Downscaling humidity with Localized Constructed Analogs (LOCA) over the conterminous United States

D. W. Pierce[1,*] and D. R. Cayan[1,2]

[1]Division of Climate, Atmospheric Sciences, and Physical Oceanography, Scripps Institution of Oceanography, La Jolla, California, 92093-0224

[2]U.S. Geological Survey, La Jolla, California

[*]Correspondence to: David W. Pierce, dpierce@ucsd.edu, 858 534-8276.

Final version

16 September 2015

Climate Dynamics, accepted

Abstract

Humidity is important to climate impacts in hydrology, agriculture, ecology, energy demand, and human health and comfort. Nonetheless humidity is not available in some widely-used archives of statistically downscaled climate projections for the western U.S. In this work the Localized Constructed Analogs (LOCA) statistical downscaling method is used to downscale specific humidity to a 1/16° grid over the conterminous U.S. and the results compared to observations. LOCA reproduces observed monthly climatological values with a mean error of ~0.5% and RMS error of ~2%. Extreme (1-day in 1- and 20-years) maximum values (relevant to human health and energy demand) are within ~5% of observed, while extreme minimum values (relevant to agriculture and wildfire) are within ~15%. The asymmetry between extreme maximum and minimum errors is largely due to residual errors in the bias correction of extreme minimum values. The temporal standard deviations of downscaled daily specific humidity values have a mean error of ~1% and RMS error of ~3%. LOCA increases spatial coherence in the final downscaled field by ~13%, but the downscaled coherence depends on the spatial coherence in the data being downscaled, which is not addressed by bias correction. Temporal correlations between daily, monthly, and annual time series of the original and downscaled data typically yield values > 0.98. LOCA captures the observed correlations between temperature and specific humidity even when the two are downscaled independently.

1 Introduction

The current generation of Global Climate Models (GCMs) archived as part of the Climate Model Intercomparison Project version 5 (CMIP5, Taylor et al. 2012) typically have spatial resolutions on the order of 1 to 2.5 degrees latitude-longitude. Many practical applications of climate data in the areas of hydrology, agriculture, ecology, and energy impacts require finer spatial resolution than afforded by GCMs. To address this need, GCM data are often transformed from the relatively coarse-resolution original GCM grid to a finer resolution (O(10 km)) regional grid using information about topography and the land surface through a process termed downscaling.

Downscaling methods fall into two broad classes. Dynamical methods generate the fine-resolution data using a computational model similar to the GCM itself but implemented over a much more limited domain and with higher spatial resolution (for a recent review of dynamical downscaling, see Rummukainen 2009). Since they use a full computational model, dynamical methods generate a full suite of climate variables and can simulate future climate states that are not dependent on direct guidance from historical observations (although even then, dynamical methods incorporate parameterizations tuned to reproduce historical conditions, and so might yield a less satisfactory simulation of future conditions). The downside is that dynamical methods are computationally expensive and require a large amount of data as boundary conditions. Dynamic models also generally produce data that are biased with respect to observations, so bias correction is required for many climate impact studies, which means that the final product may not, in fact, be independent of observations. Statistical methods infer plausible fine spatial resolution patterns based on historically observed relationships between large-scale and fine-scale information in the climate variable being downscaled (for recent reviews, see Fowler et al. 2007 and Maraun et al. 2010).

Statistical methods assume that certain historically observed statistical relationships between climate variables will persist into the future, but are orders of magnitude faster than dynamical methods (although reproducing spatial and multivariate dependencies can make the computational burden non-trivial).

Constructed analog (CA) techniques (van den Dool, 1994) implement statistical spatial downscaling by identifying a set (typically 30) of historically observed “analog” days that are similar to the GCM day being downscaled, then combining those analog days in different ways (depending on the exact CA method) to produce the final fine-resolution downscaled field. Standard CA (Hidalgo et al. 2008), Bias Correction with Constructed Analogs (BCCA, Maurer et al. 2010), and Multivariate Adaptive Constructed Analogs (MACA, Abatzoglou and Brown, 2012) combine the analog days by performing a weighted average of the analog days to produce the downscaled field. Pierce et al. (2014) show that this averaging introduces some undesirable properties, such as too much spatial coherence, a reduction in extremes, and the production of extraneous “drizzle” days with small amounts of precipitation. MACA (Abatzoglou and Brown, 2012) addresses some of these problems by performing an additional bias correction step after the constructed analogs are averaged together.

Localized constructed analogs (LOCA; Pierce et al. 2014) is a spatial statistical downscaling technique that avoids these problems by treating the analog days in a different way: at each point on the fine-resolution grid (1/16° latitude-longitude here), the single analog day of the 30 that is the best match in the local ~1° latitude-longitude box around the point being downscaled is used as the analog day for that point. LOCA’s multi-scale matching (both in the local region around the point and in the wider region used to select the 30 analog days), and the selection of a single analog at each point instead of 30, avoids some of the issues noted above that arise from averaging multiple analog days to obtain the downscaled field (see Pierce et al. 2014 for details).

This work focusses on downscaling humidity. Humidity variations are important for many climate impacts; for instance, a misrepresented humidity field can degrade future projections of runoff in critical water resource regions (Pierce et al. 2013). Humidity affects plants through the tendency of stomata to close under dry conditions (e.g. Friend 1995) and influences the chance of having wildfires (e.g. Brown et al. 2004). Humidity is an important factor in perceived human comfort as a function of temperature (e.g., Thom 1959), since sweat cannot cool the body as easily in humid conditions. Therefore humidity affects energy use through air conditioning demand (e.g., Mirasgedis et al. 2006). Many of these factors are important in the western U.S., yet a major archive of statistically downscaled climate simulations for the western U.S., hosted at the Lawrence Livermore National Laboratory's "green data oasis" (http://gdo-dcp.ucllnl.org/downscaled_cmip_projections/dcpInterface.html; Maurer et al. 2014), does not include humidity as one of the downscaled variables. To our knowledge only MACA (Abatzoglou and Brown, 2012) provides a publically available archive of statistically downscaled humidity data for the western U.S. One reason MACA is able to do so is because MACA can downscale multiple variables simultaneously, in this case daily temperature and humidity. Such an approach can produce more coherent downscaled fields than found when downscaling the variables separately (Abatzoglou and Brown, 2012).

The purpose of this work is to demonstrate the process of downscaling humidity using LOCA and to evaluate the quality of the resultant downscaled humidity field. LOCA is a spatial downscaling processes, and does not, in and of itself, incorporate bias correction. Accordingly, our primary evaluation is based on downscaling an observed field that has first been coarsened to typical GCM spatial resolution. This avoids the necessity for a bias correction step that is extrinsic to LOCA, and so provides the clearest picture of LOCA's capabilities. However, spatial downscaling is typically applied to GCM data, so we additionally show results from downscaling GCM data that has first been bias corrected. Our evaluation includes examining monthly means, the temporal standard deviation of daily values by season, 1-in-1 and 1-in-20 year extreme values, and measures of spatial coherence. The selection of these validation measures is dictated by the application areas of interest noted above. Monthly means and extreme high humidity values affect energy use, extreme low values affect wildfires, and the spatial coherence of the humidity field can influence the simulation of runoff in geographically adjacent basins, influencing a hydrological simulation of drought and flooding.

The rest of this paper is structured as follows. In section 2 we describe the data and methods used in this study, including a fuller explanation of the LOCA process. Section 3 shows the results, including LOCA downscaling performance for specific humidity using univariate downscaling and the issue of multivariate downscaling. The results are discussed in section 4, including comments on downscaling relative versus specific humidity, and a summary and conclusions are given in section 5.

2 Data and methods

2.1 Observed data

Statistical downscaling schemes such as LOCA require observations to train the model. There are limited options available for observed gridded humidity data sets over the continental U.S. Two products that we are aware of are Daymet (Thornton et al., 2012) and the University of Idaho gridded surface meteorological dataset (Abatzoglou, 2012; http://-nimbus.cos.uidaho.edu/METDATA/). The former supplies vapor pressure, while the latter supplies specific humidity. The CMIP5 GCMs differ in the humidity variable(s) that are archived, but generally speaking many models have near-surface specific humidity, while few supply relative humidity and almost no models supply vapor pressure. Accordingly, we use the University of Idaho specific humidity data for this work. Daily specific humidity values are given on a 4 km spatial grid over the period 1980-2012. For consistency with previous LOCA downscaling work (Pierce et al. 2014) the data were aggregated to the 1/16° latitude-longitude (~6 km) grid used by Livneh et al. (2014).

As described in Abatzoglou (2012), the University of Idaho data set combines data and techniques from the PRISM project (Daly et al. 2008), which provides good spatial coverage, with data from the NLDAS-2 (Mitchell et al. 2004) reanalysis, which provides good temporal coverage. It should be noted that due to the sparsity of station humidity observations, the humidity data should probably be viewed more as a topographically aware interpolation augmented with reanalysis rather than a gridded version of a directly observed, well-sampled field. However the data compare well to the station observations that are available (Abatzoglou 2012), and for our purpose (evaluating LOCA downscaling) minor errors in the training data set are irrelevant, as the downscaling should reproduce in the downscaled data the statistics of whatever observations it is supplied with, correct or not.

In section 3 where specific humidity is related to daily minimum and maximum temperatures, we use the University of Idaho temperature data in preference to other data sources so that the temperature fields are as consistent as possible with the humidity fields.

2.2 Global climate model data

As our example GCM we use data from the Community Climate System Model, version 4 (CCSM4; Gent et al. 2011), produced by the National Center for Atmospheric Research (NCAR) and archived in the CMIP5 data base. Native model resolution is 1.25°x0.94° longitude-latitude, using a finite volume dynamical core, and the model includes an updated version of the Community Land Model (CLM version 4; Lawrence et al. 2011). CCSM4 shows much better El Nino/Southern Oscillation (ENSO) statistics compared to its predecessor, likely due to an improved representation of deep convection (Gent et al. 2011), which is advantageous for our comparison since ENSO affects the climate over parts of the U.S. The model provides daily near-surface specific humidity values for run 6 ("r6i1p1" in CMIP5 jargon), which we used over the period 1950-2005.

2.3 The Localized Constructed Analogs (LOCA) spatial downscaling process

The basic physical assumption of constructed analog spatial downscaling techniques is that meteorological processes produce cyclostationary statistical relationships between area-averaged (0.5-2 degrees latitude/longitude) and point measurements of a climatological field. Global climate model outputs are then considered to be estimates of area-averaged quantities, and the observed relationships between area-averaged fields and point values appropriate to the time of year being downscaled are used to infer a plausible distribution of point values from the model output. In this sense the LOCA spatial downscaling technique in and of itself is a "perfect prog" approach (e.g., Klein and Glahn 1974), although in practice a climate model field is typically bias corrected before being passed to LOCA.

The standard constructed analogs process is conceptually straightforward: to spatially downscale a variable for a particular model day, the 30 observed days that best match (smallest spatial RMSE) the model day over the entire domain are found, then optimal weights for the 30 observed days are computed such that the weighted linear combination best reproduces the model day. Finally, the downscaled field is obtained by combining the original fine-resolution observed fields using those same optimal weights.

LOCA is nearly as straightforward: to spatially downscale a model day, the 30 observed days that best match the model day in the wider region around the point being downscaled are found (these are termed the analog days). The wider region is determined by examining the spatial map of historically observed temporal correlations, by season, between the variable at the point being downscaled and that same variable at all other locations in the domain; locations where the correlation is positive are included in the wider region. Elsewhere, the agreement or disagreement between the model field and observed day is not considered. This approach lends a natural domain-independence to LOCA that is not found in earlier constructed analog methods, which require matching the analog days over the entire domain being downscaled, leading to challenges as the domain size increases.

Next, the single one of the 30 analog days that best matches (least RMSE) the model day in the local neighborhood of the point being downscaled is identified. For the local matcing we use a 1 degree box; tests with boxes of different sizes showed little difference when the box was half or twice this size. This multi-scale matching (over both the wider region, so synoptic scale patterns are matched, and locally around the point being downscaled) is one of the key aspects of LOCA, and ensures that the final downscaled field is consistent with the day being downscaled on both local and synoptic length scales.

The final downscaled value is the value from the best-matching single analog day, scaled so that its amplitude matches the amplitude of the model day being downscaled. For example, if temperature is being downscaled and the model gridcell has a 5°C temperature anomaly, but the best matching observed day shows only a 4°C anomaly when averaged over the model gridcell, then the value at the point being downscaled is increased by 1°C. For full details on the LOCA method, see Pierce et al. 2014.

2.4 Experimental design

2.4.1 Downscaling approach

Our goal is to evaluate the LOCA spatial downscaling scheme, which, like in other constructed analog approaches, is a separate and independent step from bias correction. We downscale only over the historical period since our purpose is to validate against observations. Bias correction is not of direct interest here (cf. Pierce et al. 2015), in contrast to some other statistical downscaling schemes that combine spatial downscaling and bias correction in a single step. To assess the quality of the LOCA spatial downscaling step itself, distinct from problems in the GCM data or an independently implemented bias correction processes, two separate evaluations are conducted.

First, the observations are aggregated to the same 1x1 degree latitude-longitude grid used for the GCM data, then downscaled to 1/16° resolution (details in section 2.3.3). The advantage of downscaling the 1x1 degree coarsened observations is that it is known in advance that the statistics of the coarse resolution data are correct, therefore no bias correction is needed, and the results can be directly compared to the original fine-resolution observations. Any deficiencies in the final downscaled fields must be due to shortcomings of the spatial downscaling process alone.

For comparison, we also show results from the coarsened observations downscaled with an earlier constructed analog technique, BCCA (Maurer et al. 2010). BCCA was chosen for comparison because LOCA was developed to mitigate some of the problems that arise when multiple analog days are averaged together, as BCCA does.

Although examining the results of downscaling coarsened observations is an ideal way to evaluate a spatial downscaling scheme, in practice downscaling is used on GCM output. Therefore, our second evaluation uses data from an example GCM that is first bias corrected and then spatially downscaled by LOCA, and statistical properties of the result are compared to observations (details in section 2.3.2). Although this conflates errors in the bias correction with the results of interest here--the LOCA spatial downscaling--it nonetheless represents the typical use case for downscaled climate model data.

In LOCA, temperature is downscaled as an anomaly and precipitation is downscaled as an absolute value. In other words, when temperature is downscaled, the long-term climatic mean for that day is first computed, then the departure from the mean for any particular location and time (i.e., the anomaly) is calculated. This field of anomalies is then spatially downscaled. When precipitation is downscaled, no such transformation is performed. The reason for the distinction is so that no final downscaled precipitation value is less than zero. Were precipitation downscaled as an anomaly, this could not be guaranteed except by discarding negative values after the final downscaled field was constructed. Since specific humidity likewise cannot be less than zero, it was downscaled as an absolute value as well.

2.4.2 GCM bias correction on a common 1x1 degree grid

Downscaling typically starts with GCM data, which often fail to reproduce the statistics of the observed field. Because of this the GCM data need to be bias corrected before a spatial downscaling step such as LOCA is applied. It is advantageous if GCMs, which have a diversity of native grid resolutions, are bias corrected on a common grid rather than their native grids. This prevents a relatively high resolution GCM from being informed by more observed information during the bias correction process than a low resolution GCM. For conformity with the extensive archive of downscaled data from Maurer et al. (2014), we regrid the CCSM4 GCM data to the same 1x1 degree latitude-longitude employed by Maurer et al. (2014) and bias correct it at the 1x1 degree resolution.

Although our interest here is in evaluating the LOCA spatial downscaling scheme rather than the independently applied bias correction, for completeness we examined the effects of LOCA spatial downscaling after applying three different bias correction methods to the GCM data: 1) quantile mapping (e.g., Panofsky and Brier 1968; Wood et al. 2002; Thrasher et al. 2012) followed by adjustments of the specific humidity time series in frequency space to make the model's variance spectrum better match that observed (details in Pierce et al. 2015). Quantile mapping has been widely used as a bias correction technique for climate model data downscaled to the western U.S. (e.g., the Maurer et al. 2014 archive noted above). 2) Equidistant quantile matching (EDCDFm; Li et al. 2010), which bias corrects by adding a model-estimated change in distribution to the observed distribution of a climate variable. 3) The Cumulative Distribution Function transform method (CDF-t; Michelangeli et al. 2009), which applies to the GCM data functions that transform the model's CDF to that observed. The results with all three methods are very similar. Results from EDCDFm and CDF-t, in particular, are almost indistinguishable, and typically showed only small differences from QM on most measures. Accordingly, results from only QM and EDCDFm and displayed below; equivalent plots for CDF-t are given in the supplementary material. Note that the differences between these bias correction methods become more important in future projections; however, since we want to compare to observed data in this work, results are only considered over the historical era.

As is the case when bias correcting precipitation (Pierce et al. 2015), a minimum value is specified below which the spectral adjustment is not applied; we use a threshold of 0.0015 kg/kg for specific humidity.

2.4.3 Downscaling coarsened observations using cross validation

Bias correction is not perfect, so any residual differences between the statistics of the observations and final downscaled model field could arise from uncorrectable problems with the original GCM data, limitations in the bias correction process, or flaws in the downscaling technique itself. This makes it hard to evaluate the quality of a downscaling technique using GCM data alone. To get a better idea of the quality of the LOCA downscaling step itself, distinct from problems in the original GCM data or bias correction processes, we also downscale the observations aggregated to the same 1x1 degree latitude-longitude grid used for the GCM data.

In order to fairly compare the downscaled version of the coarsened observations to the original fine-resolution observations, the statistical downscaling model must be trained on a different data set than is used to verify the downscaling quality. This is often accomplished by partitioning the observations into independent training and verification periods. Here we only have ~30 years of observed data, so this approach is not practical. Instead we downscale using cross-validation. Cross-validation is achieved by requiring the selected analog days to be far removed from the target day being downscaled. I.e., when downscaling the coarsened representation of an observed day, no analog day can be chosen within +/- 320 days of the day being downscaled. This prevents information from any observed day from being used simultaneously in both training and validation. The autocorrelation of daily specific humidity in the CONUS typically drops into the noise around zero after ~100 days, so the 320 day separation requirement should be adequate for the cross-validation. Furthermore, 320 days is the *minimum* allowed separation; on average, it will be many years since we use a 30-year training data set.

Evaluating the coarsened observations downscaled with cross-validation gives the clearest picture of the qualities of the downscaling process itself, independent of problematic GCM data or deficiencies in the bias correction. This is complementary to evaluating the bias corrected and downscaled data from a GCM, which gives the clearest picture of the quality of the final downscaled product as it is typically generated in practice. We use both approaches in an attempt to provide a complete picture of the downscaling's quality.

3 Results

3.1 Downscaling the coarsened observations with cross-validation

Figure 1 shows selected monthly mean specific humidity fields from the observations (left column), and the error after the coarsened observations are downscaled to 1/16° using cross-validation with BCCA (middle column) and LOCA (right column). Errors using LOCA are generally small, on the order of 0.02% in the mean and < 1% RMS, where RMS errors are calculated over the spatial extent (here and in the following figures). Errors and RMS values in the other months and annual mean are similar (not shown). In sum, LOCA preserves the monthly and annual mean humidity values accurately in the downscaled fields. In the monthly means, BCCA and LOCA do a comparable job.

The temporal standard deviation of observed daily specific humidity values, by season, is shown in the left column of Figure 2. Variability is lowest in the Rocky Mountain region in winter, and highest in the Gulf Coast region during autumn. Errors after downscaling with BCCA are shown in the middle column, and errors after downscaling with LOCA are shown in the right column. LOCA captures the variability well, with mean errors of ~0.3% and RMS errors of about 1%. Although BCCA did well with the monthly means, it does a poorer job reproducing the variability, with significant underestimation of the values in the western part of the U.S., and the Southeastern seaboard in the summer.

In addition to monthly means and daily variability, extreme daily values are also relevant to evaluating the quality of the downscaling. Both minimum and maximum humidity are of interest, since low humidity affects wildfire risk and agriculture, while high humidity affects human health/comfort and air conditioning energy demand. Figure 3 shows the 1-day in 1- and 20-year maximum value of humidity from observations (left column), and the error (%) after downscaling with BCCA (middle column) and LOCA (right column). The main observed large-scale feature evident in the maximum humidity pattern is a pronounced east-west divide separating the higher maximum humidity eastern and lower maximum humidity western states, with extreme high values typically exceeding 0.018 kg/kg in the eastern U.S. versus less than 0.013 kg/kg in the western U.S. The highest values in the conterminous U.S. are found along the Gulf Coast in the 1-day-in-1-year maximum humidity fields, while the 1-

day-in-20-years maximum field shows appreciable penetration of maximum values inland as far as southern Minnesota. LOCA reproduces the extreme maximum patterns and values quite well in both cases, with mean errors of less than 0.5% and RMS errors on the order of 1-2%. BCCA shows larger errors, especially in the western half of the U.S.

Minimum extreme specific humidity values are shown in the same format in Figure 4 (note the change in units, which are 1/10 those for the maximum specific humidity in Figure 3). The lowest observed values are found in the upper Midwest, associated with outbreaks of cold, dry winter air. Interestingly, in comparing Figures 3 and 4, the southern parts of California and Arizona sometimes experiences extremely low specific humidity events (presumably these occur in fall and winter Santa Ana events) as well as extremely high specific humidity events (likely during monsoon moisture pulses). Errors after downscaling with LOCA (right column) are larger in the 1-day-in-20-year minimum extreme specific humidity field than found in the other fields, especially in central California, where errors can exceed 30% in spots. However errors in BCCA are much larger still, especially in the 1-in-20 year extreme.

The spatial coherence of the downscaled field can be a sensitive indicator of the quality of the downscaling. Pierce et al. (2014) demonstrated how averaging together multiple analog days, as done in CA and BCCA, tends to inflate the spatial coherence of the downscaled field. Spatial coherence here is evaluated in terms of how quickly the standard deviation of daily values declines as progressively more gridcells are combined into a regional average (cf. Gutmann et al. 2014). A field that has high spatial coherence will show a more gradual decline in standard deviation as progressively more surrounding gridcells are incorporated into the regional average than a field with low spatial coherence. The slope of the relationship between the number of points included in the spatial average and the standard deviation then becomes a metric of spatial coherence, with more negative values indicating low spatial coherence and less negative values indicating high spatial coherence. Values are nondimensional since they are evaluated as the best-fit power law relationship between the standard deviation and number of points included in the averaging (e.g., Lovejoy et al. 2008). The result of this analysis is shown in Figure 5 for the observations (left), and the error after downscaling the coarsened observations with BCCA (middle) and LOCA (right). The errors are shown as a percentage ((model – observations)/abs(observations)*100). LOCA generally reproduces the observed pattern reasonably well, although downscaled values are generally about 13% less negative than observed, indicating somewhat increased spatial coherence in the downscaled data compared to the original field. BCCA, by contrast, has about twice the error as LOCA.

A final, direct evaluation of the skill of the downscaling is performed by correlating, at each point, the downscaled time series with the original data (Figure 6). This is done with and without the annual cycle and for three levels of temporal averaging to show how the results vary as a function of time scale. The left column shows results using BCCA, and the right column shows results using LOCA. The top row is computed using the original daily time series, including the annual cycle; the second row using daily anomalies; the third row using monthly anomalies; and the bottom row using yearly anomalies. Correlations tend to be highest in the Midwest, with lower values along the west coast and (particularly for the yearly anomalies) along the eastern seaboard. However with LOCA all values are high, with all locations having correlations of 0.86 or greater and generally greater than 0.94. Values are somewhat lower using BCCA, especially in the western half of the U.S.

3.2 Downscaling the CCSM4 GCM

As described in section 2.3.2 we use three different bias correction methods (quantile mapping [QM], EDCDFm, and CDF-t) on the CCSM4 GCM data before the LOCA downscaling. Results from EDCDFm and CDF-t are almost indistinguishable, and differ little from results using QM, so the CDF-t results are relegated to the supplementary information.

Selected monthly mean values obtained using LOCA to downscale the QM and EDCDFm bias corrected specific humidity model data to the 1/16° grid are shown in Figure 7. (Results with CDF-t bias correction are shown in supplementary information figure S1.) Mean errors are less than 0.6%, while RMS errors are about 2%. Although these values are small, it is worth noting that they are nonetheless considerably larger than the errors found when downscaling the coarsened observations (Figure 1), where the mean error was less than 0.03% and the RMS error less than 1%. As discussed in section 2.3.2, the difference between these two cases arises from errors in the GCM data that are not completely corrected by the bias correction schemes. To put this result into context, Figure 8 shows the CCMS4 GCM data before bias correction. Mean errors are on the order of 25%, and RMS errors up to ~30%. So the bias correction methods greatly reduce both the mean and RMS errors, but do not completely eliminate them.

The seasonal standard deviation of daily specific humidity values obtained when downscaling the QM and EDCDFm bias corrected CCSM4 output is shown in Figure 9. (Results with CDF-t bias correction are shown in supplementary information figure S2.) Mean errors in the variability are on the order of 3%, and RMS errors similar. Again, both are considerably larger than seen when downscaling the coarsened observations using cross-validation (Figure 2), yet the original CCSM4 GCM has errors on the order of 50% in Dec-Jan-Feb and 25% in the other months (not shown). So the bias correction methods used here give large, but not complete, improvements.

The 1-day-in-1- and 20-year extreme maximum specific humidity values obtained when downscaling the CCSM4 GCM data are shown in Figure 10. (Results with CDF-t bias correction are shown in supplementary information figure S3.) Errors are shown as a percentage with respect to the observed value; supplementary information figure S5 shows errors as actual values. Downscaled extreme values are in reasonable agreement with the observations, with mean errors of ~0.5% and RMS errors of ~4%, although again the errors are about twice those seen when downscaling the coarsened observations (Figure 3). The Midwest through the upper Midwest region is especially prone to higher values in the downscaled CCSM4 GCM field than observed.

Downscaled extreme minimum humidity values are shown in Figure 11. (Results with CDF-t bias correction are shown in supplementary information figure S4; errors shown as actual values rather than as percentages are given in figure S6.) The RMS errors in the extreme minimum field are considerably higher than in the extreme maximum field, with a RMS error of 10-30% (compared to < 5% for extreme maximum values). The bias correction used also affects this field much more than the maximum daily specific humidity. Because the LOCA downscaling step generates notably better extreme minimum values when supplied with known-correct coarse data to downscale (Figure 4), we can infer that the reason for the relatively larger errors in Figure 11 is because of the bias correction step. I.e., bias correction of extreme minimum specific humidity values is more problematical than the bias correction of the extreme maximum values. The error measure shown here is accentuated because it is computed as a ratio, so there is a larger discrepancy between the errors of the extreme minimum and maximum values than would be seen if the errors were evaluated as a difference. In other words, a constant specific humidity error (as measured in kg/kg) represents a larger percentage error of a small base value (arid conditions) than of a large base value (humid conditions).

The spatial coherence of the daily specific humidity fields downscaled from CCSM4 is shown in Figure 12. The metric is the same as that described previously for Figure 5. Unlike the case when downscaling coarsened observations (Figure 5), where the spatial coherence was found to be increased by about 13% everywhere, when downscaling CCSM4 the coherence is lower than observed through the central part of the U.S. This is a feature of the CCSM4 GCM itself, and is little affected by the bias correction, which is implemented independently at each point and so does not take spatial patterns into account.

3.3 Multivariate downscaling

The previous results have been obtained with univariate downscaling, i.e., independently downscaling specific humidity to the 1/16° grid without regard to any relationship with other variables, such as temperature. This is different from the treatment in MACA (Abatzoglou and Brown, 2012), where humidity is downscaled in a multivariate scheme in conjunction with temperature. Abatzoglou and Brown (2012) focus on relative humidity, which has a stronger link to temperature than specific humidity, which is used here. However specific humidity in non-arid locations is still linked to temperature through the Clausius-Clapeyron relationship, so this leaves open the question of whether LOCA-downscaled specific humidity fields reproduce observed correlations with temperature.

The relationship between LOCA-downscaled temperature and LOCA-downscaled specific humidity is evaluated in Figure 13, which shows the correlation of daily specific humidity anomalies with daily maximum temperature anomalies (by season), and Figure 14, which shows the same result but using daily minimum temperature anomalies. In both figures the left column shows the observed correlations (observed temperature is taken from the University of Idaho; see section 2.1), while the right column shows the result using coarsened (to the 1x1 degree grid) temperature and specific humidity that were subsequently downscaled using LOCA. The correlation patterns are very similar between the original observations and downscaled data; spatial pattern correlations between the two are > 0.99.

We conclude that LOCA can adequately capture the observed relationship between temperature and specific humidity even when both variables are downscaled independently, at least in the case that the supplied coarse resolution data being downscaled have the correct multivariate relationships in the first place. In the event that GCM data are being downscaled, and the GCM itself does not simulate correct temperature-humidity relationships, LOCA downscaling will not fix that problem. This does not mean that multivariable downscaling is never of value in LOCA downscaling; Pierce et al. (2014) demonstrate that it is needed for correctly computing daily minimum temperature as the difference between downscaled daily maximum temperature and downscaled diurnal temperature range.

4 Discussion

Our analysis shows that LOCA does a credible job downscaling daily fields of specific humidity over the continental U.S. Compared to BCCA, an earlier constructed analog method that has been widely used in applications across the western U.S., LOCA has a very similar representation of the monthly mean fields, but shows improvements in the simulation of variability (e.g., monthly standard deviation of daily values, 1-day-in-1 and -20 year extreme values). Since many application areas are concerned with the impact of extreme events, we consider this to be a valuable improvement.

Different applications can require different humidity variables, such as relative vs. specific humidity. In this work we have examined specific humidity since this is the humidity variable most consistently available on a daily time step in the CMIP5 model data archive. From a downscaling point of view, one of the key differences between relative and specific humidity is relative humidity's stronger relationship with temperature and how that can be preserved in the downscaling process. We showed here that LOCA captures the observed correlations between temperature and specific humidity, at least at the seasonal timescale, so a fuller examination of the ability to downscale relative humidity would be warranted. The LOCA process itself can accomplish multi-variate downscaling (temperature and humidity together) by choosing analog days that are the best weighted match to multiple fields simultaneously.

Although no full multivariate-downscaled estimate of relative humidity has been attempted here, the societal importance of some applications that are dependent on relative humidity motivates us to do a preliminary evaluation using the downscaled data currently available, which are specific humidity and daily maximum temperature. The heat index (Steadman 1979; Rothfusz, 1990) makes a good application to evaluate, since it is a widely-known combination of temperature and relative humidity into an index that has bearing on human health and comfort.

We evaluated LOCA's ability to spatially downscale the heat index in two ways. First, we used LOCA's downscaled fields of coarsened observed specific humidity and daily maximum temperature to estimate relative humidity and thus the heat index using a constant sea level pressure (SLP) of 1020 hPa, and compared this to the heat index calculated from observed estimates of daily maximum temperature and minimum relative humidity (Abatzoglou 2012). This approach directly compares the LOCA results to our best estimate of the observed heat index (calculated with relative humidity), but has the drawback that it convolves errors in estimating relative humidity from specific humidity with errors in spatial downscaling, where only the latter are of interest here. We therefore also compared the LOCA-downscaled heat index to the observed heat index estimated from observed specific (instead of relative) humidity and daily maximum temperature computed with the same assumed fixed SLP as used in the LOCA calculation. This approach more clearly isolates errors due to the LOCA spatial downscaling scheme.

Results are shown in Figure 15 as the average number of days per year that the heat index exceeds thresholds of either the 27°C (top two rows) or 35°C (bottom two rows). LOCA reproduces the spatial pattern of exceedance using the 27°C quite well, with errors of only a few days/year, or typically less than 10% (although percentage values can become larger in mountainous western locations where the actual number of days/year is small). At the higher threshold errors are larger, especially in the great plains and southeastern U.S. (third row). However, when the downscaled field is compared to the observed field estimated using specific humidity in the same way as the LOCA field is estimated, errors are again small (bottom row). This suggests that the LOCA spatial downscaling step is not introducing errors, but that at the higher heat index threshold the assumption of fixed atmospheric pressure in the relative humidity estimation becomes problematical.

Like all statistical downscaling methods, LOCA makes stationarity assumptions that may be violated in a future that is subject to anthropogenic climate change. The main assumption of constructed analog-based techniques such as BCCA and LOCA is that the characteristic spatial patterns of climate variables remain unchanged in the future (although changes in the amplitude, frequency, or duration of spatial patterns can be captured). As a contrived example, imagine a domain that is smaller than a GCM gridcell, and historically has always experienced low humidity in the southern part of the domain and high humidity in the northern part of the domain. However, climate change alters the typical wind patterns such that in the future, the northern and southern parts of the domain generally have the same humidity. LOCA would have trouble reproducing that change if there were few or no historical analog days that showed that spatial pattern of humidity values.

5 Summary and Conclusions

The purpose of this work has been to evaluate the ability of the Localized Constructed Analogs (LOCA) statistical downscaling technique to spatially downscale specific humidity over the conterminous U.S. to a 1/16° spatial resolution, and examine whether a multivariate approach is required when downscaling specific humidity with LOCA.

Humidity is an important variable for applications to wildfire, agriculture, air conditioning energy demand, and human health and comfort, but is not always included in publically available archives of statistically downscaled climate simulations. We have evaluated the quality of the downscaling by comparing observed estimates of daily specific humidity (Abatzaglou 2012) to a downscaled version of the observations first coarsened to a 1x1 degree latitude-longitude grid, which isolates the effect of the LOCA downscaling step on the quality of the final downscaled result, and to downscaled specific humidity fields from the CCSM4 GCM, which illustrates how well the entire bias correction/downscaling process works for a typical GCM. We find:

- LOCA reproduces the observed monthly mean climatology of specific humidity with a mean error typically less than ~0.5% and a RMS error of typically ~2%. About half of the final error is attributable to residual errors in the GCM data after bias correction. (Before bias correction, GCM errors are on the order of ~25%, on average across the domain.)
- The temporal standard deviation of daily specific humidity values matches the observed value reasonably well, with mean errors of ~1% and RMS errors of ~3%.
- Extreme (1-day in 1- and 20-years) maximum specific humidity values, which are relevant to human health and comfort and air conditioning energy demand, are typically within ~5% of observed. Extreme minimum values, which are relevant to wildfire and agriculture, are typically within ~15% of observed. The relatively worse performance of the minimum extremes compared to the maximum extremes is largely attributable to residual errors in the bias correction.
- LOCA increases the spatial coherence of the downscaled specific humidity field by ~13%, using the metric described in this work. We found that this was reversed

in the downscaled GCM results in the central U.S. (less spatial coherence than observed), since the GCM simulated more spatial variability there than observed.

- Correlations between observed and downscaled time series of specific humidity typically are greater than 0.98, although values tend to be slightly lower in the western third of the conterminous U.S. then in other locations.
- LOCA accurately reproduces observed correlations between daily temperature (minimum and maximum) and specific humidity, even when temperature and specific humidity are downscaled independently.

Overall, these results show that the LOCA downscaling technique can provide useful high spatial resolution specific humidity fields from global climate model data, fields that can be applied to problems in hydrology, ecology, and energy demand.

6 Acknowledgements

We would like to thank a reviewer who made valuable suggestions that improved this work. This work was made possible by support from California Energy Commission, agreement #500-10-041, which is gratefully acknowledged. Additional support was provided by the NOAA California Nevada Applications Program (CNAP) RISA award NOAA NA11OAR4310150, and the Department of Interior's (U.S. Geological Survey) Southwest Climate Science Center, grant USGS G12AC20518. Computational resources in partial support of this work were provided by the NASA Earth Exchange (NEX) collaborative through the NASA Advanced Supercomputing (NAS) Division at Ames Research Center.

7 References

Abatzoglou, J. T. and T. J. Brown, 2012: A comparison of statistical downscaling methods suited for wildfire applications. Int. J. Climatol., 32, 772-780.

Abatzoglou J. T., 2012: Development of gridded surface meteorological data for ecological applications and modelling. Int. J. Climatol., doi: 10.1002/joc.3413.

Brown, T. J., B. L. Hall, and A. L. Westerling, 2004: the impact of twenty-first century climate change on wildland fire danger in the western United States: An applications perspective. Climatic Change, 62, 365-88.

Daly, C., M. Halbleib, J. I. Smith, W. P. Gibson, M. K. Doggett, G. H. Taylor, J. Curti,s, and P. A. Pasteris, 2008: Physiographically-sensitive mapping of temperature and precipitation across the conterminous United States. Int. J. Climatol., doi:10.1002/joc.1688.

Fowler, H. J., Blenkinsop, S., and Tebaldi, C., 2007: Linking climate change modelling to impacts studies: recent advances in downscaling techniques for hydrological modelling. Int. J. Climatol., 27, 1547-1578.

Friend, A. D., 1995: PGEN – An integrated model of leaf photosynthesis, transpiration, and conductance. Ecol. Model., 77, 233-55.

Gent, P. R., G. Danabasoglu, L. J. Donner, M. M. Holland, E. C. Hunke, S. R. Jayne, D. M. Lawrence, R. B. Neale, P. J. Rasch, M. Vertenstein, P. H. Worley, Z. -L. Yang, and M. Zhang, 2011: The Community Climate System Model version 4. J. Clim., 24, 4973-4991, doi: 10.1175/2011JCLI4083.1.

Hidalgo, H.G., M.D. Dettinger and D.R. Cayan, 2008: Downscaling with Constructed Analogues: Daily Precipitation and Temperature Fields over the United States. CEC PIER Project Report CEC-500-2007-123, 48 pp, January 2008.

Klein, W. H., and Glahn, H. R., 1974: Forecasting local weather by means of model output statistics. Bull. Amer. Meteorolol. Soc., 55, 1217-1227.

Lawrence, D. M., and coauthors, 2011: Parameterization improvements and functional and structure advances in version 4 of the Community Land Model. J. Adv. Model. Earth Syst., 3, M03001, doi:10.1029/2011MS000045.

Li, H., J. Sheffield, and E. F. Wood, 2010: Bias correction of monthly precipitation and temperature fields from Intergovernmental Panel on Climate Change AR4 models using equidistant quantile matching. J. Geophys. Res. Atmos., 115 (D10101), doi:10.1029/2009JD012882.

Livneh, B., E.A. Rosenberg, C. Lin, B. Nijssen, V. Mishra, K. Andreadis, E.P. Maurer, and D.P. Lettenmaier, 2014: A long-term hydrologically based data set of land surface fluxes and states for the conterminous United States: Updates and extensions, J. Clim. (in review).

Lovejoy, S., d. Schertzer, and V. C. Allaire, 2008: The remarkable wide range spatial scaling of TRMM precipitation. Atmos. Res., 90, 10-32.

Maraun, D., F. Wetterrhall, A. M. Ireson, R. E. Chandler, E. J. Kendon, et al., 2010: Precipitation downscaling under climate change: Recent developments to bridge the gap between dynamical models and the end user. Rev. Geophys., 48, RG3003, doi:10.1029/2009RG000314.

Maurer, E. P., H. G. Hidalgo, T. Das, M. D. Dettinger and D. R. Cayan, 2010: The utility of daily large-scale climate data in the assessment of climate change impacts on daily streamflow in California. Hydrol. Earth Syst. Sci., 14, 1125-1138, doi:10.5194/hess-14-1125-2010.

Maurer, E. P., Brekke, L., Pruitt, T., Thrasher, B., Long, J., Duffy, P. B., Dettinger, M. D., Cayan, D., and Arnold, J., 2014: An enhanced archive facilitating climate impacts analysis, Bulletin of the American Meteorological Society, doi:10.1175/BAMS-D-1113-00126.00121 in press.

Michelangeli, P.-A., M. Vrac, and H. Loukos, 2009: Probabilistic downscaling approaches: Application to wind cumulative distribution functions. Geophys. Res. Lett., 36, L11708, doi:10.1029/2009GL038401.

Mirasgedis, S., Y. Sarafidis, E. Georgopoulou, D. P. Lalas, M. Moschovits, F. Karagiannis, and D. Papakonstantinou, 2006: Models for mid-term electricy demand forecasting incorporating weather influences. Energy, 31, 208-227.

Mitchell, K. E., and co-authors, 2004: The multi-institution North American Land Data Assimilation System (NLDAS): Utilizing multiple GCIP products and partners in a continental distributed hydrological modeling system. J. Geophys. Res., D109, D07S90, doi:10.1029/2003JD003823.

Panofsky, H. A. and Brier, G. W.: Some Applications of Statistics to Meteorology, The Pennsylvania State University, University Park, PA, USA, 224 pp., 1968.

Pierce, D. W., A. L. Westerling, and J. Oyler, 2013: Future humidity trends over the western United States in the CMIP5 global climate models and variable infiltration capacity hydrological modeling system. Hydrol. Earth Syst. Sci.., v. 17, 1833-50.

Pierce, D. W., D. R. Cayan, and B. L. Thrasher, 2014: Statistical downscaling using Localized Constructed Analogs (LOCA). J. Hydromet, 15, 2558-85.

Pierce, D. W., D. R. Cayan, E. P. Maurer, J. Abatzoglou, and K. Hegewisch, 2015: Improved bias correction techniques for hydrological simulations of climate change. J. Hydromet., *in press*.

Rothfusz, L. P., 1990: The heat index ‘equation’ (or, more than you ever wanted to know about heat index). Technical Attachment SR 90-23, Scientific Services Division, NWS Southern Region Headquarters, Fort Worth, TX.

Rummukainen, M., 2009: State-of-the-art with regional climate models. WIRES Climate Change, 1, 82-96.

Steadman, R. G., 1979: The assessment of sultriness. Part I: A temperature-humidity index based on human physiology and clothing science. J. Appl. Meteorol., 18, 861-873.

Taylor, K. E., R. J. Stouffer, and G. A. Meehl, 2012: An Overview of CMIP5 and the experiment design. Bull. Am. Met. Soc., 93, 485-498, doi:410.1175/BAMS-D-1111-00094.00091.

Thom, E. C., 1959: The discomfort index. Weatherwise, 12, 57-61.

Thornton, P. E., M. M. Thornton, B. W. Mayer, N. Wilhelmi, Y. Wei, and R. B. Cook, 2012: Daymet: Daily surface weather on a 1 km grid for North America, 1980-2012. Available online at http://daymet.ornl.gov.

Thrasher, B., E. P. Maurer, C. McKellar, and P. B. Duffy, 2012: Technical Note: Bias correcting climate model simulated daily temperature extremes with quantile mapping. Hydrol. Earth Syst. Sci., 16, 3309-3314. Doi:10.5194/hess-16-3309-2012.

Van den Dool, H. M., 1994: Searching for analogues, how long must we wait? Tellus, 46A, 314-324.

Wood, A. W., E. P. Maurer, A. Kumar, and D. P. Lettenamier, 2002: Long-range experimental hydrologic forecasting for the eastern United States. J. Geophys. Res. D., 107, 4429, doi:10.1029/2001JD000659, 2002.

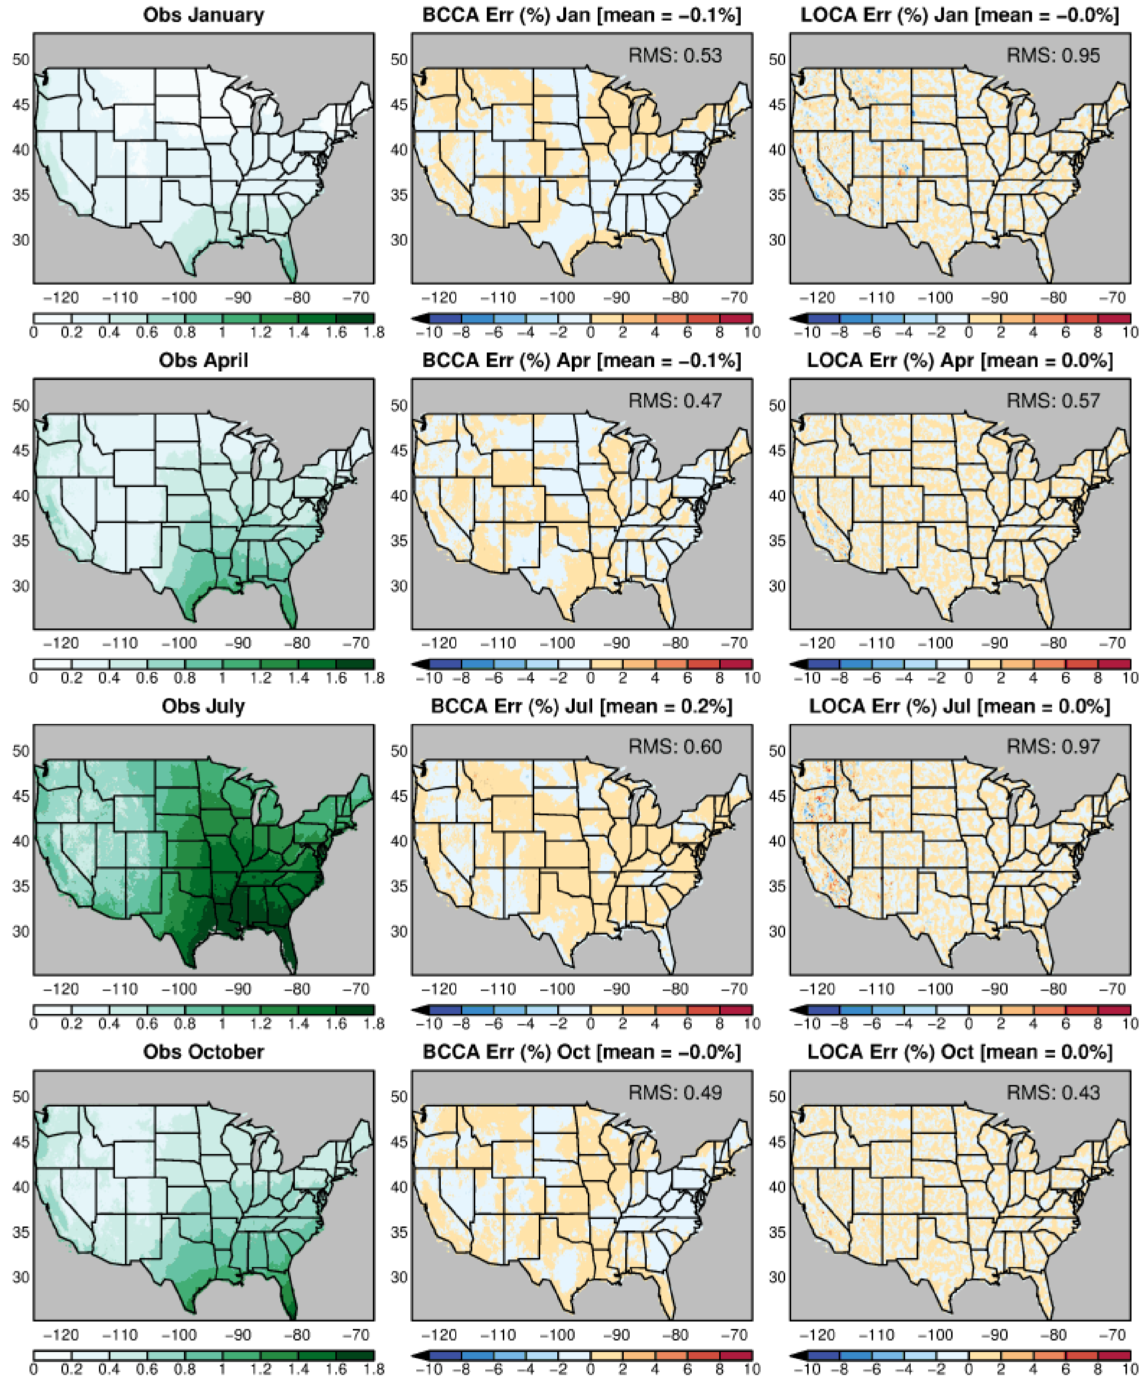

Figure 1. Left column: monthly mean specific humidity (kg/kg * 100) from the Abatzoglou observations. Middle column: error (%) in the monthly mean specific humidity after the coarsened observations are downscaled with BCCA. Right column: error (%) after downscaling with LOCA.

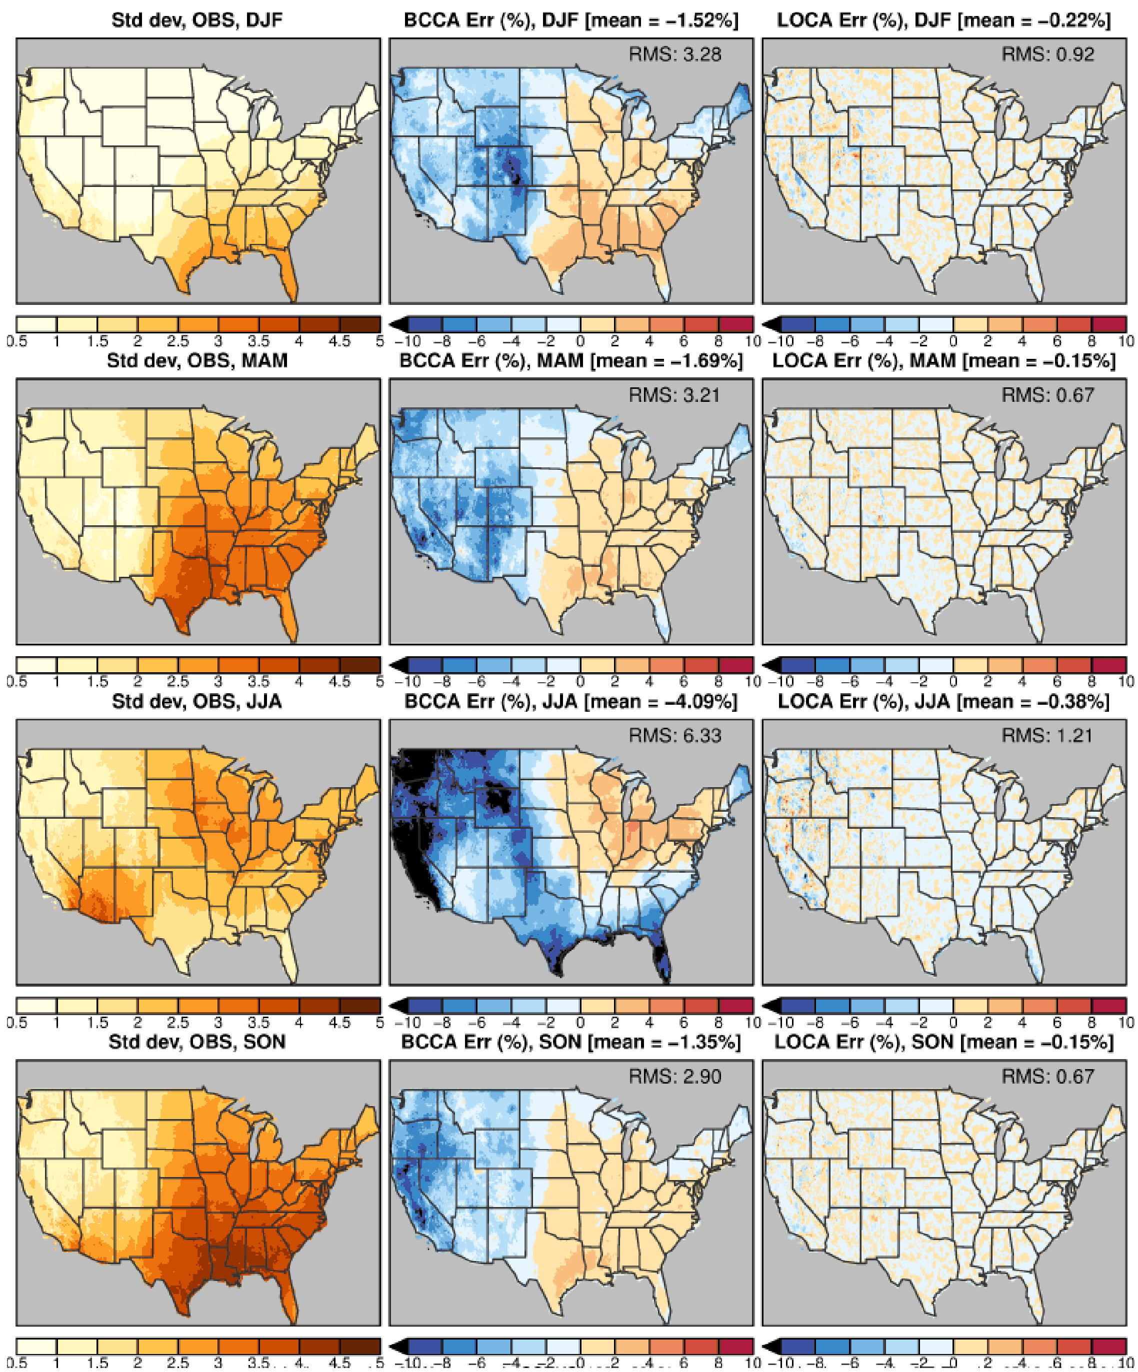

Figure 2. Left column: temporal standard deviation (kg/kg * 1000) of 1/16° daily specific humidity, by season, from the Abatzoglou observations. Middle column: error (%) in downscaled values with respect to observations after downscaling with BCCA. Right column: error (%) after downscaling with LOCA.

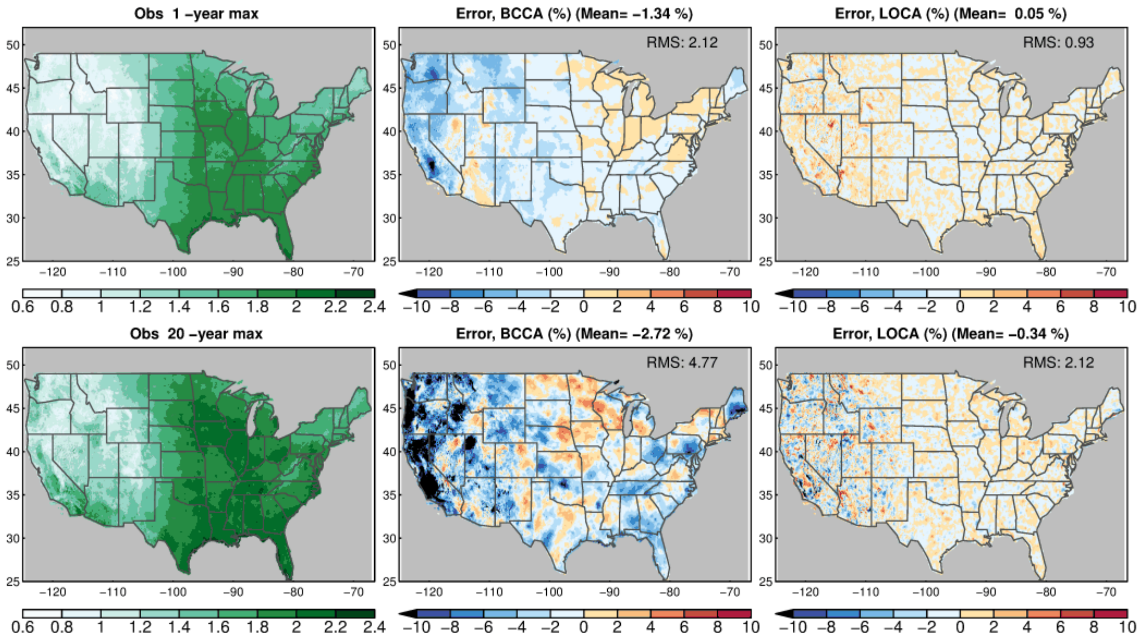

Figure 3. Left column: maximum 1-day-in-1-year (top row) and 1-day-in-20-years (bottom row) value of specific humidity (kg/kg * 100) from the Abatzoglou observations. Middle column: the error (%) in the downscaled value with respect to observations after downscaling with BCCA. Right column: error (%) after downscaling with LOCA.

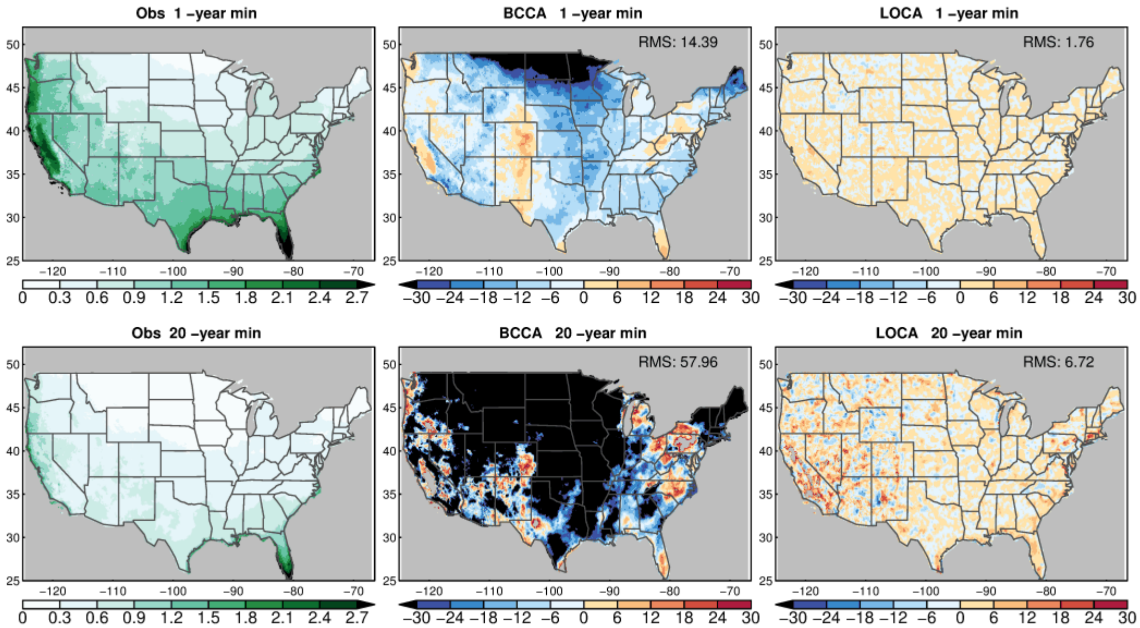

Figure 4. As in Figure 3, but for minimum daily specific humidity (kg/kg * 1000) rather than maximum. Note that units are 1/10 those for maximum daily specific humidity (Figure 3).

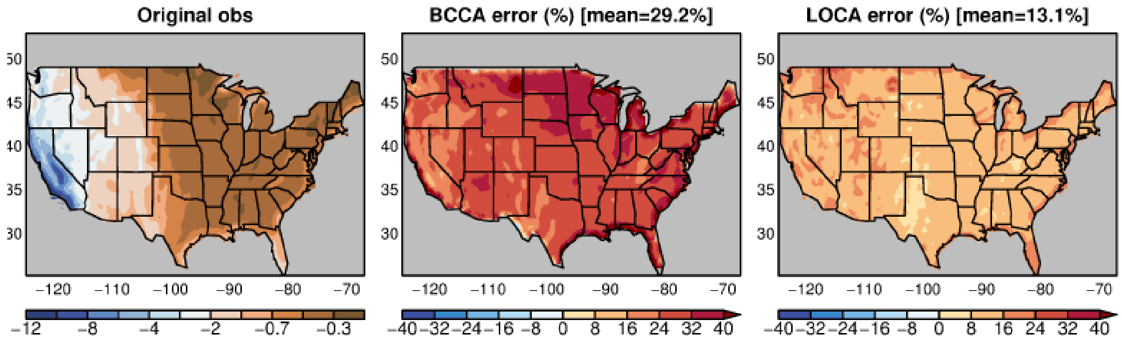

Figure 5. Left column: a metric of spatial coherence (nondimensional) for daily values of specific humidity in the Abatzoglou observations More negative values (blue) indicate low spatial coherence; less negative values (brown) indicate high spatial coherence. See text for definition of the metric plotted. Note nonlinearity of the color spacing. Middle column: error (%) with respect to the observed value after the coarsened observations are downscaled with BCCA. Right column: error (%) after downscaling with LOCA.

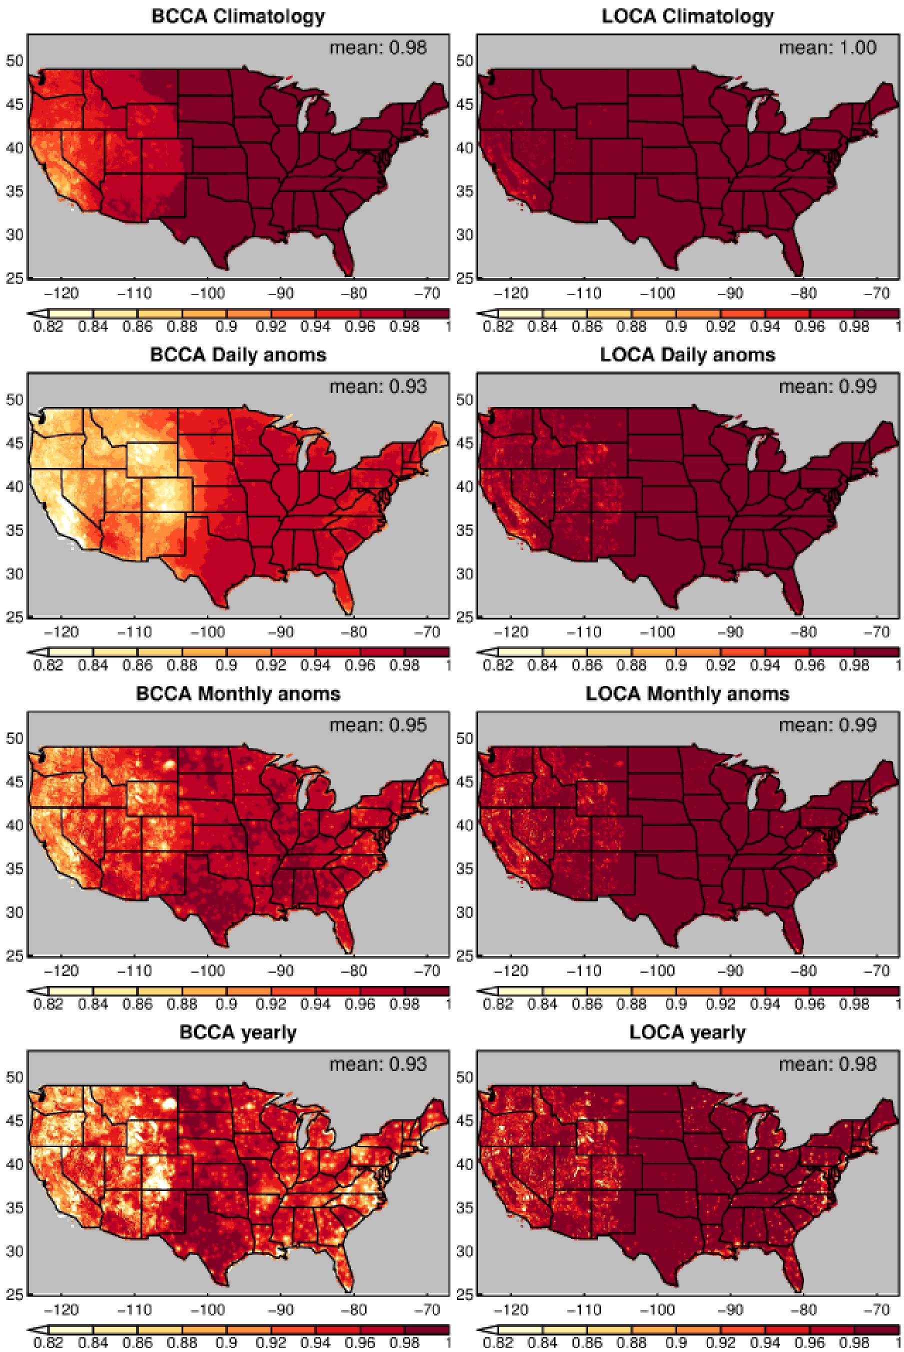

Figure 6. Temporal correlation at every point between the original observed specific humidity fields and the coarsened observations downscaled with either BCCA (left column) or LOCA (right column). Top row: evaluated with full time series of daily values including the annual cycle. Second row: evaluated with daily anomalies. Third row: evaluated with monthly anomalies. Bottom row: evaluated with yearly anomalies.

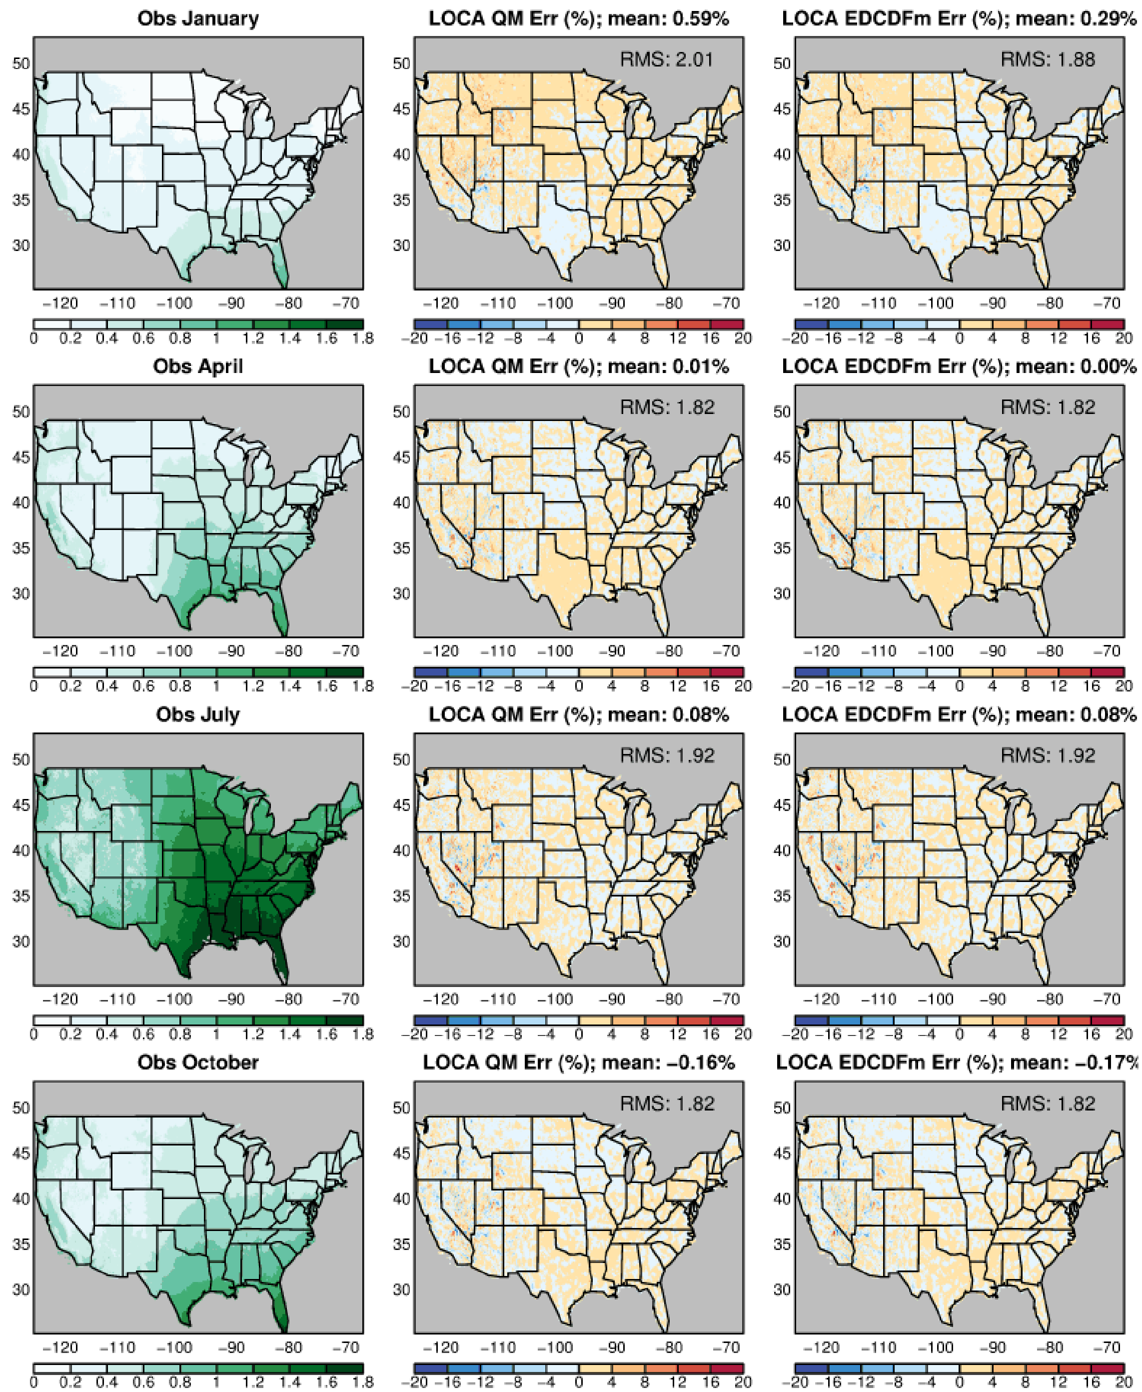

Figure 7. Left column shows the monthly mean specific humidity (kg/kg * 100) from observations for selected months. Middle column shows the error (%) after bias correcting the CCSM4 data with quantile mapping and spatial downscaling with LOCA. Right column shows similar, but using EDCDFm bias correction.

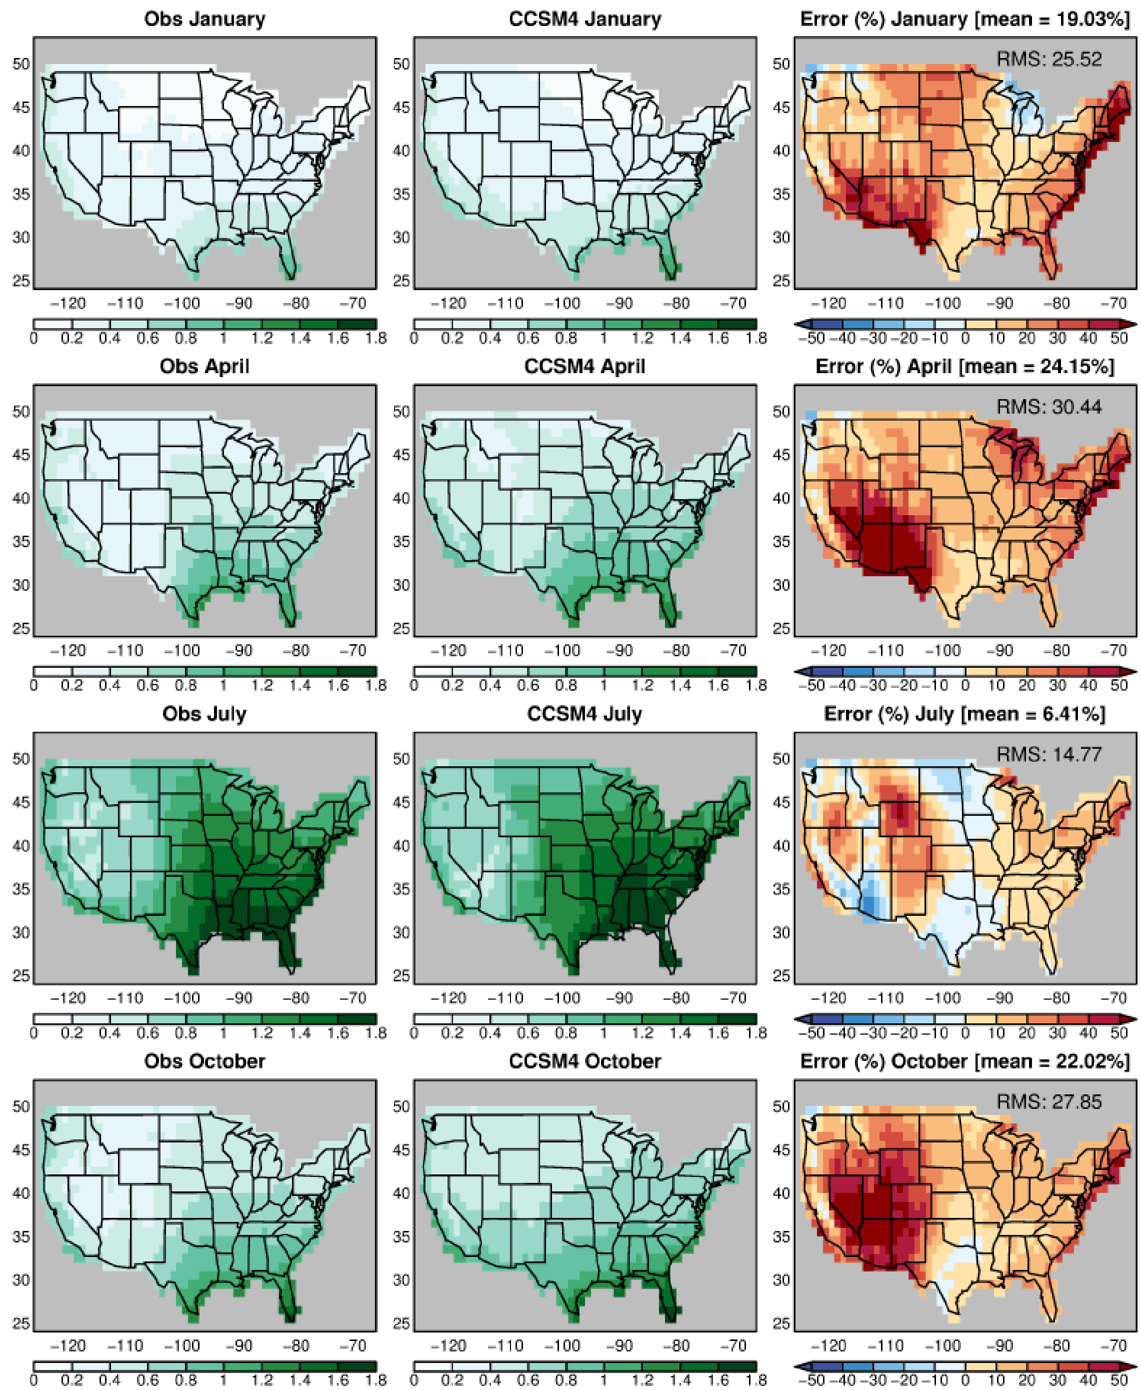

Figure 8. As in Figure 7, but for the original CCSM4 GCM data on the 1x1 latitude-longitude grid before any bias correction is applied.

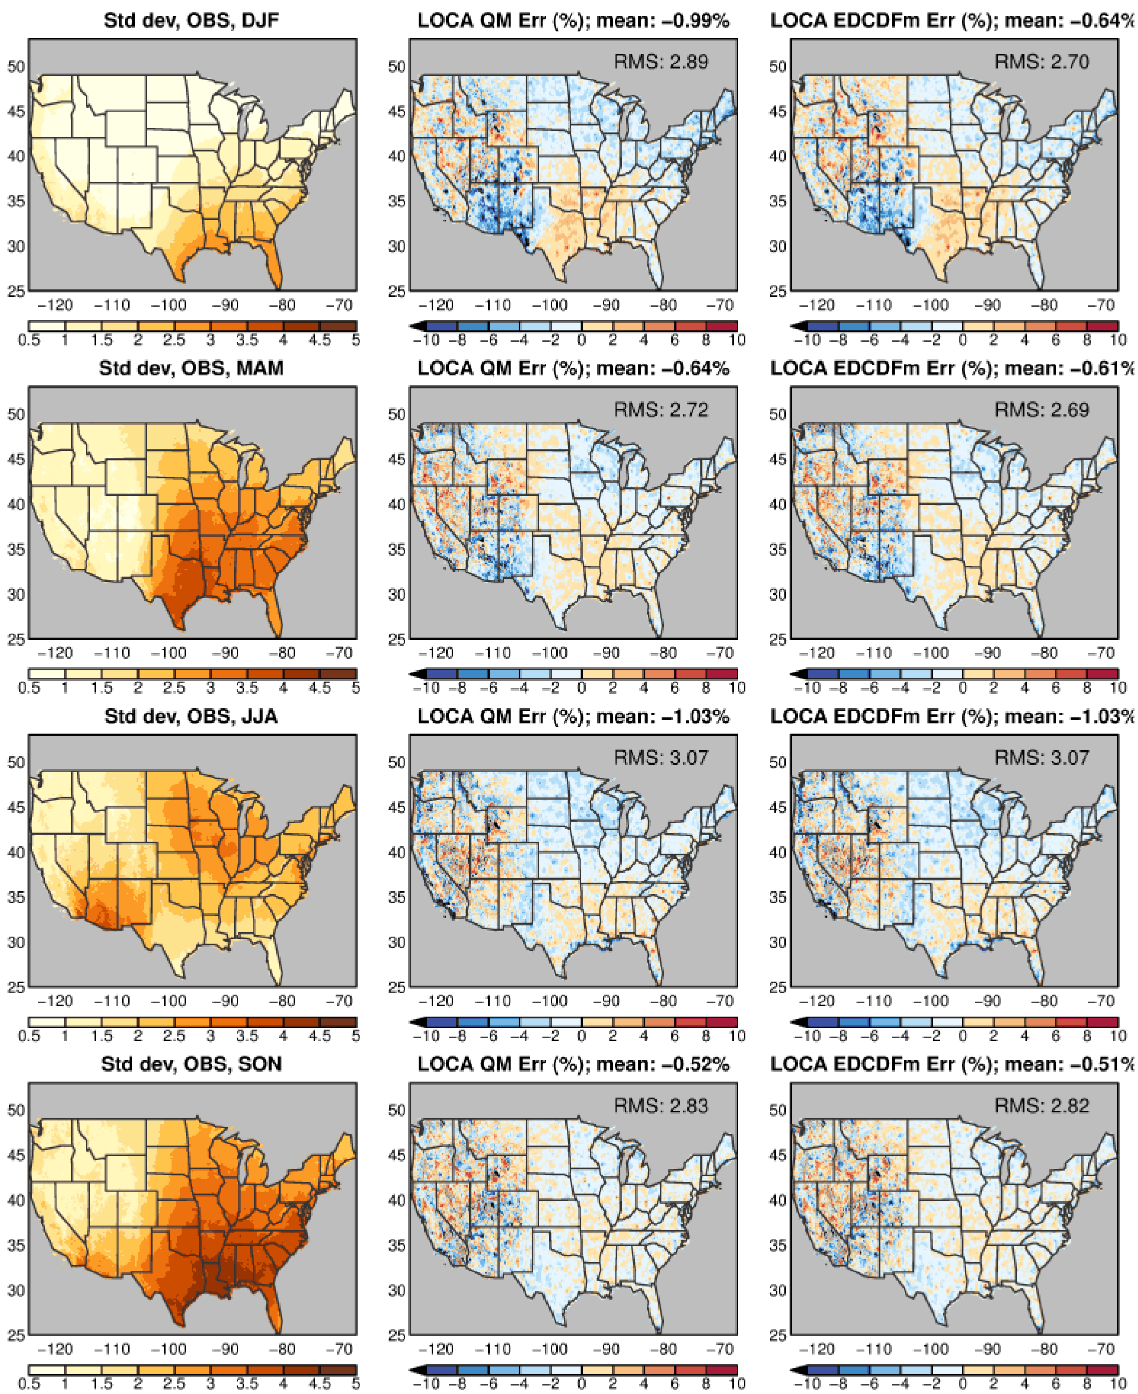

Figure 9. Left column shows the observed temporal standard deviation (kg/kg * 1000) of daily specific humidity, by season. Middle column shows the error (%) after bias correcting the CCSM4 data with quantile mapping and spatial downscaling with LOCA. Right column shows similar, but using EDCDFm bias correction.

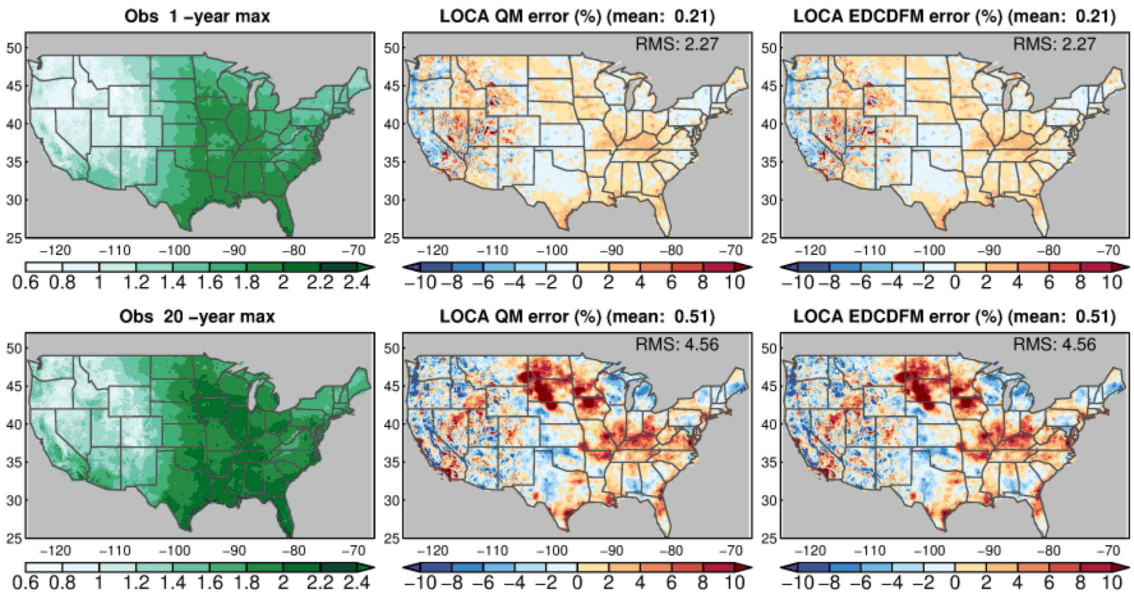

Figure 10. Left column shows the observed maximum 1-day-in-1-year (top row) and 1-day-in-20-years (bottom row) value of specific humidity (kg/kg * 100). Middle column shows the error (%) after bias correcting the CCSM4 data with quantile mapping and spatial downscaling with LOCA. Right column shows similar, but using EDCDFm bias correction.

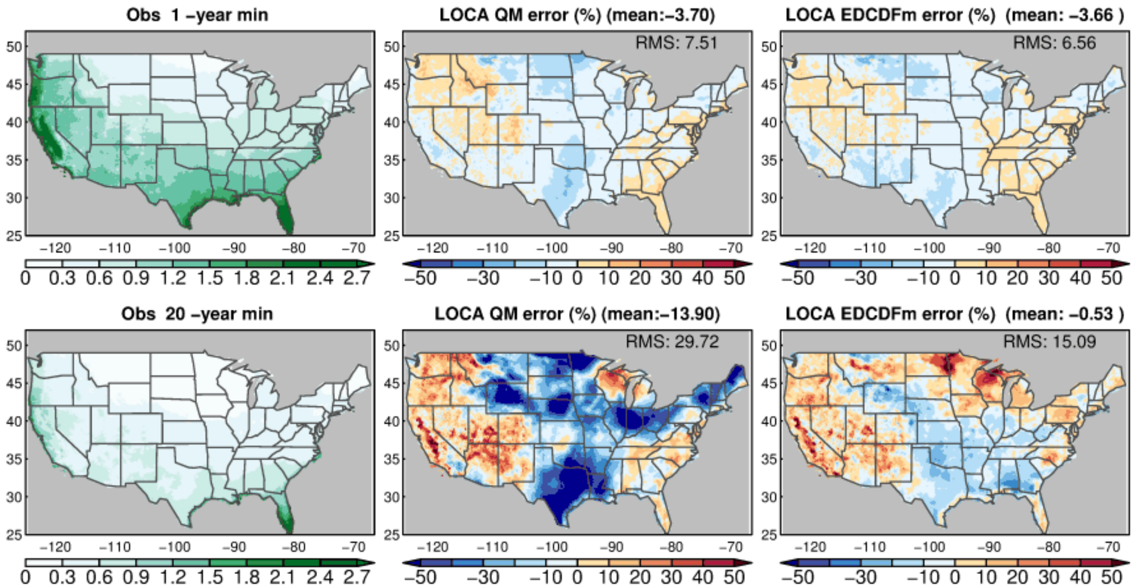

Figure 11. As in Figure 10, but for minimum daily specific humidity (kg/kg * 1000) rather than maximum. Note that units are 1/10 those for maximum daily specific humidity (Figure 10).

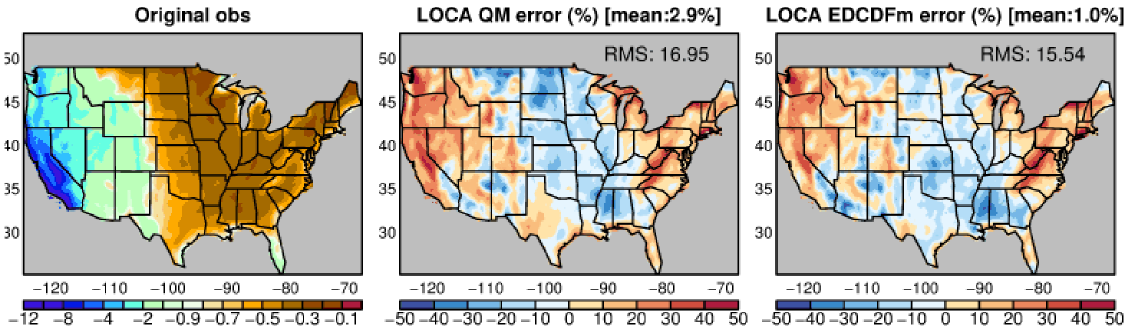

Figure 12. Left: Spatial coherence (nondimensional) for daily values of specific humidity in the original observations. See text for definition of the metric plotted. More negative values (blue) indicate low spatial coherence; less negative values (brown/red) indicate high spatial coherence. Note nonlinearity of the color spacing. Middle: Error (%) in the representation of spatial coherence in the CCSM4 GCM data after bias correction with quantile mapping and downscaling with LOCA. Right: same as middle, but for EDCDFm bias correction.

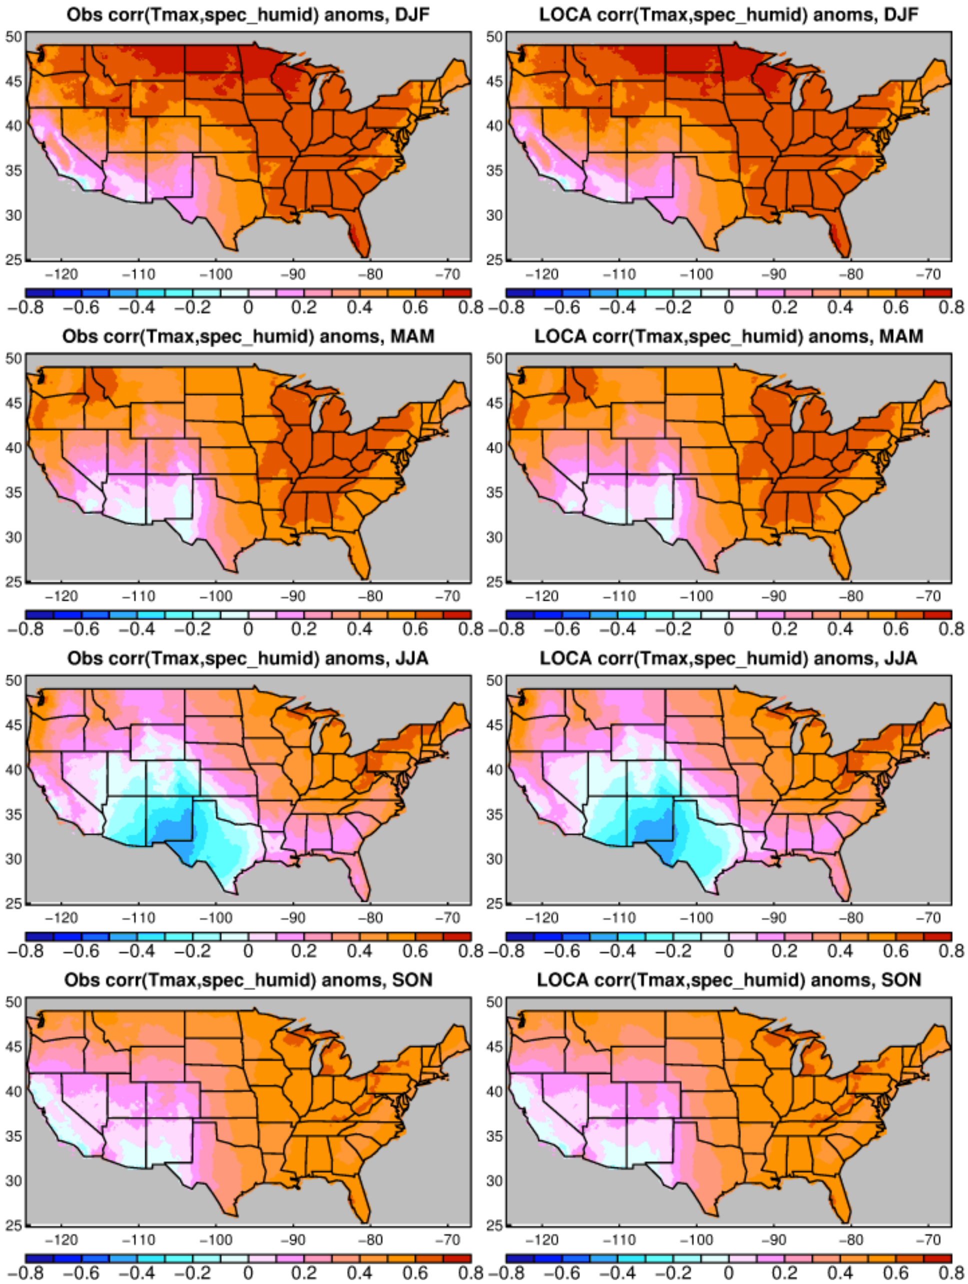

Figure 13. Point-by-point correlations, by season, between daily maximum temperature anomalies and specific humidity anomalies, computed using the observations (left column) and downscaled data (right column).

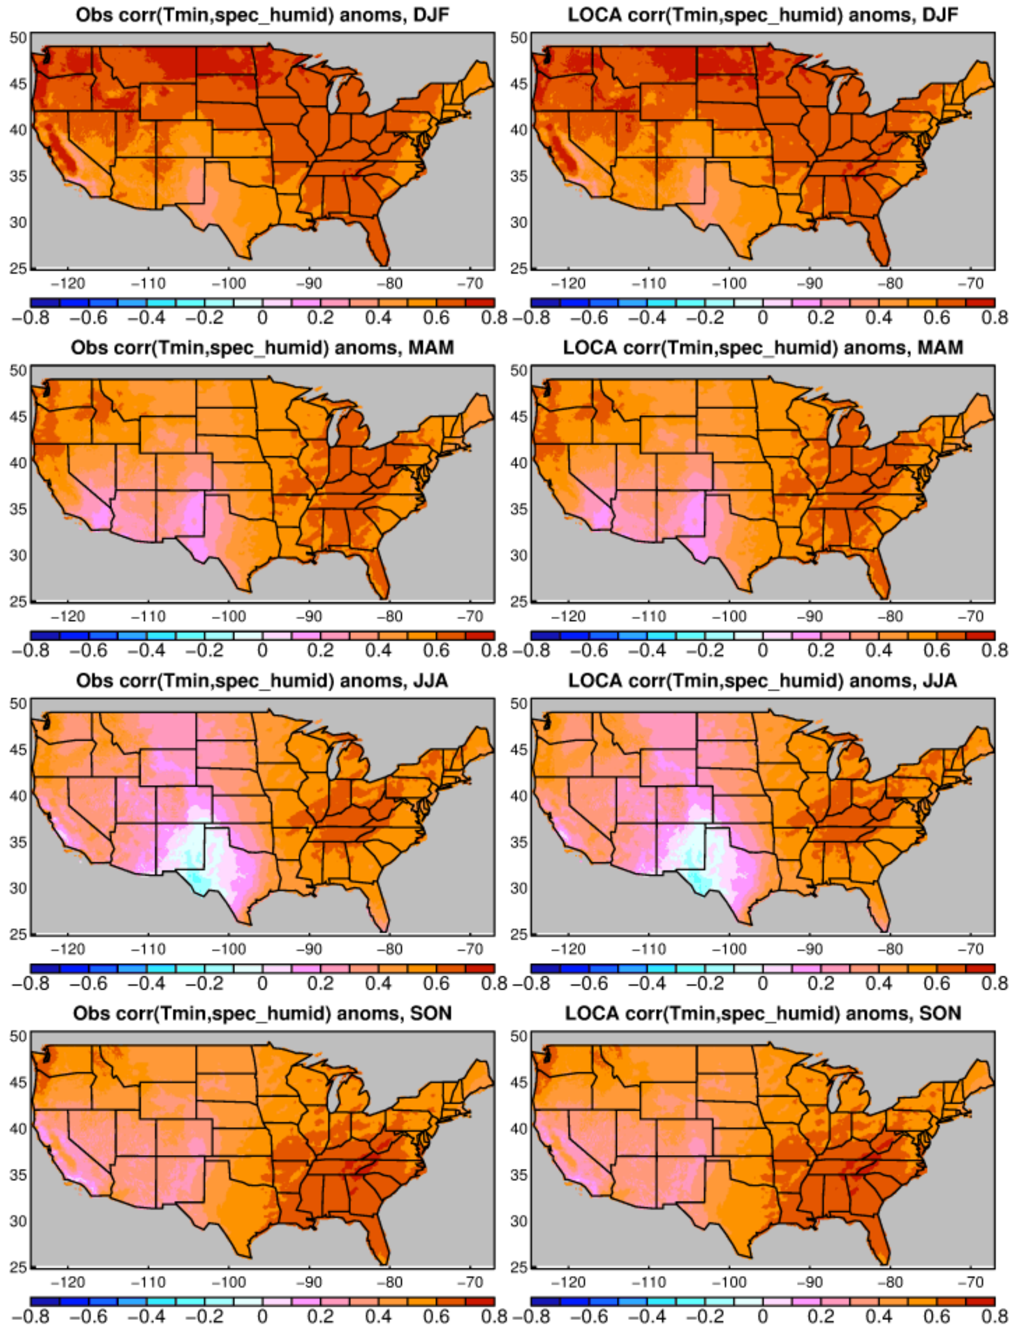

Figure 14. Same as Figure 13, but for daily minimum temperature anomalies correlated with specific humidity anomalies.

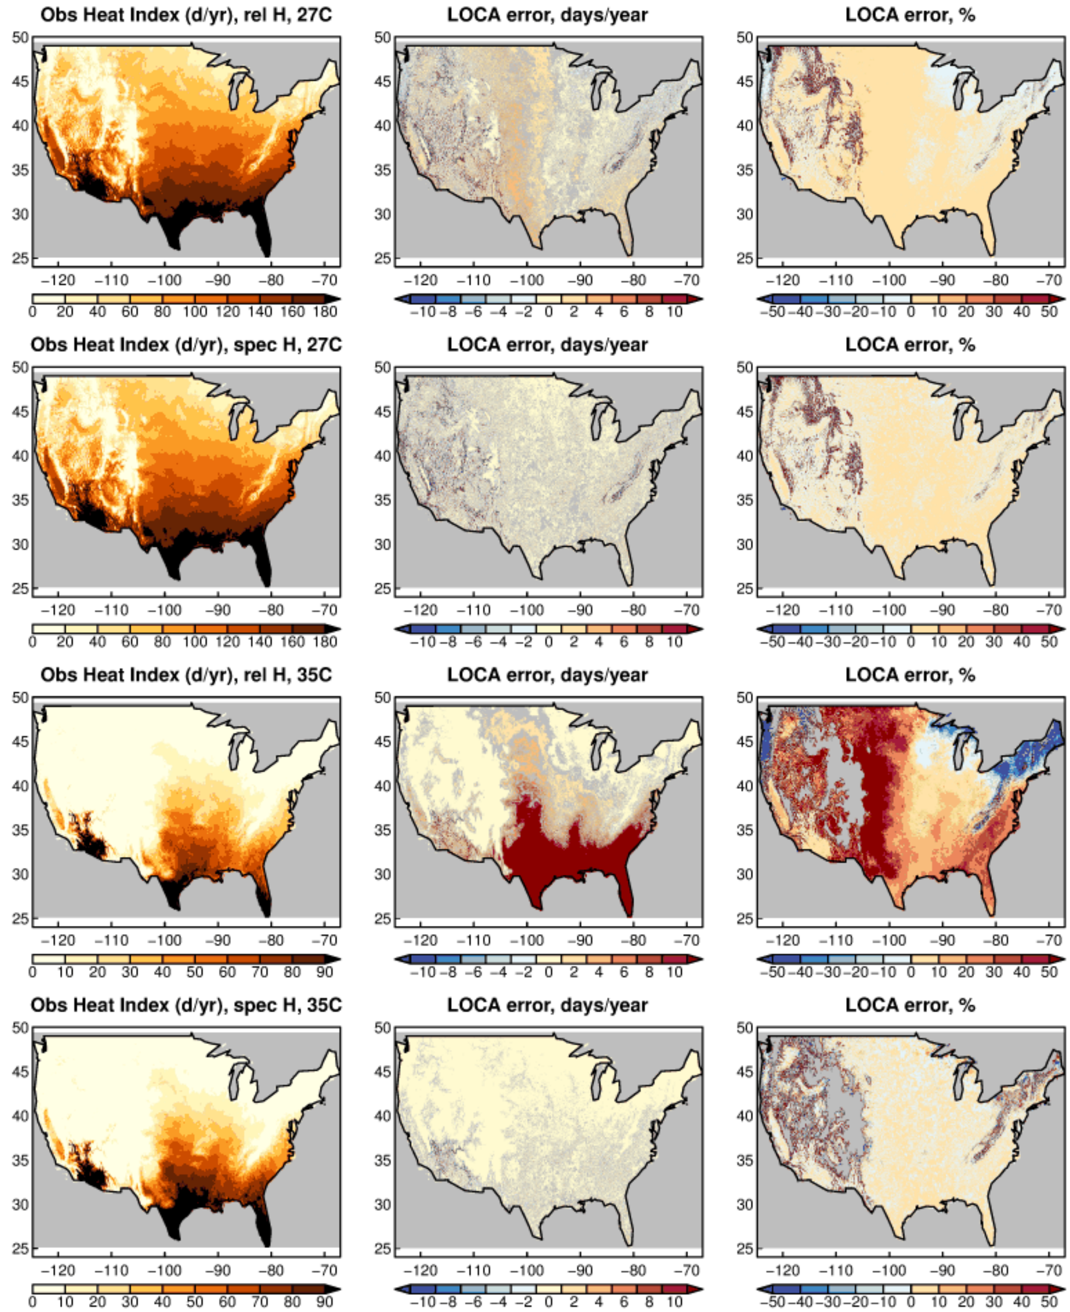

Figure 15. Comparison of the observed and downscaled heat index. The observed heat index is calculated two ways: using relative humidity (first and third rows), or specific humidity ignoring atmospheric pressure variations (second and fourth rows). The LOCA downscaled version is always calculated with the latter method. Observed values are the average number of days per year that either the 27°C (top 2 rows) or 35°C (bottom two rows) heat index thresholds are exceeded. LOCA errors are shown in days/year (middle column) and as a percentage (right column).