Chapter 3
An Application of Unified Computational Intelligence

3.1 Overview

The previous section described a unified computational intelligence learning architecture based on Adaptive Resonance Theory neural networks. In this chapter, this architecture is used in an application that was briefly introduced in Chapter 1.

The content of this chapter is adapted from a paper appearing in the *Neural Networks* journal (Brannan, Seiffertt, Draelos, & Wunsch, 2009) and a preliminary version appearing as (Brannan, Conrad, Draelos, Seiffertt, & Wunsch, 2006).

In this chapter, the unified computational intelligence algorithm is referred to as CARTMAP, for Coordinated ARTMAP. This name was determined by Sandia National Laboratories collaborators since they, being government officials, possess a certain *je ne sais quoi* for the use of acronyms. This application, in the area of situation awareness, proved a testing bed for a unified learning architecture of the type described in Chapter 2 of this book. The results indicate that the task described, if performed using only a single mode of learning, would not have achieved the same level of effectiveness as it did using all three modes in combination.

3.2 Introduction

Modern information sources to support decisions in domains such as force protection are diverse. Ground, air, and space-based sensors continue to increase in capability. Information fusion algorithms can help integrate a variety of sensor data into meaningful forms (Hall & Llamas, 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, Seiffertt, & Wunsch, 2006). Our approach coordinates multiple learning mechanisms to accommodate environments where ground-truth and feedback may not be available consistently, and it uses Adaptive Resonance Theory (ART)-based networks, which are based on understanding cognition. This ties the work into

J. Seiffertt & D.C. Wunsch: Unified Computational Intell. for Complex Sys., ALO 6, pp. 33–48.
springerlink.com

other such computational architectures seeking not only solutions to engineering problems but also an understanding of the function of the brain and mind as discussed by Werbos, Perlovsky, and others.

3.2.1 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 within a dataset. Supervised learning (learning with a teacher) provides a means of using 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 a system's performance. 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 networks 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 engine. This reveals a challenge for the application of machine learning techniques, which are data-driven and require training, whether 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.

3.2.2 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. Our system includes, as a core design principle, the use of a human-in-the-loop operator to provide reinforcement signals as well as to ensure a level of quality control.

3.3 Approach

3.3.1 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, but also addresses the needs of situation awareness and human-in-the-loop decision support.

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 of 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.

3.3.2 Information Fusion Engine

Intelligent creatures exhibit an ability to switch seamlessly among unsupervised, supervised, and reinforcement learning as needed. 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.

A well designed sensor fusion algorithm, like an intelligent creature, can make informed use of all three types of learning on the data set given. Certain information fusion paths may be pre-trained 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.

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 update as per 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 value accumulated in its column of the table. In the simulations, the values of the parameters delta and gamma are 0 and 1, respectively.

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 pre-trained 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.

3.4 Application

The situation awareness system was designed 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), using commands to turn sensors on or off, dispatching forces, and/or notifying 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 pre-trained with any 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 pre-trained 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 pre-trained 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 of 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 so low 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 Figure 3.1. This feature of the system allows it to maintain relevance in a changing environment.

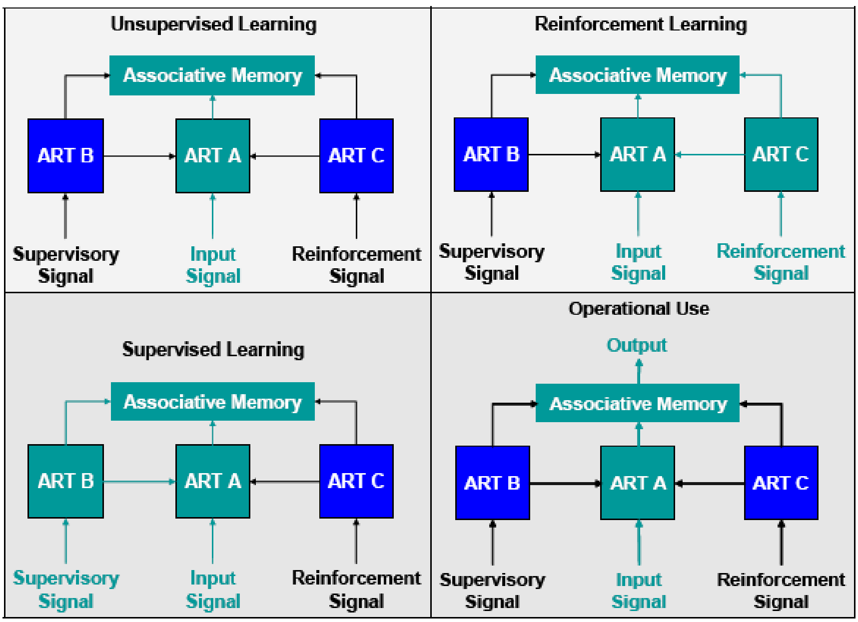

Fig. 3.1 CARTMAP Input and System Activity. Associated with unsupervised learning, supervised learning, reinforcement learning, and standard operational use. Available inputs to the system are shown in green, as are the active elements involved in learning

As shown in Figure 3.1, unsupervised learning occurs using a single ART unit. The cluster that forms is the one that maximizes the signal strength of the input with respect to a match criterion. Many forms of both the signal and the match criterion are used in various implementations of an ART architecture. Amis and Carpenter (2007) provide default values that 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 3.3.2.

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 off-line 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 of the match. A reinforcement lookup table (RLUT) is used to track an 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 Figure 3.2). The second screen displays a log of temporal track data (see Figure 3.3). 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 Figure 3.2).

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 has motivations similar to the current research in that more content needs to be devoted to supporting an understanding of 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 of and 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 such as the license plate returning 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 assessments are shown along with their respective evidence in support of or against them.

The operator can investigate various assessments along with corresponding courses of action. For example, a patrol vehicle in the vicinity of the unauthorized vehicle could be directed closer to the possible threat. Further, other types of sensors can be activated to generate additional points of reference and work towards higher levels of assessments, such as possible intent.

3.4.1 Vehicle Tracking

The situation awareness technology was applied to tracking vehicles in the vicinity of a facility under force protection. A data set suitable for testing and demonstrating our technology was collected during a DARPA SensIT program in November, 2001 at Twenty-Nine Palms, CA and exists at the University of Wisconsin (UW) (Duarte & Hu, 2004). The data set consists of raw time series (acoustic and seismic) and binary detection decisions from 23 sensor nodes distributed along three intersecting roads as one of two vehicles travels along a road. Figure 3.2 includes a map illustrating the force protection scenario, with a fence line and an Entry Control Point (ECP) providing protection for a facility on the North Road. The two vehicles used in the scenario are a light armored vehicle (AAV) and a heavier, tracked transport vehicle (DW). A scenario was developed whereby a facility under protection is assumed to exist along one of the roads, and binary sensor data processed by our fusion and situation assessment algorithms are used to inform a human decision maker.

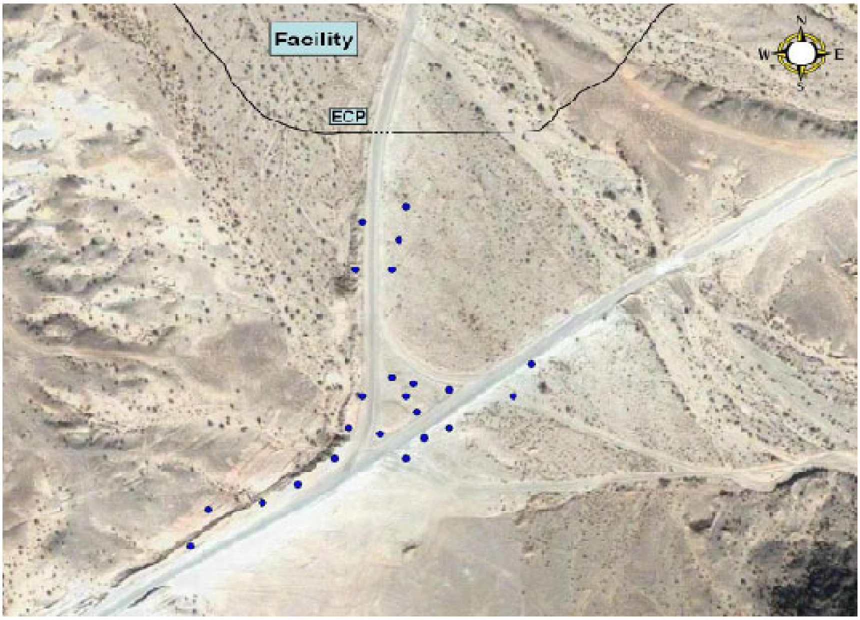

Fig. 3.2 Vehicle Tracking Scenario Map. Blue dots are seismic/acoustic sensor nodes. The speed, heading, location, and vehicle type are estimated by independent CARTMAP networks using binary data from all sensor nodes as input

3.4.2 Analysis

This section provides analysis of the experimental results.

3.4.2.1 Force Protection Experiments

In order to demonstrate the capabilities of the situation awareness system, neural networks were trained to perform sensor fusion, a situation assessment formula was constructed/calculated, and a GUI was developed, all to increase the awareness of a human decision maker of the situation around the facility under their protection. The scenario consists of a virtual checkpoint partway up the north road on the way to a sensitive facility with 23 sensor nodes scattered along three intersecting roads. Each sensor node outputs a binary detection decision at fixed time intervals (0.75 seconds in the original test set). The sensor detections derive from seismic, acoustic, and passive infrared energy levels. The (AAV and DW) vehicles move from one end of a road, through the intersection, and to the end of another road. The total number of runs is 40, which includes 20 original data sets from the SensIT experiment. An additional twenty runs were created by artificially reversing the direction of the vehicle. This is possible by simply presenting the data in reverse. In other words, the sensor record from the last time step would be presented to the information fusion system first, the first time step would be presented last, and so on for all the time steps in the run. It is plausible that the information is accurately represented in these runs because the data consists of binary decisions and the ground is relatively flat, so the engine speed and noise are presumably similar in both directions.

The primary piece of information that a decision maker wants to know is the current threat level around his facility. The threat level is a function of the location, speed, heading, and type of vehicle detected by the sensor array and other variables that are independent of the sensor array, such as Department of Homeland Security (DHS) advisory level, wind speed, average batter level of the sensors, time of day, and day of week.

The system used to produce the threat level is illustrated in Figure 3.3. The system consists of three modules: 1) Information Fusion, 2) Situation Assessment, and 3) a Graphical User Interface (GUI) focused on human decision makers in force protection applications. Multiple time steps of binary sensor data serve as input to the Information Fusion module, which implements the CARTMAP algorithm. This introduces an element of relative time, which is a necessary component in estimating speed and heading. The output from the Fusion module consists of vehicle type, speed, location, and heading, each with a corresponding confidence level, and will serve as input to the Situation Assessment module. This module consists of rules that represent the conditions under which a threat is defined. The output of the assessment module will feed the graphical user interface (GUI) with a threat level (low, medium, high), an associated confidence level, a suggested response, and evidence in support of or against its output. The GUI will also have access to the output from the fusion module, maps, and other available data, such as time, date, and environmental data. All elements of the situation awareness system were implemented in Matlab and tested with the vehicle tracking data from UW in the force protection scenario just described.

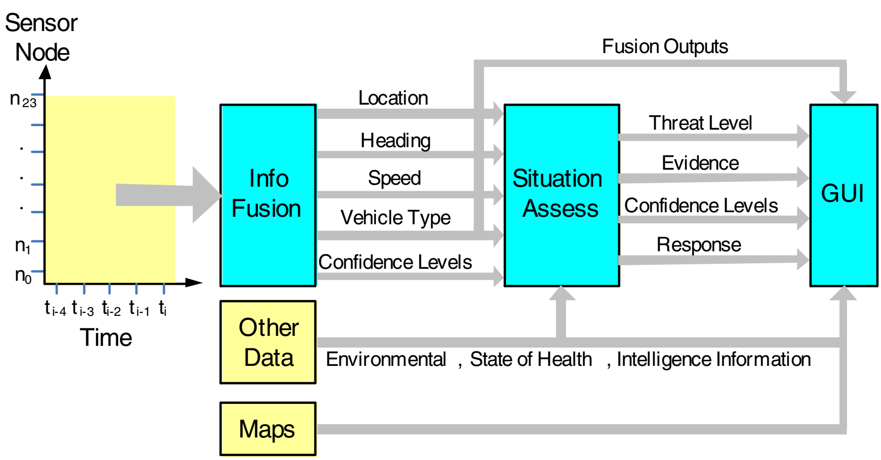

Fig. 3.3 Force Protection Experiment Using UW Vehicle Data. Multiple time steps of binary sensor data are used as input to the CARTMAP Information Fusion module. Vehicle information from the Fusion module and other additional data are used as input to the Situation Assessment module, which outputs actionable information to the user

3.4.2.2 Results of Training the Fusion Model

The fusion model consists of four different CARTMAP networks, one for each fusion output (location, heading, speed, and vehicle type). The output of a network will be of a categorical type or class except for the confidence levels, which will be real numbers. Table 3.1 presents the classes for each information fusion network. Note that for each network, if the input is all zeros, the output will be zero by virtue of a simple fixed rule (i.e., no learning is involved).

Out of the 40 total runs available for the force protection experiments, 70% were used for training and the remainder for testing. Table 3.2 shows the number of runs used in the six experiments. In real-world applications, it is expected that the amount of supervised training data is limited. In the force protection experiments, only 2 of the 28 training runs are used for supervised learning.

Table 3.1 Information Fusion Output Classes for the Four CARTMAP Networks. (Vehicle Type, Location, Heading, and Speed)

Vehicle Type Classes	Location Classes	Heading Classes	Speed Classes
0: zero input 1: AAV 2: DW	0: zero input 1: West Road 2: North Road 3: East Road 4: Intersection	0: zero input 11: N 14: NE 13: E 8: SE 4: S 1: SW 2: W 7: NW	0: zero input 1: < 10 km/hr 2: 10-20 km/hr 3: 20-30 km/hr 4: 30-40 km/hr 5: 40-50 km/hr 6: 50-60 km/hr 7: 60-70 km/hr 8: 70-80 km/hr 9: 80-90 km/hr 10: > 90 km/hr

Table 3.2 Distribution of Vehicle Runs Used to Experiment with Different Learning Modes. Experiments 1/2/3 and 4/5/6 use the same data, but use learning modes in a different order

Experiment #	# Supervised Runs	# Unsupervised Runs	# Reinforcement Runs	# Test Runs
1 & 4	2	26	0	12
2 & 5	2	13	13	12
3 & 6	2	0	26	12

Experiments 1-3 use the same runs as Experiments 4-6, but the order of training is reversed. In Experiments 1-3, supervised learning is conducted first, followed by reinforcement learning, and finally unsupervised learning. Experiments 4-6 use the opposite order of learning, using the data with the least amount of information first and finishing with supervised learning, which utilizes training data with the most information. In this case, one expects the richer data sets and training modes to correct errors and refine the classification performance of previous learning modes.

For each force protection experiment conducted, the same test set was used, consisting of 12 runs with 1755 input/output pairs. The performance (% correct classification) was computed based on this test set. For some sensor modes, such as speed and heading, a classification error may not necessarily indicate poor performance. For example, if the ground truth heading of a vehicle is North and the fusion module output is Northeast, it would be counted as a classification error even though the output is quite satisfactory. Experiments 1-6 were conducted using various combinations of learning modes for each of the information fusion networks. The best results for each network are presented in Table 3.3.

In the Classified Correct (%) column of the tables, there are three numbers separated by colons (e.g., 1 : 2 : 3). The numbers in position one represent the percentage of test samples that have a target value exactly matching the output value from a CARTMAP network.

The numbers in the second position represent the percentage of test samples that have a target value exactly or partially matching the output value from a CARTMAP network. An exact match increments the total number of correct classifications by 1, whereas a partial match increases the number by 0.5. Partial matches are possible only with the Heading and Speed networks, where the class adjacent to the target class is considered a partial match. For example, if the target class is N, then a network output of NW or NE would result in a partial match. Note that for the Vehicle Type and Location networks, no partial matches exist, so the first and second numbers in the Classified Correct column should be the same.

Table 3.3 The Best Fusion Test Results of the Four CARTMAP Networks. Reinforcement learning followed by supervised learning worked best for estimating vehicle type and location, while supervised learning followed by unsupervised learning and then reinforcement learning worked best for vehicle heading and speed. In the Classified Correct (%) column of the table, there are three numbers separated by colons (e.g., 1 : 2 : 3). The numbers in position one represent the percentage of test samples that have a target value exactly matching the output value from a CARTMAP network

Sensor Mode	Experiment #	Learning Mode	Vigilance	# Categories	Classified Correct (%)
Vehicle Type	6	Reinforcement Supervised	0.7 0.65	36 : 108 44 : 112	92.6 : 92.6 : 91.7 92.7 : 92.7 : 91.7
Vehicle Location	6	Reinforcement Supervised	0.7 0.65	22 : 58 31 : 61	96.8 : 96.8 : 98.0 96.9 : 96.9 : 98.0
Vehicle Heading	2	Supervised Reinforcement Unsupervised	0.9 0.7 0.7	39 45 : 59 45 : 59	68.4 : 69.6 : 69.6 66.6 : 79.9 : 80.3 62.7 : 75.8 : 81.7
Vehicle Speed	2	Supervised Reinforcement Unsupervised	0.9 0.7 0.7	46 53 : 77 53 : 77	72.4 : 79.9 : 79.9 74.3 : 82.1 : 82.0 73.4 : 81.3 : 81.8

The numbers in the third position represent correct classification percentages of networks that have had two passes through the training set. During the first pass, the reinforcement lookup table is updated during reinforcement learning. The updated table may be an advantage for second pass unsupervised and reinforcement learning. Correct classification percentages are computed using partial matches. Each network was trained using vigilance parameters that resulted in a reasonable number of categories.

In the next section, a weighted rule for determining the threat level of the situation awareness system is discussed. The rule combines the outputs of the fusion module and environmental conditions, and its output is categorized into High, Moderate or Low threat based on human judgment. Ground truth exists for the threat level, so performance of trained fusion networks with specified environmental conditions can be measured. Two environmental conditions are specified: 1) Benign – each environmental condition is set to its lowest value, and 2) Severe – each environmental condition is set to its highest value. For each of the learning modes, the correct classification percentage is measured against ground truth. The results are given in Table 3.4.

In practice, if only unlabeled data is available, then machine learning is typically not used at all. Machine learning is most often used when some labeled data are available and supervised learning is then used to its maximum extent, while other learning techniques are not employed. The advantage of using a variety of machine learning techniques is evident in Table 3.3 and Table 3.4 above, but a single set of networks (possibly a different network for each sensor mode) must be chosen since one cannot generally anticipate the environmental conditions. Table 3.5 summarizes the performance results of using the best combination of supervised (SL), unsupervised (UL), and reinforcement learning (RL) in comparison to the more common use of supervised learning alone. Table 3.6 lists the machine learning approaches used by each CARTMAP network to produce the best situation assessment threat level performance averaged over benign and severe environmental conditions.

Table 3.4 Best Test Results of Situation Assessment Threat Level Performance. Using a combination of learning modes under benign and severe environmental conditions. Different learning modes for different CARTMAP fusion networks are necessary to produce the best situation assessment results

Environment Condition	Vehicle Exp #	Location Exp #	Heading Exp #	Speed Exp #	Reinforcement Iterations	Classified Correct (%)
Benign	1	2	3	3	1	88.9
Benign	1	3	2	3	2	89.5
Severe	1	2	2	6	1	86.8
Severe	3	2	5	3	2	87.7

An important conclusion drawn from the experimental results is that utilizing multiple training approaches that can take advantage of additional and different data produces superior results for situation awareness compared to supervised training alone. The reason performance decreases with UL after SL is that with SL alone, all test patterns get encoded by a labeled category, whereas after UL, there are now unlabeled categories that may encode test patterns producing classification errors. Even though these unlabeled categories sometimes lowered the performance, they may eventually add value after subsequent labeling during SL or RL. Unsupervised input patterns that get encoded by existing categories with a class label can contribute to the quality of the category in representing the class in feature space. In addition, since the CARTMAP has access to a reinforcement lookup table (RLUT), if an unlabeled pattern matches a pattern in the RLUT, the corresponding class label from the RLUT can be assigned to the unlabeled pattern. This feature is used during unsupervised learning. Originally, the RLUT is generated from the supervised training data. It expands when new unlabeled patterns are encoded by categories with class labels and the pattern and its label are added to the RLUT.

Table 3.5 CARTMAP Fusion Performance Results. Using multiple machine learning modes in comparison to supervised learning alone

Learning Approach	Vehicle %	Location %	Heading %	Speed %	Avg. Threat %
SL	81.8	95.6	69.6	79.9	78.5
SL with UL and/or RL	92.7	98.0	81.7	81.9	87.6

Table 3.6 The Combination of Learning Approaches. The combinations that produced the best threat level performance. Three different combinations were used for the four different fusion modules (vehicle type, location, heading, and speed)

	Vehicle	Location	Heading	Speed
Learning Approach	SL, UL	SL, RL	SL, UL, RL	SL, RL
Reinforcement Iterations	1	2	2	2

Results for Experiment 3 (SL followed by RL) reveal a strong relationship between the hints that RL provides and partial matching in scoring the classification performance. When exact classification matches are required, hints may not be good enough. However, if a “close enough” match is sufficient, then RL hints improve performance. Even though multiple vigilance values were used in the force protection experiments, it is expected that performance will improve when the vigilance is optimized for the type of fusion mode and the type of learning.

It is important in RL to have data representing all classes that a network is designed to classify. If a class is not represented in the data, RL will not be able to establish a label for this class.

Vehicle location is the easiest piece of information to learn with binary sensor data. Location is inherent in the sensors themselves because their position is fixed.

Since 54.4% of the input patterns are all zeros, if a correct classification percentage of greater than 54.4% is achieved after UL only, then the reinforcement lookup table is being used to correctly label some patterns. During reinforcement learning, an input pattern is submitted to a network, and a reinforcement signal is generated. This signal offers negative or positive feedback on the output of the network. The following steps are taken at this point of reinforcement learning. When a reinforcement signal is received, the RLUT is updated, and SL is performed if the input pattern is found in the RLUT (the action associated with the input pattern with the highest value is used as the target). Unsupervised learning is performed if the input pattern is not found in the RLUT and the reinforcement signal is positive.

In general, SL should be used to create as many categories as possible within reason, while subsequent non-supervised training should take advantage of these existing categories and enrich them without corrupting them. The coordination of three machine learning modes therefore offers potential benefit from every sample of data available in an application.

3.5 Future Work

Arguably the most immediate area of future work lies in establishing principles and practices for employing the three learning modes. There are different ways of combining three modes of machine learning and many options for how and when to employ each mode. The current research offers a preliminary perspective on leveraging each learning mode for greatest system performance. It stands to reason that a CARTMAP network can be tailored for each information fusion mode (vehicle type, speed, heading, and location). The vigilance parameter may be different for each mode and may also require adjustment based on the type and ordering of the learning modes.

The core of our machine learning approach is an ART neural network. Other algorithms and architectures should be explored with the same goal in mind, that of integrating multiple learning modes. Reinforcement learning is a general area of research worth pursuing in the area of situation awareness where there is often not a clear win or lose outcome by which to measure success. There are also many ways of performing reinforcement learning, some closer to supervised learning, with stronger hints, and others that provide rare but consistent hints about the system's performance. How many iterations to use in reinforcement learning on this problem is a legitimate research question, as is how best to acquire feedback from human decision makers or the overall force protection system, either directly or indirectly.

Another avenue of future machine learning research is to explore the use of ensembles or bagging for supervised learning (Dietterich, 2000). The use of ensembles employs multiple "experts" that train the same network using a different sampling with replacement from the original supervised training set. The combination of the experts' solutions results in higher performance than the use of a single network trained on the original data set.

3.6 Conclusion

The coordination of the three major machine learning approaches in a single architecture, using ARTMAP at its core, is an innovation that should prove valuable in addressing real-world problems. Many domains offer a limited amount of information with ground truth that can be used with supervised learning algorithms. More available is data with hints from the environment that can be used with reinforcement learning. Almost always, data is available without labels that can be used with unsupervised learning. Allowing these three modes of learning to be used in the same framework is an important contribution. Interesting advantages emerge when these three approaches leverage one another. For example, reinforcement learning can utilize supervised learning when enough information about class labels is available from the environment. Unsupervised learning can take advantage of stored reinforcement learning information to go beyond mere clustering. There is potential for interplay between the learning modes that does not exist with a single mode.