period. After the classification process, a Markov model was built by using Algorithm 1. Figure 7 shows a Markov model built from the data shown in Figure 5 and Figure 6. Table I shows the mean and standard deviation values of the time the environment remained in each category/state before transiting to a different category/state.

TABLE I
TIME DURATION IN EACH STATE

Category	C1	C2	C3
Mean time (s)	571	555	538
Standard deviation (s)	62	75	43

Three different testing suites with four trials were run for each testing suite. In test suite 1, the environment started from "light and quiet" (category 1), and remained in that state for 600 s. Then, it transited to "light and noisy" (category 2), and remained in that state for 600 s. Lastly, it transited to "dark and noisy" (category 4), and remained in that state for 600 s. Note that "dark and noisy" had never occurred before during our training phase. This was an abnormal event. This testing suite only contained a new abnormal state; however, it did not include any temporal-related changes.

In test suite 2, the environment started from "light and noisy" (category 2), and remained in that state for 600 s. Then, it transited to "dark and noisy" (category 4), and remained in that state for 300 s. Lastly, it transited to "dark and quiet" (category 3), and remained in that state for 600 s. The environment started with abnormal transitions to state 2, then the abnormal state 4 was detected. Lastly, an abnormal transition occurred from abnormal state 4 to state 3. This testing suite contained both abnormal events of a new abnormal state and abnormal time transitions.

In test suite 3, the environment started from "light and quiet" (category 1), and remained in that state for 300 s. Then, it transited to "dark and quiet" (category 3), and remained in that state for 900 s. The environment abnormally remained at state 1 too briefly and in state 3 for too long. This testing suites only contained time-related abnormal changes.

TABLE II
PERFORMANCE EVALUATION BETWEEN THE BASIC FUZZY ART AND ENHANCED FUZZY ART

		False Alarm	Miss Rate	Sensitivity	Specificity
Original fuzzy ART	mean	6%	59%	41%	94%
	stdev	12	4	38	12
Enhanced fuzzy ART	mean	6%	14%	86%	94%
	stdev	12	2	2	1

We used these testing suites to compare the performance of the basic fuzzy ART system (Kulakov and Davcev's implementation) and our enhanced fuzzy ART system. The experimental results are shown in Table II, which are averaged over 3 testing suites (for a total of 12 trials). Approximately 1500 observations were made from each sensor for each trial. The experimental results illustrate that our enhanced fuzzy ART system is able to detect more anomalies than the original fuzzy ART system (i.e., approximately 86% vs. 41%). This is due to the fact that our enhanced system learns a time series, and is able to detect time-related anomalies, whereas the original fuzzy ART cannot. Both the fuzzy ART system and the enhanced fuzzy ART system have a low false alarm rate (approximately 6%). To determine the significance of these results, we applied the Student's t-test to the miss rate and sensitivity results for the original fuzzy ART and our enhanced fuzzy ART. This test confirms that the differences in these results are statistically significant, with a confidence level of 99.5%. Thus, our enhanced fuzzy ART approach provides a significant improvement over the original fuzzy ART approach.

After a change is detected in the environment, it does not necessarily mean that an intruder caused the anomaly. Thus, a mobile robot is sent to investigate using an additional sensor (i.e., camera) to determine if the anomaly is caused by an intruder. As future work, we plan to compare the new intruder detection system with the mobile robot's feedback against the original fuzzy ART system and the enhanced fuzzy ART system with time-series extension.

6. Conclusion

We have presented an intruder detection system by using a wireless sensor network and mobile robots. To our best knowledge, this is the first intruder detection system that uses a sensor network to detect intruders and a mobile robot for traveling to the location where the intruder was detected. We have implemented and tested our system on physical motes and robots. The sensor network uses a fuzzy ART neural network to detect intruders. We have enhanced the original fuzzy ART system to detect time-related changes. Our experimental results show that our detection system has high accuracy and is able to detect time related changes. In our ongoing work, we are investigating how to use mobile robots to save energy in a wireless sensor network.

7. Acknowledgments

This research was sponsored in part by the ORNL SensorNet program. We also thank Michael Bailey for help with implementing the fuzzy ART algorithm, the operator control program and integrating the Deluge system.

References

[1] http://www.xbow.com/.

[2] G. Carpenter and S. Grossberg. Fuzzy art: Fast stable learning and categorization of analog patterns by an adaptive resonance system. *Neural Networks*, 4:759–771, 1991.

[3] E. Catterall, K. Van Laerhoven, and M. Strohbach. Self organization in ad hoc sensor networks: An empirical study. In *The 8th International Conference on the Simulation and Synthesis of Living Systems*, Sydney, Australia, 2002. MIT Press.

[4] B. P. Gerkey, R. T. Vaughan, K. Stoy, A. Howard, G. S. Sukhatme, and M. J. Mataric. Most valuable player: a robot device server for distributed control. In *The IEEE/RSJ International Conference on Intelligent Robots and Systems*, volume 3, pages 1226–1231, 2001.

[5] W. Heinzelman, A. Chandrakasan, and H. Balakrishnan. Energy-efficient communication protocols for wireless microsensor networks. In *International Conference on System Sciences*, Maui, HI, January 2000.

[6] A. Kulakov and D. Davcev. Tracking of unusual events in wireless sensor networks based on artificial neural-network algorithms. *Information Technology: Coding and Computing (ITCC)*, 2:534–539, 2005.

[7] M. Laibowitz and J. A. Paradiso. Parasitic mobility for pervasive sensor networks. In *Pervasive Computing, Proceedings of the Third International Conference*, volume 2, pages 255–278, Munich, Germany, May 2005.

[8] A. LaMarca, W. Brunette, D. Koizumi, M. Lease, S. B. Sigurdsson, K. Sikorski, D. Fox, and G. Borriello. Plantcare: An investigation in practical ubiquitous systems. In *The 4th International Conference on Ubiquitous Computing (UbiComp)*, pages 316–332, London, UK, 2002. Springer-Verlag.

[9] Y. Li and L. E. Parker. Classification with missing data in a wireless sensor network. *To appear in IEEE SoutheastCon*, April 2008.

[10] R. C. Luo, T. Y. Lin, H. C. Chen, and K. L. Su. Multisensor based security robot system for intelligent building. In *The IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems*, Heidelberg, Germany, September 2006.

[11] L. Reznik, G. V. Pless, and T. A. Karim. Signal change detection in sensor networks with artificial neural network structure. In *The IEEE International Conference on Computational Intelligence for Homeland Security and Personal Safety (CIHSPS)*, pages 44–51, Orlando, Florida, March 2005.

[12] E. Sapojnikova. *ART-based Fuzzy Classifiers: ART Fuzzy Networks for Automatic Classification*. PhD thesis, University of Tubingen, Tubingen, Germany, April 2003.

[13] S. M. Schaffert. *Closing the Loop: Control and Robot Navigation in Wireless Sensor Networks*. PhD thesis, EECS Department, University of California, Berkeley, September 2006.