

2009 Special Issue

Coordinated machine learning and decision support for situation awareness

N.G. Brannon^a, J.E. Seiffert^{b,*}, T.J. Draelos^c, D.C. Wunsch II^b^a Reliability Assessment and Human Systems Integration Department, Sandia National Laboratories, Albuquerque, NM 87185, USA¹^b Applied Computational Intelligence Laboratory, Department of Electrical and Computer Engineering, Missouri University of Science and Technology, Rolla, MO 65409, USA^c Cryptography and Information Systems Surety Department, Sandia National Laboratories, Albuquerque, NM 87185, USA

ARTICLE INFO

Article history:

Received 24 January 2009

Received in revised form 6 March 2009

Accepted 13 March 2009

Keywords:

Neural networks

Situation awareness

Reinforcement learning

Adaptive resonance theory

ABSTRACT

Domains such as force protection require an effective decision maker to maintain a high level of situation awareness. A system that combines humans with neural networks is a desirable approach. Furthermore, it is advantageous for the calculation engine to operate in three learning modes: supervised for initial training and known updating, reinforcement for online operational improvement, and unsupervised in the absence of all external signaling. An Adaptive Resonance Theory based architecture capable of seamlessly switching among the three types of learning is discussed that can be used to help optimize the decision making of a human operator in such a scenario. This is followed by a situation assessment module.

© 2009 Elsevier Ltd. All rights reserved.

1. Introduction

Modern information sources to support decisions in domains such as force protection are diverse. Ground, air, and space based sensors are continuing to increase in capability. Information fusion algorithms can help integrate a variety of sensor data into meaningful forms (Hall & Llinas, 1997). Applications with a complex assortment of data continue to challenge machine learning approaches to information fusion, which normally utilize a single type of learning algorithm and therefore limit the use of all available data (Brannon, Conrad, Draelos, Seiffert & Wunsch, 2007). Our approach coordinates multiple learning mechanisms to accommodate environments where ground-truth and feedback may not be consistently available and it uses Adaptive Resonance Theory (ART) based networks which are based on understanding cognition. This ties the work into other such computational architectures seeking not only the solution of engineering problems but also understand the function of the brain and mind (Werbos, 2009; Perlovsky, 2009).

1.1. Machine learning

Machine learning involves programming computers to optimize a performance criterion using example data or past experience (Alpaydin, 2004). Artificial neural networks are commonly

used in machine learning and utilize supervised, unsupervised, and reinforcement learning approaches to achieve predictive properties based on example (training) data. Unsupervised learning (clustering) can be effective when ground-truth is not available with a dataset. Supervised learning (learning with a teacher) provides a means to use experience (examples with ground-truth) to correctly classify yet unseen situations. Reinforcement learning offers promise for machine learning in difficult learning environments by taking advantage of feedback about the performance of a system. The challenge addressed by the current work is to coordinate all of these learning mechanisms and utilize the appropriate one based only on available information, not human intervention.

Neural nets offer an excellent assortment of high-performance, low-cost, distributed processing options. In particular, they can be embedded into appropriate sensors for operation at the lowest levels of information fusion with effective, but low-complexity designs. At the highest levels of information fusion and situation assessment, reinforcement learning can be used with a human in the loop to provide operational feedback. Dealing with multiple sensor modalities and extracting meaningful information from massive datasets is a natural fit for these adaptive methods. Although neural networks have been applied to sensor fusion, their use in situation awareness has been limited, possibly because of the lack of rich training data for this problem.

Automated (computational) information fusion continues to suffer from very specific, ad hoc solutions (i.e., there appears to be very little general-purpose technology to apply to this problem) (Kokar, Tomasik, & Weyman, 2004). For many applications, there is also a dearth of data to use for training a computational

* Corresponding author.

E-mail address: seiffert@ieee.org (J.E. Seiffert).

¹ Sandia National Laboratories is a multi-program laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under Contract DE-AC04-94AL85000.

engine. This reveals a challenge for the application of machine learning techniques, which are data-driven and require training—via supervised, unsupervised, or reinforcement learning. On the other hand, because they are data-driven, the advantage of machine learning techniques is that they can learn solutions to problems that are difficult for humans to codify with explicit rules or models. In other words, they can represent rules/decisions that are implicit in the training data.

1.2. Information fusion

The fusion of information has been likened to the ability of animals to utilize multiple senses to derive a better understanding of a situation (Hall & Llinas, 1997). For example, one may hear a noise and, based on the sound pressure discrepancy between each ear, localize the area of the sound source. Vision can then be used to further define and understand the source of the sound. The analogy is helpful because fusion, and more generally situation assessment, is a process rather than simply a discrete event. The process leads one from raw data to understanding and actionable knowledge. Fusion can occur over various information (sensor) modalities, over geographic space, and over time.

The sources of information potentially available to decision makers continue to expand in depth and breadth. Sensor capabilities in particular are maturing rapidly, but a valid concern is that the pace of sensor development has not necessarily been consistent with advances in human effectiveness which the sensors must ultimately support (Paul, 2001). Fusion algorithms will better support human-in-the-loop system effectiveness when the decision maker is a central and balanced design element.

2. Approach

2.1. System architecture

The design of the computational engine for information fusion and situation awareness takes advantage of the diverse utility of neural networks and integrates elements of supervised, unsupervised, and reinforcement learning. The design not only advances machine learning research, it addresses the need of situation awareness and human-in-the-loop decision support as well.

Key design attributes of our system include accepting various inputs such as binary, categorical, and real-valued data. With respect to situation assessment outputs, attributes include confidence levels as well as evidence in support or against the assessment. In the context of missing or noisy inputs, the system exhibits graceful performance degradation.

In order to address the desired design attributes of our situation awareness system, neural networks are employed for information fusion, followed by a situation assessment module. ARTMAP is based on Adaptive Resonance Theory (ART), a widely implemented approach to modeling the learning capabilities of the brain (Carpenter & Grossberg, 1988). Architectures based on ART have been used successfully in a variety of areas requiring a self-organizing pattern recognition neural network. The basic ART element supports unsupervised learning and binary inputs. Fuzzy ART is an extension to accommodate categorical and real-valued inputs. ARTMAP supports supervised learning and can accommodate real-valued inputs using fuzzy logic (Carpenter, Grossberg, Markuzon, Reynolds, & Rosen, 1992). ARTMAP can also support reinforcement learning, for example, by adding a mechanism to implement actor–critic methods. Coordinated ARTMAP (CARTMAP) is the name given to the current approach and involves the integration of all three learning mechanisms in the same architecture.

The situation assessment module receives state information from the information fusion module and possibly other sources and outputs a threat assessment or action to be taken.

2.2. Information fusion engine

Intelligent creatures exhibit an ability to switch seamlessly among unsupervised, supervised, and reinforcement learning as the needs arise. However, machine learning architectures, including artificial neural networks, have not yet achieved this goal. The current research contends that it is advantageous to develop this capability in a computational framework and that the ART architecture is an excellent choice for such an implementation. In the following text, we motivate, design, and give examples of this capability.

A well-designed sensor fusion algorithm, like an intelligent creature, can make informed use of all three types of learning on the dataset given. Certain information fusion paths may be pretrained prior to deployment, thus granting the human operators license to verify that the most obvious sensor patterns will be classified successfully. During operation, a reinforcement signal provided either by the environment or by the human operator acting off of the fusion algorithm's recommendations can adjust the current adaptive weight profile to curtail or retrain a faulty clustering (negative reinforcement) or to promote successful clustering (positive reinforcement) in the ART algorithm. Finally, in the absence of any external signal, the algorithm will learn in an unsupervised manner, comparing current inputs to what it already knows.

Our core algorithm—ARTMAP—already handles both unsupervised and supervised learning problems. To augment this, we need to optimize

$$J(s) = r(s) + \gamma \sum_{s'} P(s', a) J(s', a). \quad (1)$$

This is the discounted expected reward optimality criterion, a form of the Bellman equation of dynamic programming. In this equation, s represents the current state of the system, $r(s)$ represents the current reward, a the action to be taken, $J(s)$ represents the current value of a given state, s' signifies the next states, $P(s', a)$ is the transition probability matrix for the system's evolution, and a discount factor, γ , is applied to future rewards. This equation is to be maximized over all actions.

The Bellman equation (1) states that the current value of a state is equal to the immediate reward of taking an action plus the discounted future reward that accrues from that state. Other optimality criteria are possible to account for infinite horizon or non-discounted models. The task at hand is to solve this equation given an appropriate reinforcement signal.

Sutton and Barto (1998) discuss a wide variety of solution methods for these problems. Our algorithm combines ART with Q-learning. The Q-learning algorithm iteratively updates the value of each state–action pair. The appropriate modification is calculated based on the difference between the current and realized valuations, when maximized over all possible next actions. This is a key concept that establishes the foundation for the more advanced techniques discussed in the following paragraphs.

The Q-learning algorithm utilizes a lookup table to store the Q-values for each state–action pair. As the scale of the simulation grows, the amount of memory required to catalogue these values can grow at a burdensome rate. A more computationally intensive but less memory-demanding version, called Heuristic Dynamic Programming, uses function approximators in place of the table (Werbos, 1990). However, for our purposes in this architecture, the Q-learning approach will suffice.

With the ARTMAP unit taking the place of the Actor in the actor–critic implementation, the Coordinated ARTMAP (CARTMAP) algorithm behaves according to the following steps:

1. Upon receipt of an unsupervised signal, the system uses its exemplar classification scheme (the ART unit) to output an action choice, as usual. No updating of the lookup table will be necessary.
2. When presented with a supervised signal, the internal adaptive weights updates as per our normal ARTMAP rules, and the output action is set equal to the supervised training signal. Furthermore, the values in the lookup table for actions not associated with the supervisory signal are zeroed out.
3. When a reinforcement learning input signal is received, it will be interpreted according to the Q-learning algorithm. The appropriate entry in the lookup table is augmented with the new reinforcement value, and the action selected is the one with the most values accumulated in its column of the table. In our simulations, the values of the parameters delta and gamma are 0 and 1, respectively (Puterman, 1994).

In summary, the information fusion engine accepts raw data from sensors and other information sources and processes/transforms/fuses them into inputs appropriate for the Situation Awareness Assessment engine.

The information fusion system utilizes appropriate elements of its architecture based on the data presented to it. The three ART networks are linked together by an inter-ART module (Associative Memory). One ART unit handles the inputs, another ART unit processes the supervisory (or target) signal, and the other processes the reinforcement signal as an adaptive critic. This architecture is capable of online learning without degrading previous input–target relationships.

There are times when unsupervised learning is satisfactory, such as in the presentation of new input vectors to a pretrained network. Supervised learning is appropriate and desired for initial training on fixed data. However, these two types of learning do not cover every possible complication. There are times when the human operator does not know the correct classification, yet some feedback on the decision can be provided. These situations fall into the reinforcement learning category. One aspect of developing this information fusion engine, therefore, is adding the reinforcement learning capability to the ARTMAP neural network.

3. Application

We designed our situation awareness system to operate in an environment involving distributed sensors and a central collection site for protection of a facility. Information sources in such an environment can include seismic, magnetic, acoustic, passive infrared (PIR), and imaging sensors as well as weather, time/day information, various intelligence information, local/regional/federal threat levels or law enforcement bulletins, and any other information that might be relevant to the security of a particular facility, such as current traffic situations or health issues.

Conditions of interest to force protection decision makers include: no activity, severe weather, unauthorized people or vehicles in certain locations, and certain types of unauthorized vehicles or humans with weapons in any areas. Actions include: doing nothing, identifying the type and location of a moving object (vehicle or human), commands to turn sensors on or off, dispatching forces, and/or notification of higher authorities. The information sources can include binary data, such as motion detection, categorical data, such as the type of day (weekend, holiday, etc.), and real-valued time-series data, such as seismic, acoustic, and magnetic energy levels.

Before being deployed, the system must be pretrained with information the human operator knows about the system. For example, if the data signature of a thunderstorm is easy to demonstrate (due to specific acoustic, magnetic, etc. levels), then that information can be included in the supervised training portion

of the system. The information fusion engine will adaptively learn many more data-observation relationships during online operation, but having basic readings pretrained will aid in initial operation.

When an intruder, be it an unauthorized vehicle or a human with a weapon, breaches the sensor range of a protected facility, the triggered sensor data stream into the information fusion engine. The CARTMAP network then maps these data into observations, such as a vehicle heading north at high speed. These pairings represent novel data readings that were not anticipated, which are then categorized via the CARTMAP algorithm in relation to the pretrained data.

The observation is then sent to the situation assessment engine, which follows the partially observable Markov decision process (POMDP) formulation to calculate a probability distribution over the state space. This information represents a confidence level that the system is in any given state. The state with the highest confidence from this calculation represents the system's choice for the current state. All this probability information is then passed to the human operator, who uses this evidence in making a final decision about how to respond to the situation.

Adapting online is an important element of the system and is accomplished through reinforcement signals that can be sent through the system in two ways. First, if the probabilities of each state are too low, so that the human operator would not be able to distinguish the state from simple background noise, then the situation assessment engine may issue a command to gather more information from additional sensors. Second, the human operator may disagree with the system's assessment of the current state. A reinforcement signal is then sent to the information fusion engine and the data-observation mappings will adapt online. Both of these reinforcement signal loops are noted functionally in the block diagram in Fig. 1. This feature of the system allows it to maintain relevance in a changing environment.

The operation of Fig. 1 is as follows. Unsupervised learning occurs using a single ART unit. The cluster that forms is the one which maximizes the signal strength of the input with respect to a match criterion. Many forms of both the signal and the match criterion are in use in various implementations of an ART architecture. Amis and Carpenter (2007) provide default values which work in general scenarios. Supervised learning occurs when the clusters formed by the unsupervised learning unit are given labels through interaction with supervisory inputs. This interaction is mediated by an associative learning field as explained in Amis and Carpenter (2007). This process forces a reset in the input cluster if the label does not match the supervisory signal closely enough. Finally, Reinforcement learning is handled in a similar manner. The RL signal can update the associate weights following the Q-learning explained in Section 2.2.

The CARTMAP algorithm was implemented in Matlab and applied to information fusion in a vehicle tracking scenario that is described in more detail below. ART is at the core of the fusion engine. During offline training, an input pattern is presented to the CARTMAP network and, depending on its similarity to existing category templates, it is either assigned to a current winning category or a new category is created for it. Categories may exist indefinitely without an assigned class. However, if a supervisory signal accompanies the input, the target class is immediately associated with the category. During offline reinforcement learning, an input pattern is presented to the CARTMAP network and a winning category is determined. A reinforcement signal is computed based on the class of the winning category and the ground-truth class. For example, if the category's class matches the ground-truth class, the reinforcement signal is assigned a positive reward; if not, then a penalty is assigned. A range of reinforcement values are assigned based on the quality

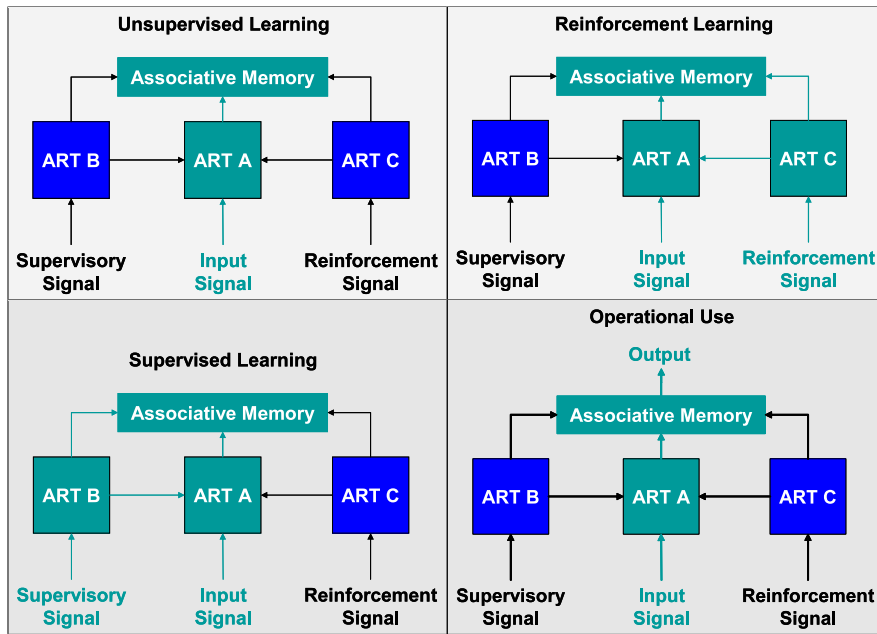


Fig. 1. CARTMAP input and system activity associated with unsupervised learning, supervised learning, reinforcement learning, and standard operational use. Available inputs to the system shown in green are the active elements involved in learning. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

of the match. A reinforcement lookup table (RLUT) is used to track input pattern’s relationship with possible classes. The RLUT stores input patterns and an accumulated reinforcement signal for each possible class. CARTMAP weights are updated according to the following criteria.

1. If no category encodes the input pattern, then a new category is created without a class assignment.
2. If the winning category has an unassigned class, then the RLUT is searched for the input pattern. If the pattern is found in the RLUT, then the reinforcement signal is applied to the class of the winning category and the class with the highest reinforcement is used as the target in supervised learning. If the pattern is not found in the RLUT, then nothing is done to the CARTMAP weights.
3. If the winning category has an assigned class, then this class and reinforcement signal are used by a critic function to determine how to update CARTMAP weights. The RLUT is searched for the input pattern. If the pattern is not found, unsupervised learning is performed and the pattern is added to the RLUT along with the reinforcement signal. If the pattern is found in the RLUT, then the reinforcement signal is applied to the RLUT for the class of the winning category and the class with the highest reinforcement is used as the target in supervised learning.

The decision support graphical user interface (GUI) consists of three screens. The center screen is primarily imagery (i.e., from cameras, photography augmented with graphics, and/or fully synthetic renderings) (see Fig. 2). The second screen displays a log of temporal track data (see Fig. 3). The log reflects temporal features, such as how long ago an unauthorized vehicle breached a sensor field and how soon another track might reach a key threshold (e.g., a fence or different sensor field). The third and most detailed screen provides track detail and assessment bases (see Fig. 4).

The log screen and track detail screen utilize features found in the Tactical Decision Making Under Stress (TADMUS) system (Morrison, Kelly, Moore, & Hutchins, 1997). The TADMUS system had similar motivations to the current research in that more content need to be devoted to supporting an understanding of

Fig. 2. Vehicle tracking scenario map. The dots represent seismic/acoustic sensor nodes. The speed, heading, location, and vehicle type are estimated by independent CARTMAP networks using binary data from all sensor nodes as input.

a given context. In both TADMUS and our situation awareness approach, less emphasis is placed upon evaluating possible courses of action.

The track detail GUI provides typical track parameters such as an object’s course and speed, but significant detail is provided with respect to the basis for assessment. Evidence in support and evidence against a given assessment is displayed. The machine learning algorithms share the evidence used to derive assessments with the operator. Such an approach provides greater transparency and allows the operator to interrogate assessments.

For the example scenario of an unauthorized vehicle, the assessment could be a “threat”. Evidence in support of such an assessment includes sensor data such as explosives detected, but also local law enforcement data that the license plate returns as a stolen vehicle. Evidence against the assessment could include a relatively slow speed and the use of the vehicle for construction when there has been ongoing construction activity. Alternative

