Proposed method for calculating area-averages for NCMPs

January 28, 2014

1 Intro

The proposed method for calculating area averages of precipitation and temperature anomalies and counts of warm days and cold nights is Ordinary Kriging (it should also work for SPI, although this was not tested). Ordinary Kriging is simple to implement. It takes a set of station data as an input and provides an estimate of the geophysical field (such as temperature) at any point. It can be used to create maps of various geophysical variables and area averages can be calculated from these. The method is widely used in geophysics and it has a number of agreeable properties. It is demonstrably a BLU estimator - that is, Best, Linear and Unbiased - as long as the assumptions made in its derivation hold. Even if the assumptions aren't quite right, it gives a reasonable estimate in many cases. It is easy to implement and code and each of the steps is relatively easy to learn and optimise. It performs better than other commonly-used methods such as simple gridding, near-neighbour interpolation and inverse distance weighting.

I have tested the method for temperature, precipitation, warm days and cold nights using station data from the United States, Australia, the United Kingdom and some other countries in Europe. Where possible, I have compared to available monitoring products as a validation of the method. This choice was not systematic; it was based on data that were readily accessible.

When Kriging, there are two principal steps: calculating the variogram and then kriging. The variogram gives an idea of the spatial correlations in the field being assessed. The kriging step then uses that variogram to interpolate the fields.

2 The Variogram

We have a set of station data at locations, x_i, where we have measured a particular geophysical variable (such as temperature) $Z(x_i)$. To calculate a variogram we first calculate the dissimilarity, γ^*, between all pairs of points This is simply half the squared difference of the temperatures (or whatever) at that point:

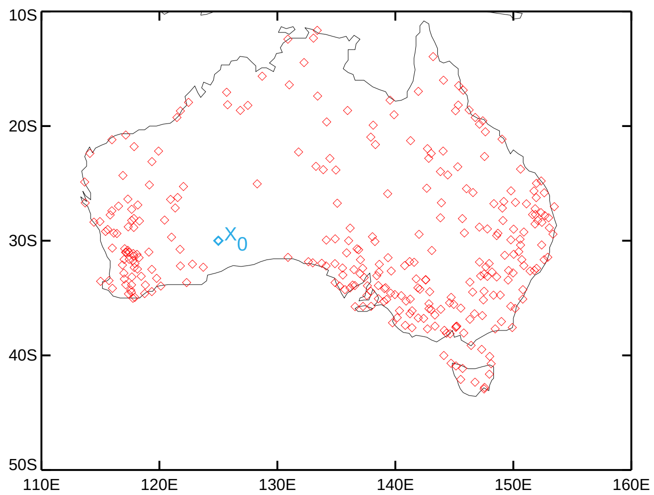

Figure 1: Station locations x_i

$$\gamma_{ij}^{*} = \frac{\left(Z\left(x_i\right) - Z\left(x_j\right)\right)^2}{2}$$

We can calculate this for all pairs of points and then graph them as a function of the distance between the points (dots in Figure 2a). This relationship is typically noisy, but if we divide the distance axis into bins and average the dissimilarity in each bin, a relationship can be seen (solid line in Figure 2a). This relationship between the bin-averaged dissimilarity and distance is the empirical variogram.

The next step is to find a theoretical function that neatly (or nearly) fits the empirical variogram. Typically these have a portion that rises smoothly and then there is a plateau. Three functional forms are widely used and fit well to the majority of data tested. These are the spherical, exponential and gaussian functions. Here h is just the distance between two points (on earth this is usually the great-circle distance) and a, b, and c are three parameters of the function which we can tune to get it to fit the empirical variogram. These are the functional forms:

$$\text{Spherical}: \gamma(h) = b\left(\frac{3h}{2a} - \frac{h^3}{2a^3}\right) + c \text{ for } h < a \; and \, b + c \; for \;\; h > a$$

$$\text{Exponential}: \;\; \gamma(\text{h}) = b\left(1 - \exp\left(-\frac{h}{a}\right)\right) + c$$

$$\text{Gaussian}: \;\; \gamma(\text{h}) = \text{b}\left(1 - \exp\left(-\frac{\text{h}^2}{\text{a}}\right)\right) + c$$

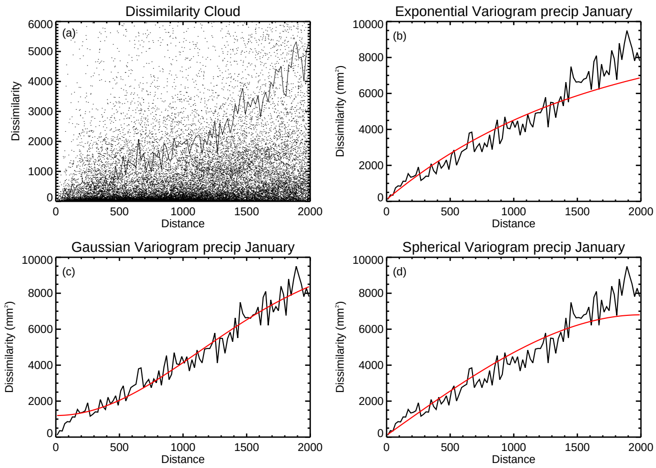

Figure 2: (a) Dissimilarity for all pairs of monthly precipitation measurements for Australia in January 1961 (points) and empirical variogram (solid line). (b)-(d) empirical variogram (black line) for Januarys betweem 1961 and 1990 and, in red, the best fit functional variogram for an exponential function (b), Gaussian function (c) and Spherical function (d).

Figure 2 b-d show the best fit of each of these functions to the same variogram. a is an effective length scale and b is the dissimilarity at large distances, the height of the plateau. The parameter c is the y-intercept of the variogram. There are various ways to interpret this. As it measures the dissimilarity at zero separation, one interpretation is that this naturally represents measurement error although there are versions of Kriging that explicitly include a measurement error term.

3 Kriging

We have the same set of station data as before. At locations, x_i, we measure a geophysical variable (such as temperature) $Z(x_i)$. We want to estimate what the geophysical variable is at some other point x_0. i.e. We want to know $Z(x_0)$ and we can estimate this as a linear weighted sum of the measurements $Z(x_i)$.

$$Z_k(x_0) = \sum w_i Z(x_i)$$

Ordinary Kriging provides the values of the weights, w_i. The weights are calculated based on the variogram. If we have a variogram function$\gamma(x_i - x_j)$ which tells us what the average dissimilarity is as a function of distance between any two locations then we can calculate the weights from this relationship

$$
\begin{pmatrix}
\gamma(x_1 - x_1) & \dots & \gamma(x_1 - x_n) & 1 \\
\vdots & \ddots & \vdots & \vdots \\
\gamma(x_n - x_1) & \cdots & \gamma(x_n - x_n) & 1 \\
1 & \cdots & 1 & 0
\end{pmatrix}
\begin{pmatrix}
w_1 \\ \vdots \\ w_n \\ \mu_{OK}
\end{pmatrix}
=
\begin{pmatrix}
\gamma(x_1 - x_0) \\ \vdots \\ \gamma(x_1 - x_0) \\ 1
\end{pmatrix}
$$

or, more succinctly,

$$AW = B$$

The 1's and 0's in the matrix ensure that the weights sum to one. The weights can then be estimated by calculating

$$W = A^{-1}B$$

Which can be rather a slow process if there are a lot of stations because it involves calculating the inverse of A. The inverse of A is a matrix which encodes the relationships between the stations. B is vector that encodes the relationships between each of the stations and the point at which we are trying to estimate the geophysical variable.

4 Calculating area-averages

Once we have the regular lat-lon grids of the variable we are interested in, the calculation of area-averages is simple. For these prototypes, I selected gridboxes that wholly or partly lay within the country and weighted them by their area (proportional to the cosine of the latitude). Because the grid boxes are rectangular, this method incoporates areas outside the country of interest. I tried to minimise this effect by ensuring that the grid boxes were sufficiently small (an iterative process in some cases). A more sophisticated approach would be to estimate what fraction of each grid box fell within the country, but I can't as yet think of a way to automate that process.

5 Implementation

Kriging is usually performed on a single field, but with climatic data we often have information from multiple fields. This makes estimating a variogram easier, because we can draw on information from more than one month to get a more reliable estimate of the empirical variogram. I calculated a separate variogram for each calendar month (jan, feb, mar...) because variability changes with the time of year. To calculate a variogram for use in January, I calculated an empirical variogram separately for each January in the climatology period. I then averaged the 30 empirical variograms to get a single empirical variogram for January.

There are a number of ways to calculate the great-circle distances between stations, h, that are needed to calculate the variogram. I found that the most robust method was to use the formula

$$h = r.\arctan\left(\frac{\sqrt{(\cos\phi_2\sin\Delta\lambda)^2 + (\cos\phi_1\sin\phi_2 - \sin\phi_1\cos\phi_2\cos\Delta\lambda)^2}}{\sin\phi_1\sin\phi_2 + \cos\phi_1\cos\phi_2\cos\Delta\lambda}\right)$$

where $\Delta\lambda$ is the absolute difference between the station longitudes and ϕ_1 and $\phi_2$are the station latitudes. r is the radius of the Earth, which is approximately 6371.009 km. Care needs to be taken that arctan returns the correct value. If the bottom half of the fraction is negative then some implementations of arctan will give the wrong value. Some programming languages have a function atan2 that does this correctly, others allow atan to receive either one argument or two and the latter allows the passing of the numerator and denominator separately.

To find the functional form and parameters that best fit the empirical variogram, I separately fit each of the three functional forms described above (spherical, gaussian, exponential) by finding the values of a, b and c that minimised the mean absolute difference between the empirical variogram and the variogram function. I then choose the variogram function that minimised the mean absolute difference overall. Sometimes the empirical variogram rises to some maximum value and then falls again. This can happen when the variability within the region is not constant, or when the distance approaches the size of the domain being analysed (i.e there are fewer pairs of points at the greatest separations). For interpolation this is not a great problem because the kriged estimate is relatively insensitive to this choice. On the other hand, if one were estimating uncertainties for the kriged fields, the problem would be more severe because the uncertainty in a sparsely observed region depends on the estimated variance. One option is to set a maximum distance for which the empirical variogram is calculated. As defaults, I chose 3000 km for temperature anomalies and 2000 kmfor precipitation anomalies, but for small regions I used smaller separations. e.g. when I looked at the Netherlands, I chose a maximum separation of 500 km as a first guess and even then I had to use information from neighbouring countries.

To krige the data I first selected the variogram function and parameters that are appropriate for the month. I then select those stations in the month that have data. I calculate A for all the pairs of stations. I then calculate the inverse of this. This is the most time consuming step, particularly if there are many stations.

I then prepare a regular latitude, longitude grid which has a sufficiently high resolution that the area of interest (this will be a country and its surroundings for NCMPs) contains several hundred to a few thousand grid points. For each grid point I then compute B. I then calculate the weights w_i from $A^{-1}B$ and the kriged value at that grid point from

$$Z_k(x_0) = \sum w_i Z(x_i)$$

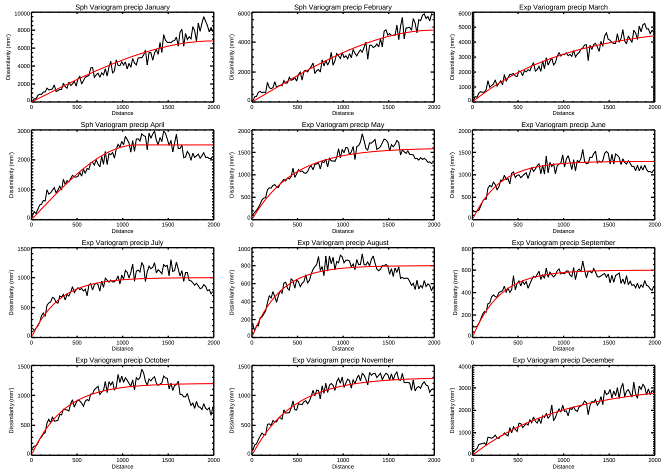

Figure 3: Empirical (black) and functional (red) variograms for Austrlian monthly total precipitation (1961-1990) for each calendar month.

Because we are taking the inverse of a matrix, there are a number of things to be wary of. First, if two of the columns of the matrix are identical, it is not possible to calculate the inverse. This can happen if two stations have the same latitude and longitude. One can avoid this by removing duplicate stations, or by moving stations slightly so that they are in different locations. Second, there can also be difficulties when stations are very close together because of numerical instabilities in the computation of the inverse. This is exacerbated by the shape of the gaussian function which is relatively flat for small separations. This can lead to columns being almost identical and the calculation of the inverse can fail (spectacularly, beautifully). One way to fix this is to ensure that the parameter c is not zero. In practice it rarely is, but it can help to constrain c to be larger than some minimum value.

6 Results

6.1 Australia

I show first results for Austrlian precipitation data. I used data taken from the Bureau of Meteorology web site which has monthly precipitation totals for 307 stations whose locations are shown in Figure 1. Figure 2 shows the creation of the variogram for January. In this case the Gaussian function best fit the variogram. Figure 3 shows variograms for each calendar month and their best fit functions. In some cases the empirical variogram turns over at large distances and starts to fall.

The monthly and annual Australian average total precipitation is shown in Figure 4. The annual total anomaly is compared to the official Bureau of

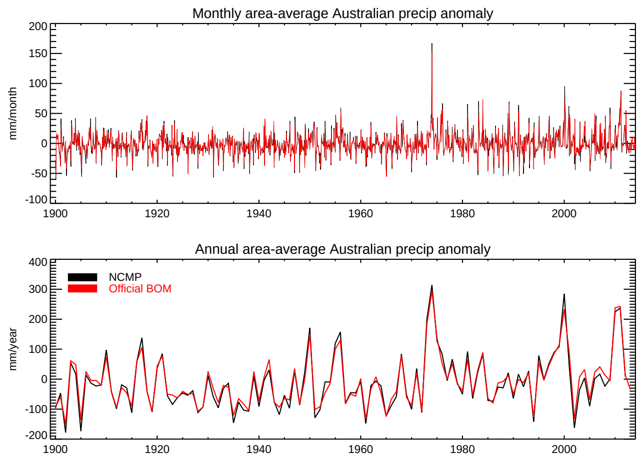

Figure 4: Monthly (top) and annual (bottom) total precipitation anomalies for Australia. The black lines are the kriged estimates and the red line is the official Bureau of Meteorology estimate.

Meteorology annual total time series. The large scale features and much of the year to year variability is captured by the simple Kriged estimate and much of the month to month variability. The standard deviation of the differences between the two estimates is 4.5 mm whereas the total variability is around 17 mm. Figure 5 shows a precipitation map for a single month and an official Bureau of Meteorology map for the same month. There are differences betweent the two, but the broad scale features are captured well by kriging.

Daily max and min temperatures were taken from the ACORN data set which contains 112 stations. The max and mins were used to calculate daily mean temperatures and monthly mean temperature anomalies. Monthly and

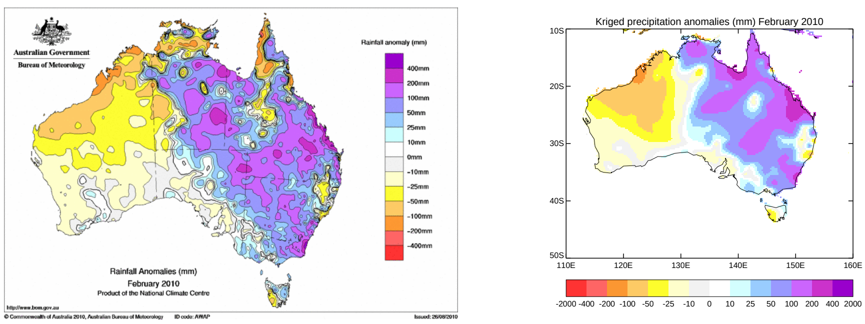

Figure 5: February 2010 precipitation anomalies from (left) Bureau of Meteorology and (right) kriged estimate using freely available stations.

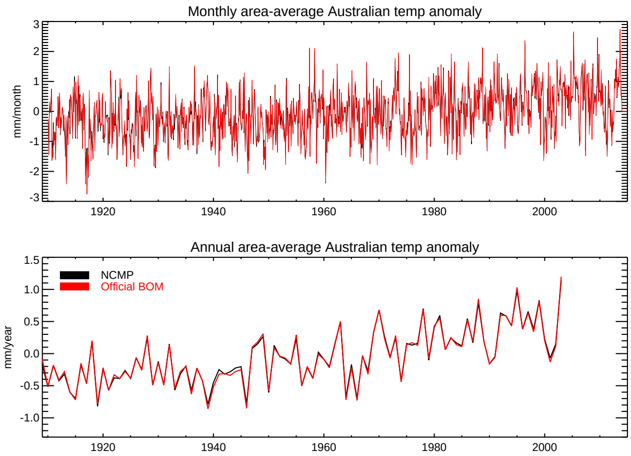

Figure 6: Monthly (top) and annual (bottom) average mean temperature anomalies for Australia. The kriged estimate is shown in black and the official Bureau of Meteorology estimate is shown in red.

annual mean temperature anomaly (NCMP 1) series are shown in Figure 6. A map for a single month is shown in Figure 7. Once again, the broad scale features are picked up by the kriged estimate, but some of the detail is lost. Nonetheless, the kriged estimate is usually within 0.1°C of the official value (standard deviation of the difference is $0.09°\ C$).

Figure 8 shows the counts of warm days and cold nights for each year and month averaged across Australia. The data were taken from the ETCCDI website.

6.2 United States

I downloaded data for several hundred USHCN stations from the NCDC website. These contain monthly mean temperatures and precipitation totals. Station coverage over the US is very high, so the calculation took a bit longer than for Australia. Figures 9 and 10 show the kriged estimates of the US average temperature and total precipitation anomalies. The differences between the Kriged estimate and the official temperature series have a standard deviation of 0.08°C. For the precipitation series the standard deviation is 1.8 mm. Figure 11 shows the annual and monthly counts of warm days and cold nights for the US (again taken from the ETCCDI web site). A map of temperature anomalies is shown in Figure 12.

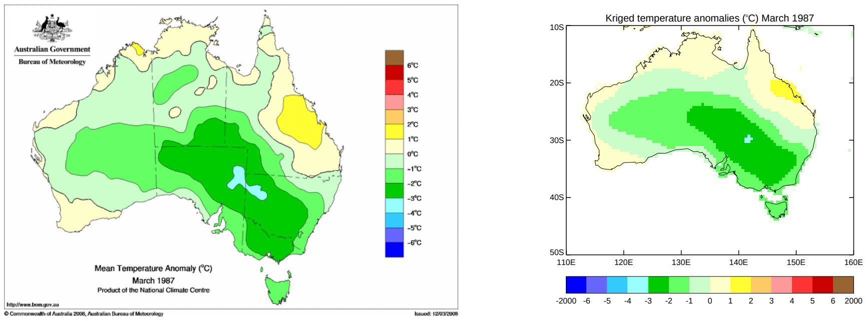

Figure 7: Maps of monthly mean temperature anomalies for March 1987 for Australia. (left) official Bureau of Meteorology map. (right) kriged estimate using freely available ACORN stations.

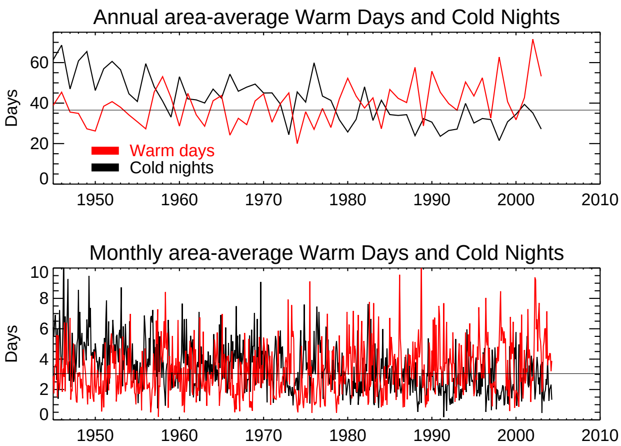

Figure 8: (top) annual and (bottom) monthly counts of warm days (Tmax > 90th percentile) and cold nights (Tmin < 10th percentile) avergae over Australia. The horizontal line indicates approximately 10% of days. During the climatology period the average numbers of warm days and cold nights should be around this line.

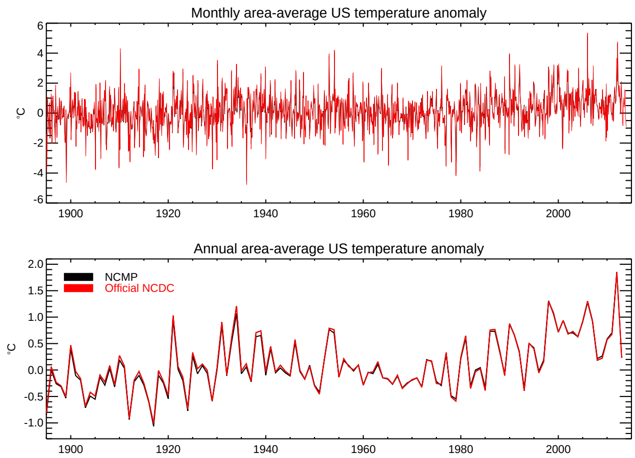

Figure 9: (top) monthly and (bottom) annual average temperature anomalies for the US. The black line shows the kriged estimate based on USHCN stations and the red lines shows the official NCDC estimate.

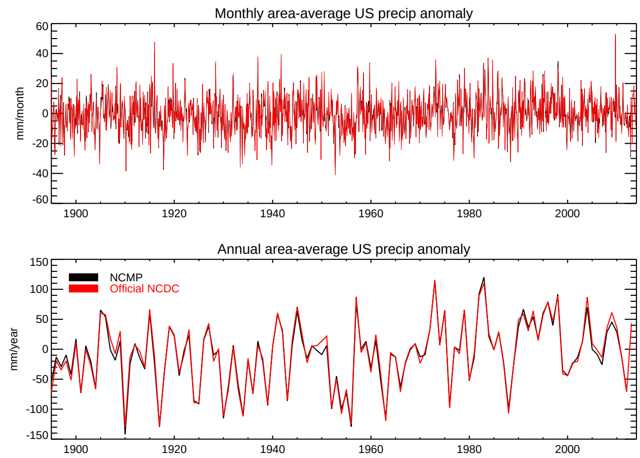

Figure 10: (top) monthly and (bottom) annual total precipitation anomalies (mm) for the US. The kriged estimate is shown in black and the official NCDC estimate is shown in red.

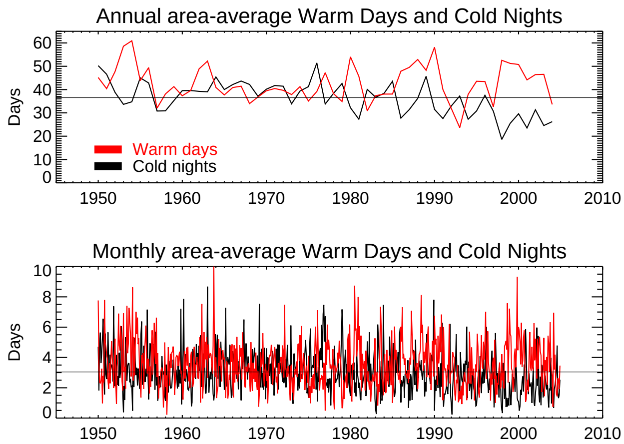

Figure 11: (top) annual and (bottom) monthly counts of warm days (Tmax > 90th percentile) and cold nights (Tmin < 10th percentile) avergae over Australia. The horizontal line indicates approximately 10% of days. During the climatology period the average numbers of warm days and cold nights should be around this line.

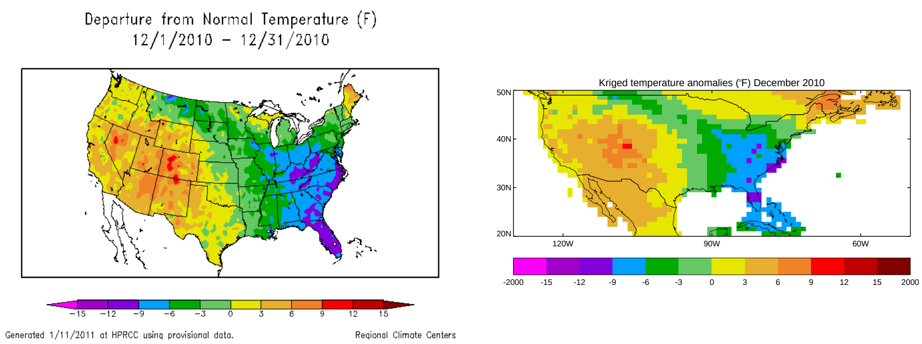

Figure 12: Monthly temperature anomalies for the US. The official US HPRCC product is shown on the left and the kriged estimate is shown on the right. Note that the climatology periods are different on the left and right.

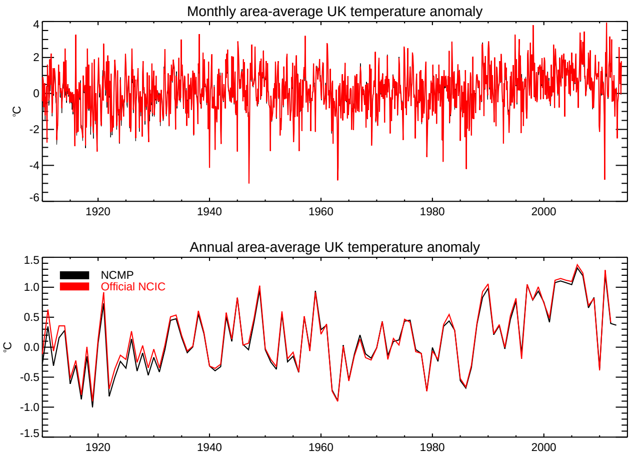

Figure 13: (top) monthly and (bottom) annual average temperature anomalies for the UK. The black line shows the kriged estimate based on public stations and the red lines shows the official NCIC estimate.

6.3 United Kingdom

I downloaded data for 38 stations from the Met Office website. These stations contain long term records of monthly mean temperature and precipitation. Monthly and annual mean temperature and total precipitation anomalies (expressed as percentages of normal) are shown in Figures 13 and 14. Although temperatures are reasonably well captured, the sparse network means that annual and monthly temperature and (particularly) precipitation averages are somewhat less well captured than they were for the US or Australia where the available network was denser. The UK does have more than 38 stations - the official figures are based on data from 100s of stations - but these are the stations that I had to hand.

7 Discussion

Implementing the NCMPs from existing indices was relatively simple. The kriging method worked well for the three example countries presented here. These were chosen because the data were easily accessible. I have gathered other data sets which I hope to try out in the near future, but I would also welcome suggestions of other data sets to try. For the countries I looked at, station density was relatively high. This means it was possible to estimate the variograms from the data although the UK with only 38 stations was close to being inadequate for this purpose. For someone producing NCMPs for a country with only a small number of stations, estimating the variogram would be either difficult, or impossible. One possible solution to this would be to provide default

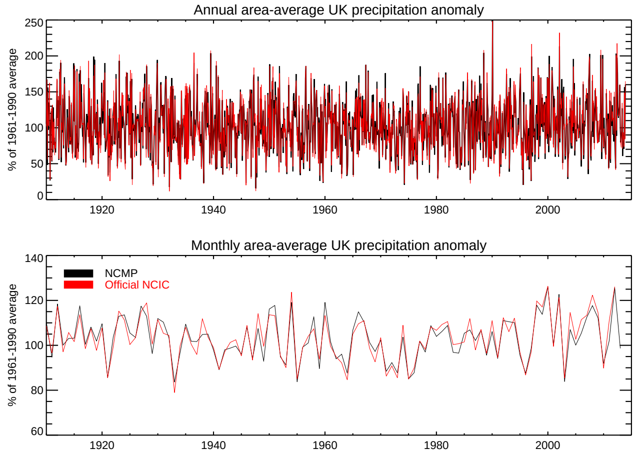

Figure 14: (top) monthly and (bottom) annual total precipitation anomalies (expressed as a percentage of 1961-1990 average) for the UK. The kriged estimate is shown in black and the official NCIC estimate is shown in red.

variograms or use data from neighbouring countries to help with the estimation.

Some of the NCMPs have properties that make them more flexible as monitoring tools. For example the monthly mean temperature can be averaged in time to give seasonal and annual means. The monthly counts of warm days and cold nights can be summed to give seasonal and annual totals.

7.1 Producing and updating NCMPs

Having gone through this process for four variables for a number of different countries, it's worth reflecting on what would be involved to produce and update the NCMPs. Once the code was written to do the Kriging, the bulk of the work was in organising the input data - finding it, checking for missing data, getting it into an appropriate format.

The calculation of the variograms involves a bit of judgement. It is normally necessary to first calculate the empirical variograms and visually inspect them. The aim is to assess whether there are sufficient stations to estimate the variogram, what the maximum reasonable distance for the variogram is and what the best functional forms are. Sometimes none of the functional forms listed above gives a perfect fit. In fact, the greater the station density, the smoother the empirical variogram will be. Consequently, when station density is great, none of the functional forms will give a perfect fit.

Once the data have been ingested and the variograms and parameters have been chosen, Kriging is simply a matter of running the code. I visually inspected the output fields and time series to check for outliers and other peculiarities. Sometimes missing data flags get mistyped or mangled in the input data and appear in the output fields. In a few cases it was necessary to go back to the variograms and recompute.