
Radar rainfall image repair techniques

Stephen M. Wesson and Geoffrey G.S. Pegram

Civil Engineering, University of KwaZulu-Natal, Durban, 4041, South Africa

Email for corresponding author: wessons@ukzn.ac.za and pegram@ukzn.ac.za

Abstract

There are various quality problems associated with radar rainfall data viewed in images that include ground clutter, beam blocking and anomalous propagation, to name a few. To obtain the best rainfall estimate possible, techniques for removing ground clutter (non-meteorological echoes that influence radar data quality) on 2-D radar rainfall image data sets are presented here. These techniques concentrate on repairing the images in both a computationally fast and accurate manner, and are nearest neighbour techniques of two sub-types: Individual Target and Border Tracing. The contaminated data is estimated through Kriging, considered the optimal technique for the spatial interpolation of Gaussian data, where the 'screening effect' that occurs with the Kriging weighting distribution around target points is exploited to ensure computational efficiency. Matrix rank reduction techniques in combination with Singular Value Decomposition (SVD) are also suggested for finding an efficient solution to the Kriging Equations which can cope with near singular systems. Rainfall estimation at ground level from radar rainfall volume scan data is of interest and importance in earth bound applications such as hydrology and agriculture. As an extension of the above, Ordinary Kriging is applied to three-dimensional radar rainfall data to estimate rainfall rate at ground level.

Keywords: ground clutter, data infilling, Ordinary Kriging, nearest neighbours, Singular Value Decomposition, border tracing, computation time, ground level rainfall estimation

Introduction

Raingauges have traditionally been used for the recording of rainfall over catchment areas and are often regarded as providing the 'true' rainfall at ground level. They provide a direct measurement of rainfall depth at a point and are relatively easy and cheap to maintain. Limitations do however exist with the raingauge data. Raingauges provide only a point accumulation measurement of the rainfall and fail to capture its spatial variability; even with a dense network of raingauges it is difficult to interpolate and extrapolate gauge data in any significant detail or accuracy, especially at short time scales.

Radar images of instantaneous rain-rate address some of the problems and limitations associated with raingauge data. Rainfall can vary greatly in both space and time. Radar rainfall images can show the instantaneous rainfall rate in far greater spatial detail and complexity than is possible with raingauge data. With this type of information it is also possible to observe and predict severe weather patterns in a far more timely and efficient manner. At present in South

Africa, the weather radar network of 11 radars provides instantaneous rainfall images of rain rate at approximately five-minute intervals and with one minute of arc (about 1.5 kilometres over South Africa) resolution.

There is a down side: there are various data quality issues associated with the radar rainfall data. Most importantly, radar provides an indirect measurement of the instantaneous rainfall rate. The reflectivity (dBZ) values returned to the radar are typically converted to a rainfall rate (mm hr^{-1}) with a variant of the Marshall Palmer formula (Marshall and Palmer, 1948). Errors in the radar rainfall data can include the following: ground clutter, beam blocking, bright band and anomalous propagation, to name a few. Ground clutter (non-meteorological echoes that influence radar data quality) is caused by the radar beam colliding with the earth's surface yielding abnormally high reflectivity values. This can result in large portions of the rainfall image being contaminated with highly reflective non-rainfall information; ground clutter is the main problem addressed herein.

The location of the ground clutter is identified by the South

African Weather Services Research Division, METSYS (Meteorological Systems and Technology), using an auto-correlation technique. In this technique the variability of reflectivity at ground clutter locations is considered to be much lower than that of precipitation. To detect this difference, the temporal auto-correlation is calculated at each pixel and once a determined threshold value of persistence is reached, the pixel is flagged as containing ground clutter. In this way the likely non-rainfall information can be marked and separated from the actual rainfall values (Visser, 2003). Figure 1 shows a typical radar image of widespread relatively low intensity rainfall at Bethlehem, South Africa, in summer. Where there was no rain recorded, the image background is white; the maximum observed reflectivity is given as yellow and does not exceed 30 dBZ, which is equivalent to approximately 3 mm hr⁻¹. The estimated ground clutter pixels are flagged, appearing black in this image.

To remove the ground clutter and estimate the missing rainfall data in real time, a computationally fast and accurate technique is required. Kriging was chosen as the method for estimating the missing data. However, for Kriging to be used in a real time application, advantage needs to be taken of various factors such as the pattern of the Kriging weights in the ‘good’ or control data pixels surrounding the ground clutter. Hence, an effective and efficient method for estimating the missing target rainfall values can be provided.

Standard techniques to remove ground clutter include subtracting a ‘clutter map’, derived from a period of no precipitation, from the observed reflectivity values. Unfortunately ground clutter is not totally stationary and its location can vary depending on atmospheric conditions (Vezzani, 1994).

The Kriging technique mentioned briefly above, and

described in more detail later, is proposed as an alternative. The methodology is presented in this paper along with relevant examples of its application. Preliminary results are encouraging, however the technique still needs to be fully validated on different types of rainfall.

Brief overview of Ordinary Kriging theory

Kriging was chosen as the computational method for infilling the contaminated rainfall data. Kriging is used to estimate the missing data as it is the optimal technique for the spatial interpolation of Gaussian data (Cressie, 1993: 106) if the sum of estimation errors squared is to be minimised. Reflectivity values of rainfall have been shown to be approximately normally distributed, because the rain rate on an image follows a truncated lognormal distribution (Bell, 1987; Pegram and Clothier, 2001). Thus, in this application, the estimation of the missing data by Kriging is carried out in the reflectivity rather than the rain rate domain. Kriging has several advantages over other interpolation techniques. The basis function used in Kriging is determined by the data set and its spatial structure; in this way the actual nature of the data set is taken into account. In other well known interpolation techniques such as multiquadrics, Fourier series or splines the basis function is merely chosen for computational convenience.

Ordinary Kriging is used here where the mean is assumed unknown and is implicitly estimated from the data. The Ordinary Kriging equations used herein are given in Eqn. (1).

$$z(s_0) = \lambda^T(s_0) \cdot z \quad (1)$$




Fig. 1. Typical instantaneous radar rainfall reflectivity image from the Bethlehem weather radar (20 February 2003). Black areas indicate contaminated data. The square surrounding the radar image has sides of length 270 kilometres.

where $z(s_0)$ is the estimated data value at the target location $s_0 = (x_0, y_0)$, $\lambda(s_0)$ is the vector of Kriging weights associated with the location s_0 and z is the vector of known or control data. This computation is carried out to estimate the reflectivity at each target pixel in turn. The weights are computed from Eqn. (2):

$$\begin{bmatrix} G & u \\ u^T & 0 \end{bmatrix} \cdot \begin{bmatrix} \lambda(s_0) \\ \mu(s_0) \end{bmatrix} = \begin{bmatrix} g(s_0) \\ 1 \end{bmatrix} \quad (2)$$

where u is a unit vector of ones and μ a Lagrange multiplier. In Eqn. (2) $g(s_0)$ is the vector of covariances or semi-variogram values between the target location (data point to be estimated) and the control points (known data values z) and G is the matrix of covariances or semi-variogram values between the control points (the covariance or semi-variogram function is estimated from the data). Equation (2) is derived by minimising the estimation error variance and the Lagrange multiplier is included to constrain the sum of the weights to unity. The variogram is defined as the expectation of the square of the differences of the field variables as shown in Eqn. (3):

$$2 \cdot \gamma(x, h) = E \left\{ [Z(x) - Z(x+h)]^2 \right\} \quad (3)$$

the semi-variogram is defined as $\gamma(x, h)$ (Journel and Huijbregts, 1978: 11).

In a stationary random field, Ordinary Kriging can be carried out by using either a covariance or semi-variogram function, where one is the complement of the other with reference to the field variance. Mathematically they are equivalent in this context, however in some applications the variogram is more robust (Cressie, 1993: 70–73), so was used throughout in preference to the covariance. The isotropic semi-variogram model chosen is the two-parameter exponential model specified in Eqn. (4) and is used in all Kriging computations herein.

$$g(s) = 1 - \exp[-(s/L)^\alpha] \quad (4)$$

where s is the Euclidian distance between data points, L is a scale parameter called the correlation length and α is the exponent parameter, which lies in the range $0 < \alpha \leq 2$; the model is thus Gaussian when $\alpha = 2$.

The advantages of the Kriging technique can be listed as follows: it is considered to be the best linear unbiased spatial estimator (Journel and Huijbregts, 1978: 57, 304); in extrapolation the values converge to the mean of the field whereas other techniques do not necessarily exhibit this desirable behaviour; the technique can be easily used on 1-D, 2-D and 3-D data sets; the precision of the estimated

data can be obtained. The precision is termed as the Kriging variance which, when using a semi-variogram function, is computed by Eqn. (5) (Cressie, 1993: 122–123):

$$\sigma_k^2(s_0) = \lambda^T(s_0) \cdot g(s_0) + \mu(s_0) \quad (5)$$

The main disadvantage of the Kriging technique is that it relies on finding the solution to a set of linear equations so that in large data sets the matrix algebra can become computationally burdensome and time consuming. Not as well recognised, and one of the main thrusts of this paper, is that the coefficient matrix can be highly ill-conditioned depending on the chosen parameters in the semi-variogram function and the data configuration. This can lead to inaccurate solutions and numerical instability. These matters will be dealt with in more detail in the sequel.

Methods for improving computational efficiency

In weather radar applications in South Africa where data sets are processed at a frequency of 5 minutes, the missing rainfall data need to be estimated in real time, thus a computationally fast and accurate technique is needed. The classical Kriging technique requires the solution of a linear system of equations whose size is determined by the number of control points. This can be time consuming, as the coefficient matrix becomes large. A radar image typically contains 120 000 pixels, leading to a coefficient matrix with dimensions greater than 100 000 by 100 000. Currently it is impractical to work with systems this size online.

There is an unexpected saving which can be exploited to ease the computational burden in a dense data set, as in radar data which typically occur on a lattice, not all of the control points need to be used in estimating the missing data points. A ‘screening effect’ occurs where the significant Kriging weights associated with the controls are concentrated around the target points (Chiles and Delfiner, 1999: 202–206). This can be used advantageously to reduce computation time significantly. For each individual target point, or set of target points in a cluster, a Kriging neighbourhood in its vicinity can be selected to reduce the dimensionality of the problem.

THEORETICAL JUSTIFICATION OF THE SCREENING EFFECT IN ONE DIMENSION

To give some theoretical justification to the ‘screening effect’, a one-dimensional analysis is presented which shows this very nicely. This does not easily extend itself to two dimensions or more, but intuitively and computationally the idea is justified by extension from a single dimension. It

turns out that, in one dimension, the screening effect is limited by the order of the model in cases where the co-variogram is that of an autoregressive process of order p . In fact only the first p intact (control) data nearest the point to be infilled (target) will have non-zero weights λ in the Kriging equation $z(s_0) = \lambda^T(s_0) \cdot \underline{z}$. A small one-dimensional example will suffice to explain the ideas.

One-dimensional AR(1) and AR(2) models

Suppose there is a sample $\{z_i, i = 1, 2, \dots, 6\}$ with two missing values at positions $i = 3$ and 4 which are to be estimated. Assume the sample is a realisation of a zero-mean unit-variance Gaussian AR(1) process with parameter ϕ : $z_i = \phi z_{i-1} + a_i$ and (because of the reversibility of the AR(1) process) $z_i = \phi z_{i+1} + e_i$ where $\{a_i\}$ and $\{e_i\}$ are independent white noise processes with the same variance. To attempt to infill the missing data using a forecast from z_2 , would give $\{\hat{z}_3, \hat{z}_4\} = \{\phi z_2, \phi^2 z_2\}$ and hindcasting from z_5 , $\{\hat{z}_3, \hat{z}_4\} = \{\phi^2 z_5, \phi z_5\}$, with no prospect of agreement. Any compromise would be subjective and *ad hoc*, unless, for example, an iterative scheme embodying the EM algorithm were to be employed.

As an optimal alternative, use Kriging to estimate the missing values. The correlation matrix between the known values is then formed from the usual 6×6 matrix and removing the two central rows and columns:

$$G_{11} = \begin{bmatrix} 1 & \phi & \phi^4 & \phi^5 \\ \phi & 1 & \phi^3 & \phi^4 \\ \phi^4 & \phi^3 & 1 & \phi \\ \phi^5 & \phi^4 & \phi & 1 \end{bmatrix}$$

whose inverse is:

$$G_{11}^{-1} = \begin{bmatrix} (1-\phi^6) & -\phi(1-\phi^6) & 0 & 0 \\ -\phi(1-\phi^6) & (1-\phi^8) & -\phi^3(1-\phi^2) & 0 \\ 0 & -\phi^3(1-\phi^2) & (1-\phi^8) & -\phi(1-\phi^6) \\ 0 & 0 & -\phi(1-\phi^6) & (1-\phi^6) \end{bmatrix} \div [(1-\phi^2)(1-\phi^6)]$$

$$= \begin{bmatrix} c & a & & \\ a & \beta & \alpha & \\ & \alpha & \beta & a \\ & & a & c \end{bmatrix}$$

To get the matrix of Kriging weights Λ^T one needs the correlation matrix between the two missing and four known values:

$$G_{21} = \begin{bmatrix} \phi^2 & \phi & \phi^2 & \phi^3 \\ \phi^3 & \phi^2 & \phi & \phi^2 \end{bmatrix}$$

so that because of the symmetry in this case

$$\Lambda^T = \begin{bmatrix} 0 & \phi(1-\phi^4) & \phi^2(1-\phi^2) & 0 \\ 0 & \phi^2(1-\phi^2) & \phi(1-\phi^4) & 0 \end{bmatrix} \div [1-\phi^6]$$

The first thing to note is that only the observations contiguous to the gap have non-zero weights, so are the only observed data involved in the estimation. The second thing to note is that the inverse of G_{11} is a tridiagonal matrix, whose form can be exploited to find explicit expressions for its elements in the general 1-D case.

In general, for a sequence of n complete observations from an AR(1) process, the correlation matrix is $R_n = \{\phi^{|i-j|}, i, j = 1, 2, \dots, n\}$. Its inverse is a tridiagonal matrix whose upper and lower diagonals are filled with equal elements $a = -\phi/(1-\phi^2)$, the main diagonal has equal elements $b = (1+\phi^2)/(1-\phi^2)$, except for the first and last which equal $c = 1/(1-\phi^2)$.

It is not difficult to show that when a block of data of width w is missing after the m th observed value, then the correlation matrix of the surviving variables becomes

$$R^*(n, m, w) = \begin{bmatrix} R_m & Q \\ Q^T & R_{n-m-w} \end{bmatrix}$$

where

$$Q = \begin{bmatrix} \phi^{m+w} & \phi^{m+w+1} & - & \phi^{n-1} \\ \phi^{m+w-1} & \phi^{m+w} & - & \phi^{n-2} \\ - & - & - & - \\ \phi^{w+1} & \phi^{w+2} & - & \phi^{n-m} \end{bmatrix}$$

and where R_m has the same form as R_n . The inverse of $R^*(n, m, w)$ will be a tridiagonal matrix of the same form as R_n^{-1} with identical diagonal elements except for the four elements either side of the gap, where instead of α and β , the elements will be a and b , which after some tedious but straightforward algebra are found to be equal to:

$$\alpha = (1-\phi^{2w+4})/[(1-\phi^2)(1-\phi^{2w+2})]$$

$$\beta = -\phi^{w+1}/(1-\phi^{2w+2})$$

which are seen to collapse to a and b when $w = 0$ (no gap).

When G_{11}^{-1} is pre-multiplied by G_{21}^T to give Λ^T , only the weights of the elements either side of the gaps will be non-zero. This happens in general and not only in the AR(1) example above. The behaviour extends to AR(p) models, where now the size of the screen is p elements deep either side of the gaps.

Other combinations of gaps in the data still produce a tridiagonal inverse whose elements either side of the gaps depend on the width of the gaps as well as the width of the intact elements between the gaps. A small-dimensional numerical example will demonstrate the ideas. In a sequence

of 11 AR(1) values with $\phi = 0.5$, the targets are at $i = 3, 4, 6$ and 7. The G_{11} matrix, its inverse, G_{21} and the resulting matrix of Kriging weights Λ^T (where the solid lines indicate the gaps) is:

$G_{11} =$

element	1	2	5	7	8	10	11
1	1	0.5	0.0625	0.015625	0.007813	0.001953	0.000977
2	0.5	1	0.125	0.03125	0.015625	0.003906	0.001953
5	0.0625	0.125	1	0.25	0.125	0.03125	0.015625
7	0.015625	0.03125	0.25	1	0.5	0.125	0.0625
8	0.007813	0.015625	0.125	0.5	1	0.25	0.125
10	0.001953	0.003906	0.03125	0.125	0.25	1	0.5
11	0.000977	0.001953	0.015625	0.0625	0.125	0.5	1

$G_{21} =$

element	1	2	5	7	8	10	11
3	0.25	0.5	0.25	0.0625	0.0313	0.0078	0.0039
4	0.125	0.25	0.5	0.125	0.0625	0.0156	0.0078
6	0.0312	0.0625	0.5	0.5	0.25	0.0625	0.0312
9	0.0039	0.0078	0.0625	0.25	0.5	0.5	0.25

$\Lambda^T =$

element	1	2	5	7	8	10	11
3		0.4761	0.1904				
4		0.1904	0.4761				
6			0.4	0.4			
9					0.4	0.4	

The only non-zero weights are those corresponding to the observations contiguous to the gaps, as asserted.

The ideas carry over to an AR(p) model as expected. This is useful because an AR(2) model is quite flexible for the purposes of describing the correlation models encountered in some random fields, particularly rainfields measured by radar. For example, given 14 values (with gaps at $i = 2, 8$ and 9) sampled from an AR(2) model: $z_i = \phi_1 z_{i-1} + \phi_2 z_{i-2} + a_i$ with $\phi_1 = 1.10$ and $\phi_2 = -0.32$, the inverse of G_{11} is tridiagonal as asserted and the Kriging weight matrix Λ^T is:

It will be noted that the weights in the second spaces from the gap are negative for the AR(2) case with these parameters. It is also noted that this behaviour (negative weights) is found when Kriging with Gaussian shaped correlograms. Nevertheless, even in the Gaussian case, the non-zero weights are confined to a zone close to the gaps.

An interim conclusion is that in the particular case of equally spaced data in one dimension, *only the observations contiguous to the gaps influence the infilling of the missing data*. This is likely to have greater economies in higher dimensions, as will now be demonstrated computationally.

element	1	3	4	5	6	7	10	11	12-14
2	0.4977	0.6570	-0.1448						
8					-0.2285	0.8932	0.4225	-0.1435	
9					-0.1435	0.4225	0.8932	-0.2285	

DEMONSTRATING THE TWO DIMENSIONAL SCREENING EFFECT BY COMPUTATION

An example of the ‘screening effect’ is illustrated in Fig. 2, which depicts a small dimensional example on a set of 8×9 pixels with four (marked by an asterisks) missing in the interior. The semi-variogram model used had an $\alpha = 1.5$ and a correlation length of $L = 11$ pixels. (Note in passing that the value of $L = 11$ pixels, of dimension 1.5 kilometres, is an average value determined from observed radar rainfield intensities for a wide range of weather types and this value will be used in the sequel; it is not the purpose of this paper to deal with estimation of co-variograms). The weights associated with the 68 control points for each of the four target points were computed one set at a time. These weights were then summed for all four targets and their totals appear in the shaded pixels. It is seen that the significant Kriging weights are concentrated in the immediate vicinity of the target points and rapidly diminish in magnitude with distance. At a range of greater than two cells from a target point the Kriging weights are nearly zero. As can also be seen in Fig. 2, negative weighting values do occur some distance from the target. The Kriging weights could be forced to be non-negative but a disadvantage in doing this is that the Kriged estimates are then constrained to lie within the minimum and maximum values of the selected control data (Chiles and Delfiner, 1999: 224), (Cressie, 1993: 143). In the rainfield infilling application it is more appropriate to allow for the Kriged estimates to range outside of the minimum and maximum values of the control data used.

By considering only the control points with significant Kriging weights a considerable reduction in size of the coefficient matrix can be made in cases where the target points are in clusters contained within an otherwise complete set of controls. With this reduction in size of the control set a concomitant decrease in computation time can be achieved with little or no loss in the accuracy of the final results.

0	0	0	0	0	0	0	0	0
0	0	-0.03	-0.07	-0.01	0.02	0.01	0	0
0	-0.04	-0.02	0.38	-0.17	-0.18	-0.03	0	0
0	-0.14	0.54	*	1.33	0.40	0.00	-0.03	0
0	-0.15	0.62	*	*	*	0.39	-0.08	0
0	-0.04	-0.04	0.61	0.75	0.49	-0.01	-0.03	0
0	0	-0.04	-0.16	-0.20	-0.14	-0.04	0	0
0	0	0.01	0.02	0.02	0.02	0.01	0	0
0	0	0	0	0	0	0	0	0

Fig. 2. Sum of all Kriging weights concentrated around four unknown data points. The variogram used had an $\alpha = 1.5$ with $L = 11$ pixels.

ILL-CONDITIONING OF THE COEFFICIENT MATRIX

In the Ordinary Kriging computational procedure, a set of linear equations must be solved to determine the weights $\lambda(s_0)$. This solution is given by Eqn. (6), derived from Eqn. (2):

$$\begin{bmatrix} \lambda^T(s_0) & \mu(s_0) \end{bmatrix} = \begin{bmatrix} g^T(s_0) & 1 \end{bmatrix} \cdot \begin{bmatrix} G & u \\ u^T & 0 \end{bmatrix}^{-1} \quad (6)$$

Solving for $\lambda(s_0)$ requires inversion of the coefficient matrix (usually the equations are solved using an efficient method based on LU decomposition). It was found that the coefficient matrix can be highly ill-conditioned depending on (i) its size, (ii) the chosen parameterisation in the semi-variogram function (iii) the layout of the data set on a lattice and (iv) the ratio (s/L) as used in the semi-variogram model, Eqn. (4), where a decrease in the ratio results in the coefficient matrix becoming more ill-conditioned. The ill-conditioning is most sensitive to the α -parameter, the exponent in the semi-variogram function defined in Eqn. (4). As α increases from a value close to zero towards a value of two (converging towards a Gaussian semi-variogram function) the coefficient matrix becomes increasingly ill-conditioned. When $\alpha = 2$ in the semi-variogram function, and the coefficient matrix is above a certain size, (as small as 68×68 , from the 8×9 region of points shown in Fig. 2) the coefficient matrix is essentially singular. As indicated in Fig. 3 the determinant of the coefficient matrix (plotted on a log scale) decreases steadily as $\alpha \rightarrow 2$. The determinant was computed for a relatively small 50 by 50 coefficient matrix, with the semi-variogram function having a correlation length of $L = 11$ pixels.

Although the Gaussian semi-variogram results in a near singular coefficient matrix, it is still commonly used in practical Kriging applications (e.g. Todini, 2001), albeit for sparsely distributed control points. It is difficult to determine at what size the coefficient matrix does become noticeably unstable since this appears to be largely determined by factors such as the layout pattern of the control data points. However for a coefficient matrix in excess of 40 by 40 in size, and with $\alpha = 2$, the matrix can be ill-conditioned enough to return nonsensical weights. In such cases, conventional methods of computing the inverse coefficient matrix such as Gauss-Jordan and its derivatives cannot be used. The method of Singular Value Decomposition (SVD) can be employed with advantage, however, to determine an accurate and meaningful solution for the weights, despite the ill-conditioning problem described above.




Fig. 3. Semi-variogram exponent α and corresponding determinant of coefficient matrix (size 50×50) for $L = 11$ pixels.

THEORETICAL JUSTIFICATION OF MATRIX RANK REDUCTION TECHNIQUE

Singular Value Decomposition involves the decomposition of a rectangular matrix into unique matrix ‘factors’; an overview of the method is readily available in Press *et al.* (1992: 59–67) which is briefly summarised here. The matrix is decomposed into a column orthogonal matrix U , a diagonal matrix W that contains the non-negative singular values along its diagonal and the transpose of an orthogonal matrix V . The decomposition of the Ordinary Kriging coefficient matrix is given in Eqn. (7).

$$\begin{bmatrix} G & u \\ u^T & 0 \end{bmatrix} = \begin{pmatrix} U \end{pmatrix} \cdot \begin{pmatrix} w_1 & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & w_n \end{pmatrix} \cdot \begin{pmatrix} V^T \end{pmatrix} \tag{7}$$

To compute the inverse, one simply needs to invert the diagonal elements of W and multiply out the matrices in the order indicated by Eqn. (8):

$$\begin{bmatrix} G & u \\ u^T & 0 \end{bmatrix}^{-1} = V \cdot \left\{ \text{diagonal} \left(\frac{1}{w_j} \right) \right\} \cdot U^T \tag{8}$$

When the coefficient matrix is ill-conditioned, the diagonal values along the W matrix need to be carefully considered because the singular values which are near zero define the (near) null space of the matrix. To obtain a meaningful pseudo-inverse, the diagonal elements of W^{-1} are set to $1/w_j$ where w_j is above some chosen near-zero threshold and to zero otherwise. This reduces the rank of the matrix by the size of the (near) null space.

To demonstrate the ideas, a coefficient matrix of size 150 by 150 was computed from a sample reflectivity field. The

coefficient matrix was then decomposed as shown in Eqn. (7). This was done several times for α values ranging from one to two. The singular values obtained in the W -matrix were ranked then plotted in Fig. 4, where it is seen that the w_j values rapidly diminish in magnitude along the diagonal and become close to zero as α goes from one to two (Exponential to Gaussian semi-variogram); the threshold of 10^{-15} reflects machine precision.

Since many of the singular values are significantly close to zero relative to w_1 which is approximately 60, the most appropriate action is to eliminate these values and replace them (and their inverses) with zeros instead. This solves the problem of ill-conditioning in the inverted coefficient matrix, yields useful accurate solutions and, incidentally, reduces computation time.

The Euclidian norm of the coefficient matrix is equal to the sum of the squares of the singular values; if the coefficient matrix were a covariance matrix this sum would also be equal to the variance, as indicated by Eqn. (9):

$$\sum_{j=1}^n w_j^2 = \sigma_x^2 \tag{9}$$

The variance is predominantly contained in the first few terms of the W matrix with the others contributing very little. Only the largest of the singular values need to be considered in most cases since their sum closely approximates the variance.

Once the appropriate number of w_j values has been trimmed and replaced with zero values, only the corresponding rows and columns need to be retained in the U and V matrices; as shown in Eqn. (10) the columns of the U matrix and the rows of the V^T matrix can be trimmed. This speeds up the computation dramatically with little or no change in the accuracy of the final results.

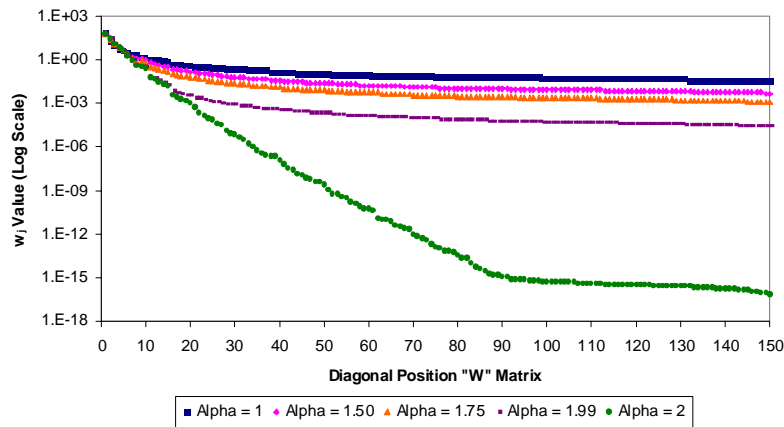


Fig. 4. Singular values in W matrix for different values of semi-variogram exponent α ; note the logarithmic scale. The coefficient matrix was 150 square and L was kept constant at 11 pixels.

$$\begin{bmatrix} G & u \\ u^T & 0 \end{bmatrix} = \begin{pmatrix} U \\ \vdots \\ \vdots \\ \vdots \end{pmatrix} \cdot \begin{pmatrix} w_1 & 0 & \dots & 0 \\ 0 & \ddots & & 0 \\ \vdots & & \ddots & \vdots \\ 0 & 0 & \dots & w_n \end{pmatrix} \cdot \begin{pmatrix} V^T \\ \vdots \\ \vdots \\ \vdots \end{pmatrix} \quad (10)$$

DEMONSTRATION OF ADVANTAGES OF MATRIX RANK REDUCTION BY COMPUTATION

As an example, a sample reflectivity field was selected and portions removed to simulate a ground clutter scenario. To infill the missing 57 target points, a 140 by 140 coefficient matrix had to be inverted. The semi-variogram parameters used were $\alpha = 1.5$ and $L = 11$ pixels. Initially, with none of the w_j values removed, the Kriging process (calculation of weights and estimation for each of the 57 target values) took approximately 5.6 seconds. The W matrix was then trimmed incrementally along with the corresponding rows and columns of the U and V^T matrix. It was found that up to 70% of the w_j values could be removed with no significant change in the final Kriged results (accords to the sum of the difference squared of hidden and computed targets) with a consequent decrease in time to 0.6 seconds as shown in Figs. 5 and 6. This corresponds to a nine-fold decrease in computation time with no concomitant change in accuracy achieved with the matrix rank reduction technique. The results would be more dramatic for a larger α .

An example of the numerical instabilities that can occur in the Kriging process is given in Fig. 7. A spatially correlated Gaussian random field (the upper part of Fig. 7) was generated with a known semi-variogram function on a 25 by 25 grid field. The semi-variogram parameters used were $\alpha = 1.5$, sill (ω) = 1, $L = 1000$ pixels (the one exception from $L = 11$, to emphasize the instability), nugget (p) = 0 and a variance (σ^2) = 1, as shown in Eqn. (11).

$$g(s_0) = \sigma^2 - \left\{ p + \omega [1 - \exp(-(s/L)^\alpha)] \right\} \quad (11)$$

Forty randomly scattered points were then selected from the correlated field as control points. Simple Kriging was then used to estimate the remaining points of the field from




Fig. 5. Sum of errors squared and corresponding percentage w_j removal, $L = 11$ pixels, $\alpha = 1.5$, 140 controls (140 \times 140 coefficient matrix).




Fig. 6. Computation time and corresponding percentage removal of singular values in a 140 \times 140 coefficient matrix with $L = 11$ pixels and $\alpha = 1.5$.




Fig. 7. Simple Kriging without removal of near zero singular values. For definition of the variogram model, see text.

the control points. As shown in comparing the two parts of Fig. 7, the Kriged field does not resemble the original correlated field at all due to the highly ill-conditioned coefficient matrix.

Figure 8 shows exactly the same scenario as in Fig. 7, however in this instance the trimming of the near zero singular values from the W matrix was carried out, as illustrated in Eqn. (10). The Kriged field is now a smoothed replica of the correlated field as desired. In Figs. 7 and 8 the 40 out of 250 points selected to perform the Kriging were

randomly chosen. The α -value was midway between exponential and Gaussian, yet still the numerical instabilities occurred. This result is surprising and was thought to be important enough to publish as a warning to the unwary and to provide an antidote.

For typical radar reflectivity fields of diameter 250 pixels (about 50 000 data points), the above two techniques, nearest neighbours and matrix rank reduction need to be used in conjunction in order to achieve efficiency and stability. Image processing techniques also need to be taken advantage




Fig. 8. Simple Kriging with removal of near zero singular values.

of so as to be able to manage the large data sets. Two techniques for ground clutter removal exploiting nearest neighbours will be compared in the next two sections; one works with individual target points in turn, the other by border tracing around a cluster of target points and solving the Kriging equations simultaneously.

Individual target's nearest neighbours (ITNN) technique

In the nearest neighbours technique applied to individual target points in turn (ITNN), the weighting distribution associated with Ordinary Kriging is taken into account in the sense that the 'screening effect' described in the previous section (where the significant weights at control points are concentrated in the immediate vicinity of the selected target data point) is exploited. In the algorithm suggested here, only the neighbours nearest to the target point are identified. This results in a considerable decrease in computation time and with no significant loss in accuracy. The SVD matrix rank reduction process is also carried out to reduce the computation time in finding the Kriging weights.

The ITNN algorithm can be described as follows for a two-dimensional data set. The radar reflectivity image is searched, pixel by pixel, starting at the top left-hand corner and working across each row then down the image, one row at a time, the ground clutter having been already flagged as -10 dBZ values. Once a ground clutter point is located, a search outwards from that point is initiated to locate the 20 nearest valid control points – the maximum would be 24 in a 5×5 square region if the target was an isolated pixel. Thus control values within two or so pixels of the target are selected. The search strategy proceeds as follows. Initially the rows directly above and below the point are searched, and then the columns on either side of that point. The search then moves incrementally outwards searching in turn rows and then columns until 20 uncontaminated neighbours are located. This strategy is illustrated in Fig. 9. The pixel marked with an asterisk is the target point initiating the first search on the image; other targets (yet to be infilled) are masked in grey.

Once the nearest neighbour controls are located they are tested to determine whether they are all zero: if they are, the target is set to zero. If the controls are all non-zero, Ordinary Kriging is carried out to estimate a value for the target.

Once the target point has been identified, an estimate of its value is calculated and then stored in a separate vector. The search resumes where it left off and continues searching across each row and then down the image, one row at a time, until the next ground clutter pixel is located. In this




Fig. 9. Nearest neighbours search strategy. Ground Clutter pixels marked as grey squares.

way the rainfall at each ground clutter pixel is estimated in a point by point manner unaffected by previous estimates until all the ground clutter pixels have been infilled. After the entire data set has been scanned the estimated reflectivity values, stored sequentially in a vector, are then inserted in place of the flagged ground clutter pixels.

As examples, two radar-rainfall reflectivity images in South Africa have been selected for repair, Fig. 10 from the Durban weather radar and Fig. 11 from the Bethlehem weather radar. On each of them the ground clutter pixels have been flagged with -10 dBZ values, which appear as black segments in Figs. 10 and 11. The number of ground clutter targets is approximately 1400 (for both the Bethlehem and Durban data) and the time taken to estimate the missing data for one 400 by 400 pixel image is typically 0.5 to 2 seconds computed on a Pentium(R) 4 with a 2,40GHz CPU and 512MB of RAM. The images on the right of each pair are those which have been 'repaired', i.e. where the ground clutter values, or contaminated pixels, have been replaced with reflectivity values estimated by Kriging.

Border tracing nearest neighbours (BTNN) technique

An alternative for identifying segments which have to be infilled is the border tracing with nearest neighbours (BTNN) technique, where advantage is again taken of the distribution of the Kriging weights to reduce the computational load and improve computational efficiency by concentrating on the nearest neighbours around a cluster. In this method image processing algorithms were investigated to provide an alternative approach to infilling the missing data. In the ITNN technique the clutter was infilled point by point whereas in the BTNN technique the infilling is done a whole segment at a time, where a segment is one or more target pixels that are grouped together in a connected region.

Once again, the image is searched from the top left hand corner, row by row, moving downwards until a flagged

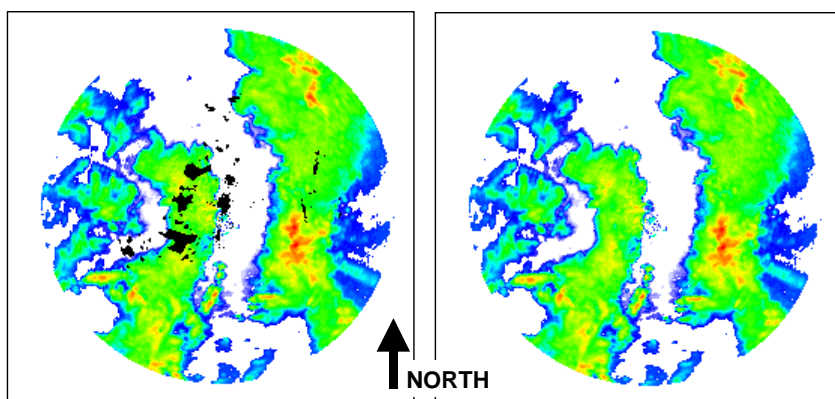


Fig. 10. Durban radar rainfall reflectivity data before and after ground clutter removal (18 November 2000). Black areas indicate contaminated data. The square surrounding each radar image has sides of length approximately 270 by 270 kilometres.




Fig. 11. Bethlehem radar rainfall reflectivity data before and after ground clutter removal (25 February 2003). Black areas indicate contaminated data. The square surrounding each radar image has sides of length 270 kilometres.

ground clutter pixel is located. An eight-connectivity border tracing technique (Sonka *et al.*, 1999: 142–145) is then implemented from that local origin. From that pixel the eight immediate surrounding neighbours are searched in an anticlockwise direction until another new ground clutter pixel is located. This becomes the new local origin and the process is repeated. The border is then traced in this manner until the original starting point is reached. Figure 12 (left) illustrates how the border is traced in an anticlockwise direction around the ground clutter segment, with the dashed lines indicating pixels tested and identified as controls in the process.

While the border is being traced, a Kriging neighbourhood two pixels deep from each border target is selected simultaneously. As indicated in Fig. 12, the control points within a range of 2.5 pixels are selected — this strategy neglects the corner pixels on a 5×5 square grid surrounding the target. All of the pixels selected in this manner are then

identified as the Kriging control neighbourhood for the estimation of the cluster of target points in the particular segment. Once the border tracing has been completed, and the Kriging neighbourhood has been selected, an algorithm




Fig. 12. Border tracing technique indicating trace direction and pixels tested in process (left image). Selection of Kriging neighbourhood for infilling of missing data (right image) (Sonka *et al.*, 1999: 143).

identifies the positions of the remainder of the target points that are located in the interior of the traced contour.

The Kriging control neighbourhood is then tested to determine whether the controls are all zero. If this is the case, all of the target points in the segment are set to zero. If the data points are not all zero, ordinary Kriging is used to estimate values for all of the target points that make up that segment.

The solution to the Kriging equations is then computed by an inversion by partitioning technique (Aitken, 1967: 148) in conjunction with the SVD matrix rank reduction technique in an effort to decrease the computation time. The inversion by partitioning technique as employed here splits the coefficient matrix into four equal sub matrices, given by Eqn. (12).

$$A = \begin{bmatrix} P & Q \\ Q^T & S \end{bmatrix}$$

$$A^{-1} = \begin{bmatrix} \tilde{P} & \tilde{Q} \\ \tilde{Q}^T & \tilde{S} \end{bmatrix}$$

where:

$$\begin{aligned} \tilde{P} &= (P - Q \cdot S^{-1} \cdot Q^T)^{-1} \\ \tilde{Q} &= -(P - Q \cdot S^{-1} \cdot Q^T)^{-1} \cdot (Q \cdot S^{-1}) \\ \tilde{Q}^T &= -(S^{-1} \cdot Q^T) \cdot (P - Q \cdot S^{-1} \cdot Q^T)^{-1} \\ \tilde{S} &= S^{-1} + (S^{-1} \cdot Q^T) \cdot (P - Q \cdot S^{-1} \cdot Q^T)^{-1} \cdot (Q \cdot S^{-1}) \end{aligned} \quad (12)$$

Further partitioning of the matrix into nine sub-matrices was tested but this did not result in a decrease in computation time. The inversion by partitioning technique was also found to be most effective for coefficient matrices with a size in excess of 30 by 30; for matrices below this size the SVD matrix rank reduction technique was computationally faster.

Once the target values for a cluster have been estimated

they are inserted into the working data set immediately and the pixels flagged as known. The algorithm then carries on searching from the initial local origin that identified the now infilled target segment, row-by-row, moving across then down the image until the next Ground clutter pixel is located, at which point the process is repeated. Figure 13 is an example where the BTNN technique has been used to remove ground clutter, as before the left image contains the marked ground clutter and the right image the estimated reflectivity data.

To estimate the missing values in a typical radar rainfall image, the BTNN technique takes approximately 5 to 20 seconds. However, the time is highly dependent on the range used for selecting the Kriging neighbourhood for each segment. The greater the number of control points the greater the size of the coefficient matrix that needs to be inverted, where the computation time is dominated by the matrix inversion.

Comparison of the two nearest neighbour infilling techniques

An extended formal and detailed comparison of the two methods, ITNN and BTNN has yet to be done. However, as a preliminary example of such a comparison, on 18 November 2000, a radar reflectivity image from the Durban weather radar with no ground clutter but containing convective rainfall was selected. The Bethlehem ground clutter map (not coinciding with the Durban one) was then superimposed on this image, as indicated in Fig. 14. The missing reflectivity data were then estimated using both the ITNN and BTNN techniques. In this application of the ITNN technique the 20 nearest neighbours for each target were used and for the BTNN technique pixels within a range of 2.5 from the ground clutter border were selected. Ordinary




Fig. 13. Bethlehem radar rainfall reflectivity image before and after ground clutter removal (25 February 2003). Black areas indicate contaminated data. The square surrounding each radar image has sides of length approximately 270 by 270 kilometres.

Kriging was used to compute the missing reflectivity data with a semi-variogram function having $L = 11$ pixels and $\alpha = 1.5$ being used in both cases. The estimated values were compared to the real reflectivity values and the square of the errors computed at each target point; these sets are shown in the lower pairs of diagrams in Fig. 14. The computation time was recorded for each technique; Table 1 gives a summary of the results.

For this particular radar reflectivity image, the ITNN technique performed better than the BTNN technique. Not only was the sum of the errors squared lower than the BTNN technique but the computation time was faster by a factor of ten. However, as stated earlier, a more comprehensive testing process needs to be undertaken to determine the most appropriate technique on a variety of rainfall types.

Limitations may exist for very large areas of ground clutter

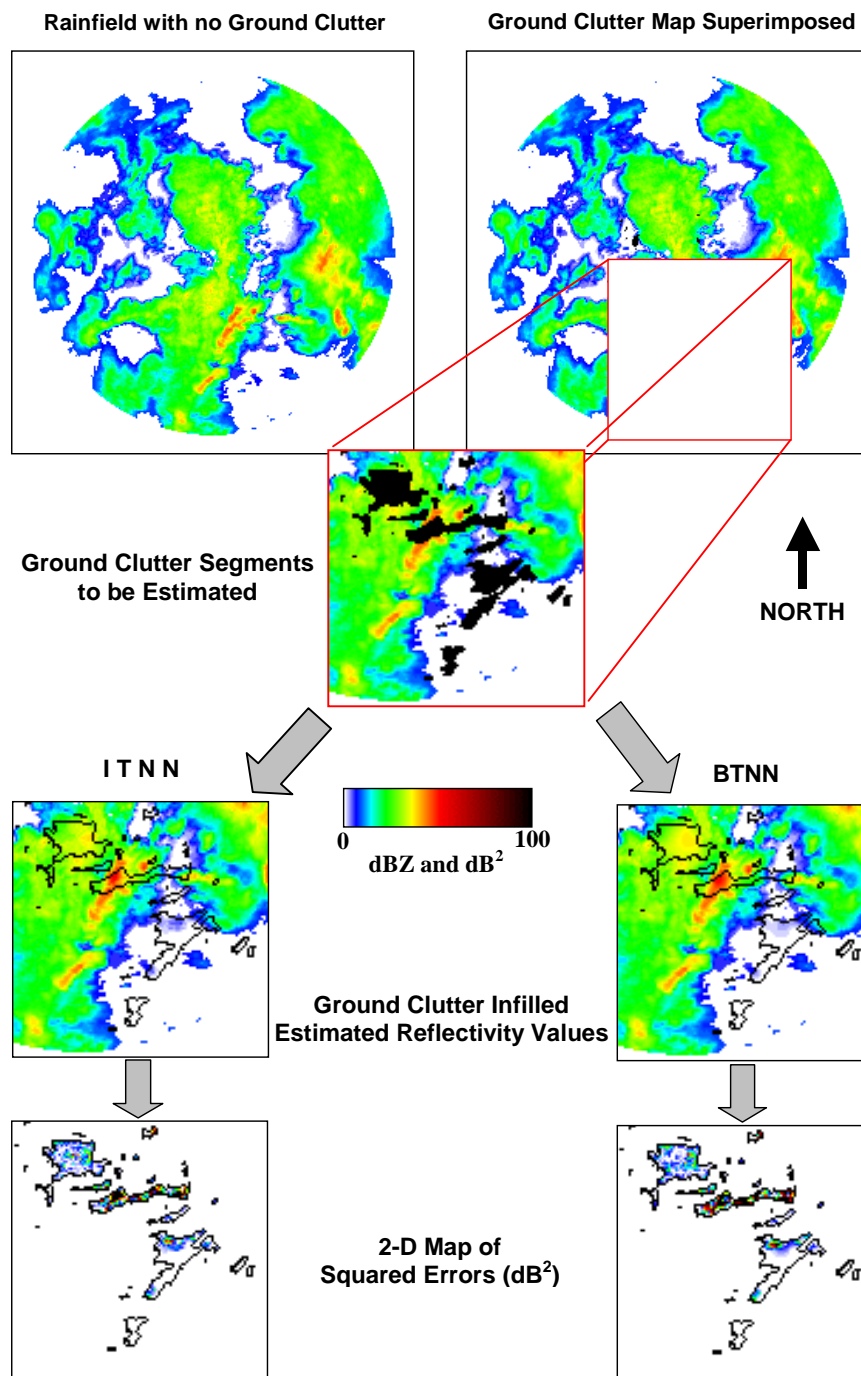


Fig. 14. Comparison of ITNN and BTNN techniques.

Table 1. Errors and computation time for techniques tested in the example in Fig. 14.

	Sum of squared errors (dB ²)	Computation time (seconds)
ITNN	14 700	1.66
BTNN	19 000	16.8

where the best estimate of the rainfall near the centre of the ground clutter will be the mean value of the control points (if the distance exceeds the correlation length) which may result in errors of high magnitude. Large errors may also occur where ground clutter completely or partially hides highly convective rainfall which is surrounded by stratiform rain, this phenomenon occurs in one of the images shown in Fig. 14.

Rainfall estimation at ground level

The rainfall at ground level is of more interest than measurements aloft in Earth-bound applications such as hydrology and agriculture. Accurate rainfall estimates at ground level will also assist in the current Water Research Commission (WRC) research project in South Africa concerned with daily rainfall mapping over South Africa (Sinclair and Pegram, 2003). For ground level rainfall estimates, the Kriging technique can be extended easily from two dimensions into a three dimensional space, as illustrated by the following example.

In this algorithm the control points are the information in the pixels aloft which are Kriged onto target points in a 200-kilometre diameter circle at ground level. The Kriging is carried out with an ITNN approach, as suggested by Seed and Pegram (2001), with the 25 nearest control points to each target being selected.

An SVD matrix rank reduction technique was once again used to ensure computational efficiency. However, when Kriging in a 3-D data set, a 3-D semi-variogram model is used, as given by Eqn. (13), developed by and reported in Seed and Pegram (2001).

$$\gamma(h) = \sigma^2 \cdot [1 - \exp(-d^{2H})] \quad (13)$$

where:
$$d^2 = \left(\frac{r}{r_0}\right)^2 + \left(\frac{h}{h_0}\right)^2 \quad \rightarrow$$

$$d^2 = \left(\frac{x^2 + y^2}{r_0^2}\right) + \left(\frac{h}{h_0}\right)^2$$

where σ^2 is the field variance, r is the distance in the

horizontal plane, r_0 the horizontal correlation length, h is the distance in the vertical plane, h_0 the vertical correlation length and H the scaling exponent.

An example application of this technique is illustrated in Fig. 15. The image is from the Durban weather radar, 9 November 2000. The following semi-variogram values were used in the Ordinary Kriging process: $r_0 = 15$ km, $h_0 = 12$ km and $H = 1,25$ (Seed and Pegram, 2001).

Summary and conclusion

Radar rainfall data have importance in many fields of endeavour. However, the fact that errors persist in these data means that the rainfall estimates are not as accurate as they might be. These errors, noticeably ground clutter, can be detected and identified in the radar images. The technique proposed for estimating the missing data in this paper was Ordinary Kriging. Various methods to improve computational efficiency such as using nearest neighbours to a target point, matrix rank reduction techniques and inversion by partitioning have been implemented in order to make the Ordinary Kriging process computationally efficient in real-time applications.

Two techniques for removing ground clutter contamination have been developed which exploit the screening effect in Kriging. These are Individual Target Nearest Neighbours (ITNN) and Border Tracing Nearest Neighbours (BTNN) techniques. These methods still need to be tested thoroughly to determine which is optimal in terms of accuracy and appropriateness, but preliminary testing on typical sample sizes indicates that ITNN is ten times faster than BTNN and is more accurate. Both algorithms exhibit sufficient computational speed for on-line data cleansing, when supported by computationally efficient methods: (i) rank reduction using Singular Value Decomposition and (ii) matrix partitioning. It is anticipated that a further advantage of the ITNN Kriging approach would be to exploit local semi-variogram models evaluated dynamically; this has yet to be implemented but will be an advantage where radars sample convective and stratiform rainfall simultaneously.

The problem of rainfall estimation at ground level is of importance. The Ordinary Kriging technique was used in three dimensions to provide an estimate of rainfall at ground level. An ITNN approach was employed to extrapolate from control points aloft to targets at ground level. Further work needs to be carried out on the ground level rainfall estimation procedure to improve computational efficiency and validate the resulting rainfall estimates. This technique will be particularly useful where ground clutter targets are to be infilled, by importing information from aloft as well as laterally.




Fig. 15. 3D illustration of measured rainfall reflectivity data at 1, 2 and 3km above ground level and Kriged reflectivity estimate at ground level.

The techniques presented in this paper show considerable promise and, as the ideas mature, should prove valuable in hydrometeorology.

Acknowledgments

This work is currently supported financially by the South African Water Research Commission (WRC) and the National Research Foundation (NRF). Thanks go to Scott Sinclair for Figs. 7 and 8 and work relating to them, also to Deon Terblanche and Pieter Visser of SAWS: METSYS at Bethlehem, South Africa for invaluable methodological input and the ground clutter maps. Finally, Alan Seed of the Australian Bureau of Meteorology is thanked for drawing attention to the reference to the ‘screening effect’ in Kriging and continual encouragement in the endeavour.

References

- Aitken, A.C., 1967. *Determinants and Matrices*. Oliver and Boyd Ltd. London, UK.
- Bell, T.L., 1987. A Space-Time Stochastic Model of Rainfall for Satellite Remote Sensing Studies, *J. Geophys. Res.*, **92**, D8, 9631.
- Chiles, J.P. and Delfiner, P., 1999. *Geostatistics – Modelling Spatial Uncertainty*. Wiley, New York, USA.
- Cressie, N. 1993. *Statistics for Spatial Data*. Wiley-Interscience Publication. New York, USA.
- Journel, A.G. and Huijbregts, Ch. J., 1978. *Mining Geostatistics*. Academic Press. London, UK.
- Marshall, J.S. and Palmer, W.M., 1948. The Distribution of Raindrops with Size. *J. Meteorol.*, **5**, 165–166.
- Pegram, G.G.S. and Clothier, A.N., 2001. High resolution space-time modelling of rainfall: the “string of beads” model. *J. Hydrol.*, **241**, 26–41.
- Press, W.H., Teukolsky, S.A., Vetterling, W.T. and Flannery, B.P., 1992. *Numerical Recipes in C – The Art of Scientific Computing*. 2nd Edition. Cambridge University Press. Cambridge, UK.
- Seed, A.W. and Pegram, G.G.S., 2001. Using Kriging to infill gaps in radar data due to ground clutter in real time. *Proceedings Fifth International Symposium on Hydrological Applications of Weather Radar – Radar Hydrology*. Kyoto, Japan. 73–78.
- Sinclair, D.S. and Pegram, G.G.S., 2003. Combining traditional and remote sensing techniques of rainfall measurement as a tool for Hydrology, Agriculture and Water Resources Management. *Proceedings 11th SA National Hydrology Symposium*, Port Elizabeth, South Africa.
- Sonka, M., Hlavac, V. and Boyle, R., 1999. *Image Processing, Analysis, and Machine Vision*. 2nd Edition. PWS Publishing. Pacific Grove.
- Todini, E., 2001. A Bayesian technique for conditioning radar precipitation estimates to rain-gauge measurements. *Hydrol. Earth Syst. Sci.*, **5**, 187–201.
- Vezzani, G., 1994. Ground clutter Suppression in Atmospheric Phenomena Observation. *COST 75 Weather Radar Systems International Seminar*. Brussels, Belgium. 281–295.
- Visser, P., 2003. *Spatial interpolation and Mapping of rainfall: 2. Radar and Satellite products*, Draft final report (K5/1152) to the Water Research Commission for the period March 2000 to March 2003. Water Research Commission. Pretoria. 10–19.