A Cognitive Radar Network: Architecture and Application to Multiplatform Radar Management

Felix Smits, Albert Huizing, Wim van Rossum, Peter Hiemstra

TNO Defence, Security and Safety
Oude Waalsdorperweg 63, 2594 AK, The Hague
felix.smits@tno.nl
albert.huizing@tno.nl
wim.vanrossum@tno.nl
peter.hiemstra@tno.nl

***Abstract*—The objective of a cognitive radar network is to optimise radar performance in the highly variable mission environments that current operational systems encounter, while minimising its interference with other systems and its vulnerability to countermeasures such as jamming and anti-radiation missiles. A cognitive radar network may achieve these challenges by fully exploiting the available radar resources, sharing data among network components and taking into account prior environmental and situational knowledge as well as experience accumulated during operations. This knowledge can vary from high level information such as intelligence about the threat to low level information such as clutter maps. This paper presents a cognitive radar network architecture that supports this functionality and the application of (self)-learning methods. In this paper reinforcement learning is used to maximise the survivability of naval surface ships in a littoral scenario by managing the modes of an air surveillance radar.**

I. Introduction

In the past decades there has been a growing uncertainty about the missions and the threat environments in which radars have to operate. For example, air surveillance radars that have been designed to detect and track low flying hostile aircraft and helicopters in a land environment may suddenly be deployed in scenarios where rockets and mortars are the main threat. It may take skilled personnel days or weeks to modify an existing air surveillance radar in such a way that it is (partially) capable of detecting and tracking rockets and mortars instead of aircraft and helicopters. If humans would have had radar senses, the adaptation to the new threat environment would have been much quicker than what is currently feasible with radar systems. Owing to their cognitive abilities, humans can quickly learn from their experiences, they can incorporate knowledge obtained from previous experiences to improve their behaviour and adapt their behaviour adequately to new situations. It is the objective of this paper to investigate how the cognitive abilities of humans can be implemented in artificial cognitive systems to improve the performance of radar systems in unexpected situations.

Simon Haykin was the first to use the term cognitive in conjunction with radar [1]. According to Haykin there are three ingredients that are basic to the constitution of a cognitive radar: (i) intelligent signal processing, which builds on learning through interactions of the radar with the surrounding environment; (ii) feedback from the receiver to the transmitter, which is a facilitator of intelligence; and (iii) preservation of the information content of radar returns, which is realized by a Bayesian approach to target detection through tracking.

Cognitive principles have already been applied to artificial vision systems for a considerable time. The European Union has established the ECVision research network to promote research, education, and application systems engineering in cognitive computer vision. According to the definition by David Vernon on the associated website [2], a cognitive vision system can achieve four levels of generic visual functionality: detection, localisation, recognition and understanding.

A cognitive vision system exhibits purposive goal-directed behaviour, is adaptive to unforeseen changes, and can anticipate the occurrence of objects and events. In adddition to this, an artificial cognitive system should be able to explain what it is doing and why it is doing something [3].

A model of an observation system with different levels of functionality has also been identified in the JDL[1] model which is popular in the data and information fusion domain. The revised JDL model as defined by Steinberg and Bowman identifies four data fusion levels [4]:

- Level 0 : signal assessment
- Level 1 : object assessment
- Level 2 : situation assessment
- Level 3 : impact assessment

These levels roughly correspond with the levels defined by Vernon for a cognitive vision system. In the same paper, Steinberg and Bowman also extended the JDL model to the management of resources at different levels. In a further extension of the JDL model, Llinas et al. [5] identify the need for co-processing of abductive, inductive and deductive reasoning. While deductive reasoning already plays an important role in radar processing, abductive reasoning (i.e. discovery of patterns) and inductive reasoning (i.e. generalisation of patterns) are not yet widely applied but are an essential part of the self-learning and anticipation capability of a cognitive system.

[1] JDL: Joint Directors of Laboratories is a US DoD government committee. The Data Fusion Group of the JDL created the original JDL Data Fusion Model

978-2-87487-009-5 © 2008 EuMA

This paper presents a high level architecture of a cognitive radar network that allows information to be exchanged at various information abstraction levels and describes how various inference processes may be implemented in this architecture. The paper demonstrates how reinforcement learning (RL) techniques may be used to improve the performance of a network of air surveillance radars in a littoral environment.

This paper is organised as follows. Chapter II introduces a model of an artificial cognitive system. Chapter III describes a high level architecture of a cognitive radar network and the different levels at which cognitive radars in a network may communicate. Chapter IV demonstrates how a cognitive radar network can be used to perform multiplatform radar management using RL. Finally, conclusions are drawn and suggestions for future work are presented.

II. A MODEL OF AN ARTIFICIAL COGNITIVE SYSTEM

Figure 1 shows a hierarchical model of an artificial cognitive system that is based on the JDL model for data fusion and resource management. The branch on the left-hand side represents the fusion of information from various sensors and other information sources. The end-product of this branch is an awareness and understanding of the situation. The branch on the right-hand side represents the resource management process that is executed as a consequence of the mission objectives and the situation awareness.

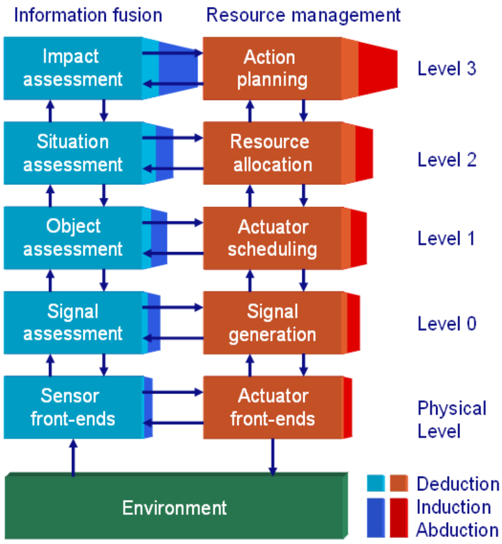

Fig. 1 Hierarchical model of a cognitive system with an information fusion branch on the left-hand side and a resource management branch on the right-hand side. At each level deductive reasoning and abductive and inductive reasoning may take place with the depth indicating the time and spatial horizon.

At the lowest (or physical) level, analogue sensor and actuator front-ends interact with the environment. Sensor front-ends may include radar receivers, wireless communications receivers, cameras, microphones, thermometers, etc. Actuator front-ends may include radar transmitters, wireless communications transmitters, antenna drives, engines, guns, etc. At level 0, digital sensor signals are processed (signal assessment) and digital actuator signals are generated. The estimation of object properties (e.g. location, velocity, and class) (object assessment) and the scheduling of actuator resources occur at level 1. Level 2 is involved with the estimation of relations between objects (situation assessment) and the allocation of resources to the different actuators. Finally, the highest level involves the prediction of the impact of relations between objects (including the system itself) and the planning of actions to reach the mission goals.

At each level of the cognitive system model, the following reasoning processes may be performed:

- abduction: discovery of patterns and generation of hypotheses
- induction: generalisation of patterns and validation of hypotheses
- deduction: detection of patterns and testing of hypotheses.

Abduction and induction are generally off-line processes while deduction is an on-line process. At each step in the model hierarchy, the scale of the spatial and temporal processing is roughly increasing with an order of magnitude, while the spatial and temporal resolution is decreasing with an order of magnitude.

III. COGNITIVE RADAR NETWORK ARCHITECTURE

This section describes, as a specific case of an artificial cognitive system, the architecture of a cognitive radar network. Figure 2 shows the architecture of a cognitive radar network in which two (or more) cognitive radars exchange information at various levels.

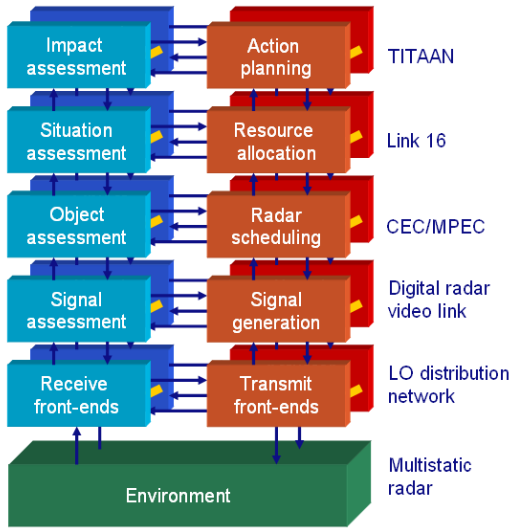

Fig. 2 Cognitive radar network architecture with examples of communication networks at each level.

At the highest level (level 3), information is exchanged about the current situation and the planned course of action. An example of a network supporting the information exchange at this level is TITAAN (Theatre Indepependent Tactical Army and Air Force Network) [8]. Deductive reasoning (on-line processing) at this level includes the evaluation of threats and in-situ radar performance prediction while the off-line processing (abduction and induction) may involve the analysis of radar performance during the mission and the forecasting of the impact of the weather on the mission.

At level 2, target tracks are exchanged and tasks are distributed to the radars in the network. Link 16 is a typical example of a data link at this level [9]. Deductive processing at this level includes the detection of hostile intent. Abductive and inductive processes at this level involves the analysis of track behaviour and RL for multiplatform radar management.

At level 1, radar plots are exchanged and information about the scheduled radar frequencies and beam directions. A typical example of a network for the exchange of radar plot data is the Data Distribution System of the Cooperative Engagement Capability (CEC) [8]. Target tracking and classification are examples of deductive processing while adaptive tracker settings and analysis of the radar time budget are examples of abductive and inductive processing at this level.

At level 0, digital radar signals may be distributed to provide a multiperspective detection, tracking and recognition capability. Very high data rates are required here and broadband data links such as the tactical common data link (TCDL) should be used. The detection of targets and adaptive beam steering are examples of deductive processing at this level while calibration of the receiver and transmitter front-ends is an example of off-line processing.

At the physical level, analogue signals such as local oscillators (LO) may be distributed through optical fibers. Matched filtering is a form of deductive processing at this level. Off-line processing generally does not occur at this level. Also, multistatic radar can be viewed as a radar network that exchanges radar information through EM field propagation.

IV. Cognitive Multiplatform Radar management

A. *Multiplatform radar management*

In this section we use level 3 and 2 of the cognitive radar network architecture to prepare and support a mission with three ships in a littoral environment where weapons such as anti-ship cruise missiles (ASCM), anti-radiation missiles (ARM), rockets and artillery may be deployed against them. Each ship is fitted with an air surveillance radar and an air defence system.

The air surveillance radars support a short range and a long range mode, see figure 3. The two surveillance modes can be interchanged from one 45° azimuth sector to the other. The air defence system consists of three weapon systems, each with different range. The hit probablity of the missile correspondingly decreases stepwise with detection range.

As shown in Fig.3, short range mode will detect high missiles at medium range and with correspondingly higher hit probablity than long range mode. In contrast, long range mode will not detect high missiles hiding behind high terrain during launch. In varying terrain, the best choice is not obvious: a balanced trade-off between certain but late detection and high hit probability against less certain early detection and low hit probability must be made.

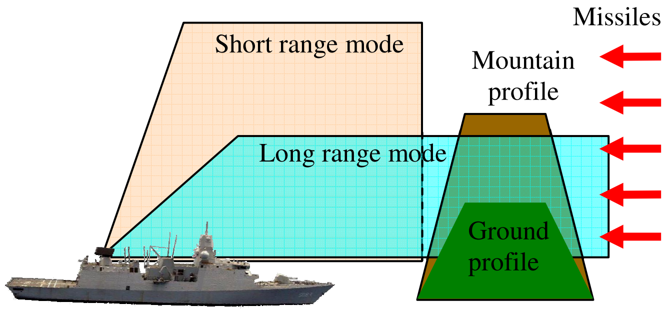

Fig. 3 Vertical coverage of the long range and short range modes of a volume surveillance radar, and the effects of terrain on the effectiveness against missiles of each mode.

Optimum combination of the defence coverage of three vessels in a complex littoral environment is even more difficult. The multiplatform radar management problem is to find the combination of long range and short range modes on board of the three ships that provides the best combined survivability in a littoral air defence scenario. The result could be of the type displayed in Fig. 4. These settings can be found by a self-learning technique as RL that uses vessel hit probability data collected during many Monte Carlo runs of the scenario to optimise the settings of the three radars.

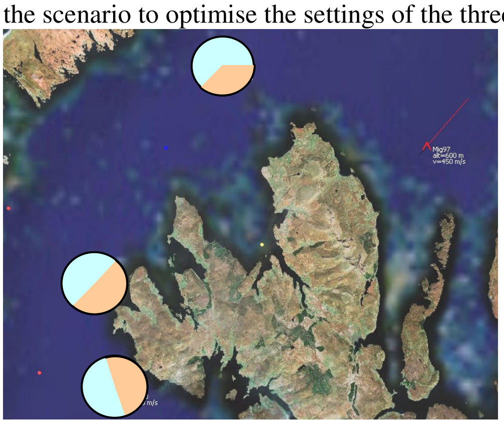

Fig. 4 Possible air surveillance radar mode settings as a function of azimuth for a three vessel mission in a littoral environment. The short-range mode is indicated in orange while the long-range mode is in cyan.

B. *Reinforcement learning*

RL is a form of machine learning that allows a cognitive system to take actions in an environment that maximizes some long-term reward [9,10]. RL algorithms try to find a policy that maps states of the world to the optimum actions the cognitive system should take in those states. The environment of the cognitive system is typically formulated as a finite-state Markov decision process (MDP). State transition probabilities and reward probabilities in the MDP are typically stochastic but stationary over the course of the problem. RL differs from supervised learning in that correct input/output pairs are not required, and sub-optimal actions are not explicitly corrected.

The reinforment learning problem is formally defined with: discrete sets of environmental states *S,* and cognitive system actions *A*, and a set of scalar reward signals *R.* The learning problem consists of finding a sequence of time-dependent decision rules δ_t, t = 1,…,T, that maps the sequence of environmental states to a sequence of actions. A sequence of decision rules is known as a policy P. A policy that minimizes the expected sum of the direct and future costs at any time t is called the optimal policy P^*. An optimal policy always exists for an MDP, and sometimes more than one exist.

To find an optimal policy, the RL algorithm requires a value function that gives for every policy P and time t the expected sum of the direct and future costs. The value function used for the multiplatform radar management problem is the risk function as defined by Bolderheij en van Genderen [11]. The risk at time t is defined as:

$$R(t) = \sum_{i=1}^{L} V(i) \cdot \left(1 - \prod_{k=1}^{K} \left(1 - L(i,k) \cdot P_{occ}(i,k,t)\right) \right)$$

where L is the number of assets, K is the number of threat objects, $V(i)$ is the value of an asset i, $L(i,k)$ is the lethality of threat object k against asset i, and $P_{occ}(i,k,t)$ is the probability of threat object k reaching asset i undetected at time t.

C. *Results*

Reinforcement learning requires large numbers of simulations. To allow for sufficient learning runs in a reasonable time span, a one-vessel problem is investigated by fast and simple models containing the major problem characteristics (Fig.5). The ship follows a 9-step trajectory along two coast lines with fixed mountain heigths (grey). Attacks are uniformly distributed over azimuth and heigth, except for the coast where they appear only above mountains. The human 'best guess' for the radar settings chooses short range towards the coast and long range towards open sea. In contrast, the RL optimum chooses short range towards the coasts up to much larger distances.

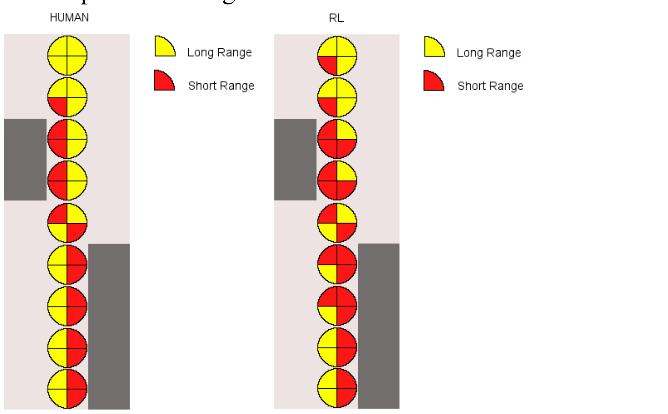

Fig. 5 Air surveillance radar mode settings as a function of vessel position as determined by a human, and as determined bij RL.

The importance of this result is that RL provided a policy which is not foreseen by a human, and has better performance. In Fig. 6, it is shown that the RL solution outperforms the human solution after about 20.000 simulation runs.

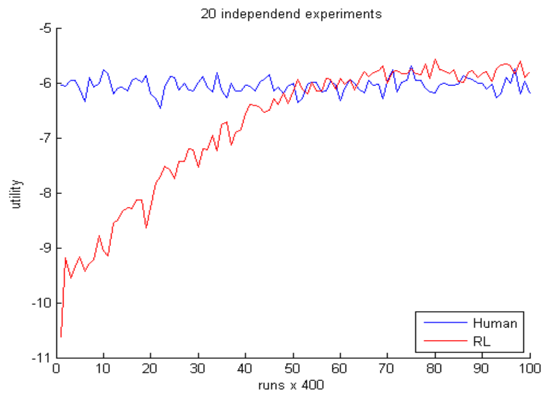

Fig. 6 Value function (utility) of an RL strategy as a function of vessel position as determined by a human, and as determined bij RL.

V. Conclusions

A new cognitive radar network architecture is presented building on the JDL model and its recent extensions, and incorporating cognitive capabilities. The architecture allows the exchange of information at different levels and examples of cognition are given for each level. An example of cognitive radar network usage for multiplatform radar management is given in a littoral scenario with three air surveillance radars. In a single vessel simulation, RL provided an unforeseen policy which outperforms the human policy.

Future work will concentrate on multi-sensor networked cognition and the impact of cognitive capabilities at other architectural levels.

VI. Acknowledgement

This work is performed as part of the Radar Research Programme V703 of the Dutch Ministry of Defence

References

[1] Haykin, S., "Cognitive radar" in *IEEE Signal Processing Magazine*, Vol 30 January 2006, pp. 30-40.
[2] Vernon, D. "Cognitive Vision – The Development of a Discipline" *ECVision research network*, www.ecvision.org
[3] Vernon, D.; Metta, G.; Sandini, G.; " Survey of Artificial Cognitive Systems: Implications for the Autonomous Development of Mental Capabilities in Computational Agents" *IEEE Tr. on Evolutionary Computation*, Volume 11, Issue 2, April 2007, pp. 151 – 180
[4] Steinberg, A., Bowman, C., "Rethinking the JDL Data Fusion Levels," in *Proc .National Symposium on Sensor Data Fusion, JHUAPL,* June 2004
[5] Llinas, J., Bowman, c., Rogova, g., Steinberg, A., Waltz, e., and White, F., "Revisiting the JDL Data Fusion model II," in *Proc. 7th Int. Conf. on Information Fusion* (Jun 2004), pp1218-1230.
[6] Royal Netherlands Army C2 Support Centre, TITAAN, www.c2sc.org/pages/products/titaan.htm
[7] Wikipedia, Link 16, http://en.wikipedia.org/wiki/Link_16.
[8] Applied Physics Laboratory, "The Cooperative Engagement Capability", *Johns Hopkins APL Technical Digest*, Volume 16, No. 4, 1995, pp. 377-396
[9] L. Pack-Kaelbling, M.J. Littman, A.R. Moore, "Reinforcement learning", *Journal of Artificial Intelligence Research*, Volume 4, 1996, pp. 237-285
[10] R. Sutton, A.G. Barto, "*Reinforcement Learning: An Introduction*" M.I.T. Press, Cambridge MA, 1998.
[11] Fok Bolderheij, Piet van Genderen, "Mission driven Sensor management", in *Proc. 7th Int. Conf. on Information Fusion* (Jun 2004), pp 799-804.