

# Spatial-Temporal Extreme Modeling for Point-to-Area Random Effects (PARE)

CARLYNN FAGNANT<sup>1</sup>, JULIA C. SCHEDLER<sup>1</sup>, AND KATHERINE B. ENSOR<sup>1,\*</sup>

<sup>1</sup>*Department of Statistics, Rice University, Houston, TX 77005, U.S.A.*

## Abstract

One measurement modality for rainfall is a fixed location rain gauge. However, extreme rainfall, flooding, and other climate extremes often occur at larger spatial scales and affect more than one location in a community. For example, in 2017 Hurricane Harvey impacted all of Houston and the surrounding region causing widespread flooding. Flood risk modeling requires understanding of rainfall for hydrologic regions, which may contain one or more rain gauges. Further, policy changes to address the risks and damages of natural hazards such as severe flooding are usually made at the community/neighborhood level or higher geo-spatial scale. Therefore, spatial-temporal methods which convert results from one spatial scale to another are especially useful in applications for evolving environmental extremes. We develop a point-to-area random effects (PARE) modeling strategy for understanding spatial-temporal extreme values at the areal level, when the core information are time series at point locations distributed over the region.

**Keywords** *CAR model; change-of-support; extended-Hausdorff distance metric; geospatial modeling; spatial-temporal extremes*

## 1 Introduction

Spatial data can be observed over different spatial extents, leading to different types of spatial data and different statistical models. One type is point referenced data, or observations of the feature of interest indexed by a coordinate location. For example, whether or not a household experienced adverse outcomes during a major storm indexed by the latitude and longitude coordinates corresponding to the household. Modeling of point referenced data incorporates spatial dependence amongst the observations. For example, the covariance between two observations can be defined as a function of the distance between those two observations, and predictions of observations at new locations accounts for this spatial dependence. The consequences of ignoring spatial dependence includes biased estimates of standard errors, which can lead to incorrect inference. In this paper, we investigate rainfall data collected from fixed-location rain gauges in the Greater Houston area, revealing discernible spatial dependencies.

Even when spatial dependence in point-level data is properly accounted for, understanding regional spatial phenomena remains valuable for informed decision-making. Areal data, observed spatially across regions, often involves aggregating point-level data. For instance, in the case of the Hurricane Harvey Registry, responses collected at the point level (addresses) regarding adverse outcomes from the storm can be aggregated to the areal level (neighborhoods) to offer a city-wide summary of the storm's impacts (Miranda et al., 2021). From a modeling perspec-

---

\*Corresponding author. Email: [ensor@rice.edu](mailto:ensor@rice.edu).

tive, incorporating spatial dependence in areal data requires specifying an adjacency matrix, indicating which regions are considered neighbors and thus inducing dependence between them. Various neighbor structures have been explored (Getis and Aldstadt, 2004), with a common approach being to define neighbors as regions sharing a boundary point with a given region. The adjacency matrix can be weighted to form a spatial weight matrix, allowing a region's neighbors to exert spatially informed influences on its estimate in the model. For example, if region A has two neighbors, regions B and C, the impact of B and C on A need not be identical. If region B is larger than region C, it can be given a higher weight in the spatial weight matrix, enhancing its influence (Getis, 2009). Similar to point-referenced data, neglecting positive spatial dependence in areal data results in underestimation of parameter uncertainty and hence flawed statistical inference. In this study, we focus on analyzing data at the hydrologic region level, as these regions are scientifically meaningful in the context of rainfall (ACECHouston, 2019).

For our purposes, we observe time series of rainfall levels at fixed location gauges, but produce rainfall estimates for flood modeling and management decisions that are made at well-defined hydrologic areal regions. Our fundamental question is, how can the underlying distribution of rainfall be aggregated to the hydrologic level? This is called the change-of-support problem in spatial statistics (Gelfand et al., 2001).

Much of the change-of-support literature focuses on obtaining estimates of the mean or median, in other words measures of center or "typical" values. In the context of flood risk, the interest is instead in estimating extreme events. When estimating risk, accounting for high-impact but lower probability tail events is of interest rather than the most likely scenario. Utilizing the hierarchical modeling framework advanced in Craigmire (2014) and highlighted in Cressie and Wikle (2011), we incorporate extreme value modeling, namely the generalized Pareto distribution and peak-over-threshold model (GPD+POT), as the data model to capture the extreme value nature of time series of point-level measured rainfall. We rely on asymptotic normality of the parameter estimates of the GPD for the Gaussian characterization of the process model in our hierarchy (Hosking and Wallis, 1987).

Our unique contribution is the synthesis of an end-to-end approach for estimating important regional hydrologic parameters in flood modeling, making use of ideas of hierarchical spatial modeling, the change-of-support framework as described in Cressie and Wikle (2011), and conditional autoregressive (CAR) modeling to create a unique point-to-area-random-effects (PARE) model. Of particular importance is that our modeling is performed using spatial weights determined by the extended Hausdorff distance (Schedler, 2020b).

We apply three methods to move from point-referenced geostatistical data to achieve extremal inference at the areal level, namely our proposed PARE method, block kriging (Cressie, 2006), and a simple regional maximum approach. Block kriging is regularly used to address change-of-support issues in spatial modeling and the regional max model is often used by hydrologist studying extreme rainfall. We do not provide an extensive comparison of our PARE method with other strategies to address the spatial change-of-support issue. We apply our modeling paradigm to extreme rainfall for a large geographic region, specifically Houston, TX, with the goal of understanding the temporal evolution of extreme rainfall levels by hydrologic regions and accounting for spatial dependence. The methodology is adaptable to any problem where areal aggregation of point-level data is required.

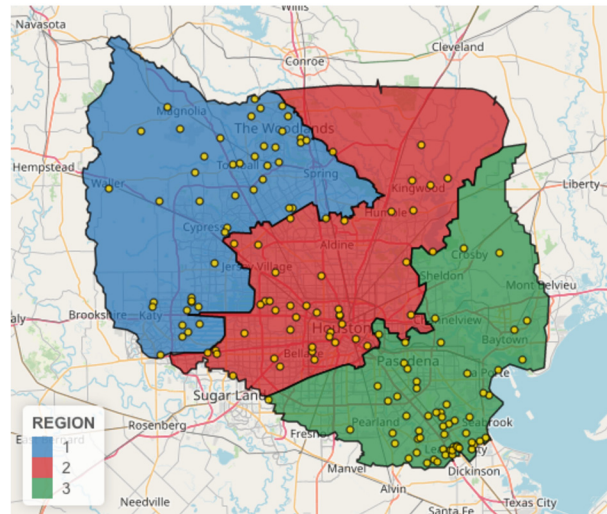


Figure 1: The three Harris County hydrologic regions are depicted by blue (region 1), red (region 2), and green (region 3) areas. Rain gauge locations for measured rainfall are represented by the yellow dots on the map.

## 2 Understanding Our Case Study, Hydrologic Regions, and Return Levels

The findings of this paper hold significant implications for flood hazard mapping in the broader Houston, Texas area as well as subsequent analyses aimed at reducing flood risk. For example, they offer valuable insights into the fundamental design criteria for various hydraulic structures such as culverts, roadway drainage systems, bridges, and small dams. This study also complements recent literature that delves into attributing flood risk in Southeast Texas and coastal Louisiana to factors like climate change and urbanization (Risser and Wehner, 2017; Van Der Wiel et al., 2017; Van Oldenborgh et al., 2017; Wang et al., 2018; Zhang et al., 2018; Sebastian et al., 2019), and it is particularly relevant given the recent release of NOAA Atlas 14 results for the State of Texas (Perica et al., 2018).

To contribute to the scientific conversation in support of the resilience of the greater Houston, Texas, area we estimate rainfall return levels for three hydrologic regions depicted in Figure 1.

### 2.1 Return Levels

Return levels represent extreme rainfall statistics, serving to link a specific duration of rainfall with a defined time span. While the common duration of interest is daily or 24-hour rainfall, assessments of flood hazards can involve other durations such as 1-hour, 12-hour, 2-day, or 3-day intervals.

An  $N$ -year return level signifies the amount of rainfall expected to be equaled or surpassed, on average, once every  $N$  years, where  $N$  denotes the return period. This essentially indicates the rainfall magnitude with a probability of  $\frac{1}{N}$ , of being exceeded in any given year.

It is important to recognize that the  $N$ -year return level is influenced by the duration of analysis. Thus, one would anticipate a 100-year 24-hour rainfall to exceed a 100-year 12-hour

rainfall at a specific location. For our purposes, we consider 25-year, 100-year, and 500-year 24-hour rainfall return levels. These levels are often referred to as the 25-, 100-, and 500-year flood plane levels.

Our estimated return levels, with uncertainty bounds, are derived from our spatial hierarchical modeling, coupled with extreme value models as the data model for the time series of rainfall measurements at each rain gauge. These estimated return levels are key inputs for the large scale flood models in operation for the greater Houston area (Fagnant et al., 2020).

## 2.2 Data Collection and Processing

We derive historical precipitation data from NOAA’s Global Historical Climatology Network (GHCN)-Daily product version 3.22 (Menne et al., 2018), a comprehensive database aggregating daily observations from land surface stations worldwide. These data undergo stringent quality assurance procedures and are procured from various National Meteorological and Hydrological Centers (NMHCs) globally. Station records range from less than a year to over 175 years. Our analysis focuses on a study area delineated within coordinates approximately 28.1°N to 31.0°N, 92.7°W to 97.0°W, coinciding with specified hydrologic regions shown in Figure 1. Within this area, we identified 149 stations providing data from January 1, 1900, to December 31, 2020. We compiled a dataset comprising daily precipitation totals from these 149 stations during this period, accessible at [doi.org/10.25612/837.XVGJ30NMA45X](https://doi.org/10.25612/837.XVGJ30NMA45X) [United States National Oceanic And Atmospheric Administration et al. (2020)].

Subsequent to initial data collection, the rainfall time series undergo processing to identify independent rainfall events. As the extreme value theory we employ assumes the independence of each occurrence, preprocessing is necessary to mitigate correlations between events. Such correlations arise from single storm events generating multi-day precipitation, wherein more than one day surpasses the specified threshold for our models. In such cases, despite the threshold being exceeded multiple times, there exists only one storm occurrence; hence, we retain only the highest 24 hour total from that storm. In simpler terms, if two or more consecutive days exhibit nonzero rainfall, only the day with the highest rainfall is retained, while the others are set to zero. Following this declustering process, the data are rendered independent, thus suitable for subsequent peaks-over-threshold analysis. This data adjustment, known as declustering, is a widely used technique in rainfall time series analyses (Gilleland and Katz, 2016; Fagnant et al., 2020).

## 3 Methods

The models presented are hierarchical, consisting of three steps: an extreme value modeling step, a spatial modeling step, and repeating the previous two steps for 3-different time epochs. In the data-process-parameter model approach advocated in Cressie and Wikle (2011), the extreme value modeling serves as our data model, and then the spatial structure of step two is captured in our process model. Although the data will have a spatial index for all steps, spatial structure is only modeled in the second step.

### 3.1 Extreme Value Modeling

Our starting point is the time series observed at each of  $n$  spatial locations. A first step is to characterize each of these time series in an extreme value modeling context appropriate for

the application. For this paper, the application concerns rainfall in Southeast Texas, which has been successfully modeled using the generalized-Pareto peak-over-threshold distribution (GPD+POT) to a contiguous segment of each time series. A primary reason for the choice of GPD+POT is the fact that the Houston region experiences multiple extreme rain events within a year. The GPD+POT applied to 24-hour rainfall events allows for this unique feature for the region. A detailed characterization of the choice of this particular extreme value model for the problem at hand is the focus of the paper Fagnant et al. (2020).

The peaks-over-threshold (POT) approach selects all events above a threshold for better return period estimation, incorporating a broader range of rainfall events compared to say the block maxima method, which is another extreme value modeling approach. POT treats independent rainfall occurrences as identically distributed random variables and focuses on observations exceeding a specified threshold  $u$  to analyze extremes, resulting in a conditional distribution:

$$P(X > u + y | X > u) = \frac{1 - F(u + y)}{1 - F(u)}, \quad y > 0. \tag{1}$$

The conditional distribution converges to the generalized Pareto distribution (GPD) of the form:

$$P(X > x | X > u) = \left[ 1 + \frac{\xi(x - u)}{\tilde{\sigma}} \right]^{-1/\xi}, \quad x > u, \tag{2}$$

where  $\xi$  is the shape parameter. The scale parameter for the GPD+POT is a function of the generalized extreme value parameters and the threshold, namely ( $\tilde{\sigma} = \sigma + \xi(u - \mu)$ ). Here  $\mu$  is the location parameter and  $\sigma$  is the scale parameter for the Generalized Extreme Value (GEV) distribution. These are asymptotic relationships that hold for appropriate thresholds.

The GPD+POT approach, widely adopted in recent extreme rainfall analyses, owes its popularity to its capacity to encompass temporally precise rainfall data and its theoretical linkage to the GEV distribution.

Return levels for each rain gauge time series can be calculated directly from the fitted distribution parameters ( $\sigma$ ,  $\xi$ , and threshold  $u$ ) obtained during the peaks-over-threshold modeling, along with the estimated rate of occurrences over the threshold. We utilize the R package `extRemes` (Gilleland and Katz, 2016) to fit our GPD+POT and to calculate return levels and corresponding confidence intervals. A detailed characterization of the choice of GPD+POT, plus selection of the threshold to one-inch of rainfall is the provided in Fagnant et al. (2020).

For our spatial model rainfall case study we focus on three time epochs each of forty years, namely 1920 through 1960, 1950 through 1990, and 1990 through 2020. For each temporal window and each location, we obtain estimates for the parameters of shape, scale and rate of the GPD+POT. These GPD+POT model estimates can be viewed in the hierarchical modeling paradigm advanced in Cressie and Wikle (2011) as well as Cooley et al. (2012), where we consider the data model, process model and predictive distribution. In this case, the GPD+POT serves as the data model for our rainfall time series.

### 3.2 Spatial Models

For each of our three temporal windows we model the point-to-area spatial structure using three different spatial paradigms. These paradigms include our contribution of Point-to-area Random Effects (PARE) incorporating the extended Hausdorff distance, traditional Block Kriging, and the Regional Max model. The statistical literature on change of support is extensive. In addition,

we compare our PARE model to classic approaches used by engineers working in flood modeling (see ACEC estimates described in 5). Each paradigm is described below.

### 3.2.1 Model 1. Point-to-Area Random Effects (PARE)

For the first model, we make use of ideas behind the change of support framework as described in Cressie and Wikle (2011) and conditional autoregressive (CAR) modeling Besag (1974) to create a unique random effects model to move from the point-level to the area-level. Of particular importance is that our modeling is performed using weights determined by the extended Hausdorff distance.

A CAR model is a popular choice for accounting for spatial correlation when working with areal or lattice data, and can easily be extended to find the relationship between covariates measured at the same areal level. The CAR model accounts for spatial dependence between the areal units by specifying an  $r \times r$  spatial weight matrix, where  $r$  is the number of regions. If the  $ij^{th}$  entry of the spatial weight matrix is nonzero, regions  $i$  and  $j$  are considered “neighbors”, and the CAR model will use observations at region  $i$  to inform estimates at region  $j$  and vice-versa. Determining which regions are neighbors is often based on shared boundary points (contiguity) or centroid distance (e.g.,  $k$  nearest neighbors or inverse distance weighting). Instead, we apply a method derived from the Hausdorff distance, a distance metric defined on sets. Specifically, we pursue a specification for the weight matrix which uses the inverse of the median Hausdorff distance between regions, yielding neighbor weights based on the maximum of the median distances between each pair of regions. The median Hausdorff distance accounts for irregularities in the geometry or orientation of the regions by only considering the closest 50% of a region when computing pairwise distances. The median was chosen to give distances comparable to the commonly used centroid distance. (Schedler, 2020a).

We make use of the covariance structure defined in a CAR model for areal data, but instead bring everything to the dimension of the stations (point-level observations). In order to create an  $n \times n$  weight matrix, we must define distances from each point to the others. However, since we are interested in moving to the areal level, we instead define point-to-point distances by the associated area-to-area distance of the regions the points fall within. Particularly, we use the extended Hausdorff distance with  $f = 0.5$  (in other words, the median Hausdorff distance) between our three regions to define these distances. For example, the distance between a point in region 1 to a point in region 2 is defined as the median Hausdorff distance between regions 1 and 2. This means that there will be many repeat values in our distance matrix for every pair of points that fall within the same pair of regions. The matrix will look like a block version of the region-to-region distance matrix where each block is made up of a single repeated value.

Specifically, for the distance matrix, let  $D$  be the  $3 \times 3$  matrix defining the median Hausdorff distance  $d$  between the three regions.  $D$  can be found using the `hausMat` function in the `hausdorff` package (Schedler, 2020a). Let  $n_j$  be the number of stations in region  $j$ . Then  $D$  and  $D_{block}$  are defined as follows,

$$D = \begin{bmatrix} d[1, 1] & d[1, 2] & d[1, 3] \\ d[1, 2] & d[2, 2] & d[2, 3] \\ d[1, 3] & d[2, 3] & d[3, 3] \end{bmatrix}, \quad (3)$$

$$D_{block} = \begin{bmatrix} C_{n_1 \times n_1} & D[1, 2]_{n_1 \times n_2} & D[1, 3]_{n_1 \times n_3} \\ D[1, 2]_{n_2 \times n_1} & C_{n_2 \times n_2} & D[2, 3]_{n_2 \times n_3} \\ D[1, 3]_{n_3 \times n_1} & D[2, 3]_{n_3 \times n_2} & C_{n_3 \times n_3} \end{bmatrix}, \quad (4)$$

where each block in  $D_{block}$  is a constant value times a matrix of 1's of the specified dimension. For example,  $D[1, 2]_{n_1 \times n_2} = d[1, 2] * \mathbf{1}_{n_1 \times n_2}$ . The diagonal elements of  $D_{block}$  are identically a constant  $c$ , which we now define.

While the distance from a region to itself is technically zero, we define the distance between points within the same region to be a specified positive constant  $c$  that is less than the other region-to-region distances (i.e.  $c < d[1, 2], d[1, 3]$ ). We do this because there is spatial variability of points within each region. This construction ensures the weight matrix is invertible and still weights values within a region more than values from other regions. For our application, we choose  $c = 1$  mile.<sup>1</sup> The distance matrix is converted into a weight matrix  $W$  by taking the reciprocal of each element in the matrix in order to create inverse-distance weights. Overall, with the block structure of our weight matrix  $W$ , the observations within each region are weighted equally, but less than those stations within regions closer to the target region or within the target region itself. When using  $c = 1$ , the inverse distance weight matrix has all elements  $\leq 1$ . If using a different value for  $c$ , one could scalar normalize the matrix by dividing all entries of the matrix by the value of the maximum entry.

We bring in the change-of-support concept in the mean structure of the model through covariates by creating variables which are indicator functions describing which region each station falls within. Let  $\mathbf{Z}$  be the vector of point-level observations (extreme value parameter estimates for each rain gauge location) at stations  $\mathbf{s} = (s_1, \dots, s_n)$ . Let  $\mathbf{Y}$  be the process of interest measured at the spatial areas/regions  $\mathbf{r} = (r_1, r_2, r_3)$ . The  $n \times 3$  matrix  $\mathbf{H}$  represents indicator variables for the regions the station falls within, in other words,

$$h_{ij} = \begin{cases} 1 & \text{if } s_i \in r_j \\ 0 & \text{otherwise.} \end{cases} \tag{5}$$

The traditional areal CAR model with covariates is defined through the following structure:

$$Y_j | y_{(-j)} \sim N\left(x_j \beta + \sum_{k \neq j} \rho w_{jk} (y_k - x_k \beta), m_{jj}\right), \tag{6}$$

where  $y_{(-j)}$  represents all areal values  $y$  except for  $y_j$ ,  $\beta$  is the coefficient for covariate  $x$ ,  $\rho$  is a spatial dependence parameter,  $w_{jk}$  are entries from the weight matrix  $W$ , and  $m_{jj} = \tau_j^2$ , the conditional variance of region  $j$ .

We will use a similar structure, but instead define it on the point-level. As such, we use it to model the observations  $\mathbf{Z}$  instead of the area-level process values  $\mathbf{Y}$ . Our covariates  $\mathbf{x}_i$  are represented by the matrix  $\mathbf{H}$ , so we can set  $X = \mathbf{H}$ . Our proposed point-to-area random effects (PARE) model takes the following structure:

$$Z_i | z_{(-i)} \sim N\left(\mathbf{x}_i' \boldsymbol{\beta} + \sum_{j \neq i} \rho w_{ij} (z_j - \mathbf{x}_j' \boldsymbol{\beta}), m_{ii}\right), \tag{7}$$

which can also be written as  $Z \sim N(X\boldsymbol{\beta}, \Sigma_{CAR})$  where  $\Sigma_{CAR} = (I - \rho W)^{-1}M$  and  $M = \text{diag}(\tau_1^2, \dots, \tau_n^2)$ . This model can be run in R using the `spautolm` function from the `spatialreg` package (Bivand et al., 2013) with family = ‘‘CAR’’. The required inputs include observations  $\mathbf{Z}$ , covariates  $X = \mathbf{H}$ , and the weight matrix  $W$ . The model output provides maximum likelihood

---

<sup>1</sup>The choice of  $c$  does not influence our final return level estimates. A sensitivity analysis is provided in Fagnant (2021).

estimates for  $\rho$  and  $\beta$ . We interpret the coefficient estimates  $\hat{\beta}$  to be the estimates of our process  $\mathbf{Y}$  at the areal level, i.e. the GPD scale, shape, and rate parameter estimates for the three regions.

One assumption the CAR model requires is that the covariance matrix  $\Sigma_{CAR}$  be symmetric. Additionally, since  $\Sigma_{CAR}$  is defined as  $(I - \rho W)^{-1}M$ , our weight matrix  $W$  must be invertible. With the current block setup for the weight matrix, our  $W$  is indeed symmetric. However, by construction our  $W$  has many rows and columns with the same exact values, thereby making the matrix singular and non-invertible. We address this issue through jittering, in other words, adding a normal random variable with small standard deviation of 0.1 to the extended Hausdorff distance matrix which has values 7.7, 7.9, and 27.6 for each of the three pairwise distances. The jittered matrix is made symmetric by replacing the upper triangular values with the lower triangular values, or vice versa.

### 3.2.2 Model 2. Block Kriging

For the second model, ordinary kriging is performed on the point-level extreme value estimates to obtain estimates of the extreme value parameters on a uniform fine grid. Integrating over the gridded points within each region reveals the “block” average, or the overall parameter estimate for each region. This approach is an established technique in the change-of-support literature to address the transition of spatial data from point-to-area (Gotway and Young, 2002; Craigmile, 2014), and is known more generally as block kriging (Cressie, 2006). We present this more established approach from the spatial literature so that we may compare it against our proposed PARE model.

We start by defining this block kriging model in terms of an underlying hierarchical model framework, and later describe how we estimate it in practice. We follow the notation of Cressie and Wikle (2011), but present it in terms of our application.

We summarize the hierarchical framework behind kriging of our GPD+POT parameter estimates below, where the predictive distribution gives the kriging estimates for new location  $s_0$ . In particular,  $Y^*(s_0)$  is the ordinary kriging predictor and  $\sigma_Y^2(s_0)$  the associated kriging variance, which are given below in equations (9).

In this paradigm, we let  $\mathbf{Z}$  be the vector of point-level observations (extreme value parameter estimates) at stations  $\mathbf{s} = (s_1, \dots, s_n)$ . Let  $\mathbf{Y}$  be the process of interest assessed on a uniform grid of points  $G$  spanning the regions. The data process  $Z$  is considered a noisy version of the underlying process  $Y$  with measurement error, i.e.  $Z = Y + \epsilon$  where  $\epsilon \sim N(0, \sigma_\epsilon^2)$ . Therefore we have the data model as  $[Z(s_i)|Y(s_i), \sigma_\epsilon^2] \sim N(Y(s_i), \sigma_\epsilon^2)$ .

Next we model the process at point  $s$ , or  $Y(s)$  with constant but unknown mean  $\mu_Y(s) = \mu$  and covariance  $C_Y(u, v) = \text{Cov}(Y(u), Y(v))$  for  $s, u, v \in D_s$ . Then the process model is:  $[Y(\cdot)|\mu, C_Y] \sim N(\mu, C_Y)$ .

Finally, the predictive distribution for a new point  $s_0$  is given by:

$$\begin{aligned} Y(s_0)|\mathbf{Z}, \mu, C_Y, \sigma_\epsilon^2 &\propto [Z(s_i)|Y(s_i), \sigma_\epsilon^2][Y(\cdot)|\mu, C_Y] \\ &\sim N(E[Y(s_0)|\mathbf{Z}], \text{var}(Y(s_0)|\mathbf{Z})) \\ &\sim N(Y^*(s_0), \sigma_Y^2(s_0)), \end{aligned} \tag{8}$$

where

$$\begin{aligned} Y^*(s_0) &= \mu + \mathbf{c}_Y(s_0)'C_Z^{-1}(\mathbf{Z} - \mu\mathbf{1}), \\ \sigma_Y^2(s_0) &= C_Y(s_0, s_0) - \mathbf{c}_Y(s_0)'C_Z^{-1}\mathbf{c}_Y(s_0) + \\ &\quad (1 - \mathbf{1}'C_Z^{-1}\mathbf{c}_Y(s_0))^2 / (\mathbf{1}'C_Z^{-1}\mathbf{1}). \end{aligned} \tag{9}$$



Defining the mean and covariance elements for  $Y$  we have:

$$\begin{aligned}
 C_Y(u, v) &= cov(Y(u), Y(v)) \text{ for } u, v \in D_s, \\
 \mathbf{c}_Y(s_0)' &= cov(Y(s_0), \mathbf{Z}) = (C_Y(s_0, s_1), \dots, C_Y(s_0, s_n)) \\
 &\text{and} \\
 \mu &= \hat{\mu}_{gls} = (\mathbf{1}'C_Z^{-1}\mathbf{Z}) / (\mathbf{1}'C_Z^{-1}\mathbf{1}).
 \end{aligned}
 \tag{10}$$

Finally,  $C_Z$  is an  $n \times n$  matrix where

$$C_Z(u, v) = \begin{cases} C_Y(v, v) + \sigma_\epsilon^2 & \text{if } u = v \\ C_Y(u, v) & \text{if } u \neq v \end{cases}
 \tag{11}$$

These equations for the ordinary kriging predictor follow the notation of Cressie and Wikle (2011). Note that the constant mean  $\mu$  is taken to be the generalized least squares estimate,  $\hat{\mu}_{gls}$ .

The next step in block kriging is to predict at a uniform grid  $G$  across the three hydrologic regions, giving a (discretized) predicted spatial surface of the entire region. We then average the predicted values over the respective regions  $\hat{y}(r_j) = \frac{1}{n_{r_j}} \sum_{g_k \in G \cap r_j} y(g_k)$  where  $n_{r_j}$  is the number of grid points in region  $j$ .

Turning to our specific application, we are interested in bringing extreme value analysis from the point observation level to the level of our hydrologic regions. Similar to the PARE model, we do this by bringing the extreme value parameter estimates to the region level and then calculate return levels based on the region level parameter estimates. Therefore, the kriging model above is performed where the observations are taken to be the extreme value parameter estimates for each station based on the univariate extreme value modeling described in 3.1.

One added benefit to the block kriging model not yet present in the PARE model is that kriging can be performed on multivariate data, meaning that the parameters can be modeled jointly if appropriate. In exploratory data analysis we find that the shape and log(scale) parameters are negatively correlated with each other, while the rate parameter does not show a relationship to either. In order to capture these relationships, we estimate the model above by cokriging the shape and log(scale) parameters jointly and kriging the rate parameter separately. The `gstat` package in R provides functions to fit a cross-variogram and perform cokriging (Pebesma, 2004).

### 3.2.3 Model 3. Regional Max

We discuss one final model, called the regional max model, which is a more simplified approach towards obtaining return level estimates for each region. This model is a data analytics approach one might take in order to summarize extreme data for a region. In particular, it takes the maximum daily rainfall value across all stations within a region to create a consolidated series of these maximum daily rainfall totals for each region, hence the name regional max. From the consolidated series, we can then perform traditional univariate extreme value analysis using the GPD+POT to obtain return levels estimates for each region.

We note that this model is a data analytic approach for which the mathematical theory has not been developed. Additionally, the Regional Max model is only spatial in the sense that maximums are taken over spatial subsets (the regions)– the spatial structure is not modeled directly. We include it as a simple comparison against our proposed PARE model, as well as

the traditional Block Kriging approach to change-of-support. We predict that since the regional max model uses the daily maximum values in combination with the GPD+POT, it will likely produce higher return level estimates than the previous models.

## 4 Simulations

Recall that the purpose of the models above is to bring information and extreme value modeling estimates from the point-referenced observations to the areal level of the three specific hydrologic regions in the greater Houston area. In order to test the performance of the models proposed, we simulate multiple data sets of similar spatial structure to run the models on. We simulate data assuming we know the true extreme value parameters at the hydrologic region level. This simulation paradigm is consistent with the requirements from the hydrologists, in that the goal is to understand the 25-, 100-, and 500-year return levels for each of the three regions. From these simulated data sets, in which we set the true values for each hydrologic region, we can test how well each spatial model paradigm recovers these true values.

In order to construct simulated data, we start from the end with our estimated true values, and work backwards step-by-step to create data that results in the proper structure. The first step taken is to observe the mean structure of the extreme value parameters of the data that we already have. Taking the GPD fits of the last 40 years of data (1981–2020) for the stations within the three hydrologic regions, we find the mean values of the shape and scale parameter estimates by region. We set these as the true parameter values for the regions, and aim to recover similar values when we run our models on the simulated data.

In particular, the simulated rainfall data are generated on a uniform grid with a resolution of three miles. This produces 314 points of daily rainfall data within the three hydrologic regions. The simulation grid is visible in Figure 2. All points falling within a region are assigned the same true shape and scale parameters, which are the average values calculated from the last 40 years of observed data.

The rate parameter is held constant in our simulation, which is supported by the result of

3-Mile Grid for Simulations




Figure 2: Map of the three hydrologic regions in Harris County (outlined in black), the station locations with data available for the 1981–2020 period (blue dots), and the grid locations where the simulated rainfall data are generated to (+ symbols).

fitting our three models to the rainfall data in Section 5. In particular, the rate parameter has the least deviation between regions or models, staying fairly constant across the larger area. The constant rate is set to be 0.0544, the observed average of the rate parameter across all stations in the last 40-year period (1981–2020).

For the next step, we use the parameter estimates at each location to simulate daily rainfall data from the GPD. Since we use the GPD to model independent exceedances above the threshold, we can similarly only generate data values that are independent observations above the threshold. As such, we only generate a proportion of the daily data— a proportion which is equal to our constant rate. For a 40-year period, we generate 795 daily values. In order to preserve the same rate parameter estimates for our fit, we fill the remaining daily values with zero representing a rain value below the threshold. We generate rainfall values (exceeding the threshold) from the GPD using the `revd` function from the `extRemes` package in R (Gilleland and Katz, 2016).

The generated rainfall data lacks time-reference since we did not assign the data values to specific dates. We can only guarantee that the generated observations are independent. The lack of time-referencing only poses a problem for the regional max model, since it works off of the raw time-referenced daily data and not the GPD fits at the station level. In the interest of evaluating the performance of the regional max model in capturing the true parameter values, we propose an additional step to reorder the simulated rainfall values to impose a pseudo-time-referencing. This approach makes use of the idea that for one day, if there is a large rainfall amount at one station, the other stations are more likely to also have larger rainfall amounts that day. In general, rainfall totals across stations in a single day will be correlated, though not perfectly. To impose temporal considerations, we rank each simulated series and then perturb the ranks with uniform random noise, and then re-rank each series to assign a temporal structure. This strategy induces temporal correlation in the simulated series. Therefore, the regional max approach can be implemented, taking the maximum across each “day” based on the region each station falls within.

### 4.1 Results from Simulation Study

Fifty iterations of the simulation are performed and the results displayed in Table 1. Model performance is evaluated using two metrics: the Root Mean Square Error (RMSE), and the Mean Absolute Error (MAE). If  $\hat{\xi}_1, \hat{\xi}_2, \dots, \hat{\xi}_{50}$  are the 50 estimates of the shape parameter for a given model,  $\xi$  is the true value of the shape parameter for the simulation, then the RMSE and MAE are calculated as follows:

$$\begin{aligned}
 RMSE &= \sqrt{\frac{\sum_{i=1}^{50} (\hat{\xi}_i - \xi)^2}{50}} \\
 MAE &= \frac{\sum_{i=1}^{50} |\hat{\xi}_i - \xi|}{50}.
 \end{aligned}
 \tag{12}$$

We obtain values of both the RMSE and MAE the shape and scale parameters for each combination of the three regions and each of the three models.

While the first two models both performed well at recovering the true parameter values, our PARE Model performs almost uniformly better than the block kriging model according to the RMSE and MAE. The block kriging model only performed better in estimating the scale parameter for region 3, according to the RMSE. The regional max model performed much worse

Table 1: Extreme value parameter estimates from our proposed models tested on the simulated data for 50 iterations. Bolding denotes the smallest value of either RMSE or MAE for estimating the given parameter for a given region. Note that the rate parameter is held constant, so rate estimates are not evaluated here. Included are the true parameter values used for simulation (truth), as well as the mean parameter estimates (mean), root mean squared error (RMSE), and mean absolute error (MAE) for each region and each model from the 50 iterations. Smaller values of RMSE and MAE indicate better performance of a model at recovering the true parameter values. Our proposed PARE model performs almost uniformly better according to these metrics.

		Region 1			Region 2			Region 3		
		Mean	RMSE	MAE	Mean	RMSE	MAE	Mean	RMSE	MAE
Scale	Truth	233.64			246.78			229.38		
	PARE	233.96	<b>1.3466</b>	<b>1.0536</b>	247.07	<b>1.3125</b>	<b>1.0537</b>	229.70	1.1994	<b>0.9160</b>
	Block Kriging	235.85	2.5194	2.2470	245.11	2.1282	1.8337	229.61	<b>1.1697</b>	0.9225
	Regional Max	250.84	17.5993	17.1950	264.27	17.8354	17.4853	247.01	17.9624	17.6309
Shape	Truth	0.2044			0.2319			0.1641		
	PARE	0.2013	<b>0.0056</b>	<b>0.0044</b>	0.2295	<b>0.0050</b>	<b>0.0041</b>	0.1621	<b>0.0044</b>	<b>0.0035</b>
	Block Kriging	0.1984	0.0072	0.0065	0.2261	0.0073	0.0061	0.1700	0.0071	0.0060
	Regional Max	0.3326	0.1291	0.1282	0.3676	0.1363	0.1357	0.2775	0.1143	0.1134

at capturing the true parameter values, especially that of the scale, just as we expected after seeing the large scale estimates in the results on observed data (see Table 2). The RMSE and MAE values for the regional max were around 10–30 times larger than those of the other two models.

## 5 Case Study Results

We apply the proposed models to our rainfall data for the last 40 years (1981–2020) and compare the regional extreme value estimates produced by each. The comparison of parameter estimates and return levels are available in Tables 2 and 3, respectively.

One discovery of note that makes the implementation of the block kriging model in R fast and simple is the option to input a *SpatialPolygonsDataFrame* as the “newdata” argument in the kriging functions of the `gstat` package (Pebesma, 2004). In particular, we can input the spatial polygons for our three hydrologic regions, and get the averaged estimates directly. When polygons are input as the new data to kriging, the function uses `sp::spsample` to randomly sample points uniformly across each polygon and calculates a block average (Pebesma, 2004; Bivand et al., 2013). This is intuitively the same calculation we are doing, without specifying a grid of points beforehand.

This built-in function implementation runs much faster than kriging to a fine grid and then averaging over the points within each region. For example, when cokriging the shape and  $\log(\text{scale})$  parameters, kriging to a fine grid of 19,791 points and then averaging the gridded points by region took over 15 minutes while kriging to the regions took only 3.7 seconds, a substantial improvement to computation time. One could definitely use a coarser grid to speed things up, but the `gstat` functions (`gstat::predict` and `gstat::krige` for multivariate and univariate kriging, respectively) with “newdata” set to the region polygons performs quickly and produces very similar estimates to the gridded version with 19,791 points. We compare the averaged parameter estimates for each region in Supplemental Table 1, where the values obtained from using the `gstat` function shortcut are displayed in brackets. Due to the similarity of the estimates (with the largest deviation being 0.16 in the scale parameter), we evaluate all

Table 2: Comparison of extreme value parameter estimates (standard errors in parentheses) across our proposed models for the last 40 years of data, 1981–2020. As we predicted, the regional max results differ more than the other two models with its scale parameters being much higher, which may lead to larger return level estimates.

Method	Parameter	Region 1	Region 2	Region 3
1. PARE Model	scale	215.08(4.3232)	236.68(4.8949)	223.99(3.6994)
	shape	0.2(0.0164)	0.16(0.0168)	0.19(0.0134)
	rate	0.06(0.0023)	0.04(0.0023)	0.06(0.0019)
2. Block Kriging	scale	208.83 (16.521)	228.30 (18.557)	229.41 (17.530)
	shape	0.2205 (0.0696)	0.1771 (0.0715)	0.1121 (0.0672)
	rate	0.0559 (0.0058)	0.0554 (0.0058)	0.0557 (0.0057)
3. Regional Max	scale	291.21 (14.897)	343.46 (16.778)	331.38 (16.747)
	shape	0.1523 (0.0377)	0.1397 (0.0355)	0.1351 (0.0366)
	rate	0.0565 (0.0019)	0.0601 (0.0020)	0.0557 (0.0019)

Table 3: Comparison of return level estimates (in inches) across the proposed models for the last 40 years of data, 1981–2020. Standard error estimates are displayed in parentheses. Additionally, the model estimates are compared to region estimates produced by the American council of engineering companies (ACEC) in Houston (ACECHouston, 2019). The ACEC estimates for regions 2 and 3 are marked with asterisks because the hydrologic regions are slightly different than the HCFCD designation. Our proposed PARE model produces very similar estimates to the ACEC study for region 1, which is the same as our region 1.

Method	Return Period	Region 1	Region 2	Region 3
1. PARE Model	25-Year	11.634(0.6874)	10.442(0.5857)	11.399(0.5461)
	100-Year	16.48(1.2435)	14.291(1.0239)	15.856(0.9666)
	500-Year	24.131(2.3097)	20.002(1.8186)	22.7(1.7504)
2. Block Kriging	25-Year	12.01 (2.27)	11.21 (2.06)	9.14 (1.44)
	100-Year	17.28 (4.48)	15.46 (3.89)	11.87 (2.54)
	500-Year	25.80 (8.97)	21.91 (7.36)	15.61 (4.44)
3. Regional Max	25-Year	12.96 (1.33)	14.68 (1.44)	13.75 (1.35)
	100-Year	17.54 (2.39)	19.68 (2.53)	18.37 (2.39)
	500-Year	24.23 (4.32)	26.83 (4.48)	24.94 (4.22)
ACEC Estimates	25-Year	10.90	11.50*	12.30*
	100-Year	16.30	16.90*	18.00*
	500-Year	24.20	25.00*	27.20*

future results using the faster method.

Additionally, since the PARE and block kriging models both perform modeling on the log of the scale parameter, the estimates from these models must be translated back to the regular scale. Since these estimates are derived from the idea of taking the mean of the log(scale), we use approximations to bring it back to the mean of the scale as opposed to just taking the

exponential value of the estimate. We use Taylor series approximations to transform back to the scale, which are outlined and derived in the appendix of Fagnant (2021).

After fitting each of our proposed models to the last 40 years of data, we compare results for the estimates across the three regions by looking at both the extreme value parameter estimates as well as the estimated return levels. In Table 2, we see that the scale parameter estimates for the PARE model and block kriging are similar, whereas those of the regional max model are much larger, by roughly 50%. For all three models, the shape parameter for region three is the smallest. Similarly, region three usually has the largest scale parameter. We also note that the rate parameter estimates vary the least across regions or models, which provides the justification for holding the rate parameter constant in the simulations in Section 4.1. The approximately constant estimates suggest the modeling of this parameter might be weighed against the added complexity of including it. If we want to simplify modeling, we could hold this parameter constant across regions.

From the estimated extreme value parameter estimates, we also calculate the 25-, 100-, and 500-year return level estimates (equivalent to the 4%, 1%, and 0.2% precipitation events, respectively) and display the results in the first 9 rows of Table 3. The regional max model produces the largest return level estimates, but has somewhat similar results to the PARE model. For the PARE model, region one exhibits the largest return levels. The return levels for region three in the block kriging approach are lower than expected.

An important distinction in the return level estimates is that the PARE model produces smaller standard errors, illustrating the viability of our approach. However, in order to calculate standard errors and confidence intervals for the return levels, one piece of information we need is the covariance matrix of the shape and scale parameters. Due to the nature of modeling the parameters separately in the PARE model, we set the covariance to be zero, despite the parameters being negatively correlated. Future development of this model will address joint models the two parameter estimates, so that the covariance can be used to improve our knowledge of the precision of the estimates.

We extend our modeling of return levels for each region to the temporal setting by incorporating the moving window approach used in Fagnant et al. (2020). Specifically, we repeat our modeling on data subset to different 40-year periods. In order to get a general idea of how these regional return level estimates have changed over time, we choose three 40-year windows spaced evenly across our period of record. The overlapping windows chosen are 1921–1960, 1951–1990, and 1981–2020. We provide tables of the parameter estimates for the last window in Table 2, and the first and second windows in Supplemental Tables.

We display return level estimates along with standard errors for the three models for each window of time. See Supplemental Table 4 for 1921–1960 and Supplemental Table 5 for 1951–1990. To visualize clearly how return levels have changed over time for the three hydrologic regions, we also plot the return level estimates by window as a simple time series, along with 95% pointwise confidence intervals. Time series plots of the 25-, 100-, and 500-year return levels estimates with 95% pointwise confidence intervals, for each of the three regions, are provided in Figures 3–5.

Region one has similar return levels for the final period (ending in 2020) across all models. The return level is around 24–25 inches for the 500-year and around 17 inches for the 100-year. The uncertainty associated with the block kriging and regional max models, makes interpretation of trends more difficult. For example, region two saw slight decreases in return levels from 1960 to 1990 for the block kriging model and the regional max model, but the uncertainty is large.

Considering the PARE model results, regions one and two exhibit a steady increase in return

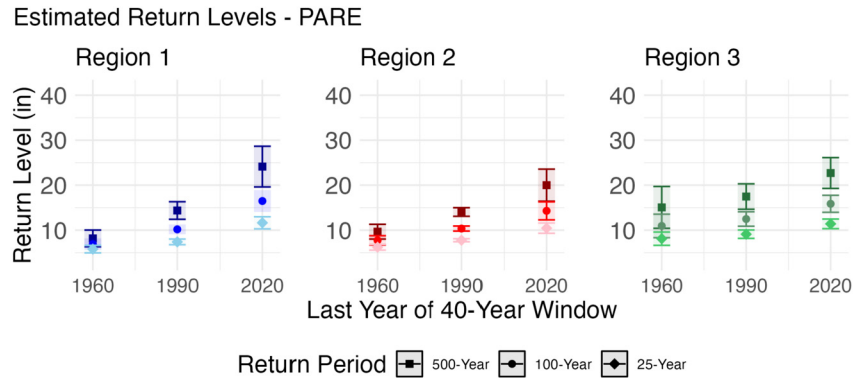


Figure 3: Return level estimates and 95% CIs for moving windows - PARE model. Regions are identified in Figure 1, and color coded in the same fashion.

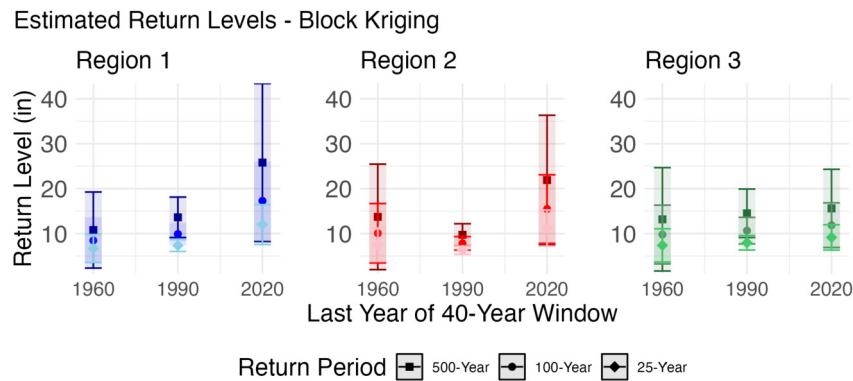


Figure 4: 24-hour return level estimates and 95% CIs for moving windows - block kriging. Regions are identified in Figure 1, and color coded in the same fashion. Values for visualization extracted from Fagnant (2021).

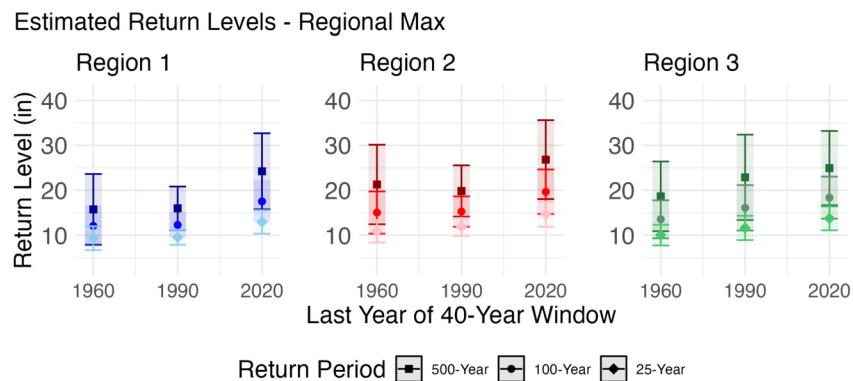


Figure 5: Return level estimates and 95% CIs for moving windows - regional max. Regions are identified in Figure 1, and color coded in the same fashion. Values for visualization extracted from Fagnant (2021).

levels since 1921. The increase in return levels from 1990 to 2020 was greater than that from 1960 to 1990. Region one represents the area to the northwest of Houston, TX, and region two covers the center part of the city. Region three saw the most gradual increase in extreme rainfall over time. It is worth noting, that region three is coastal and includes Galveston Island, which is regularly inundated with heavy rains.

Two interesting results are that the return level estimates for the PARE model demonstrate consistent trends, comparable to the findings in Fagnant et al. (2020), and also provide precise confidence intervals for the 25-, 100- and 500-estimated return levels for each region. For all regions and return levels, the PARE model has lower uncertainty estimates than the regional max and the block kriging approach. The regional max model in has very wide confidence intervals across all windows and regions. Similar to how we suspected the construction of the data using the maximum increases the return level estimates for the regional max model, it also increases the variability of these estimates and therefore the confidence interval widths.

There has been an effort in recent years to update extreme rainfall data in the Houston area. In particular, NOAA Atlas 14 released updated rainfall frequency estimates for all of Texas in 2018 (Perica et al., 2018). This updated data set provides gridded precipitation frequency estimates (equivalent to return levels) across Texas for multiple durations (ranging from 5 minutes to 60 days) and multiple return periods (from 1 to 1000 years). We wish to compare our regional return level estimates to updated region estimates for the Harris County hydrologic regions using this updated NOAA Atlas 14 data. As of the time of writing, the only such study found which also aims to bring return level estimates to the region level is a report produced by the American Council of Engineering Companies (ACEC) in Houston in 2019 (ACECHouston, 2019). The goals of this study slightly differed from ours, as they used the NOAA Atlas 14 frequency estimates directly, and also use these estimates to determine ideal regions for Harris County (i.e. grouping the watersheds into regions themselves). This resulted in almost the same designation as that of the Harris County Flood Control District (HCFCD) manual (Storey and Talbott, 2009) which we use, except that ACEC placed the Sims Bayou watershed in region 2 as opposed to region 3.

We compare our estimated regional return levels for the last 40-year window (1981–2020) to the ACEC study estimates in the last three rows of Table 3. We note that since regions 2 and 3 were slightly different between our analysis and theirs, we cannot compare these values directly, and so mark them with an asterisk. However, region 1 is designated exactly the same for both studies. We discover that our proposed PARE model matches the ACEC estimates more closely than the other models, indicating that the PARE model run on raw rainfall data performs similarly to the most recent advanced modeling of rainfall frequency estimates by NOAA. It is promising that estimates from our PARE model are consistent with these estimates derived from NOAA Atlas 14, because the NOAA report is the official source of precipitation frequency estimates provided for the U.S. government (Perica et al., 2018) and is widely used and accepted.

## 6 Discussion

We brought forward an end-to-end approach to estimating the return level, focusing on 25-, 100-, and 500-year levels for Houston, Texas. Our approach made use of hierarchical spatial modeling, and introduced point-to-area-random-effects or PARE to address the change of support. Our multi-tiered approach provides return level estimates for key regions in flood modeling along with uncertainty quantification. The results are consistent over our three time epochs, representing a



span of 100 years. When compared with the block kriging and regional max models, we obtain a clear picture of the rainfall dynamics in Houston, Texas. A key advantage of the PARE approach is the use of standard software, ensuring ease of use for future problems. The method also easily incorporates additional fixed effects.

## Supplementary Material

A pdf file containing the Supplemental Tables is included in the Supplemental Materials. Code to reproduce the PARE analysis for the three windows considered in this paper is provided as a Quarto project available for download on Github ([https://github.com/juliaSchedler/ST\\_PARE](https://github.com/juliaSchedler/ST_PARE)), or viewing as a rendered vignette at [https://juliaschedler.github.io/ST\\_PARE/](https://juliaschedler.github.io/ST_PARE/). Due to large file sizes, additional resources on the GPD fitting process for all 80 years' worth of rolling windows and models used for comparison to PARE are provided in another Github Repository ([https://github.com/carly-fagnant/Spatial\\_Extreme\\_Value\\_Modeling](https://github.com/carly-fagnant/Spatial_Extreme_Value_Modeling)).

## References

- ACECHouston (2019). Recommendation for: Rainfall Depths and Intensities in Harris County, *Technical Report*, ACEC Houston.
- Besag J (1974). Spatial interaction and the statistical analysis of lattice systems. *Journal of the Royal Statistical Society, Series B, Methodological*, 36(2): 192–236. Publisher: [Royal Statistical Society, Wiley].
- Bivand RS, Pebesma E, Gómez-Rubio V (2013). *Applied Spatial Data Analysis with R*, 2nd edition. UseR! Series. Springer.
- Cooley D, Cisewski J, Erhardt RJ, Mannshardt E, Jeon S, Omolo BO, et al. (2012). A survey of spatial extremes: Measuring spatial dependence and modeling spatial effects. *REVSTAT Statistical Journal*, 10(1): 135–165. Number: 1.
- Craigmile PF (2014). Spatial change-of-support and misalignment problems.
- Cressie N (2006). Block Kriging for lognormal spatial processes. *Mathematical Geology*, 38(4): 413–443. <https://doi.org/10.1007/s11004-005-9022-8>
- Cressie N, Wikle CK (2011). *Statistics for Spatio-Temporal Data*. John D. Wiley & Sons, Inc., Hoboken, NJ.
- Fagnant C (2021). Spatiotemporal Extreme Value Modeling with Environmental Applications, Thesis, Rice University. Accepted: 2021-10-06T16:17:51Z.
- Fagnant C, Gori A, Sebastian A, Bedient PB, Ensor KB (2020). Characterizing spatiotemporal trends in extreme precipitation in Southeast Texas. *Natural Hazards*, 104(2): 1597–1621. <https://doi.org/10.1007/s11069-020-04235-x>
- Gelfand AE, Zhu L, Carlin BP (2001). On the change of support problem for spatio-temporal data. *Biostatistics*, 2(1): 31–45. <https://doi.org/10.1093/biostatistics/2.1.31>
- Getis A (2009). Spatial weights matrices. *Geographical Analysis*, 41(4): 404–410. [\\_eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1538-4632.2009.00768.x](https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1538-4632.2009.00768.x)  
<https://doi.org/10.1111/j.1538-4632.2009.00768.x>
- Getis A, Aldstadt J (2004). Constructing the spatial weights matrix using a local statistic. *Geographical Analysis*, 36(2): 90–104. [\\_eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1538-4632.2004.tb01127.x](https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1538-4632.2004.tb01127.x) <https://doi.org/10.1111/j.1538-4632.2004.tb01127.x>

- Gilleland E, Katz RW (2016). extRemes 2.0: An extreme value analysis package in R. *Journal of Statistical Software*, 72(8): 1–39. <https://doi.org/10.18637/jss.v072.i08>
- Gotway CA, Young LJ (2002). Combining incompatible spatial data. *Journal of the American Statistical Association*, 97(458): 632–648. Publisher: [American Statistical Association, Taylor & Francis, Ltd.].
- Hosking JR, Wallis JR (1987). Parameter and quantile estimation for the generalized Pareto distribution. *Technometrics*, 29(3): 339–349. <https://doi.org/10.1080/00401706.1987.10488243>
- Menne MJ, Durre I, Korzeniewski B, McNeal S, Thomas K, Yin X, et al. (2018). Global Historical Climatology Network - Daily (GHCN-Daily), Version 3.
- Miranda ML, Callender R, Canales JM, Craft E, Ensor KB, Grossman M, et al. (2021). The Texas flood registry: A flexible tool for environmental and public health practitioners and researchers. *Journal of Exposure Science and Environmental Epidemiology*, 31(5): 823–831. Number: 5. Publisher: Nature Publishing Group.
- Pebesma EJ (2004). Multivariable geostatistics in S: The gstat package. *Computers & Geosciences*, 30(7): 683–691. <https://doi.org/10.1016/j.cageo.2004.03.012>
- Perica S, St Laurent M, Trypaluk C, Unruh D, Wilhite O (2018). Precipitation-Frequency Atlas of the United States. Volume 11, Version 2.0. Texas. *NOAA Atlas*. Publisher: United States, National Weather Service.
- Risser MD, Wehner MF (2017). Attributable human-induced changes in the likelihood and magnitude of the observed extreme precipitation during Hurricane Harvey. *Geophysical Research Letters*, 44(24): 12,457–12,464. <https://doi.org/10.1002/2017GL075888>
- Schedler JC (2020a). <https://github.com/juliaSchedler/hausdorff>.
- Schedler JC (2020b). Advances in the Analysis of Spatially Aggregated Data. Thesis, Rice University. Accepted: 2020-04-27T19:14:17Z.
- Sebastian A, Gori A, Blessing RB, Van Der Wiel K, Bass B (2019). Disentangling the impacts of human and environmental change on catchment response during Hurricane Harvey. *Environmental Research Letters*, 14(12): 124023. <https://doi.org/10.1088/1748-9326/ab5234>
- Storey A, Talbott M (2009). Harris County Flood Control District Hydrology and Hydraulics Guidance Manual, *Technical Report*, Harris County Flood Control District.
- United States National Oceanic And Atmospheric Administration, Fagnant C, Gori A (2020). Southeast Texas Precipitation.
- Van Der Wiel K, Kapnick SB, Van Oldenborgh GJ, Whan K Philip S, Vecchi GA, et al. (2017). Rapid attribution of the August 2016 flood-inducing extreme precipitation in south Louisiana to climate change. *Hydrology and Earth System Sciences*, 21(2): 897–921. <https://doi.org/10.5194/hess-21-897-2017>
- Van Oldenborgh GJ, Van Der Wiel K, Sebastian A, Singh R, Arrighi J, Otto F, et al. (2017). Attribution of extreme rainfall from Hurricane Harvey, August 2017. *Environmental Research Letters*, 12(12): 124009. <https://doi.org/10.1088/1748-9326/aa9ef2>
- Wang SYS, Zhao L, Yoon JH, Klotzbach P, Gillies RR (2018). Quantitative attribution of climate effects on Hurricane Harvey’s extreme rainfall in Texas. *Environmental Research Letters*, 13(5): 054014. <https://doi.org/10.1088/1748-9326/aabb85>
- Zhang W, Villarini G, Vecchi GA, Smith JA (2018). Urbanization exacerbated the rainfall and flooding caused by hurricane Harvey in Houston. *Nature*, 563(7731): 384–388. <https://doi.org/10.1038/s41586-018-0676-z>