

Adaptive Spatio-Temporal CFAR & Multiple-Hypothesis Tracking System.

E. J. Hughes, M. Lewis
Cranfield University at The Defence Academy of the UK
Shrivenham, Swindon, 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.

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. A table of confirmed and potential tracks is then used to classify the target returns into valid detections for existing tracks, possible targets worth investigating or noise.

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.

With low observable and low flying targets (where multi-path can cause significant fading), many returns will be below the detection threshold and there may be many missing detections along the track, resulting in targets being classified as noise if re-investigated, tracks never being initiated, tracks being deleted early or each track being maintained for an extended period. In order to increase the probability of detection of weak targets, the detection threshold must be lowered with a consequent increase in the number of false alarms.

Medium PRF radar systems allow all-round measurements of both the range and waveforms that are ambiguous in range, Doppler or both. Existing techniques that resolve these ambiguities require the number of detections input to the ambiguity resolution process to be kept to a small number, as otherwise the number of false correlations ('ghosts') becomes unworkably large.

As the information from the received signal is limited, a false alarm or a 'ghost' must be treated as a true target, until it can be established as false. The increased false alarm rate causes problems with the association of returns with tracks and leads to an excessive number of false tracks being reported with the consequent risk of the tracking system becoming overwhelmed.

Philosophy of the CFAR and Tracker

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 reacts to features in the environment according to simple rules and modifies the areas over which the statistics gathering processes are performed accordingly such that the spatio-temporal data gathering is more effective. In particular the statistics are gathered over regions of homogenous clutter.

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

Doppler of targets in high clutter environments to be made. Such radars use 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.

The Pre-Tracker Architecture

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.

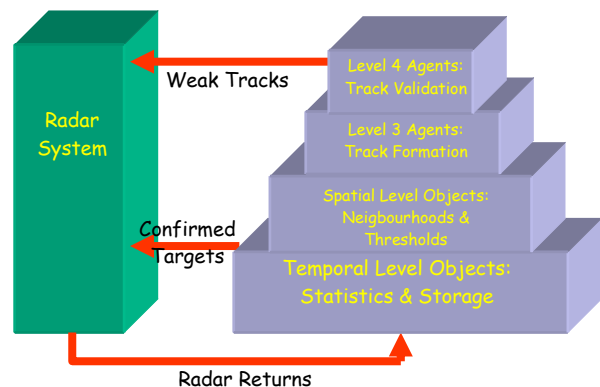


Figure 1 Functional Arrangement of the System

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.

Two thresholds are applied within the system. An upper threshold derived in a manner similar to that of cell averaging

CFAR and a lower threshold controlled by the Temporal and Spatial Level objects. All detections that cross the upper threshold are passed to the radar for processing as likely targets using the existing track algorithms. This guarantees that performance is no worse than conventional CFAR.

The Temporal and Spatial Level objects form the Spatio-Temporal CFAR Subsystem whilst Levels 3 and 4 function as a multiple hypothesis track forming subsystem. 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.

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 range-azimuth cells are 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.

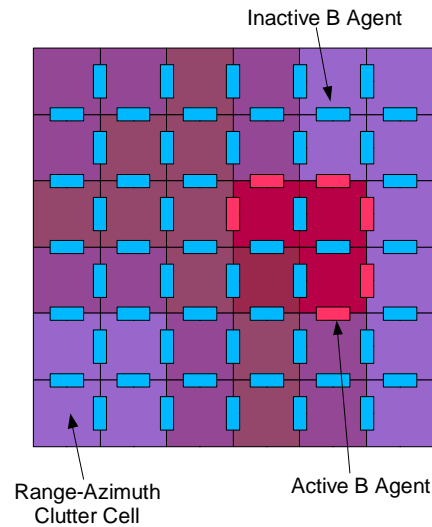


Figure 2: Layout of cells and agents

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:

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

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.

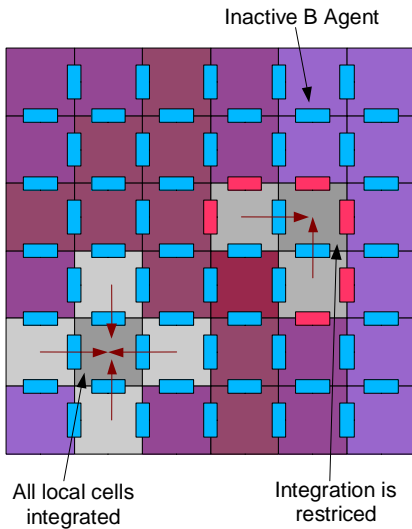


Figure 3: Agent operation in restricting spatial integration

The decision as to whether a B agent should block or not is generated by first identifying the B agents which separate cells having the greatest dissimilarity (one B agent for T_{μ} data and one for T_{σ}). Thus the B 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. Figure 3 is an illustration of the blocking action.

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

$$B(R, \theta+, t) \Leftrightarrow \neg \left(\begin{array}{l} B_{\mu}(R, \theta+, t) > 0.7B_{\mu}^{(\max)} \\ |B_{\sigma}(R, \theta+, t)| > 0.7B_{\sigma}^{(\max)} \end{array} \right)$$

Where TRUE and FALSE correspond to 1 and 0 respectively.

The integration of the means is then described by:

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

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

The spatial and temporal integration may also be expressed in three dimensional, R, θ and t, z -transform form and for the means, ignoring the effect of the blocking agents, is

$$\frac{S_{\mu}(z_R, z_{\theta}, z_t)}{T_{\mu}(z_R, z_{\theta}, z_t)} = \frac{0.7z_R^{-1}z_{\theta}^{-1}}{5.6z_R^{-1}z_{\theta}^{-1} - \left(\begin{array}{l} 1 + 0.9z_R^{-1}z_{\theta}^{-1} + z_R^{-2} \\ + z_R^{-2}z_{\theta}^{-2} + z_{\theta}^{-2} \end{array} \right) z_t^{-1}}$$

For a single impulse the recursive form of the expression suggests that the effect of the impulse will decay exponentially the greater the distance from the cell in which it originated.

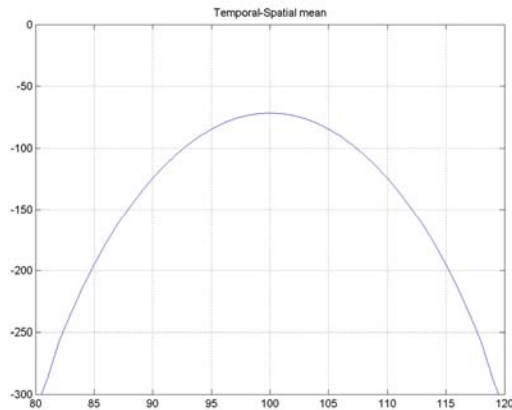


Figure 4. Spatial Impulse Response of CFAR

Figure 4 shows the Spatial Impulse response of the CFAR. The vertical scale is db, the horizontal scale is cell no. The impulse has little effect due to the spatial integration process utilising uncorrupted data from the cell's neighbours.

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.

Potential Track Formation, Level 3 Agents

Conceptually Level 3 agents are formed with each agent being associated with a target return. 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.

If the new agent is within the feasible region, the importance of the link is calculated. This value can be used to prune the link set of the agent to reduce storage requirements.

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. The Level 4 agent may also interrogate and analyse the target returns along the track in order to aid the track assessment by identifying possible missed detections. The Spatial and Temporal Object system is interrogated to see if a 'near miss' occurred when the data was thresholded. If a return is classified as belonging to a valid track at any time the Spatial and Temporal Level return may be promoted, the detection classification held in the temporal record being updated and the statistics describing the clutter updated accordingly. This process allows crisp

tracks to be confirmed, some noise to be rejected, and areas of uncertainty to be identified.

As the number of agents reaches the upper limit of the processing capability, the life of the agents can be managed to allow a maximum population size to be maintained, whilst performance is allowed to degrade gracefully. This contrasts with conventional track formation where track overload can be catastrophic.

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.

Project Status

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 Levels of processing required to implement the Spatial and Temporal systems has been investigated and the

capabilities of the full self-adaptive spatio-temporal CFAR system demonstrated. Multi-agent code has been written which has allowed the full dynamic threshold control system to be integrated with the Level 3 process and tested. The results, when compared against conventional methods including Cell Averaging CFAR, indicate that the multi-level system has the potential to provide a very significant pre-track capability.

Further Potential Applications

The ability to classify areas of returns is seen as having potential ECCM applications. The technique could be extended to IR and EO systems. It also has the potential for processing images in particle physics and astronomy.

Conclusion

Many existing CFAR approaches will produce very good results if the clutter statistics are known exactly, but can perform badly if there is even a small error in the estimated parameters. The result is that current CFAR techniques, by attempting to provide an optimal solution, can create a very fragile process.

In contrast the MISA process is, in effect, a simplified multiple hypothesis tracker, tightly coupled to a self-adaptive, context sensitive, spatio-temporal CFAR system. In environments with diverse clutter characteristics, the self-adaptive nature of the agent system self-organises using simple processing and by assuming that there will be too few data measurements to establish the clutter statistics accurately, a robust sub-optimal solution is formed.

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 requirement 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

References

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*