Detected Changes in Precipitation Extremes at Their Native Scales Derived from In Situ Measurements

MARK D. RISSE

Lawrence Berkeley National Laboratory, Berkeley, California

CHRISTOPHER J. PACIOREK

University of California, Berkeley, Berkeley, California

TRAVIS A. O'BRIEN, MICHAEL F. WEHNER, AND WILLIAM D. COLLINS

Lawrence Berkeley National Laboratory, Berkeley, California

(Manuscript received 29 January 2019, in final form 12 August 2019)

ABSTRACT

The gridding of daily accumulated precipitation—especially extremes—from ground-based station observations is problematic due to the fractal nature of precipitation, and therefore estimates of long period return values and their changes based on such gridded daily datasets are generally underestimated. In this paper, we characterize high-resolution changes in observed extreme precipitation from 1950 to 2017 for the contiguous United States (CONUS) based on in situ measurements only. Our analysis utilizes spatial statistical methods that allow us to derive gridded estimates that do not smooth extreme daily measurements and are consistent with statistics from the original station data while increasing the resulting signal-to-noise ratio. Furthermore, we use a robust statistical technique to identify significant pointwise changes in the climatology of extreme precipitation while carefully controlling the rate of false positives. We present and discuss seasonal changes in the statistics of extreme precipitation: the largest and most spatially coherent pointwise changes are in fall (SON), with approximately 33% of CONUS exhibiting significant changes (in an absolute sense). Other seasons display very few meaningful pointwise changes (in either a relative or absolute sense), illustrating the difficulty in detecting pointwise changes in extreme precipitation based on in situ measurements. While our main result involves seasonal changes, we also present and discuss annual changes in the statistics of extreme precipitation. In this paper we only seek to detect changes over time and leave attribution of the underlying causes of these changes for future work.

1. Introduction

In spite of the increasing prevalence of high spatial resolution climate models and gridded daily products, it is important to utilize measurements from weather stations to characterize changes in the climatology of precipitation. This is especially true when considering changes in extreme precipitation, for which gridded data products may contain biases due to the fact that daily precipitation exhibits fractal scaling (e.g., Lovejoy et al. 2008; Maskey et al. 2016, and numerous references

therein) and any smoothing technique will minimize extreme values. A recent thread of research documents these inherent problems with gridded data products and their use for characterizing climatological extremes (see, e.g., King et al. 2013; Gervais et al. 2014; Timmermans et al. 2019; Risser et al. 2019b). As a result, changes in extreme precipitation calculated from gridded data products are potentially misleading. Without using gridded products, however, it is difficult to characterize spatially coherent changes in extremes on the scales relevant for assessing impacts and designing infrastructure.

There is a large existing literature on changes in extreme precipitation based on observations; these results fall into three general categories: 1) changes derived from daily gridded precipitation products; 2) changes

---

Denotes content that is immediately available upon publication as open access.

Corresponding author: Mark D. Risser, mdrisser@lbl.gov

based on station data, but only at the station locations; or 3) changes based on station data, but aggregated into large areal averages. Across these different methods for presenting results, the literature explore changes in a variety of extremal metrics, for example the annual total precipitation divided by frequency of rainy days, the number of days that exceed a particular threshold or percentile, or return levels for a fixed duration.

The Intergovernmental Panel on Climate Change (IPCC) Fifth Assessment Report (AR5; [Hartmann et al. 2013](#)) presents annual changes in extreme precipitation based on the HadEX2 dataset ([Donat et al. 2013](#)), which is a coarsely gridded (2.5° latitude by 3.75° longitude) dataset produced by calculating metrics of extreme precipitation and temperature at individual stations and then interpolating to a common grid. Figure 2.33 in the AR5 ([Hartmann et al. 2013](#)) shows changes in the annual amount of precipitation from days greater than the 95th percentile as well as daily precipitation intensity. Both of these metrics exhibit increases over the central and eastern United States, some of which are *very likely* significant, as well as slight decreases over the west and southwest United States. Overall, the AR5 finds that there are more regions with increases in the number of heavy precipitation events than decreases with the caveat that there are strong regional and seasonal differences, but in general it is difficult to ascribe a high degree of confidence in the magnitude and direction of the change.

While the AR5 analyses are based upon a gridded product, [Kunkel \(2003\)](#) explores changes in the United States using daily weather station measurements from the Global Historical Climatology Network (GHCN; [Menne et al. 2012](#)) from the late nineteenth through twentieth centuries. He concludes that overall the United States experienced an upward trend in extreme precipitation events since the 1920s and 1930s, but that extreme events in the late 1800s were as frequent as in the 1980s and 1990s. However, [Kunkel \(2003\)](#) notes that natural variability in the climate system could be the cause of the observed increase over the last century. [Balling and Goodrich \(2011\)](#) also use the GHCN data to describe annual changes in total precipitation over 1975–2010 at individual weather stations, identifying a general reduction in precipitation in the western United States and a general increase in the central and northeastern portions of the country. However, [Balling and Goodrich \(2011\)](#) find that trends in precipitation display a large degree of spatial entropy or disorder wherein nearby stations show different or even opposite changes, and that the changes can differ for the same location depending on the metric and period of study applied to the station data. [Kunkel et al. \(2013\)](#) use GHCN

measurements to estimate changes during 1948–2010 in the 20-yr return value at the locations of the GHCN stations, finding that approximately three-fourths of all stations experience an increase in extreme precipitation but only 15% experience a statistically significant increase. The central United States and North Atlantic show spatially coherent upward changes, while the northwest United States shows less spatially coherent downward changes.

Finally, many studies evaluate changes for large areal regions, where the changes are estimated from daily weather measurements that are first aggregated spatially. [Groisman et al. \(2005\)](#) use daily total precipitation datasets compiled at the National Climatic Data Center (including the GHCN) to evaluate changes in intense precipitation globally and find areas with both increases and decreases across the globe. Interestingly, they note that “to obtain statistically significant estimates, the characteristics of heavy precipitation should be areally averaged over a spatially homogeneous region. Otherwise, noise at the spatial scale of daily weather systems masks changes and makes them very difficult to detect” ([Groisman et al. 2005](#), p. 1327). [Janssen et al. \(2014\)](#) explore trends in the GHCN data over 1900–2012 for seven large subregions of the contiguous United States (CONUS) and obtain an overall increase in extreme precipitation. Statistically significant increasing trends in extreme precipitation at the 95% confidence level occur over the Midwest and Eastern regions, while most decreasing trends occur over Western regions, consistent with the results of [Kunkel \(2003\)](#) and [Groisman et al. \(2005\)](#). [Easterling et al. \(2017\)](#) evaluate seasonal changes in the 20-yr return value based on the GHCN measurements over 1948 to 2015 and find major regional differences in the changes with the largest increases occurring in the northeastern United States. [Easterling et al. \(2017\)](#) identify the largest positive changes in fall and winter when all but one of seven large CONUS subregions exhibit increases but do not indicate where and when a degree of confidence might be ascribed to the changes.

In summary, we reiterate that the literature contains a diverse set of data analyses focusing on observed changes in extreme precipitation over CONUS and that most of these come to common conclusions about the sign of the changes. Of the aforementioned studies, only [Donat et al. \(2013\)](#) present spatially continuous trends based on irregularly spaced weather station observations; however, they do so on a very coarse grid via a distance-based weighting algorithm ([Donat et al. 2013](#)). Among the remaining studies, many fail to ascribe a confidence statement to the results. Furthermore, many of these analyses only evaluate annual changes, which ignores important seasonal differences in extreme

precipitation that should be accounted for when characterizing changes.

In this paper, we characterize changes in observed extreme precipitation for a fine 0.25° spatial grid over the contiguous United States for 1950–2017 using station data only, utilizing spatial statistical methods to derive gridded changes based on irregularly observed measurements. Furthermore, we use a robust statistical technique to identify significant pointwise changes in the climatology of extreme precipitation while also carefully controlling the rate of false positives. The use of spatial statistics yields both an increased signal-to-noise ratio for the changes as well as insight into the physical behavior of changes in the climatology of extreme precipitation. While we conduct separate analyses for each season, we also derive annual changes in extreme precipitation based on the seasonal results.

To be clear, the results presented here only seek to detect the presence of changes in extreme precipitation and specifically do *not* seek to characterize temporal trends or attribute the underlying causes of these changes. In characterizing smooth (linear) trends over time only, we do not intend to account for sources of interannual variability (e.g., El Niño–Southern Oscillation) or decadal to multidecadal sources of variability (e.g., the Pacific decadal oscillation or the Atlantic multidecadal oscillation) that are known to influence precipitation extremes in the contiguous United States. Given the relatively short time period considered in this analysis (1950 to the present), it would be essentially impossible to account for these low-frequency oscillations using observational data only because any changes we identify over the 68 years considered could be due to, for example, the Atlantic multidecadal oscillation (which is roughly monotone over this period). However, given that our methodology provides a better characterization of precipitation extremes and allows for a high-resolution assessment of changes based on station data, we argue that our results are nonetheless an important addition to the literature on empirical changes in extreme precipitation. For future work, we plan to pursue a formal attribution of the underlying causes of any changes in extreme precipitation that may appear in the historical record.

The paper proceeds as follows: in [section 2](#), we describe the weather station dataset used in our analysis, and in [section 3](#) we describe our statistical methodology. Results are presented in [section 4](#), and [section 5](#) concludes the paper.

2. Data

The data used for the following analysis consist of daily measurements of total precipitation (in millimeters)

FIG. 1. The spatial distribution of the $n = 5202$ GHCN stations with a minimum of 66.7% of daily precipitation measurements over December 1949–November 2017.

obtained from the Global Historical Climatology Network (GHCN; [Menne et al. 2012](#)) over the contiguous United States (CONUS). We refer the reader to [Risser et al. \(2019b\)](#) for details on the quality control procedure. After preprocessing the daily values, we select the subset of stations that have a minimum of 66.7% nonmissing daily precipitation measurements between 1 December 1949 through 30 November 2017. This procedure yields a high-quality set of daily precipitation measurements spanning $T = 68$ years for $n = 5202$ stations shown in [Fig. 1](#) down-selected from the entire GHCN database of over twenty thousand weather stations. All subsequent analysis in this paper is based on seasonal maxima $\{Y_t^{(j)}(\mathbf{s}): t = 1950, \dots, 2017; j \in \{\text{DJF}, \text{MAM}, \text{JJA}, \text{SON}\}; \mathbf{s} \in \mathcal{S}\}$, where $\mathcal{S} = \{\mathbf{s}_1, \dots, \mathbf{s}_n\}$ denotes the $n = 5202$ stations shown in [Fig. 1](#). Note that the year index t represents a “season year” such that, for example, the 1950 DJF is defined as December 1949 to February 1950.

3. Statistical methods

The various components of the statistical methods used in our analysis are now described. Furthermore, the corresponding workflow is displayed schematically in [Fig. 2](#) with specific pointers for each of the following steps.

a. Spatial extreme value analysis

To characterize the spatially complete climatological distribution of extreme precipitation based on measurements from irregularly observed weather stations, we apply the spatial data analysis outlined in [Risser et al. \(2019b\)](#). In the workflow shown in [Fig. 2](#), this analysis is represented by the green diamond labeled “Spatial GEV analysis.” An important feature of the analysis is

FIG. 2. A workflow diagram for the analysis described in section 3b. In the diagram, dark blue boxes represent raw input data; purple boxes represent derived datasets; green diamonds represent operations (e.g., “Spatial GEV analysis” refers to section 3a); red circles represent an output. The numbered light blue boxes are pointers to specific components of the analysis and are referenced throughout section 3b. In the gray boxes, steps I–III correspond to the subsections of section 3b.

that it allows one to estimate the distribution of extreme precipitation even for locations where no data are available. Furthermore, the methodology proposed in Risser et al. (2019b) can be used for a large network of weather stations over a heterogeneous spatial domain like CONUS, which is critical for the problem at hand.

For a full description of the methodology used, we refer the reader to Risser et al. (2019b). In short, the essence of the method is to first obtain estimates of the climatological features of extreme precipitation based on measurements from the weather stations via the generalized extreme value (GEV) family of distributions, which is a modeling framework for the maxima of a process over a prespecified time interval or “block” (e.g., the 3-month seasons used here). Coles [2001; Theorem (3.1.1), p. 48 therein] shows that when the number of daily measurements is large (e.g., when considering the ~ 90 daily measurements in a given season), the cumulative distribution function (CDF) of $Y_t(\mathbf{s})$ (which is the seasonal maximum daily precipitation measurement in year t at station $\mathbf{s}$) is a member of the GEV family

$$G_{s,t}(y) \equiv \mathbb{P}[Y_t(\mathbf{s}) \leq y] \\ = \exp \left(- \left\{ 1 + \xi_t(\mathbf{s}) \left[\frac{y - \mu_t(\mathbf{s})}{\sigma_t(\mathbf{s})} \right] \right\}^{-1/\xi_t(\mathbf{s})} \right), \quad (1)$$

defined for $\{y: 1 + \xi_t(\mathbf{s})[y - \mu_t(\mathbf{s})]/\sigma_t(\mathbf{s}) > 0\}$. The GEV family of distributions (1) is characterized by three space–time parameters: the location parameter $\mu_t(\mathbf{s}) \in \mathcal{R}$, which describes the center of the distribution; the scale parameter $\sigma_t(\mathbf{s}) > 0$, which describes the spread of the distribution; and the shape parameter $\xi_t(\mathbf{s}) \in \mathcal{R}$. The shape parameter $\xi_t(\mathbf{s})$ is the most important for determining the qualitative behavior of the distribution of daily rainfall at a given location. If $\xi_t(\mathbf{s}) < 0$, the distribution has a finite upper bound; if $\xi_t(\mathbf{s}) > 0$, the distribution has no upper limit; and if $\xi_t(\mathbf{s}) = 0$, the distribution is again unbounded and the CDF (1) is interpreted as the limit $\xi_t(\mathbf{s}) \rightarrow 0$ (Coles 2001).

As described in section 1, recall that the goal of this analysis is to characterize how the climatological distribution of extremes changes from 1950 to 2017.

As such, we use the following time trend models for the GEV parameters:

$$\mu_t(\mathbf{s}) = \mu_0(\mathbf{s}) + \mu_1(\mathbf{s})t, \quad \sigma_t(\mathbf{s}) = \sigma(\mathbf{s}), \quad \xi_t(\mathbf{s}) = \xi(\mathbf{s}) \quad (2)$$

(following, e.g., Westra et al. 2013; Risser et al. 2019b, and others). We henceforth refer to $\mu_0(\mathbf{s})$, $\mu_1(\mathbf{s})$, $\sigma(\mathbf{s})$, and $\xi(\mathbf{s})$ as the *climatological coefficients* for location $\mathbf{s}$, as these values describe the climatological distribution of extreme precipitation in each year. Of course, using only four parameters to model $G_{\mathbf{s},t}(y)$ over time for each location is an admittedly simplistic representation given that the true smoothed trend could be nonlinear (due to, e.g., the nonlinear trends in greenhouse gas concentrations). However, we argue that a linear approximation is reasonable, particularly since we evaluate estimated differences in $G_{\mathbf{s},t}(y)$ between the endpoints of the time series (1950 and 2017; see section 3b). We explored the possibility of (i) including nonlinear trends in the location parameter and (ii) further modeling the scale and/or shape parameter as changing linearly over time; however, we found that the model described by (2) performed as well [in a statistical sense (i.e., using the Akaike information criteria); see

Tables B3 and B4 in appendix B for these comparisons] than any of these alternative characterizations. We further discuss nonlinear (quadratic) trends in section 5. For all of these reasons, we argue that (2) is a simplistic but acceptable model for characterizing smooth changes over 1950–2017 in the climatological distribution of extremes over CONUS.

Conditional on station-specific estimates of the climatological coefficients, we next apply a spatial statistical approach using second-order nonstationary Gaussian processes to infer the underlying climatology over a fine grid via kriging. This approach yields fields of best estimates of the climatological coefficients, denoted $\{\hat{\mu}_0(\mathbf{s}'), \hat{\mu}_1(\mathbf{s}'), \hat{\sigma}(\mathbf{s}'), \hat{\xi}(\mathbf{s}'); \mathbf{s}' \in \mathcal{G}\}$, where $\mathcal{G}$ is the 0.25° grid of $M = 13\,073$ grid cells over CONUS. These best estimates can be used to calculate corresponding estimates of the r -yr return value, denoted $\hat{\phi}_t^{(r)}(\mathbf{s}')$, which is defined as the seasonal maximum daily precipitation total that is expected to be exceeded on average once every r years. In other words, $\hat{\phi}_t^{(r)}(\mathbf{s}')$ is an estimate of the $1 - 1/r$ quantile of the distribution of seasonal maximum daily precipitation in year t at grid cell $\mathbf{s}'$, that is, $P[Y_t(\mathbf{s}') > \hat{\phi}_t^{(r)}(\mathbf{s}')] = 1/r$, which can be written in closed form in terms of the climatological coefficients:

$$\hat{\phi}_t^{(r)}(\mathbf{s}') = \begin{cases} [\hat{\mu}_0(\mathbf{s}') + \hat{\mu}_1(\mathbf{s}')t] - \frac{\hat{\sigma}(\mathbf{s}')}{\hat{\xi}(\mathbf{s}')} \{1 - [-\log(1 - 1/r)]^{-\hat{\xi}(\mathbf{s}')}\}, & \hat{\xi}(\mathbf{s}') \neq 0 \\ [\hat{\mu}_0(\mathbf{s}') + \hat{\mu}_1(\mathbf{s}')t] - \hat{\sigma}(\mathbf{s}') \log[-\log(1 - 1/r)], & \hat{\xi}(\mathbf{s}') = 0 \end{cases} \quad (3)$$

(Coles 2001). Note that $\hat{\phi}_t^{(r)}(\mathbf{s}')$, the best estimate of the return value in year t , does not represent a return value estimated from a single year of data but is instead calculated using climatological coefficients estimated from the complete time record with the specific year $[t$ in (3)] plugged in. In other words, the time-varying return value estimates represent temporally smoothed quantities.

b. Quantifying changes in extreme precipitation and significance testing

The spatial extreme value analysis described in section 3a is used to both quantify changes in the climatology of extreme precipitation and conduct significance testing to determine where meaningful changes in the distribution of extreme precipitation have occurred. The workflow described in the next three subsections is displayed schematically in Fig. 2.

1) CHANGE METRICS, TEST STATISTICS, AND NULL HYPOTHESES

First, we determine the metrics for evaluating changes in the return values, the relevant test statistics, and the

null hypotheses of interest. Using the results of the fitted and smoothed GEV distributions, we are interested in identifying areas that have experienced significant changes in the extreme statistics of precipitation (e.g., the return values) over 1950–2017. To quantify these changes, we consider the relative change in 20-yr return values,

$$\Delta^R(\mathbf{s}') \equiv \frac{\phi_{2017}^{(20)}(\mathbf{s}') - \phi_{1950}^{(20)}(\mathbf{s}')}{\phi_{1950}^{(20)}(\mathbf{s}')},$$

where $\phi_t^{(20)}(\mathbf{s}')$ is the 20-yr return value at grid cell $\mathbf{s}' \in \square$ in year t . Alternatively, we also consider the absolute change in 20-yr return values

$$\Delta^A(\mathbf{s}') \equiv \phi_{2017}^{(20)}(\mathbf{s}') - \phi_{1950}^{(20)}(\mathbf{s}').$$

Recall that the return values in $\Delta^R(\mathbf{s}')$ and $\Delta^A(\mathbf{s}')$ refer to temporally smoothed quantities evaluated in 1950 and 2017, not return value estimates specific to each year. Using the linear trend model in (2), note that $\Delta^R(\mathbf{s}')$ depends on the return period r , which is a necessary

implication of only allowing the center of the distribution to shift over time: the relative change depends on where you are in the distribution. On the other hand, $\Delta^A(\mathbf{s}')$ does *not* depend on r : given a constant shift per time in the center of the distribution only, all quantiles change by the same magnitude (in mm). Note that under the alternate parameterization of [van der Wiel et al. \(2017\)](#) (or any other parameterization that allows the scale parameter σ_t to change over time, e.g., using a trend or physical covariates), both the relative and absolute changes depend on the return period r (see [appendix A](#) for further details).

Let $\Delta(\mathbf{s}')$ represent an arbitrary change metric, which we assume to have some fixed but unknown value. Note the distinction between the true but unknown quantity $\Delta(\mathbf{s}')$, a function of the unknown $\phi_t^{(20)}(\mathbf{s}')$, and our estimate of this quantity $\hat{\Delta}(\mathbf{s}')$ based on estimated return values $\hat{\phi}_t^{(20)}(\mathbf{s}')$. (These estimates are represented in the flow diagram in [Fig. 2](#) with pointer 1.) To test for the presence of any change in $\Delta(\mathbf{s}')$, the null and alternative hypothesis for each location is

$$H_0(\mathbf{s}'): \Delta(\mathbf{s}') = 0 \text{ vs. } H_1(\mathbf{s}'): \Delta(\mathbf{s}') \neq 0.$$

The test statistic we use to assess the statistical significance of each $H_0(\mathbf{s}')$ is

$$z(\mathbf{s}') = \frac{\hat{\Delta}(\mathbf{s}')}{\text{se}[\hat{\Delta}(\mathbf{s}')]}, \quad (4)$$

where $\text{se}[\hat{\Delta}(\mathbf{s}')] is the standard error of our best estimate $\hat{\Delta}(\mathbf{s}')$. To estimate this uncertainty, we use the block bootstrap (following, e.g., [Risser et al. 2019b](#)), which, as described in [Risser et al. \(2019b\)](#), is important for quantifying uncertainty here because the two-stage spatial GEV analysis in [section 3a](#) does not explicitly account for the spatial dependence in the daily measurements of precipitation, or the so-called storm dependence [dependence due to the spatial coherence of storm systems; see [Risser et al. \(2019b\)](#) for further details]. Instead of explicitly accounting for this systematic source of error, we rely on the block bootstrap to implicitly characterize the resulting uncertainty, since any real spatial signal will show up in most of the resampled datasets. Exploratory analyses in [Risser et al. \(2019b\)](#) verify that this approach appropriately accounts for storm dependence.$

The standard error $\text{se}[\hat{\Delta}(\mathbf{s}')] is obtained as follows: recall $\mathbf{Y}(\mathbf{s}_i) = \{Y_t(\mathbf{s}_i); t = 1950, \dots, 2017\}$ is the observed vector of the seasonal maxima weather station $\mathbf{s}_i$ (for now we suppress the j subscript denoting the specific season). Then, for $b = 1, \dots, B = 250$, the bootstrap sample is obtained by drawing $T = 68$ years *with*$

replacement from $\{1950, \dots, 2017\}$, denoted $\{a_1, a_2, \dots, a_T\}$, so that the b th bootstrap sample for station $\mathbf{s}$ is $\mathbf{Y}_b(\mathbf{s}_i) = [Y_{a_1}(\mathbf{s}_i), \dots, Y_{a_T}(\mathbf{s}_i)]$ (note that the temporal covariate values are resampled correspondingly). Furthermore, note that for a specific bootstrap sample, all stations use the same sequence of years in order to maintain the spatial coherence of the underlying seasonal maxima. (Bootstrap sampling with replacement is represented in the flow diagram in [Fig. 2](#) with pointer 2.) For each bootstrap sample, we then use the spatial extreme value analysis outlined in [section 3a](#) to obtain the bootstrap estimates of the return values $\hat{\phi}_{b,t}^{(r)}(\mathbf{s}')$ and change metric $\hat{\Delta}_b(\mathbf{s}')$ for all $\mathbf{s}' \in \mathcal{G}$. (This step is represented in the flow diagram in [Fig. 2](#) with pointer 3.) The resulting standard error is then

$$\text{se}[\hat{\Delta}(\mathbf{s}')] = \sqrt{\frac{1}{B-1} \sum_{b=1}^B \left[\hat{\Delta}_b(\mathbf{s}') - \frac{1}{B} \sum_{b=1}^B \hat{\Delta}_b(\mathbf{s}') \right]^2}.$$

[This standard error is used to calculate the test statistic in (4); shown in the flow diagram in [Fig. 2](#) with pointer 4.]

2) DETERMINATION OF NULL DISTRIBUTION FOR TESTING

Next, we determine the null distributions against which to test the null hypotheses defined in the previous section. In general, there are a variety of ways to use resampling-based procedures, for example the bootstrap, and the resulting z -scores (4) to conduct significance testing (see, e.g., [Efron and Tibshirani 1994](#)). However, in this case the test statistic (4) must be compared with corresponding z -scores under the assumption that the null hypothesis $H_0(\mathbf{s}')$ is true: in other words, we require a null distribution for the z -scores based on data that are known to have a trend of zero. Therefore, we cannot use the bootstrap estimates, because these are based on the original weather station data, which may or may not have a temporal trend. To navigate this issue, we can instead use a non-parametric permutation framework, in which we reshuffle the years in order to build up a data-driven null distribution for the z -scores based on datasets that have no trend by construction [see chapter 15 of [Efron and Tibshirani \(1994\)](#)—as is also done in, for example, [Lehmann et al. \(2015\)](#)]. This null distribution explicitly accounts for uncertainty arising from multiple sources: a limited sampling period, noisy measurement of the seasonal maxima, and the estimation procedure outlined in [section 3a](#). The permutation test proceeds similarly to the bootstrap:

1) For replicate $h = 1, \dots, H = 250$.

(i) *Without* replacement, sample $d_t \in \{1, \dots, T\}$, for $t = 1, \dots, T$. For each $\mathbf{s}_i \in \mathcal{S}$, collect the corresponding seasonal maxima into a vector, $\mathbf{Y}_h(\mathbf{s}_i) = [Y_{d_1}(\mathbf{s}_i), \dots, Y_{d_T}(\mathbf{s}_i)]$. As with the bootstrap, all stations use the same sequence of years to preserve the spatial structure of each reshuffling. (This step is represented in the flow diagram in Fig. 2 with pointer 5.)

(ii) Using the analysis outlined in section 3a, obtain the best estimate of the field of test statistics, denoted $\{\hat{\Delta}_h(\mathbf{s}'): \mathbf{s}' \in \mathcal{G}\}$, based on the permuted dataset $\{\mathbf{Y}_h(\mathbf{s}_i): \mathbf{s}_i \in \mathcal{S}\}$. Note that we treat these seasonal maxima for each weather station as if they arose from a consecutive T -yr period. In other words, the covariate values are not shuffled in the temporal trend, unlike the bootstrap. (This step is represented in the flow diagram in Fig. 2 with pointer 6.)

To be clear, note that these permutation estimates $\{\hat{\Delta}_h(\mathbf{s}'): h = 1, \dots, H\}$ define the null distribution for our best estimate of the change in the original (unshuffled) data at $\mathbf{s}'$, denoted $\hat{\Delta}(\mathbf{s}')$.

2) The permutation z -score for location $\mathbf{s}'$ is calculated as

$$z_h(\mathbf{s}') = \frac{\hat{\Delta}_h(\mathbf{s}')}{\text{se}[\hat{\Delta}_h(\mathbf{s}')]}.$$

In principle, one could bootstrap each reshuffled dataset to estimate the standard error $\text{se}[\hat{\Delta}_h(\mathbf{s}')]$. Instead, to avoid an unnecessary computational burden, we approximate the standard error for permutation sample h using the standard deviation of the estimates from all $H = 250$ permuted datasets

$$\text{se}[\hat{\Delta}_h(\mathbf{s}')] \approx \sqrt{\frac{1}{H-1} \sum_{j=1}^H \left[\hat{\Delta}_j(\mathbf{s}') - \frac{1}{H} \sum_{j=1}^H \hat{\Delta}_j(\mathbf{s}') \right]^2}, \quad (5)$$

which we verified gives quantitatively similar results to the brute-force approach of bootstrapping each permutation dataset.

3) ADDRESSING MULTIPLICITY AND FIELD SIGNIFICANCE

Finally, we deal with issues of multiplicity in hypothesis testing and field significance. (The steps that follow are represented in the flow diagram in Fig. 2 with pointer 7.) Given that we are conducting a hypothesis test for each of $M = 13073$ grid cells in $\mathcal{G}$ (for each season), testing each hypothesis individually at a significance level of α (e.g., often $\alpha = 0.05$ or 0.10) means that we would incorrectly reject $M\alpha$ of these hypotheses on

average even if all M null hypotheses were true. We must therefore think carefully about a multiple testing adjustment in order to avoid making a large number of type I errors (i.e., false positives or false discoveries). Following arguments described in Risser et al. (2019a), we seek to control the “false discovery rate” (FDR), or the proportion of type I errors from the set of spatial locations for which we conclude there is a nonzero change.

Usually, we might argue from a Bayesian or shrinkage perspective that the spatial smoothing underlying the calculation of the z -values would automatically account for the multiple testing aspect (Gelman et al. 2012), and we could simply compare $z(\mathbf{s}')$ with the corresponding $\{z_h(\mathbf{s}')\}$ to determine our decision regarding $H_0(\mathbf{s}')$. However in this case, our best estimates of the return values (and therefore change estimates) include both the true spatial signal as well as spatially correlated error due the presence of storm dependence (see Risser et al. 2019b), and so we require a multiple testing adjustment even after borrowing strength over space. Any such adjustment must account for the fact that there is strong spatial correlation among the M hypotheses. While most traditional multiple testing approaches are defined for independent hypotheses (see, e.g., Benjamini and Hochberg 1995), the literature also contain a number of approaches specifically designed for correlated hypotheses (see, e.g., Benjamini and Yekutieli 2001; Perone Pacifico et al. 2004; Sun and Cai 2009; Schwartzman and Lin 2011; Sun et al. 2015; Shu et al. 2015; Risser et al. 2019a). As discussed in Risser et al. (2019a), many of these approaches are designed for specific statistical modeling frameworks and are not robust to deviations from the statistical model; on the other hand, we cannot use the robust Bayesian methodology of Risser et al. (2019a) since we are in a frequentist setting.

While Benjamini and Hochberg (1995) was originally derived for independent hypotheses, their procedure is also valid (in the sense of controlling the FDR) under positive regression dependency (PRD) of the test statistics (Benjamini and Yekutieli 2001). Intuitively, PRD is satisfied if the test statistics are positively correlated (see section 2.2 of Benjamini and Yekutieli 2001); one can argue that this condition holds here when dealing with spatially smooth change estimates. Therefore, we can use an argument similar to the one outlined in section 3.3.3 of Benjamini and Hochberg (1995) to derive our testing procedure: here, the authors frame their FDR-controlling procedure as a post hoc maximization where one finds the value $\alpha \in [0, 1]$ that maximizes the number of rejections $R(\alpha)$ subject to the constraint that

$$\alpha M / R(\alpha) \leq q, \quad (6)$$

where q is the desired rate of false discoveries. Note that this procedure is implicitly based on a set of P values, for example, $\{p(\mathbf{s}'): \mathbf{s}' \in \mathcal{G}\}$, in that $R(\alpha) = \sum_{\mathbf{s}' \in \mathcal{G}} I\{p(\mathbf{s}') \leq \alpha\}$ (here, $I\{\cdot\}$ is an indicator function) as described in [Benjamini and Hochberg \(1995\)](#). Therefore, after observing the outcome of the procedure, the value α can be chosen by maximizing the number of rejections subject to the constraint in (6). Framing the FDR procedure in this way is based on the fact that $V(\alpha)$, the expected number of false rejections for a particular α , can be bounded above by αM because $V(\alpha) = \alpha M$ when all of the null hypotheses are true.

However, given the permutation results, we can adapt this framework slightly to yield a more appropriate FDR procedure. First, unlike [Benjamini and Hochberg \(1995\)](#), we define a rejection cutoff in terms of z -scores (instead of P values) because we do not have a defensible method for deriving appropriate P values from the permutation estimates. As such, the number of rejections and false rejections are now defined in terms of a z -score cutoff c , that is, $R(c)$ and $V(c)$. Second, we do not need to assume a theoretical result for the number of false rejections (in the previous paragraph, this was $V(\alpha) = \alpha M$) because we can derive an estimate empirically from the permutation z -scores $\{z_h(\mathbf{s}'): h = 1, \dots, H\}$. These z -scores automatically account for both the spatial correlation among the tests and all sources of uncertainty in the analysis. For each h , the number of false rejections for a given z -score cutoff $c \geq 0$ is estimated by

$$\hat{V}_h(c) = \sum_{\mathbf{s}' \in \mathcal{G}} I\{|z_h(\mathbf{s}')| > c\} \quad (7)$$

because the z -scores $\{z_h(\mathbf{s}')\}$ are calculated from a reshuffled dataset with no change and hence none of the tests should be rejected. Then, averaging over all permutation datasets, an estimate of the overall number of false rejections for a given c is $\hat{V}(c) = H^{-1} \sum_{h=1}^H \hat{V}_h(c)$, which is an improvement on the theoretical upper bound used in [Benjamini and Hochberg \(1995\)](#) in that it accounts for the spatial correlation of the collective set of hypotheses. Our testing procedure that includes a multiple testing adjustment proceeds as follows (these three steps are represented in the flow diagram in [Fig. 2](#) with pointer 7):

- 1) For each $c \geq 0$, calculate $\hat{V}(c)$ following (7) as well as the number of rejections in the full data, $\hat{R}(c) = \sum_{\mathbf{s}' \in \mathcal{G}} I\{|z(\mathbf{s}')| > c\}$.
- 2) Find the smallest value c such that $\hat{V}(c)/\hat{R}(c) \leq q$, denoted c^* .
- 3) Reject all $H_0(\mathbf{s}')$ with $|z(\mathbf{s}')| > c^*$.

Applying this procedure ensures that the rate of false discoveries (equivalently, the rate of type I errors or false positives) is bounded by q ([Benjamini and Hochberg 1995](#)). This yields a systematically different characterization of statistical significance than the concept of field significance (FS; [Livezey and Chen 1983](#)), which is an alternative multiple testing approach that seeks to evaluate the collective significance of a set of statistics. As such, FS provides no specific information about which individual grid cells are significant. A procedure that controls the FDR, on the other hand, does provide this local information: we can be confident that any grid cell flagged as significant is indeed significant. In any case, while we argue that FDR control provides a better characterization of statistical significance, we also present the FS of the changes for comparison, as this provides an alternate and complementary line of evidence for the collective significance of the changes. Here, the field significance is calculated based on the original data z -scores $\{z(\mathbf{s}')\}$ and the permutation z -scores $\{z_h(\mathbf{s}')\}$ as follows: first, calculate the cumulative distribution functions (CDFs) of the absolute value of the z -scores, denoted $F(z)$ and $F_h(z)$. Then, for each z -score cutoff c , the field significance is calculated as

$$\text{FS}(c) = \frac{1}{H} \sum_{h=1}^H I\{F_h(c) < F(c)\}; \quad (8)$$

that is, the proportion of the permutation CDFs that are less extreme than the original data CDF (in other words, the permutation probability density function has a heavier upper tail relative to the full data density).

4. Results

a. Seasonal results

Applying the analysis procedure described in [section 3](#) to the seasonal maxima, our best estimates of the relative change in the 20-yr return value [i.e., $\Delta^R(\mathbf{s}')$] are shown in [Fig. 3a](#) with best estimates of the absolute change [i.e., $\Delta^A(\mathbf{s}')$] in [Fig. 3b](#). Note that part of California is masked out for JJA because the dry season in this region means that an extreme value distribution for the seasonal maxima at these stations is inappropriate. In general, the largest spatially coherent changes appear in DJF, SON, and possibly MAM. The dominant direction of the observed changes is positive, although there are also large regions with negative changes in each season.

[Figure 3](#) also displays the statistical pointwise significance of these changes corresponding to “low” (i.e., controlling the rate of false discoveries at $q = 0.33$) and

(a) Relative change in 20-year return value (%/ha)

(b) Absolute change in 20-year return value (mm/ha)

FIG. 3. Gridded best estimates of the (a) relative change (percent ha^{-1}) and (b) absolute change (mm ha^{-1}) in the 20-yr return value (1950 vs 2017). Grid cells hashed with a “—” contain a change with low statistical significance (i.e., the proportion of type I errors for these grid cells is $\leq 33\%$); those with a “+” have high statistical significance (i.e., the proportion of type I errors for these grid cells is $\leq 10\%$).

“high” (i.e., controlling the rate of false discoveries at $q = 0.1$) confidence statements. These thresholds for significance are chosen as reasonable limits for bounding the proportion of type I errors, with $q = 0.33$ yielding a less conservative statement and $q = 0.1$ yielding a more conservative statement. The proportion of the map falling into each significance category for each season is given in Table 1 (in the “Gridded” columns). With regards to the relative change in return value, very few of the grid cell changes are significant even at the low level for DJF, MAM, and JJA; in SON, on the other

hand, a large proportion of CONUS (21.9%) contains changes with a low level of significance, with approximately 19.5% of these being increasing changes and the remaining 2.4% being decreasing changes. Turning to the absolute changes in return value, all seasons now exhibit at least some areas with changes that are at least of low significance, with spatially coherent increasing changes in the central United States in DJF and the New England states in MAM. As with the relative change in return value, a large proportion of CONUS displays changes with a low significance level in SON (33.3% of

TABLE 1. Proportion of independent station estimates (“Independent”) and spatially smoothed grid cells (“Gridded”) falling into each confidence category for (top) relative change in return value and (bottom) absolute change in return value in each season.

	Low significance (only)		High significance (only)	
	Independent	Gridded	Independent	Gridded
Relative change in return value				
DJF	0.0058	0.0000	0.0000	0.0000
MAM	0.0004	0.0000	0.0000	0.0000
JJA	0.0010	0.0058	0.0000	0.0000
SON	0.0031	0.2185	0.0000	0.0000
Absolute change in return value				
DJF	0.0686	0.0400	0.0258	0.0000
MAM	0.0336	0.0121	0.0083	0.0000
JJA	0.0125	0.0169	0.0038	0.0000
SON	0.0677	0.2655	0.0136	0.0675

the grid cells), with 6.8% of the grid cells further showing highly significant changes. Of those with low significance in SON, there are spatially coherent increases in the southeast United States, New Mexico, and northern Great Plains; on the other hand, there are highly significant decreases in California (specifically the northern coast and a cluster in the Sierra Nevada).

Clearly, when considering both the relative and absolute change in return values, there are many grid cells with a large change (in absolute value) that are not statistically significant even at the lower level. In principle, this could be due to several factors: noise in the station data itself, an overly conservative false discovery rate procedure, or the statistical interpolation procedure used throughout. The lack of significance is not due to the choice of return level $r = 20$: the significance maps are approximately the same when considering other return levels, both more and less extreme (not shown). In this analysis, however, the predominant causes are 1) noise in the station data and 2) unaccounted-for storm dependence in the underlying precipitation measurements. First, consider [appendix B \(Fig. B1\)](#) as well as [Table 1](#) (in the “Independent” columns), which show the results of a corresponding analysis of changes in precipitation extremes at each station that does not borrow strength spatially (but still applies the multiple testing adjustment). Very few stations exhibit even low significance in the relative change in return value, and while there are a moderate number of stations with high significance in the absolute change in return value (approximately 3%–6% for each season), the changes do not exhibit the spatial coherence that might be expected, consistent with the analysis of [Balling and Goodrich \(2011\)](#) indicating that station-based changes exhibit a high degree of entropy. This suggests that noise in the station data is the culprit, as opposed to the

interpolation procedure itself. (Note that here “noise” refers both to inherent uncertainty in the estimation of parameters in our GEV model and to measurement errors that induce errors in the change estimates.)

A second contributing factor to the relatively few significant changes is the unaccounted-for dependence in the underlying fields of daily precipitation, or storm dependence. Returning to the gridded results shown in [Fig. 3](#), recall that the significance assessment is based on reshuffling or permuting the seasonal maxima. Even though these permuted datasets have no time trend by construction, it is possible that the estimated change is as large (or larger) than the change estimates shown in [Fig. 3](#). For example, [appendix B \(Fig. B2\)](#) shows the standardized relative change in return value (as a z -score) in DJF for the unshuffled (original) data versus the corresponding quantity for five permutation datasets with particularly large changes. Clearly, the z -score for the change estimates in the permuted datasets can be of the same order as the z -score from the original data. As discussed in [Risser et al. \(2019b\)](#), the spatial extreme value analysis conducted here cannot distinguish between spatial correlation in the true underlying change and spatially correlated error due to storm dependence. This is likely the reason why we estimate such large changes in the permutation datasets: the random reshuffling of individual years with an individual storm that impacted a large geographic area (and hence a large number of stations) could result in large (but spurious) positive or negative changes.

To consider this more generally for all permutation datasets, consider [appendix B \(Figs. B3 and B4\)](#), which shows the cumulative distribution functions (CDFs) for the relative and absolute change in return values (respectively). Specifically, these figures show the CDF of the absolute value of the z -scores for both the original (unshuffled) data as well as each of the individual permutation datasets. The unshuffled CDF is systematically larger than the permutation mean CDF across all z -score values for both the relative and absolute change, but (aside from SON) the unshuffled CDF is not extremely unusual relative to the individual permutation CDFs. Therefore, we have good reason to believe the false discovery rate procedure is not overly conservative because our testing results are in line with what we might intuitively expect (from [Figs. B3 and B4](#)).

Finally, the field significance of these changes is shown also shown in [appendix B \(Fig. B5\)](#) [note that the calculations for $FS(c)$ in (8) are based on the CDFs shown in [Figs. B3 and B4](#)]. For the relative change, SON has a strong signal emerging regardless of the significance threshold; otherwise, only JJA has a small degree of field significance. This matches the FDR results, which

FIG. 4. Estimated 20-yr return values (mm) for each season over 1950–2017, with the annual “trend” shown with the dashed black line. The top (bottom) panel corresponds to a selected grid cell in Idaho (Mississippi). Note that the annual trend for the grid cell in Mississippi coincides exactly with the MAM seasonal trend.

show a large number of significant grid cells in SON and only a few significant cells in JJA. For the absolute change, on the other hand, all seasons have strong signals emerging in terms of field significance, with SON showing the strongest signal. Again, this is in agreement with the FDR-based significance statements, for which all seasons have at least some grid cells with pointwise significance. We reiterate that the FS results here (and in general) do not inform *where* the significant changes are, and furthermore FS tends to imply an overconfidence in the spatial extent of significant changes: as an example, for the absolute change, the JJA map is field significant even though only a small proportion of the grid cells (approximately 1%) are pointwise statistically significant.

b. Annual trends derived from estimated seasonal return values

While all results so far have been presented seasonally, we can also use the results from our fitted statistical model to make a statement about annual changes in extreme precipitation. Note that we do *not* reapply our spatial extreme value analysis to the annual maxima at each station: our analyses are conducted on the seasonal maxima because extreme precipitation results from different physical processes in different seasons, and aggregating over these differences confounds the important characteristics of the resulting extremes. Nonetheless, annual changes can be useful to water resource managers.

Here, we first conduct seasonal analyses, then calculate the largest 20-yr return value amongst the four

seasonal values in each year (which implicitly accounts for distributional differences) and then estimate a trend over time from the derived (artificial) time series. Given that we have specified a linear trend, this means that we can simply compare the largest seasonal return value in 2017 with the largest seasonal return value in 1950. In other words, we can assess

$$\hat{\Delta}_{*}^R(\mathbf{s}') \equiv \frac{\max_u [\hat{\phi}_{u,2017}^{(20)}(\mathbf{s}')] - \max_u [\hat{\phi}_{u,1950}^{(20)}(\mathbf{s}')] }{\max_u [\hat{\phi}_{u,1950}^{(20)}(\mathbf{s}')] } \quad (9)$$

and

$$\hat{\Delta}_{*}^A(\mathbf{s}') \equiv \max_u [\hat{\phi}_{u,2017}^{(20)}(\mathbf{s}')] - \max_u [\hat{\phi}_{u,1950}^{(20)}(\mathbf{s}')], \quad (10)$$

where $u = \{\text{DJF}, \text{MAM}, \text{JJA}, \text{SON}\}$ are the different seasons and $\hat{\phi}_{u,t}^{(20)}(\mathbf{s})$ is the *estimated* 20-yr return value in season u and year t at grid cell $\mathbf{s}$. As an illustration, consider Fig. 4, which shows the estimated seasonal trends for two sample grid cells as well as the annual “trend.” Clearly, the annual change can be no larger than the largest seasonal change: note that for the grid cell in Idaho, the annual trend is nearly flat in spite of relatively large changes in MAM, JJA, and SON; on the other hand, the annual change can coincide with one of the seasonal changes, as in the grid cell in Mississippi.

Applying the same significance testing procedure as outlined in section 3b to the annual changes $\hat{\Delta}_{*}^R(\mathbf{s}')$ and $\hat{\Delta}_{*}^A(\mathbf{s}')$, the best estimates of the annual relative and absolute changes are shown in Fig. 5. Note that none of the annual changes are statistically significant even at the low significance level (nor are the maps field

FIG. 5. Annual gridded best estimates of the (a) relative change (percent ha^{-1}) and (b) absolute change (mm ha^{-1}) in the 20-yr return value (1950 vs 2017). Grid cells hashed with a “-” contain a change with low statistical significance (i.e., the proportion of type I errors for these grid cells is $\leq 33\%$); those with a “+” have high statistical significance (i.e., the proportion of type I errors for these grid cells is $\leq 10\%$). Note that no cells in either panel are significant even at the lowest level.

significant). This is not entirely surprising given the way these changes were constructed: again looking at Fig. 4, note that a large SON change in Mississippi is disguised by a nearly flat trend in the wetter MAM season.

5. Discussion

In this paper, we have characterized changes in observed extreme precipitation for a fine spatial grid over CONUS using station data only and derived a statistical technique to robustly identify statistically significant pointwise changes in the climatological distribution of extreme precipitation. Our results constitute a novel contribution to the literature on observed changes in extreme precipitation because 1) we characterize changes at a point-based scale that is the native spatial scale for precipitation observations, 2) we account for seasonal differences, and 3) we report uncertainty by ascribing statistical significance. No other study to our knowledge simultaneously addresses these three components [although, for example, Zhang et al. (2010) explore seasonal changes in extremes for individual stations]. Furthermore, we describe an approach for constructing an annualized change based on the seasonal estimates that is of interest to, for example, water resource managers.

The maps in Fig. 3b exhibit a strong correspondence with the observed changes in 20-yr return values over 1948 to 2015 given in Easterling et al. (2017) for seven large aggregated geographic regions [see Fig. 7.2 in Easterling et al. (2017)]. In winter, Easterling et al. (2017) find large positive changes in the southern Great Plains, with smaller increases in the Midwest, Southeast, and Northeast; our results identify the same changes. There is a similar agreement for the other seasons (e.g., large increases in the Northeast for spring, the southern Great Plains in summer, and the Southeast in fall).

However, note that Easterling et al. (2017) show roughly zero change in the Southwest for fall, which is likely the result of aggregating over individual station trends (see Fig. B1) that show large decreases in Northern California and correspondingly large increases in New Mexico. These similarities provide support for the accuracy of our results while illustrating the impact (and importance) of resolving changes spatially.

The technique developed to produce this dataset provides a natural method for evaluating changes in historical simulations at the native scale of the simulations. The Gaussian process method used to model the spatial statistics of the station-based changes can be used to estimate the changes at arbitrary spatial locations (e.g., locations of model grid cells). Model-based precipitation is typically a cell-averaged quantity, and so the magnitudes of extreme values will generally be lower than the point-based values used here. Because of this, Chen and Knutson (2008) advocate for utilizing equivalently area-averaged observations when evaluating the statistics of simulated precipitation. However, as model resolution increases, the values of precipitation should asymptote to point-like values, and so the changes from this dataset (specifically, the significant changes) represent a quantity to which a suitable historical model should converge as resolution increases. These two perspectives are not mutually exclusive: the Chen and Knutson (2008) perspective suggests that the most fair evaluation of model output would utilize precipitation data that is aggregated in a way that represents an appropriate area average, whereas our perspective suggests that point-based values provide statistics to which a model should converge as grid spacing approaches zero. This effectively represents an upper bound for the statistics of extremes from simulations. This dataset could therefore form the basis for a complementary metric of model performance in

simulating temporal changes in extremes, and this metric could be tracked over model development cycles. More broadly, it is worth considering whether the methodology employed here could be merged with that of Donat et al. (2013) to produce spatially complete estimates of a variety of metrics of extremes: at the native resolution of models against which the data would be compared.

Of course, our analysis involves several important limitations. While we have carefully derived a valid statistical procedure for detecting statistically significant changes, note that our results appear to be conservative since we emphasize the pointwise significant results and not the global assessment of field significance. As mentioned in section 4, the changes might be clarified (or “more significant”) if we accounted for the spatial dependence in the underlying fields of daily precipitation, but there is no existing methodology to do so appropriately for a large network of weather stations over a heterogeneous domain like CONUS. A potentially clearer way to handle the multiple testing would be to use the methodology in Risser et al. (2019a)—in a Bayesian setting, this would allow us to answer questions like “is the relative change in return value greater than 5% per century?”—but again the lack of appropriate methodology prevents us from using such an approach in this paper.

It should also be noted that the generalized extreme value (GEV) distribution used throughout this paper is the asymptotic (limiting) distribution of the sample maximum of a large sample of n independent draws from a parent population; this theoretical result applies exactly as the sample size $n \rightarrow \infty$ (see, e.g., Coles 2001, chapter 3). Of course, in practice, one only has a finite sample size; furthermore, in a block maxima framework, there is a trade-off between choosing large block sizes (leading to a better approximation, i.e., less bias, but larger variance due to fewer block maxima) versus small block sizes (an increase in the bias but smaller variance). A standard choice when dealing with climate data is to choose block sizes of one year; in this paper, we have chosen to apply the GEV distribution to seasonal maxima, where $n = 90$ (in which case we have $T = 68$ blocks or years of data). While a block size of 90 (in general) might be considered sufficient for the asymptotic results to hold, there are two important considerations when analyzing real measurements of daily precipitation:

- 1) In some locations over CONUS and in some seasons, there are a large number of days with zero precipitation. Technically, the GEV approximation applies to draws from any underlying parent population (e.g., a distribution with a point mass at zero and a

long right tail), but the point remains that if we are considering the “precipitation process” (i.e., days with nonzero rainfall), we may have fewer than 90 measurements of this process in a 90-day season.

- 2) Furthermore, the GEV approximation applies to a *statistically independent* sample of size n . While daily precipitation may not exhibit strong temporal autocorrelation, the effective sample size of independent measurements of daily precipitation from a 90-day season is certainly less than 90.

In light of these considerations, appendix C contains the results of a simple perfect data experiment (or simulation study) to evaluate the quality of the GEV approximation as a function of sample (block) size. In summary, there is indeed a nonnegligible effect of block size on the quality of the GEV approximation and resulting bias/uncertainty quantification for estimating return values, but when estimating return values for return periods that are not too extreme (within the range of the data) this effect is minimal. Certainly for the 20-yr return values estimated in this paper, we can be confident that our estimates contain minimal bias and the standard errors are well calibrated. This result holds even when the effective number of measurements of daily precipitation we have from each season (due to zero rainfall days, autocorrelation, or missing data) is as low as 25 or even 10. It should be noted that this result is absolutely not true when considering much longer-term (e.g., 500- or 1000-yr) return values: in this case there is significant bias and the standard errors are much too large (i.e., conservative).

Returning to the potential inappropriateness of using a linear trend to characterize nonsmooth and possibly nonlinear trends over the last 70 years, we reiterate that the simple trend model used here nonetheless yields important insights into changes in the climatology of extreme precipitation. As a sensitivity analysis, we analyzed the same data using a quadratic trend in the location parameter, that is, $\mu_t = \mu_0 + \mu_1 t + \mu_2 t^2$ (although still using $\sigma_t = \sigma$ and $\xi_t = \xi$). The resulting relative and absolute changes using a quadratic trend correspond very closely to the changes using a linear trend (see Fig. B6 in the appendix), which provides support for the appropriateness of characterizing changes using a linear trend. We reiterate that when attempting to explain (i.e., attribute) these changes it will be necessary to account for interannual and decadal (natural) variability in precipitation extremes, which we will explore in future work.

Acknowledgments. The data supporting this article are based on publicly available measurements from the

TABLE B1. Fixed quantities used throughout the analysis.

Quantity and fixed value	Definition
$T = 68$	The number of years (1950–2017) used in the analysis, also the number of seasonal maxima used to estimate extreme value statistics over time for each station.
$n = 5202$	The number of reliable GHCN stations and their geographic locations.
$M = 13\,073$	The number of 0.25° grid cells over CONUS.

TABLE B2. Specified quantities used throughout the analysis.

Quantity and fixed value	Definition
$B = 250$	The number of bootstrap datasets used to quantify all sources of uncertainty in the original (unshuffled) data.
$H = 250$	The number of reshuffled or permuted datasets used to approximate the null distribution and quantify all sources of uncertainty.
$q^* = 0.33$ and 0.1	The acceptable rates of false discoveries, which also define the statistical significance thresholds. These correspond, respectively, to low and high confidence statements.

National Centers for Environmental Information (at <ftp://ftp.ncdc.noaa.gov/pub/data/ghcn/daily/>). This research was supported by the Director, Office of Science, Office of Biological and Environmental Research of the U.S. Department of Energy under Contract DE-AC02-05CH11231 and used resources of the National Energy Research Scientific Computing Center (NERSC), also supported by the Office of Science of the U.S. Department of Energy, under Contract DE-AC02-05CH11231. The authors would also like to acknowledge three anonymous referees whose feedback greatly improved the quality of this paper. This document was prepared as an account of work sponsored by the U.S. government. While this document is believed to contain correct information, neither the U.S. government nor any agency thereof, nor the Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the U.S. government or any agency thereof, or the Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the U.S. government or any agency thereof or the Regents of the University of California.

APPENDIX A

Changes in Return Values

To simplify notation in the following, note that when the shape parameter is constant over time, that

is, $\xi_t(\mathbf{s}) \equiv \xi(\mathbf{s})$, we can write the r -yr return value in year t as

$$\phi_t^{(r)}(\mathbf{s}) = \mu_t(\mathbf{s}) - \sigma_t(\mathbf{s})f_{\xi(\mathbf{s})}(r),$$

where

$$f_{\xi(\mathbf{s})}(r) = \begin{cases} \frac{1}{\xi(\mathbf{s})} \{1 - [-\log(1 - 1/r)]^{-\xi(\mathbf{s})}\}, & \xi(\mathbf{s}) \neq 0 \\ \log[-\log(1 - 1/r)], & \xi(\mathbf{s}) = 0 \end{cases}.$$

Recall from [section 3a](#) that the time-varying model we use in this paper is

$$\mu_t(\mathbf{s}) = \mu_0(\mathbf{s}) + \mu_1(\mathbf{s})t, \quad \sigma_t(\mathbf{s}) \equiv \sigma(\mathbf{s}), \quad \xi_t(\mathbf{s}) \equiv \xi(\mathbf{s}).$$

For this model, the relative change in return value is

$$\Delta^R(\mathbf{s}) = \frac{\phi_{t_2}^{(r)}(\mathbf{s}) - \phi_{t_1}^{(r)}(\mathbf{s})}{\phi_{t_1}^{(r)}(\mathbf{s})} = \frac{\mu_1(\mathbf{s})(t_2 - t_1)}{\mu_0(\mathbf{s}) + \mu_1(\mathbf{s})t_1 - \sigma(\mathbf{s})f_{\xi(\mathbf{s})}(r)},$$

which depends on r ; on the other hand, the absolute change in return value is

$$\Delta^A(\mathbf{s}) = \phi_{t_2}^{(r)}(\mathbf{s}) - \phi_{t_1}^{(r)}(\mathbf{s}) = \mu_1(\mathbf{s})(t_2 - t_1),$$

which is independent of r . An alternative four-parameter representation of the time-varying model is

$$\mu_t(\mathbf{s}) = \mu_0 \exp\left[\frac{\alpha(\mathbf{s})t}{\mu_0(\mathbf{s})}\right], \quad \sigma_t(\mathbf{s}) = \sigma_0(\mathbf{s}) \exp\left[\frac{\alpha(\mathbf{s})t}{\mu_0(\mathbf{s})}\right], \\ \xi_t(\mathbf{s}) \equiv \xi(\mathbf{s})$$

([van der Wiel et al. 2017](#)), which allows both the location and scale to vary over time such that their ratio is a constant, that is, $\mu_t(\mathbf{s})/\sigma_t(\mathbf{s}) = \mu_0(\mathbf{s})/\sigma_0(\mathbf{s})$. Under this parameterization, the relative change in return value is

TABLE B3. Summary of the statistical trend models for quantifying changes in the GEV distribution.

Model label	Location trend	Scale trend	Shape trend	No. of parameters (N_{par})
M0	Linear in time	Constant	Constant	4
M1	Quadratic in time	Constant	Constant	5
M2	Linear in time	Linear in time	Constant	5
M3	Linear in time	Constant	Linear in time	5
M4	Linear in time	Linear in time	Linear in time	6

$$\frac{\mu_0(\mathbf{s}) \exp \left[\frac{\alpha(\mathbf{s})t_2}{\mu_0(\mathbf{s})} \right] - \sigma_0(\mathbf{s}) \exp \left[\frac{\alpha(\mathbf{s})t_2}{\mu_0(\mathbf{s})} \right] f_{\xi(\mathbf{s})}(r)}{\mu_0(\mathbf{s}) \exp \left[\frac{\alpha(\mathbf{s})t_1}{\mu_0(\mathbf{s})} \right] - \sigma_0(\mathbf{s}) \exp \left[\frac{\alpha(\mathbf{s})t_1}{\mu_0(\mathbf{s})} \right] f_{\xi(\mathbf{s})}(r)} - 1,$$

which again depends on r ; note that the absolute change also depends on r .

APPENDIX B

Supplemental Tables and Figures

Supplemental tables (Tables B1–B4) and figures (Figs. B1–B9) related to the results in section 4 are presented here.

APPENDIX C

Sample Size Considerations for Asymptotic Theory

As mentioned in the discussion section (section 5), the generalized extreme value (GEV) distribution used throughout this paper is the asymptotic (limiting) distribution of the sample maximum of a large sample of n independent draws from a parent population; this theoretical result applies exactly as the sample size $n \rightarrow \infty$ (see, e.g., Coles 2001, chapter 3). Of course, in practice, one only has a finite sample size; furthermore, in a block maxima framework, there is a trade-off between choosing large block sizes (leading to a better approximation, i.e., less bias, but larger variance due to fewer block maxima) versus small block sizes (an increase in the bias but smaller variance). Here, we attempt to evaluate the quality of the GEV approximation for finite n .

A standard choice when dealing with climate data is to choose block sizes of one year; in this paper, we have chosen to apply the GEV distribution to seasonal maxima, where $n = 90$ (in which case we have $T = 68$ blocks or years of data). Two considerations are important when analyzing the extreme values of daily precipitation:

- 1) In some locations over CONUS and in some seasons, there are a large number of days with zero precipitation. Technically, the GEV approximation applies to draws from any underlying parent population (e.g., a distribution with a point mass at zero and a long right tail), but the point remains that if we are considering the “precipitation process” (i.e., days with nonzero rainfall), we may have fewer than 90 measurements of this process in a 90-day season.
- 2) Furthermore, the GEV approximation applies to a *statistically independent* sample of size n . While daily precipitation may not exhibit the temporal autocorrelation of, e.g., temperature, the effective number of independent measurements of daily precipitation from a 90-day season is certainly less than 90.

In light of these considerations, we conduct a simple perfect data experiment (or simulation study) to evaluate the quality of the GEV approximation as a function of sample (block) size. For now, we set aside the spatial and temporal aspect of the analysis in this paper and focus on analyzing the extreme values of a stationary time series from a single station.

For the purposes of this illustration, we generate artificial data from a known underlying distribution that has a point mass at zero and otherwise an exponential distribution with rate 1 (the exponential distribution with rate 1 has mean 1 and variance 1). In other words,

TABLE B4. Comparison of models in Table B3 using the Akaike information criterion (AIC). This AIC can be thought of as quantifying the predictive skill of the various models, since we use smoothed coefficient estimates to “predict” the original seasonal maxima. Formally, we calculate the log-likelihood of the original seasonal maxima using the spatially smoothed GEV coefficients for each weather station, and then averaged this log-likelihood over all stations for M0–M4 (denoted $\log L$). The predictive AIC is then calculated as $\text{AIC} = -2\log \bar{L} + 2N_{\text{par}}$, where N_{par} is the number of statistical parameters in each model. Smaller AIC values indicate a better statistical model; best model is in bold.

Season	M0	M1	M2	M3	M4
DJF	465.18	467.41	466.53	467.0881	468.64
MAM	499.18	501.59	500.38	500.8088	502.43
JJA	522.08	529.5582	523.43	524.11	528.83
SON	522.97	525.07	524.39	524.83	526.30

(a) Relative change in 20-year return value (%/ha)

(b) Absolute change in 20-year return value (mm/ha)

FIG. B1. Best estimates of the (a) relative change and (b) absolute change in the 20-yr return value (1950 vs 2017) at the station locations *without* spatial smoothing. Station locations hashed with a “-” contain a change with low statistical significance (i.e., the proportion of type I errors for these grid cells is $\leq 33\%$); those with a “+” have high statistical significance (i.e., the proportion of type I errors for these grid cells is $\leq 10\%$).

let $\{X_1, X_2, \dots\}$ denote an independent sequence of values representing daily precipitation, where the probability distribution function of each X_i is

$$f_X(x) = p \times I\{x = 0\} + (1 - p) \times I\{x > 0\} \times f_E(x),$$

where $I\{\cdot\}$ is an indicator function and $f_E(x) = \exp\{-x\}$ is the exponential distribution with rate 1. Here, p is the probability of experiencing zero rainfall on a particular day. Let $F_X(x) \equiv P(X \leq x) = \int_{-\infty}^x f_X(y) dy$ denote the corresponding cumulative distribution function of X , and define $M_n = \max\{X_1, \dots, X_n\}$ to

be the maximum value from a sample of size n . The cumulative distribution function of the sample maximum M_n can be derived as follows:

$$\begin{aligned} F_{M_n}(y) &\equiv P(M_n \leq y) = P(X_1 \leq y, \dots, X_n \leq y) \\ &= P(X_1 \leq y) \times \dots \times P(X_n \leq y) \\ &= F_X^n(y) \end{aligned}$$

The extremal types theorem (Fisher and Tippet 1928) states that if there exist normalizing constants $a_n > 0$ and b_n such that, as $n \rightarrow \infty$,

FIG. B2. Gridded best estimates of the standardized relative change in the 20-yr return value (1950 vs 2017) for (top left) the full unshuffled data and five selected permutation datasets.

$$P\left(\frac{M_n - b_n}{a_n} \leq y\right) nG(y),$$

for some nondegenerate distribution G , then G is the generalized extreme value distribution. Without loss of generality, for now suppose $p = 0$; if we let $a_n = 1$ and $b_n = n$, it can be shown that

$$\begin{aligned} P\left(\frac{M_n - b_n}{a_n} \leq y\right) &= F_X^n(y + \log n) \\ &= [1 - \exp(y + \log n)]^n, \\ &= [1 - n^{-1} \exp(y)]^n \\ &\rightarrow \exp[-\exp(-y)] \end{aligned}$$

as $n \rightarrow \infty$ for each $y \in \mathcal{R}$ (Coles 2001, chapter 3, example 3.1 therein). Note that $\exp[-\exp(-y)]$ is the GEV distribution with shape parameter $\xi = 0$ (i.e., the Gumbel distribution). In this simulated example, given that we know F_X , the sequences $\{a_n\}$ and $\{b_n\}$, and the limiting Gumbel distribution, we can compare the approximation for various fixed values of n : see Fig. B7. In this case, the convergence happens quickly, with $F_X^n(y + \log n)$ and $\exp[-\exp(-y)]$ being visually indistinguishable by the time $n = 50$. This relatively fast convergence holds even when the underlying distribution has a much heavier tail [e.g., a Pareto distribution; see Davison and Huser (2015) or the video (<http://www.annualreviews.org/doi/story/10.1146/multimedia.2015.02.23.354>)]. Considering the case where p (the proportion of zero rainfall days) is not zero, we could simply rescale the number of measurements

according to p to give an “effective” sample size from the nonzero component of F_X .

To consider the appropriateness of the GEV approximation more quantitatively, we now evaluate estimated return values for the simulated data. In this case where we know the form of F_X , we can derive exactly the distribution of the seasonal maximum:

$$f_M(y) = n(1 - p) \times F_E^{n(1-p)-1}(y) \times f_E(y),$$

where $F_E(y) = 1 - \exp\{-y\}$ is the CDF of the exponential distribution; recall n is the block size and p is the proportion of zero rainfall days. The r -yr return value is then simply the upper $1 - 1/r$ quantile of f_M . In a general case where F_X is unknown (as is the case in this paper), one might use the T block maximum values to estimate the r -yr return value. As an illustration, consider Fig. B8, where we show histograms of three simulated datasets of daily precipitation values generated from f_X , representing $T = 68$ years of $n = 90$ values with a proportion of $p = 0.3$ zeros. The red density is the true density f_M of the seasonal maximum, while the blue density is the estimated GEV distribution for the block maxima. In some cases the GEV approximation is good (Fig. B8a), while in other cases the GEV either overestimates (Fig. B8b) or underestimates (Fig. B8c) the upper tail. These differences simply represent sampling variability in the simulated datasets.

More generally, we can explore the performance of using the GEV to estimate return values across a large number of simulated datasets. For now we maintain a fixed number of $T = 68$ blocks or years, corresponding to the number of years in our actual study, but we now

FIG. B3. Cumulative distribution functions (CDFs) for the relative change in return value (absolute value of the z -score) for each season. The red line refers to the CDF for the original (unshuffled) data, while the light gray lines refer to the individual permutation CDFs and their mean (in green). The unshuffled CDF lies well within the range of the permutation CDFs except for SON.

generate 10 000 simulated datasets for a variety of block sizes $n \in \{5, 10, 25, 50, 100, 200\}$ with a fixed proportion of zero rainfall days $p = 0$ (note that there is a redundancy when considering *both* variable block size and variable proportion of zero rainfall days). Also, in the simulated examples, note that the data are also statistically independent. We also explore estimating the $r \in \{10, 20, 50, 100, 500, 1000\}$ year return value from these 68 blocks of data with varying block size. As in the main text, we use resampling-based techniques (the block bootstrap) to estimate the resulting standard errors in the return value estimates.

To evaluate the return value estimates and corresponding uncertainty quantification, we use two metrics:

- 1) Root-mean-square error (RMSE): the square root of the mean squared differences (taken across the 10 000 replicates) between the true and estimated return values. RMSE evaluates the bias of the estimation procedure; smaller RMSE indicates less bias, and vice versa.
- 2) Relative error (RE): the standard deviation of the return value estimates across the 10 000 replicates

divided by the average bootstrap standard error of the return value estimates from the 10 000 replicates, where we subtract one from this ratio and multiply by 100 to convert to a percent error. The RE evaluates the uncertainty quantification of the estimation procedure: if the RE is zero, then the bootstrap standard errors are (on average) equal to the “true” (Monte Carlo) standard error (i.e., the standard errors are well calibrated). If the RE is less than 0, this indicates that the bootstrap standard errors are too large (i.e., the uncertainty is conservative); if the RE is greater than 0, this indicates that the bootstrap standard errors are too small (i.e., the uncertainty is anticonservative).

Figure B9 shows the results for the exponential data previously described (the exponential distribution with rate 1 has a mean of one and a variance of one) as well as data generated from a heavier-tailed gamma distribution with shape 1/3 and scale 3 (this distribution still has mean 1 but a variance of 3). Again recall that the block sizes given here correspond to $p = 0$ (i.e., a zero proportion of nonrainy days) and are also statistically independent. Real daily precipitation data

FIG. B4. Cumulative distribution functions (CDFs) for the absolute change in return value (absolute value of the z -score) for each season. The red line refers to the CDF for the original (unshuffled) data, while the light gray lines refer to the individual permutation CDFs and their mean (in green). As with the relative change, the unshuffled CDF for the absolute change lies well within the range of the permutation CDFs except for SON.

both (i) have zero rainfall days and (ii) are not statistically independent; however, the block sizes here could be considered “effective” block sizes (i.e., in a 90-day season, we may only have a much smaller

number of nonzero, independent measurements of precipitation).

There are several important messages from this simulation study. When estimating return values for return

FIG. B5. Field significance for the seasonal changes for each of a range of z -score cutoffs: (left) relative change in 20-yr return value; (right) absolute change in 20-yr return value. For the relative change, SON has a strong signal emerging for essentially all z -score cutoffs; otherwise, only JJA has a slightly lesser degree of field significance. For the absolute change, SON again has the strongest signal, but there appears to be an emerging signal in the other seasons as well.

FIG. B6. Comparison of relative and absolute changes using a linear trend in the location parameter (x axis) vs the corresponding change using a quadratic trend in the location parameter (y axis). Scale and shape parameters are time-invariant for both analyses.

periods within the range of the data (up to approximately 50-yr return periods), the bias of the estimates is independent of the block size: the RMSE is essentially constant across the block sizes considered. Similarly,

the standard errors are well calibrated for shorter-term return periods (certainly 10- to 20- and maybe even 50-yr return values) and are approximately constant across block sizes. For the 20-yr return value, standard

FIG. B7. Comparison of the cumulative distribution function (CDF) of the normalized sample maximum $F_X^n(y + \log n)$ (solid line) vs the GEV (Gumbel) limiting distribution (dashed red line) for several values of n . In this case, the distribution of the maximum converges quite quickly to the limiting GEV distribution.

FIG. B8. Histograms of three simulated datasets of daily precipitation values generated from f_X , representing $T = 68$ years of $n = 90$ values with a proportion of $p = 0.3$ zeros (note that the y axis is truncated). The red density is the true density f_M of the block (seasonal) maximum, while the blue density is the GEV approximation. In some cases the GEV approximation (a) appears to be good, while in other cases the GEV either (b) overestimates or (c) underestimates the upper tail.

errors are on average about 4% (2%) too large for the exponential (gamma) data across the block sizes considered. This means that for shorter-term return values, the standard errors are actually slightly anti-conservative, but in any case this effect is minor. However, block size is much more important for longer-term return periods: considering the estimated 1000-yr return values, the bias decays monotonically with increasing block size and is an order of magnitude larger for a block size of 5 versus a block size of 200. Similarly, the relative error improves monotonically with block size, such that

standard errors are much too large (i.e., conservative) for small block sizes relative to the larger block sizes. These results hold for both the exponential and heavier-tailed gamma data, although the relative bias and lack of calibration are worse for the heavier-tailed data.

In summary, there is a nonnegligible effect of block size on the quality of the GEV approximation and resulting bias/uncertainty quantification for estimating return values, but when estimating return values for return periods that are not too extreme (within the range of the data) this effect is minimal. Certainly for the 20-yr

FIG. B9. Root-mean-square error (RMSE) and relative error (RE; the true Monte Carlo standard deviation divided by the average bootstrap standard error minus 1) of estimated return values, averaged over 10 000 simulated datasets from an underlying exponential or gamma distribution for a variety of return periods and block sizes. Smaller RMSE indicates less bias; RE greater than zero indicates standard errors that are too small (i.e., the uncertainty is anticonservative) while RE less than zero indicates that the standard errors are too large (i.e., the uncertainty is conservative).

return values estimated in this paper, we can be confident that our estimates contain minimal bias and the standard errors are well calibrated. This result holds even when the effective number of measurements of daily precipitation we have from each season (due to either zero rainfall days, autocorrelation, or missing data) is as low as 25 or even 10. Of course, real daily precipitation data likely do not perfectly follow either an exponential or gamma distribution, but these results provide an assurance that the GEV approximation is a safe one to make for the return periods considered in the paper.

REFERENCES

- Balling, R. C., and G. B. Goodrich, 2011: Spatial analysis of variations in precipitation intensity in the USA. *Theor. Appl. Climatol.*, **104**, 415–421, <https://doi.org/10.1007/s00704-010-0353-0>.
- Benjamini, Y., and Y. Hochberg, 1995: Controlling the false discovery rate: A practical and powerful approach to multiple testing. *J. Roy. Stat. Soc.*, **57B**, 289–300, <https://doi.org/10.1111/J.2517-6161.1995.TB02031.X>.
- , and D. Yekutieli, 2001: The control of the false discovery rate in multiple testing under dependency. *Ann. Stat.*, **29**, 1165–1188, <https://doi.org/10.1214/aos/1013699998>.
- Chen, C.-T., and T. Knutson, 2008: On the verification and comparison of extreme rainfall indices from climate models. *J. Climate*, **21**, 1605–1621, <https://doi.org/10.1175/2007JCLI1494.1>.
- Coles, S., 2001: *An Introduction to Statistical Modeling of Extreme Values*. Springer, 208 pp, <https://books.google.com/books?id=2nugUEaKqFEC>.
- Davison, A. C., and R. Huser, 2015: Statistics of extremes. *Annu. Rev. Stat. Appl.*, **2**, 203–235, <https://doi.org/10.1146/annurev-statistics-010814-020133>.
- Donat, M., and Coauthors, 2013: Updated analyses of temperature and precipitation extreme indices since the beginning of the twentieth century: The HadEX2 dataset. *J. Geophys. Res. Atmos.*, **118**, 2098–2118, <https://doi.org/10.1002/JGRD.50150>.
- Easterling, D., and Coauthors, 2017: Precipitation change in the United States. *Climate Science Special Report: Fourth National Climate Assessment, Volume I*, 207–230, <https://doi.org/10.7930/J0H993CC>.
- Efron, B., and R. J. Tibshirani, 1994: *An Introduction to the Bootstrap*. CRC Press, 456 pp.
- Fisher, R. A., and L. H. C. Tippett, 1928: Limiting forms of the frequency distribution of the largest or smallest member of a sample. *Mathematical Proceedings of the Cambridge Philosophical Society*, Vol. 24, Cambridge University Press, 180–190.
- Gelman, A., J. Hill, and M. Yajima, 2012: Why we (usually) don't have to worry about multiple comparisons. *J. Res. Educ. Eff.*, **5**, 189–211, <https://doi.org/10.1080/19345747.2011.618213>.
- Gervais, M., L. B. Tremblay, J. R. Gyakum, and E. Atallah, 2014: Representing extremes in a daily gridded precipitation analysis over the United States: Impacts of station density, resolution, and gridding methods. *J. Climate*, **27**, 5201–5218, <https://doi.org/10.1175/JCLI-D-13-00319.1>.
- Groisman, P. Ya., R. W. Knight, D. R. Easterling, T. R. Karl, G. C. Hegerl, and V. N. Razuvaev, 2005: Trends in intense precipitation in the climate record. *J. Climate*, **18**, 1326–1350, <https://doi.org/10.1175/JCLI3339.1>.
- Hartmann, D., and Coauthors, 2013: Observations: Atmosphere and surface. *Climate Change 2013: The Physical Science Basis*. T. F. Stocker et al., Eds., Cambridge University Press, 159–254, <https://doi.org/10.1017/CBO9781107415324.008>.
- Janssen, E., D. J. Wuebbles, K. E. Kunkel, S. C. Olsen, and A. Goodman, 2014: Observational- and model-based trends and projections of extreme precipitation over the contiguous United States. *Earth's Future*, **2**, 99–113, <https://doi.org/10.1002/2013EF000185>.
- King, A. D., L. V. Alexander, and M. G. Donat, 2013: The efficacy of using gridded data to examine extreme rainfall characteristics: A case study for Australia. *Int. J. Climatol.*, **33**, 2376–2387, <https://doi.org/10.1002/joc.3588>.
- Kunkel, K. E., 2003: North American trends in extreme precipitation. *Nat. Hazards*, **29**, 291–305, <https://doi.org/10.1023/A:1023694115864>.
- , and Coauthors, 2013: Monitoring and understanding trends in extreme storms: State of knowledge. *Bull. Amer. Meteor. Soc.*, **94**, 499–514, <https://doi.org/10.1175/BAMS-D-11-00262.1>.
- Lehmann, J., D. Coumou, and K. Frieler, 2015: Increased record-breaking precipitation events under global warming. *Climatic Change*, **132**, 501–515, <https://doi.org/10.1007/s10584-015-1434-y>.
- Livezey, R. E., and W. Y. Chen, 1983: Statistical field significance and its determination by Monte Carlo techniques. *Mon. Wea. Rev.*, **111**, 46–59, [https://doi.org/10.1175/1520-0493\(1983\)111<0046:SFSaid>2.0.CO;2](https://doi.org/10.1175/1520-0493(1983)111<0046:SFSaid>2.0.CO;2).
- Lovejoy, S., D. Schertzer, and V. Allaire, 2008: The remarkable wide range spatial scaling of TRMM precipitation. *Atmos. Res.*, **90**, 10–32, <https://doi.org/10.1016/j.atmosres.2008.02.016>.
- Maskey, M. L., C. E. Puente, B. Sivakumar, and A. Cortis, 2016: Encoding daily rainfall records via adaptations of the fractal multifractal method. *Stochastic Environ. Res. Risk Assess.*, **30**, 1917–1931, <https://doi.org/10.1007/s00477-015-1201-7>.
- Menne, M. J., I. Durre, R. S. Vose, B. E. Gleason, and T. G. Houston, 2012: An overview of the Global Historical Climatology Network–Daily database. *J. Atmos. Oceanic Technol.*, **29**, 897–910, <https://doi.org/10.1175/JTECH-D-11-00103.1>.
- Perone Pacifico, M. P., C. Genovese, I. Verdini, and L. Wasserman, 2004: False discovery control for random fields. *J. Amer. Stat. Assoc.*, **99**, 1002–1014, <https://doi.org/10.1198/0162145000001655>.
- Risser, M. D., C. J. Paciorek, and D. A. Stone, 2019a: Spatially dependent multiple testing under model misspecification, with application to detection of anthropogenic influence on extreme climate events. *J. Amer. Stat. Assoc.*, **114**, 61–78, <https://doi.org/10.1080/01621459.2018.1451335>.
- , —, M. F. Wehner, T. A. O'Brien, and W. D. Collins, 2019b: A probabilistic gridded product for daily precipitation extremes over the United States. *Climate Dyn.*, **53**, 2517–2538, <https://doi.org/10.1007/S00382-019-04636-0>.
- Schwartzman, A., and X. Lin, 2011: The effect of correlation in false discovery rate estimation. *Biometrika*, **98**, 199–214, <https://doi.org/10.1093/biomet/asq075>.
- Shu, H., B. Nan, and R. Koeppe, 2015: Multiple testing for neuroimaging via hidden Markov random field. *Biometrics*, **71**, 741–750, <https://doi.org/10.1111/biom.12329>.
- Sun, W., and T. T. Cai, 2009: Large-scale multiple testing under dependence. *J. Roy. Stat. Soc.*, **71B**, 393–424, <https://doi.org/10.1111/j.1467-9868.2008.00694.x>.

- , B. J. Reich, T. Cai, M. Guindani, and A. Schwartzman, 2015: False discovery control in large-scale spatial multiple testing. *J. Roy. Stat. Soc.*, **77B**, 59–83, <https://doi.org/10.1111/rssb.12064>.
- Timmermans, B., M. Wehner, D. Cooley, T. O'Brien, and H. Krishnan, 2019: An evaluation of the consistency of extremes in gridded precipitation data sets. *Climate Dyn.*, **52**, 6651–6670, <https://doi.org/10.1007/s00382-018-4537-0>.
- Westra, S., L. V. Alexander, and F. W. Zwiers, 2013: Global increasing trends in annual maximum daily precipitation. *J. Climate*, **26**, 3904–3918, <https://doi.org/10.1175/JCLI-D-12-00502.1>.
- van der Wiel, K., and Coauthors, 2017: Rapid attribution of the August 2016 flood-inducing extreme precipitation in south Louisiana to climate change. *Hydrol. Earth Syst. Sci.*, **21**, 897–921, <https://doi.org/10.5194/hess-21-897-2017>.
- Zhang, X., J. Wang, F. W. Zwiers, and P. Ya. Groisman, 2010: The influence of large-scale climate variability on winter maximum daily precipitation over North America. *J. Climate*, **23**, 2902–2915, <https://doi.org/10.1175/2010JCLI3249.1>.