Estimating Radar Cross Section using Bayesian Image Restoration

Richard O. Lane

QinetiQ Malvern[†] and University College London

Abstract: This paper describes a numerical Bayesian technique for computing the two-dimensional radar cross section (RCS) of a vehicle given a radar image of the vehicle. A Markov chain Monte Carlo technique known as the Metropolis-Hastings algorithm is used to generate a set of statistical samples that characterise the probability distribution of the RCS. Given a model for the sensor measurement process, the method may be applied to any type of radar image such as those produced by a synthetic aperture radar (SAR), inverse SAR (ISAR) or a real beam imaging radar. In principle the resolution of the recovered RCS may be higher than that of the initial image, which indicates a super-resolution ability. Properties of samples generated by the radar image of an idealised military vehicle are analysed and the proposed algorithm is shown to have potential to work under realistic scenarios.

1 Introduction

In the digital age there are a wide variety of sensor technologies available to image military vehicles. Radar has many advantages over other technologies due to its ability to operate day or night and in all weather conditions. However, there are several types of radar such as synthetic aperture radar (SAR), inverse SAR (ISAR) and real beam imaging radar and the logistics of comparing information from these various radar sensors is non-trivial.

Images gathered from each of these sensors will in general have differing resolutions and could be collected with different signal-to-noise ratios, making it difficult to compare images. Consider two ways of attempting a comparison between sensor images. The first is to convert both images to a representation of the vehicle by its two-dimensional plan view radar cross section (RCS). The advantage of this representation is that it is independent of the imaging platform. The second method takes the RCS extracted from one image and using knowledge of the imaging process of the other sensor converts this RCS into a radar image of the appropriate type. If an automatic target recognition (ATR) algorithm has been trained previously on a certain type of radar image the second method has the advantage that all target images could be converted to this image type and run through the algorithm.

The problem of RCS estimation is an inverse one. When the imaging process is well understood it is relatively simple to apply an imaging operator to the RCS of a vehicle to obtain the equivalent radar image. However, obtaining the RCS, given an image, is not trivial and there may be many possible RCS realisations that give rise to the same image. A Bayesian approach to the inversion process provides a principled way to recover the two-dimensional RCS. Bayesian methods are probabilistic, able to cope with the possibility of multiple solutions and can readily incorporate other available information as prior knowledge.

The rest of the paper is organised as follows. Section 2 outlines the model used to describe the imaging process. Section 3 describes the algorithm to perform Bayesian image restoration. Results are presented in section 4 and conclusions are drawn in section 5.

2 Imaging Model

Due to the nature of the radar scattering process a target will reflect energy by differing amounts depending on the geometry of the imaging platform and the target. Each type of target is assumed to have a two-dimensional plan view RCS, which is therefore a function of imaging geometry. The Bayesian image restoration technique could in theory also be applied to a three-dimensional representation using interferometric techniques, for example. However, this significantly increases the computational burden so this paper addresses the two-dimensional problem only.

High resolution radar imagery can be difficult to interpret since military vehicles tend to have complex surfaces with many scattering positions. In addition, the image produced by a radar has a resolution dependent on the bandwidth and beamwidth of the transceiver and antenna and there may be many scatterers contributing to a particular pixel in the image.

The radar image may be thought of as being produced by a point spread function applied to each scatterer and sampled at the image resolution with the addition of thermal noise (although other models are possible). If certain parameters of the radar such as power, frequency and temperature are known then the noise power can be calculated. Each type of radar has its own point spread function. For example, real beam radar images have a high range resolution but a low cross-range resolution so the point spread function is long in the cross-range direction and thin in the range direction.

Due to the random nature of noise the final image x is only probabilistically dependent on the RCS σ. On the assumption of independent Gaussian noise with a variance ψ^2, the likelihood of the image conditional on the RCS is

$$p(x|\sigma) = \prod_{j=1}^{n} \frac{1}{\sqrt{2\pi\psi^2}} \exp\left[-\frac{[x_j - f(psf,\sigma)_j]^2}{2\psi^2}\right], \tag{1}$$

where n is the number of pixels in the image and $f(psf, \sigma)$ is some arbitrary function that produces a noiseless image given the point spread function psf and the RCS σ. This function could be implemented using a computer model that accurately reflects the image formation process. Such computer models were not readily available at the time of writing so the results in this paper use a simple linear system $f(psf, \sigma) = psf * \sigma$, with $*$ denoting convolution. However, it is important to note that the numerical Bayesian technique described in the next section is fully able to cope with non-linear and non-Gaussian systems by replacing equation (1) with the appropriate form.

So far it has been assumed that the two-dimensional RCS and the radar image have the same resolution. This need not be the case and it is desirable, especially for low resolution images, to increase the resolution. This is performed by evaluating the function $f(psf, \sigma)$ at the required high resolution. The high resolution image is then re-sampled to the actual image resolution before the addition of noise. This super-resolution technique is possible within the Bayesian framework of the next section but it does introduce non-linearities and is not considered further in this paper. An alternative super-resolution technique may be found in [1].

3 Bayesian Image Restoration

The Bayesian image restoration approach to RCS estimation is a probabilistic way of modelling uncertainty in the RCS. It is possible that different combinations of RCS could give rise to the same image due to the combination of several scatterers in a single pixel and the addition of noise. This uncertainty is described by a probability distribution $p(\sigma|x)$ of the RCS conditional on the image under consideration. Bayes' theorem gives the distribution as

$$p(\sigma|x) = \frac{p(x|\sigma)p(\sigma)}{p(x)}. \tag{2}$$

In some simple situations it may be possible to calculate the distribution directly. However, in general this will not be the case as the form of the point spread function could be non-trivial and indeed non-linear. For the general case a numerical technique such as Markov chain Monte Carlo is required. One specific algorithm that has gained popularity in the statistical community is the Metropolis-Hastings algorithm [2].

The Metropolis-Hastings algorithm is a method for generating samples of a probability distribution. For the case considered here the samples represent the probability distribution $p(\sigma|x)$ of the RCS of the object of interest. One advantage of the algorithm is that it is only necessary to know the shape of the distribution hence there is no need to calculate the normalising factor $p(x)$. The likelihood of the image, given the RCS, was shown in equation (1). With choice of a suitable prior $p(\sigma)$ for the RCS the quantity of interest is then

$$\pi(\sigma|x) = p(x|\sigma)p(\sigma). \tag{3}$$

At the ith iteration a proposed new sample is generated from a proposal distribution $q(\sigma^{i+1}|\sigma^i)$. The proposal distribution may take a wide variety of forms each having its own advantages and disadvantages - see [2] for details. The proposed sample is accepted with a probability $\alpha(\sigma^i, \sigma^{i+1})$, where

$$\alpha(\sigma^i, \sigma^{i+1}) = \min\left[\frac{\pi(\sigma^{i+1}|x)q(\sigma^i|\sigma^{i+1})}{\pi(\sigma^i|x)q(\sigma^{i+1}|\sigma^i)}, 1\right]. \tag{4}$$

In other words, at each step a new sample is generated; if it is more likely[1] than the current one it is always accepted but less likely ones are also accepted with a certain probability. This avoids the problem of getting trapped in local maxima and is analogous to simulated annealing [3]. If the proposed sample is rejected then the current sample is used in the next iteration step. Initial samples generated during the so-called burn-in period depend on the starting position and must be discarded. The remaining samples are distributed from $p(\sigma|x)$ as required.

[1]Taking "likely" to mean the product of the target distribution and the proposal distribution.

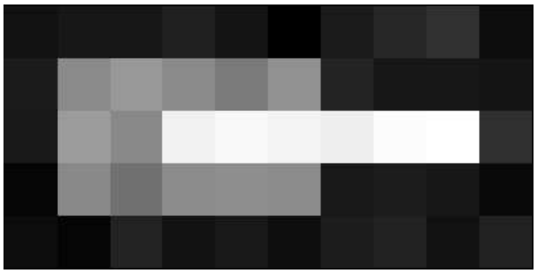

Figure 1: Two-dimensional RCS of a military vehicle

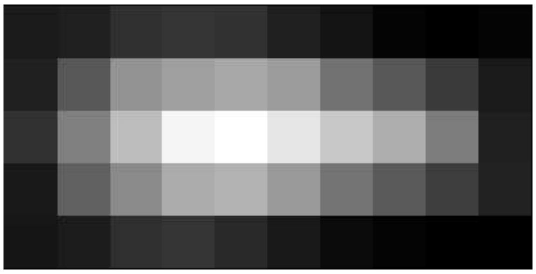

Figure 2: Image of the military vehicle

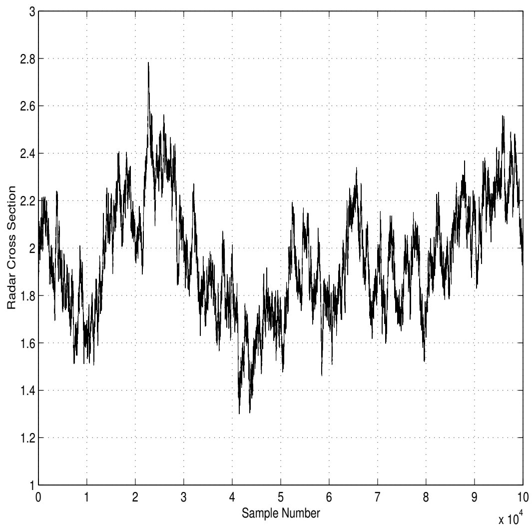

Figure 3: Series of RCS samples

4 Results

The Metropolis-Hastings algorithm described in the previous section was applied to synthetic data. The data consisted of an idealised military vehicle with a two-dimensional RCS shown in figure 1. A simple point spread function was used to filter the RCS and independent identically distributed Gaussian noise was added. The resulting image is shown in figure 2, where blurring due to the point spread function is particularly apparent. A uniform prior distribution for the RCS was used and the proposal distribution was a Gaussian random walk centred on the current sample with a width selected to achieve a 60% acceptance rate.

The series of samples of a single element of the RCS produced by the Metropolis-Hastings algorithm is shown in figure 3. The first 10,000 samples from the burn-in period were discarded and the remaining 100,000 are shown. The true RCS value in arbitrary units is 2 and it can be seen from the figure that most of the samples are close to this value. A notable feature of interest in the figure is that the samples are correlated. This could be important when estimating statistics of the samples because for the statistics to be reliable a reasonable number of independent samples are required. A different proposal distribution may produce samples with a lower correlation than the random walk [2], which could imply a necessity for fewer dependent samples than those produced here. The auto-correlation function of the graph in figure 3 was formed and the correlation decay time was estimated from the derivative of this function at the origin. The decay time was found to be 1720, which means there are approximately 58 independent samples. The same calculation for other elements produced similar results and this is considered a reasonable number of samples to estimate statistics of interest.

The mean RCS sample is shown in figure 4. The overall structure is very similar to the true RCS shown in figure 1 and it is easy to pick out the shape of the target. This mean RCS sample could be passed on to later stages of an automatic target recognition program but in keeping with the Bayesian philosophy it would be better to pass on all the samples and let any later processing stages use the full probability distribution characterised by the samples.

As an aside, it is of interest to see the best RCS sample produced by the Metropolis-Hastings algorithm in terms of the Euclidean distance between the sample and the true underlying RCS. This is shown in figure 5, where the shape of the target is clearer than the mean RCS shown in figure 4.

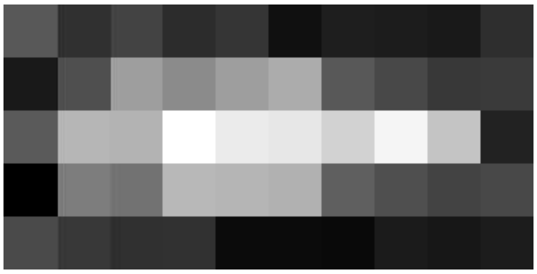

Figure 4: Mean RCS Sample

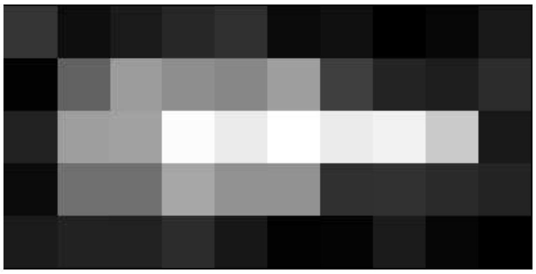

Figure 5: “Best” RCS Sample

5 Conclusion

An algorithm for estimating the two-dimensional plan view radar cross section (RCS) of a vehicle from a radar image has been implemented. The algorithm is based on a numerical Bayesian technique and produces a series of samples that characterise the probability distribution of the RCS. A simple problem involving synthetic data was considered where an idealised military vehicle was imaged using a linear system with additive Gaussian noise. Summary statistics of the RCS distribution were calculated and the mean image was found qualitatively to match well with the true underlying RCS. The complete ability of the algorithm has not fully been tested since there are no assumptions that preclude the use of non-linear and non-Gaussian systems. It is expected the algorithm will perform equally well under these conditions. The possibility of using this algorithm for super-resolution has also been raised but its potential performance has not yet been assessed. Future research should concentrate on these issues and also test the technique on real data.

Acknowledgements

This work was sponsored by the EPSRC Engineering Doctorate programme and the UK Ministry of Defence Corporate Research Programme “Communications, Information and Signal Processing”.

References

[1] S. P. Luttrell, “A Bayesian derivation of an iterative autofocus/superresolution algorithm”, *Inverse Problems*, 6:975-996, 1990.

[2] S. Chib and E. Greenberg, “Understanding the Metropolis-Hastings algorithm”, *The American Statistician*, 49(4):327-335, 1995.

[3] W. H. Press, S. A. Teukolsky, W. T. Vetterling, B. P. Flannery, *Numerical Recipes in C*, 2nd Edition, Cambridge University Press, 1992.

†Address: QinetiQ, Malvern Technology Centre, Worcs. WR14 3PS UK. Email: rolane@sar.qinetiq.com