Continuous Maritime Domain Awareness

an evolving maritime security solution

Abstract

Achieving continuous maritime domain awareness is an essential requirement of national defence and security. Maritime domain awareness is the understanding of activities that impact maritime security, safety, economy or environment. It enables quick threat identification, informed decision making and effective action by responding units. The core component of maritime domain awareness, persistent surveillance, involves multiple systems collaborating in order to detect, classify, identify, track and assess situations within an area of interest. In this white paper, we propose a solution to improve maritime domain awareness based on a sensor exploitation architecture capable of incorporating a vast array of data sources. The architecture addresses persistent surveillance challenges by combining well established concepts of information fusion with novel approaches to the problem. We first introduce the conventional techniques and their drawbacks, discuss the contemporary data ecosystem and present a potential solution which learns to closely match the dynamic internal structures present in the data. This solution, developed and patented by Larus Technologies, performs behavior analysis through predictive modeling, is capable of dealing with heterogeneous (i.e. multi-source, multi-sensor) data, is automated yet human-centric and resolves many of the challenges presented in maritime domain awareness.

Maritime Domain Awareness

In a world where more than 40% of the population lives within 100 kilometers of a coast [1] and where traditional and asymmetric threats to physical and cyber infrastructures and borders continue to rise each year, countries are becoming increasingly aware of the gaps that exist in their ability to achieve persistent surveillance and continuous awareness of their maritime domains. Persistent surveillance is an essential component in a global system to ensure *Territorial Security.* The latter is defined as the prevention, detection and response to unauthorized persons and/or goods crossing a physical or virtual perimeter, making this problem a security concern of individual, corporate, national and international scope.

In a vast and mostly uninhabited country such as Canada, which borders the Atlantic, Pacific and Arctic oceans, a major sector of Territorial Security is Maritime Domain Awareness (MDA), which provides awareness of potential threats from maritime approaches and cueing to military and interagency responders. MDA is defined as the situational understanding of activities that impact maritime security, safety, economy or environment [2]. MDA involves a system of people, processes and technological tools that discover, sense, analyse and react to events and perform physical and virtual defence of the country's borders. It includes the capture and storage of domain knowledge obtained along with the actions, effects and outcomes for use in planning future surveillance operations. The outcome expected from MDA is the effective tasking of joint and interagency forces to respond to offensive/illegal activities, disasters and rescue scenarios in the maritime domain. This complex process is depicted in Figure 1.

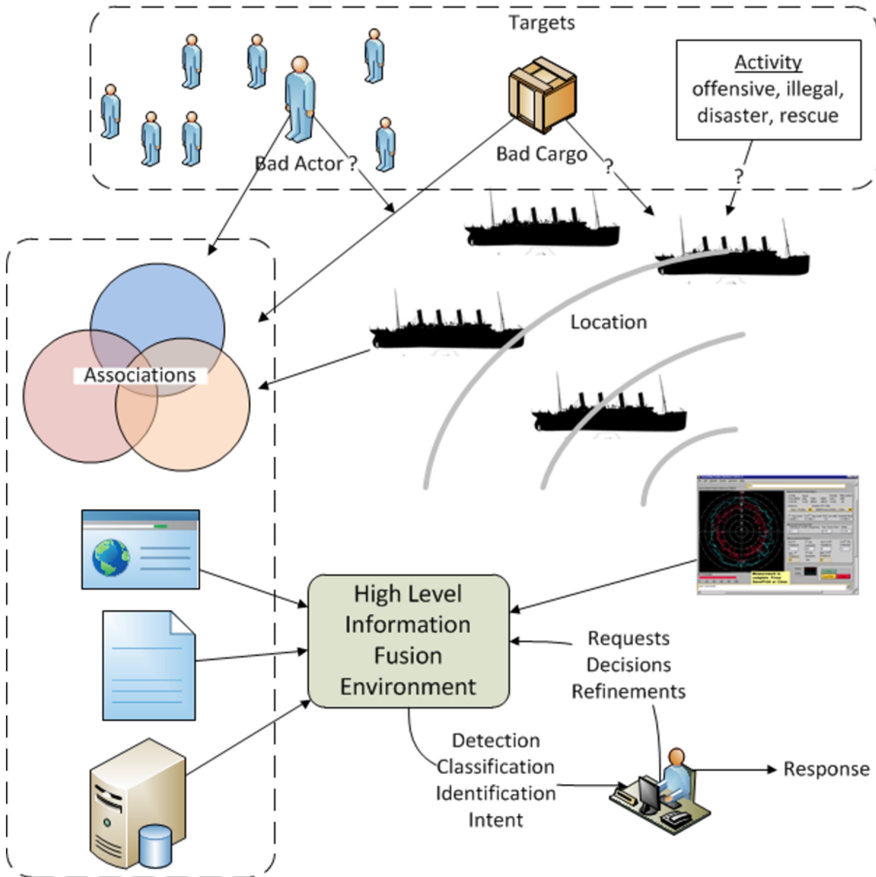

Figure 1. A model of maritime domain awareness

In Canada, MDA requires the surveying of 10 million km2 across the Pacific, Atlantic and Arctic oceans, over 200 thousand km of coastline and 5 million km2 of Arctic landmass (refer to Figure 2) and the inherent challenge of monitoring and controlling the vast amount of data and information that will be generated. This activity falls within the jurisdiction of the Marine Security Operations Centres (MSOCs) and the Canadian Forces' (CF) Regional Joint Operations Centres (RJOCs). These organizations are responsible to detect and assess Canadian marine security threats and provide support to responders. Threats include individuals, vessels, cargo and infrastructure involved in any activity that could pose an injury to the safety, security, environment or economy of Canada.

At present, there are many loosely connected surveillance and exploitation systems used to monitor maritime areas resulting in the existing disjointed maritime surveillance architecture. These intra-connected (i.e. linking the sensors that make up a surveillance system) and inter-connected (i.e. linking the surveillance systems themselves) systems have been both inflexible and expensive to setup while not being interoperable from the start (i.e. knowledge sharing between authorized users and systems is not a design consideration).

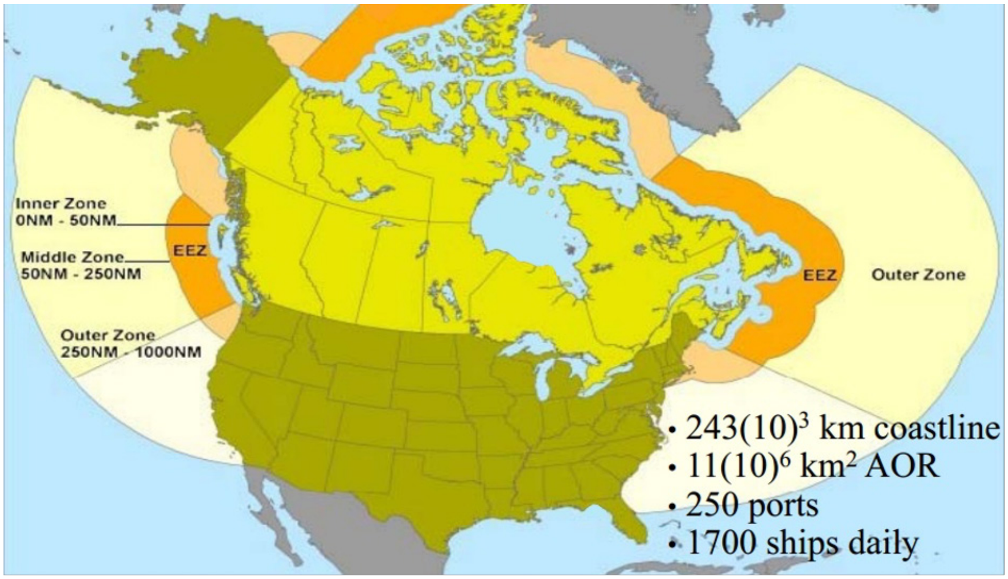

Figure 2. Canada's areas of responsibility and surveillance zones (extracted from [3])

Additionally, operators and analysts have been overwhelmed by the tide of incoming data, including sensor outputs, databases, reports and other sources of information. This situation has led to operator/analyst fatigue, overload, stress and inattention which, in turn, have led to human errors. We have seen that on a limited basis, surveillance solutions have been effective, particularly where the regions of interest were well delineated, data sources structured and precise, events-of-interest few and far between, and response requirements neither time-critical nor calculated. However, this level of performance is not sustainable over time and on a global scale. Hence, any proposed solution to these challenges will need to feature continuous awareness of the environment unconstrained by data parameters or geographical boundaries, i.e. persistent surveillance.

Persistent Surveillance

To enable effective continuous awareness, threat mitigation and response to territorial breaches, persistent surveillance is needed and must be instituted in a systematic way. Persistent surveillance systems incorporate multiple collection, exploitation and dissemination capabilities that cooperatively detect, classify, identify, track, corroborate and assess situations within maritime areas. This cooperative approach has two significant positive effects: (i) it permits the creation of fused information and intelligence products for use by decision and policy makers and (ii) it results in effectiveness and efficiency benefits due to the systems being coordinated, widely dispersed, remotely controlled and intelligent.

Looking more closely at the Canadian maritime surveillance challenge, the Arctic has been the subject of much discussion in the past few years. As Chair of the Arctic Council back in 2013 Canada led the advance of Arctic foreign policy and strongly promoted Canadian Northern interests [4]. The Canadian Northern strategy includes MDA, which is central to two of its priorities, namely: (i) securing international recognition for the full extent of our continental shelf and (ii) addressing Arctic governance and related emerging issues, such as public safety. MDA solutions that are in the works to mitigate Arctic MDA challenges are the CF Aurora patrol aircraft, the development of unmanned aerial vehicle (UAV) platforms, and introduction in 2011 of the Polar Epsilon RADARSAT-2 capability [5],a space-based radar that augments surveillance of Canada's Arctic and maritime approaches. Within the Polar Epsilon context, MDA is sequentially defined as Detect → Classify → Identify → Track → Intent [6]. Future efforts in this direction will concentrate on the "Intent" phase, including the development of additional exploitation and assessment capabilities, as well as better utilization of the upcoming RADARSAT Constellation Mission (RCM) which is scheduled for launch in 2018 and will initially include three satellites with capacity to support up to six satellites within the constellation (see Figure 3). RCM's three main uses will be maritime surveillance, disaster management and ecosystem monitoring. RCM recently received Government approval to proceed to its next and final stage of development [7].

The Polar Epsilon, RADARSAT-2 and RCM systems offer one of the best opportunities to demonstrate enhanced intra and inter connection within and between MDA systems needed to create a truly persistent surveillance capability for Canada.

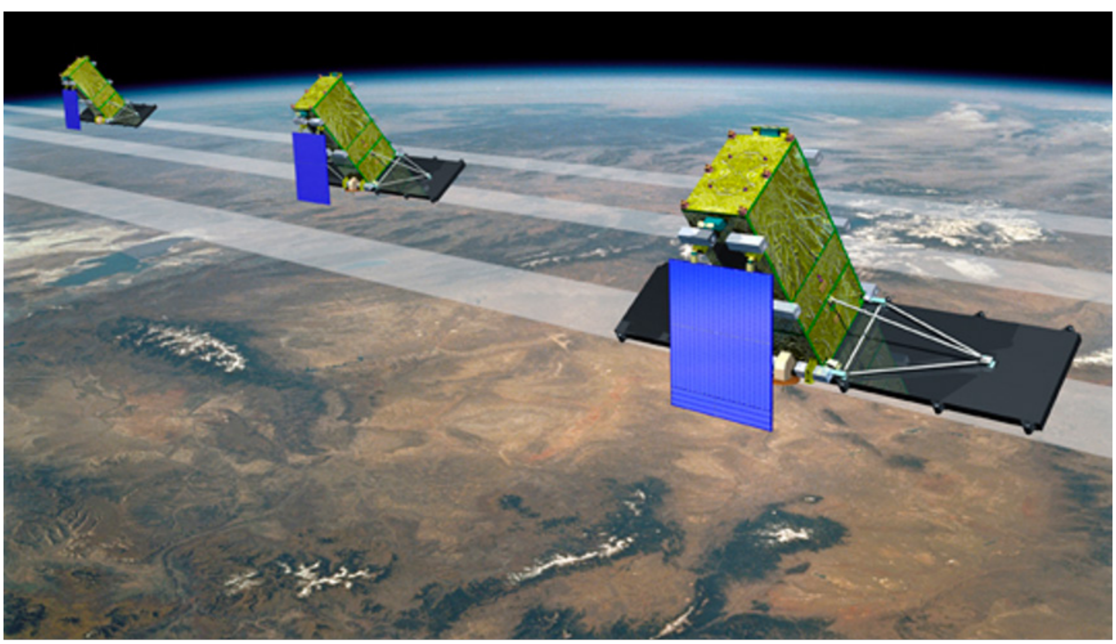

Figure 3. RCM's three satellites (Credit: MacDonald, Dettwiler and Associates Ltd.)

Data Sources

There are many potential data sources that can be inputted into MDA. These sources fall in two categories: structured vs. unstructured (sometimes referred to as hard vs. soft) and sensed vs. unsensed. Structured or "hard" indicates data that is calibrated and precise such as data from imagery and radar sensors, while unstructured or "soft" indicates data that is uncalibrated and imprecise such as operator reports and open source intelligence available from internet web pages.

Hard data typically has a high observational sampling rate, is easily repeatable and provides attributes for single objects that allow system integrators to easily interface and extract features for further processing. Examples of hard MDA data sources include, but are not limited to:

- Radar-based (e.g. Synthetic Aperture Radar (SAR), Automatic Radar Plotting Aid (ARPA));
- Tracking-based (e.g. Ground Moving Target Indicator (GMTI), LINK 11/16/22, Over The Horizon (OTH)-Gold);
- Contact-based (e.g. Automatic Identification System (AIS), Global Positioning System (GPS), National Marine Electronics Association (NMEA 0183);
- Electro-Optical-based (e.g. day/night cameras, thermal sensors, infrared cameras);
- Environmental-based (e.g. temperature, humidity, pressure, precipitation, dew, smoke);
- Ranging-based (e.g. sonar, Light Detection and Ranging (LIDAR), laser);
- Orientation-based (e.g. magnetic compass, gyroscope); and
- Ontology-based (e.g. Wikipedia, Linking Open Data Project [8]).

Soft data provides relations between discovered entities but typically has a low observational sampling rate, is not easily repeatable and is less precise. This lack of structure forces system integrators to develop techniques for feature extraction and data source ingestion. Examples of soft MDA data sources include, but are not limited to:

- Meteorological-Oceanographic-based (e.g. weather prognosis/forecast reports and ocean features reports);
- Human observation-based (e.g. field reports, interviews, intelligence reports, logs);
- Map-based (e.g. navigational charts, climate maps);
- Web-based (e.g. web sites/pages, forums, Rich Site Summary (RSS) feeds); and
- Social-based (e.g. Facebook pages, Twitter feeds, personal blogs).

Issues with soft data sources that need to be better defined include source and report credibility, handling of uncertainty, natural language processing, fusion point delineation (i.e. determining which layer is the best for hard and soft data sources to be integrated) and extraction/integration of contextual information – as soft data sources typically contain limited inferential knowledge. Hard-soft fusion has become a hot topic of research with possible solutions emerging that include the introduction of soft data exploitation within existing hard data fusion systems and the introduction of separate hard and soft streams at the ingestion point moving towards a harmonized situational understanding as the model gains experience.

Information Fusion

In order to accurately and effectively monitor a maritime area, the vast depth and breadth of incoming data must be interpreted and managed. Often referred to as the "Big Data Problem", this state is best handled through the creation and maintenance of a real-time representative model of the world. Early solutions attempted to resolve this challenge through low level Information Fusion (IF) modules that used complex mathematical formulations or brute force number crunching; however, these solutions were inadequate because the complexity created by the 4-dimensional vector (variety, volume, velocity and veracity) quickly increased to the point where low level IF modules were overwhelmed. Low level IF was only capable of performing fusion when the data itself was limited in volume, involved few types (low variety), did not frequently change in mission-critical applications (low speed) and was somewhat trustworthy (high veracity).

As data complexity continued to grow exponentially researchers realized that at some point a new approach to the Big Data Problem would be needed. That point is today, where we see data expressed in terabytes when it comes to its size, in millions per second when it comes to speed, in tens, if not hundreds of types when it comes to diversity and in jams and interferences per second when it comes to trustworthiness. A new computational paradigm is required.

To address the challenges of Big Data, *High-Level Information Fusion (HLIF)*, which in the Joint Director of Laboratories (JDL) model is defined as Fusion Level 2 and above (see Figure 4), has become the focus of research and development efforts. HLIF uses a mixture of numeric and symbolic reasoning techniques running in a distributed fashion, over a secure underlying backbone while presenting internal functionality through an efficient user interface. HLIF allows the system to learn from experience, capture human expertise and guidance, analyze contextually and semantically, lower computational complexity, automatically adapt to changing threats and situations, and display inferential chains and fusion processes graphically.

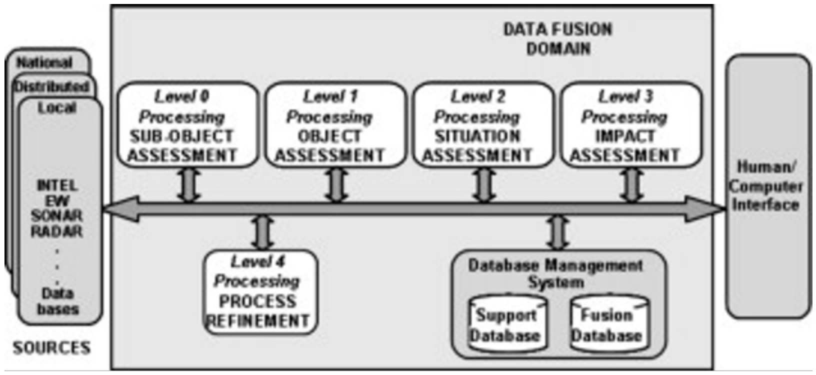

Figure 4. Information Fusion process within the JDL fusion framework (extracted from [9])

Instead of attempting to keep up with the ever-increasing complexity of the 4-dimensional data streams, HLIF, aided by Artificial Intelligence (AI) allows one to model and therefore, better understand the data stream sources and better adapt to the dynamic structures that exist within the data. Thus, the way to create an effective persistent surveillance system is to apply HLIF techniques and algorithms to the problem. So, let us now take a look at algorithms, based on AI, that furnish an HLIF system with its reasoning, inference and learning capabilities, after which we will introduce the Larus HLIF solution.

Artificial Intelligence

Artificial Intelligence (AI) involves the design of computational architectures, methodologies and processes to address complex real-world problems using nature-inspired approaches. There are three main divisions within AI, namely Neural Networks (NNs), Evolutionary Computation (EC) and Fuzzy Systems (FS), with a few more emerging trends. NNs, EC, FS and distributed algorithms, which are included in the proposed HLIF solution of this paper, are introduced in the following paragraphs.

Neural Networks

The first theory on the fundamentals of neural computing was described by W. McCulloch & W. Pitts [10] in 1943 as all-or-none threshold device that made up the basic processing unit called a neuron. When a collection of neurons was connected via weighted links, the result was a **Neural Network (NN)**, where the activity of one neuron was amplified or reduced and summed with the activity of other neurons to affect the behavior of yet another. NNs replaced the centrally executed, symbolic logical system of artificial intelligence (AI) and offered distributed processing based on sub-symbolic continuous activation levels. See Figure 5 for a depiction of one typical neuron as well as the generic architecture of a NN consisting of three layers: one input, (at least one) hidden and one output layer.

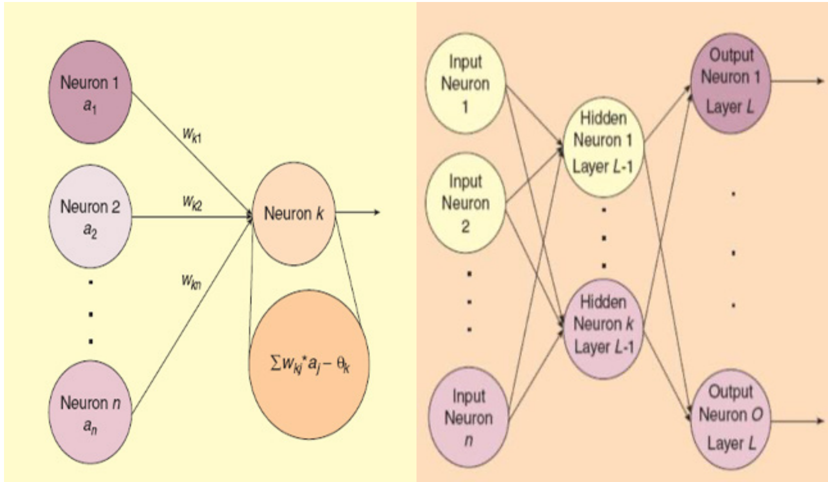

Figure 5. Neural network generic architecture

There are many types of NNs that have been devised over the years; some of the most popular and useful ones include feed-forward networks, such as the Multilayer Perception (MLP) networks, where information flow is strictly unidirectional and recurrent networks, such as the Hopfield and NARX networks, where information is allowed to feed back to nodes in earlier layers of processing.

Evolutionary Computation

NNs were found to perform successful distributed processing; however, information flow between the subcomponents was completely fixed and predetermined by the network topology. Along came evolutionary algorithms, loosely based on the interpreted operation of natural evolution, which essentially represented a distributed system of simple agents with no a priori designed communication flow pattern. In the 1960s, E**volutionary Computation (EC)** became the field of investigation into all evolutionary algorithms (EAs), including Evolution Strategies (ES), Evolutionary Programming (EP), Genetic Algorithms (GA) and Genetic Programming (GP).

Contingent on agents constructing new hypotheses about a solution to the problem, EC uses a random variation and recombination of the information about the old/previous hypothesis and performance-related evolutionary pressure which is biased towards retaining better hypotheses in the next cycle/generation of operation. EC is typically applied to problems where heuristic solutions are not available or generally lead to unsatisfactory results, where, through iterations of random variation and selection, the population can be made to converge asymptotically to optimal solutions (derived from schemata theory).

Of particular importance due to their popularity, Genetic Algorithms (GAs) were search techniques modeled after natural selection, including the associated genetic operators and were developed by John Holland at the University of Michigan in the early 1970s [11]. GAs are stochastic algorithms with very simple operators that involve random number generation, and copying and exchanging string structures. The three major operators are: selection, mutation and crossover, with fitness evaluation acting as a control factor in the feedback path [12]. GAs fare well in large search space problems because better solutions tend to "grow old with time". See Figure 6 for a depiction of a GA process as well as a pictorial of the genetic pipeline present at the heart of every GA.

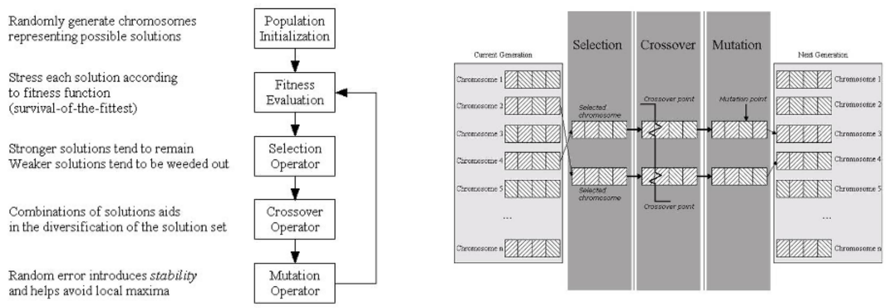

Figure 6. Genetic Algorithm process flow and the genetic pipeline

Fuzzy Systems

The mathematic notion of fuzzy sets was introduced by Lotfi Zadeh [13] in 1965 based on the concept of imprecision. Instead of presenting precise rules or instructions, the system was guided by *fuzzy rules* that described tasks more easily, such as, "when you are close to the door, open it". This became the foundation of fuzzy computation which stipulated that the interaction between computers and humans can be greatly facilitated by the use of words. **Fuzzy Systems (FS)** or Fuzzy Inference Systems (FIS) became the physical manifestations of fuzzy computation.

By crafting rules or describing data in terms that are easily understood, a system designer can simplify the design of a very complex system where measurements need only be described using fuzzy terms such as "very often" or "quite high" while membership functions can be intricately designed for fuzzification of crisp inputs. The defuzzification of fuzzy output variables into crisp values uses methods such as the center of gravity or mean of maxima methods. See Figure 7 for the typical process flow of a FS as well as a sample membership function which represents the degree of truth of an element to a particular fuzzy set. For example, a value of 0 indicates that the element does not belong to the fuzzy set, a value of 1 indicates that the element fully belongs to the fuzzy set, and a value in between indicates that the element partially belongs to the fuzzy set. This powerful concept aids in the processing of imprecise data in order to arrive at adaptive, yet rigorous, systems that yield human-assisted and interpretable solutions.

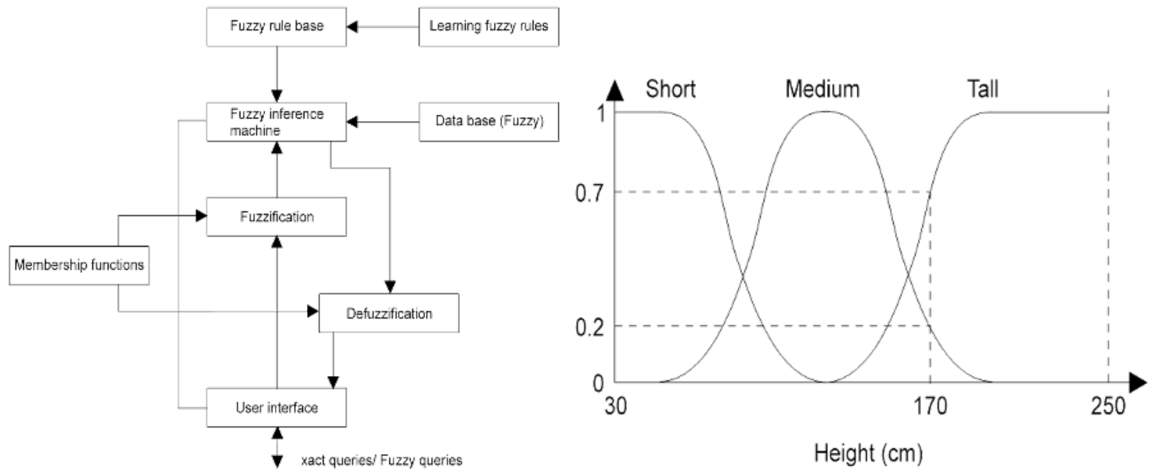

Figure 7. Fuzzy System process flow and sample membership function

Distributed Algorithms

Recently, **Distributed Artificial Intelligence (DAI)** has become viable for certain applications due to reasons, including:

- It is costly to spend all your efforts on one entity;
- Problems are physically distributed;
- Problems are complex and require local points of view; and
- Systems must be able to adapt to environmental changes.

An example of a distributed algorithm is swarm intelligence [14], which includes "ant" algorithms, *Particle Swarm Optimization (PSO)* and diffusion search. This method is based on the operation of a population of simple agents, each of which explores the space-possible solutions, until an overall solution emerges from the interactions between the agents.

Multi-agent systems (MAS) [15] consist of individual agents coordinating their activities and cooperating with each other to avoid duplication of effort as well as to exploit other agents' capabilities. MAS are typically applied to areas such as spacecraft control, social simulations, ecommerce and industrial systems management. Distributed sensing and sensor networks are a major application area of multi-agent systems [16]. Finally, *hierarchical networks* are another type consisting of probabilistic learning networks that are used to deal with problems of uncertainty and complexity. These are complex systems that are typically built by combining simpler parts. Examples include Bayesian networks and Hidden Markov Models.

Finally, it is important to mention that there are two ways to extract regularities from presented patterns, namely (i) supervised learning, where networks are provided with quantitative information on their performance, the latter being used to adjust the weights to achieve better performance and (ii) unsupervised learning, where no provision of feedback is provided to the network and the process is mostly based on an appropriately defined cost function which uses local interactions between the processing elements to arrive at a desired solution.

HLIF Capabilities

HLIF, which deals with fusion at JDL level 2 and above (refer back to Figure 4), has become the focus of recent research and development efforts to reduce the stress on surveillance operators/analysts and the burden being placed on computational systems dealing with Big Data streams.

HLIF capabilities are continuing to evolve to alleviate the challenges presented by Big Data including *(i) anomaly detection*, a process by which patterns are detected in a given dataset that do not conform to a pre-defined typical behavior (e.g. outliers), *(ii) trajectory prediction*, a process by which future positions (i.e. states) and motions (i.e. trajectories) of an object are estimated, *(iii) intent assessment*, a process by which object behaviors are characterized based on their purpose of action, and *(iv) threat assessment*, a process by which object behaviors are characterized based on the object's capability, opportunity and intent.

Additionally, real-time adaptive learning becomes an imperative feature of any MDA solution deployed in the field. *Situational learning* (shaping future responses to already seen situations based on human feedback) and *procedural learning* (minimizing the error between predicted and actual events) are two methods that enable a system to better understand its real-world dynamics.

Larus HLIF Architecture

Larus Technologies has developed a patent-pending HLIF architecture targeted as an MDA solution. The Larus HLIF architecture performs behavior analysis through predictive modeling, fuses heterogeneous (i.e. multi-source, multi-sensor) data and is automated yet human-centric. While other HLIF solutions provide separate frameworks for sensing and acting, the Larus architecture allows for bidirectional in-network processing (i.e. sensing and acting are performed within a unified framework), makes applications data-centric and reduces the state-estimation errors by closely matching the world model to the real world. The Larus architecture can be summarized by the two-way relation shown in Figure 8.

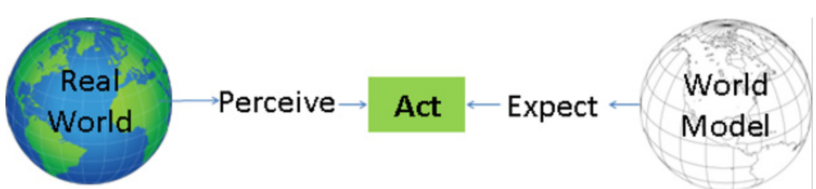

Figure 8. Larus HLIF architecture based on real world state estimation

The Larus architecture is made up of perception, validation, expectation and action modules. The architecture combines AI algorithms situated within a persistent surveillance environment to support decision makers by classifying, identifying, tracking and assessing targets. The perception module, which is responsible for data consumption, processes and analyses sensor inputs extracted from data sources, including the environment. The validation module, which is responsible for *information consumption*, performs multi-source multi-sensor data fusion to extract common patterns and parameters from heterogeneous data. The expectation module, which is responsible for *decision support*, diffuses commands to actual tasks through predictive modeling. Finally, the action module, which is responsible for sensor tasking, provides effectors in the environment by performing tasks. Each action changes the state of the environment, after which the entire cycle repeats. This architecture includes a *world model* that represents the knowledge base attained by the system. Figure 9 depicts the entire flow.

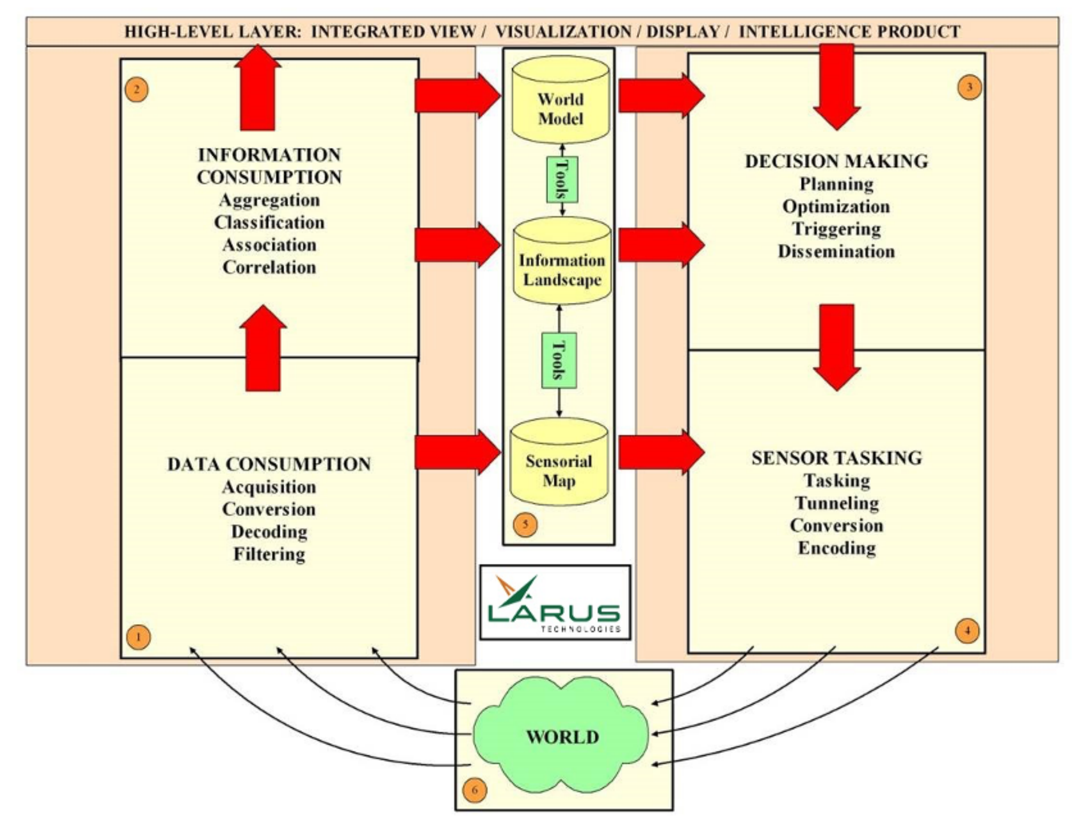

Figure 9. Data flow overview

Larus Solution Background

The Larus Technologies HLIF architecture originates in the field of intelligent agent architectures. Intelligent agent architectures define approaches to building intelligent systems, including the internal structures and operations of the agents. Typically these architectures include proactive and reactive architectures. Proactive agent architectures are intuitive and easily decomposed into subsystems; however, they suffer from the problem of calculative rationality and usually are difficult to realize. Reactive agent architectures are simple, economical and computationally tractable; however they suffer from the problems of having a short-term view and being difficult to implement when they contain many layers.

The novel *retroactive agent* architecture was designed and developed by Larus to combine and improve on both proactive and reactive characteristics within a Decision Support System (DSS). In the retroactive architecture, an event in the environment occurs, causing momentum in a behavior to increase, which eventually causes a behavior to fire and react to the event by executing the behavior's plan. The events of importance are changes in the environmental features that the world model is concerned with; i.e. ones that conflict with its goals. The reactive characteristic of the retroactive architecture involves knowing which behavior to fire through *momentum resolution*, while the proactive characteristic involves knowing how to behave through plan execution.

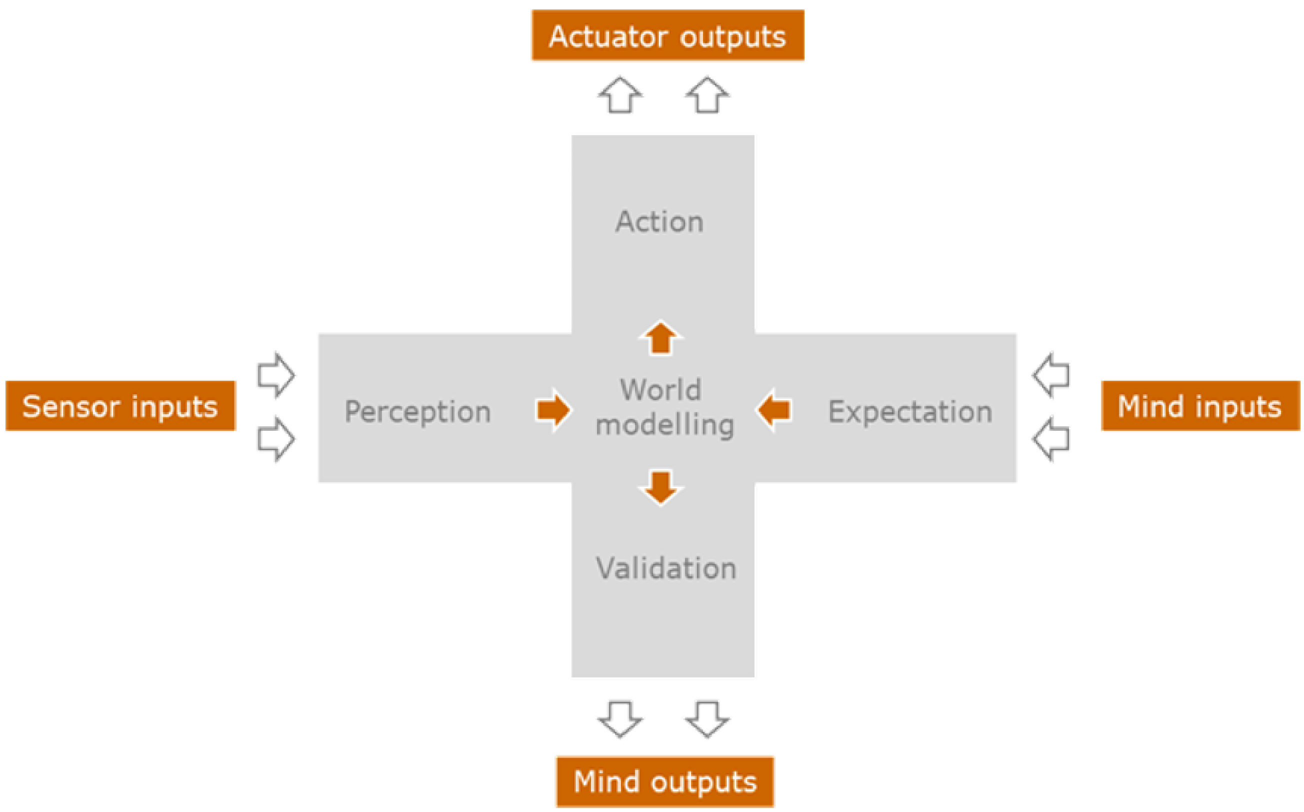

Figure 10. Retroactive agent architecture

The retroactive HLIF architecture builds on the importance of synaptic feedback, where, for example, every sensed event elicits a reaction by the DSS onto the world to handle the event, a retro-assertion (i.e. future validation) onto the DSS' world model to compartmentalize the new or old data, and subsequently a possible pro-action onto the world again to seek more information about the event. Larus has identified six synaptic acts, namely reactions, pro-actions, retroactions, reassertions, pro-assertions and retro-assertions, which if observed from the outside, ascribe intelligent behavior to the DSS. To better visualize the synaptic acts, refer to Figure 10. A retroactive architecture allows the DSS to learn over time, solving the short-term view of reactive systems, and to respond in real-time to world events, solving the calculative rationality of proactive systems.

Larus HLIF Architecture Scenario

The Larus HLIF architecture is presented in the context of the MDA scenario shown in Figure 11. In this scenario, the Larus HLIF architecture automatically fuses multiple hard and soft data sources to determine location, destination and intent of a target vessel (i.e. Level 2 information fusion). The data sources are then used to refine the situational understanding by identifying potential consequences and postulating courses of action (i.e. Level 3 information fusion). At the highest information fusion level (i.e. Level 4), the architecture offers response options to the decision maker who then orders units to respond and tasks sensors and sources to further refine situational understanding and support those responding units.

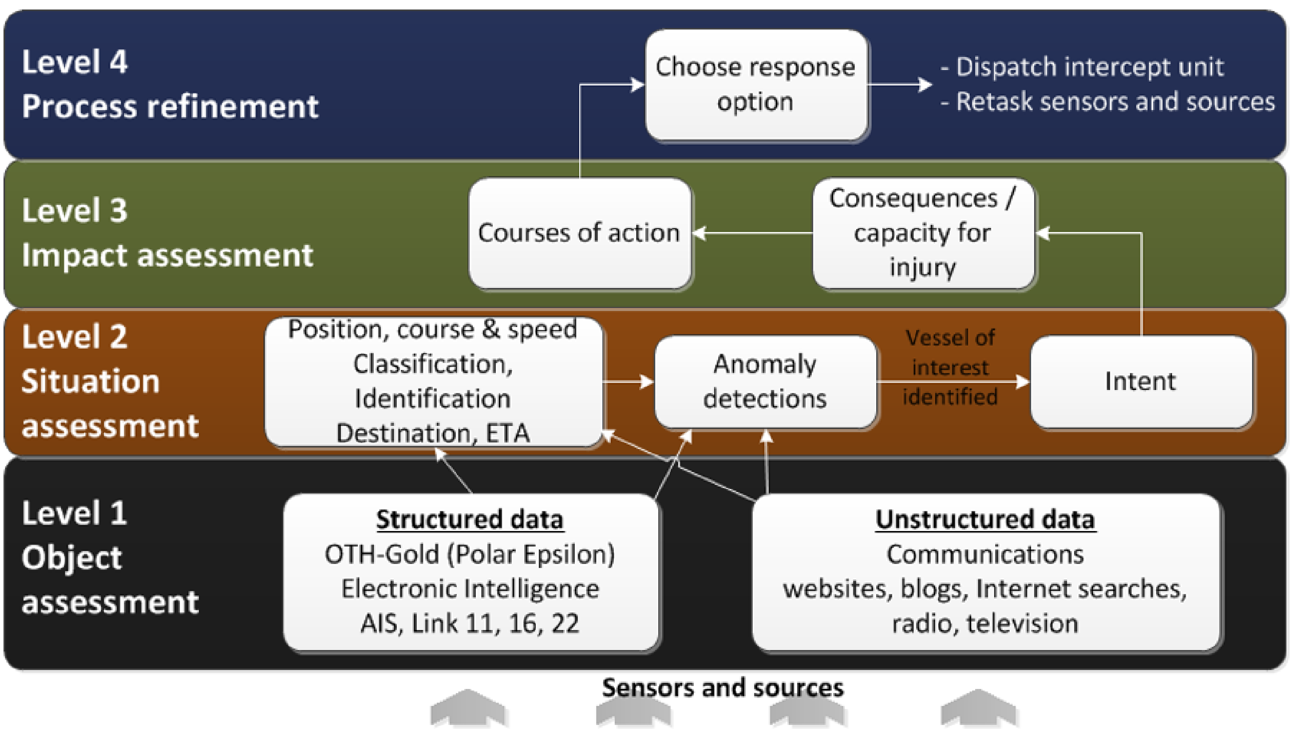

Figure 11. Application of HLIF architecture to a maritime domain scenario

The fusion process described in Figure 11 operates over a wide area of interest continuously assessing information, highlighting anomalies, refining situational understanding and learning to improve its own performance over time.

Conclusions

Continuous Maritime Domain Awareness provides situational understanding and support to decision makers through the ingestion, processing and presentation of a vast array and volume of data about the maritime environment, objects and actors in that environment and intentions of those actors. To be successful, continuous Maritime Domain Awareness requires the collaboration of multiple systems to provide persistent detection, classification, identification, tracking and assessment of situations within the maritime domain. The Larus Maritime Domain Awareness solution, through its retroactive HLIF agent architecture, improves on existing persistent surveillance methods by generating understanding of the objects, actions and intentions. It adds automation to the surveillance process by fusing a multitude of hard and soft data sources through Artificial Intelligence and Machine Learning algorithms and behavior analysis into a Decision Support System. The Larus solution learns and continuously improves upon itself in real-time to provide true and timely information on maritime activities, reduce operator workload, provide accurate and reliable world model and enable interoperability and knowledge sharing.

References

[1] NASA Socioeconomic Data and Applications Center (SEDAC). "Percentage of Total Population Living in Coastal Areas." **http://sedac.ciesin.columbia.edu/es/papers/Coastal_Zone_Pop_Method.pdf**, New York, NY, USA. 25 July 2007.

[2] International Maritime Organization, Amendments to the International Aeronautical and Maritime Search and Rescue (IAMSAR) Manual, pp. 1, last accessed on February 24, 2013.

[3] P.W. Vachon, "New RADARSAT Capabilities Improve Maritime Surveillance", October 2010.

[4] Government of Canada. "Canada's Northern Strategy: Our North, Our Heritage, Our Future", 2009.

[5] A., Clive, "Vision for the Arctic," FrontLine Security, 6(1), June 2011.

[6] R.J. Quinn, "Project Polar Epsilon: Joint space-based wide area surveillance and support capability", Proc. 8th International Conference on Remote Sensing for Marine and Coastal Environments, 17-19 May, 2005, Halifax, Nova Scotia, Canada, 2005.

[7] Canadian Space Agency, "The Government of Canada Launches Final Stage of RADARSAT Constellation Project", last accessed on February 23, 2013.

[8] Linking Open Data Project, **http://www.w3.org/wiki/SweoIG/TaskForces/CommunityProjects/LinkingOpenData**, last accessed on February 23, 2013

[9] D. A. Lambert, "A blueprint for higher-level fusion systems," Information Fusion, vol. 10, pp. 6-24, January 2009.

[10] W. McCulloch, and W. Pitts, "A logical calculus of the ideas immanent in nervous activity," *Bull. Math. Biophysics,* vol. 5, pp. 115–137, 1943.

[11] J. Holland, Adaptation in Natural and Artificial Systems, The University of Michigan Press, Ann Arbor, Michigan, 1975.

[12] D.E. Goldberg, *Genetic Algorithms in Search, Optimisation and Machine Learning.* Reading, MA: Addison-Wesley, 1989.

[13] L.A. Zadeh, "Fuzzy sets," Inform. Control, vol. 8, no. 3, pp. 338–352, 1965.

[14] J. Kennedy and R.C. Eberhart, Swarm Intelligence. San Mateo, CA: Morgan Kauffmann, 2001.

[15] Gerhard Weiss (editor), "Multiagent Systems," The MIT Press, 2001.

[16] R. Abielmona, E.M. Petriu, and T. Whalen, "Multi-Agent System Environment Mapping by Entropy Reduction" Proc. ROSE'2005, IEEE Intl. Workshop on Robotic Sensing, pp. 30-35, Ottawa, ON, Canada, Oct. 2005.

.

About Larus

Through our culture of innovation and research, Larus Technologies has developed the next generation of embedded technology for developers of mission-critical C4ISR Systems and Security Systems.

With a solid foundation pioneering high level information fusion (HLIF) for the ever-changing defense and security industries, Larus is perfectly positioned to help Original Equipment Manufacturers (OEMs) make a world of difference. Working at the higher levels of the US Department of Defense's Joint Director of Laboratories (JDL) information fusion model, our technology not only delivers more knowledge, its adaptive learning algorithms deliver more accurate and more predictive information—faster.

170 Laurier Ave West, Suite 310
Ottawa, Ontario K1P 5V5
Canada

Phone +1 (613) 244-8916
Fax +1 (613) 244-9941

www.larus.com