Cognitive Radar Information Networks for Security – Enhancing Risk Mitigation for Reducing Operator Overload

Tim J. Nohara, P.Eng, Ph.D, *Member IEEE*
Accipiter Radar Technologies Inc.
Niagara, Canada
tnohara@accipiterradar.com

Simon Haykin, Ph.D, *Fellow IEEE*
Distinguished University Professor
McMaster University, Hamilton, Canada
haykin@mcmaster.ca

***Abstract*—The paper presents a practical foundation for building and deploying affordable, cognitive radar information networks (CRINs) to secure large, unmanned borders such as the 3,700 km Canada/U.S. border that runs through the Great Lakes. The proposed systems and methods are novel and necessary if radar's future in the 21st century is to embody the risk management principles needed to optimize system and operator resources for detection of suspicious target behaviour. Human cognitive abilities of attention and intelligence are proposed as cardinal characteristics that are built into CRINs to provide force multiplication and reduce operator overload.**

I. INTRODUCTION

The events of September 11, 2001 focused the efforts of various public and private stakeholders on homeland security. Identified threats include terrorist or criminal activities, accidents or natural disasters. Threats occurring on or alongside water are particularly challenging as waterways are vast in extent with large numbers of recreational and commercial vessels [1].

Terrorist or criminal activities can be carried out using low-flying general aviation aircraft, and vessels of all sizes from large container ships down to zodiacs and jet-skis. When the water is frozen over, snow-mobiles and vehicles add to the target mix. Awareness of what these uncooperative targets are doing at any given time and understanding whether particular target behaviour is suspicious and requires closer examination is what we mean by situational awareness. Protecting people and property from such threats requires situational awareness that provides authorities and citizens with timely information to prevent, respond to, and mitigate them.

Threats can occur anywhere over vast waterways such as the 3700 km long Great Lakes St. Lawrence Seaway System (GLSLSS); and they can occur at any time, day or night, and are infrequent; therefore situational awareness is needed 24/7/365. Furthermore, because threats can unfold in just seconds (e.g. a vessel crosses a narrow waterway such as the St. Lawrence River and lands on the shoreline of another country violating an international border, or a vessel enters a marine exclusion zone on the waterside of a nuclear power plant on Lake Ontario), persistent surveillance is needed to provide adequate situational awareness.

Manufacturers have responded with the development of affordable, wide-area surveillance radar information networks which are in the early stages of deployment to provide the required situational awareness to stakeholders [2]. New operator issues referred to herein as the operator overload problem, arise from the vastness of the areas covered, the large number of friendly targets present, and the large number of radars to be controlled. These 21st century radar information networks would benefit significantly if the operator's cognitive abilities of attention and intelligence could be built into them.

The wide-area radar information networks we have in mind, typically, consist of a number of inexpensive ground-based radar sensors mounted on structures around the large waterways they intend to provide surveillance coverage for. The structures include towers, roof-tops, tripods on the ground, tethered aerostats, and even mobile structures such as trucks on land or vessels on the water. Such a network is illustrated in Figure 1 and described in detail in [2].

The arrangement of radar nodes and the coverage afforded by each are designed to overlap so that seamless coverage is available across the wide area of interest. In this way, targets moving anywhere through-out the entire area will be picked up by at least one radar sensor node. At each sensor node, the radar typically scans 360^o, and uses its transmitter to interrogate the environment with a selected waveform. Its receiver picks up reflections or echoes from targets and clutter in the environment, and processes them, using a particular receiver processing mode, to automatically extract information from each target in the environment. The extracted information is referred to herein as target data, and typically consists of both detections and tracks.

Target data typically include an estimate (also referred to as a measurement) of target parameters updated every couple of seconds. These parameters include {latitude, longitude, altitude, speed, heading, size (i.e. radar cross section)}. At any instant in time, the current target data represent the locations, dynamics and sizes of all targets seen by the radar

The authors greatly acknowledge Defence Research Development Canada Centre for Security Sciences for financially supporting this work.

network. Over time, complete trajectories are extracted or formed for each target indicative of target behaviour. Target data are typically sent over standard computer or cellular network links (wired or wireless) to an information system or repository that stores forever, organizes, and relays desired target data to operators/users in real-time. For the purpose of this paper, the information system will be considered centralized without loss of generality.

Remote users are typically provided with a number of applications (software) which query the information system for both real-time or historical target data, including post-processed target data. A host of target data processors can access the information system over the network, generate a variety of information products (e.g. traffic patterns, border-crossing statistics, suspicious behaviour alerts, marine exclusion zone violations around critical infrastructure, etc.), and make these available to users as well. Information sent to users support a variety of user missions including surveillance, automated alerting of suspicious activity or activity of interest, interdiction, intelligence, investigations, analysis, prosecution and research.

Radar networks, in accordance with the architecture described in [2], allow owners of radar nodes to share any subset of their target data with others, creating virtual radar networks that span political boundaries, for example. Standardized, open interfaces allow authorized users the ability to access shared target data and generate their own information products behind their own respective firewalls for information privacy and policy purposes. Such radar networks have been shown to be particularly valuable in joint law enforcement operations [3].

The modular nature of the radar information network includes flexibility in selection of platforms, transceivers, and antennas to address coverage and performance requirements; and software definable radar processing algorithms suitable for target extraction of surface and air targets of interest. Typical radar sensors are X or S-band marine radars with a rotation rate typically in the 24-48 RPM range. Conventional magnetron radars typically include three waveforms: a short-pulse (SP), high-range resolution waveform (e.g. 10 m), a medium-pulse (MP), medium range resolution waveform (e.g. 50 m), and a long-pulse, low-resolution waveform (e.g. 150 m). Solid-state, Doppler marine radars are also available. A SP waveform is typically available; and rain and sea clutter suppression is afforded by Doppler processing. Radar remote controllers are available which allow transmitter illumination and receiver processing mode to be changed under remote software control [2].

The commercial-off-the-shelf (COTS) marine radars all come with non-agile array antennas that provide a horizontal beam width typically between 0.3^{o} to 3^{o} and a vertical beam width typically 20^{o} wide and oriented as +/- 10^{o}. These 2D radars can not provide altitude information in their target data. Marine radar transceivers coupled with custom-developed, agile pencil-beam antennas are also available [2], provide altitude information for airborne targets, and can be commanded under software control to follow an airborne target such as a low-flying aircraft.

Specialized, military or coherent, 3D radars are also available in some cases and can integrate with the radar networks considered here. In such cases, phase information may be available in the target data as well.

In summary, the wide-area radar information networks of the 21st century will involve a heterogeneous mixture of radars, mostly marine radars, which are not synchronous, but operate independently as described above and form a virtual radar information network (VRIN). They will provide target information in a common format to a target information system so that integrated situational awareness can be provided. New radar technologies will build upon such radar information networks to provide new capabilities.

II. WHY COGNITIVE RADAR INFORMATION NETWORKS

With conventional military or air traffic control (ATC) radars, operator overload is mitigated by employing dedicated and highly-trained radar operators. Furthermore, targets of interest (TOIs) (i.e. aircraft) are cooperative making them relatively easy to assign to operators and manage. In military scenarios, identify friend or foe (IFF) transponders installed on friendly targets allow operators to automatically distinguish friendly radar target tracks from enemy tracks or false tracks caused by clutter. Civil aircraft carry devices that allow air traffic controllers to do the same. The air space or maritime operating areas are also controlled so that targets that should not be there stay out. Dedicated radar operators ensure their respective radars are optimally tuned for detecting and tracking their respective TOIs under different environmental conditions.

This is not the case for 21st century radar information networks. It is the uncooperative targets (that do not transmit their location to make them identifiable) that are of most interest, and especially the small ones which are most difficult to detect and track with radar at further distances. Furthermore the area of interest is not controlled, so thousands of commercial and recreational targets are present. Finding suspicious targets is like looking for a needle in a haystack. In addition, the vastness of the area covered by the radar information network, and the number of (dissimilar) radars involved make optimizing radar sensitivity to specific developing situations particularly challenging. It is no longer a simple task of a dedicated operator adjusting the gain-control dial on one radar in order to increase its sensitivity. Now, a network of remote radars needs to be intelligently adjusted to focus attention in a particular localized area where it is needed at a particular time. If ignored, the above issues will lead to either operator overload or result in a limit in system performance that is otherwise achievable. By addressing these issues, we can significantly enhance the price, performance, and complexity advantages of radar information networks; and help operators to be more productive by focusing on tasks best suited to their skills.

No one today would deny the informational, communication, and joint / common situational awareness that the Internet brings to people around the world, including law enforcement personnel. This wide-area computer information network, while organizing information on a scale never before imagined, brings with it a major information overload problem. How do different users find the information they are interested in when they need it? The answer lies in algorithms that continually crawl the web, identifying new content and network behaviour, indexing and organizing the information so that it is searchable, and providing search engine tools and an ecosystem of applications to help users quickly get what they need when they need it. The same is true for radar information networks, albeit on a much smaller scale. To maximize situational awareness and situation understanding, the information content (including targets and disturbances such as clutter) of the environment must be continuously analysed and indexed to automatically learn from the environment, with tools provided to assist operators in exploring, discovering and finding what they need when they need it.

If we had enough of them, and if they had the time, experienced and dedicated radar operators have the cognitive abilities to focus attention where it is needed, and the intelligence to learn from the environment and trade-off how best to keep the radar network optimized for changing security situations. These innate capabilities of the human brain lead us to examine the potential of cognitive dynamic systems to further improve the performance of radar information networks. In this context, there is much that we can learn from the visual brain in designing a new generation of cognitive radar networks [4–6].

Radar information networks would typically be deployed and tuned to baseline operating conditions. If conditions never changed, there would be reduced need (for an operator) to change radar network settings. However, in practice, changes in the environment occur which are referred to herein as unexpected or uncertain events. The events drive our need to focus the radar information network's attention to localized areas, intelligently, i.e. without compromising the overall performance of the radar network.

A first class of events leads to an automatic CRIN response not requiring operator intervention, thereby mitigating operator overload. Based on the location and nature of the event and past experience gained by the CRIN, the operator is alerted to the event and appropriate radar sensor nodes are automatically adjusted (e.g. a particular transmitter waveform and/or associated receiver processing mode is selected for use) to robustly focus system attention or optimize performance where needed. Suspicious targets as well as environmental disturbances fall into this first class of events.

The CRIN will automatically detect suspicious targets around particular areas of interest such as border crossings, or marine exclusion zones (MEZ) around critical infrastructure such as a nuclear power plant. A rendezvous, as well as other abnormal behaviour such as deviation from regulated routes, will be detected by the CRIN; then the operator will be alerted, and the CRIN will automatically adapt itself to bring attention to such areas to reduce operator overload and/or enhance performance. Detecting such suspicious behaviours is particularly difficult for operators, especially as they occur within dense traffic environments, and may take considerable time and concentration to observe.

Environmental disturbances include weather (e.g. precipitation) which can lead to the appearance of "false" targets in localized areas, and sea/lake clutter variations which result in reduced radar sensitivity and/or increased probability of false alarm (PFA). Such disturbances are typically isolated to relatively small areas in comparison to the total coverage area; and they typically move or change with the predominant winds, which for Lake Ontario are westerly.

A second class of events is operator driven. In this case, the operator has knowledge of an unfolding situation and wants the CRIN to robustly focus attention in a designated area or areas for some period of time. The following are examples of operator-driven events:

- Intelligence indicates that an illegal transaction is likely to take place in a certain area and extra sensitivity is need for evidence and prosecution;
- law enforcement personnel are conducting a covert surveillance operation in an area and want increased sensitivity;
- there has been an accident on the water with a small vessel and search and rescue personnel require focused attention to find the drifting vessel or wreckage; and
- a particular target, e.g. an LNG (liquefied natural gas) tanker, is moving through the wide-area and a high-sensitivity region or protective bubble-zone around the target is desired throughout its journey.

For both classes of events, the CRIN must continue operating robustly in its primary surveillance mission, notwithstanding the fact that it is making changes to the baseline operating state of the network, by changing in a localized manner, the transmitter illumination and receiver operating mode of one or more radars. Operators will be able to define global and regional performance figures of merit (FOMs) that are to be maintained during the application of attention to designated areas of interest. Areas or cells available for increased attention can be pre-defined as illustrated in Figure 1, allowing the CRIN to gain experience by learning from its environment. The FOMs might be based on a multi-target tracking continuity measure applied to a given attention cell, for example. The FOMs ensure that global target sensitivity does not degrade below some specified level, while particular regional areas may have different performance thresholds that must be maintained. It is all about exploiting the available system resources associated with the network of radars in the best manner possible to meet mission requirements. For suspicious targets or disturbances that move over large areas, the attention cells will change over time while the CRIN maintains its surveillance mission.

For CRINs that include agile radar sensors, the CRIN could also be tasked with locking onto a designated, high-risk TOI such as a low-flying aircraft headed towards an urban area, and sacrificing surveillance for a short period of time to maintain target lock.

III. Cognitive Radar Information Network Concept

A detailed presentation of cognitive radar principles, theory and application to single radars and the single target tracking problem is found in [4]. An extension and adaptation to multi-target radar surveillance networks is provided in [7], which should be consulted for a theoretical treatment of the CRIN presented here. This CRIN is based on Fuster's paradigm for cognition [8] and a two-state, cognitive radar environment model that includes the usual, deterministic target state combined with a specially-adapted, information-dynamical entropic state. The adoption of entropic state as the feedback information from the receiver to the transmitter opens the door for using "reinforcement learning" [9] as the approximate dynamic programming algorithm for designing the cognitive controller, whereby computational complexity of the transmitter is reduced dramatically.

A. *Concept of Operation*

The wide-area covered by the CRIN is divided up into a number of virtual, geographical, attention cells as illustrated in Figure 1. The CRIN operates with a baseline radar configuration and performance (which may change with season). Attention is applied to one or more requested cells while intelligence ensures that surveillance in the remaining cells continues to function robustly. Robustness includes both global and local cell metrics – i.e. global performance is controlled so that it will not degrade below a certain threshold; and individual local cells are controlled so as not to degrade below respective performance thresholds.

A cell may be designated for added attention by: (1) the operator, or (2) automatically by the CRIN's cognitive controller. Both the cognitive controller and the operator provide intelligent feedback to the system. In one case, the operator may see a situation first that requires attention; in another, the cognitive controller may deduce a situation (e.g. due to a disturbance such as weather or a potentially suspicious behaviour detected) that results in attention being designated for one or more cells.

Using pre-defined attention cells allow the CRIN to build up knowledge and experience in relation to each cell so that the cognitive controller can act quickly and robustly, on-the-fly, when attention is needed in a certain area.

An area of focused interest may move with a target as the target journeys (like a protective bubble), causing cells needing attention to change automatically in time. Even weather cells move – hence, the cognitive controller can use its intelligence to automatically adjust the cognitive radar information network on-the-fly in response to such TOIs or weather disturbances.

Cognition may also be applied to the problem of detecting suspicious target behaviour, in order to assist the operator and reduce information overload, especially in dense target environments. The system could also automatically propose new scenarios not described previously by the operator, but determined by the CRIN to be anomalous, through its gained experience. For example, the CRIN could learn from past behaviour that a particular type of target has never been observed in a cell or area where it is now detected.

Even a single target of interest will occupy a region over a period of time (because it moves). Furthermore, situations of interest such as border crossings, loitering in a marine exclusion zone, or a target deviating from a regulated laneway all occur in geometric regions. These regions have physical extent and shape in space. We call them geometries. Hence, activity around such a geometry can be emphasized by applying attention to the cell or cells containing it. Methods for detecting suspicious targets such as hypothesis testing are described in [7], where thresholds are behavioural in nature and soft-computing algorithms are exploited.

B. *Cognitive Radar Basics*

A single Transmitter and Receiver that incorporates a perception-action cycle similar to that used in the brain is at the heart of a cognitive radar [4,7]. During each cycle of the perception-action cycle, the Transmitter transmits a waveform to illuminate the environment, the Receiver receives measurements from the environment, and provides feedback information about the environment back to the Transmitter, which is used by the Transmitter's Cognitive Controller to adaptively (based on learning) select a new (potentially different) illumination for the next perception-action-cycle. Traditional radar (without cognition) does not have the feedback path between the Receiver and Transmitter and does not learn from experience how to best select a new illumination.

The radar environment or world includes two sources: unknown targets and disturbances such as clutter (e.g. precipitation, sea/lake clutter) and thermal noise. Targets are deterministic in the following sense. There is a certain number of them, and each one has a particular location, speed, heading and RCS at each time instant. What we have described here are the unknowns. Though targets are deterministic, since they are unknown, they represent a source of uncertainty. On the other hand, disturbances are stochastic in nature. Disturbances and uncertainties are responsible for state estimation errors.

The Receiver extracts information from the measurements it receives about the world. In particular, it computes or estimates two states associated with the world at each particular time update t: (1) the target state of the world represents its knowledge of the multiple targets present at time t, and the entropic state of the world which represents the disturbances in the world, along with the imperfections associated with its target knowledge (as represented by the target state). The dynamical entropic state is associated with the target state estimation error. Entropy is a metric for assessing the quantity of information we are lacking – hence the name entropic state. The Receiver is assumed to have memory to store the target state and dynamical entropic state forever, or alternatively, sends this information to a central location or repository for storage. In [7], algorithmic details for computing the dynamical entropic state are provided.

The Transmitter does not see (i.e. sense) the environment directly; rather, it illuminates the environment. As a result, if it is to learn from the environment in order to select an appropriate illumination in an intelligent manner, it must see the environment through the Receiver's eyes. This seeing is implemented through the feedback path from the Receiver to the Transmitter, which is a key characteristic of cognitive radar systems.

Humans learn how to make good, robust decisions from past experience; also they remember their consequences. The Transmitter (or some central computer where the Transmitter's Cognitive Controller lives) is afforded with memory so that it can learn from its past illumination-selection decisions and their impact on each attention cell. The Feedback Information is based on the current Receiver measurements (typically predicted forward to the next time step) and/or the entropic state which the Cognitive Controller uses in its algorithms to make a robust illumination selection.

C. Implementation Details

The single cognitive radar can be expanded and extended into a cognitive radar information network (CRIN). If one takes advantage of the ability of modern RINs to organize their Receiver information in a central repository (by having each Receiver send its information in real-time there), then a simplified Figure 2 results where the feedback paths are indirect through the central repository.

The CRIN block diagram in Figure 2 shows Receiver and Transmitter memory organized in a Central Repository which attempts to isolate the cognitive elements so they can be spiraled into existing radar networks that do not presently have cognition. This is essential to manage cost and risk. Furthermore, cognition can be introduced incrementally, in one or more nodes at a time, with increasing levels of cognition added as software upgrades once the underlying R&D and implementations mature. Receiver information from the collection of radars is organized in the Central Repository providing direct access (or feedback) to the Transmitter via the Cognitive Controller.

While shown centralized, it should be understood that the Cognitive Controller can provide local cognitive control to individual Transmitters in the network based on respective Receiver feedback, or alternatively, cognitive control could make use of feedback from multiple Receivers or globally from the entire network. Computer network interfaces are assumed to exist between all system components allowing information to be easily shared.

Requests for attention come directly from the Operator, or automatically through the Automatic Behaviour Analysis and Detection (ABAD) processor illustrated in Figure 2. It should be noted that the ABAD processor can generate on the fly post-processed Receiver information that can help characterize the environment for future rapid use by the Cognitive Controller.

Transmitter modes are defined for each radar node. A transmitter mode is a particular set of transmitter illumination parameters which include waveform (e.g. SP, MP, LP), each of which affect range and resolution performance; RPM and elevation beam in some cases (e.g. surface beam and air beam with particular elevation angle).

The Cognitive Controller controls each Receiver by changing the Receiver (processing) mode in combination with the selected and associated Transmitter mode. Receiver mode parameters include threshold for sensitivity (PD, PFA), masking for areas of processing interest, tracking filters for air versus surface targets, etc. The Cognitive Controller also balances system constraints in its decisions, such as target data rate limits due to bandwidth constraints, and computational loading.

We introduce a new environmental information quantity called adaptive target maps. Adaptive target maps can be generated from targets of opportunity observed over short-term and seasonal time frames in each attention cell, and versus each radar node's available Transmitter and Receiver modes. In this way, the CRIN can learn from the environment target behaviour, as well as the CRIN's own performance in terms of localized, cellular coverage maps for each target type. Adaptive target maps can be further organized by their attributes (e.g. small or large RCS, slow or fast speed, acceleration, etc.) into subgroups for efficient cognitive processing. These post-processed data characterize the knowledge learned from the environment over time; and they can be organized in the Central Repository for permanent storage and exploitation by the system and the operator.

These new adaptive target maps (analogous in a loose sense to adaptive clutter maps), provide meaningful surveillance performance information as a function of all Transmitter/Receiver modes and all attention cells for all radars in the network, allowing the Cognitive Controller to manage system robustness when selecting new Transmitter modes for illumination and associated Receiver modes in response to attention requests.

For more discussion, block diagrams, algorithmic methods and examples concerning the Cognitive Controller, the ABAD processor, the perception-action-cycle, attention, intelligence, dynamical entropic state, adaptive target maps, synchrony and communications between asynchronous radar nodes, along with a variety of other related topics, please consult [7].

D. Policy & Security Implications

The Beyond the Border Agreement [10] between Canada and the United States is a timely policy blueprint the radar community can use to foster innovation for improvement in security while respecting privacy, civil liberties and human rights. Radar, by its very nature, is the surveillance technology of choice not only because of its wide-area, all weather, performance, but also because personal identity is respected in the processes of radar information sharing, common threat understanding, risk management, and joint, intelligence-led operations called for by the Agreement.

Using data from real information networks overlooking Lake Ontario and the Canada/US Border that runs through it, we have already seen how intelligence-led border enforcement can be supported and security enhanced with high-value and timely understanding of suspicious vessel activity, border crossings and low-flying aircraft [11]. Building cognition into these radar information networks as described herein will further enhance risk mitigation and reduce operator overload, providing even greater force multiplication.

IV. CONCLUSIONS AND FUTURE WORK

In the 20th century, radar had its genesis and major advances were made to radar engineering in response to the physics of remote sensing of the environment. A skilled operator was coupled to large, scalar radar sensors to provide interpretation and intelligence.

This paper presents a vision for the 21st century. Homeland security, rather than theatre defence, will drive radar innovation. Homeland security requires situational understanding over vast areas, continuously, forever, and the areas are target rich, with mostly friendly targets, and the odd threatening one. Because it is the homeland, inexpensive, ground-based radars will be exploited. This is in stark contrast to defence environments which are typically limited in space and time, where there are only friends (who identify themselves) and foes who are relatively sparse, and where expensive airborne or space-based radars are required for theatre surveillance.

Convergence between radar sensing, communications, and information technology during the past 20 years has lead us to spatially distributed, heterogeneous radar information networks that will grow in dominance and utility in the 21st century. From a technology stand point, we are now in striking distance of being able to introduce powerful, cognitive capabilities into our surveillance machines. The vision for cognitive radar information networks presented here will hopefully excite stakeholders into action and drive innovation across human and machine sciences during the early part of this century.

ACKNOWLEDGMENT

The authors thank members of the Royal Canadian Mounted Police who have exchanged numerous ideas and challenges with Accipiter's engineering staff over the past decade that have provide the impetus for this work, and DRDC Centre for Security Sciences for their support and encouragement.

Figure 1: Model of wide-area radar information network for the Great Lakes with pre-defined attention cells indicated by polygons.

REFERENCES

[1] Nohara, T.J., "Affordable, Multi-mission, Radar Surveillance Networks for Marine and Port Security", Journal of Ocean Technology, Vol.4, No.2, July 2009, pp.29-38.

[2] T.J. Nohara, "A Commercial Approach to Successful Persistent Radar Surveillance of Sea, Air and Land Along the Northern Border", 2010 IEEE International Conference on Technologies for Homeland Security, pp.276-282, 8-10 November 2010, Waltham, MA.

[3] Pierre Meunier and Andrew Vallerand, "Border Integrity Capability: Enhancements of Multi-jurisdictional Situation Awareness on Lake Ontario during the G20", Defence R&D Canada – CSS, Technical Memorandum DRDC CSS TM 2011-12, June 2012.

[4] S. Haykin, "Cognitive Dynamic Systems, Perception-action cycle, radar and radio". Cambridge University Press, England, March 2012.

[5] S. Haykin, "Cognitive radar: a way of the future," IEEE Signal Processing Magazine, January 2006.

[6] S. Haykin, Y. Xue and P. Setoodeh, "Cognitive Radar: Step Toward Bridging the Gap Between Neuroscience and Engineering". Proc. IEEE,Vol.100,No.11, pp.3102-3130, November 2012.

[7] S.Haykin and T.J. Nohara, "Cognitive Radar Information Networks for Security Along Canada/US Border", White Paper, Prepared for DRDC CSS, 15 Oct 2012. www.accipiterradar.com/media/pdf/CRIN-Whitepaper.pdf.

[8] J.M. Fuster, "Cortex and Mind: Unifying Cognition". Oxford University Press, 2003

[9] R.S. Sutton and A. G. Barto, "Reinforcement Learning: An Introduction", MIT Press, 1998.

[10] Beyond the Border: A Shared Vision for Perimeter Security and Economic Competitiveness, A Declaration of the Prime Minister of Canada and the President of the United States, 4 February 2011, Washington, D.C..

[11] T.J. Nohara, "Radar Information Networks with Target Data Persistence for Intelligence-Led Border Enforcement", 2012 IEEE International Conference on Technologies for Homeland Security, 13-15 November 2012, Waltham, MA.

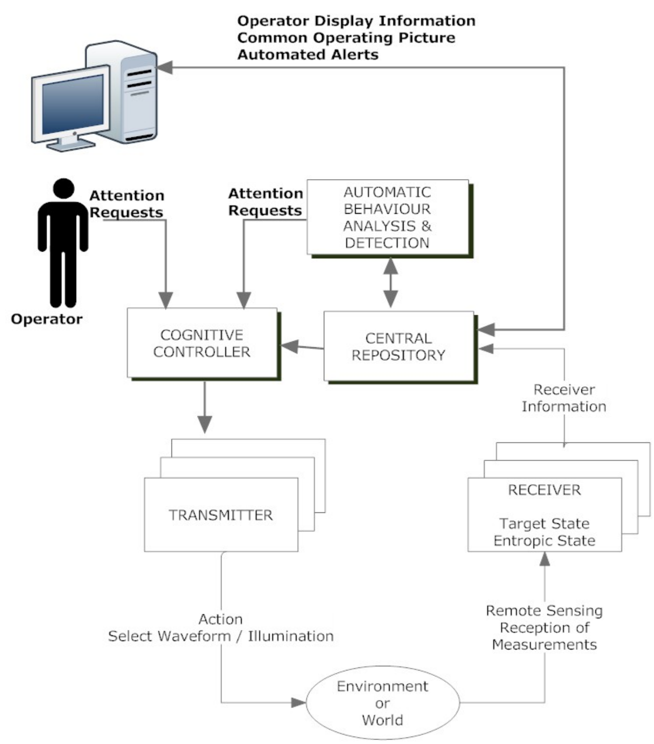

Figure 2: Practical Cognitive radar information network block diagram