Application of ANUSPLIN to produce new Intensity-Frequency-Duration (IFD) design rainfalls across Australia

Cynthia The
Hydrologist, Environment and Research Division, Bureau of Meteorology, Sydney, Australia

Michael Hutchinson
Professor of Spatial and Temporal Analysis, Fenner School of Environment and Society, Australian National University, Canberra, Australia

Fiona Johnson
Lecturer, School of Civil and Environmental Engineering
University of New South Wales, Sydney Australia

Catherine Beesley
Hydrologist, Environment and Research Division, Bureau of Meteorology, Melbourne, Australia

Janice Green
IFD Revision Project Manager, Environment and Research Division, Bureau of Meteorology, Canberra, Australia

The Bureau of Meteorology has recently released new Intensity-Frequency-Duration (IFD) design rainfall information as a component of the broader revision of Australian Rainfall and Runoff. New IFDs are available for any location in Australia. However, because of the sparseness of the rain gauge network especially for sub-daily durations, an integral part of deriving the new IFDs was the gridding the available rainfall data. In the IFD Revision project, gridding was undertaken using the software package ANUSPLIN which applies thin plate smoothing splines to interpolate and smooth multivariate data. The splines were fitted using three independent variables; latitude, longitude and elevation. The elevation scale was exaggerated by a factor of 100 to reflect the importance that elevation has on precipitation patterns.

A considerable number of iterations was required to achieve an optimum outcome that represented the observed rainfalls but which did not place too much significance on short rainfall records or from poorly located rainfall stations. The appropriate degree of smoothing of the fitted functions was determined through generalised cross validation with the number of knots and data transformation varied to achieve optimal results. In addition to the statistical tests to determine the appropriate degree of smoothing, qualitative assessments were also conducted by preparing maps which compared the index rainfall derived from at-site frequency analysis of rainfall records, the length of record available at each station, and the spatial density of the rain gauge network to the gridded index rainfalls produced by ANUSPLIN for daily durations.

The final IFD grids were produced by the application of ANUSPLIN using 3570 knots with no transformation of the data.

1. INTRODUCTION

As part of the revision of Australian Rainfall and Runoff, the Bureau of Meteorology (BoM) has released new Intensity-Frequency-Duration (IFD) data corresponding to selected Annual Exceedance Probabilities (AEPs). The estimates are based on statistical analysis of historical rainfall at gauge locations. The point data are then interpolated using the software package ANUSPLIN (Hutchinson 2007). The performance of the interpolation is affected by a number of factors, including the number of data points used in the gridding; or knots, and transformations applied to the data as discussed in The *et al* (2012). Since then further investigations have been done to further optimise the number of knots and the data transformation. Furthermore additional stations have been included in the daily dataset to improve spatial coverage of data since the previous study. These investigations aim to improve the performance of the model and reduce predictive error. Both quantitative and qualitative assessments are conducted. The following section provides a discussion of the approach used for the gridding of daily rainfall data for the project. This is followed by the results and a discussion of their implications. Finally the adopted approach used to produce the IFD 2013 estimates is described.

2. METHOD

This paper discusses the gridding of the location parameter of the Generalised Extreme Value (GEV) distribution. The GEV distribution has been used to model Annual Maximum Series (AMS) data for the IFD revision project (Green *et al*, 2012). Due to the regionalisation of the point data that is undertaken as part of the IFD revision methodology (Johnson et al, 2012), the location parameter has been termed the Index Rainfall. Thin plate splines have been adopted as the analysis method to interpolate spatial point data. The rationale for using thin plate smoothing splines and initial details of the method are described in The *et al* (2012). Since the previous study an additional 545 stations have been included to improve spatial density (Figure 1). This consists of 543 Water Regulations data and 2 BoM daily stations with more than 20 years of data, to supplement the 8074 BoM daily stations with at least 30 years of data used by The *et al* (2012). This made a total of 8619 daily sites available for the analyses.

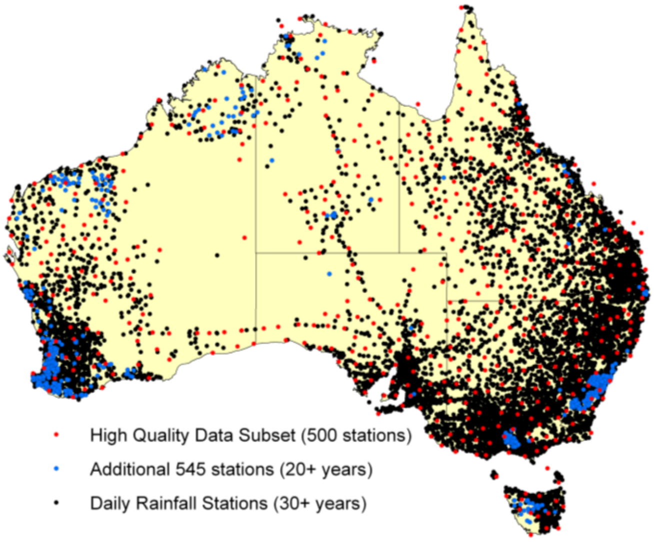

Figure 1 Location of additional stations added to the analysis and high quality data subset used in cross validation. A total of 8619 stations used in the daily analysis.

The previous assessments had primarily focused on the performance of the spline model across Australia. However a review of the results indicated possible areas of regional biases which needed to be better considered. Therefore a regional scale analysis was conducted to identify any issues that could be of concern, and how well the model performs at a local scale. The regional scale analysis was conducted for 29 key river regions and rainfall districts (Figure 2 and Table 3). The selected areas included regions with a significant difference between the draft new IFD grids and AR&R87, with a relatively high density of data, in high gradient areas or associated with key population centres.

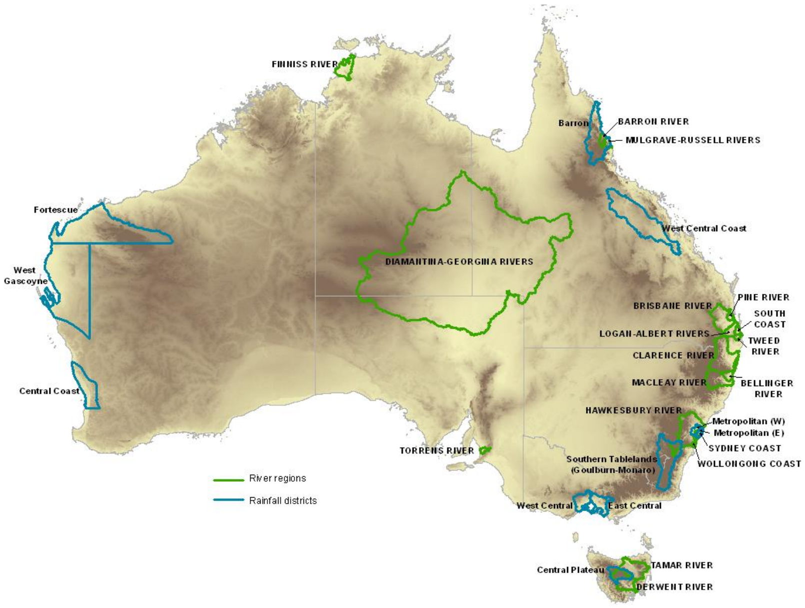

Figure 2 Selected river regions and rainfall districts

2.1. Model optimisation and evaluation

The optimisation of the thin plate spline fits and the evaluation of the different modelling strategies have been achieved by monitoring several summary statistics described in detail in The *et al* (2012). The spline surface for each duration was fitted independently and the ANUSPLIN log file reports separate results for each surface. Summary statistics averaged across all durations Australia wide and on a regional scale were also reported. For all the analysis a 0.025 degree DEM is used based on The *et al* (2012). The key statistics discussed here are:

- Cross Validation (CV) Statistics – the ANUSPLIN package permits each data point to be implicitly left out of the analysis to calculate the individual residual from the fitted surface at each withheld station. The Mean Absolute Error (MAE) and Root Mean Square (RMS) of these individual unweighted residuals are used to evaluate the overall predictive error of the fitted spline surface.
- CV statistics on a high quality (HQ) data subset – the cross validation statistics are also reported for a spatially representative set of 500 stations across Australia, all of which have at least 50 years of data (Figure 1). This removes the bias associated with CV statistics calculated from the complete, but unevenly spaced, data set.
- Surface Residuals – these assess the differences of the data values from the fitted surface values across all stations.

The following methodological questions were addressed in this investigation:

- Optimisation of knot choice for the thin plate smoothing spline
- Appropriate transformations for the index rainfall data
- Appropriate scaling of elevation in the trivariate spline

3. RESULTS AND DISCUSSION

The *et al* (2012) found that the Australia wide predictive error was approximately stabilised when 2550 knots were chosen and the log transformation was applied to the data. However, close inspection of data residuals indicated that this analysis over-smoothed the data in complex areas, such as the high relief topography inland from Tweed Heads. This appeared to be due to both the application of the log transformation to the data, which smoothed large data values significantly more than small data values, as well as an insufficient number of knots. Moreover, an additional 545 daily sites were included in the data set to improve the spatial coverage of data. With these additional stations, further testing has been conducted to optimise both the chosen data transformation and the number of knots.

3.1. Optimisation of knots

Tests were carried out with 3000 and 3500 knots. For both knot sets, the number of knots was increased by 2% taken from the data points with the largest residuals in the initial fits (e.g. initially 3000 then 3060 knots and initially 3500 then 3570 knots). The results are shown in Table 1. These results are based on using a log transformation of the data as recommended by The *et al* (2012).

Table 1 Summary statistics for optimising number of knots (1-7 day duration).

Log Transformation	**All Stations**	**Surface Residuals**		**CV Residuals**		**HQ CV Residuals**	
No.of knots	**GCV**	**RMS**	**MAE**	**RMS**	**MAE**	**RMS**	**MAE**
2500	0.0842	10.80	5.67	13.70	7.18	18.20	8.21
2550 (2500 +2%)	0.0804	9.48	5.29	13.10	7.03	18.40	8.25
3000	0.0823	10.20	5.32	13.30	7.00	17.89	8.10
3060 (3000 +2%)	0.0789	8.98	4.96	13.00	6.90	18.17	8.13
3500	0.0812	9.76	5.09	13.30	6.93	17.98	8.02
3570 (3500 +2%)	**0.0778**	**8.36**	**4.66**	**12.80**	**6.80**	**17.57**	**7.93**

Results clearly indicate that the addition of knots to a total of 3570 improves the performance of the model resulting in lower cross validated and surface residuals for the set of all stations and for the high quality stations. Further testing investigated if additional improvements could be achieved by increasing the knots by 5% rather than 2%. However it was found that this did not add any value with little improvement in CV statistics.

3.2. Appropriate transformations for index rainfall

ANUSPLIN offers a choice of transformations for the dependent variable. Three options include a natural logarithm transformation, a square root transformation or no transformation to be applied to the data. The log transformation applies the largest data smoothing and no transformation applies the least data smoothing in the thin plate smoothing spline fits. All three transformation options were tested with the addition of 2% of the number of knots as in the above. The square root and no transformation were expected to result in less data smoothing and to improve the fit in areas with high topographic relief. Results are presented in Table 2.

The 3570 knots with no data transformation provides the best performance in the surface residuals across Australia, with the log and square root producing slightly lower error statistics for the cross validated dataset. With a larger number of knots the square root and no transformation performed better than the log transformation.

As the 3570 knot set with square root and no transformation options came out as the best strategies for further consideration, these options were compared at a regional scale. Results are shown in Table 3 for the 1 day duration. Error statistics (mean error (ME), mean absolute error (MAE) and root mean square error (RMSE)) for each of the index rainfall surfaces have been calculated where there were more than three rainfall stations. Results clearly indicate the 3570 knot set with no transformation gave the lowest residuals in the majority of river regions and rainfall districts.

Table 2 Summary statistics of different transformations for index rainfall (1-7 day duration)

Knots and Transformation	All Stations	Surface Residuals		CV Residuals		HQ CV Residuals	
	GCV	RMS	MAE	RMS	MAE	RMS	MAE
2500 - log	0.0842	10.8	5.67	13.7	7.18	18.20	8.21
2550 - log	0.0804	9.48	5.29	13.1	7.03	18.40	8.25
3000 - log	0.0823	10.2	5.32	13.3	7	17.89	8.10
3000 - sqrt	0.4740	10.1	5.29	13.3	7.02	17.78	8.10
3060 - log	0.0789	8.98	4.96	13	6.9	18.17	8.13
3060 - sqrt	0.4320	8	4.68	12.7	6.88	17.50	8.15
3500 - log	0.0812	9.76	5.09	13.3	6.93	17.98	8.02
3500 - sqrt	0.4690	9.68	5.07	13.3	6.96	17.87	8.01
3500 - none	13.1000	9.69	5.09	13.4	7.02	17.82	8.04
3570 - log	0.0778	8.36	4.66	12.8	**6.8**	17.57	**7.93**
3570 - sqrt	0.4280	7.61	4.42	**12.6**	**6.8**	**17.32**	7.96
3570 - none	10.8000	**7.03**	**4.16**	12.8	6.99	17.64	8.25

Table 3 Summary 1 day residual statistics for selected river regions and rainfall districts. Smallest residuals for each rainfall and river region are highlighted in blue.

Reg No	Region Name	No. Sites	Sqrt 3570			None 3570		
			ME	MAE	RMSE	ME	MAE	RMSE
River:29	Tweed River	16	-6.26	11.67	13.53	-4.92	10.05	12.05
37	Barron River	18	2.19	9.88	12.33	2.07	6.89	8.11
44	Mulgrave-Russell R	14	-4.34	14.20	17.36	-2.28	12.18	15.93
52	Hawkesbury River	202	0.03	4.90	6.97	0.04	5.25	7.11
60	Sydney Coast-Georges R	107	-0.10	6.44	8.27	0.19	6.54	8.29
64	Wollongong Coast	35	-1.02	8.33	10.97	-0.99	7.60	10.13
71	Finniss River	18	-0.48	6.75	8.62	-0.29	6.31	8.31
83	Tamar River	82	-0.11	3.33	5.14	-0.01	3.18	4.81
97	Diamantina-Georgina R	139	0.02	1.48	2.14	0.02	1.25	1.93
161	Clarence River	81	-0.11	5.37	7.14	0.20	5.08	6.68
163	Torrens River	57	0.23	2.43	3.04	0.25	2.44	3.05
164	Derwent River	79	0.29	3.64	4.60	0.27	3.55	4.62
172	Pine River	26	-0.44	5.56	6.66	-0.40	5.80	6.81
173	Brisbane River	132	-0.05	4.63	5.97	0.11	4.49	5.70
175	Logan-Albert Rivers	48	1.19	4.63	5.56	1.01	4.25	5.13
176	South Coast	19	-3.14	14.50	19.82	-1.48	13.56	17.64
177	Normanby River	8	1.08	4.81	5.63	0.92	4.68	5.47
178	Bellinger River	21	-1.30	6.43	8.29	-1.02	6.12	7.73
Rain: 5	Fortesque	49	0.65	3.72	5.27	0.73	3.49	5.19
6	West Gascoyne	63	0.07	1.46	1.78	0.05	1.25	1.58
9	Central Coast	284	-0.45	2.29	3.11	-0.44	2.26	3.07
31	Barron	79	-0.44	10.28	13.38	0.17	8.53	11.31
34	West Central Coast	41	0.26	3.84	4.88	0.09	3.70	4.70
66	Metropolitan (E)	80	-0.60	6.33	7.91	-0.46	6.41	7.96
67	Metropolitan (W)	36	-0.47	4.66	6.75	-0.26	5.58	7.41
70	Southern Tablelands	146	-0.01	3.13	4.17	-0.04	3.15	4.15
86	East Central	139	-0.89	2.77	4.04	-0.83	2.70	3.92
87	West Central	77	-0.03	2.15	2.80	-0.06	2.12	2.78
96	Central Plateau	17	0.40	4.03	5.01	0.40	3.96	4.96

The Tweed River catchment includes the 2 largest residuals (site 40192 Springbrook Forestry and site 40550 Numinbah) (Figure 3). The map highlights large residuals overlaid on topography. The contours represent the 1 day index rainfall surface for the 2550 knot log transformation, and 3570 knot no transformation for comparison of performance in complex areas (Figure 3). The site labels are displayed in black and sized according to the length of the AMS for each site (small = <50yrs, medium = 50 – 80 yrs, large = >80yrs). The relative errors are coloured and sized according to the scale displayed below the maps. The 2 largest residuals from the previous study are circled. They are most noticeable in the 2550 knot log case and is significantly reduced in the 3570 knot no transformation case (Figure 3). With fewer knots and more data smoothing the log transform was superior, but as the number of knots increased and the smoothing reduced, the no data transformation became the best option. With a larger number of knots the fitted surfaces are allowed more flexibility to follow various local patterns, and appear to have a better alignment along the escarpment in Northern NSW (Coffs Harbour) area. In areas with a flatter terrain, there is lower variability in the index rainfall values and there is little difference between the different options.

It is clear that the 3570 knot case with no data transformation produced the best results with closer fits to large index rainfall values in complex topography.

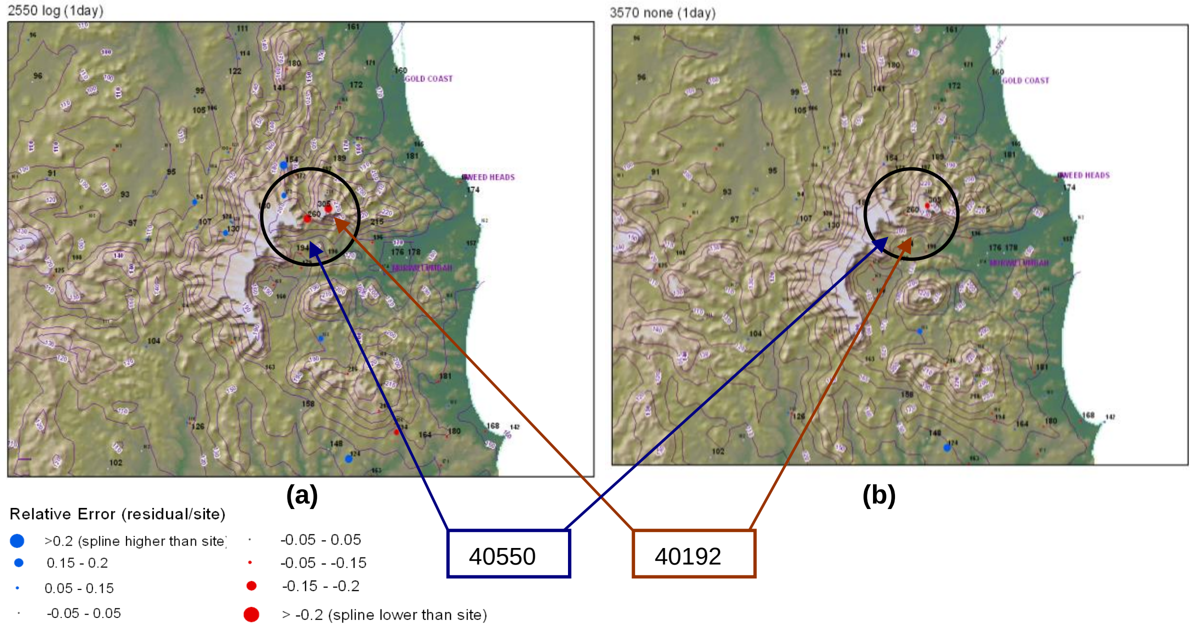

Figure 3 - Gold Coast Area 1 day index rainfall. (a) 2550 knots; log transformation. (b) 3570 knots; no transformation

3.3. Final model optimisation

A final iteration was performed to further optimise the selected knot set with respect to minimising large residuals in particular problem areas. In the Tweed area presented in Figure 3 site 40192, which showed a significant improvement by modelling the area with the no transformation strategy, could be further improved. This site has particularly high rainfall values compared to its neighbours, but is located on the escarpment edge which made it one of the more difficult areas to model.

First the effects of adding extra knots were tested by increasing the number of knots by another 1%. This led to a large increase in the model signal, and an increase in overall predictive error, indicating that the data were being fitted too closely because the additional knots were sensitive to short range correlated errors. A further increase in the number of knots is not recommended.

The second test increased the number of knots in two stages rather than the single step reported in Section 3.1; (i.e. modelling initially 3500 knots, then 3535 knots and 3570 knots compared to using a single jump from 3500 knots to 3570 knots). This approach is recommended by Hutchinson and Xu (2013) for larger datasets. The addition of knots in two stages is expected to assist in picking up sites with large residuals (e.g. site 40192), without adding more knots (comparison of results is given in Table 4).

Table 4 Summary statistics for optimising the number of knots. Comparison of the addition of knots in one versus two stages for daily durations.

All Stations Total 3570 knots	All Stations	Surface Residuals		CV Residuals		HQ CV Residuals	
No. of Knots	**GCV**	**RMS**	**MAE**	**RMS**	**MAE**	**RMS**	**MAE**
3500 +2%	10.8	7.03	4.16	***12.8***	***6.99***	***17.64***	***8.25***
3500+1%+1%	***10.6***	***6.74***	***4.05***	***12.8***	7.07	18.22	8.59

In comparing the 3570 (3500+2%) case to 3570 (3500+1%+1%) case, there is a decrease in the GCV, a slight increase in the cross validated MAE while the cross validated RMS remains the same. The MAE and RMS surface residuals also both improved. Although the reduction in GCV is marginal (the GCV reduced by 2%), and the surface residuals reduced by around 3% to 4%, there is still an overall improvement following this strategy. Comparison of the two strategies, indicate the addition of knots in 2 stages has reduced the residuals at site 40192 by 20-30% and at its closest neighbour site 40607 by 40-60% (Table 5). Maps were also created to visualise the effect of the application of the different strategies. It is clear the addition of knots in 2 stages improves the results most significantly for the daily duration, improving the fit of the spline to the site values on a national and regional scale at the rainfall districts and river regions assessed earlier.

Based on these results it is recommended that a total of 3570 knots increased in two stages be used with no transformation applied to the data. This strategy improves the performance of the model, particularly at stations in isolated elevated regions that were previously difficult to fit.

Table 5 Comparison of site values to the ANUSPLIN strategy of addition of knots in one and two stages.

No. knots	1 day	2 day	3 day	4 day	5 day	6 day	7 day
Site - 40192							
Site Value	304.9	437.4	507.3	546.1	573.8	600.1	617.8
3570 - (3500 +2%)	255.8	359.2	416.2	450.3	480	502.6	517.4
3570- (3500 +1%+1%)	264.8	376.5	439.9	477.1	506.2	530.8	547.7
Site - 40607							
Site Value	210.4	279.2	315.2	337.9	366.9	383.1	391.4
3570 - (3500 +2%)	233.5	324.7	372.5	399.2	422.6	441.7	453.8
3570 - (3500 +1%+1%)	223.6	304.3	344.1	366.8	391.5	408.3	417.5

3.4. Impact of elevation as a predictor of index rainfall

Rainfall is known to have a significant level of topographic dependence. This relationship has been useful in interpolating gridded rainfall data, where the inclusion of elevation as a predictor of index rainfall for the IFD grids led a significant reduction in error statistics by as much as 10% (The *et al.*, 2012). Hutchinson (1998,1995) has recommended an elevation exaggeration factor of 100 relative to the scaling of the horizontal co-ordinates. This accords with the known relative distance scales of atmospheric dynamics (Daley, 1991). To test that this elevation scaling is appropriate for rainfall extremes, tests were run with the elevation scaling systematically varied. For this exercise the tests were based on 3060 knots (3000+2%) using the square root transformation. This was sufficient to use as a base test as the overall response to elevation would be very similar using the other variants (eg. no transformation, 3570 knots). The scaling factors tested are rounded approximate powers of 2 from 2^3 to 2^17.The current base case is 1000 which is the transformation coefficient for elevation (ie. x/a where a = 1000). Figure 4 shows plots of the RMS validation residual and the MAE for the tri-variate spline fitted to the data. The plots clearly show that the current factor of a = 1000 leads to minima in the GCV, MAE and RMS. This translates to an equivalent vertical scaling factor of 100 as recommended by Hutchinson (1995) when the latitude and longitude are in units of degrees because a = 1000 leads to elevations in units of kilometres.

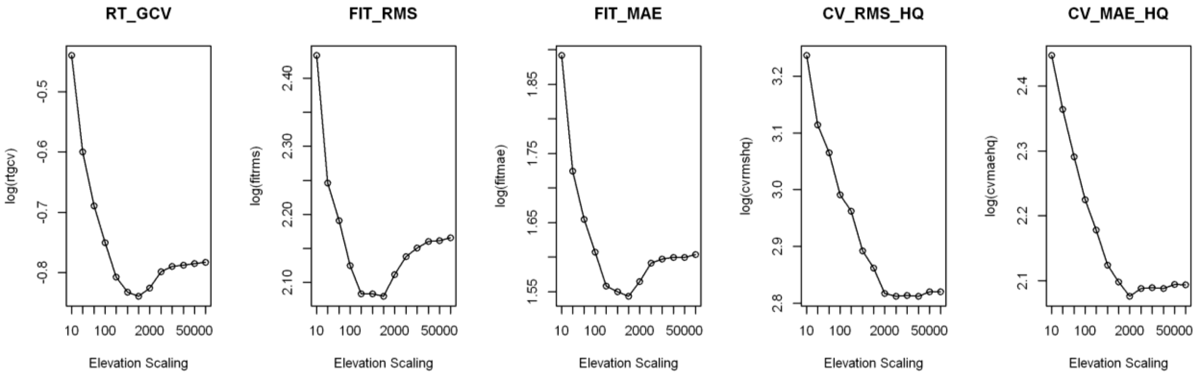

Figure 4 Validation statistics versus scaling of elevation in powers of 2.

4. CONCLUSIONS

The investigations in ANUSPLIN reported in this paper have been used to recommend a strategy to provide grids of rainfall parameters for the IFD Revision Project. It is recommended that the 0.025 degree DEM resolution should be adopted using 3570 knots. The final number should be obtained by increasing the number of knots in 2 stages (ie. 3500+1%+1% = 3570). This has been shown to reduce the GCV and surface residuals across Australia for both the daily and subdaily analysis and also gave the best performance on a regional scale with a clear improvement evident in a number of catchments. The revised spline analysis has improved the fitted surfaces with a significant reduction in predictive error and in residuals in areas with complex topography. No transformation for the index rainfall using the 3570 knot set also yielded an improvement in the fitted gradients on the Queensland coast. This appears to be more consistent with known atmospheric processes than the other options tested. In accordance with previous studies, elevation scaled by a factor of 100 was appropriate to reflect the important influence elevation has on precipitation patterns and optimised the performance of the elevation dependant thin plate spline analyses. Adopting these parameters will allow for improved definition in the final IFD 2013 grids particularly in urban areas where practitioners are likely to need IFD's derived for small catchments.

5. REFERENCES

Daley, R. (1991), *Atmospheric Data Analysis*, Cambridge University Press, 457 pp.

Hutchinson, M. F. (1995), *Interpolating mean rainfall using thin plate smoothing splines,* International Journal of Geographical Information Systems 9(4): 385-403.

Hutchinson, M.F. (1998), *Interpolation of rainfall data with thin plate smoothing splines – part 2: Analysis of topographic dependence*, Journal of Geographic Information and Decision Analysis 2(2): 152-167.

Hutchinson, M.F. (2007), *ANUSPLIN Version 4.37 User Guide*, The Australian National University. Centre for Resources and Environmental Studies, Canberra, Australia.

Hutchinson, M.F and Xu, T (2013), *ANUSPLIN Version 4.4 User Guide*, The Australian National University. Fenner School of Environment and Society, Canberra, Australia

The, C, Johnson, F, Hutchinson, M and Green, J. (2012) *Gridding of Design Rainfall Parameters for the IFD Revision Project for Australia*, Hydrology and Water Resources Symposium 2012, Sydney, Australia.