Downloaded from rsta.royalsocietypublishing.org

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

# The Role of Prior Knowledge in Coherent Image Processing [and Discussion]

S. P. Luttrell, C. J. Oliver, P. L. C. Jeynes, S. Quegan and J. A. Allan

*Phil. Trans. R. Soc. Lond. A* 1988 **324**, 397-407
doi: 10.1098/rsta.1988.0028

**Email alerting service**

Receive free email alerts when new articles cite this article - sign up in the box at the top right-hand corner of the article or click **here**

To subscribe to *Phil. Trans. R. Soc. Lond. A* go to: **http://rsta.royalsocietypublishing.org/subscriptions**

This journal is © 1988 The Royal Society

Downloaded from rsta.royalsocietypublishing.org

*Phil Trans. R. Soc. Lond.* A **324**, 397–407 (1988)
*Printed in Great Britain*



# The role of prior knowledge in coherent image processing

By S. P. Luttrell and C. J. Oliver
*Royal Signals and Radar Establishment, St Andrews Road, Great Malvern, Worcester WR14 3PS, U.K.*

[Plates 1 and 2]

Models that encode prior knowledge about a scene provide a means for interpreting image data from that scene in more detail than would otherwise be so. Information about both background clutter and target characteristics should be included in this prior knowledge. We demonstrate the use of a generalized noise model to represent a variety of naturally occurring random terrain clutter textures observed in high-resolution synthetic aperture radar (SAR) images. In addition a similar approach is adopted for the simulation of such textures. Having established the background properties we next introduce prior knowledge about any target within the scene and exploit this in achieving a cross-section reconstruction having improved resolution compared with the original image. Examples of such a super-resolution method based on singular value decomposition are demonstrated and the limits of the technique are indicated.

## Introduction

In this paper we shall be concerned with aspects of the interpretation of synthetic aperture radar (SAR) images. In common with all image interpretation processes we shall imbue the raw-image data with some 'meaning' based on relevant past experience in a similar context. This 'prior knowledge' is cast in the form of a model that underlies our image interpretation.

In many radar applications we are concerned with detecting the presence of 'targets' against background 'clutter'. This statement assumes more information about the scene beyond that found in the image alone. We have introduced prior knowledge that (typically) targets are small bright blobs embedded in a weak clutter background. To detect the presence of such targets we must develop an understanding of the properties of both targets and clutter. In this paper we describe techniques to increase our understanding of both classes by the application of prior knowledge.

## Clutter models as prior knowledge

In figure 1, plate 1, we show a SAR image of a region of fields and woods, obtained with the RSRE X-band airborne system. It is apparent that the intensity in the image shows strong speckle behaviour. This is a well-known phenomenon in coherent imaging and arises where the total intensity is made up by the interference between large numbers of randomly positioned scatterers within each resolution cell. The presence of this speckle over the entire image suggests that a *physical model* for scattering and imaging of the surface might describe each resolution cell as containing many randomly positioned scatterers. No attempt is made to represent the configuration within the cell. The physical processes of scattering and imaging can then be analysed theoretically. Where the cross section is spatially non-varying the image statistics



The Royal Society is collaborating with JSTOR to digitize, preserve, and extend access to *Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences.* www.jstor.org

Downloaded from rsta.royalsocietypublishing.org



should correspond to classical Rayleigh speckle. In table 1 we compare the higher-order single-point moments of the detected intensity distribution for a region of open field, for which the cross section is expected to be essentially constant, with the theoretical values for pure speckle.

TABLE 1. COMPARISON OF THE INTENSITY MOMENTS FOR THE FIELD TEXTURE

| | | model | |
|---|---|---|---|
| moment | data | mean | SD |
| second | 1.970 | 2.000 | 0.014 |
| third | 5.68 | 6.00 | 0.14 |
| fourth | 21.3 | 24.0 | 1.3 |

It is apparent that the statistics are indeed consistent with such a model. In general the detected intensity of each pixel of an image will then take the form of a convolution of the underlying cross-section fluctuation from pixel to pixel with a negative exponential speckle contribution. This process is clearly visible qualitatively in figure 1. One would therefore expect that the intensity probability distribution will no longer be negative exponential but will have larger higher-order moments indicating the presence of excess fluctuations.

Having proposed a physical scattering model consistent with the observed speckle it is next necessary to provide a model for the cross-section fluctuations. In general, no physical description of the surface appears feasible so we are forced to consider a *phenomenological model* for the surface cross-section fluctuations. The aim of such a model is to represent the surface in such a way that the scattered radiation is consistent with the properties of the observed image; there is no attempt to reproduce the exact surface. We shall show that natural clutter textures can be treated as homogenous noise processes. In terms of figure 1, therefore, we are concerned that the model should be capable of reproducing the statistical and correlation properties of clutter textures such as that shown enclosed by the white rectangle.

One such model that has been proposed is that the surface cross section, $\sigma$, should be represented by a random process with a gamma-distributed probability-density function (PDF)

$$p(\sigma) = \frac{1}{\langle\sigma\rangle}\left(\frac{\sigma}{\langle\sigma\rangle}\right)^{\nu-1}\frac{1}{\Gamma(\nu)}\exp\left[-\frac{\sigma}{\langle\sigma\rangle}\right], \tag{1}$$

where $\Gamma(\nu)$ is the gamma function of order $\nu$ (Jakeman & Pusey 1976, 1977, 1978; Jakeman 1980*a*). This is the continuum analogue of the discrete negative-binomial distribution. On scattering the observed intensity PDF would then be the convolution of the gamma-distribution with the negative exponential distribution, namely the $K$-distribution given by

$$p(I) = \frac{1}{\langle I\rangle}\left(\frac{I}{\langle I\rangle}\right)^{\frac{1}{2}(\nu-1)}\frac{1}{\Gamma(\nu)}K_{\nu-1}\left[2\left(\frac{I}{\langle I\rangle}\right)^{\frac{1}{2}}\right], \tag{2}$$

where $K_{\nu-1}$ is the modified Bessel function ($K$ distribution) of order $\nu-1$. From (1) it is apparent that this noise model is a generalization of thermal noise (the case when $\nu = 1$). The order parameter $\nu$ serves to control the contrast and higher-order moments of the distribution. Such a noise model would result from a physical model in which the local cross section, or density of scatterers, depended on a birth–death-migration process (Jakeman 1980*b*). Thermal noise describes a process in which only migration (diffusion) occurs. This noise process is the result of including an additional source and sink of scatterers. A variety of evidence from scattering experiments at both microwave and optical wavelengths has indicated that such a model is



PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES
PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

Downloaded from rsta.royalsocietypublishing.org



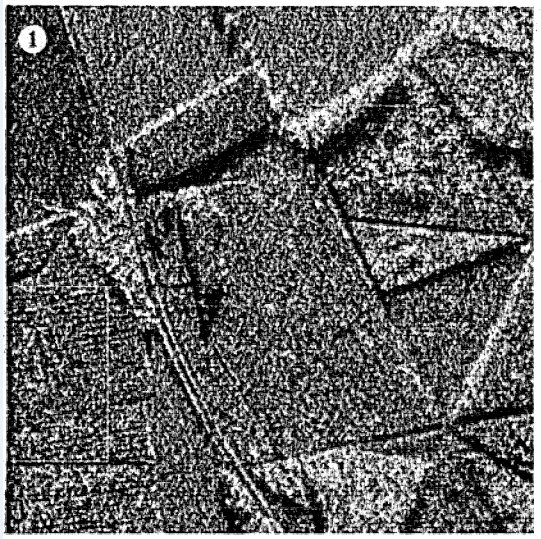


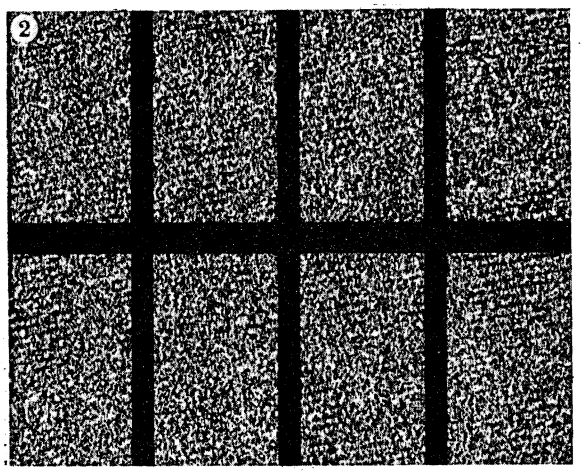


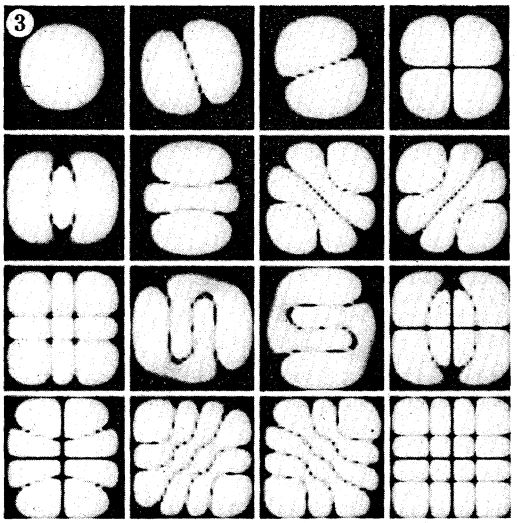


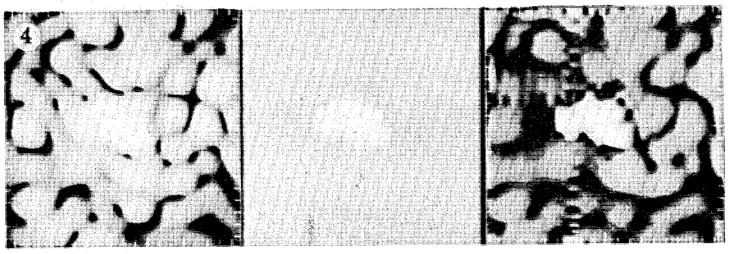


FIGURE 1. SAR image of a region of fields and woods.

FIGURE 2. Comparison of original (top right) and several simulated textures for the region indicated in figure 1.

FIGURE 3. First 16 of the weighted object eigenvectors when no prior knowledge is introduced; logarithmic display.

FIGURE 4. Stages in the first iteration of the reconstruction process for the simulated target. Original image (left), prior knowledge weight (centre), reconstruction (right); logarithmic display.

(*Facing p.* 398)

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

Downloaded from rsta.royalsocietypublishing.org



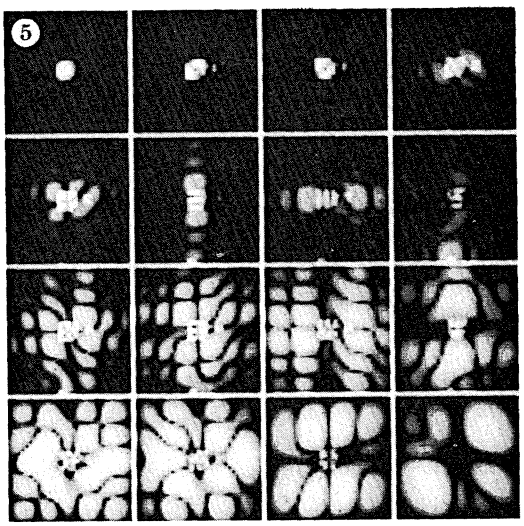


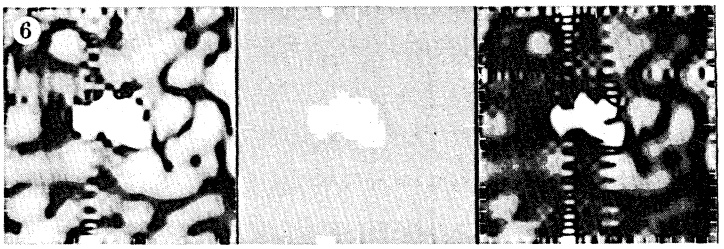


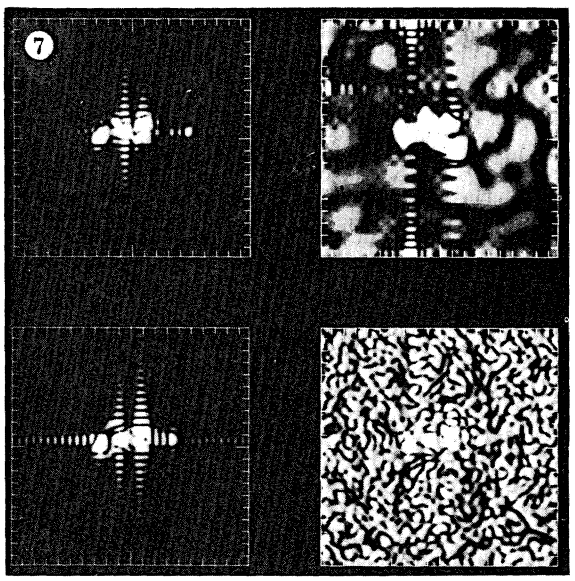


FIGURE 5. First 16 of the weighted object eigenvectors when localized target prior knowledge is introduced; logarithmic display.

FIGURE 6. Stages in the second iteration of the reconstruction process for the simulated target. First reconstruction (left), prior knowledge weight (centre), reconstruction (right); logarithmic display.

FIGURE 7. Comparison of the final reconstructions with high resolution images of the same object. Reconstruction without background (top left), reconstruction with background (top right), high-resolution image without background (bottom left), high-resolution image with background (bottom right); logarithmic display.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

Downloaded from rsta.royalsocietypublishing.org



consistent with the data (Jakeman & Pusey 1975; Jakeman *et al.* 1976; Parry *et al.* 1977; Ward 1981, 1982; Oliver 1984, 1986*a*, *b*, 1988) though no unambiguous physical justification is possible.

Before embarking on a detailed comparison of the statistics it is essential to consider the correlation properties within the texture which often convey more information than the contrast fluctuations. Gamma-distributed noise can be regarded as the resultant of a random walk in $\nu$ dimensions of a complex gaussian variable, i.e. thermal noise. Thus all the results for classical thermal noise can be generalized to the gamma-distributed equivalent by exploiting the appropriate factorization properties (Oliver 1984). For distributions with high contrast it is necessary to use fractional values of the order parameter $\nu$. Provided that the model truly represents the texture, it then contains fundamentally all the information available about that texture because gaussian processes are described entirely in terms of their mean and covariance. Various simple theoretical models of this kind have been proposed and compared with image data (Oliver 1984, 1985, 1986*b*, 1988).

Let us next apply this approach to the texture of the plantation of young trees shown in figure 1 enclosed by a white rectangle. The first stage is to derive a reasonable model for the surface autocorrelation function (ACF) which, when imaged, gives an intensity ACF that closely represents the observed image result. For the texture shown the model requires three components for an approximate correspondence. A method also exists for obtaining even closer agreement by regarding the ACF of the original data as a perturbation of the model ACF (Oliver 1987). Having obtained the form of the surface ACF the higher-order moments can then be calculated following the methods established earlier (Oliver 1984, 1985, 1986*b*). A comparison of the observed single-point moments with the predicted values for the first four moments is given in table 2. In addition the predicted values for a simple *K*-distribution having the observed second moment are included.

TABLE 2. COMPARISON OF THE INTENSITY MOMENTS FOR THE WOODED TEXTURE DENOTED IN FIGURE 1

| | | full model | | simple model | | simulation | |
|---|---|---|---|---|---|---|---|
| moment | data | mean | SD | mean | SD | mean | SD |
| second | 2.64 | 2.64 | 0.04 | 2.64 | 0.06 | 2.55 | 0.11 |
| third | 13.1 | 13.1 | 0.7 | 13.0 | 1.4 | 14 | 3 |
| fourth | 100 | 108 | 10 | 102 | 32 | 160 | 90 |

Both theoretical predictions are consistent with the original data within the theoretical errors included in the table. Strictly speaking an integrated *K*-distribution is not itself *K*-distributed so the simple model should not fit as well. However, for distributions with lower contrast than thermal noise the difference is often insignificant. The simple model, of course, is not capable of predicting how the statistics vary with the imaging resolution; for this the full theory is required.

We have so far demonstrated that random textures in coherent images can be represented in terms of a correlated gamma-distributed surface cross section, which when imaged yields *K*-distributed correlated intensity fluctuations. Although only a single example is included here many other clutter textures in both radar and sonar images have been demonstrated to be capable of representation with the same fundamental model.



PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

Downloaded from rsta.royalsocietypublishing.org



As a final test of the model it is possible to simulate textures having the same statistical properties as the original image by using a linear filter method (Oliver & Tough 1986; Oliver 1986*b*). This method seeks to reproduce the two-point statistics, the covariance function, of the original image. The surface is first simulated by passing uncorrelated gamma-distributed noise through a linear filter. The filter weights are chosen so that their ACF is determined by the surface ACF (Oliver & Tough 1986). The complex scattered field is then simulated by speckling the cross section with a random zero-mean gaussian variable for both the in-phase (I) and quadrature (Q) field components. This field is then convolved with the imaging response function of the system. A set of such simulations is compared with the original texture (top right) in figure 2, plate 1. Visually the agreement is excellent. As a further test the mean value and standard deviation of the higher-order single-point statistics of the simulations is included in the last two columns of table 2. Again the agreement is within the statistics. This ability to simulate textures is of great importance in deriving statistical estimates of, for example, target detection probabilities obtained with various algorithms against such a clutter background.

The highest level of prior knowledge, the *contextual model*, describes the relation of defined textures within a scene. In terms of figure 1, therefore, it would contain map information on the position of the boundaries of the woods and fields, statistical information about the correlation properties and order parameter of the texture in each region, plus other information, such as tree height, ground slope and aircraft altitude, to enable one to predict shadowing and other geometrical effects. In principle, the combination of all three stages in the model should enable us to predict coherent images that are statistically equivalent to a SAR image of the same region.

Once the model for the clutter is established it can be used in the detection of targets against this clutter background. The first stage in the process is to form a composite clutter model for the area created from combinations of map data, previous SAR images of the same area and other sensor data. This clutter model (prior knowledge) then defines the predicted statistical properties that are expected at each point in the SAR image. To detect 'targets' we then start with the assumption that the scene is entirely described in terms of the clutter map and search for anomalies which we define as targets. These anomalies may not be simply related to the image intensity but may depend on higher-order statistical properties of the image. We might, for instance, wish to detect changes of texture correlation length where the mean intensity and contrast remain the same. Whatever criterion is selected for the discrimination the clutter model is used to define a 'target detection threshold map' for that measure. Target detection then occurs when the value of the measure exceeds this threshold.

## Target resolution-enhancement from prior knowledge

Once a target has been detected the next stage is to seek to classify that target. In many examples of SAR image the resolution of the system is of the same order as the size of targets of interest making any type of geometrical classification based on the distribution of cross section within the target essentially impossible. The second part of this paper now concentrates on the use of prior knowledge about targets to improve the effective resolution of the sensor. This is achieved by introducing a model to represent the target. There are therefore two key issues that have to be considered: (i) what is an appropriate form for the prior knowledge model and (ii) how is the model applied to yield the maximum additional information about the target?



PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

Downloaded from rsta.royalsocietypublishing.org



Before addressing these issues directly let us consider a particular example of a prior-knowledge model and examine the implications. The Rayleigh resolution criterion states that two point sources are just resolved when the first null of one point spread function (PSF) lies on the main lobe of the second. However, in the 'monopulse' technique of radar processing strong prior knowledge about the target is introduced, namely that it consists of two scatterers. We can then inspect the detailed shape of the image to determine the amplitudes, phases and positions of the two scatterers with greatly enhanced resolution, limited only by the noise in the system, not the PSF directly. This type of approach is common in many experimental fields where one defines a model to represent the process under study and then determines the parameters of that model. The approach is effective provided that the prior knowledge introduced in the model is in fact correct. It is therefore advisable to use as weak a form of prior knowledge as possible because there is less chance of introducing false information. Let us now consider what general properties of targets might be appropriate as prior knowledge in the context of SAR images.

Obviously it would be possible to generalize the monopulse technique by assuming that one had a number of point sources within the target and searching for the best fit as that number is varied. As the number increases the ambiguity introduced by the noise means that a large number of alternative solutions exist and one is forced to introduce some arbitrary criterion for choosing a particular solution, for example, that with the smallest number of scatterers. This form of prior knowledge is generally speaking too strong because we cannot be sure that targets consist of point-like scatterers. A weaker form of prior knowledge that we can be confident about is that targets are of limited spatial extent. Furthermore we may assert that the target's position and dimensions are well represented by the original image. The fundamental premise of all radar target detection algorithms is that a target exists where a large return is recorded. This form of prior knowledge is based on the same assumption and may be viewed with the same level of confidence. It is obviously true that pathological scatterer configurations can violate the assumption and so invalidate the result. One could only have complete confidence in the prior knowledge if one knew precisely what the target was, in which case there is no more information to be gained from the image. This concentration of the target cross section into a localized continuum form makes no assertions about the scattering centres within the target itself; this information will be derived by applying the prior knowledge to the image data.

In applying resolution-enhancement methods on an SAR image a small region enclosing a detected target will be selected. Thus the prior knowledge must combine background and target models. For the total model we assert that it is made up of a background region of uniform cross section containing many random scatterers per resolution cell. Imposed on this background at a position and with dimensions determined from the original image is a localized target within which the scatterer distribution is unknown. Comparative strengths of target and background can be deduced from the image data.

Let us now examine how this prior knowledge may be exploited in the reconstruction of a target at an enhanced resolution. We shall follow the methods previously proposed by Luttrell (1985*a*) and Luttrell & Oliver (1986) based on a bayesian reconstruction scheme. Suppose that our object function, $f(=f(x)$, a continuous complex variable denoting the scattering amplitude and phase), is imaged through a response function $T$ to give a complex image vector $\boldsymbol{g}$. The imaging process can be denoted by the shorthand notation

$$\boldsymbol{g} = T f + \boldsymbol{n}, \tag{3}$$



PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

Downloaded from rsta.royalsocietypublishing.org



where the vector $\boldsymbol{n}$ is the noise introduced in the process. Bayesian reconstruction then provides a means for relating the *a posteriori* PDF, $P[f(x)\,|\,\boldsymbol{g}]$ (the conditional probability of a particular object given the image data), to the prior knowledge of the object, expressed in terms of the *a priori* PDF $P[f(x)]$ (the predicted probability of a particular object) and the imaging process via the expression

$$P[f(x)\,|\,\boldsymbol{g}] \propto P[\boldsymbol{g}\,|\,f(x)]\,P[f(x)], \tag{4}$$

where $P[\boldsymbol{g}\,|\,f(x)]$ describes the imaging properties in terms of the conditional probability of the image given a particular object. Note that a one-dimensional form is used for simplicity. This *a posteriori* PDF is the basic construct from which inferences about the data, in particular the target reconstruction, are to be drawn.

To progress we need to make assumptions about the PDFs of the object and the noise. Let us assume that the object distribution, the *a priori* PDF, is a zero-mean gaussian. This will not accurately represent the non-gaussian statistics that actually occur but will suffice to specify the scale of the amplitude of $f(x)$. In principle we cannot specify what the statistics of the object should be without knowing in advance what the target is. Thus the gaussian assumption is a reasonable compromise. The properties of the object are then all contained in its covariance matrix, $W$. In the imaging process in the absence of noise the conditional PDF $P[\boldsymbol{g}\,|\,f(x)]$ will be a delta function because the mapping is determined precisely by the imaging function $T$. When noise is introduced it broadens the conditional PDF. Let us assume that the measurement noise is also a zero-mean gaussian process specified completely in terms of its covariance matrix, $N$. Note that it is not necessary to assume that the noise is uncorrelated; correlated noise can be treated correctly. It can then be shown (Luttrell 1985*a*; Luttrell & Oliver 1986) that the *a posteriori* PDF can be expressed in the form

$$P[f(x)\,|\,\boldsymbol{g}] \propto \exp\left[-(f-\langle f\rangle)^{\dagger}\,(T^{\dagger}N^{-1}T+W^{-1})\,(f-\langle f\rangle)\right], \tag{5}$$

where

$$\langle f\rangle = WT^{\dagger}(TWT^{\dagger}+N)^{-1}\boldsymbol{g} \quad (\dagger \text{ denotes hermitian conjugate}). \tag{6}$$

This *a posteriori* PDF contains the available information about the object. In practice we require to select some particular member of this distribution as representing the 'best' reconstruction. A logical choice would be to choose the average value of $f(x)$, defined in (6).

Let us now make the simplifying assumptions that both the object and noise covariances are delta-correlated. The 'weight' $W$ is now a real function which can be related to a profile function $P(x)$ such that

$$W(x_1, x_2) = |P(x_1)|^2\delta(x_1-x_2). \tag{7}$$

It can be shown (Luttrell 1985*a*; Luttrell & Oliver 1986) that the object reconstruction may then be expressed as

$$f_{\mathrm{rec}}(x) = P(x)\sum_{k=1}^{m}\frac{g_k}{(\lambda_k+N)^{\frac{1}{2}}}u^k(x), \tag{8}$$

where $g_k$ are the components of the image $\boldsymbol{g}$ in the image basis vectors and $u^k(x)$ are the equivalent eigenvectors in object space. These object and image eigenvectors correspond to singular value decomposition (SVD) of the imaging process and are characterized by the eigenvalues $\lambda_k$. If the signal to noise ratio is large enough the results of this bayesian reconstruction method are identical to those obtained elsewhere using an SVD technique (Bertero & Pike 1982; Bertero *et al.* 1984*a*, *b*).



PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

Downloaded from rsta.royalsocietypublishing.org



Let us next examine the *a posteriori* distribution in more detail to determine where the dominant information is contained. It can be shown that any object can be expanded in terms of object space vectors in the form

$$f(x) = P(x) \sum_{k=1}^{m} \frac{f_k}{(\lambda_k + N)^{\frac{1}{2}}} u^k(x) + f_c(x), \tag{9}$$

where the expansion coefficients $f_k$ are the components of the image of the object $f$ in the image space. The second term in (9), $f_c$, denotes the sum of the components of $f$ that do not project onto the image space basis. Because these components have no contribution in the image (data) space they can be said to be 'invisible' to the data. Thus we cannot infer anything about these components from a measurement of the image $\boldsymbol{g}$. All that we can say about them is contained in the form of the *a priori* PDF that we have defined.

If we next consider the implication of the first term in (9) we see that the contribution to the reconstructed object of each of the vectors $f_k$ is characterized by an eigenvalue $\lambda_k$, which defines the visibility of that contribution. Where these eigenvalues are smaller than the noise $N$ the corresponding data have insufficient energy to dominate the noise and so cannot convey information about the object reconstruction. In this situation also the only information available is that introduced in the prior knowledge. Only those contributions $f_k$ whose eigenvalues exceed the noise provide additional information beyond that represented by the prior knowledge.

A more detailed discussion of the information content of these images, expressed in terms of mutual information, is to be found in Luttrell (1985*b*, *c*) and Luttrell & Oliver (1986).

If we next substitute for $f$ and $\langle f \rangle$ $(= f_{\text{rec}})$ from (8) and (9) into (5) we obtain for the *a posteriori* PDF

$$P[f(x)\,|\,\boldsymbol{g}] \propto \exp\left[f_c^{+}|P|^{-2}f_c\right] \exp\left[-\sum_{k=1}^{m} |f_k - g_k|^2 \left(\frac{1}{N} - \frac{1}{\lambda_k + N}\right)\right]. \tag{10}$$

The first factor corresponds to a piece of the *a priori* PDF that is not changed in passing to the *a posteriori* PDF, since $f_c$ is invisible in the data space. The second factor depends on the data and expresses the influence that the data, $\boldsymbol{g}$, have on the *a posteriori* PDF. Each mode carries an independent piece of useful information about the object, which vanishes as the noise increases.

Let us now illustrate the process by using (8) to calculate reconstructed objects for a simulated target embedded in clutter. We shall simulate a two-dimensional imaging system whose PSF is that for a square aperture. The data acquisition system measures complex field Nyquist-spaced samples (I and Q) on a square lattice. Because the target is localized we shall consider an image space of $9 \times 9$ samples. This must be sufficiently large that no sidelobes of the target are visible outside the sampled image. In the object space each image sample is subdivided into $4 \times 4$ subsamples without lying outside the bounds of the $9 \times 9$ array. This sampling is deemed adequate to provide at least Nyquist sampling on the enhanced-resolution object reconstructions. For the sake of this demonstration a typical small target is simulated which leads to the observed image shown on the left of figure 4, plate 1.

Having defined the sampling spaces we next require estimates of $W$ and $N$. In typical SAR images the measurement noise and the background clutter speckle are both essentially delta-correlated. Furthermore the noise is negligible compared with the clutter speckle fluctuations and may be ignored. If we also assume that the target is delta-correlated then the target and clutter prior knowledge are expressed in terms of the profile weighting function $W(x) = |P(x)|^2$

Downloaded from rsta.royalsocietypublishing.org



which has the form of a cross section. Having determined the weighting function we next compute the reconstructed object $f_{\text{rec}}(x)$ from the image $\boldsymbol{g}$ as given in (8). In the analysis it can be shown (Luttrell 1985*a*) that the reconstruction vectors in the object space include the profile function $P$. It is instructive to examine the consequence of introducing no prior knowledge into the reconstruction. This is equivalent to setting $P$ constant; we are totally non-committal about where the scatterers are located within the image. In figure 3, plate 1, we show the first 16 of the 81 possible object space eigenvectors arranged in order of decreasing eigenvalues. Note that as the eigenvalue decreases the spatial frequency content of the eigenvectors increases. In general higher spatial frequency terms are more attenuated by the imaging process leading to the ordering of the functions. However, for a uniform weighting function all eigenvalues are approximately equal so that all 81 are required to perform the reconstruction. As one might expect, a reconstruction based on this prior knowledge is essentially the same as the original image and can be shown (Luttrell 1985*a*) to provide an interpolated form of the data. We may use this interpolated form to make cross-section estimates for the combined clutter and target prior-knowledge weighting function.

Initially we use threshold detection to separate those pixels that may be defined as 'target' from those we define as 'clutter'. The mean intensity of those pixels that lie below the threshold is taken as the background cross-section estimator whereas the intensity of the pixels that exceed the threshold is taken as the estimate of the local target cross section. This process whereby the data is used to generate 'prior knowledge' may seem incestuous. It is, however, important to appreciate that the prior knowledge does not reside in the image data but in the rules that are used in separating the image into the two components and setting up the weighting function. In particular we are asserting, in common with nearly all other radar systems, that locally bright pixels are extremely likely to have been derived from local concentrations of cross section rather than from the myriad of other possible configurations that would give the identical image. This rule is not a fundamental constraint so much as a consequence of empirical observation of typical SAR images.

Let us now apply these methods to a particular simulated image, shown on the left of figure 4, plate 1. This image has been formed from an array of four dominant scatterers. However, the resolution is so poor that the image shows merely some evidence of asymmetry. The signal to background ratio is 20 dB. Performing threshold detection we derive the prior knowledge cross-section weighting function shown in the centre of the figure. Because the method uses threshold detection there is a statistical probability that pixels which might actually belong to the background class will emerge as targets, and vice versa. Two examples of the former effect can be observed. This weighting function now takes the form of the *a priori* PDF in the bayesian reconstruction scheme. The resulting reconstructed object is shown on the right of figure 4. Comparison with the original image reveals that the apparent resolution in the image has indeed been improved. The degree of enhancement depends strongly on the nature of the data and it is not possible to define a global enhancement factor. Some understanding of how the enhancement is achieved may be gained from studying the object space eigenvectors for the case where a localized target leads to a strongly peaked weighting function. The first 16 of these modes are shown in figure 5, plate 2. Again they are presented in order of decreasing eigenvalue. Note how the first few modes concentrate almost exclusively on a strongly localized region contained within the central peak of the data. No such effect was observed with the uniform weight (figure 3). Unlike the uniform example also, the eigenvalues of the modes decrease comparatively rapidly as the order increases. Thus a small region of the object space contains



PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

Downloaded from rsta.royalsocietypublishing.org



a large amount of the information present in the data. Only a comparatively small amount of the information, contained in the weak higher-order modes, is used to reconstruct the background clutter. This is the means whereby resolution enhancement is achieved in this approach.

Having applied the resolution-enhancement technique once it is tempting to attempt to apply it again. There are two main objections to such a step. Firstly, one has to vindicate the choice of the prior knowledge in an iterative procedure. Regardless of the true nature of the target this type of processing will tend to concentrate the object cross section into a number of localized scatterers. If the image were of an object such as a tree then this progressive concentration of the cross section would be inappropriate since the scatterers are distributed in a comparatively diffuse manner. On the other hand, man-made objects, such as vehicles, can usually be represented in terms of a small number of point scatterers. Thus, if it were possible to introduce further prior knowledge that the object in question is indeed man-made, we may proceed with further iterations of the reconstruction process treating the output of the first reconstruction as if it were itself an image.

However, proceeding with iterated reconstruction introduces a second limiting factor into consideration. The presence of the random clutter speckle within the image introduces a noise component into the reconstruction so that, eventually, the result is dominated by the noise. It is clearly essential to iterate the process only for as long as this noise component permits. In this instance a second iteration can be justified. Exactly the same procedure is adopted as before with the previous reconstruction treated as the new image (figure 6, plate 2, left).

The weighting function, shown in the centre of figure 6, plate 2, is derived in the same way as before. Note that the threshold detection algorithm still passes the two false alarm targets detected in the original data. The reconstruction shows still more resolution enhancement and even serves to split the original image into four dominant scatterers (figure 6, right).

Let us conclude by comparing the reconstruction with an image of the original object where the resolution is increased by a factor of four in each direction. In figure 7, plate 2, reconstructed images after two iterations with (top right) and without (top left) a clutter background are compared with equivalent high-resolution images (bottom).

The example on the left without background gives a very high-quality reconstruction, closely similar to the image, because the only limiting factor is the noise introduced by rounding errors in the computer. The example on the right, with an initial target to background ratio of 20 dB, shows the distortion that has been introduced by the random speckle. Nevertheless the information content of the reconstruction has been considerably enhanced by the introduction of prior knowledge when compared with the original image (top left, figure 4). Note also that the process has suppressed the background clutter compared to the peaks of the reconstructed target. Although this reconstruction may not in itself be sufficient to classify the target it does represent a significant step in that direction. This has been achieved entirely by post-processing the original data and does not require any modification of the existing SAR sensor.

## Conclusions

The task of interpreting SAR images is very difficult. In addition to the problems associated with the complexity of the image and the sheer quantity of data, which it shares with many other sensor systems, there are the further limitations of speckle and low resolution. It would be extremely useful to be able to summarize the content of a SAR image in terms of a few



PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

Downloaded from rsta.royalsocietypublishing.org



significant features, 'targets', set against a 'clutter' background. One can then concentrate one's efforts on the targets. In the first main section of the paper we presented a model of SAR terrain that could be used to define a target detection threshold and to suggest suitable detection algorithms. Ensuing target detections, which are essentially a result of the first stage of the reconstruction process, can then be studied in more detail including, in particular, resolution enhancement as outlined in the second main section of the paper. Both steps are vital tools in the total problem of SAR image interpretation.

## References


Bertero, M. & Pike, E. R. 1982 *Opt. Acta* **29**, 727–746.
Bertero, M., Brianzi, P., Parker, P. & Pike, E. R. 1984*a* *Opt. Acta* **31**, 181–201.
Bertero, M., De Mol, C., Pike, E. R. & Walker, J. G. 1984*b* *Opt. Acta* **31**, 923–946.
Jakeman, E. 1980*a* *J. Phys.* A **13**, 31–48.
Jakeman, E. 1980*b* *Opt. Acta* **27**, 735–741.
Jakeman, E., Pike, E. R. & Pusey, P. N. 1976 *Nature, Lond.* **263**, 215–217.
Jakeman, E. & Pusey, P. N. 1975 *J. Phys.* A **8**, 369–410.
Jakeman, E. & Pusey, P. N. 1976 *IEEE Trans. Antennas Propag.* **AP-24**, 806–814.
Jakeman, E. & Pusey, P. N. 1977 *Radar 77 (IEEE Conf. Publ.* 155), pp. 105–109.
Jakeman, E. & Pusey, P. N. 1978 *Phys. Rev. Lett.* **40**, 546–550.
Luttrell, S. P. 1985*a* *Opt. Acta* **32**, 701–714.
Luttrell, S. P. 1985*b* *Opt. Acta* **32**, 255–257.
Luttrell, S. P. 1985*c* *Inv. Problems* **1**, 199–218.
Luttrell, S. P. & Oliver, C. J. 1986 *J. Phys.* D. **19**, 333–356.
Oliver, C. J. 1984 *Opt. Acta* **31**, 701–722.
Oliver, C. J. 1985 *Opt. Acta* **32**, 1515–1547.
Oliver, C. J. 1986*a* In *Wave propagation and scattering* (ed. B. J. Uscinski), pp. 155–174. Oxford: Clarendon Press.
Oliver, C. J. 1986*b* *Inv. Problems* **2**, 481–518.
Oliver, C. J. 1988 In *Proc. IMA Conf. on Applications of Mathematics to Remote Sensing*. Oxford University Press. (In the press.)
Oliver, C. J. 1987 *Inv. Problems*. **3**, 463–475.
Oliver, C. J. & Tough, R. J. A. 1986 *Opt. Acta* **33**, 223–250.
Parry, G., Pusey, P. N., Jakeman, E. & McWhirter, J. G. 1977 *Opt. Commun.* **22**, 195–201.
Ward, K. D. 1981 *Electron. Lett.* **17**, 561–565.
Ward, K. D. 1982 *Radar 82* (*IEE Conf. Publ.* 216), pp. 203–207.


## *Discussion*

P. L. C. Jeynes (*Oxford Computer Services Ltd, 244 Barnes Road, Cowley, Oxford*). I would have thought that in general the concept of 'individual scatterers' is rather meaningless in wave scattering from rough surfaces.

C. J. Oliver. Firstly, the concept of a surface is not really applicable for describing radar scattering from the terrain. Secondly, we could not solve scattering from such a surface even if we could represent it. Accordingly, as described in the paper, we adopt a phenomenological model which reproduces, as far as possible, the observed image texture. This model enables us to make assertions about image properties as the operating conditions change and has a reasonable, though non-rigorous, physical interpretation.

S. Quegan (*Department of Applied and Computational Mathematics, The University, Sheffield*). As a proviso to Dr Oliver's last reply, texture described by the autocorrelation function is present only in some land cover types. It is not present in, for example, arable crops, and for these cover types the only information available is the mean backscatter.

Downloaded from rsta.royalsocietypublishing.org



J. A. Allan (*School of Oriental and African Studies, University of London, U.K.*). For those of us interested in 'backgrounds' rather than in objects disturbing backgrounds, could Dr Oliver tell us whether it is possible to distinguish different backgrounds?

C. J. Oliver. There are two independent types of information about the background clutter in SAR images: (i) the average cross section and (ii) the texture, described by the surface autocorrelation function. In my presentation I concentrated on the structural (texture) information rather than the cross section. However, both of these properties can be used to distinguish between clutter classes though some texture information will only be available in images from high-resolution systems. One problem is that the mean cross section of fields, for example, may be greater or less than that of the woods. Another is that the texture variation of woods overlaps with typical field textures. Thus it is not practical to classify the different regions unambiguously. However, we can separate images into regions of the same intensity within the statistics (i.e. segmentation) more reliably. It seems probable that the next important stage is to combine SAR images with map-like prior knowledge to address questions such as 'given that this region is a field what information can one deduce about the crop growing in it?' No convincing correlation between images and crop identification has as yet been generally demonstrated. Even more specific prior knowledge may be required such as defining the crop and asking how well it is growing, for example.

J. A. Allan. Do these procedures require massive computing power?

C. J. Oliver. The type of computing power needed for reasonable response times of about a second in the analysis that I have been describing would be provided by, for example, a transputer-based architecture using about 30 transputers delivering about 30 Mips (million instructions per second). Whether this is massive depends on perspective! The methods I outlined were ideally suited to this type of parallel-architecture implementation that is currently advancing dramatically, particularly in terms of unit cost.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

IGURE 1. SAR image of a region of fields and woods.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES
PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

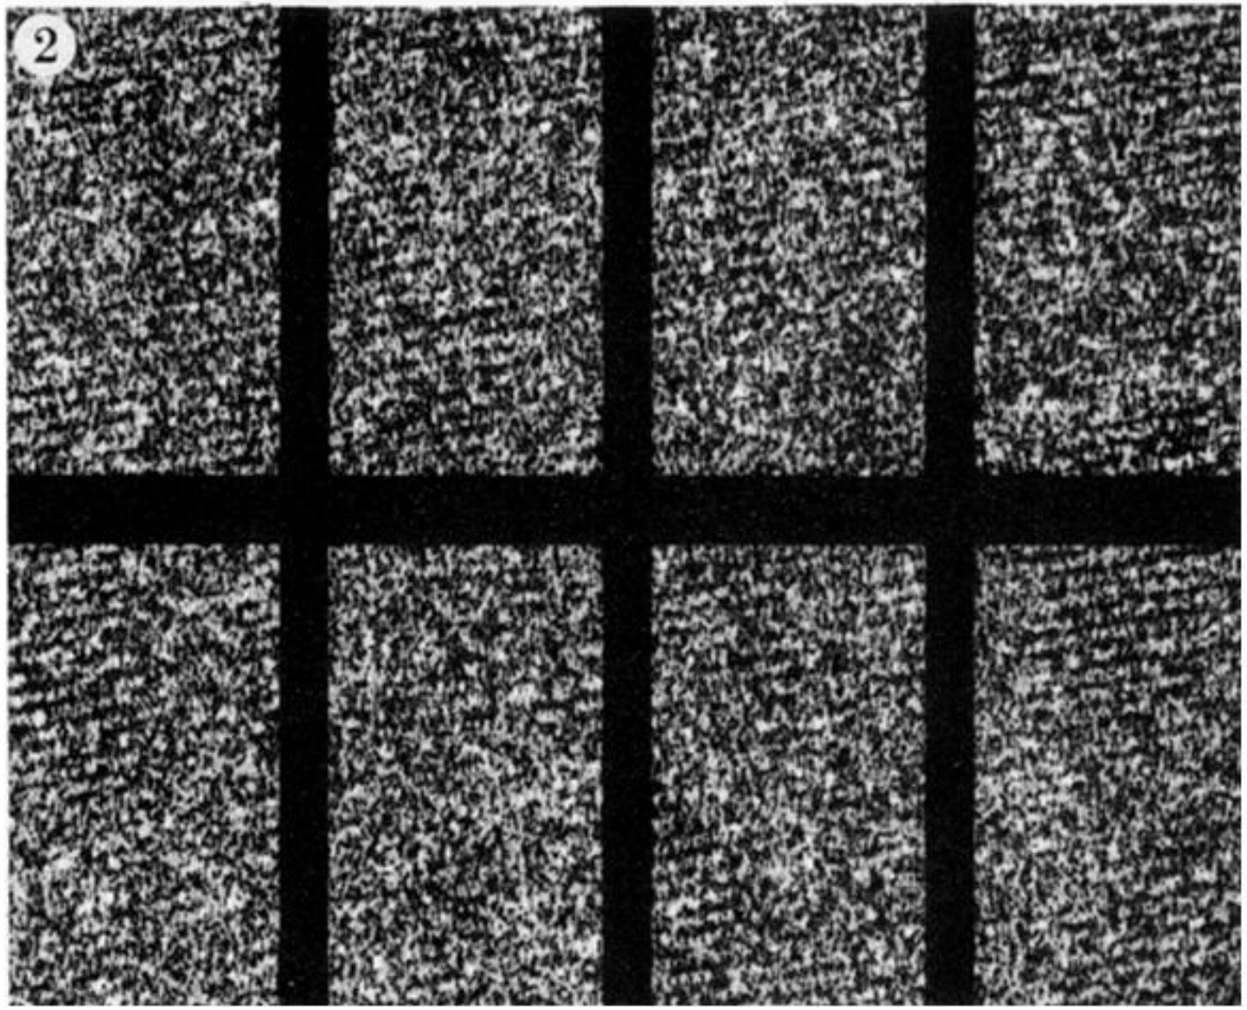


GURE 2. Comparison of original (top right) and several simulated textures for the region indicated in figure 1.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

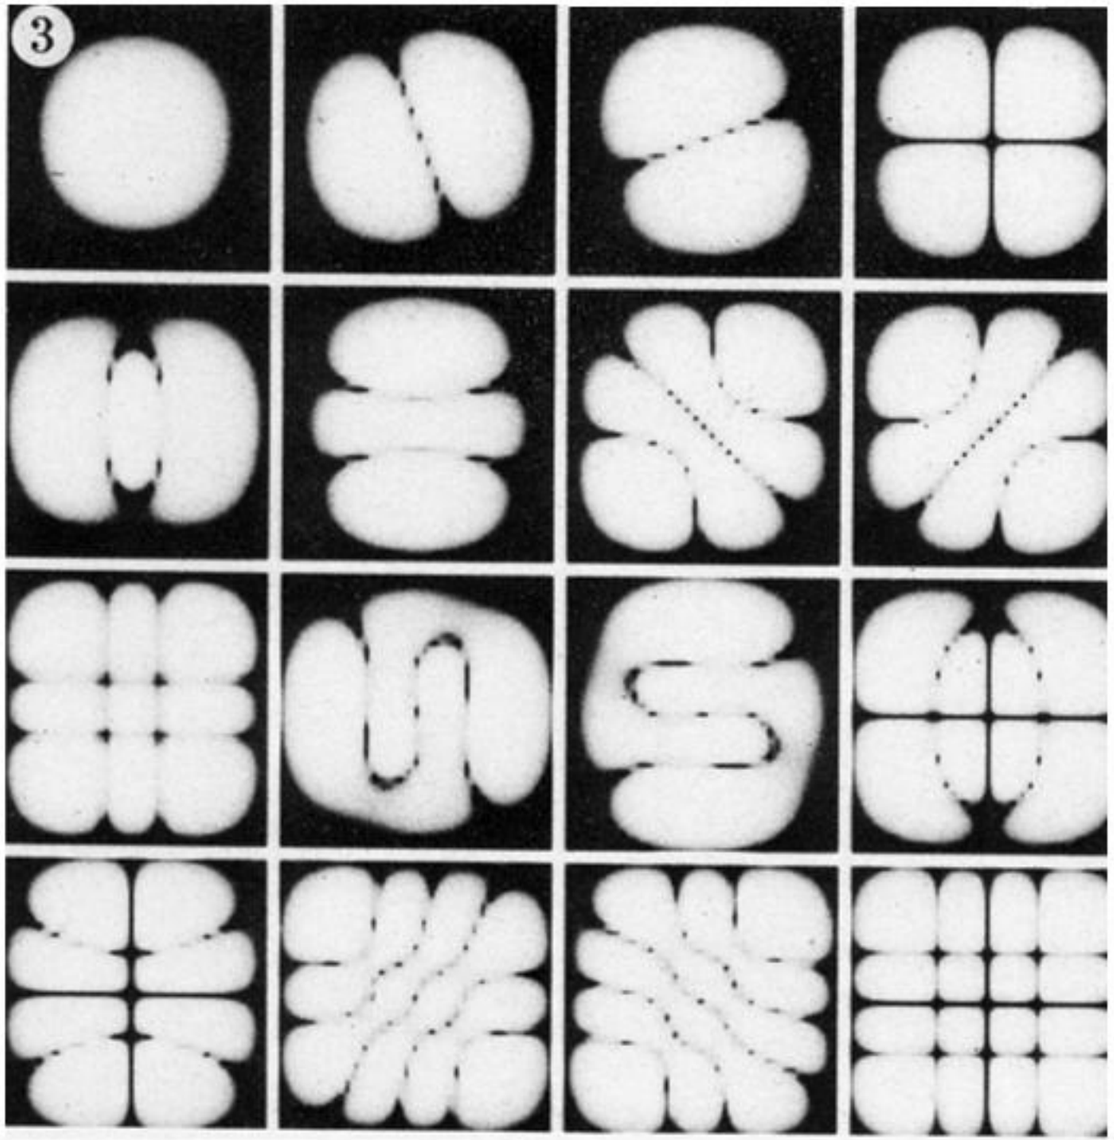


IGURE 3. First 16 of the weighted object eigenvectors when no prior knowledge is introduced; logarithmic display.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A
MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

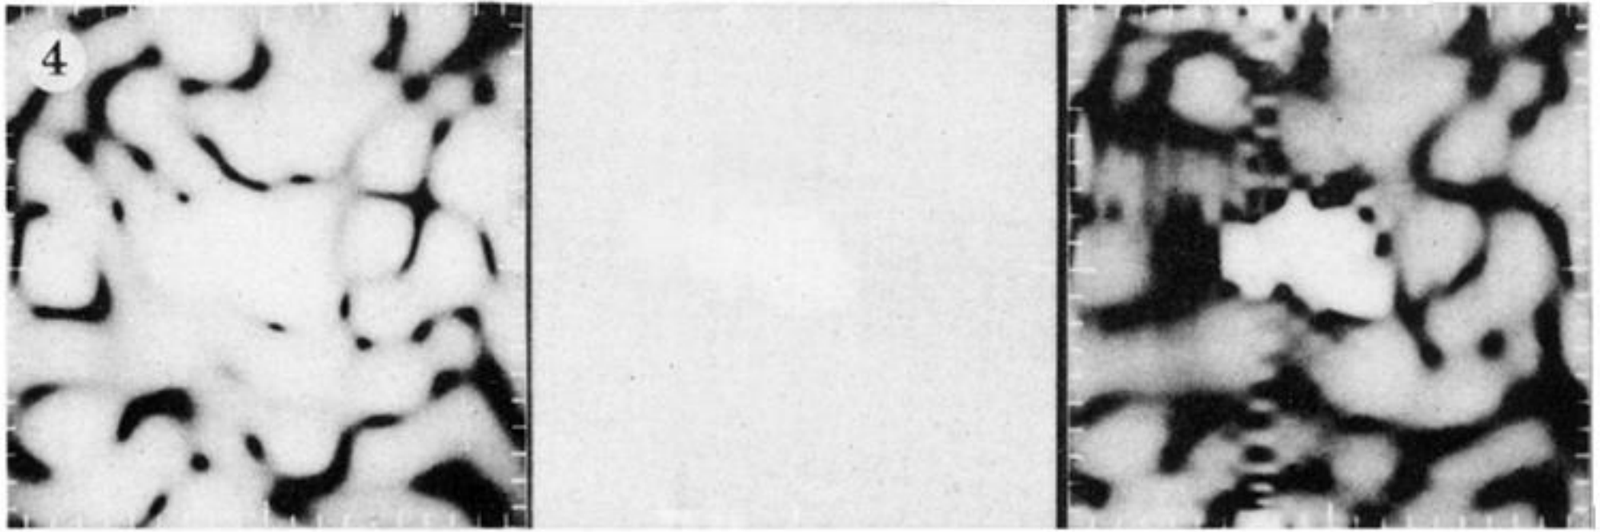


IGURE 4. Stages in the first iteration of the reconstruction process for the simulated target. Original image (left), prior knowledge weight (centre), reconstruction (right); logarithmic display.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

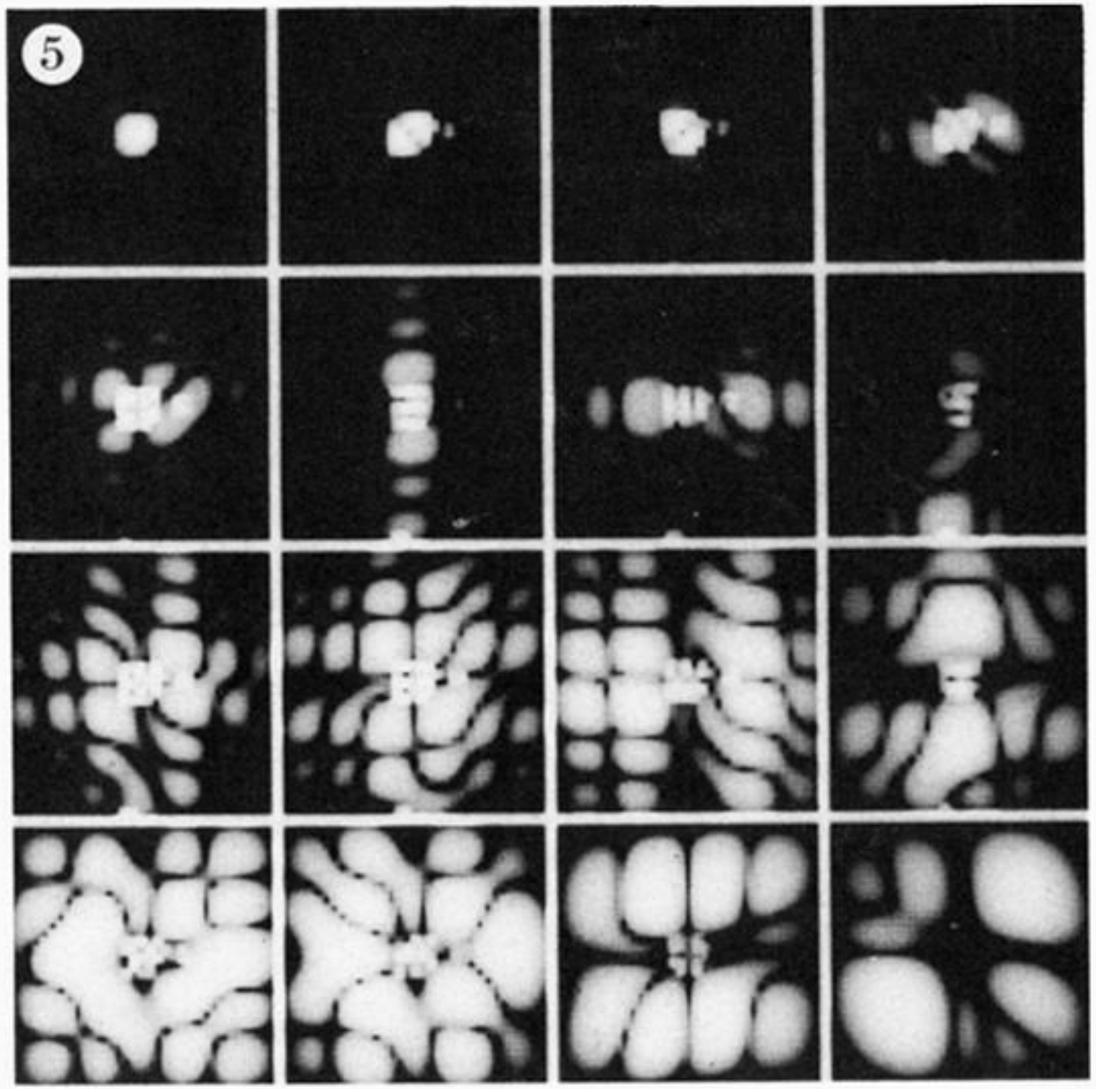


FIGURE 5. First 16 of the weighted object eigenvectors when localized target prior knowledge is introduced; logarithmic display.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

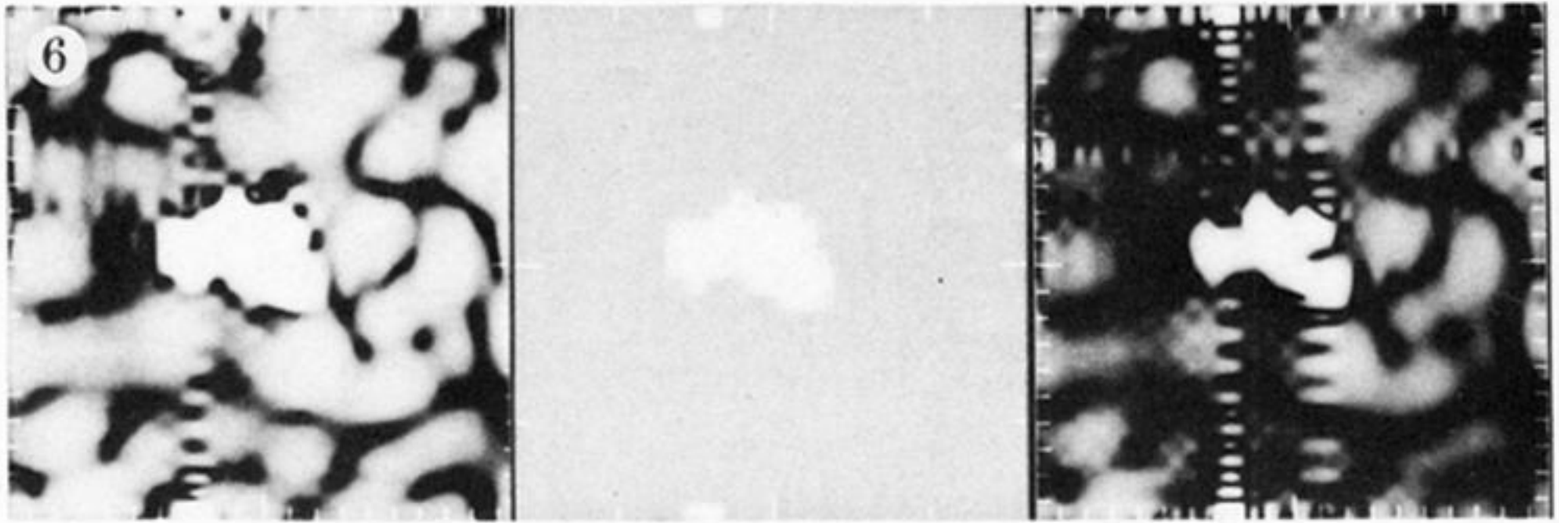


FIGURE 6. Stages in the second iteration of the reconstruction process for the simulated target. First reconstruction (left), prior knowledge weight (centre), reconstruction (right); logarithmic display.

Downloaded from rsta.royalsocietypublishing.org

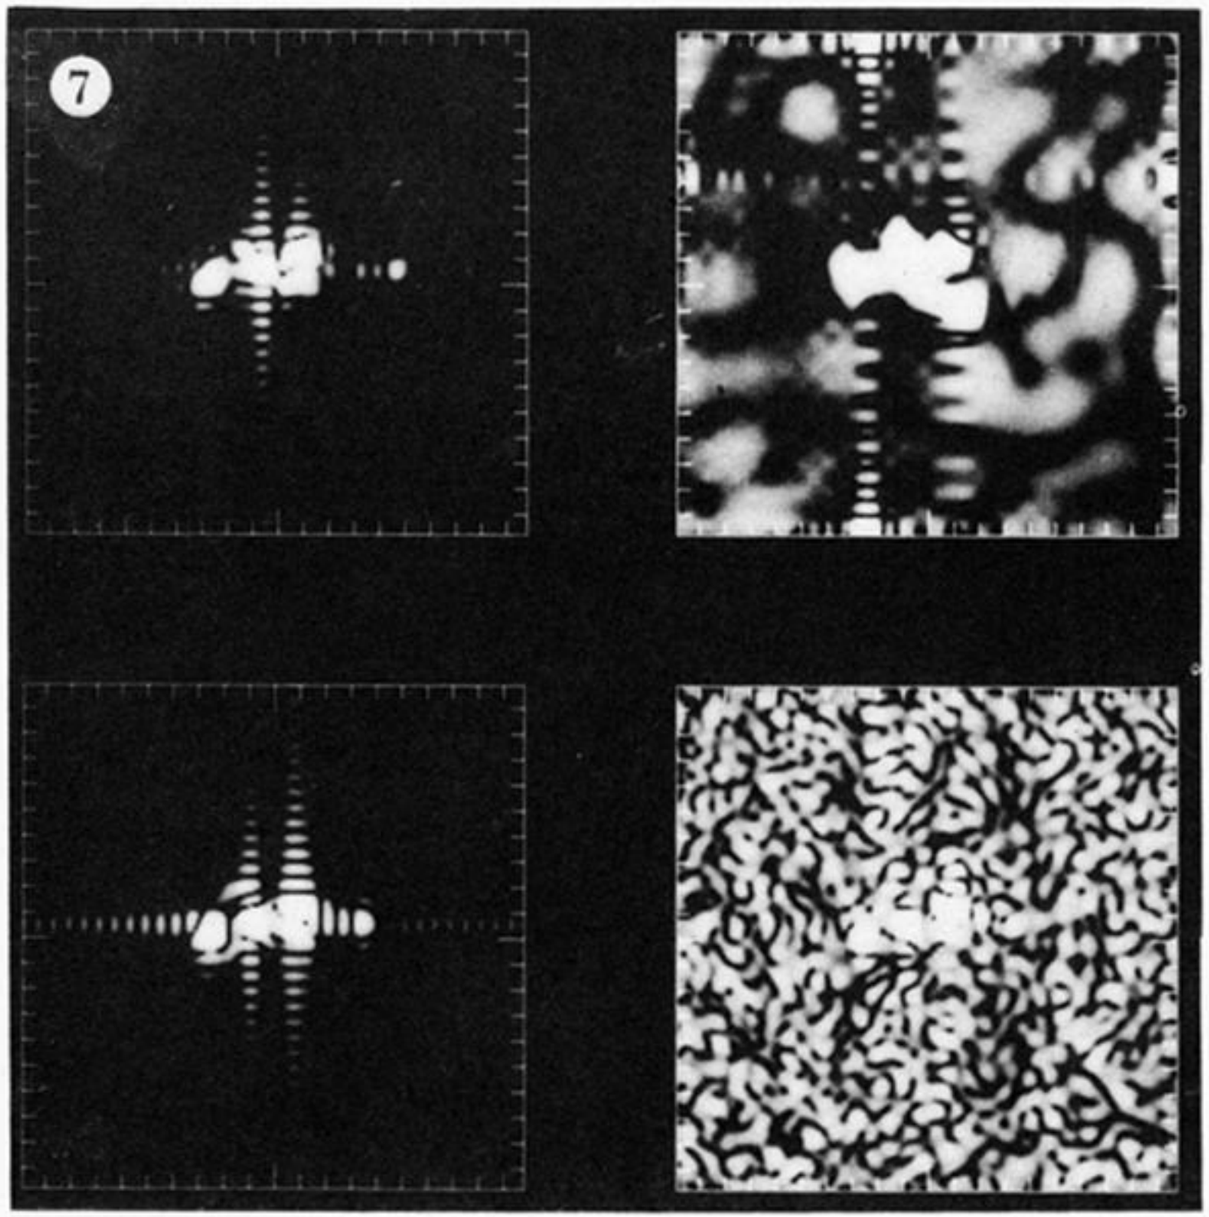


FIGURE 7. Comparison of the final reconstructions with high resolution images of the same object. Reconstruction without background (top left), reconstruction with background (top right), high-resolution image without background (bottom left), high-resolution image with background (bottom right); logarithmic display.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES