Advances in Radar Tracking for Small Target Detection

Presented at 12th International Radar Symposium (IRSI-19),

Bengaluru, 10-14th December 2019

Dr D G Johnson, Dr A D Wild

Cambridge Pixel, Cambridge, UK

Abstract

With growing interest in the detection of small targets as potential threats, advances in tracking algorithms for surveillance radars are described. A multi-layered approach is presented, which combines multiple hypothesis tracking (MHT), model-based tracking (MBT) and an ability to switch the track processing from plots to video to extend the tracking range.

Introduction

Surveillance radars may be required to detect a wide range of target types including surface and air targets. Where the sensors have a primary mission for vessel traffic or air surveillance, for example, they may increasingly be used to provide a secondary role in security. This includes the detection of small surface vessels or small air targets, including drones, which demonstrate behaviour of concern. Where an existing radar sensors can be dual-purposed to provide enhanced capabilities, there are obvious cost savings associated with the equipment, installation, maintenance and processing. This paper presents recent developments in target tracking technology that is particularly suited to dual-use radar sensors looking for small targets. The approach permits data to be analysed as separate processing channels that can be configured for different requirements. The paper is primarily concerned with rotating surveillance radars that typically provide 360 degrees of coverage. However, the techniques are also applicable to electronically scanned radars with controlled scan angles.

Radar Tracking – Terminology and Data Flow

The target tracking process starts with radar video which is analysed to identify potential target returns called plots. The plots are characterised by a position, size, intensity and, where appropriate, a Doppler value. Plots are then used to update existing tracks and, where there is no existing track, to start a new track.

The configuration of the processing is a key stage in the set-up of the tracker for a given installation. Some degree of user adaption is necessary to tune the track processing for specific types of target or situations. Additionally, some degree of automatic adaption is needed so that the track processing automatically adjusts to changing environmental conditions. For example, this includes dynamic adjustment of the detection threshold to maintain a constant false alarm rate (CFAR).

The key components of the processing are summarised by the following stages:

Radar acquisition

Radar video may be acquired as a network stream or as radar signals. The ASTERIX CAT-240 standard is a network format that is supported by a growing number of mid and high-end surveillance radars. Other network formats are proprietary, such as those from Simrad, Furuno and Raymarine. These can be accommodated with a dedicated translator to a common format. For radars that provide signals (video, trigger, azimuth) a hardware digitiser card provides the interface and conversion to digitised data packets.

Pre-processing

The pre-processing of the radar video may be split between the radar's own processor and the front-end of the track processing. Low-level functions such as filtering, sea clutter removal may be handled by the radar's own processor.

Area masking

An area of interest may be identified and used to restrict the processing to a geographic area. This area may be defined statically, for example to remove fixed clutter from land-based radar applications, or may be defined dynamically to remove coastlines and land areas in applications of a ship-borne naval radar. For dynamic masking, the lat/long of the ship is used with a world map data base to compute the position of land relative to the platform.

Detection

The detection process compares the amplitude of each radar sample against a locally-derived threshold. A sample is passed on if it exceeds the threshold. The calculation of the threshold is sensitive to the statistics of the local area (mean and standard deviation), thereby permitting automatic adjustment of the level in the face of changing background noise.

Plot Extraction

Radar samples that pass the detection test are passed through to the plot extraction stage. The plot extraction process combines connected radar samples to form plots which are defined by a weighted centroid, size, amplitude and Doppler value. Plots are stored in a plot database. A plot is extracted if the connected video (which passes the detection test) passes a set of tests based on size and amplitude. These are called qualifying plots. Plots that do not pass the size tests and still created, but they are marked as non-qualifying. The tracking process can access to both qualifying and non-qualifying plots.

Track Processing

An existing track is updated by considering new plots in the plot database. The update process incorporates a filter since plots are incomplete or imprecise observations of the true target position.

Plots that are unused for update may be used to create new tracks. An example display of video, plots and tracks is shown in Figure 1.

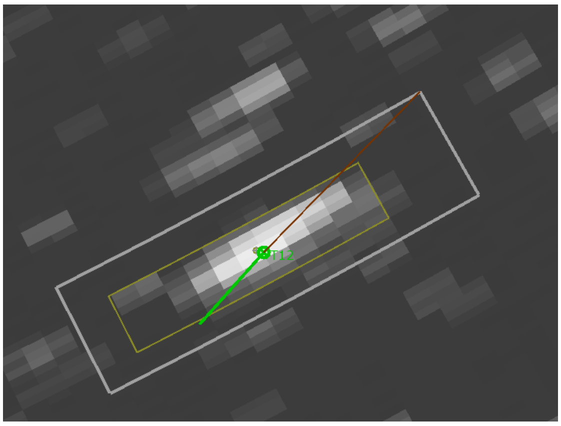

Figure 1 - A screen shot from Cambridge Pixel's SPx Server radar processor shows radar video with the extracted plot (yellow box), the track symbol (green) and the track's search gate (white rectangle)

Tracking Problems for Small Targets

Detection and tracking of small targets is of growing interest as these may represents threats, but there are a number of challenges, which may include some of all of the following:

Small radar cross section

A low radar cross section target results in a small echo, which in the limit may be comparable to the background noise or clutter. Ensuring that the processing is able to detect these weak plots means that many background noise plots will also pass through as a detection. The separation of targets from clutter can then be performed by subsequent stages of processing, but the generation of potentially tens of thousands of plots risks overloading the processing.

Intermittent Detection

A key tool in the separation of targets from clutter is correlation of the detections from scan to scan to look for detections in similar positions that are consistent with permitted target motions. It is expected that a true target will provide a sequence of detections in related positions, whereas random clutter has no regular pattern. By considering detections over sufficient time it should always be possible to build up enough confidence to distinguish real from false tracks. In practice, a real target, especially a small one, cannot be assumed to be observed on every scan. The probability of detection can be somewhat less than 1, with a value of 0.5, for example, being typical and indicative of a small target that is visible on average only 50% of the time. This intermittent detection makes it especially difficult to identify a small weak target in the face of high levels of background clutter.

Ambiguous Interpretation

Radar, like any sensing process, is imperfect. The observed radar video from a target can change considerably according to the environmental conditions, sea-state and orientation of the target. This is exaggerated for very small targets, which may be seen only in a single pulse and over a small number of range samples. The resultant plot can therefore differ from scan to scan from the true position. An apparent change in direction of a target may be no more than a change in its orientation to the radar. Care must be taken to interpret errors in expected plot positions with care. It is highly desirable to detect a target manoeuvre as quickly as possible to ensure accurate tracking, but being over-sensitive to errors can mean that manoeuvres are declared when it is simply a spurious measurement that would have been best filtered out.

Strategies for Small Target Tracking

Several strategies can work together to alleviate some of the problems above. The key strategies are defined in the following sections.

Multi-hypothesis tracking (MHT)

In multi-hypothesis tracking several interpretations of the measurement are allowed to be considered in parallel, with the expectation that future data will help resolve the uncertainty. In contrast, a single hypothesis tracker will need to make a firm decision on how to interpret data. For example, if an observation is some distance from the expected position it may be suggestive of a target manoeuvre, which would imply increasing the filter gains to improve dynamic response. The single hypothesis tracker would need to decide whether or not to declare a manoeuvre. There is a cost in not declaring a manoeuvre when there really is one – the target may move away from the expected position and hence not be observed. Equally, there is a cost in declaring a manoeuvre when there isn't one – the track filtering will become more erratic as a result of the higher filter gains. Being forced to make a decision may mean the wrong decision is made. In contrast, the multi-hypothesis tracker will consider several possibilities in parallel. One of the two will be preferred for the purposes of reporting a track position, but internally the tracker will process all options until it becomes clear what the correct interpretation is.

Model-Based Tracking (MBT)

The parameters of the tracking process must be tuned to match the types of targets being detected. The more restricted the tuning, the better the ability of the tracker to distinguish true targets from noise. This is similar to the effect of having a narrow-pass filter that searches for specific frequencies in a complex waveform. If it is desired to find different specific frequencies, then multiple narrow-pass filters may be the preferred option. This is case with tracker models. The MBT approach defines a number of disparate tracking configurations, or models, that independently process the radar data with their own set of parameters to look for specific conditions. Restricted models can successfully process very complex data sets looking for very specific conditions of interest. As an example, a model may be configured to search for targets moving only in a specific direction (eg towards the radar location) and a certain speed range. This condition defines a restricted search space so that when a new candidate target is first acquired, the expected position of the target for the next detection is a small area. If nothing is seen in that area the candidate track can be deleted. The more restricted the conditions for the targets of interest, the smaller the search area. In this situation, even tens of thousands of plots are easily handled because almost none of them are in the correct location for a target of interest.

Radar Video Tracking (RVT)

As described earlier, the normal processing chain for target tracking has a set of video processing stages followed by plot extraction. The track update process uses plots as the basis of the updates. As the radar echo for the target becomes weaker it will become progressively more difficult to reliably detect a plot. The thresholds for plot detection can always be lowered, but then it becomes easy to be swamped by an excess number of plots. Where a target is already being tracked and has some confidence, another approach to updating the track is to bypass the plot extraction process and refer directly back to the video. The raw radar video can be analysed to estimate the best position of the target. A simple approach is to compute the centroid of the video in the search area and assume that, however weak, any non-zero video in the gate will give an estimated position for the target which is better than no estimate. However, the centroiding must ensure that there is video in the window that is distinct from the background. A confidence level is calculated to decide whether there is sufficient video, and in a sufficiently compact form, to resemble a target. It is important not to generate an update on background noise as this would simply cause the tracker to track the background noise. In practical terms, it has been observed that switching the track processing from plot updates to video updates can significantly extend the ability of the tracker to follow weak targets beyond their normal range of detection with plot updates only.

Tracking Strategy

Individually, the MHT, MBT and RVT modules provide important benefits to assist with small target detection and tracking. However, using them together multiplies the benefit and significant performance improvements have been observed in the acquisition of small fast-moving targets.

As an example we consider a data set derived from a surveillance radar (from Kelvin Hughes) with a small target - see Figure 2- moving towards the radar. The small target is weak and barely visible in the video. It has a probability of detection of around 0.7, which we use to mean that for 30% of the scans there is no video derived from the target. When the target is visible, the video level is very low and it is no larger than the general level of background noise.

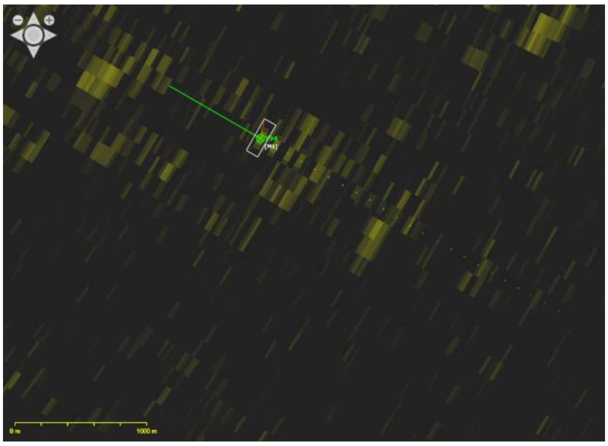

Figure 2 - Weak, fast moving target. The video is very low level and small with a Pd of about 0.7. A suitable MBT model is tuned to detect targets moving in a specific direction so the search space is greatly reduced.

When the plot extractor is configured to extract plots for the small size of the candidate target a large number of plots are observed, the vast majority corresponding to the background noise. In Figure 3, the plots are displayed.

Figure 3 - Approximately 10,000 plots are processed per scan to look for inbound targets. The above screen shot shows the plots drawn as a simple cross (non-qualifying plots) or a square (qualifying plots).

The plots are input to an MBT model which is tuned to look for targets moving in a specific direction (in this case towards the radar). Provisional tracks are created and maintained for each plot, with each track then looking for evidence of a plot in a small search space that would be consistent with the motion of interest. Most candidate tracks are quickly deleted, but for genuine targets of interest, albeit weak and intermittent, confidence grows and eventually reaches a threshold where a new established track can be created – Figure 4**.**

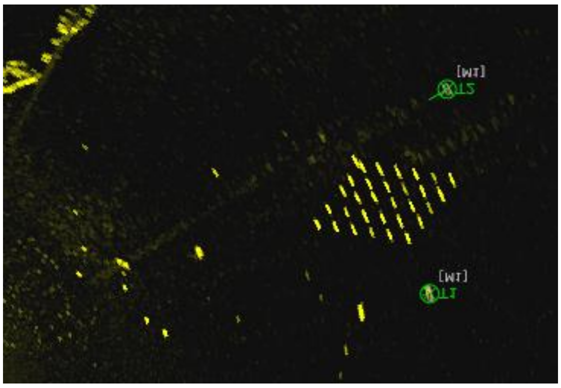

Figure 4 - Radar video, showing superimposed tracks that are moving in the direction of interest (towards the radar). These tracks are extracted with the MBT process using a model that is tuned for directional targets. Approximately 5,000 possible tracks are considered at any one time.

Once a track is acquired, the use of multiple hypotheses provides a robust mechanism that can handle the measurement errors, which are notably larger for very small targets. Finally, the ability to switch between plot-based and video-based tracking permits the small target to be continually tracked as the probability of detection diminishes at longer ranges.

Conclusions

There is a growing demand for increased awareness in detecting small surface and airborne targets, which may represent threats. Radars that are deployed for maritime or air surveillance purposes may be dual-purposed to permit the same radar video to be analysed for small targets. By combining techniques of multi-hypothesis tracking (MHT) and model-based tracking (MBT), as in Cambridge Pixel's SPx Server radar processor, specific target profiles can be detected and tracked in the video.

About the Authors

Dr David Johnson graduated from the University of Hull with a BSc and Phd in Electronic Engineering. As a software engineer and engineering manager, he has worked in industry developing image and radar processing solutions for over 20 years. Dr Johnson is the CEO or Cambridge Pixel and is actively involved in the development of the company's cost-effective software solutions for radar processing and tracking applications.

Dr Adrian Wild graduated from Salford University with a B.Sc. in Physics and a Ph.D. in Radar Meteorology. His experience of radar systems spans 20 years in military, commercial and private sectors, spanning aerospace, naval defence, offshore oil and gas, and unmanned surface vessel markets. Dr Wild is a member of Cambridge Pixel Technical Sales team.