

Improved target detection and tracking in littoral environments using a self-organising spatio-temporal CFAR.

Evan J. Hughes and Michael B. Lewis

Cranfield University, DCMT, Shrivenham,

Swindon, Wiltshire, SN6 8LA

ejhughes@iee.org , m.lewis@computer.org

Abstract

This paper describes a self organising spatio-temporal radar CFAR system that uses multiple intelligent software agents to detect and adapt the processing to features in the environment. By combining both temporal and spatial data gathering sufficient samples can be collected to allow both the first and second order moments of the clutter distribution to be approximated for each cell. By gathering higher order statistics to a useful accuracy, more stable thresholds may be produced.

1. Introduction

This paper describes an improved method of target detection applicable to littoral environments where a wide range of clutter characteristics are present. Classic detection methods, such as cell averaging CFAR systems and clutter maps, attempt to gather a small number of spatial or temporal samples from around the range-azimuth cell of interest in order to estimate the local clutter and noise statistics. A threshold level can then be calculated against which the amplitude of the return in the cell of interest can be compared to determine the presence or absence of a potential target.

In general the homogeneity and stationarity of the clutter in the littoral environment is poor. If a large number of spatial samples is gathered, implying that the statistics are gathered over a wide area, the region around the cell-under-test must be clear of artefacts such as buoys, harbour walls, cliffs etc. When only a few samples are gathered, the resulting estimate of the mean will be poor and the calculated higher central moments, such as variance and skewness, will be highly inaccurate and often biased. The resulting poor statistical estimates mean that the detection threshold must be placed higher than the ideal to prevent excessive false alarms with the result that small targets are not detected. If a moderate spatio-temporal region is used to gather data for the statistical analysis, more points can be gathered and the estimates of the statistics will be more accurate, however there is also a risk of undesirable fixed targets falling within the region and corrupting the estimates of the statistics.

To overcome these problems a novel self-organising system based on the use of multiple intelligent software agents (MISA) has been developed and is an improved version of the system described in [1]. The key concept is the exploitation of the spatio-temporal coherence of true target tracks, but with practical levels of processing. The agent system detects features in the environment and modifies the areas over which the statistics gathering processes are performed accordingly such that the spatio-temporal data gathering is more effective. The system has been further coupled to an agent-based pre-tracker which allows a depressed threshold to be used and therefore low-observable targets to be detected and tracked in a complex littoral environment, whilst also extracting information on the location of fixed targets etc. The key design philosophy has been to recognise that as the statistics of the scene are changing too rapidly to allow calculation to sufficient accuracy, any processing that is applied can only ever be sub-optimal. Thus a tracking system has been designed where sub-optimality is assumed, but the effects of

sub-optimal processing have been carefully considered and controlled, leading to a highly effective, robust algorithm.

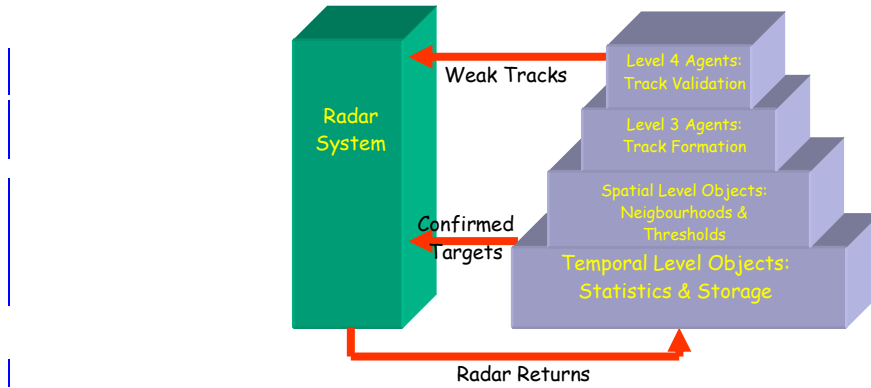


Figure 1 Functional Arrangement of System

The system architecture is based on a hierarchical structure of layers of objects and intelligent agents. Each agent or object represents an individual radar cell that is allowed, in conjunction with other cells, to self-organise into target tracks. **Figure 1** shows a functional block diagram of the system. The radar system is shown on the left, feeding the radar returns into the lowest levels of the hierarchy. The radar returns at this point will have had all necessary processing applied prior to the application of a CFAR system and a threshold.

The Temporal and Spatial Level objects form the Spatio-Temporal CFAR Subsystem whilst Levels 3 and 4 function as a multiple hypothesis track forming sub-system. The radar returns traverse the hierarchy, with high-confidence target detections being fed to the main radar tracker as track segments.

The Level 3 and 4 pre-track system attempts to associate the returns with previous returns according to a set of simple rules that define the likely feasible region that previous returns could lie in. The pre-track system does not make any explicit track predictions, unlike conventional multiple hypothesis trackers, but relies on associations between returns producing ‘virtual’ tracks within the data.

Section 2 describes the operation of the self-organising spatio-temporal CFAR algorithm, section 3 presents example results of CFAR processing of simulated radar data. Sections 4 and 5 give a brief description of the tracking levels, and section 6 concludes.

2. The Self-Organising Spatio-Temporal CFAR Subsystem

The Temporal, or T , Level cells are arranged as elements of a range-azimuth map. Each cell contains two identical IIR filters that perform temporal integration of successive target returns and its square from the point represented by the co-ordinates. The IIR filter that calculates the mean is described by the following recurrence relationship

$$T_{\mu}(R, \theta, t) = \frac{0.9T_{\mu}(R, \theta, t-1) + I(R, \theta, t)}{1 + 0.9}$$

Where $T_{\mu}(R, \theta, t)$ is the temporal mean at each range, azimuth and time, $I(R, \theta, t)$ are the new raw input data. The filters produce the sum of exponentially decaying contributions from previous radar returns.

A similar filter, $T_{\sigma}(R, \theta, t)$, that sums the squares of the input voltages is also applied with $I(R, \theta, t)$ replaced by its square. Thus the variance (and therefore standard deviation) may be calculated as $T_{\sigma}(R, \theta, t) - T_{\mu}(R, \theta, t)^2$. The temporal IIR filters can also be described by the following z -transform transfer function:

$$\frac{T_{\mu}(R, \theta, z)}{I(R, \theta, z)} = \frac{0.526}{1 - 0.474z^{-1}}$$

The integration algorithm may be implemented in efficient ways, for example to exploit multiply-accumulate instructions within digital signal processing devices.

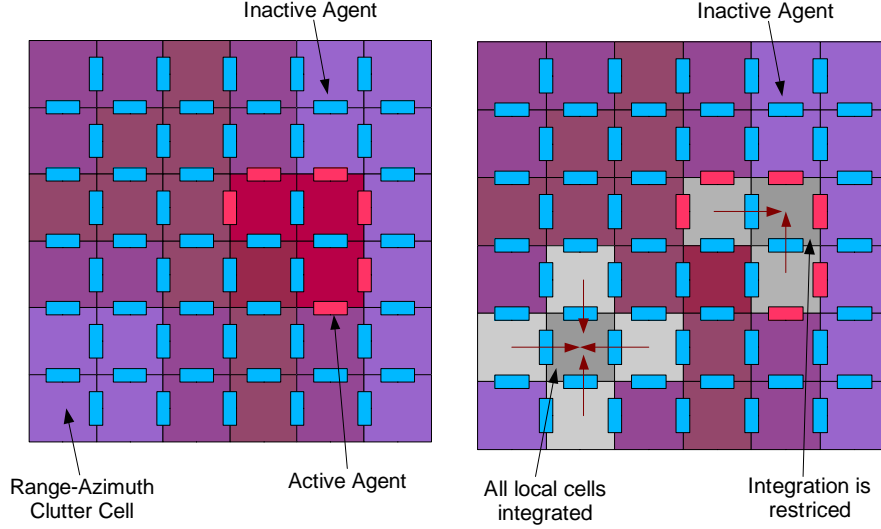


Figure 2: Layout of cells and agents, and agent operation in restricting spatial integration

The range-azimuth cells are also part of the Spatial layer. The purpose of the spatial layer is to perform a spatial integration across regions of homogenous clutter. A means of adapting the regions over which spatial integration is performed is incorporated within the layer.

Each range-azimuth cell has 4 intelligent agents around its borders, the bridging or B agents, shared with its neighbours, as shown in Figure 2. The B agents prevent the spatial integration from being disturbed by fixed targets. Each B agent monitors the $T_{\mu}(R, \theta, t)$ and $T_{\sigma}(R, \theta, t)$ values of the cells on either side of it, and if either $T_{\mu}(R, \theta, t)$ or $T_{\sigma}(R, \theta, t)$ are consistently different, it switches to a blocking state and prevents spatial integration occurring across the boundary. Each agent maintains μ and σ values, the μ value being:

$$B_{\mu}(R+, \theta, t) = 0.9B_{\mu}(R+, \theta, t-1) + \text{sgn}(T_{\mu}(R, \theta, t) - T_{\mu}(R+1, \theta, t))$$

Where the notation $B(R+, \theta, t)$ denotes the agent that lies between cells (R, θ) and $(R+1, \theta)$ etc. The agent $B(R, \theta+, t)$ is the equivalent in the orthogonal grid direction. The process may also be extended to include Doppler and Elevation dimensions. The use of the signum function rather than the raw difference results in an indication of the median rate of dissimilarity rather than the mean of the difference between the agents.

The decision as to whether the agent should block or not, $B(R, \theta+, t)$ etc., is generated by first identifying the B agents which separate cells having the greatest dissimilarity (one agent for T_{μ} data and one for T_{σ}). Thus the agent with the largest magnitude for the difference between means, and similarly the agent with the value with the largest magnitude for the difference between the squared returns are identified. The magnitudes of these two values are then used to set a threshold to determine the bridging agent's activity. The agent will record $B(R, \theta+, t) = 0$ if either the value of $|B_{\mu}|$ or $|B_{\sigma}|$ is greater than 70% of the respective maximal values. It will record a 1 otherwise.

Expressed in formal logic the truth value for the blocking action, for a single azimuth B agent is

$$B(R, \theta, t) \Leftrightarrow \neg \left(B_\mu(R, \theta, t) > 0.7 B_\mu^{(\max)} \mid B_\sigma(R, \theta, t) > 0.7 B_\sigma^{(\max)} \right)$$

Where TRUE and FALSE correspond to 1 and 0 respectively.

The state of the B agents surrounding each range-azimuth cell can also be used to infer which range-azimuth cells may be fixed targets or other discontinuities such as harbour walls, coastline etc.

The integration of the means is then described by:

$$S_\mu(R, \theta, t) = \frac{0.9 S_\mu(R, \theta, t-1) + \sum_4 \left(S_\mu(R \pm 1, \theta \pm 1, t-1) B(R \pm 1, \theta \pm 1, t) \right) + 0.7 T_\mu(R, \theta, t)}{0.9 + \sum_4 B(R \pm 1, \theta \pm 1, t) + 0.7}$$

The integration of the squared returns is performed in a similar manner.

A threshold is calculated based on the S results and used to threshold the input data in I . To prevent moving targets from disrupting the mean and standard deviations, target detections are censored. The censoring process simply prevents T level updates for any cells in which detections have been made.

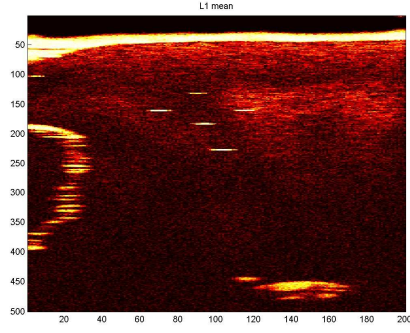
The controlled spatial integration allows more samples to be gathered and more stable and accurate estimates of mean and variance to be obtained with edges in the scene preserved as sharp discontinuities. This process allows accurate thresholds to be determined to within a few cells of features within the environment.

3. Example Results

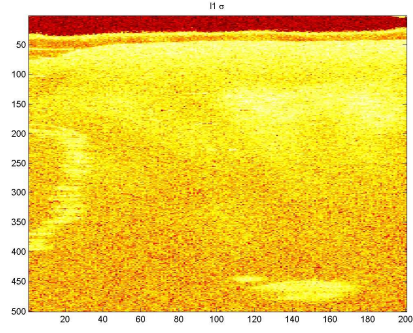
The processing has been applied to simulated radar data modelled to resemble the output from a low-cost non-coherent marine radar. The scene model is a realistic simulation containing radial, crossing and spiralling targets moving amongst fixed targets and through heavy sea clutter regions.

In the real marine radar used as a basis for the simulation, the radar returns pass through a logarithmic input amplifier. In the simulation it has been assumed that the underlying clutter power distribution is a Weibull distribution (the simulation is actually a compound noise distribution, not true Weibull) which the logarithmic amplifier transforms to a Fisher-Tippett distribution. This has proved to be a good general assumption when applied to the real radar data. The threshold level for detecting targets is calculated as the S mean plus a scaling factor times the S standard deviation. The scaling factor is adjusted dynamically to maintain a reasonably stable false alarm rate.

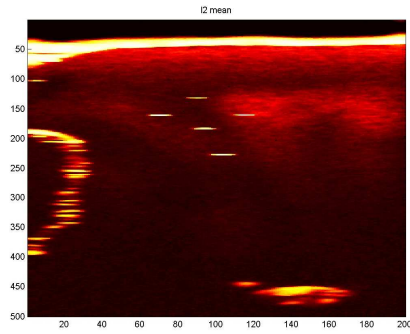
The figures below show the algorithm behaviour on one example scan of the processing.



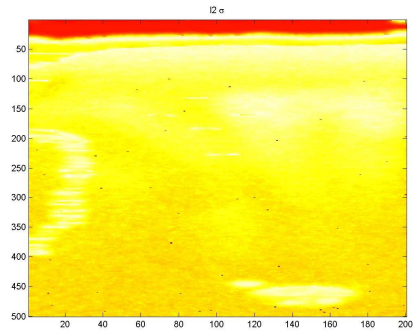
T_{μ}



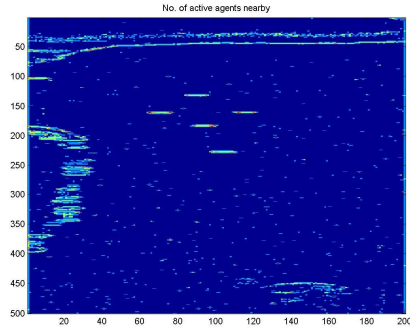
T_{σ}



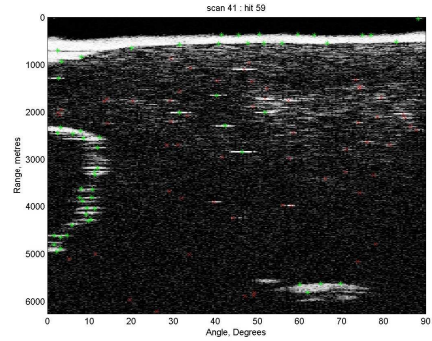
S_{μ}



S_{σ}



B Agent Activation Profile



Likely Fixed Targets

Figure 3. Examples of Returns at the Various Levels

The T_{μ} contains a significant amount of noise as only temporal filtering has been applied, but sharp edges in the scene are preserved exactly. The T_{σ} is extremely noisy and could not be used directly for calculating the threshold level. The S_{μ} has had restricted spatial integration applied and is far more stable. The action of the agents has restricted the integration and preserved the sharp edges in the scene. The S_{σ} is now stable enough to be used for setting useful threshold levels. The agent activation profile shows where strong edges were identified and if the active cells are clustered, the results can be used to identify likely regions of fixed targets.

4. Potential Track Formation, Level 3 Agents

Conceptually Level 3 agents are formed with each being associated with a target detection. When a Level 3 agent is created, it strives to form links with existing Level 3 agents that represent *virtual* tracks within the multi-agent system.

The ‘Agent is a detection’ approach allows many track hypotheses to be formed for each return and it is assumed that many tracks could pass through each Level 3 Agent. If Doppler information is available it may be incorporated easily.

Agents marked as having the potential to be part of a track are scanned to see if any previous links are recorded. If links exist they are checked to determine if the speed and direction changes are within a reachable set. The calculation of the reachable set for association of agents to allow links to be formed whilst keeping processing to an absolute minimum is one of the cornerstones of this research.

5. Track Validation, Level 4 Agents

The primary function of a Level 4 agent is to assess the most likely path through a series of Level 3 agents and report the track to the main track database if it appears to be a true target. Level 4 agents are created when potential tracks are identified as a sequence of links formed between Level 3 agents. The Level 4 agent scans the track, looking for all the necessary correlations between stages that indicate a valid track is likely and eliminates unlikely tracks in the process. This process allows crisp tracks to be confirmed, some noise to be rejected, and areas of uncertainty to be identified.

Once a track has been validated the track’s elements are passed to the main radar tracker and the corresponding Level 3 agents notified that the track has been validated

6. Conclusions

The self-adaptive spatio-temporal CFAR is proving to be very effective at gathering large numbers of statistically homogeneous data samples from complex and difficult environments. The ability to gather large sample sizes means that robust estimates of threshold locations can be generated, reducing fluctuations in false alarm rates and allowing depressed thresholds to be used in combination with a pre-track system. Even though the approach is essentially cell-averaging CFAR, the performance is proving to be extremely reliable in complex environments and processing losses are small as accurate threshold locations can be calculated.

The system has a low memory and processor overhead and runs easily on a desktop PC.

Acknowledgements

The prototype MISA pre-track system was developed as part of CRP programme contract no. RCSC/2/040. The extension to MPRF radar was funded by the EMRS DTC under contract no. EMRS/DTC/2/53.

In addition we wish to thank Dr., Sam Lycett, Qinetiq, Malvern and Dr., Andy Stove, Thales, Crawley for supplying radar return data that enabled the prototype MISA pre-track system to be evaluated

Reference

1. E.J. Hughes and M. Lewis, “An Intelligent Agent Based Track-Before-Detect System Applied to a Range and Velocity Ambiguous Radar”, *EMRS DTC 1st Annual Technical Conference*, Edinburgh, May 2004

Improved target detection and tracking in littoral environments using a self-organising spatio-temporal CFAR

Hughes, Evan J.

2009-07-14T14:58:07Z

<http://hdl.handle.net/1826/3467>

Downloaded from CERES Research Repository, Cranfield University