

# Release notes for the updated North American LOCA2 precipitation projections, version v20240915

David W. Pierce

Daniel R. Cayan

Daniel Feldman

Mark Risser

2024-10-15 Version 1.0

## Introduction

In this note we describe the updated version of Localized Constructed Analogs version 2 (LOCA2) downscaled CMIP6 precipitation: version v20240915, released in September 2024.

The first version of the LOCA2 North American precipitation dataset, which is based on the CMIP6 global climate models (GCMs), was released in spring of 2022<sup>1</sup>. Subsequently, an analysis of the LOCA2 product (Ullrich, 2023) showed that precipitation extremes in LOCA2 in some locations and seasons were 5-10% weaker than corresponding extremes found in the training data at very high precipitation values (e.g., at the 99.9<sup>th</sup> percentile). The largest discrepancies are found in summer and in a band stretching from the Gulf of California northwards through Nevada and eastern Oregon and Washington. Although not all applications will be affected by a modest deficiency at the far tail of the precipitation distribution, assessments of flooding and other heavy precipitation impacts could be impacted. Furthermore, the non-linearity of runoff processes with respect to the magnitude of rainfall means that even annually averaged runoff can be influenced by the extreme tails of the precipitation distribution (e.g., Pierce et al. 2021).

Accordingly, we addressed the shortfall in extreme precipitation found in the first version of LOCA2 and, in September 2024, released an updated version of the precipitation data that focuses on the representation of extremes. The purposes of this note are to 1) document the magnitude, extent, and cause of the low precipitation extremes in the first version of the data; 2) describe the process

---

<sup>1</sup> In the first released version, the temperature files are identified in the filename as version v20220413 and the precipitation files as v20220519.

we undertook to create the updated precipitation dataset; 3) evaluate the updated dataset, which is denoted version v202040915.

## Magnitude and extent of precipitation deficiencies in the original data

Statistical downscaling aims to reproduce the statistics of the training data used in the downscaling process, so the proper way to assess the skill of the downscaling method is to compare statistics of the downscaled data to those of the original training data. This comparison separates evaluation of the downscaling scheme from evaluation of the training data set, which is a separate issue. In the LOCA2 North American downscaling we used Pierce et al. (2021) as the precipitation training data set. We generally refer to this data set as “unsplit Livneh” because it is an updated version of the gridded surface hydrometeorological dataset from Livneh et al. (2015) that did not have data from morning-observing meteorological stations “split” across two calendar days. As described in Pierce et al. (2021) the splitting process reduces extreme precipitation and increases both the fraction and consecutive number of wet days, and we did not want the LOCA2 downscaled data to inherit those biases from the training data.

We evaluate the new (v20240915) and old (v20220519) versions of the LOCA2 data across a number of measures, illustrated in Figure 1, against the training data set. In each panel, the title shows the mean value of a given measure in the North American domain. Additionally, the domain is split into western, central, and eastern sub-domains using longitude lines 105°W and 90°W (shown as the vertical dashed red lines in Figure 1), and the means over these three subdomains are shown along the bottom of each panel.

The top row shows the fraction of wet days in winter (Dec-Jan-Feb, or DJF), summer (Jun-Jul-Aug, or JJA), and over the entire year, in percent. In our analysis a day is considered to be wet if it experienced precipitation of 0.5 mm or more in that day. The second row shows mean DJF, JJA, and yearly precipitation, in mm/day. Even though our focus here is on precipitation extremes, it is important that the representation of mean precipitation is realistic, so we include it in our evaluation. The third through fifth rows show percentiles of wet day precipitation (computed using the same wet-day threshold of 0.5 mm/day as previously). Values are shown for winter (DJF), summer (JJA), and yearly. Finally, our focus here is on precipitation extremes, so we include in our analysis the 50-, 100- and 500-year return values of daily precipitation, which are shown in the bottom row.

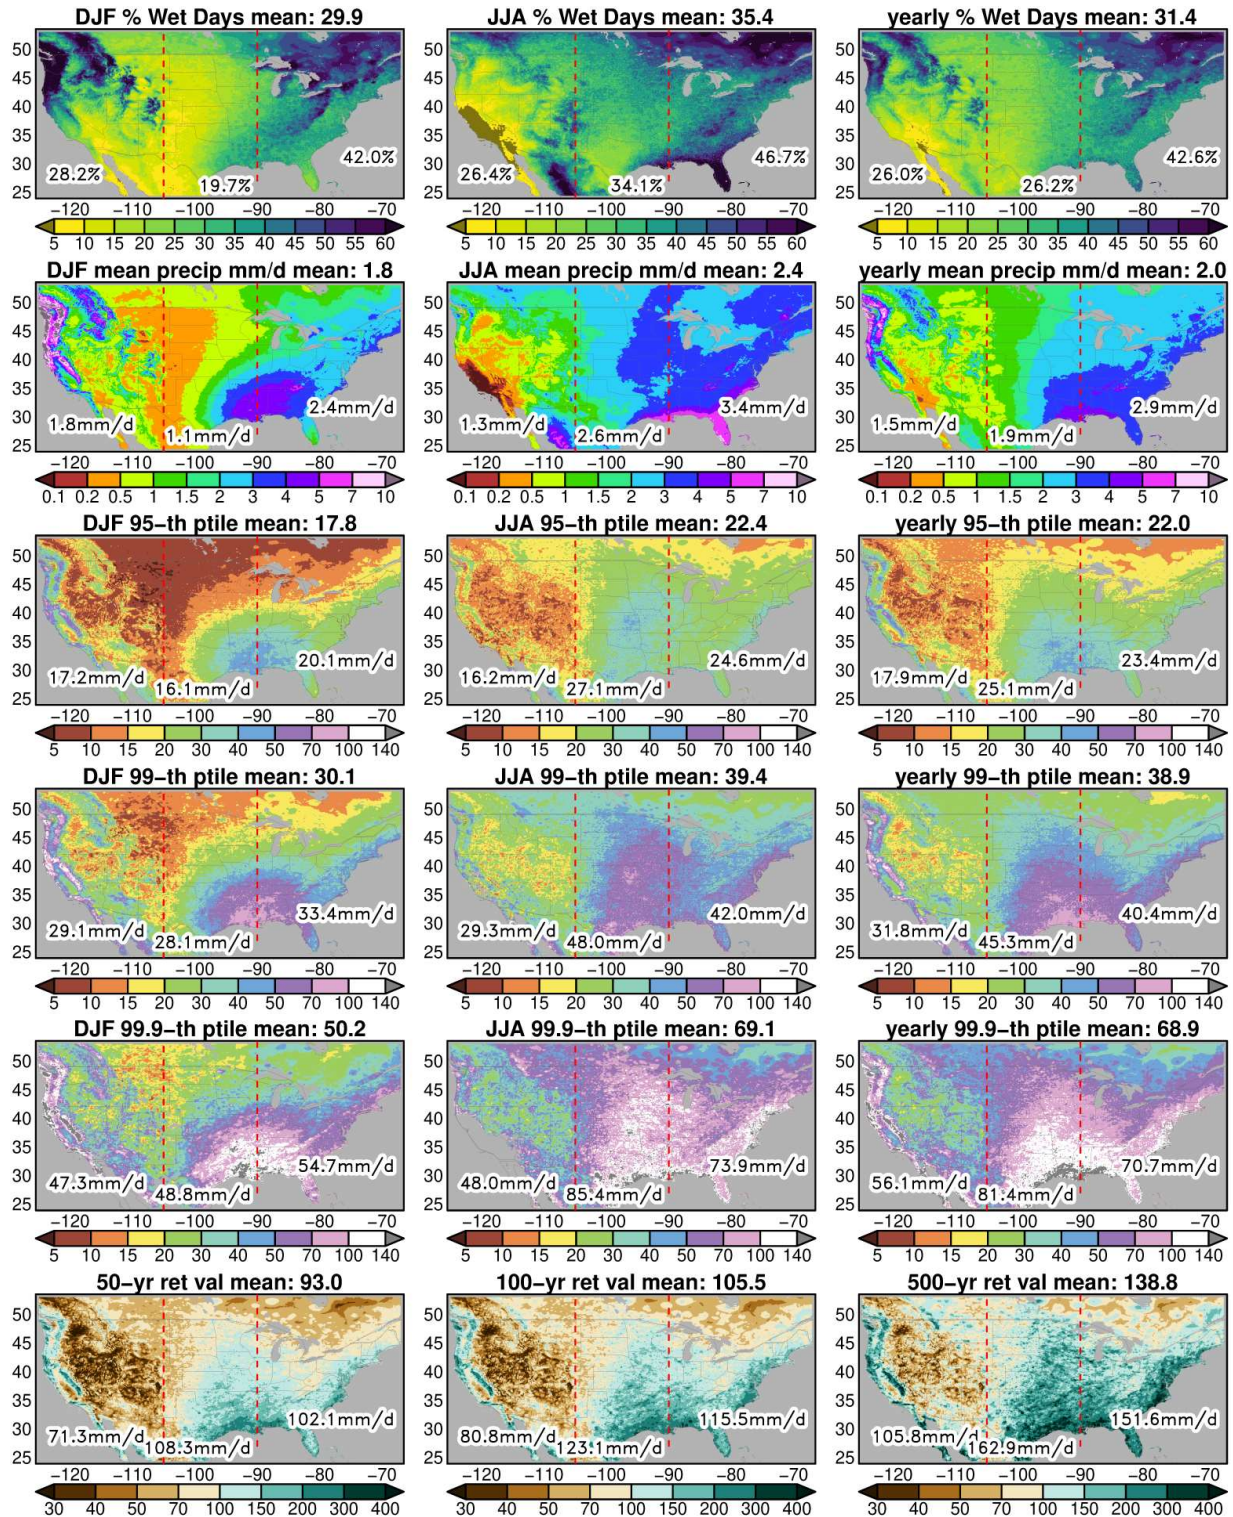
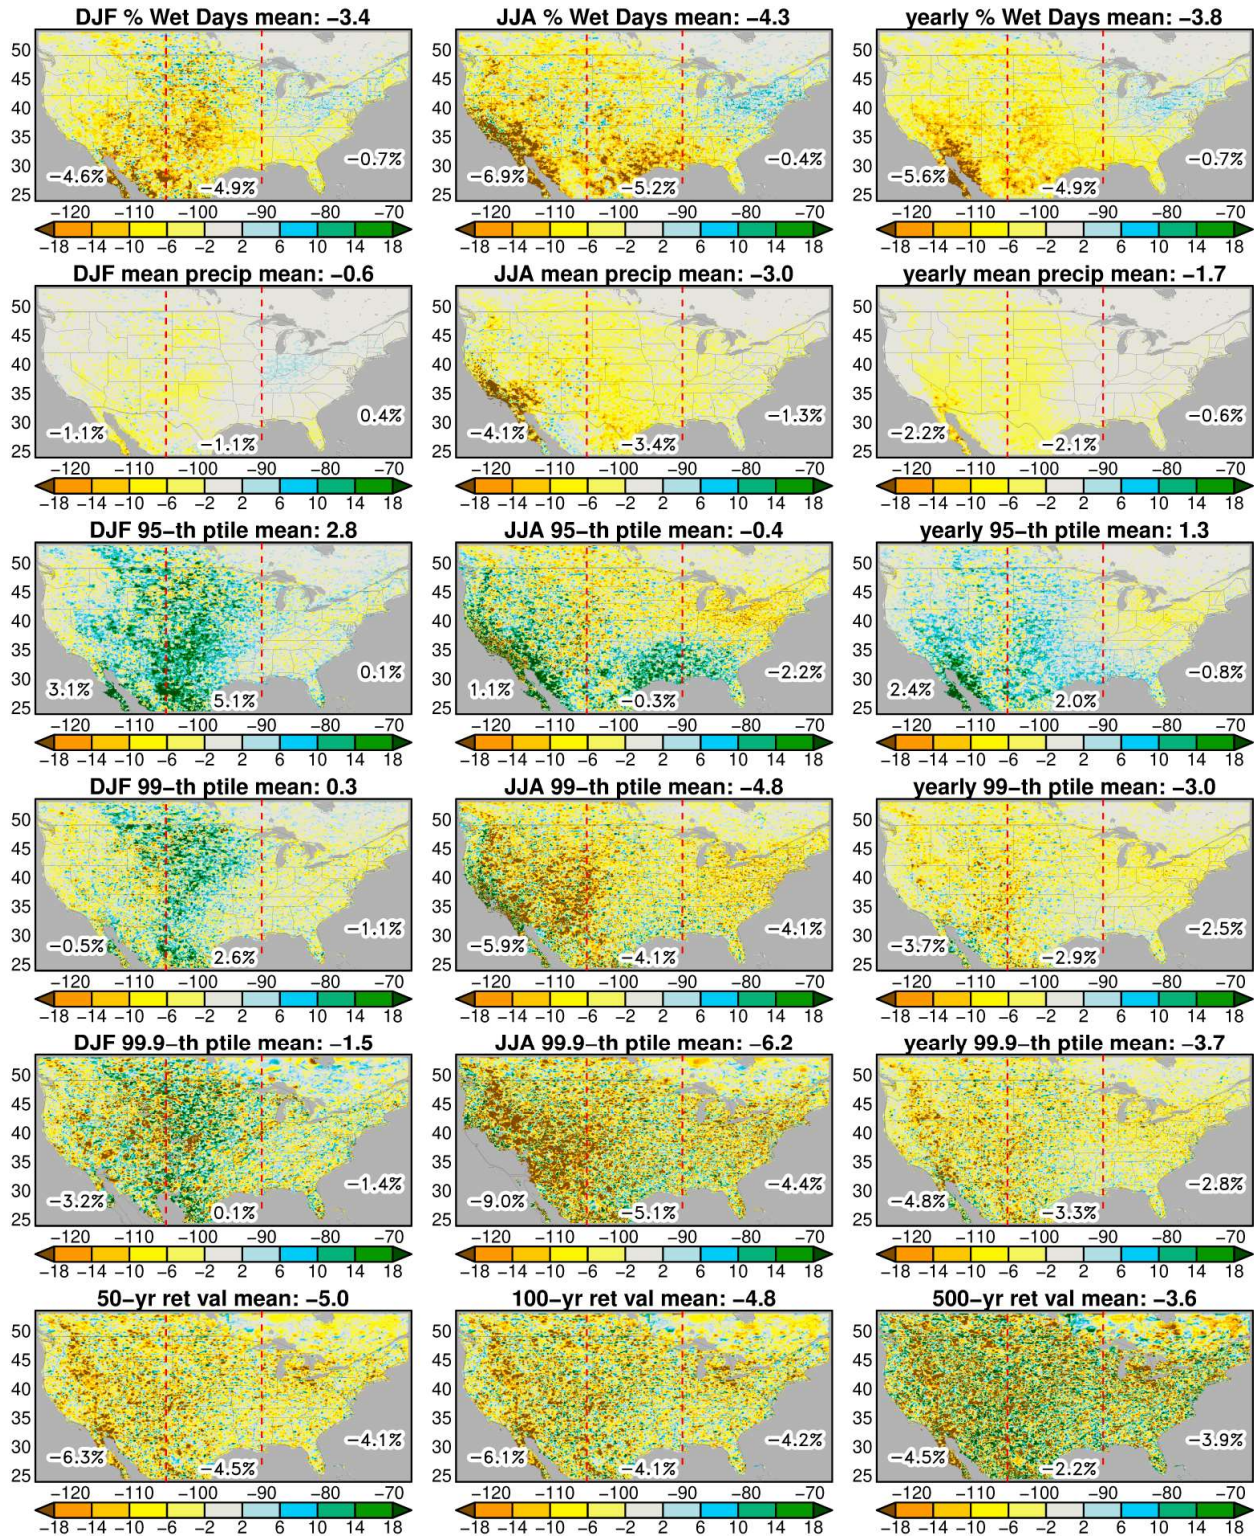


Figure 1. Measures of precipitation obtained from the training data set (Pierce et al. 2021). Top row: Percent of wet days for winter (Dec-Jan-Feb), summer (Jun-Jul-Aug), and yearly. Second row: mean precipitation in winter, summer, and yearly. Third through fifth rows: the 95<sup>th</sup>, 99<sup>th</sup>, and 99.9<sup>th</sup> percentile values of wet day precipitation in winter, summer, and yearly, respectively. Bottom row: the 50-, 100-, and 500-year return values of daily precipitation. All units are mm/day except for the percent of wet days. The numbers at the bottom of each panel show the mean in the western, central, and eastern part of the domain, as demarked by the vertical dashed red lines; the mean in the entire domain is shown in the panel title.



net/valve22/data/LOCA2\_CONUS\_analysis/all\_data\_proc/src/mos\_corrector/diag\_panels.R Tue Sep 17 14:07:13 2024

Figure 2. For measures presented in Figure 1, the error (%) with respect to the training data found in the originally released LOCA2 precipitation data set, version v20220519. All values shown are the multi-model ensemble average across all 27 models.

The multi-model ensemble average (MMEA) errors<sup>2</sup> with respect to the training data in the first version of the LOCA2 precipitation data set are shown in Figure 2. Values are in percent. The fraction of wet days is biased low in summer, particularly in the southwestern U.S. (top row). The yearly means are about 2% low over the year, with the largest shortfall found in the southwestern U.S. in the summer, most notably in a band stretching northwards from the Gulf of California, through Nevada, and into eastern Oregon and Washington. The primary reason for this is the incidence of wet days in the original LOCA2 data being lower than the observational training in the same region and season.

Percentile values are close to the observational training data in winter but somewhat dry in summer. At the most extreme percentile considered here (99.9), the average over the domain is 6.2% lower in LOCA2 than in the observational training data, and 9.0% lower than observed in the western part of the domain. The return values in LOCA2 are likewise lower than those in the observational training data, with domain average errors of about -5% for return periods of 50 and 100 years, moderating slightly to a domain average error of -3.6% for the 500 year return value. For all three of the return periods considered, the LOCA2 values are biased low relative to the observational training data in the western part of the domain, reaching -6.3% for a 50 year return value. Our objective in the new version of the data set is to ameliorate these low biases in the extremes relative to the training data without appreciably degrading the representation of the seasonal and yearly mean values.

## Origin of the low bias in fraction of wet days

Although our focus here is on the precipitation extremes, Figure 2 shows errors in LOCA2 in the representation of the fraction of wet days, particularly in summer, that are worth evaluating. The LOCA process is described in Pierce et al. 2014, 2015, and 2023, but understanding the reason for the error in fraction of wet days that can be seen in Figure 2 requires delving into some of the details of how LOCA downscaling works. As shown in Figure 3, the standard LOCA workflow involves two iterations of bias correction and spatial downscaling; the first iteration downscales from the original coarse-scale global climate model (GCM) grid, which varies by model, to a common 0.5x0.5 degree latitude-longitude grid. The second iteration downscales from the common 0.5 degree grid to the final 16<sup>th</sup>-degree grid.

The reason for having two iterations and an intermediate common 0.5 degree grid is so that the bias correction information applied to the models before the final spatial downscaling step is the same whether the original GCM's spatial resolution is very coarse or relatively fine (some GCMs use a 2.5 degree latitude-longitude grid, while others have a spatial resolution of ~0.7 degrees). If the downscaling were completed in only one iteration, then the coarse-resolution models would be receiving appreciably less bias correction information than the fine-scale models, putting the coarse-resolution models at a disadvantage in the realism of their downscaled result. On the other

---

<sup>2</sup> Errors are computed for each realization, then averaged across all the realizations for each model to form the model's ensemble average error field. The model ensemble average error fields are then averaged across all models to form the multi-model ensemble average. This approach prevents models with more ensemble members from being weighted more strongly in the multi-model ensemble average.

hand, the reason we do not simply bilinearly interpolate all models to the common 0.5 degree grid before the LOCA process, and instead have two complete iterations of the LOCA process, is because our experiments showed that bilinear interpolation of the precipitation fields gives models an unrealistically large spatial coherence, especially for models that were originally on a coarse scale grid. An overly large spatial coherence is detrimental to the simulation of flooding, since downstream flooding is sensitive to the catchment area over which rainfall is occurring simultaneously. Distorted spatial coherence in the precipitation field will therefore distort the flooding statistics.

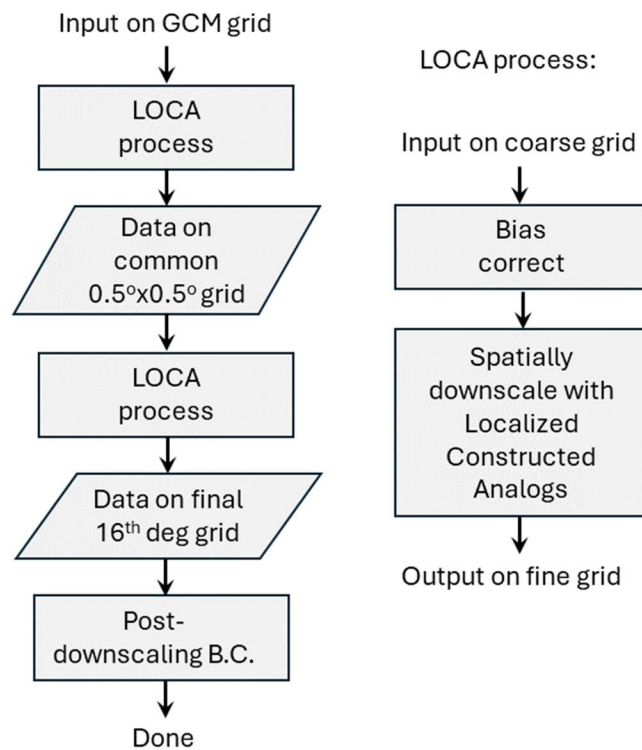


Figure 3. Flowchart of the entire downscaling workflow, transforming data from the original GCM grid to the final 16<sup>th</sup> degree output. Left column: the overall scheme includes two iterations of what we call here the “LOCA process.” The first iteration downscales from the original GCM grid to a common 0.5 degree grid. The second iteration downscales from the common 0.5 degree grid to the final 16<sup>th</sup> degree grid. Right column: the components that make up one iteration of the “LOCA process.”

When precipitation is being downscaled, an important part of the LOCA process is to correct for model biases in the fraction of wet days. It is common for GCMs to have a so-called “drizzle problem,” where too many wet days are predicted by the model (e.g., Chen et al. 2021). This model bias is removed during the bias correction step in the LOCA process shown in the right column of Figure 3 by calculating, at each point, a threshold that will make the model’s fraction of wet days match the observed fraction of wet days.

There are some subtleties involved in the wet day threshold calculation. We use a fixed precipitation threshold of 0.5 mm/day on the 16<sup>th</sup> degree grid as our definition of a “wet day” in the original observational data set. However, the entire downscaling workflow (Figure 3) considers

multiple spatial scales (16<sup>th</sup> degree, 0.5 degree, and original GCM grid), and the observed fraction of wet days depends on the spatial scale. The larger the area that is examined, the lower the chance that all gridcells in that domain have zero precipitation. To calculate the fraction of wet days at larger spatial scales we first apply the 0.5 mm/day threshold to the observations on the 16<sup>th</sup> degree grid: any observed days with precipitation less than 0.5 mm/day are set to zero. The fraction of wet days on the 16<sup>th</sup> degree grid is then the fraction with precipitation > 0. The observations are then spatially aggregated to the 0.5 degree grid, and the fraction of wet days on the 0.5 degree grid computed by simply counting the number of days that have precipitation greater than zero. The 0.5 degree data are then aggregated to the GCM grid, and the fraction of wet days on the GCM grid again obtained by simply counting the number of days with precipitation greater than zero.

At each bias correction step (right column of Figure 3), a threshold is calculated that makes the model fraction of wet days match the observed fraction of wet days in that gridcell. Because of the drizzle problem, the model threshold is typically greater than zero to make the model's fraction of wet days match the observations. However, some models produce fewer wet days than are observed even when a threshold of zero is applied to the model. This situation is rare but does occur in the Southwestern U.S. during the summer months. This is not a correctable model error, since there is no practical way of adding wet days to a model that that does not have enough to begin with.

Our examination of the weak extreme precipitation issue uncovered an implication of the LOCA workflow (Figure 3) that we had not previously appreciated. Any gridcell on the GCM grid that is set to zero precipitation from the first thresholding process means that none of the finer scale gridcells that fall within that coarse gridcell can have a precipitation value other than zero<sup>3</sup>. In other words, the fraction of wet day threshold process on the coarse grid is a one-way operation such that no enclosed finer resolution gridcell can be set to anything but zero precipitation. This biases the fraction of wet days low.

This is different from temperature or non-zero precipitation (exclusive of the thresholding issue), where bias correcting the coarse gridcell does not prevent finer-scale gridcells that fall in that domain from taking on values either greater than or less than the value of the coarse gridcell. For instance, imagine downscaling temperature, and the coarse gridcell value is bias corrected to 10°C. After downscaling, finer scale gridcells enclosed by the coarse scale gridcell will have temperatures distributed around 10°C – some warmer, some cooler. However, once a coarse scale gridcell is set to zero precipitation based on the fraction of wet day thresholding process, no finer-scale gridcell that falls in that domain can have a non-zero precipitation.

An equivalent way of describing this issue is that picking a threshold to make the fraction of wet days match observations on the coarse grid results in the *average* fraction of wet days matching the average fraction in the domain of the coarse gridcell. However, some locations within that coarse

---

<sup>3</sup> There is a caveat to this statement. A 16<sup>th</sup> degree gridcell on the edge of a 0.5 degree gridcell could be influenced by an analog match in the immediately adjacent 16<sup>th</sup> degree gridcell belonging to a different 0.5 degree gridcell. See the discussion of “edge cells” in Pierce et al. 2014. However, this is not the common case.

gridcell might require a lower threshold, while others may require a higher threshold. This is not recoverable after the coarse gridcell precipitation value is set to zero.

A bias in the fraction of wet days is a fundamental problem in a bias correction process, since one cannot simultaneously have correct percentile values of wet day precipitation and seasonal or annual mean precipitation if the number of wet days is different from observed. For example, if the number of wet days is lower than observed, bias correcting the wet day percentile values will result in a mean that is too low. Bias correcting the seasonal or annual mean to match observations will result in percentile values that are incorrect. Exactly what distortions of precipitation's observed distribution occur in any particular model dataset depend on the order of operations and implementation of the bias correction process. It can be seen from Figure 2 that the low bias in fraction of wet days in the southwestern U.S. during summer corresponds to the low bias in seasonal and, to a lesser degree, annually averaged precipitation in this region.

The description above suggests that one way to address this problem could be to omit the fraction of wet day bias correction in the early part of the LOCA workflow, so that sufficient wet days are passed through to the final bias correction step. However, three issues discouraged us from using this approach. 1) as noted above, some models have too few wet days to begin with in certain locations and seasons; nothing in the LOCA framework can bias correction a situation where there are too few wet days, since wet days cannot be created out of dry days. 2) such an approach would require implementing a new post-downscaling bias correction; the reasons for this are discussed below. 3) implementing such a fix would require an entire re-run of the LOCA workflow from the beginning, which was not possible here due to resource constraints. We therefore focused on step 2, implementing a new post-downscaling bias correction that can be applied to the current precipitation data, which is both immediately useful in addressing the weak extreme precipitation values in the existing data set and would be necessary in a future version of LOCA that implemented the wet day fraction correction described above.

## New post downscaling bias correction

Post downscaling bias correction was explored and ultimately implemented to address the weakness in extreme precipitation described above, which is our primary concern here. It is not unusual for statistically downscaled data sets to have a final bias correction step at the end of the spatial downscaling process; for example, MACA statistical downscaling (Abatzoglou and Brown, 2012) works this way. LOCA has traditionally used a mild post downscaling bias correction step (Figure 3), which nudges the final product towards the observed seasonal means on large spatial scales ( $> 2^\circ$ ). However, if the fraction of wet day thresholds were changed to allow all wet days through the bias correction and spatial downscaling process, as described above, then the resultant distortions to the distribution of precipitation would require a more complete post downscaling bias correction. Implementing a new post downscaling bias correction would therefore be both immediately useful for addressing the weak extreme precipitation values and potentially be required for a future version of LOCA that modified the treatment of wet day fraction.

Since the LOCA process already has a bias correction step (right column of Figure 3; described in Pierce et al. 2015), it is natural to think of using this existing code to perform the post downscaling bias correction. Our tests in trying this approach were not successful, for the following reasons. 1)



The existing bias correction is computed on the coarse-scale input grid, so memory issues are not a concern. For example, the volume of data on the final 16<sup>th</sup> degree grid is 64 times larger than on the 0.5 degree input grid. The low memory demands of implementing the bias correction on the coarse input grid mean that the precipitation percentile values needed for the bias correction can be computed from the entire data set that is read into memory and sorted. This is not possible on the final 16<sup>th</sup> degree output grid due to compute node memory limits on the resources we had available for this work. 2) Our attempts to limit the memory footprint of the existing bias-correction process by reading only a longitude-time slice of the data at a time proved to be too slow to be practical.

The bias correction approach used in LOCA2 is PresRat (Pierce et al. 2015), so named because it preserves the GCM-predicted ratio of future precipitation changes by quantile. Quantile mapping bias correction is used over the historical period (1950-2014), and GCM-predicted change factors, by quantile, are used for the future period. The bias correction is implemented iteratively with seasonal window widths of 3, 6, and 12 months; see Pierce et al. 2015 for further details.

For this new release of the precipitation data set we implement a post downscaling bias correction version of PresRat in an approximate way by pre-computing tables of wet-day precipitation percentile values and (at the high extremes) return values at a limited number of percentiles and return periods. Specifically, we computed 23 wet day precipitation percentile values, by month, at  $p=0.1, 0.5, 1, 2, 5, 10, 20, \dots, 80, 85, 90, 95, 97, 98, 99, 99.25, 99.75, 99.9$ . Coverage increases at the ends, and particularly at the high end, due to our interest in improving the representation of extreme high values. In locations and months with insufficient samples to calculate the high percentile values (for example, in the dry Southwestern U.S. during summer), the temporal window was expanded from 1 month to 3 months to provide more samples. When even this wider window yielded insufficient samples for computation of the percentiles, we used the GCM predicted quantile change factors at the highest percentile that had sufficient data to be computed. Annual return values of daily precipitation were calculated using the method of L-moments (Hosking 1990) based on yearly block maxima. Return values were calculated at return periods of 2, 3, 4, 5, 7, 10, 15, 20, 30, 50, 70, 100, 150, 200, 250, 300, 400, ... 1000 years. As an additional resource saving measure the post downscaling bias correction was implemented with only two iterations, using window widths of 1 and 12 months, rather than three. Return values are only bias corrected using the annual window, not the monthly window.

Precomputing the precipitation wet day percentile value and return value tables allows us to apply the bias correction using a table lookup approach with linear interpolation between the compiled points. This allows efficient implementation of the bias correction without excessive memory demand once the tables have been computed. Residual errors are apparent that arise from using a fast but approximate linear interpolation between the computed values; these are addressed by applying a final model output statistics (MOS) corrector (e.g., Glahn and Lowry, 1972; Wilks 2006; Eden and Widmann, 2014). In brief, the MOS corrector works by computing the overall downscaled error with respect to observations, then applying a (in this case) multiplicative factor to reduce this error. In our case, we compute the overall model error using the multi-model ensemble average of the error with respect to the training data, in keeping with our approach of bias correcting the model ensemble average to the training data, not each individual model realization (Pierce et al. 2023). This approach better preserves natural internal climate variability across the model realizations.

## Results

The MMEA error fields after the new post-downscaling bias correction is applied are shown in Figure 4. The winter mean precipitation is nearly the same (one percentage point lower) as found in the original data, averaged over the domain, with the greatest decrease being about one percentage point in the western U.S. in winter. However, errors are reduced in summer by 1.1, 2.2, and 0.7 percentage points in the western, central, and eastern parts of the domain, respectively. In the yearly average, the errors for the new data set are within 0.2 percentage points of the original data errors, except for in the central part of the domain, where the error reduced from -2.1% to -1.4%. Overall we consider this mix of small changes, most of which are improvements, to have satisfied our objective of not making the representation of seasonal and annual mean precipitation appreciably worse.

Of more immediate interest to the current work is the behavior at the precipitation extremes. Jumping directly to the 99.9<sup>th</sup> percentile of seasonal and annual precipitation (fifth row of Figure 4), the new post downscaling bias correction process results in a clear improvement in the representation of extreme precipitation and, in particular, elimination of the undesirable weakness in extreme precipitation seen in Figure 2. The largest 99.9<sup>th</sup> percentile shortfall in the original data set was -9.0% in the western U.S. during the summer; the same region and season in the revised data set has an error of only 0.4%. The summer domain mean was -6.2% previously, but in the new data set is much smaller at 0.6%. Improvements were also obtained in winter, although with more modest values because the errors were smaller in winter to begin with. While the original data set showed errors of -3.2%, 0.1%, and -1.4% in the western, central, and eastern parts of the domain in winter, the new data set shows errors of 1.1%, 1.6%, and 0.7%. In the yearly average, errors were originally -4.8%, -3.3%, and -2.8% (west, central, east), but are 1.3%, 0.8%, and 0.5% in the new version.

Errors in the most extreme precipitation values considered here are shown in the bottom row of Figure 4. In the original data set, the largest errors in return value were seen at a return period of 50 years, with a domain mean of -5.0% and subdomain averages of -6.3%, -4.5%, and -4.1% (west, central, east). In the revised data, the 50-yr return value domain mean error is greatly reduced to only 0.3%, and the subdomain averages are 0.6%, 0.6%, and -0.3%, indicating that the weakness in extreme precipitation is nearly eliminated at this return value. The domain averaged errors in the original data set at 100- and 500-year return periods were -4.8% and -3.6% respectively. These have been reduced to 0.5% and 1.7% in the revised data set.

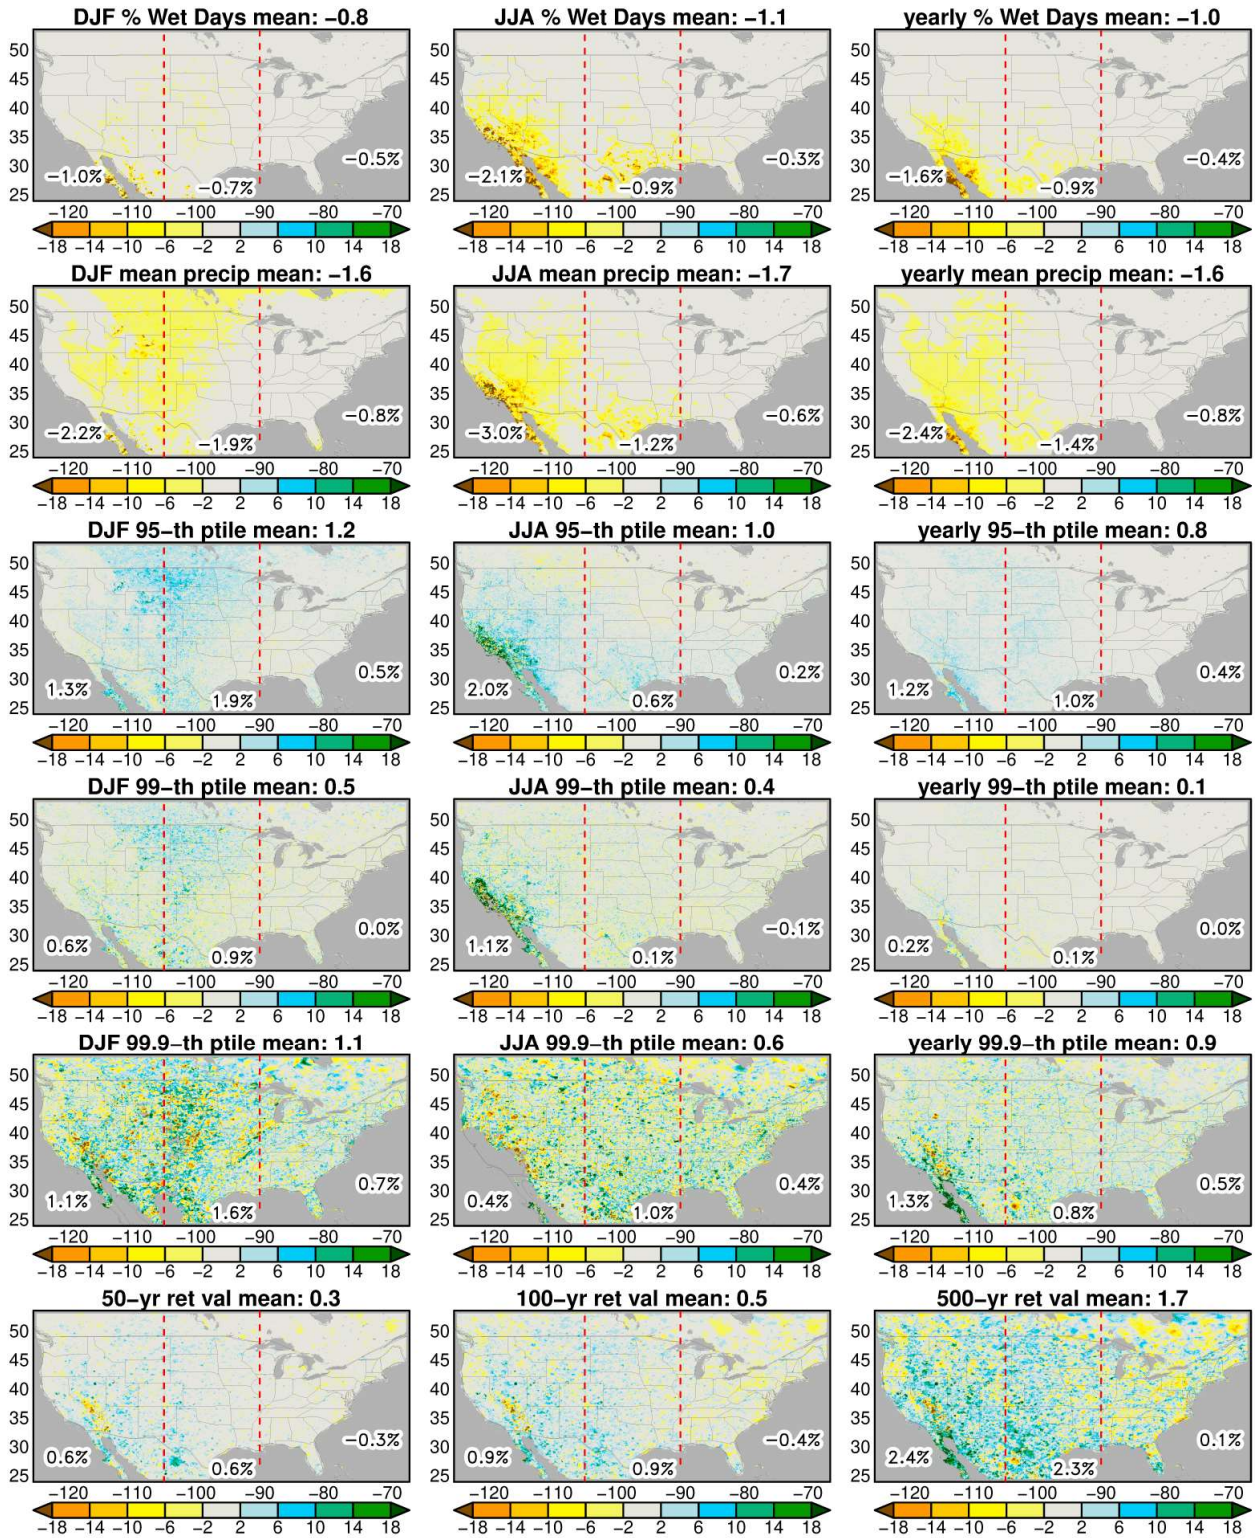
Based on these results, we consider the new approach to the post downscaling bias correction to have substantially reduced, and in many cases very slightly reversed, the original weakness in extreme precipitation values, while still maintaining the level of fidelity in the seasonal and annual means seen in the original data set.

## Summary

The first version (v20220519) of the LOCA2 precipitation data set showed weaker than observed daily precipitation extremes on the order of 5-10% in the multi-model ensemble average, with deficiencies greatest in the western U.S. and in summer, most notably in an interior region from the

Gulf of California northwards through Nevada and into eastern Oregon and Washington. The largest domain-averaged shortfalls were seen at the 99.9<sup>th</sup> percentile value of daily precipitation but were evident in the yearly data at the 99<sup>th</sup> percentile as well, and at return periods from 50 to 500 years.

To address these shortfalls we implemented a new version of the post downscaling bias correction process that was developed to accommodate the memory demands of bias correcting on the full fine spatial resolution (16<sup>th</sup> degree) output grid. It achieved this by using pre-computed model values at a range of percentiles and return values rather than by trying to compute these values in memory on the fly during processing. Evaluation of the final result shows a good improvement in the extreme precipitation values, with a near-complete elimination of low extreme precipitation values at the 99<sup>th</sup> and 99.9<sup>th</sup> percentile values and at return periods of 50- to 500-years. Domain averaged errors in the revised data set, termed version v20240915, are slightly positive and in the range of 0.1 to 1.7% over the percentiles (p95 to p99.9) and return periods (50-, 100-, and 500-years) of daily precipitation examined here.



/home6/dwpierc2/l2r/all\_data\_proc/src/mos\_corrector/diag\_panels.R Tue Sep 17 14:11:27 2024

Figure 4. Multi-model ensemble averaged error with respect to the training data for the new release of the precipitation data set, version v20240915.

## Acknowledgements

This work was funded by a contract from the Strategic Environmental Research and Development Program (SERDP), Project RC19-1391. Additional support for DWP and DRC was provided by the California Energy Commission under Contract EPC-20-006, and from the California Department of Water Resources Atmospheric River Program Phase III under Grant 4600014294. The NASA High-End Computing Capability (HECC) Program provided resources supporting this work through the NASA Earth Exchange (NEX), Earth Science Division, and the NASA Advanced Supercomputing (NAS) Division at Ames Research Center. We thank the NEX staff for their technical support. We acknowledge the World Climate Research Programme, which, through its Working Group on Coupled Modelling, coordinated and promoted CMIP6. We thank the climate modeling groups for producing and making available their model output, the Earth System Grid Federation (ESGF) for archiving the data and providing access, and the multiple funding agencies who support CMIP6 and ESGF.

## References

- Abatzoglou, J. T., and T. J. Brown, 2012: A comparison of statistical downscaling methods suited for wildfire applications. *Int. J. Climatol.*, v. 32, 772-780.
- Chen, D., A. G. Dai, and A. Hall, 2021: The convective-to-total precipitation ratio and the “drizzling” bias in climate models. *J. Geophys. Res. Atmos.*, v. 126, e2020JD034198. DOI:10.1029/2020JD034198.
- Eden, J. M., and M. Widmann, 2014: Downscaling of GCM-simulated precipitation using model output statistics. *J. Clim.*, v. 27, 312-324.
- Glahn, H. R., and D. A. Lowry, 1972: The use of model output statistics (MOS) in objective weather forecasting. *J. Appl. Met. Climatol.*, v. 11, 1203-1211.
- Hosking, J. R. M., 1990: L-moments: analysis and estimation of distributions using linear combinations of order statistics. *J. R. Statist. Soc. B*, v. 52, 105-124.
- Livneh, B., T. J. Bohn, D. W. Pierce, F. Munoz-Arriola, B. Nijssen, R. Vose, D. R. Cayan, and L. Brekke, 2015: A spatially comprehensive, hydrometeorological data set for Mexico, the U.S., and Southern Canada 1950-2013. *Scientific Data*, v. 2, article 150042 (2015). doi:10.1038/sdata.2015.42.
- Pierce, D. W., D. R. Cayan, and B. L. Thrasher, 2014: Statistical downscaling using localized constructed analogs (LOCA). *J. Hydrometeorology*, v. 15, p. 2558, doi:10.1175/JFM-D-14-0082.1.
- Pierce, D. W., D. R. Cayan, E. P. Maurer, J. T. Abatzoglou, and K. C. Hegewisch, 2015: Improved bias correction techniques for hydrological simulations of climate change. *J. Hydrometeorology*, v. 16, p. 2421-2442. DOI: <http://dx.doi.org/10.1175/JHM-D-14-0236.1>.
- Pierce, D. W., L. Su, D. R. Cayan, M. D. Risser, B. Livneh, and D. P. Lettenmaier, 2021: An Extreme-Preserving Long-Term Gridded Daily Precipitation Dataset for the Conterminous United States. *Journal of Hydrometeorology*, v. 22, p. 1883-1895. <https://doi.org/10.1175/JHM-D-20-0212.1>.

Pierce, D. W., D. R. Cayan, D. R. Feldman, and M. D. Risser, 2023: Future increases in North American extreme precipitation in CMIP6 downscaled with LOCA. *J. Hydrometeorology*. DOI: 10.1175/JHM-D-22-0194.1

Ullrich, P. A., 2023: Validation of LOCA2 and STAR\_ESDM statistically downscaled products. Technical report LLNL-TR-856094, 19 pp. Available at <https://www.osti.gov/servlets/purl/2202926>.

Wilks, D. S., 2006: Statistical methods in the atmospheric sciences. Second edition. Elsevier, 627 pp.