Copyright © 2009 Tech Science Press CMC, vol.14, no.1, pp.1-22, 2009

Anomaly Detection

Nadipuram R. Prasad[1], Salvador Almanza-Garcia[1] and Thomas T. Lu[2]

Abstract: The paper presents a revolutionary framework for the modeling, detection, characterization, identification, and machine-learning of anomalous behavior in observed phenomena arising from a large class of unknown and uncertain dynamical systems. An evolved behavior would in general be very difficult to correct unless the specific anomalous event that caused such behavior can be detected early, and any consequence attributed to the specific anomaly following its detection. Substantial investigative time and effort is required to back-track the cause for abnormal behavior and to recreate the event sequence leading to such abnormal behavior. The need to automatically detect anomalous behavior is therefore critical using principles of state motion, and to do so with a human operator in the loop. Human-machine interaction results in a capability for machine self-learning and in producing a robust decision-support mechanism. This is the fundamental concept of intelligent control wherein machine-learning is enhanced by interaction with human operators.

Keywords: Anomaly detection, soft-computing, decision-making, machine intelligence, nonlinear dynamical systems

1 Introduction

Asymmetric threats from roadside improvised explosive devices (IEDs) deployed by terrorists and insurgents in Iraq and in Afghanistan have given rise to a new class of anomalies that need to be detected. Detection must occur at the time of their deployment so that appropriate countermeasures can be undertaken to effectively defuse the IEDs. The US military estimates that 75% of troop casualties in Iraq and in Afghanistan are due to roadside bombs and fears that casualties might grow rapidly over the next few years. Mine Resistant Ambush Protected (MRAP) vehicles have been deployed to war zones as a means to protect US and coalition forces against the threat of IED attacks. Despite the protection afforded by this

[1] New Mexico State University, Las Cruces, NM 88003

[2] Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109

technologically advanced vehicle, detecting IEDs remains the number one problem for military operations against terrorist attacks. Of particular interest to the US Department of Defense is the detection of roadside car bombs and other forms of IEDs placed by insurgents and terrorists.

Anomalies by virtue of their definition are events that occur spontaneously with no prior indication of their existence or appearance. Any event deemed as being out-of-the-ordinary may be called an anomaly. Effects of anomalies are typically unknown until they actually occur, and their effects aggregate in time to show noticeable change from the original behavior. The question is how to detect any abnormal activity and to do so using efficient computer-aided techniques. At a minimum, fast and efficient algorithms that will alert human operators to react quickly to what appears as an anomaly are needed.

An artificial intelligent system is a machine that can mimic human decision-making capabilities and strategies, and aid humans towards improved understanding of the observed phenomena. Naturally advanced machine intelligence leads to better and faster human-machine decision-making abilities. Accomplishing such a goal requires seamless human-machine interactions that make the need for machine self-learning a fundamental requirement. Machine self-learning leads to autonomy and gives rise to the possibility for developing systems which provide humans with fast decision-making capabilities.

A fully autonomous decision-making system however may not always be favorable due to the possibilities for false alarms and the possible repercussions due to unwarranted machine-made decisions. A human in the loop must therefore be present to correct the machine outcomes that satisfy the human perception of the event. We are referring to the “event” as that detected by the machine as a possible anomaly. Identification of an event as an anomaly therefore must be performed under human supervision and intervened, if necessary.

The objective in this paper is to outline and discuss a revolutionary approach to detect anomalous events while simultaneously building machine intelligence through interaction with humans. The focus of this paper is on the development of a machine intelligence-based decision-support (MIND) system that can assist humans in automatic anomaly detection Prasad et al (2005); Rangamani (2006); Prasad et al (2007). To achieve such an objective, the computational platform that emulates human intelligence must represent a natural integration of human-like decision-making capabilities that include logic, memory, and evolutionary abilities that together provide sufficient decision-making attributes for decision making under uncertainties. Combining the attributes of these three fundamental decision-making abilities in human-like form has the propensity for MIND systems to parallel the decision-making strategies of the human Mind, at least from a task-oriented viewpoint. In principle, one would agree that machine learning initially occurs from training, and subsequently the machine learns from similarities with training examples. This represents a transition from supervised learning to self-learning, which is how humans learn.

The technology discussed in this paper exploits the innate human abilities for logical reasoning, learning, and adaptation to distinguish between normal and abnormal behavior in observed phenomena, and in detecting, identifying, characterizing, monitoring, and possibly in some cases exerting control to avert anomalous behavior. The fundamental issue that we must address is the recognition of an abnormal event when it occurs so we can classify the event as abnormal and then use the information in the context of deciding what to do next. This would require having a perception of what we are attempting to identify amidst "noise" in the observed system, or within the background clutter in the environment we are attempting to perceive. Specific perceptions are needed to identify the type of anomaly that needs detection.

Anomalous behavior is the emergence of disorder within an ordered set of coupled dynamical systems which under "normal" conditions generally exhibit some form of quasi-steady-state behavior. A good example of this is the societal dynamics in regions like Iraq and Afghanistan that are volatile and unstable, and are yet considered normal within their societies amidst turmoil and chaos. The fundamental nature of such dynamical systems is that they react quite rapidly to perturbations, and appear to settle to a state of quasi-normal conditions relatively quickly. The causes for such dynamic behavior are due to the embedded uncertainties that are inherent to the system as a whole and more specifically represent the adaptive capabilities of the system.

From a system-theoretic viewpoint, reaction to perturbations cause ripples within the system which transfers energy to neighboring coupled systems – a form of chain reaction. If the energy transfer is large enough then there is a potential for disorder to grow, and ultimately result in chaos. As such, order within chaos may be regarded as a flux-field that is typically in a quasi-steady-state and which enables us to speculate upon the magnitude of change required to disturb the flux that could lead to disarray and confusion.

It is well understood in chaos theory that order and chaos are not mutually exclusive. Order within chaos is a transitory process. In fact anomalous behavior has a tendency to drive a system that is initially in quasi steady-state towards instability. A system that is in "steady-state" is only marginally stable in a sense with the possibility of being moved into a chaotic state depending upon the magnitude and impact of perturbations. For a system that remains in steady-state there is a sense of robustness that is governed by the energy transfers between the states manifested by the system dynamics. The robustness is due to the ability of system super-states whose energy to withstand perturbations is greater than the energy imparted by sub-states to cause any transitions in overall state behavior. Transitions from steady-state to a chaotic state are in essence the result of anomalous behavior that represents sufficient energy build-up in the sub-states within the neighborhood of stable states to cause a shift in the dynamic behavior. Such shifts are generally the cause for dynamic instability. Anomalous behavior in systems therefore is a cause for dynamic instability and detecting such shifts is of paramount importance.

Figure 1 illustrates a conceptual characterization of "normalcy" and the shift from normalcy due to anomalous behavior.

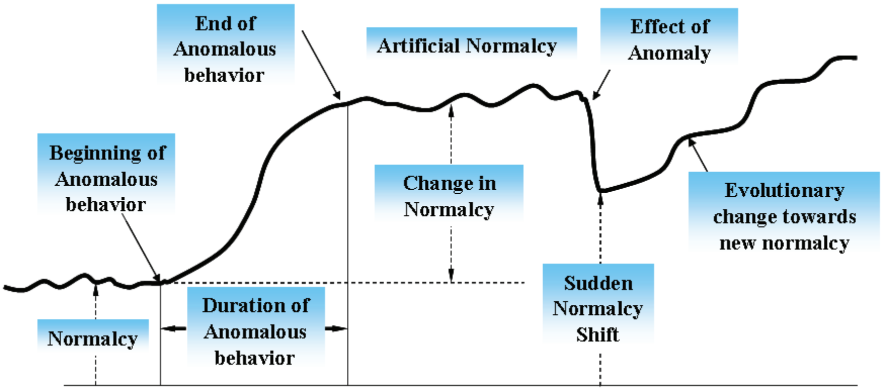
End of Anomalous behavior
Artificial Normalcy
Effect of Anomaly
Beginning of Anomalous behavior
Change in Normalcy
Evolutionary change towards new normalcy
Sudden Normalcy Shift
Normalcy
Duration of Anomalous behavior

Figure 1: State transitions of an observed system characterizing the change from normalcy due to anomalous behavior

Referring to Figure 1, we define normalcy as the degree of pattern consistency in observed states where the motion of states represented by a state vector has a relatively constant norm. One example could be that each state has a relatively constant velocity and hence the vector norm is also relatively constant. Small variations in state motion could be acceptable as long as the norm is within a threshold of perceived normalcy. Again we emphasize that human perception of what constitutes "normal" is indeed very subjective. If we observe a scene long enough, we tend to characterize the states in a scene to be normal if there is no significant deviation from what we already know is normal.

When something very unusual occurs, such as the velocity of some states suddenly becoming zero for a period of time that is longer than anticipated, or the sudden scattering of a crowd in a mass gathering for example, then there is cause for concern which points to the beginning of an anomaly. Here again the notion of normalcy in the pattern of behavior is important to note so that if indeed there is an onset of abnormal behavior then the anomaly can be detected.

Logic has an important role in identifying whether or not a specific set of observed conditions constitute anomalous behavior. Using rule-based reasoning, dynamic system states momentarily exhibiting zero velocity in the observed pattern of behavior can be identified as being normal, or as abnormal. A simple example of this is a stop sign at a road intersection where automobiles must stop momentarily. An anomaly however is when an automobile does not stop when directed to do so. Another example could be the normal traffic on a roadway when suddenly an automobile comes to a standstill along the roadside where automobiles are generally not parked. While these examples serve to demonstrate specific instances, it is indeed possible to develop rules in which the antecedents can take on substantial variations in parameters to cover the wide range of human perceptions needed to identify anomalous behavior.

The duration over which the anomaly persists causes a dramatic change in the normalcy that could have prevailed earlier. This possibility can be examined by extending the previous example wherein if more than one automobile were to pass a stop sign without stopping and the pattern continues, then a condition of artificial normalcy may come about. Similarly if the automobile which came to a standstill remains in a specific location for more than a specified amount of time with other observed states appearing to be normal, then an artificial state of normalcy exists. The duration of this artificial normalcy will last until the effect of the anomaly unfolds itself.

Suppose the standstill automobile turned out to be an IED, namely, a roadside bomb, then immediately following its detonation (intentional or unintentional) there is a shift in normalcy towards a new flow of traffic. From this moment on the state of normalcy begins to evolve and achieve a new state until the next anomaly occurs, and the cycle of transitions between normalcy and abnormality continues. State transition between normal and abnormal behavior therefore is a consequence of anomalous behavior and approaches which can transform observations to signal form are needed to capture this transition.

From a systems science perspective the idea of anomaly detection is synonymous to detecting change in motion around “singularity” points in the observed phenomena. Singularities exhibit trajectories (patterns) that are characteristic to the type of behavior known to exist in nonlinear systems, e.g., stable node, stable focus, saddle-point, limit cycles, etc. Detecting change from normalcy therefore requires the identification of change in trajectories (patterns) from “normal” behavior.

Figure 2 illustrates two examples of nonlinear motion around singularity points.

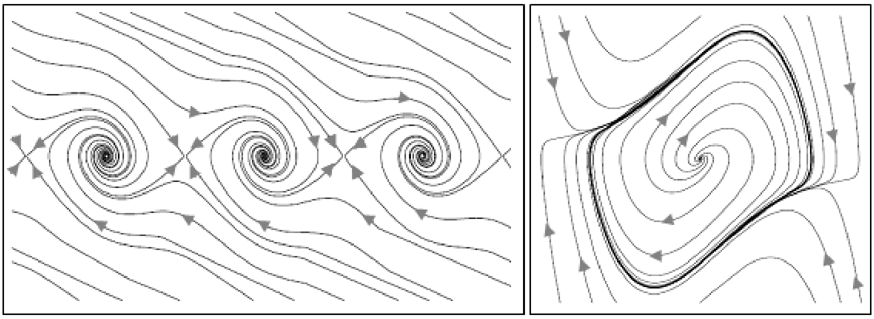

Figure 2: Phase portraits of trajectories around singularities exhibiting stable focus, saddle points, and limit cycle behavior

The trajectories in Figure 2 illustrate typical behavior of nonlinear systems which aid in the interpretation of physical systems. Referring to the phase portrait on the left, it is seen that trajectories close to the singularity points are attracted toward a stable focus, while those at a saddle-point are marginally stable and are deflected towards a stable point. The phase portrait on the right represents a limit cycle behavior with all trajectories, regardless of their initial states, converging toward a closed path

2 Modeling approach

In a multi-sensor environment such as that in the battlefields of Iraq and Afghanistan where electro-optical (EO), infrared (IR), video, synthetic aperture radar (SAR), and Moving Target Indicator (MTI) sensing systems mounted on UAVs and manned aerial vehicles are used for surveillance, the area covered is so large that a polling system is required to poll all the sensors and as such images are obtained at discrete time intervals. Fusing the information generated by these sensors effectively is a complicated task towards achieving a high degree of situational awareness Fennell and Wishner (1998). To utilize these discrete-time images requires some effective means to observe change in the images that clearly show anomalous behavior. As such, the development outlined in this paper focuses on a highly innovative, simple and yet elegant modeling approach which supports closed-loop sensor-based monitoring and control of dynamical processes.

Because anomaly detection is based upon a perceived sense for change, the motivation towards employing a fuzzy logic-based approach is clearly warranted by the need to give linguistic meaning to what anomalous behavior means Zadeh (2002). This is especially significant where the detection pertains to the nature of human behavior that is the cause of an anomaly. Our overall objective is to develop a revolutionary framework using machine intelligence that will allow the characterization, detection, identification, and modeling of anomalous behavior in observed phenomena arising from a large class of unknown and uncertain dynamical systems. A computational platform using "soft computing" to bio-mimic human decision-making in detecting anomalous behavior is proposed.

As stated previously, from a system-theoretic viewpoint it is postulated that anomalies are synonymous to singularities in chaotic/nonlinear systems. Singularities exhibit trajectories that are characteristic of the type of behavior known to exist in nonlinear systems. By exploiting this characteristic of nonlinear systems, our approach of mapping the trajectories of process states has the potential to uncover unique trajectories corresponding to possible anomalous behavior.

Figure 2 illustrates a conceptual framework for anomaly detection.

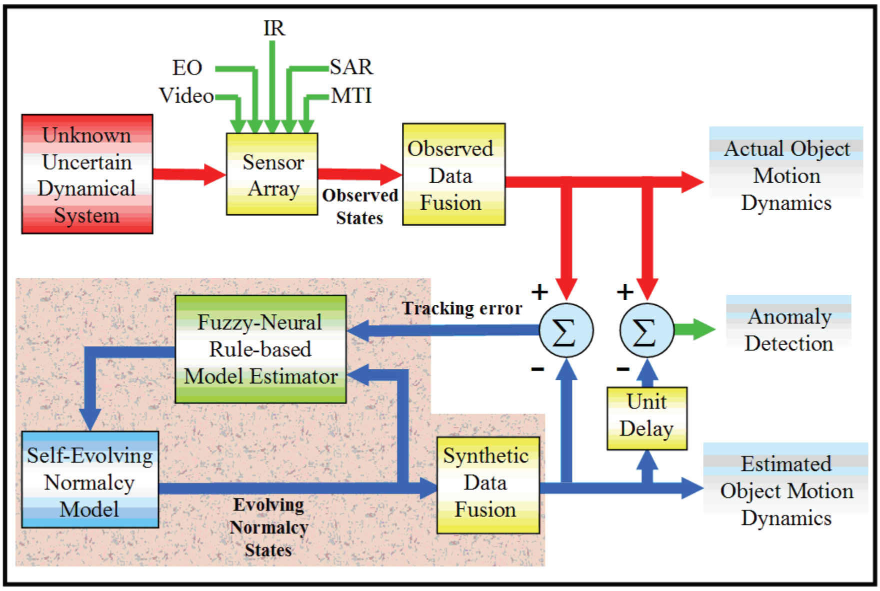
IR
EO
SAR
Video
MTI
Unknown Uncertain Dynamical System
Sensor Array
Observed States
Observed Data Fusion
Actual Object Motion Dynamics
Fuzzy-Neural Rule-based Model Estimator
Tracking error
+
+
Σ
Σ
-
-
Anomaly Detection
Unit Delay
Self-Evolving Normalcy Model
Synthetic Data Fusion
Evolving Normalcy States
Estimated Object Motion Dynamics

Figure 3: Conceptual framework for developing normalcy models and for detecting anomalous behavior in complex dynamical systems

The concept is centered around tracking the dynamics of unknown and uncertain dynamical systems, and to utilize tracking data in developing a self-evolving normalcy model of the observed phenomenon. A self-evolving normalcy model is developed using a fuzzy-neural architecture Jang et al (1996); Nguyen et al (2002).

The use of robust tracking algorithms ensure the tracking error is at a minimum so that a reliable normalcy model could be anticipated as a basis for anomaly detection Li et al (2008); Gentile (2001). The effect of noise and occlusion of objects in frames need to be minimized. Normalcy modeling is equivalent to establishing a set of initial conditions from which system state changes can occur. If the change is significant enough to cause observable perturbations then the perturbations have the potential for being characterized as anomalies.

Referring to Figure 2, tracking error is computed by comparing the actual object motion dynamics with synthesized (estimated) object motion dynamics. A fuzzy-neural rule-based model estimator minimizes the tracking error by continuously updating the normalcy model at each sampling instant.

As illustrated in Figure 3, close tracking provided by the tracking algorithms mask the occurrence of an anomaly, if one did occur. In other words, if the model-based estimator produces very accurate estimates of the object motion dynamics, then the tracking error must be very close to zero. Consequently, the anomaly is masked and hence cannot be observed.

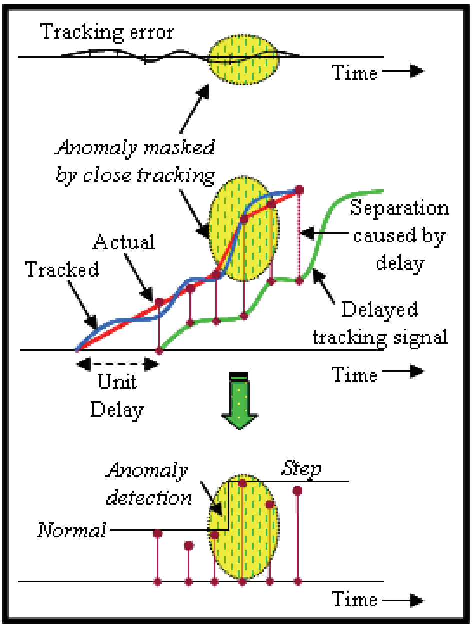

Figure 4: Detecting anomalies as a step change

For the detection of anomalous behavior, we propose a novel approach that bio-mimics human reasoning of comparing past information with present in order to distinguish change from normalcy. The idea is to compare a delayed tracking signal with a real-time signal of the actual object motion dynamics. This approach separates the signals thereby amplifying the change from normalcy. The "step change" allows a clear means for anomaly detection.

The unit delay shown in Figure 3 has the effect of providing a means to examining the sensitivity of normalcy models to change caused by anomalous events. Because the exact delay is not known *a priori* it must be treated as a sensitivity parameter – a parameter that automatically adjusts the "sensitivity" to change in patterns of the observed phenomena. Because anomalies are typically short-lived, the return to normalcy is provided by yet another indication of a step change that could be used to confirm an anomaly as having occurred. This adds reliability to the detection and towards the use of normalcy-based models for predicting anomalous behavior in complex dynamical systems.

2.1 Characterization and detection of anomalous behavior

Modeling the normal background clutter in which anomalous events could occur is indeed a major step in characterizing anomalous behavior. This is analogous to developing the necessary initial conditions which prescribe the environment or domain in which anomalous behavior could occur. The idea is to employ data products from existing and future intelligence surveillance and reconnaissance (ISR) systems such as unmanned aerial vehicles (UAVs) and Space-based imaging systems that offer both sensitivity and selectivity, respectively. High resolution images from UAVs offer the desired level of observation sensitivity (fine-granularity), while Space-based observations provide selectivity (coarse-granularity). A combination of these imaging sources provides the desired information detail to develop models of normal background activity He (2005).

While anomalous behavior is characterized as a deviation from normal observed behavior, detecting such a change at the instant of occurrence requires knowledge of what is normal to begin with. Not all changes from normal behavior are anomalies and may in fact represent evolutionary changes from what was previously considered normal. In fact, it is important to differentiate between the normal nonlinear behavior and that which constitutes an anomaly. This requires a constant "update" of normalcy models that describe gradual changes in observed information. What we seek is the so-called step change that may be characterized as anomalous and hence worthy of further investigation. In regions such as Iraq and Afghanistan where there is constant change, it is reasonable to expect normalcy models to evolve continuously throughout a 24-hour period.

Of particular concern and interest to the US military in Iraq and in Afghanistan is the occurrence of anomalies in the urban battlefield that if left undetected pose serious threat to the safety of the warfighter. An example of this is the deployment of improvised explosive devices (IEDs) by terrorists and militant insurgents whose activities constitute anomalous behavior. The most serious threat issues are indeed when, where, and how terrorists deploy such deadly force against unsuspecting US warfighters and coalition security forces. In this context, two very important criteria, namely, selectivity and sensitivity need to be considered. These two important criteria need to be integrated within the framework for anomaly detection systems.

When and where insurgents deploy IEDs is a selectivity issue as it requires a macroscopic (large area) view of potential attack sites. This is true both from the insurgent's point-of-view as well as from the military point-of-view to selectively identify/locate potential attack sites. From the insurgent's standpoint we can perceive that selectivity is based upon their mission success and ultimately the effectiveness of their planned action. The issue pertaining to "how" such deployment is performed is indeed a sensitivity issue. From the military point-of-view sensitivity is reflected in the imaging resolution of the urban ground terrain that allows a microscopic (small area) examination of insurgent activity in and around suspected and potentially viable attack sites. It leaves no doubt therefore that both sensitivity and selectivity dictate the granular details of the anomaly within the observed process.

It is clear that anomaly detection requires the ability to first track the behavior of the dynamical system. Tracking provides all the information to develop a model "on-the-fly" while also being able to reproduce the trajectories of object motion from any frame or sequence of frames on demand. This may be required for the purpose of verifying that an anomalous event occurred, and to verify the accuracy of the normalcy model itself. This is how humans would have to reconstruct the scenario after an IED event.

Detecting anomalies requires the recognition of some pattern of behavior in the observed process that has either never been seen before, or that the pattern of behavior is out-of-the-ordinary and is different from the norm. As such there are a number of questions one could ask that deal with what needs to be observed in order to see what is different from the norm. Naturally there is a compelling need towards discovering some underlying "pattern" in behavior that vividly describes the fundamental behavior of systems. From a systems science perspective these patterns are indeed those which describe the trajectories which characterize strange attractors.

In the following sections we illustrate how trajectories of motion can be extracted from image sequences that give rise to the potential for anomaly detection.

3 An example for anomaly detection

Scenarios aid in developing a general framework for machine learning while providing a means to formulate a strategy for detecting anomalies. Scenarios provide a “benchmark” towards establishing normalcy in observed information, and to detect any change in normalcy as and when they appear.

For illustrative purposes, we consider a scenario of an isolated intersection frequently patrolled by the US Army which is targeted by Al Qaeda for a possible IED site. Figures 4(a)-(e) illustrate a sequence of simulated images illustrating object motion within the horizontal field-of-view of a UAV. We assign states to each moving object. We arbitrarily establish a center-line reference such that the range of state motion is constrained within a field-of-view labeled [-X, X].

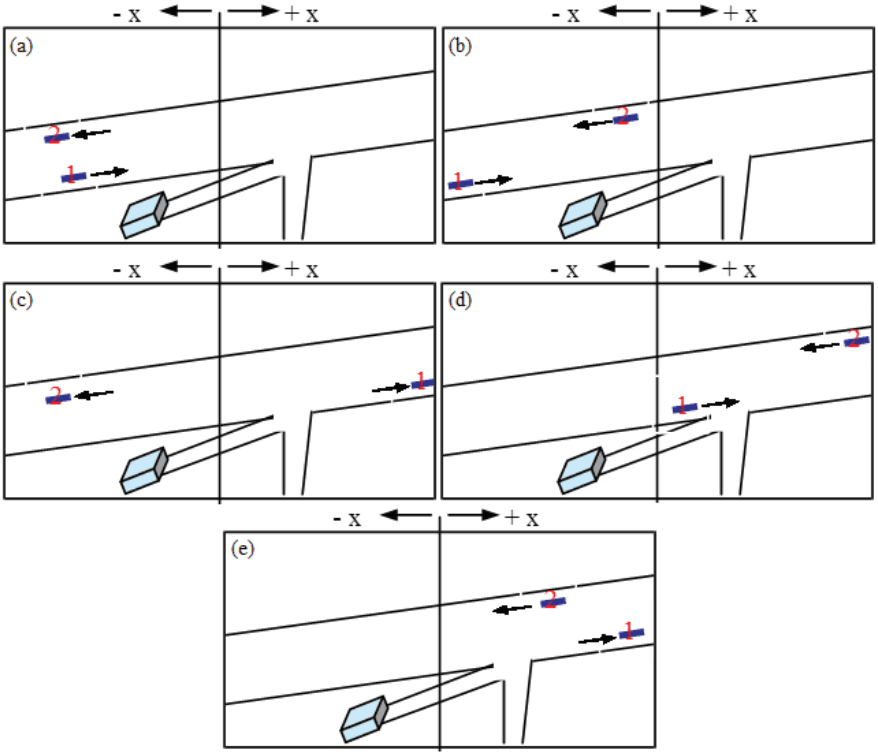

Figure 5: Image sequence (a) through (e) showing “normalcy” pattern in observed field-of-view

Although there are no restrictions to the number of states in motion, we assume two states x_1 and x_2 to represent two vehicles in the field-of-view, and at any instant of time that are travelling in opposite directions, i.e., one travelling in the direction from –X to +X and the other travelling in the direction from +X to -X. Both position and velocity are easily determined based on the imaging frame rate and the position of objects in each frame. Mapping the position and velocity of the two states provide a time-series representation of the observed phenomena. The objective then is to map the time-series data to a state-space that yields a phase portrait of the observed process. Depending upon the time of travel from either the positive reference to the negative reference, or vice versa, the phase portrait over time represents "normalcy" in the field-of-view. The time series data serves as input to a set of fuzzy rules developed in the form of a neural network structure. Over sufficient time, patterns of behavior emerge that could be used as those representing normalcy.

Figures 5(a)-(e) show a simulated image sequence to illustrate a possible anomaly.

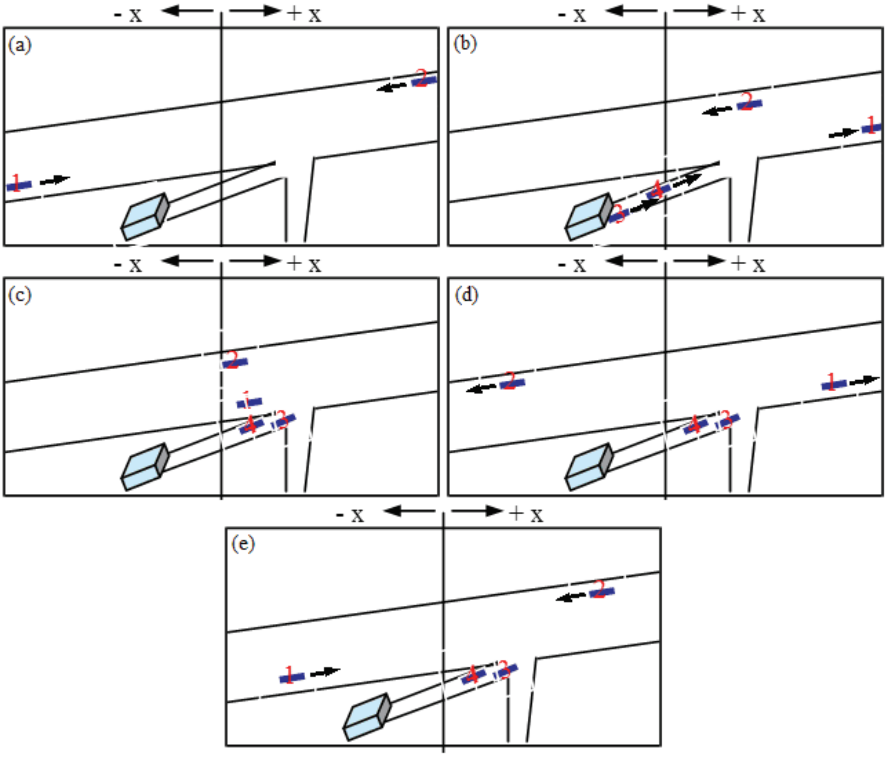

Figure 6: Image sequence (a) through (e) in which two new states have appeared

In Figure 5, two new objects suddenly appear in the field-of-view that are in fact two new additional states x_3 and x_4. Notice that the two new states emerge from a box that is offset from the center-line of the field-of-view. This is a change from the normalcy patterns observed in figure 4 which allows detecting anomalous change in behavior due to the "*possibility of IED threat*".

The following descriptions of state motion apply to the image sequence in Figure 5.

Normalcy pattern is shown with two initial states x_1 and x_2 exhibiting position and directional velocity.

Additional states x_3 and x_4 suddenly appear and exhibit position and directional velocity. This is the beginning of an anomaly, and a "step change". At this point there are four states in motion.

Since the initial anomaly where a set of four states were in motion, a new state behavior emerges wherein all states x_1, x_2 x_3, and x_4 come to rest for a finite interval, another step change.

The next change of behavior in the anomalous pattern is when states x_1 and x_2 begin motion while states x_3 and x_4 remain at rest. This would constitute a step change as well.

A new normalcy pattern emerges with two states in motion and two states at rest. Based on the scenario that this was an IED possibility by insurgents it would be wise to train a machine to recognize the pattern signatures. We assume the remainder of the imaging will show the behavior of states x_1 and x_2 in addition to stationary states x_3 and x_4 which show the evolved change in normalcy.

A phase portrait of objects in motion taken from a synthesized set of images is depicted in Figure 6. The trajectories show normal behavior when object motion has a pattern consistency. When the pattern consistency is disturbed a possible anomaly is identified. Following the anomaly the emergence of a new state of normalcy is identified.

4 A framework for automatic anomaly detection

Tracking provides a basis for developing normalcy models from which any "distinct" deviations may be characterized as being anomalous. We assume that we have very good tracking algorithms that track the object motion described in Figures 4 and 5 very closely. As such the tracking error may be assumed to be very close to zero. Given that we now have both tracking and actual motion trajectories, we separate the signals by a delay of duration "$\boldsymbol{a\tau}$" seconds, where "$\boldsymbol{a}$" is the sensitivity control parameter. In a typical manual setting, users would shift the data

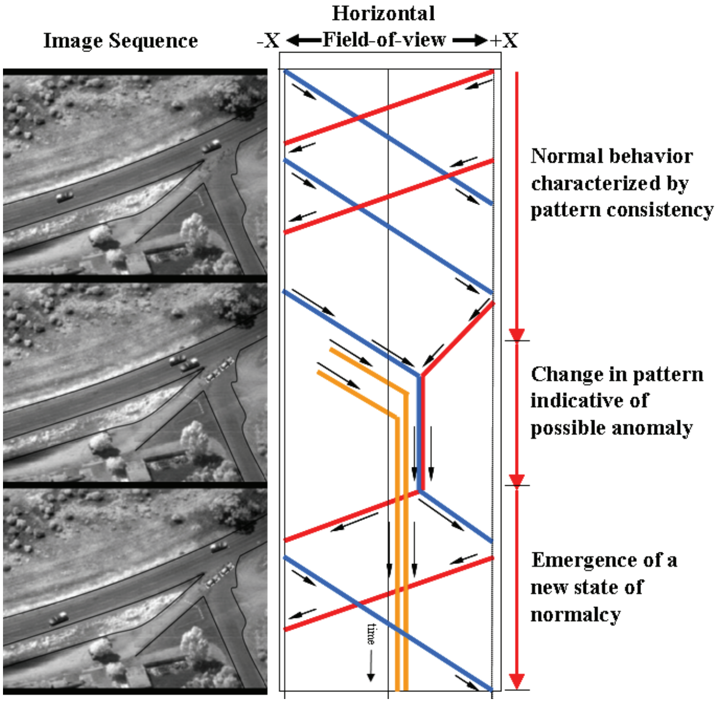
Image Sequence
Horizontal
-X Field-of-view +X
Normal behavior characterized by pattern consistency
Change in pattern indicative of possible anomaly
Emergence of a new state of normalcy
time

Figure 7: Phase portrait of the sequence of images showing object motion

based on time stamping or else use an average value for the time shift when individual sample information was unavailable. Instead, by using an adaptive delay parameter, each output sample value transparently shifts, according to its applicable delay, to form the pattern-forming data set. Training a neural network to adjust the sensitivity factor for maximum transparency therefore is an essential component to the automatic detection of anomalous behavior.

The strategy to formulate an automatic anomaly detection system requires a system identification approach to identify "normal" behavior of the observed system. Our approach is to employ a fuzzy-neural framework to identify the dynamical behavior of any unknown and uncertain system whereby the states of the unknown system can be estimated. This naturally yields a model reference of the unknown system dynamics Prasad et al (2004). Figure 7 illustrates a more detailed conceptual framework compared to Figure 2 presented earlier.

We assume that the behavior of the unknown system can be observed through an array of sensors comprising video, infrared and multispectral, electro-optical, synthetic aperture radar, MTI, and possibly other sensor systems. These observations

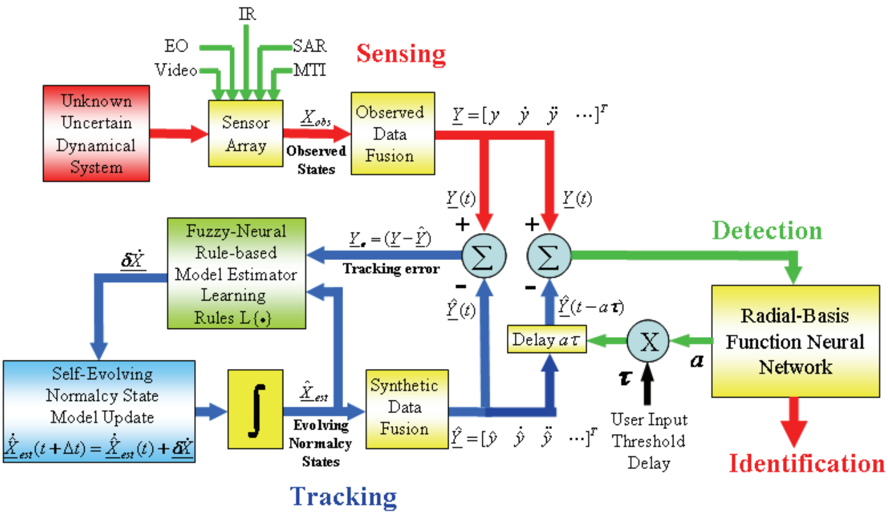
IR
EO
Video
SAR
MTI
Sensing
Unknown Uncertain Dynamical System
Sensor Array
$\underline{X}_{obs}$
Observed States
Observed Data Fusion
$\underline{Y}=[y \quad \dot{y} \quad \ddot{y} \quad \cdots]^T$
$\underline{Y}(t)$
$\underline{Y}(t)$
+
+
Fuzzy-Neural Rule-based Model Estimator Learning Rules L{•}
$\underline{Y}_e=(\underline{Y}-\hat{\underline{Y}})$
Tracking error
$\delta\dot{\underline{X}}$
Σ
Σ
Detection
-
-
$\hat{\underline{Y}}(t)$
$\hat{\underline{Y}}(t-a\tau)$
Delay $a\tau$
X
a
τ
Radial-Basis Function Neural Network
Self-Evolving Normalcy State Model Update
$\underline{\hat{X}}_{est}(t+\Delta t)=\underline{\hat{X}}_{est}(t)+\delta\dot{\underline{X}}$
∫
$\hat{\underline{X}}_{est}$
Evolving Normalcy States
Synthetic Data Fusion
$\hat{\underline{Y}}=[\hat{y} \quad \dot{\hat{y}} \quad \ddot{\hat{y}} \quad \cdots]^T$
User Input Threshold Delay
Identification
Tracking

Figure 8: Framework for an automatic anomaly detection system

correspond to some observed states of the unknown dynamical system. Position, rate, and other higher-order information may be derived by employing appropriate data fusion algorithms Klein (2007); Chatzigiannakis et al (2007); Goodman et al (1997).

Let $\underline{Y} = \begin{pmatrix} y & \dot{y} & \ddot{y} & \cdots \end{pmatrix}^T$ represent the vector of observations of the unknown dynamical system obtained as a result of data fusion.

For brevity, we can state that $\underline{Y} = C\underline{X}_{obs}$ is the observed behavior of the unknown system dynamics, where C is the matrix of data fusion coefficients that combine the state observations. The data fusion coefficients are primarily similarity coefficients which provide relative measures of similarity between states observed from each of the sensor measurements Salim et al (2003).

Background clutter can have significant effect in masking the observed behavior and is an important consideration in formulating the fusion coefficients.

We postulate the state model for the observer system dynamics as an autonomous system of the form $\dot{\hat{\underline{X}}}_{est} = f(\hat{\underline{X}}_{est}, t)$ where $\hat{\underline{X}}_{est}$ is the estimated state vector of the unknown dynamical system. The state estimates implicitly include any latency in the sensor behavior relative to each other in the modeled system dynamics.

In principle, the synthetic data fusion coefficients are similar to the data fusion coefficients that combine real observed data. So we let $\hat{\underline{Y}} = C\underline{X}_{est}$ represent the observer model response.

Ideally, if the observed output signal and the synthesized signal are approximately the same, namely,$\underline{\hat{Y}} \cong \underline{Y}$, then the estimated system states and observed system states are also approximately equal, i.e., $\underline{X}_{est} \cong \underline{X}_{obs}$, and the dynamics of the unknown system is fully identified. For this condition to be achieved within some desired tolerance, we need to minimize the error between the observed response and the modeled system response, i.e.,$\left(\underline{Y} - \underline{\hat{Y}}\right) \rightarrow \varepsilon$, a very small value.

Let $\underline{Y}_e = \left(\underline{Y} - \underline{\hat{Y}}\right) = \begin{pmatrix} y_e & \dot{y}_e & \ddot{y}_e & \cdots \end{pmatrix}^T$ represent the error vector to be minimized. We can now define each element of the error vector in terms of linguistic variables, for example small, medium, and large position error; small, medium, and large velocity error, etc., which serve as the subsets of fuzzy variables that can be used to develop a set of fuzzy rules.

Let $L\{\bullet\}$ represent a set of fuzzy "*If-Then*" learning rules that associate the elements of the error vector to obtain close estimates of the unknown dynamical system states.

Considering the unknown dynamical system as an autonomous system, RuleL_i for example may be written in Sugeno form as:

$$If\ y_e \in S^i_{1k},\ and\ \dot{y}_e \in S^i_{2l}\ and\ \ddot{y} \in S^i_{3m}\ \cdots,\ Then\ \delta\dot{x}_1 = a^i_{11}x^1_{est} + a^i_{12}x^2_{est} + \cdots$$

where the consequent is expressed as a time rate of change of the estimated states. The coefficients $\{a^i_{11},\ a^i_{12},\ \cdots\}$ are obtained from linear regression.

From classical control theory it can be easily recognized that the coefficients $\{a^i_{11},\ a^i_{12},\ \cdots\}$ are the row elements of a canonical form of the characteristic matrix. The time rate of change of the estimated states in top companion controller canonical form at time step "*j*" is given by the autonomous system:

$$\underline{\delta\dot{x}} = \begin{bmatrix} \delta\dot{x}_1 \\ \delta\dot{x}_2 \\ \delta\dot{x}_3 \\ \vdots \\ \delta\dot{x}_n \end{bmatrix} = \begin{bmatrix} a^j_{11} & a^j_{12} & a^j_{13} & \cdots & a^j_{1n} \\ 1 & 0 & 0 & \cdots & 0 \\ 0 & 1 & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \cdots & \vdots \\ 0 & 0 & 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ \vdots \\ x_n \end{bmatrix}$$

From the canonical representation above we observe that it is sufficient to obtain the derivative of one of the n states as illustrated in Rule L_i, and to directly formulate the derivatives of the remaining (n-1) states according to the canonical form. The approximation provided by the rule consequent is equivalent to the piecewise linear approximation of a nonlinear system at an operating point.

The advantage of this approach is that due to the piecewise linear approximation, the eigenvalues of the linear system can be computed at each time step and used as an indicator to track the stability of the self-evolving model.

The consequent part in Rule L_i, is an estimated deviation in state that updates the state derivative of the observer model. A recursive relationship that yields new estimates of state derivatives is of the form: $\underline{\dot{\hat{X}}}_{est}^{new} = \underline{\dot{\hat{X}}}_{est}^{old} + \underline{\delta\dot{x}}$. Integration provides the estimated states $\underline{\hat{X}}_{est}$ of the unknown dynamical system.

Near real-time anomaly detection is conceivable if a self-tuning system is integrated within the tracking system to automatically zoom-in to the beginning of an abrupt change. The best performance of any detection scheme is the speed and sensitivity in response. An approach to optimize the sensitivity of information sources towards identification of anomalous behavior can be developed using a radial basis recurrent neural network that will automatically determine the delay between the present "normalcy" state and the previous "normalcy" state. The purpose of a radial-basis recurrent neural network (RBRNN) is to adaptively determine the maximum time-delay so the instant of occurrence of an anomaly can be captured Tan (2004); Kokkinos and Maragos (2005); Sarimveis et al (2002). This is effectively a "zoom" function that is used to monitor the sensitivity of observed information to change.

By entering a threshold parameter $0 < \tau \leq 1$, and based upon the RBRNN initialization, the sensitivity parameter output "**a**" of the RBRNN is used to generate a delay "**a**τ" such that the tracking signal $\underline{\hat{Y}}(t)$ can be delayed by a maximum amount with respect to the observed state motion signal $\underline{Y}(t)$. Figure 8 shows the region of RBRNN performance relating the signal $U_s(t) = \underline{Y}(t) - \underline{\hat{Y}}(t - a\tau)$ versus the sensitivity parameter "**a**".

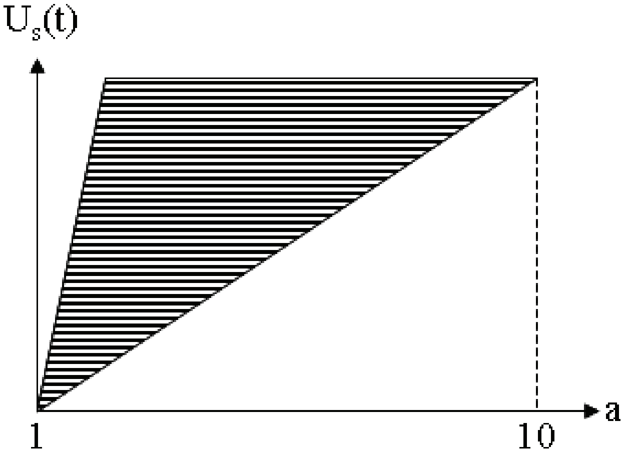

Figure 9: RBRNN performance region

Because of the feedback structure of the recurrent neural network, the sensitivity parameter is rapidly adjusted to a value which makes the amplitude of the signal $U_s(t)$ generated reach a maximum. If the signal characteristic suddenly changes, the existence of an anomaly may be suspected, and appropriate operator alerts can be generated. The range of "**a**" chosen in Figure 8 is only for illustrative purposes and is suitable for fast dynamics such as the motion of automobiles or other fast moving objects. For slow dynamics, such as tracking the motion of humans the range may be extended on a logarithmic scale to several hundred seconds.

Estimating the maximum value of the sensitivity parameter "**a**" therefore enables the identification of the anomaly at the instant of occurrence. This is shown graphically in the following section wherein the pattern signal generated from state motion looses its normal structure and shows the tendency to transform into a chaotic signal.

4.1 Graphical interpretation of anomaly detection

We consider the example scenario discussed earlier to provide a graphical interpretation of anomaly detection. This is illustrated in Figure 7.

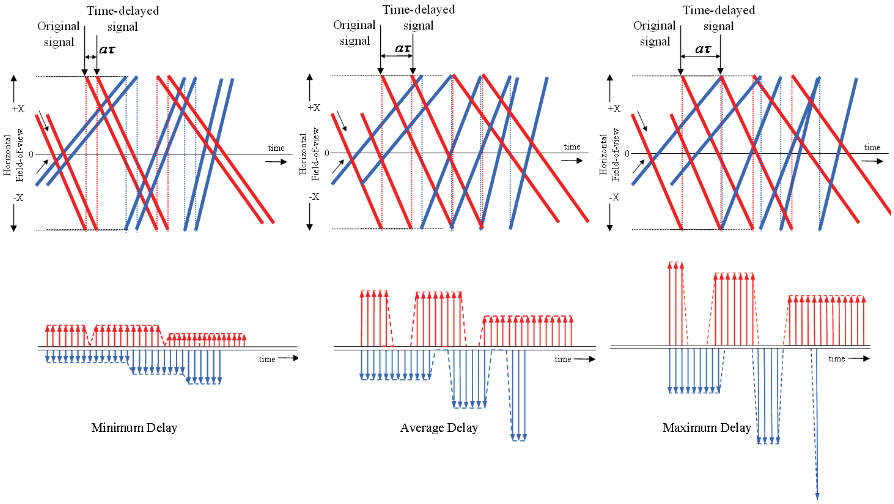

Figure 10: Effect of time-delay on normalcy pattern generation and transition to anomalous behavior

Prior to detecting any anomalies in a physical process, the idea is to first obtain a very effective tracking system that will keep the tracking error to a minimum. If the actual dynamics is then compared to a delayed tracking signal the time-shifted error signal provides “signatures” for anomaly detection.

Figure 7 illustrates three cases of pattern signal generation as a result of obtaining the difference between an observed signal and a delayed tracking signal. In the top part of the figure, blue and red lines are use to represent the trajectories of objects moving from left to right and from right to left, respectively, at any instant of time in a given field-of-view. The slope of the red or blue lines represents the velocities of moving objects. In the bottom part of the figure the difference between the actual and delayed pattern signals are plotted for cases involving minimum delay, average delay, and maximum delay.

The example shows pattern evolution as the time delay is increased from a minimum value to a maximum value. When the delay is at a minimum, the signal has low amplitude. When the delay is at a maximum, the difference shows the tendency towards transformation into chaotic outputs (note the blue signal spiking). The example shows amplification of the difference signal as the delay is increased. A sudden transition in pattern will occur when the delay is increased beyond its maximum. The time-series plot of object motion dynamics therefore is a key to determining anomalous behavior. The behavior is analogous to stretching a rubber band to the maximum until it snaps.

5 Conclusions

Identifying anomalies as and when they occur provides a means to understand the complex dynamics of the observed non-stationary stochastic process, and aids in developing adaptive models that can be used as predictive tools. Identifying anomalous behavior is based upon a perception of what constitutes abnormal behavior in the observations and how such abnormalities can be spotted within the normal occurrences of events.

We have discussed a highly innovative, simple and yet elegant systems-theoretic concept which supports closed-loop sensor-based monitoring and control of the environment. To enable the application of such an approach it is postulated that anomalies are synonymous to singularities in chaotic/nonlinear systems. Singularities exhibit trajectories that are characteristic of the type of behavior known to exist in nonlinear systems. The objective is to use state motion of objects as a means to map their behavior.

Mapping the trajectories of objects in an image sequence has the potential to uncover unique patterns that can be used to distinguish between normal and abnormal behavior and to correlate these patterns to possible anomalous behavior. The concept provides a revolutionary framework for the characterization, detection, identification, and modeling of anomalous behavior. The framework is adaptable to various applications including urban ground surveillance using combinations of ground, air and Space-based radar, electro-optical and infrared images, and MTI systems.

The human in the loop is a very important consideration as it provides a “filter” towards eliminating spurious machine identification of anomalous behavior. In a sense this provides a platform that incorporates both supervised as well as unsupervised training/learning capabilities. Human-machine interaction should be used to enhance the quality of performance and the reliability of models developed for classification and prediction of anomalous behavior.

Much work remains to be addressed pertaining to the issue of background clutter and its effect on the precision of data extracted. While in principle object motion can be determined and tracked, testing the viability of the proposed approach to human motion will greatly add to the value of the proposed approach. This can be very useful in monitoring and detecting how, when, and where insurgents physically plant IEDs under concrete sidewalks, inside garbage dumps, and other locations where there are masses of ordinary citizens who are extremely vulnerable.

To the best of our knowledge there is no public-domain literature that discusses automatic anomaly detection for possible roadside IEDs. Available literature that discuss anomaly detection ultimately rely upon humans to identify anomalous behavior. The proposed approach is therefore a radical departure from current technologies and is a novel approach using machine intelligence for automatic anomaly detection.

From a homeland security viewpoint there is a need for the detection of anomalies along international borders and border checkpoints, in and around air/water/land transportation centers, within large metropolitan shopping centers, in the vicinity of federal, state, and local government establishments, and around a wide range of other sensitive and crucial infrastructure establishments. Implementing real-time anomaly detection systems aid in preserving national security and protect the society at-large against threats to human life and safety. Detecting anomalous activities therefore is the first step towards mitigating a clear and present danger of terrorist threat to the Nation and for the safety of the US warfighter.

A soft computing approach clearly helps bio-mimic human decision-making abilities in detecting anomalous behavior. The approach exploits the innate human abilities for logical reasoning, learning, and adaptation to distinguish between normal and abnormal behavior in observed phenomena. From a computational standpoint this literally translates to a fuzzy-neural-evolutionary architecture.

We have discussed the qualitative aspects of anomalous behavior in a class of dynamical systems for which the mathematical model is unknown, and which exhibit high levels of uncertainty in their behavior. The uncertainty lies in variables and coupled-variable interactions that constitute the system dynamics. In this context we have outlined a revolutionary, and a viable systems-theoretic approach towards developing an intelligent anomaly detection system.

Acknowledgement: Work reported here was funded in part by the Defense Threat Reduction Agency under a University Associated Research Partnership Program.

References

Chatzigiannakis, V., Androulidakis, G., Pelechrinis, K., Papavassiliou, S., Maglaris, V. (2007): Data fusion algorithms for network anomaly detection: classification and evaluation, Third International Conference on Networking and Services (ICNS'07).

Fennell, M. T., Wishner, R. P. (1998): Battlefield Awareness Via Synergistic SAR and MTI Exploitation *IEEE AES Magazine*, Vol. 13, No. 2.

Gentile, C. (2001): Robust Tracking with Parts in the Presence of Severe Occlusion, Ph.D. Thesis, The Pennsylvania State University.

Goodman, I. R., Mahler, R. P. S., Nguyen, H. T. (1997): Mathematics of Data Fusion, Kluwer Academic Publishers Group, Dordrecht.

He, Z. (2005): Unusual activity detection for persistent target surveillance, *Proc. SPIE*, vol. 5778, pp. 945-952.

Jang, J.-S. R., Sun, C.-T., Mizutani, E. (1996): Neuro-Fuzzy and Soft Computing: A Computational Approach to Learning and Machine Intelligence, Prentice Hall.

Klein, L. A. (2007): Sensor and Data Fusion: A Tool for Information Assessment and Decision Making, SPIE, PM 138.

Kokkinos, I., Maragos, P. (2005): Nonlinear speech analysis using models for chaotic systems, *IEEE Transactions on Speech and Audio Processing*, Vol. 13, No. 6, pp. 1098-1109.

Li, D., Zhang, Y., Sun, Y., Yan, W. (2008): A multi-frame particle tracking algorithm robust against input noise, *Meas. Sci. Technol.*, 19.

Nguyen, H. T., Prasad, N. R. Walker, C. J., Walker, E. A. (2002): A First Course in Fuzzy and Neural Control, Boca Raton, CRC Press.

Prasad, N.R., Rangamani, A., Ross, T. J., Taha, M. R. (2005): Machine Intelligence in Decision-making (MIND) Automated Generation of CB Attack Engagement Scenario Variants, *Proceedings of 2005 S&T for Chemical & Biological Information Systems Conference*, Albuquerque.

Prasad, N. R., DiVita, J. Morris, R. (2004): A Systems Approach to Task Prioritization in Complex Dynamical Systems. *InTech'04, Conference on Intelligent Technologies*, Houston, TX.

Prasad, N. R., Ross, T. J., Taha, M. R. (2007): Mining for resource effectiveness to mitigate Chem-Bio attack consequence, *Proceedings of 2007 Chemical & Biological Information Systems Conference*, Austin, TX.

Rangamani, A. (2006): *Machine Intelligence in Decision-making* (MIND): *Mitigating Threats from Chemical and Biological Attacks*, MSEE Thesis, New Mexico State University.

Salim, N., Holliday, J., Willett, P. (2003): Combination of Fingerprint-Based Similarity Coefficients Using Data Fusion, *J. Chem. Inf. Comput. Sci.*, 43 (2), pp 435–442.

Sarimveis, H., Alexandridis, A., Tsekouras, G., Bafas, G. (2002): A fast and efficient algorithm for training radial basis function neural networks based on a fuzzy partition of the input space, *Industrial Engineering Chemistry Research*, vol. 41, pp. 751-759.

Tan, Y. (2004): Time-varying time-delay estimation for nonlinear systems using neural networks, *Int. J. Appl. Math. Comput. Sci.*, Vol. 14, No. 1, 63–68.

Zadeh, L. A. (2002): From Computing with numbers to computing with words - From manipulations of measurements to manipulation of perceptions, *Int. J. Appl. Math. Comput. Sci.*, Vol.12, No.3, 307–324.