Coverage bias in the HadCRUT4 temperature series and its impact on recent temperature trends.
UPDATE
Temperature reconstruction by domain: version 2.0 temperature series

Kevin Cowtan, Robert G. Way

January 6th, 2014

1 Temperature reconstruction by domain: version 2.0 temperature series

This update document describes a new version of the global temperature reconstruction introduced by Cowtan and Way (2014) (henceforth CW14). The principal difference between the version 2.0 temperature reconstructions reported here and the versions in CW14 is that the land and sea surface temperatures (SSTs) are now reconstructed separately. Land and sea temperatures show different levels of agreement with the satellite temperatures and required different kriging ranges. Sea ice represents a third category: however theory, the IABP/POLES data (Rigor *et al.*, 2000) and multiple weather models support the determination of air temperatures over sea ice from the land-based station data.

The new temperature series also use the latest versions of the source datasets: HadCRUT version 4.2.0 and UAH version 5.6. The former change yields a marginal increase in 16 year trends in comparison to CW14.

2 The input ensemble

To calculate temperature reconstructions from the unblended data both the HadCRUT4 land and SST ensemble data (Morice *et al.*, 2012) are required. HadSST3 is distributed, but the land data is not. The land ensemble is derived from the publicly available CRUTEM4 data by adding a UHI term, an exposure term and local terms to account for homogenisation errors - these vary from ensemble member to member. The first two are of most interest because they are global and thus influence global temperatures.

The UHI and exposure terms are controlled by two random numbers which determine the size of the adjustment according to distributions derived from the literature. The UHI term is a ramp from 1900, and exposure affects early values, with the scale and timing of the effect varying between latitude bands.

In the HadCRUT4 calculation the corresponding land and SST ensemble members are combined to produce a blended map, weighting the land and SST values according to the land proportion in the grid cell. The blended map data are publicly available. For cells where only a land or SST value is available, the cell is assigned that value. Therefore for any cell where a land value but no SST value are available, the land ensemble cell must equal the blended cell value. The only cells for which values are unknown are the ones for which land and SST values have been blended, since the exact weights are unknown.

Four different approaches to recreating the land ensemble are compared:

1. 'CRUTEM4' approach: Use the unmodified CRUTEM4 data (i.e. without UHI or exposure terms) for all the land ensemble members.

2. 'Mix' approach: Use cell values from the HadCRUT4 ensemble members for any cell for which a land temperature but no SST temperature is available, augmented by the CRUTEM4 cell values for cells containing a blend of land and SST values.

3. 'Fit' approach: Estimate the values of the UHI and exposure terms by ordinary least squares fitting the CRUTEM4 data to the HadCRUT4 data for any cell for which a land temperature but no SST temperature is available. The CRUTEM4 data is then modified using the UHI and exposure terms scaled by the resulting coefficients.

4. 'Mix+Fit' approach: Use cell values from the HadCRUT4 ensemble members for any cell for which a land temperature but no SST temperature is available, augmented by the values from the 'Fit' approach for cells containing a blend of land and SST values.

The 'Mix' reconstruction is very simple, but fails to apply the UHI term to some coastal cells which play a significant role in early years, and so underestimates the correction slightly. The 'Fit' reconstruction treats all cells in the same way although the UHI term is not very well determined and so the spread of UHI values is slightly greater than in the HadCRUT4 ensemble; in addition the resulting ensemble contains no homogenization noise. The 'Mix+Fit' reconstruction combines benefits of both approaches.

3 Uncertainty estimates

CW14 only provide estimates of the uncertainty in the global mean temperature due to coverage, using a modification of the HadCRUT4 method which captures the effect of slowly varying regional trends in the temperature data but loses the annual cycle in the uncertainty. Two improvements have been made to this procedure: the coverage bias calculation now captures some of the annual cycle in the uncertainties, and the additional uncertainty introduced by parameters of the method, represented in HadCRUT4 by an ensemble of 100 temperature realisations, is also determined.

3.1 Coverage uncertainty

The uncertainty due to limited coverage is estimated by testing the skill of the method in reconstructing the NCEP/NCAR reanalysis data (Kalnay *et al.*, 1996) after reducing the coverage of the data to match the corresponding month from the original HadCRUT4 data. A moving average over multiple months of the difference between the global mean of the original NCEP/NCAR temperature field and the masked-and-reconstructed field provides a time varying estimate of the uncertainty due to coverage alone.

To estimate the uncertainty in the mean of the blend of the land and SST reconstructions, the NCEP/NCAR data for a given month are masked twice to match the land and SST data respectively. Global reconstructions are determined for the two masked maps, and these are then blended to provide a global land/ocean reconstruction. The difference between the mean of this map and the mean of the original map is determined for each month in the data.

To turn these differences into an uncertainty, a weighted moving average is used to capture the time variation in the uncertainty, including some of the annual cycle. The weighting function is of the form $w(t) = 1/4(1+\cos(2\pi t/\delta))(1+\cos(2\pi t))$ for $|t| < \delta/2$ and zero otherwise, where $\delta = 20$ years. The first term weights differences according to distance from the month under consideration, and the second downweights differences at different times of year.

For years prior to 1960 the data for the years 1960-1969 are repeated (the period having good observational coverage but little global trend). The NCEP/NCAR data show exaggerated high latitude warming since 2000 and so the uncertainties are more likely to be overestimated than underestimted for this period. Any impact of early 20th century warming on the coverage uncertainty is lost - the 20th century reanalysis (Compo *et al.*, 2011) was investigated as an alternative to NCEP/NCAR for this purpose but polar trends were unrealistic.

3.2 Ensemble uncertainty

The uncertainty due to the parameters of the method is determined in the same way as for HadCRUT4 by producing a blend of the land and SST reconstructions for each of the 100 ensemble members. Temperature series are calculated for each ensemble member, and the median and standard deviation of the temperature among the ensemble members determined for each month. (The use of the standard deviation for this purpose is simpler than the HadCRUT4 approach of providing upper and lower quantiles, but loses any asymmetry of the ensemble distribution.)

3.3 Total uncertainty

The total uncertainty is determined from the square root of the sum of the squares of the coverage and ensemble uncertainties. Coverage uncertainty is reduced in comparison to the HadCRUT4 data, but continues to dominate the total uncertainty.

Extrapolation of unobserved regions inflates the monthly variability of the global mean temperature data - this is an inevitable result of the temperature reconstruction, however the resulting noise is captured in the uncertainty estimates. In practice the process reduces bias at a cost of increased noise. The fact that the uncertainty estimates are smaller than for the unreconstructed data illustrates that there is a net benefit to the process (and the kriging calculation serves to optimise this trade-off). The increased monthly variability in the early record should not however be mistaken for a climate signal.

4 Global blended land/ocean reconstruction

In CW14 the HadCRUT4 map series was rebaselined on the period 1981-2010 for compatibility with the satellite data. This had an additional benefit of reducing the impact of changes in coverage during the period 1981-2010, at a cost of losing some cells with limited observations during the baseline period. In practice this corrected a small bias due to improving SST coverage coupled with declining land coverage leading to a small increase in recent temperature trends.

When producing a pure kriging reconstruction no rebaselining is required, avoiding the loss of coverage. In addition reconstructing global land and ocean fields before blending addresses the changing land and SST coverage, and as a result the method is more robust against different choices of baseline period. This is an additional benefit of the separate land and ocean reconstructions.

The e-folding range of the kriging covariance function, determined from the data for the period 1981-2010 using the method of CW14, is 767km for the land data and 860km for the ocean data.

The resulting land and ocean temperature fields are then blended. A mask is constructed containing the fraction of land plus sea ice in each cell for each month of the year, however the mask remains constant from year to year to avoid a potential bias due to the transformation of cells from ice to open ocean. Sea ice coverage is determined from the HadISST data using the median of the sea ice concentration on the period 1981-2010 for each cell and each month. This underestimates ice coverage for early decades and overestimates ice coverage for recent summers, however CW14 found coverage bias to be least in the summer when the change in ice cover is most significant, so the effect on temperatures is minimal.

Temperature reconstructions were then performed for each of the four approaches to generating the land temperature ensemble from section 2. For each approach, the land temperature ensemble is generated and global land and SST reconstructions were determined by kriging. These are blended and the resulting global temperature series compared to determine the sensitivity of the results to the reconstruction approach. The results are shown in Figure (U1) using a 12 month moving average.

The temperature series for the four approaches are almost identical. The differences between each other approach and the basic CRUTEM4 approach are plotted in Figure (U2). Both the 'Mix' and 'Fit' methods give broadly similar results, differing from the CRUTEM4 reconstruction by the addition of a cooling trend starting in 1900 - this is the UHI correction term. The exposure correction is symmetric and so has limited impact on the ensemble median temperatures.

Given that the resulting global temperature reconstructions are so similar any of these approaches are probably adequate for generating starting data for the determination of global temperatures, however the 'Mix+Fit' reconstruction is favoured because it includes most of the ensemble variability and so gives a better estimate of the ensemble contribution to the uncertainties in the resulting temperature series.

The final temperature reconstruction using the 'Mix+Fit' data is compared to the original HadCRUT4 data in Figure (U3) using a 12 month moving average. For much of the reconstruction the global means are very similar, however recent years show a divergence which is unusual in terms of magnitude and duration.

The assumption in this temperature reconstruction is that unobserved land regions behave like observed land regions and unobserved ocean regions behave like observed ocean regions. Global kriging sets isolated cells to an optimal estimate of the global mean, which is equivalent to leaving them out when calculating a global mean. As a result isolated land regions are set to the global land mean and isolated ocean regions to the global ocean mean.

References

Compo GP, Whitaker JS, Sardeshmukh PD, Matsui N, Allan R, Yin X, Gleason B, Vose R, Rutledge G, Bessemoulin P, BroNnimann S, Brunet M, Crouthamel R, Grant A, Groisman P, Jones P, Kruk M, Kruger A, Marshall G, Maugeri M, Mok H, Nordli O, Ross T, Trigo R, Wang X, Woodruff S, Worley S. 2011. The twentieth century reanalysis project. *Q. J. R. Meteorol. Soc.* **137**(654): 1–28.

Cowtan K, Way RG. 2014. Coverage bias in the hadcrut4 temperature series and its impact on recent temperature trends. *Quarterly Journal of the Royal Meteorological Society* .

Kalnay E, Kanamitsu M, Kistler R, Collins W, Deaven D, Gandin L, Iredell M, Saha S, White G, Woollen J, Y Z, A L, R R. 1996. The ncep/ncar 40-year reanalysis project. *Bull. Amer. Meteorol. Soc.* **77**(3): 437–471.

Morice CP, Kennedy JJ, Rayner NA, Jones PD. 2012. Quantifying uncertainties in global and regional temperature change using an ensemble of observational estimates: The hadcrut4 data set. *J. Geophys. Res.* **117**(D8): D08 101.

Rigor IG, Colony RL, Martin S. 2000. Variations in surface air temperature observations in the arctic, 1979-97. *J. Climate* **13**(5): 896–914.

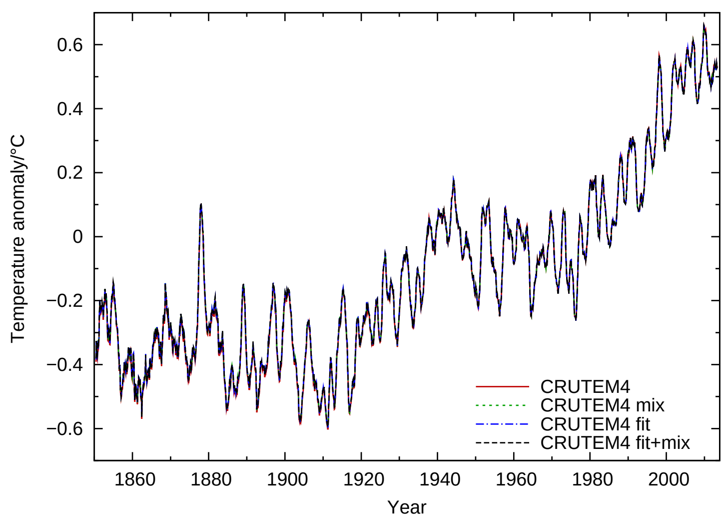

Figure U1: Temperature series determined by blending global reconstructed land and ocean fields using 4 different approaches to reproducing the HadCRUT4 land ensemble.

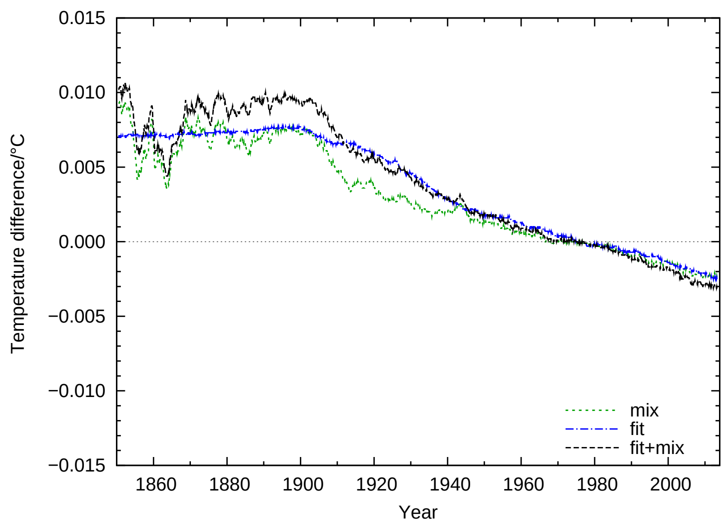

Figure U2: Differences between reconstructions from the 'Mix', 'Fit', and 'Fit+Mix' land data and the CRUTEM4-based reconstruction.

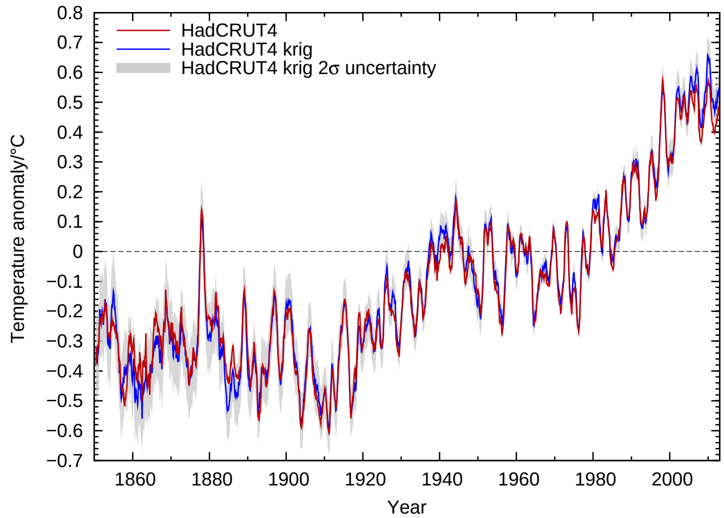

Figure U3: Comparison of the original and kriging HadCRUT4 temperature series using a 12 month moving average. Uncertainties for the kriging reconstruction are shown.