MAPPING THE N-YEAR DESIGN RAINFALL - A CASE STUDY FOR THE WESTERN CAPE

S. Khuluse[1,2], M. Dowdeswell[1], P. Debba[2] and A. Stein[3]*

[1]Department of Statistics and Actuarial Science, University of Witwatersrand, Johannesburg, South Africa
[2]Council for Scientific and Industrial Research (Built Environment), Pretoria, South Africa
[3]Faculty of Geo-Information Science and Earth Observation, University of Twente, Enschede, Netherlands

ABSTRACT

Flooding is often associated with heavy rainfall. Hence quantifying the probability associated with heavy rainfall events is useful in hydrology for design flood estimation and mapping. Rainfall varies over time and space, therefore it is anticipated that at high levels the process will also vary over time and space. The objective of this study is to regionally quantify the average size of a 24-hour winter rainfall event associated with a 2% chance of being exceeded for an area in the Western Cape, South Africa. The point process approach to extreme value theory is employed to quantify the 50-year winter rainfall return level estimate from daily rainfall data from fifteen stations in the Western Cape. Ordinary kriging is used to estimate the 50-year winter rainfall return level surface over the study area. We compare the result to those obtained when the universal kriging approach is undertaken. We present a technique used to obtain spatial correlation models used in kriging, where the return level estimates are extended spatiotemporally, to circumvent inaccurate specification of the model and its parameters as a result of the spatial sparseness of the sample.

1. INTRODUCTION

An estimate of the N-year design rainfall event is an important input for estimating the N-year design flood. Flooding is not a single site event, hence the need for design rainfall estimation over space. A common approach in hydrology for spatial modelling of extreme values is the regional frequency analysis (RFA) method (Hosking & Wallis, 1993, 1997). It involves delineating the study area into homogeneous regions through cluster analysis of the observations. The extreme rainfall frequency distribution for each homogeneous region is estimated by pooling data from sites within that region and fitting the generalized extreme value (GEV) distribution through the Probability Weighted Method of Moments (PWM) to the new series which is an average over that region. Spatial variation is a result of variation in rainfall of different regions. Smithers & Schulze (2000, 2002) used the RFA method to derive estimates of short and long duration design rainfall and design floods events for South Africa. To address the issue of instability in parameter estimates as duration changes, Smithers & Schulze (2002) augmented the regional L-Moment algorithm by incorporating a scale invariance factor. The main criticism against the RFA is the implausibility of the finite mean requirement in some practical cases (Katz et al., 2002).

Other approaches include the latent spatial process (LSP) (Casson & Coles, 1999) and the max-stable approach to extreme value theory. For the LSP approach, the parameters of the extreme value distribution are considered to be realizations of an unobserved spatial process, which can be modelled using the spatial regression framework. Spatial characteristics are incorporated as independent variables in the model. Spatial variation is the result of requiring that the parameters vary smoothly over space. An estimate of the design rainfall surface, which is a quantile of the extreme value distribution, is obtained by substituting the estimate of the parameter surface. The max-stable process approach involves modelling spatial dependence as a multivariate extreme value distribution based on the theory of max-stable processes (Coles, 2001). Theoretical developments for this approach are advance, however application is lacking due to the computational burden brought on by high dimensionality.

An approach followed in this study, involves applying classical geostatistics to rainfall return level estimates (or design values) (Szolgay et al., 2009). The rainfall return level estimates are an output of a point process extreme value model for the rainfall series at each site. These estimates become an input in the geostatistical analysis, where semivariogram models infer about local spatial variation of these values. The aim of the geostatistical analysis is to obtain an estimate of the design rainfall

surface or map. In this approach, once the semivariogram model is obtained, prediction of the design values at unsampled locations is achieved through kriging. The result is a regional map of the N-year design rainfall, which is an important input in simulation of run-off resulting from heavy rainfalls of such magnitude. While geostatistics for mapping design rainfall has been employed, the issue of small sample sites is often left

In the next section, the point process extreme value model and the geostatistical method used to arrive at an estimate of the 50-year design rainfall surface is discussed. The results and discussion follow in section 3, and the final remarks on the study are contained in section 4.

2. MATERIALS AND METHODS

The data consists of observed daily rainfall data from fifteen weather stations in the Western Cape province of South Africa. The study area is bounded between latitudes $-34.058\,°\mathrm{S}$ and $-32.463\,°\mathrm{S}$ and longitudes $18.157\,°\mathrm{E}$ and $26.493\,°\mathrm{E}$, covering an area of approximately $132000\,\mathrm{km}^2$. The Western Cape is classified as a winter rainfall region. The study is therefore restricted to rainfall events in the months June to August. Seven stations had observation periods of fifty years, whilst for other stations the periods are shorter.

Given a random sample of observations, the extreme value limit theorem states that the distribution of the largest (or smallest) member of that sample can be approximated by the Generalized Extreme Value distribution (Coles, 2001). In the point process approach to EVT (Pickands, 1971), the rainfall series at each site is reformulated as a bivariate point process $\{(i, X_i); \ i = 1, 2, \ldots, n\}$. Here first dimension refers to the position of X_i in the sequence and the second dimension is the magnitude of that rainfall event. As a consequence of the extreme value limit theorem, on bounded regions of the form $\mathcal{A} = [t_1, t_2] \times [u, \infty)$, the number and size of threshold exceedances can be approximated by the Poisson and the generalized Pareto distributions, respectively (Coles, 2001). An extensive study into the sensitivity of model parameters and assumption is done (details in Khuluse, 2010) to determine the value at each site where the observation can be assumed independent and where the generalized Pareto distribution (GPD) approximation to exceedances is plausible. This includes examining the sensitivity of the average excess, the dispersion index and return level values to a range of threshold values. Inference for the extreme value model is done by maximum likelihood. The parameter estimates are used to obtain the 25 to 50 year return levels (at 5 yearly intervals).

To determine whether there is spatial variation in the 50-year rainfall return level within the study area, classical geostatistics is applied to the estimated return levels obtained for each site. With only fifteen stations direct estimation of the variogram model leads to unreliable parameter estimates(Stein & Sterk, 1999). The technique used to increase the sample size for estimating the variogram is done by considering return level estimates for return periods 25 to 50 years (in intervals of 5 years) as temporal replicates of the return level surface. Design values at each site are assumed to be a sample from a random variable which forms continuous surface over the study area. This continuous surface or spatial random field is denoted as $\{Y(\mathbf{s}, \mathbf{t}) : \mathbf{s} \in \mathbf{D} \subset \mathbb{R}^2, \mathbf{t} \in \mathbf{T}\}$, where D is a fixed, continuous subset of a two-dimensional plane and $\mathbf{t} \in \mathbf{T}$ is an index for the temporal component. The primary objective is to make inference about the unknown underlying design rainfall process from which we assume our sample to be generated. An important property in spatial analysis is that the strength of the association between attribute values decreases as the distance between measurement locations increases, that is spatial autocorrelation.

> "...the sensitivity of geographic and other phenomena to local interactions implies that we should carefully measure and analyze relations among near things." (Miller, 2004)

While the covariance statistic is useful in quantifying spatial autocorrelation, it is common in geostatistics to use the variogram $2\gamma(\mathbf{h})$ which is also a measure of the second-order property of the underlying process and related to the covariance as follows,

$$\gamma(\mathbf{h}) = C(\mathbf{0}) - C(\mathbf{h})$$

A vector of distances between locations is denoted as $\mathbf{h}$.

In geostatistics, it is assumed that spatial correlation can be represented by some parametric form. Estimation of the parameters of the variogram model when the sample size is small may lead to unreliable results (Cressie, 1991). With design rainfall values at just fifteen sites, estimation and prediction uncertainty is likely to be high. The method of regionalisation as described in (Stein & Sterk, 1999; Sterk et al., 2004) is used to extend the observations in space and time in order to reliably

estimate the variogram. Hence, design values are considered as space-time observations, with space relating to their position on the earth's surface and time represented by a sequence of return periods.

Generally in statistics, assumptions are made to enable modelling of a phenomenon of interest. One such assumption is that of random variables being 'independent and identically distributed'. Its counterpart in geostatistics is *strict stationarity*, where if any set of locations are shifted spatially by $\mathbf{h}$, then observations from the two sets of locations will have the same probability density. Strict stationarity can be relaxed by requiring the mean and covariance to be stationary, that is a constant mean and covariance which only depends on the distance between any two locations. The weakest assumption is that of a constant mean and a stationary variogram, termed *intrinsic stationarity*. Another property that is required to be able to estimate the parameters of the model for the underlying unknown spatial process is *ergodicity*. This means that the single realization of the process must be able to wander through all possible values that the process can take. Therefore, in spatial statistics stationarity and ergodicity are necessary assumptions. The space-time field from which design rainfall values are sampled is assumed to be intrinsically stationary and ergodic (Stein & Sterk, 1999). That is

$$
\begin{aligned}
0 &= E_{\mathbf{s},\mathbf{t}}[Y(\mathbf{s}+\mathbf{h_s}, \mathbf{t}+\mathbf{h_t}) - Y(\mathbf{s},\mathbf{t})] \\
2\gamma(\mathbf{h_s}, \mathbf{h_t}) &= \mathrm{var}(Y(\mathbf{s},\mathbf{t}) - Y(\mathbf{s}+\mathbf{h_s}, \mathbf{t}+\mathbf{h_t}))
\end{aligned}
$$

For replication in time of return level estimates, consider the sequence of equally spaced return periods as pseudo-time. The return periods satisfy the requirement that $t_i \bigcap t_j = \emptyset$ for $i \neq j$. To satisfy the requirement of a constant mean for intrinsic stationarity, the data is standardized by the ratio of the overall average design rainfall (m_0) and the average design rainfall for that specific period (m_{t_j}), $j = 1, 2, \ldots, p$. The standardized design value is determined as

$$
\tilde{Y}(\mathbf{s}, \mathbf{t}) = Y(\mathbf{s}, \mathbf{t}) \times \frac{m_0}{m_{t_j}} \tag{1}
$$

This removes the effect of temporal replication, making the assumption of $\mathbf{h_t} = 0$ plausible. Hence, the space-time variogram is reduced to a spatial variogram. At each location there will be p standardized observations. These can be extended in space, by displacing the set of locations by a fixed distance c repetitively for $p-1$ instances. The displacements are as follows $\{c\mathbf{s}, 2c\mathbf{s}, \ldots, (p-1)c\mathbf{s}\}$ where $\mathbf{s}$ is the index for the original coordinates. Recall in analyzing spatial autocorrelation, emphasis is on distances between paired locations rather than the exact locations. In this way a larger data set is created for the purpose of choosing a variogram model and estimating its parameters. Three models were considered: the exponential, the spherical and the penta-spherical models. The model forms can be found in Cressie (1991). The penta-spherical model is given by

$$
\gamma(\mathbf{h}) = \begin{cases} C_0 + C\left(\frac{15}{8}\frac{\mathbf{h}}{a} - \frac{5}{4}\left(\frac{\mathbf{h}}{a}\right)^3 + \frac{3}{8}\left(\frac{\mathbf{h}}{a}\right)^5\right), & 0 \leq \mathbf{h} < a \\ C_0 + C, & \mathbf{h} \geq a \end{cases} \tag{2}
$$

where C_0 is the nugget, which is composed of micro-scale variation and measurement error. The partial sill is C and the range is a. The range is the distance beyond which there is lack of spatial correlation. Estimation in this study is by weighted least squares (WLS) (Cressie, 1991). In WLS, the weights increase with the number of pairs in each lag class and distance closer to the origin. Hence it is important for the variogram to fit well near the origin. The estimates of the partial sill and nugget for a particular return period are obtained by multiplying estimates of the variogram model for the pooled data with the square of the reciprocal of the standardization factor. The re-scaled parameters are input in the kriging procedure to obtain the desired design rainfall map. Kriging is a spatial prediction technique and details on its application in this study are discussed in the next subsection.

Standardization by the ratio of the overall average return level estimate to the one for that particular return period is done to remove the effect of different expectations for each return period. Replication in space is done by displacing the standardized observations by a fixed distance of $100\,\mathrm{km}$. The variogram is modelled using this extended sample. The next step is the prediction of design rainfall values at unobserved sites through kriging.

Kriging is a generalized least-squares technique that allows one to account for spatial dependence in the observations as given by the variogram model (Cressie, 1991). In kriging, design values at unsampled locations will be predicted by the predictor

$$
Y(s_0) = \sum_{i=1}^{p} \lambda_i y_i, \quad \text{with} \quad \sum_{i=1}^{p} \lambda_i = 1 \tag{3}
$$

Kriging is known as the Best Linear Unbiased Predictor (BLUP). This is due to the kriging weights being chosen to minimize estimation variance $\text{Var}(Y(s_0) - Y(s))$, whilst ensuring that the estimator is unbiased, i.e. $\text{E}(Y(s_0) - Y(s)) = 0$. *Linearity* is a result of the predictor being a linear combination of neighbouring observations at un-sampled locations. The weights are obtained by solving the kriging system of equations as detailed in Cressie (1991).

In *ordinary kriging* a spatial process is assumed constant over the study regions. If there is evidence of a global spatial trend, the random process at location s can be defined as

$$U(s) = X^{\text{T}}\beta + \epsilon(s) \tag{4}$$

The model for spatial variation due to a non-constant mean of the process is given by

$$\text{E}(U(s)) = \sum_{k=1}^{q} \beta_k f(s) \tag{5}$$

where q is the number of beta coefficients. For variogram estimation initially parameters of the trend surface model given in equation 5 are estimated. Once the global trend has been removed any remaining spatial correlation is detected as variation in residuals, modelled as the residual variogram. Estimation of the residual variogram continues similarly to variogram estimation for observations in the constant mean case.

For the *universal kriging* predictor, a predicted value can be expressed as a linear combination of observed values

$$U(s_0) = \sum_{i=1}^{p} \lambda_i u(s_i). \tag{6}$$

Kriging weights are obtained by minimizing

$$\left(U(s_0) - \sum_{i=1}^{p} \lambda_i U(s_i)\right)^2 = \left(X^{\text{T}}\beta + \epsilon(s_0) - \lambda^{\text{T}} X \beta - \sum_{i=1}^{p} \lambda_i \epsilon(s_i)\right)^2 \tag{7}$$

To derive the ordinary kriging map for the 50 year return level, the parameters are re-adjusted by the reciprocal of the standardizing ratio to derive the variogram model for the 50 year return level estimates. Kriging is applied to the fifteen 50 year return level estimates using the variogram model corresponding to this return period to derive the 50 year winter rainfall return level surface. We compare the results of ordinary and universal kriging in the next section.

3. RESULTS

The sample proportion of exceedances at each station is generally small (the percentage of average number of exceedances is above 3% for only three of the fifteen sites). For Cape Town International Airport, the exceedance percentage is 1.6% for the chosen threshold of 26 mm. The proportion of exceedances is smaller (at only approximately 0.01) for stations towards the east such as Jonkersberg, Plettenbergbay and Tygerhoek. This can be attributed to rainfall values being mostly low, with few instances where large amounts of rainfall were observed. The location parameter appears stable over the region whereas the scale parameter increases into the south easterly direction. The model fits the exceedances at each station adequately. The degree of uncertainty however is large for stations such as Tygerhoek, where the largest values were approximately twice as large as the other excesses. In such cases it was tempting to remove these large values, but we avoided this to minimize the risk of obtaining inaccurate return level estimates. The 25-year design value for Ladismith is the smallest at 22.42 mm. This can be attributed to the station's generally low rainfall values for the winter period. In contrast, the return level value for Tygerhoek is high, almost six times larger than that of Ladismith at 119.16 mm. Design values show variation between sites, therefore it is of interest to investigate how this can be modelled. Relation between return level estimates and spatial coordinates are investigated. Changes in longitude seem to have no influence on the 25-year design values (figure 2(a)), however, there seems to be a weak negative relation with latitude as seen in figure 2(b). With only 15 estimated return level values across the study region, it is difficult to be overly confident of the spatial trend in design values. Further investigation is done by regressing the design values to the spatial coordinates. The model reveals that the spatial coordinates do have an overall influence on the return level values ($p - \text{value} = 0.028$), but again the strength of this linear relationship is weak - only 36% of the variation in
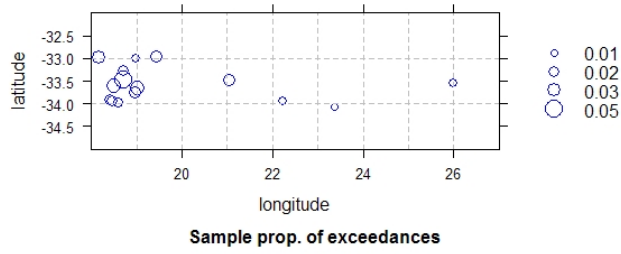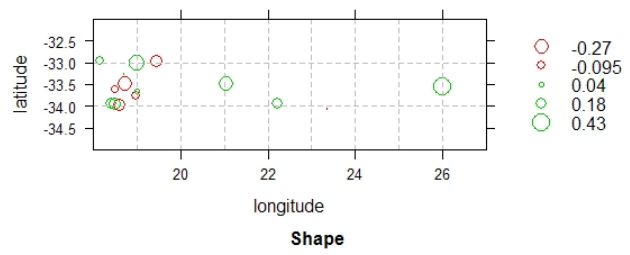
(a) Sample proportion of exceedances
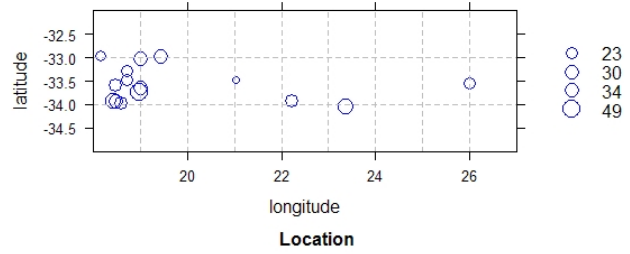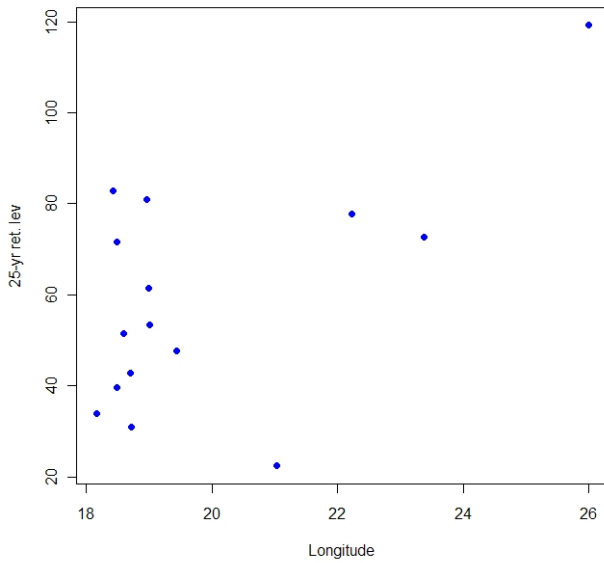
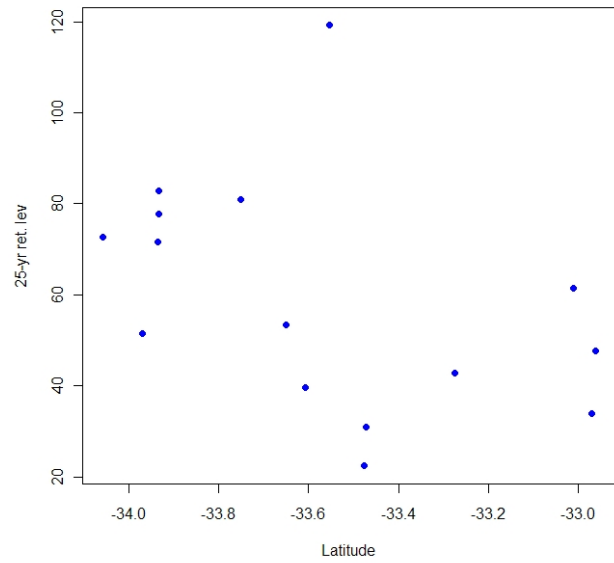(b) Location parameter estimates
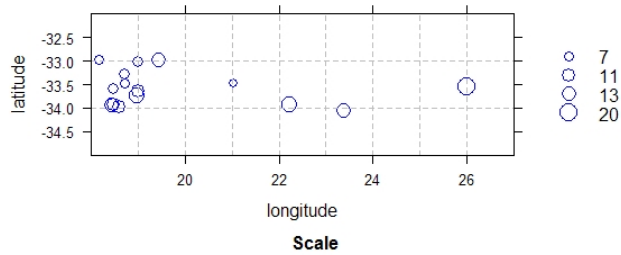
(c) Scale parameter estimates

(d) Shape parameter estimates

Fig. 1. Summary of the parameter estimates of the point process extreme value model at each site

(a) 25-year return level estimate vs longitude

(b) 25-year return level estimate vs latitude

Fig. 2. Relation of the 25-year return level estimate to the coordinates

the return level estimates is explained by spatial coordinates.

From the regression results there is nearly a 3% chance that the positive relation between design values and longitude is

due to chance only, whilst for the latitude this is high at nearly 20%. Choosing the level of significance as 5%, there is evidence (albeit weak) in support of a linear spatial trend in design values. However, this global trend explains a small proportion (0.36) of the variation in the design values. It is anticipated that the remainder of the variation can be explained by the model for local spatial variation. Diagnostic plots, namely the normal quantile-quantile plot and the plot of residuals against fitted values concur on the inadequacy of the linear spatial trend. Influential observations are highlighted as those corresponding to Ladismith, Molteno and Tygerhoek. This is understandable since the 25-year design rainfall value for Ladismith (22.42 mm) is small in comparison to other sites, whilst 119.16 mm for Tygerhoek is more than 70% greater than values obtained for the rest of the sites. The q-q plot also reveals violation of the normality assumption of the residuals. Potentially such discrepancies could be a result of the small sample size, since the normal approximation is reliable for sample sizes of at least $25 - 30$ observations.

(a) Checking for patterns between fitted and residual values
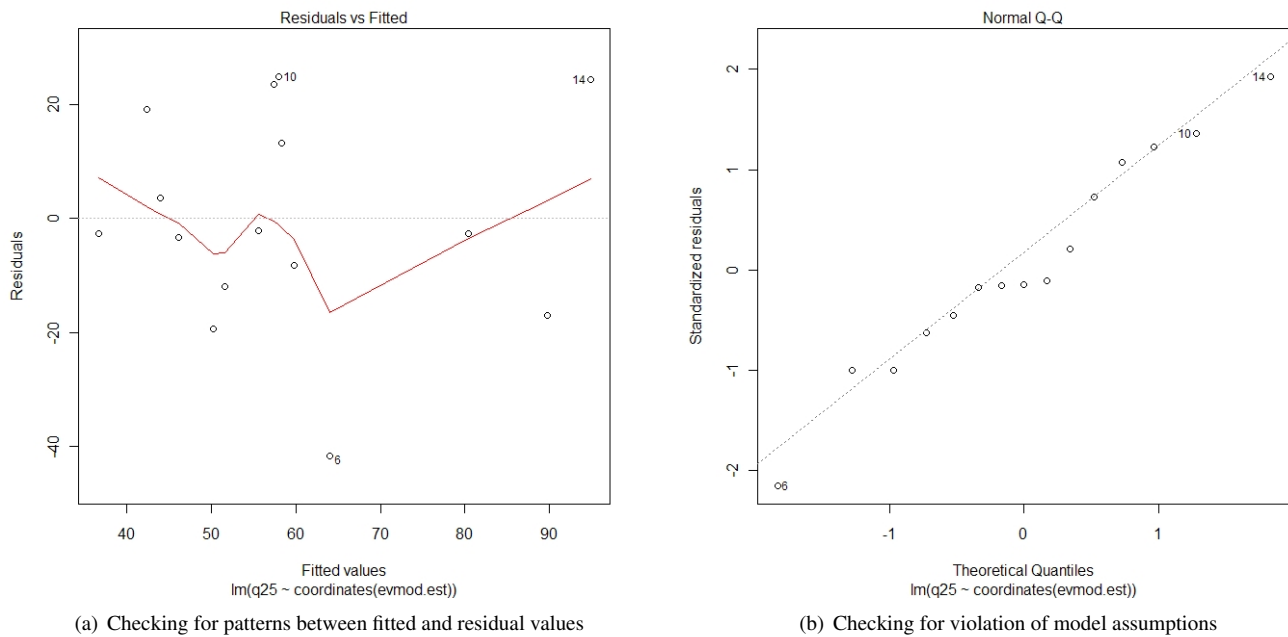
(b) Checking for violation of model assumptions

Fig. 3. Graphics for evaluating closeness of the fitted model to observed values and for checking the validity of model assumptions

Exploratory analysis for regional trend revealed that only 36% of the variation is due to a regional trend. Therefore, it is anticipated that the remaining variation in design values is due to localized spatial effects, which can be modelled through a variogram. A particular challenge in this case is that there are only fifteen sites, which is too few observations spatially. To overcome this challenge space-time replication of design values is performed, which enables derivation of parameter estimates for the variogram model. For the replication in time, the design values of a specific period are considered as realization of an unobserved return level surface for that period. Specifically, six consecutive return periods are considered from 25 to 50 years in intervals of five years (see Table 1). Thereafter, return level estimates at each site are standardized by a factor that is a ratio of the overall average design value to the average design value for a specific return period. Return levels are strictly positive values, hence the use of the ratio of averages instead of the normal standardization by a ratio of the difference from the mean and the standard deviation. The reason for the standardization is to counter the challenge of return level estimates that increase with increasing return period. Once standardized, the spatial autocorrelation property stating that the distance between points rather than their position in space determines similarities in attribute values, is used to justify replication in space. The spatial replication procedure involves creating observations at new sites by displacing sequentially five of the six observations from the original sites by a fixed distance, in our case 100 km which was chosen arbitrarily. Physical reasoning and evidence may be used to inform the choice of displacement distance, however in this case such information was not available.

A variogram cloud was computed to check whether the structure of the spatial correlation is the same across the different return

Table 1. Replicates of the design values in pseudo-time

Site	Long.	Lat.	q_{25}	q_{30}	q_{35}	q_{40}	q_{45}	q_{50}
Atlantis	18.483	-33.607	39.63	40.72	41.63	42.41	43.09	43.69
CPT Astro.	18.477	-33.935	71.56	75	77.99	80.66	83.06	85.26
CPT Int. Air.	18.597	-33.969	51.55	52.69	53.63	54.42	55.1	55.7
Excelsior	19.43	-32.963	47.58	49.34	50.77	51.98	53.01	53.92
Jonkers.	22.227	-33.934	77.86	82.49	86.52	90.09	93.31	96.24
Ladismith	21.035	-33.476	22.42	24.17	25.73	27.14	28.42	29.61
Langebaan.	18.157	-32.972	33.99	35.56	36.91	38.1	39.17	40.13
Langgewens	18.706	-33.276	42.84	44.23	45.41	46.42	47.32	48.13
Malmesbury	18.718	-33.472	30.95	31.75	32.39	32.92	33.38	33.77
Molteno	18.417	-33.933	82.8	86.66	90.02	92.99	95.68	98.12
Paarl	18.967	-33.75	80.96	83.15	84.96	86.51	87.87	89.06
Plettenberg.	23.372	-34.058	72.77	75.16	77.18	78.92	80.46	81.84
Porteville	18.994	-33.012	61.52	64.9	67.93	70.68	73.22	75.57
Tygerhoek	25.993	-33.553	119.16	130.82	141.41	151.18	160.27	168.8
Wellington	19.006	-33.651	53.45	55.58	57.38	58.96	60.35	61.61

periods. Sample variograms as shown in figure 4 were computed. In this study two cases are considered – in the first case lack of regional trend is assumed and in the second case a linear trend surface is considered, hence variogram model (Fig. 4(b)) is for residuals resulting from removal of the regional trend. As anticipated the semivariances for the residuals would be lower as some of spatial variation has been accounted for by the regional trend, but the spatial correlation structure remains similar for the two cases.
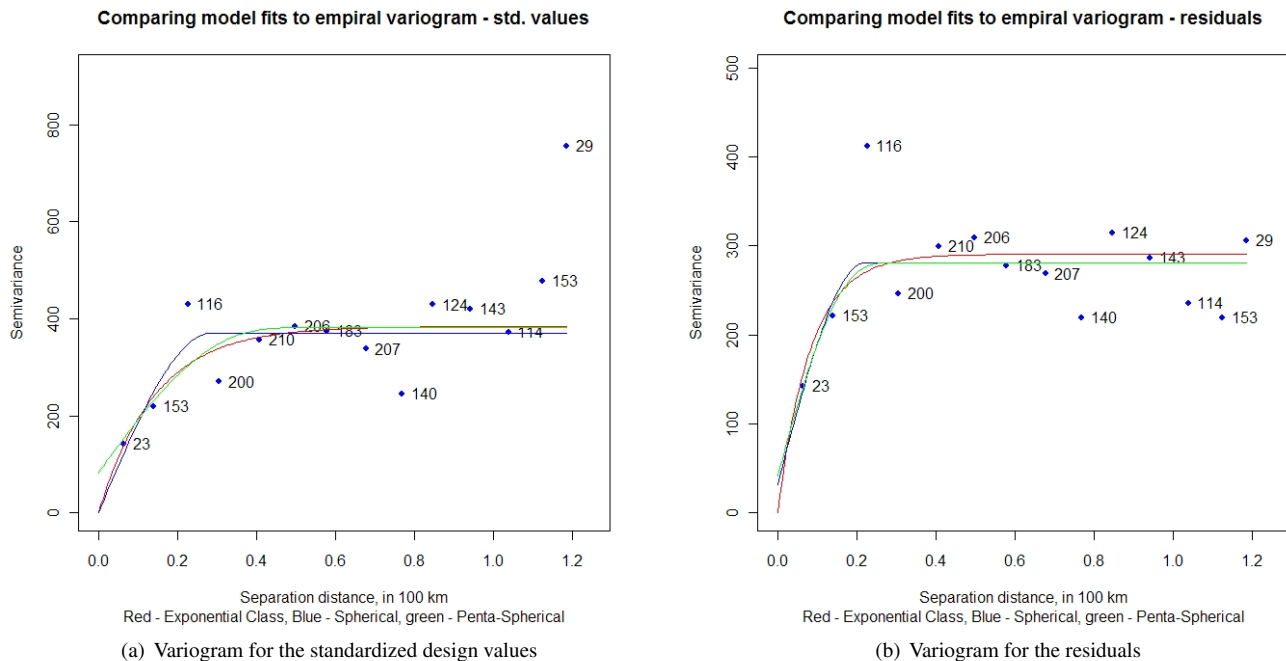

(a) Variogram for the standardized design values

(b) Variogram for the residuals

Fig. 4. Empirical variogram and potential variogram model curves superimposed. In red is the exponential model, blue is the spherical and in green the penta-spherical model.

Variogram model estimation was done by curve-fitting, minimizing the weighted sum of squared differences of the curve to points of the empirical variogram. Three variogram models were considered. From the structure of the sample variograms,

there appears to be definite sills, which are reached gradually. Therefore, the spherical class of models was considered. The exponential model was also considered, but it was the first to be eliminated. The elimination of this model was due to the nugget estimate being zero, which is impractical in this case since design values are estimates and by definition a certain degree of uncertainty is expected. From the plots in figure 4, the penta-spherical model appears to be a closer fit than the spherical model, especially the shoulder of the empirical variogram. The penta-spherical model had the least weighted sum of squares.

Similar variogram analysis was carried out for the second case, the residuals. The weighted sum of squared differences was again minimum for the penta-spherical model. To arrive at the parameter estimates for the 50-year return level, the sill and the nugget were multiplied by the square of the reciprocal of the standardization factor. The nuggets and sills for the first case are $101 \, \text{mm}^2$ and $446 \, \text{mm}^2$ and for the second case $48.21 \, \text{mm}^2$ and $328.17 \, \text{mm}^2$ respectively. The ranges for the two cases are $0.51 \times 100 \, \text{km}$ and $0.28 \times 100 \, \text{km}$ respectively.

Lastly kriging maps are obtained. In both the ordinary and universal kriging cases, due to the small number of observations and the vastness of the study region, the intensity is strong near observed sites. In ordinary kriging, the predicted values are the same as the observed in a circular area around it, whilst in universal kriging there is spread in the direction of the regional trend. The design rainfall level in most of the Western Cape is between $40 \, \text{mm}$ and $80 \, \text{mm}$, but as one moves east this becomes higher, largely due to the influence of the large design value obtained for Tygerhoek. As anticipated for ordinary
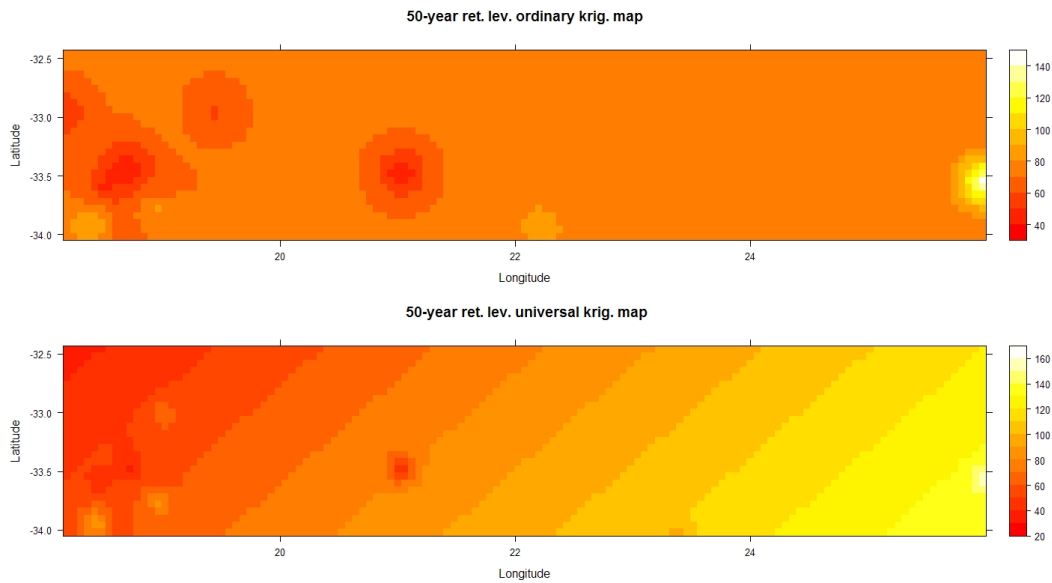

Fig. 5. Comparison of maps derived by ordinary kriging compared with universal kriging for the 50-year 24-hour rainfall return level surface

kriging, the kriging error is greatest at unsampled locations, but it is lower when compared to universal kriging. Region of minimum error for the universal kriging is in the inner spherical area containing all the sites, with the exception of Tygerhoek and Plettenbergbaai which are furthest stations from the rest.
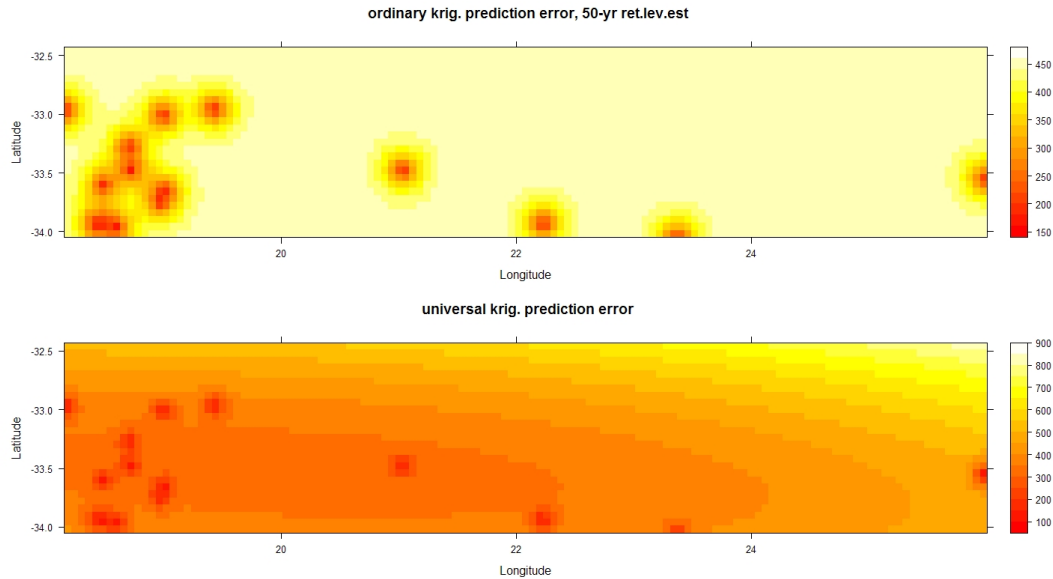
Fig. 6. Map of the uncertainty about the ordinary and universal kriging predictions for the 50-year 24-hour rainfall return level surface

4. DISCUSSION AND CONCLUDING REMARKS

In this study the point process approach to extreme value theory was used to quantify rainfall return levels for discrete sites within the Western Cape. The 50 year 24 hour winter rainfall return level is the single-day average amount of winter rainfall that can be expected to be exceeded with 0.02 probability. This gives a measure of the rarity of such events. Quantifying this statistic for this study involved taking into account short-range temporal correlations to ensure model accuracy, by declustering the series of exceedances at sites where the level of persistence was strong, as measured by extremal index values smaller than 0.80 for the chosen thresholds. The quantile-quantile and return level plots for evaluating the model's goodness-of-fit did not reveal any inadequacies, except for high uncertainty for large quantiles for sites like Porteville, Ladismith and Tygerhoek. The uncertainty was due to the influence of the largest excess being nearly twice the size of other excesses at these sites. Direct fit of the GPD to threshold exceedances could have been used to quantify the rainfall return levels. In this study the point process approach was chosen because it's parametrization in terms of the generalized extreme value distribution (GEV), ensures that the scale parameter is invariant to threshold choice.

In hydrology an estimate of the rainfall return level surface may be more informative than the discrete site statistic. Ordinary kriging was done to obtain the 50 year winter rainfall return level map from the 15 site-wise estimates. Kriging is the Best Linear Unbiased Estimator, given a model for spatial correlation. In geostatistics spatial correlation is modelled through a variogram, which requires the sample to be sufficiently large, with at least 100 observations (Stein & Sterk, 1999). The method of regionalisation as described in (Stein & Sterk, 1999; Sterk et al., 2004) was used to circumvent this. This ensured that predictions by ordinary kriging of the 50 year daily rainfall return level at unsampled sites were based on a reliable model for spatial correlation.

The main limitations of this study with respect to the spatial analysis of extreme values was that some stations had an observation period shorter than 50 years and that the sample was spatially sparse. These factors affect the quality of the parameter estimates and the resulting rainfall return level map. The spatiotemporal variogram estimation technique presented provided a way of solving the spatial sparseness problem. As further research the use of densely sampled covariate information to improve on the quality of the resulting map can be explored.

REFERENCES

Casson, E. & Coles, S. G. (1999). Spatial regression models for extremes. *Extremes*, 1(4), 449–468.

Coles, S. G. (2001). *An Introduction to Statistical Modeling of Extreme Values*. Statistics. London: Springer Verlag.

Cressie, N. (1991). *Statistics for Spatial Data*. John Wiley and Sons, Inc., New York.

Hosking, J. R. M. & Wallis, J. R. (1993). Some statistics useful in a regional frequency analysis. *Water Resources Research*, 29(2), 271–281.

Hosking, J. R. M. & Wallis, J. R. (1997). *Regional frequency analysis: An approach based on L-Moments*. Cambridge University Press.

Katz, R. W., Parlange, M. B., & Naveau, P. (2002). Statistics of extremes in hydrology. *Advances in Water Resources*, 25, 1287–1304.

Khuluse, S. (2010). Modelling heavy rainfall over time and space. MSc Research Report, University of Witwatersrand, Johannesburg, South Africa. Submitted.

Miller, H. J. (2004). Tobler's First Law and spatial analysis. *Annals of the Association of American Geographers*, 94(2).

Pickands, J. (1971). The two-dimensional Poisson process and extremal processes. *Journal of Applied Probability*, 8, 745–756.

Smithers, J. & Schulze, R. (2000). *Development and evaluation of techniques for estimating short duration design rainfall in South Africa*. Technical Report 681/1/00, Water Research Commission, Pretoria, RSA.

Smithers, J. & Schulze, R. (2002). *Design rainfall and flood estimation in South Africa*. Technical Report K5/1060, Water Research Commission, Pretoria, RSA.

Stein, A. & Sterk, G. (1999). Modeling space and time dependence in environmental studies. *International Journal of Applied Earth Observation and Geoinformation*, 1(2).

Sterk, G., Stein, A., & Stroosnijder, L. (2004). Wind effects on spatial variability in pearl millet yields in the Sahel. *Soil and Tillage Research*, 76, 25–37.

Szolgay, J., Parajka, J., Kohnova, S., & Hlavcova, K. (2009). Comparison of mapping approaches of design annual maximum daily precipitation. *Atmospheric Research*, 92, 289–307.