

Stochastic tropical cyclone precipitation field generation

William Kleiber¹  | Stephan Sain² | Luke Madaus³ | Patrick Harr⁴

¹Department of Applied Mathematics,
University of Colorado, Boulder,
Colorado, USA

²Jupiter Intelligence, Boulder, Colorado,
USA

³Jupiter Intelligence, Madison, Wisconsin,
USA

⁴Jupiter Intelligence, San Mateo,
California, USA

Correspondence

William Kleiber, Department of Applied
Mathematics, University of Colorado,
Boulder, CO, USA.

Email: william.kleiber@colorado.edu

Abstract

Tropical cyclones are important drivers of coastal flooding which have severe negative public safety and economic consequences. Due to the rare occurrence of such events, high spatial and temporal resolution historical storm precipitation data are limited in availability. This article introduces a statistical tropical cyclone space-time precipitation generator given limited information from storm track datasets. Given a handful of predictor variables that are common in either historical or simulated storm track ensembles such as pressure deficit at the storm's center, radius of maximal winds, storm center and direction, and distance to coast, the proposed stochastic model generates space-time fields of quantitative precipitation over the study domain. Statistically novel aspects include that the model is developed in Lagrangian coordinates with respect to the dynamic storm center that uses ideas from low-rank representations along with circular process models. The model is trained on a set of tropical cyclone data from an advanced weather forecasting model over the Gulf of Mexico and southern United States, and is validated by cross-validation. Results show the model appropriately captures spatial asymmetry of cyclone precipitation patterns, total precipitation as well as the local distribution of precipitation at a set of case study locations along the coast. We additionally compare our model against a widely-used statistical forecast, and illustrate that our approach better captures uncertainty, as well as storm characteristics such as asymmetry.

KEYWORDS

empirical orthogonal function, hurricane, Lagrangian coordinate system, polar coordinates

1 | INTRODUCTION

Tropical cyclones (TCs) are major drivers of coastal flooding. Due to the extreme nature of these events (both in terms of magnitude and rarity), quantifying flood risk is a difficult task that requires a combination of atmospheric, hydrologic and statistical modeling. Flood risk from TCs is primarily driven by two factors: storm surge due to intense winds over the open ocean, as well as substantial sustained precipitation over land. Storm surge causes flooding along the immediate coastline while extreme precipitation over land also causes flooding away from the coast, and the two hazards together conspire to cause extreme water levels throughout large areas that extend far from the TC center (Lai et al., 2021; Wahl et al., 2015).

Wind fields associated with tropical cyclones have been well-studied and are measured directly via aircraft- and satellite-based observing systems when TCs are located over the open ocean. However, much less is known about

the forcing, distribution, and variation in tropical cyclone-induced precipitation over water and land (Elsberry, 2002). Nevertheless, TC-induced precipitation is a key factor in defining the upper tail of flood distributions throughout the eastern United States (Villarini & Smith, 2010) and TC-induced precipitation is projected to increase in future storms due to changes in atmospheric moisture content and increasing storm intensities (Knutson et al., 2019). Regardless of the impacts due to TC-induced precipitation, the lack of long observation records, the physical complexities associated with factors that control the precipitation, and large computational resources required to model such complex processes pose particularly difficult problems to quantify the impacts of TC-induced precipitation.

One approach to counter the lack of observational data on TCs is the use of simulated or synthetic storm tracks (Emanuel, 2006; Hall & Jewson, 2007; Nakamura et al., 2015; see Section 2.1 below for more detail) in a computational pipeline that uses these synthetic tracks in conjunction with geophysical models. These synthetic tracks provide only limited information about location and direction, speed, strength, and size of storms and reflect assumptions about the distribution of TC behavior. While this limited information is often enough to enable a computational approach for surge modeling, the lack of corresponding atmospheric conditions for these synthetic tracks makes direct physical modeling of the associated weather, specifically precipitation, difficult, if not impossible.

While the use of geophysical models pose significant challenges, there is an opportunity for statistical techniques that can generate realistic fields of TC precipitation (Rogers et al., 2009), and many of the differences between these models involve the availability of covariates on which the model conditions. The simplest models, to which our ensuing proposal should be compared, rely on only a small handful of statistics such as storm track, central pressure deficit and so forth. A popular statistical emulator is the Rainfall Climatology and Persistence (R-CLIPER) statistical model, trained on large sets of satellite-based observations of TC-induced precipitation Lonfat et al. (2004). This first-order model produces symmetric precipitation patterns about a TC center with no variation due to environmental interactions. Lonfat et al. (2007) modify R-CLIPER to include asymmetries about the storm center caused by environmental interactions. While simple to implement, the R-CLIPER versions do not produce precipitation patterns with sufficient detail to capture extreme variations due to environmental interactions, extreme storm intensity, or steep topography. A first model using simplified physical principles and detail in the atmosphere boundary layer was introduced by Langousis and Veneziano (2009). While an improvement over R-CLIPER, the model of Langousis and Veneziano (2009) is only applicable over water as it does not account for topography or over-land friction. Other approaches rely on detailed information about atmospheric structure that is not available in widely-used storm track databases. For example, Zhu et al. (2013) build the tropical cyclone rainfall (TCR) model that incorporates topography, land-induced friction, and TC characteristics such as varying intensity and size. Lu et al. (2018) compare TCR with full-physics, high-resolution simulations of TC-induced precipitation from a numerical weather model and found that TCR produced detailed spatially- and temporally-varying precipitation patterns. However, the main drawback to TCR is that it requires substantial information about the atmosphere, for example wind structure, gradient wind, shear, topography as well as saturation specific humidity. Such variables are not available in common storm track databases.

To motivate the statistical approach outlined here, it is helpful to examine some partial realizations of tropical cyclones that directly impact the Texas coastline near the city of Houston. Figure 1 shows two such example tropical cyclone events from an atmospheric model detailed in the next section. Each row contains three snapshots of hourly-integrated precipitation along with the historical storm track, and estimated storm center. From simply a statistical point of view these images present some serious obstacles: precipitation exhibits a mixed discrete-continuous distribution whose spatial intensity and signature change as a function of the storm size, intensity, and track. While over open ocean, precipitation will vary with sea-surface temperatures and general atmospheric conditions such as vertical wind shear (Ueno, 2007). When the TC approaches land and makes landfall, interactions with topography and increased friction will alter precipitation patterns and intensities. Standard space-time statistical models are not well-suited for a non-Gaussian dynamical process such as those exhibited in Figure 1.

Our approach is to focus on the behavior of the storm, with respect to its estimated center, in a Lagrangian coordinate framework. The precipitation patterns seen in Figure 1 show a degree of spatial structure in a circular fashion with respect to the center of the storm but also includes an asymmetric structure around the storm track. Modeling the process in polar coordinates presents some challenges, including the requirement of polar-indexed processes. The basic structure of the model outlined in Section 3.1 is not too surprising for a reader familiar with the space-time statistics literature, including separation of mean and stochastic variation components. However, extant space-time models for precipitation are not well suited to the type of structures seen in Figure 1; this work introduces necessary novel statistical innovations including a new polar-indexed process convolution representation, as well as the incorporation of geophysical predictors in a statistical learning framework.

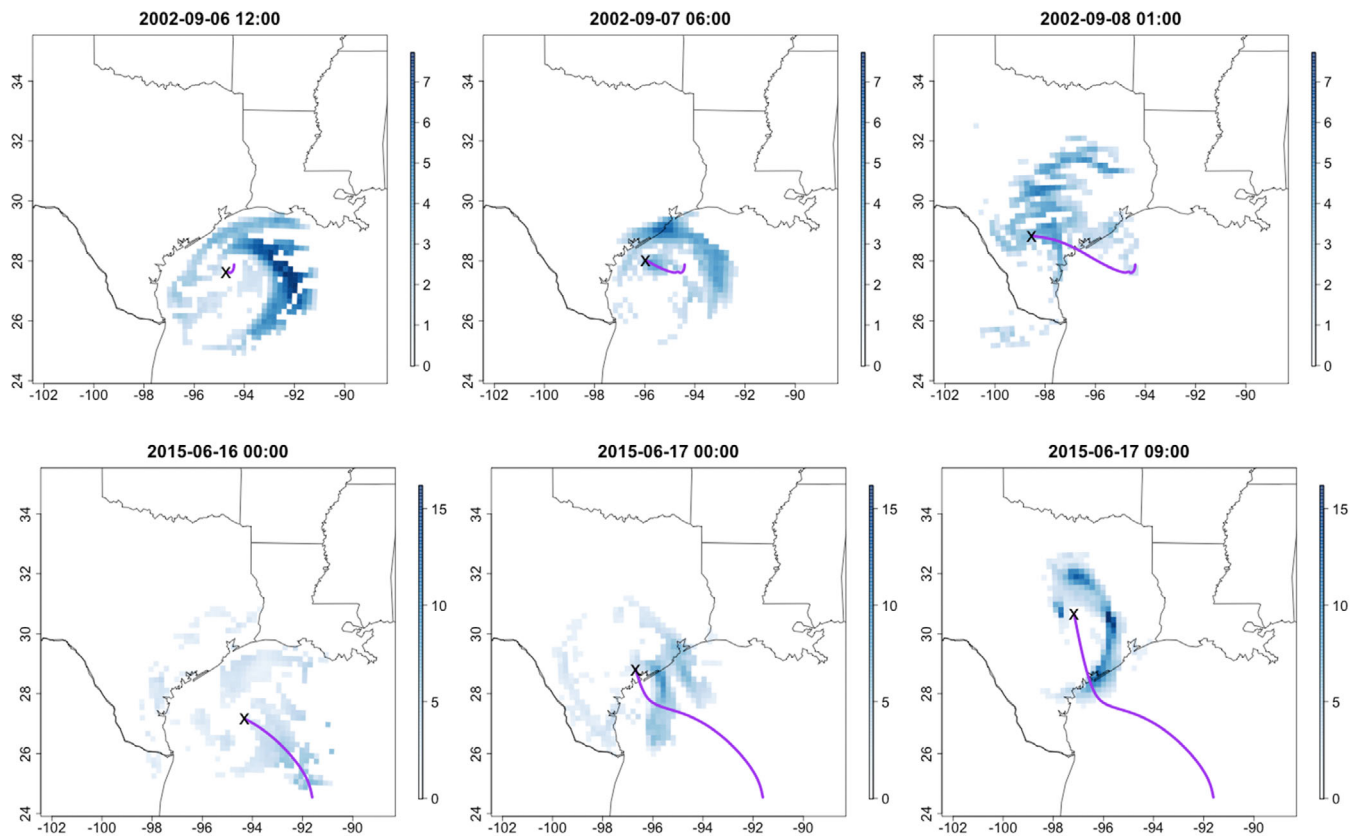


FIGURE 1 Each row contains example snapshots of hourly-integrated tropical cyclone precipitation; the purple line is the storm track to the current snapshot, while the “x” represents the storm center at the hour displayed. The top row is hurricane Fay in September 2002, while the bottom row is hurricane Bill in June 2015. Units are mm.

The next section details the data used in our study, Section 3 describes the model and estimation procedure while Section 4 contains the illustration of our approach, including a comprehensive cross-validation study as well as comparison against an extant statistical precipitation model. The final section discusses possible future routes of research.

2 | OBSERVATIONAL AND MODELED TC DATA

The limited information in TC storm tracks are the basis for inputs to the precipitation generator outlined in the previous section and are available in both historical and simulated databases. For training the precipitation generator, however, full space-time fields of precipitation that align with historical TCs are necessary. In what follows, we describe available storm track databases, observed and modeled TC precipitation data, as well as specifics on the training data used in the illustration in Section 4.

2.1 | Storm track databases

Detailed information regarding atmospheric conditions during TCs are difficult to procure due to the limited historical record and localized nature of such events. However, simple summary statistics regarding a storm’s evolution are much more readily available, and in fact are commonly logged in so-called storm track databases. Such statistics include general geophysical descriptors of a given TC behavior at any point in time, typically including location of the storm center, maximum wind speed (measure of intensity), direction and speed of motion, central pressure deficit with respect to the large-scale environment (a measure of strength) and radius of maximum winds (a measure of storm footprint). For each TC that occurs worldwide, such descriptors are archived in the International Best Track Archive for Climate Stewardship

(IBTrACS, Knapp et al., 2010). The IBTrACS database is constructed based on official TC analyses as defined by Regional Specialized Meteorological Centers (RSMCs) of the World Meteorological Organization (WMO).

While IBTrACS provides a source of actual worldwide TC activity over approximately the past 100 years (depending on ocean basin), synthetic TC data can be generated using a variety of methods (Emanuel, 2006; Hall & Jewson, 2007; Nakamura et al., 2015) and are stored in formats that mimic the IBTrACS to enable consistency among data sets and allow a more thorough study of TC behavior and impacts. Recently some authors have begun considering the problem of generating synthetic storm track data using stochastic models (Rekabdarkolaei et al., 2019).

2.2 | Existing TC precipitation data and limitations

There are four basic sources of data for TC precipitation: in situ observations, radar-derived estimates, satellite-derived estimates, and atmospheric model output. This study focuses on a statistical framework for generating complete space-time fields of TC-driven precipitation. Rain gauge data are not optimal for such a study, given their sparse spatial sampling and potential biases or measurement errors, especially in windy conditions such as during a TC (Pollock et al., 2018; Villarini et al., 2008). Radar-derived estimates such as those from National Center for Environmental Prediction (NCEP) Stage-IV precipitation analysis (Lin & Mitchell, 2005) are lacking over open water where TCs form, and hourly estimates lack quality control (Nelson et al., 2016). Satellite-derived estimates such as from the Global Precipitation Mission are gridded and available at high time frequencies, but have been found to underestimate heavy rainfall in coastal areas (Omranian et al., 2018). This leaves atmospheric model output as a prime candidate for developing and training a statistical approach.

Our approach to studying TC-driven precipitation is to exploit advanced atmospheric models such as the Weather Research and Forecasting (WRF) model developed at the National Center for Atmospheric Research (Skamarock et al., 2019). Dynamic weather models such as WRF use physically-based equations to advance an atmospheric state forward through time, simulating the likely evolution of a number of weather-relevant features including temperature, winds, pressure, moisture, and precipitation. These simulations prescribe and advance the “state” of the atmosphere as defined on a regular geospatial grid.

While global-scale weather simulations with these models are possible, they typically require enormous computational expense; for this reason scientists interested in spatially localized weather features such as TCs (Davis et al., 2008) typically use the WRF in a regional simulation mode, only simulating a small portion of the globe at a time. This still requires a fair amount of computational resources, but a reasonable-quality simulation of a single multi-day tropical cyclone event may be accomplished on a laptop within a few hours time (Hacker et al., 2017). In addition, such simulations require initial and boundary conditions for the model that are consistent with the environment necessary to develop and maintain a tropical cyclone. These initial and boundary conditions (i.e., space-time fields of temperature, wind, moisture, and pressure at all grid points in the model grid) must maintain a dynamical balance to avoid instabilities in the model. However, it is difficult to constrain all the factors in the environment that lead to a TC developing to a particular strength or taking a particular path while retaining dynamical balance, making it hard to synthetically generate initial and boundary conditions in a systematic way. Rather, these are typically extracted from other, global scale forecasts or historical “reanalyses” of weather conditions (e.g., ERA-Interim; Dee et al., 2011) which are often too coarse to well-resolve the tropical cyclone itself, requiring the WRF model to properly simulate the tropical cyclone structure.

2.3 | Training data

To provide training data for this model, simulations of tropical cyclone events were performed with the WRF model, which has been shown to produce realistic simulations of historical tropical cyclone events, including their precipitation characteristics (e.g., Davis et al., 2008, 2010). The WRF model is used to simulate seven historical TC events that impacted the Texas Gulf Coast, in particular hurricanes Fay (2002), Claudette (2003), Erin (2007), Humberto (2007), Edouard (2008), Ike (2008) and Bill (2015). The WRF model has a number of configuration options that can be tuned to change its performance, and such tuning was done here to optimize the WRF configuration to simulate these TC events. The exact WRF model configuration is described in Appendix A. Example snapshots of 1-h accumulated precipitation output on the WRF grid is shown in Figure 1.

We additionally require a set of features that match information typically contained in storm track databases such as location of storm center, radius of maximal wind, central pressure deficit, and storm direction. These quantities are calculated directly from the WRF output, and the derivations are detailed in the Appendix A.

The WRF data are thus available for $n = 7$ different TCs, or events, at hourly-spaced time points $t = 1, \dots, T_e$ the total number of which depends on the particular event. In this application, the number of hours in a TC ranges between 72 and 108 h; altogether there are $T = \sum_{e=1}^n T_e = 576$ observed realizations. The WRF data are on a 59×59 longitude-latitude grid, representing a 36 km spacing over the Gulf of Mexico region, extending into the southern United States (see e.g., Figure 6).

3 | METHODS

The past twenty years have seen a substantial increase in developing space-time statistical models. Much literature has focused on investigating valid space-time covariance functions, Porcu et al. (2021) give a nice up-to-date overview. However, space-time Gaussian random field models are not, by themselves, appropriate for the type of dynamical and nonnegative behavior seen in Figure 1. The moving nature of the storm beckons a Lagrangian framework, and indeed although some authors have investigated such models previously (Alegria & Porcu, 2017; Gneiting, Genton, & Guttorp, 2007; Rodriguez-Iturbe et al., 1987), the attention has been restricted to standard Gaussian random field models that advect in time.

To the best of our knowledge, no authors have developed a stochastic model for TC-induced precipitation fields. A related, but distinct, problem in understanding TC risk is in developing models for associated wind vector fields. Reich and Fuentes (2007) introduce a multivariate semiparametric model for surface U- and V-direction wind components driven by hurricanes. Their model conditions on satellite measurements, as well as a physical model for wind vectors. Modlin et al. (2012) use a polar coordinate representation, and exploit a standard wrapped normal circular statistical model. Our approach introduces a new polar coordinate stochastic model relying on a process convolution (Fuentes & Smith, 2001) that allows for high-dimensional space-time simulations, accommodating nonstationarity, and asymmetry in precipitation fields.

Our interest focuses on hourly integrated precipitation amount at a spatial grid cell indexed by its center, \mathbf{s} , at time point t . A standard approach would directly model such precipitation, denoted $P(\mathbf{s}, t) = P(s_{lon}, s_{lat}, t)$, where, over a limited domain such as ours, latitude and longitude can be projected to a Euclidean coordinate system (or directly treated as Cartesian near the equator). However, modeling in Euclidean coordinates is difficult due to the complex dynamic nature of tropical cyclones that involves rotation about a moving storm center with asymmetric precipitation patterns.

3.1 | Statistical model

We propose modeling P in a Lagrangian coordinate system with respect to the space-time varying storm center. That is, we model $P(r, \theta, t)$ where (r, θ) are polar coordinates from the storm center at time t . Precipitation amounts are non-Gaussian and nonnegative; to this end we exploit a common integral transformation technique (sometimes called an anamorphosis or copula),

$$P(r, \theta, t) = \begin{cases} G^{-1}(F(Y(r, \theta, t))), & Y(r, \theta, t) > 0, \\ 0, & Y(r, \theta, t) \leq 0, \end{cases} \quad (1)$$

where $Y(r, \theta, t)$ is a space-time Gaussian process, G^{-1} is a gamma quantile function, and F is a cumulative distribution function (cdf). The use of a gamma distribution to model quantitative precipitation is common in the stochastic weather generator literature (Kleiber et al., 2012; Stern & Coe, 1984). The inner transformation F is the cdf of the mixture distribution $f(y) = \int f_{r,\theta,t}(y) dr d\theta dt$ where $f_{r,\theta,t}$ is the pdf of $[Y(r, \theta, t) | Y(r, \theta, t) > 0]$. The model for Y follows a familiar mean-plus-residual decomposition that is standard in the spatial modeling literature, but with details adapted to modeling space-time processes in polar coordinates:

$$Y(r, \theta, t) = \sum_{\ell=1}^L c_{\ell}(t) \varphi_{\ell}(r, \theta) + Z(r, \theta, t), \quad (2)$$

for time-varying stochastic coefficients $\{c_\ell(t)\}_{\ell=1}^L$, fixed basis functions $\{\varphi_\ell(r, \theta)\}_{\ell=1}^L$ and a space-time residual stochastic process $Z(r, \theta, t)$. The model (2) assumes a stochastic mean function expressed as a weighted combination of spatial patterns.

Basis decompositions as in (2) are very popular due to their computational benefits, as well as abilities to approximate processes from traditional spatial statistical models. They are also widely applicable and can be used to model nonstationary and heterogeneous processes (Bandyopadhyay & Lahiri, 2010; Cressie & Johannesson, 2008; Lindgren et al., 2011; Nychka et al., 2015; Wikle, 2010). We first concern ourselves with specification of the mean function in (2). Stochasticity is imposed in both the mean function coefficients and residual field, $Z(r, \theta, t)$. In particular, we decompose ℓ th coefficient as

$$c_\ell(t) = \mu_\ell(t) + W_\ell(t), \quad (3)$$

for a deterministic trend function $\mu_\ell(t)$ and time series process $W_\ell(t)$.

Exploratory data analyses do not suggest strong simple linear relationships between the predictors and the estimated coefficients, so we opt to model each $\mu_\ell(\cdot)$ as a random forest (Breiman, 2001) with predictors discussed in Section 2. The residual processes $W_\ell(t)$ are modeled as independent (across ℓ), mean zero autoregressive Gaussian time series of order one, suggested by exploratory data analysis plots.

3.1.1 | Residual process

It is worthwhile to distinguish the problem we are tackling in the context of process models for polar data. Much research has been devoted to directional statistics problems where the response is an angle or direction (Mardia & Jupp, 2009); a common model for univariate data is the Von Mises distribution. Jona-Lasinio et al. (2012) extended such approaches to spatially-indexed data by introducing wrapped Gaussian processes. Our problem is distinct in that the response remains real-valued, but the indexing is no longer in Euclidean coordinates. The residual process at a given fixed band r and time point t must be periodic in θ . We adopt a Gaussian process specification; due to the reference system these processes have recently been termed polar Gaussian process (Padonou & Roustant, 2016). Traditionally such a model would involve defining a positive-definite function in polar coordinates. There are many routes to this end, and Gneiting (2013) and Rasmussen and Williams (2006) discuss some approaches. Due to the moderate size of our dataset, rather than specifying a functional form of a covariance kernel, we opt for a constructive approach that ensures a nonnegative definite covariance kernel.

To this end, we propose a basis representation that can be seen to be a synthesis of common low-rank ideas (Wikle, 2010) along with the polar Gaussian process idea. We model $Z(r, \theta, t)$ as

$$Z(r, \theta, t) = \int k(r, \theta, r_0, \theta_0) U(r_0, \theta_0, t) dr_0 d\theta_0, \quad (4)$$

for a space-time correlated process $U(r, \theta, t)$ and kernel function k . The kernel function is a bilinear interpolation function related to a polar coordinate grid, discussed in the next section; this choice of kernel function was motivated by its computational convenience. The model in (4) can be seen to be an adapted case of the nonstationary mixture models of Fuentes and Smith (2001), but different than the framework proposed by Higdon (1998) in that the kernel function varies over space, and U is a correlated process. The model for U is a low-rank circular Gaussian process,

$$U(r, \theta, t) = d_0(r, t) + \sum_{m=1}^M (d_{1,m}(r, t) \cos(m\theta) + d_{2,m}(r, t) \sin(m\theta)), \quad (5)$$

where $\{d_i(\cdot, \cdot)\}$ are stochastic. The specification in (5) can be interpreted as a harmonic decomposition for circular processes analogous to the spectral decomposition for stationary spatial processes (Stein, 1999). The intercept accounts for any bias incurred from the fitted random forest mean functions, and is seen to provide more realistic residual field simulations. Each process $d_0(r, t)$ or $\{d_{1,m}(r, t), d_{2,m}(r, t)\}_m$ is a radius-time process that is correlated in time according to an autoregressive process of order one, with a common autoregressive coefficient to avoid overfitting. For all time t , we model $\text{Cov}(d_0(r_1, t), d_0(r_2, t))$ as the same function of r_1 and r_2 nonparametrically (analogously for $d_{1,m}$ and $d_{2,m}$), described in the next section.

3.2 | Estimation

Our estimation strategy is intimately tied to the way in which the model is developed. For example, is it difficult to envision a cohesive likelihood-based framework for simultaneously estimating empirical orthogonal functions (our favored choice of basis functions, described next), random forests defining the principal components and polar Gaussian process covariance parameters. Moreover, our objective is directly tied to simulating realistic precipitation fields rather than robustly estimating and interpreting statistical parameters. In the next section we show the stepwise approach works well in out-of-sample testing.

3.2.1 | Stochastic mean function

We first concern ourselves with specifying and estimating the components of the stochastic mean function in (2) and (3). A common approach to modeling nonstationary processes in the atmospheric and climate sciences is to use an empirical orthogonal function representation (EOFs; von Storch & Zwiers, 2002; Wike, 2010), which is a low-rank representation in which the shape and form of the basis functions are specified by the data (we provide a technical description of EOFs below). Such an approach is equivalent to a principal components decomposition common in the statistical literature, and requires replicates of the process of interest on the same spatial grid. The native grid of our data in Euclidean space is common across all events and time points and directly lends itself to an EOF decomposition; however, the typical TC structure (e.g., the circular patterns of precipitation intensity about the center) varies widely across the domain depending on the location of the TC, which hampers an empirical method's ability to determine reasonable basis function structures. Moreover, the grid of our data in Euclidean space is common to all time points, modeling the process in a Lagrangian coordinate system results in irregularly-spaced and time-dependent coordinates in (r, θ) as the cyclone center moves across the domain.

For the purpose of developing a mean function model, we artificially grid the Lagrangian-indexed data to a regular grid in polar coordinate space spanning $\theta \in [-\pi, \pi)$ and $r \in [0, 1000]$ extending to 1000 km, resulting in a 100×100 gridded product at each time point for each storm; denote this grid by $\{(r_{k_r}, \theta_{k_\theta})\}_{k_r=1, k_\theta=1}^{100, 100}$. The choice of gridding to 100 points in each axial direction is admittedly arbitrary, but was seen to work as well as higher-resolution grids in exploratory simulations; moreover in radius this corresponds to a grid spacing of 10 km which is below the native resolution of the WRF data at 36 km. For each event e and time point t we take the available 59×59 data points $P_e(S_{lon}, S_{lat}, t) = P_e(r, \theta, t)$, fit a LatticeKrig model (Nychka et al., 2015) and krige to the regular grid in (r, θ) . To account for differences of units, we rescale the radius axis to $[0, 10]$. We use four levels of resolution and model settings that mimic an exponential covariance function, see Nychka et al. (2015) for details on the LatticeKrig model. Although the data are not Gaussian, fitting the LatticeKrig model by maximum likelihood assuming Gaussianity is one objective function approach, and as our goal is simply to produce a kriged field (rather than uncertainty estimates), LatticeKrig is an attractive option for computational reasons involving large spatial datasets (see Figure B1 for an example kriged surface). Because of the smoothing effect of kriging, the peaks of the original precipitation data will be slightly reduced in the polar-gridded dataset; in our experiments the median loss in maximal precipitation is 0.28 mm, or about 1% under the max. This slight biasing is partially addressed in the anamorphosis transformation, discussed later. An example of the polar coordinate smoothing is shown in Figure B1 in Appendix B.1. Note that the gridding of the data in polar coordinate space is necessary to estimate the EOF basis functions; indeed other bases could be imagined, but we prefer EOFs due to their widespread use in the atmospheric sciences, as well as ability to extract physically meaningful patterns in geophysical data (see example section below).

Collecting together the gridded data in a $K \times T$ matrix \mathbf{P} , where each column is a “flattened” polar image with $K = 100^2$ for all $T = 576$ available time points, the EOFs and principal components are obtained through a singular value decomposition (SVD) of \mathbf{P} . In particular, if $\mathbf{P} = \mathbf{U}\mathbf{D}\mathbf{V}^T$ is the SVD of \mathbf{P} then the columns of \mathbf{U} represent ordered EOFs, and the columns of $(\mathbf{D}\mathbf{V}^T)^T$ the principal components. These EOFs are spatial functions that define $\{\varphi_\ell(r, \theta)\}$. We note that such an approach may seem unintuitive at first, as we fit EOFs to the gridded precipitation data P , rather than Y , which is unobserved. We perform a simulation study to assess the ability of this approach to estimate EOF patterns, Appendix B.7 shows that our approach adequately reproduces true mean function patterns contained in the leading EOFs.

We use the first 13 EOFs, which explain more than 75% of variability. We also experimented with more EOFs, but found that the predictive ability of the principal component model, introduced next, suffered with higher order EOFs and likely resulted from overfitting. Instead, we opt for a parameterized residual process $Z(r, \theta, t)$.

The resulting principal components $\{c_\ell(t)\}_{\ell,t}$ are modeled as in (3). The trend functions $\mu_\ell(t)$ are random forests (Breiman, 2001; Hastie et al., 2009) with predictors of radius of maximal wind, central pressure deficit, storm center and direction (both vectors) and great circle distance of the storm center to coast. Some of these variables (e.g., storm center) are not directly available from the WRF output, but are derived to match the typical format of IBTrACS data; details can be found in Appendix A. We use 500 trees and fit the forests using the `randomForest` package in R (Liaw & Wiener, 2002) on the estimated principal components contained in $(\mathbf{DV}^T)^T$. Estimated residuals are formed from the fitted trend functions $\hat{\mu}_\ell(t)$, $\hat{W}_\ell(t) = c_\ell(t) - \hat{\mu}_\ell(t)$. The covariance parameters of the first-order autoregressive model for $\hat{W}_\ell(\cdot)$ for a given event ℓ are estimated by maximum likelihood, yielding $\{\hat{\phi}_\ell, \hat{\sigma}_\ell^2\}_{\ell=1}^n$ where ϕ_ℓ is the autoregressive coefficient and σ_ℓ^2 is the innovation variance.

3.2.2 | Residual process

We turn to the residual process $Z(r, \theta, t)$, which follows a specification of (4) and (5). Based on the prior section, we have estimated trends $\hat{\mu}_\ell(t)$ from the fitted random forest. These yield estimated residuals $\hat{U}(r, \theta, t) = P(r, \theta, t) - \sum_{\ell=1}^L \hat{\mu}_\ell(t) \varphi_\ell(r, \theta)$; note we use P here as Y is unobserved; in cross-validation experiments shown below, this approximation is seen to work well. Here it is useful to introduce additional notation: call $\hat{U}(r_{k_r}, \cdot, t)$ the estimated residuals for the k_r th radial band in the gridded polar coordinate system ($k_r = 1, \dots, 100$). Using $\hat{U}(r_{k_r}, \cdot, t)$ as data, coefficients $d_0(r_{k_r}, t), \{d_{1,m}(r_{k_r}, t), d_{2,m}(r_{k_r}, t)\}_m$ are estimated by least squares regression assuming a linear model of the form (5), yielding $\hat{d}_0(r_{k_r}, t), \{\hat{d}_{1,m}(r_{k_r}, t), \hat{d}_{2,m}(r_{k_r}, t)\}_m$. Taking $\hat{\mathbf{d}}_0(t) = (\hat{d}_0(r_1, t), \dots, \hat{d}_0(r_{100}, t))^T$ for each time point $t = 1, \dots, T$, we estimate the covariance matrix of $\text{Var } \mathbf{d}_0(t)$ for any time point empirically using $T^{-1} \sum_{t=1}^T \hat{\mathbf{d}}_0(t) \hat{\mathbf{d}}_0(t)^T$ (and analogously for $d_{1,m}$ and $d_{2,m}$); this covariance matrix controls dependence across radial bands and is common across all time points. The autoregressive coefficient governing each of $d_0(r_{k_r}, \cdot)$ and $\{d_{1,m}(r_{k_r}, \cdot), d_{2,m}(r_{k_r}, \cdot)\}_m$ are estimated independently by maximum likelihood assuming marginal Gaussianity, and are then averaged to provide a single autoregressive parameter estimate. See also Appendix B.6. One natural question is to what extent the residual component in (2) contributes variability to the simulated fields; Appendix B.4 considers the removal of the residual component with both a fixed, and stochastic mean function, and illustrates that the residual process is critical to include.

3.2.3 | Transformation function

By this point all statistical parameters are estimated, except those governing the transformation functions G and F in (1). The parameters of the gamma function G are estimated by maximum likelihood based on $\{P_e(r, \theta, t) | P_e(r, \theta, t) > 0\}$ assuming space-time independence. The assumption of independence may lead to, for example, overly narrow confidence intervals for estimated parameters; however, in our data example we only use the point estimate. Other marginal transformations could be envisioned; if focus is on the extreme tail of the distribution then options that include generalized extreme value distribution tail behavior may be adopted (Naveau et al., 2016); our initial focus is on the bulk of the precipitation distribution which is well modeled by the gamma transformation. Note the space-time varying nature of the statistical model implies that the distribution of precipitation at (\mathbf{s}_1, t) and (\mathbf{s}_2, t) are different, even with the common gamma transformation. Finally, F is estimated as follows: assume an ensemble of E realizations of space-time trajectories from the model for a given storm has been generated, $\{Y_e(\mathbf{s}, t)\}_{\ell=1}^E$ on the original Euclidean grid and set of times for a study storm. We estimate F as the empirical cumulative distribution function of the positive values $\{Y_e(\mathbf{s}, t) | Y_e(\mathbf{s}, t) > 0\}_{\ell=1}^E$, which, as $E \rightarrow \infty$, can be seen to approximate an equal weight mixture of local space-time densities; such an approach is an attempt to avoid overfitting of local probability density functions that can vary substantially over both space and time, and is seen to work well in the data study. In Appendix B.2 we compare against a standard marginal transformation, assuming F is the cdf of a standard normal (see, e.g., Kleiber et al., 2012); the results suggest a strong tendency to over-simulate precipitation.

3.2.4 | Simulation and interpolation

As modeling is done in regularly-spaced polar coordinates $\{(r_{k_r}, \theta_{k_\theta})\}$, initial simulations are quickly generated on the polar grid, yielding $\{U(r_{k_r}, \theta_{k_\theta}, t)\}$. For an arbitrary polar coordinate $(r, \theta) \notin \{(r_{k_r}, \theta_{k_\theta})\}$, the values of $Z(r_{k_r}, \theta_{k_\theta}, t)$ are bilinearly interpolated, that is, $k(r, \theta, r_0, \theta_0)$ in (4) represents a bilinear interpolation kernel. Generating $U(r, \theta, t)$ directly

is not straightforward given the empirical estimation approach for the component coefficients $d_0, d_{1,m}, d_{2,m}$. The EOFs are also bilinearly interpolated to the native grid for a given storm center, weighted by a realization of $\{c_\ell(t)\}$ and are added to $Z(r, \theta, t)$. The final transformation (1) is applied at each grid cell separately.

4 | APPLICATION TO HOUSTON DOMAIN

In this section we explore the adequacy of the statistical model for capturing the distribution of precipitation patterns in the WRF data. First we explore some of the fitted model components based on the full available dataset (an in-sample assessment), followed by detailed cross-validation study to assess the model's ability to capture out-of-sample behavior. In both cases we apply a tapering function to both the simulated and observed precipitation fields that is radially-symmetric and whose support depends on a particular time point's radius of maximal winds; the motivation for this is that our model is primarily designed to capture the local structure of tropical cyclones near the eye of the storm, whereas in WRF simulations other auxiliary precipitation events can be present due to other atmospheric conditions and forcings which our model is not designed to tackle. The taper function formula is given in Appendix B.3. Future research may combine both tropical cyclone generators with other precipitation generators for convective or frontal events.

To assess the realism of our model's space-time simulations, we examine a number of validation statistics at both the point-level, and domain-level. These include visual assessment of the simulated fields to confirm corresponding asymmetric structure of the TC, as well as statistics such as total precipitation both at a point level as well as domain level, assessment of pointwise precipitation distributions at a set of coastal locations, an assessment of probability of precipitation as well as the statistical calibration of the pointwise precipitation distributions.

4.1 | In-sample assessment

In this section we explore some details of the estimated model using the estimation procedure of the previous section applied to all seven tropical cyclones. The basis functions $\{\varphi_\ell(r, \theta)\}$ comprising the mean function in (2) are empirical orthogonal functions, estimated as described in the previous section. Figure 2 shows the first eight estimated EOFs for a storm center located in approximately the center of the study domain. The first EOF shows a classical approximately-circular pattern around the storm center, but with a skewed feature to the southeast of the center. The second EOF can be seen to be a correction to the first that can reduce cyclone asymmetry and increase precipitation to the north of the center. The combination of these two EOFs with the coastline approximately determining the separation between positive and negative regions of these fields is anticipated, as there is often a notable enhancement of precipitation as the rain bands of the storm arrive over land. In addition, the concentric "swirls" seen in other EOFs relate to the banded nature of precipitation in tropical cyclones which is enhanced as they approach land. The banding is enhanced as drier air from the continental interior gets wrapped into the center of the storm, suppressing precipitation in its path.

The coefficients of (2) are modeled including a random forest component as in (3); see also Appendix B.5. Although random forests consist of numerous easier-to-interpret trees, variable importance is one approach to assess the utility of features included in the predictor set over all trees simultaneously. Variable importance for a particular feature can be defined as the total (or average) reduction in variance explained at all nodes for all component trees over which that feature was split (Hastie et al., 2009). Table 1 includes variable importance factors for the first seven principal components' random forest models. The first three principal components' random forest models highly favor the storm's central pressure deficit as the most important predictor, which is unsurprising given the statistical relationship between such deficit and the overall strength of the storm (Chavas et al., 2017). As we move into higher frequency patterns, there are not apparent strong patterns highly favoring one particular feature over another. In the fourth principal component, the radius of maximum wind, V component of center (analogous to latitude) and distance to coast experience jumps in their variable importance factors, which indicates the importance of those variables for the EOF pattern seen in Figure 2.

4.2 | Cross-validation

In this section we consider the predictive ability of our approach. In particular, the model depends on only a handful of predictors: radius of maximal winds, central pressure deficit, storm center and direction of movement. Recall that these

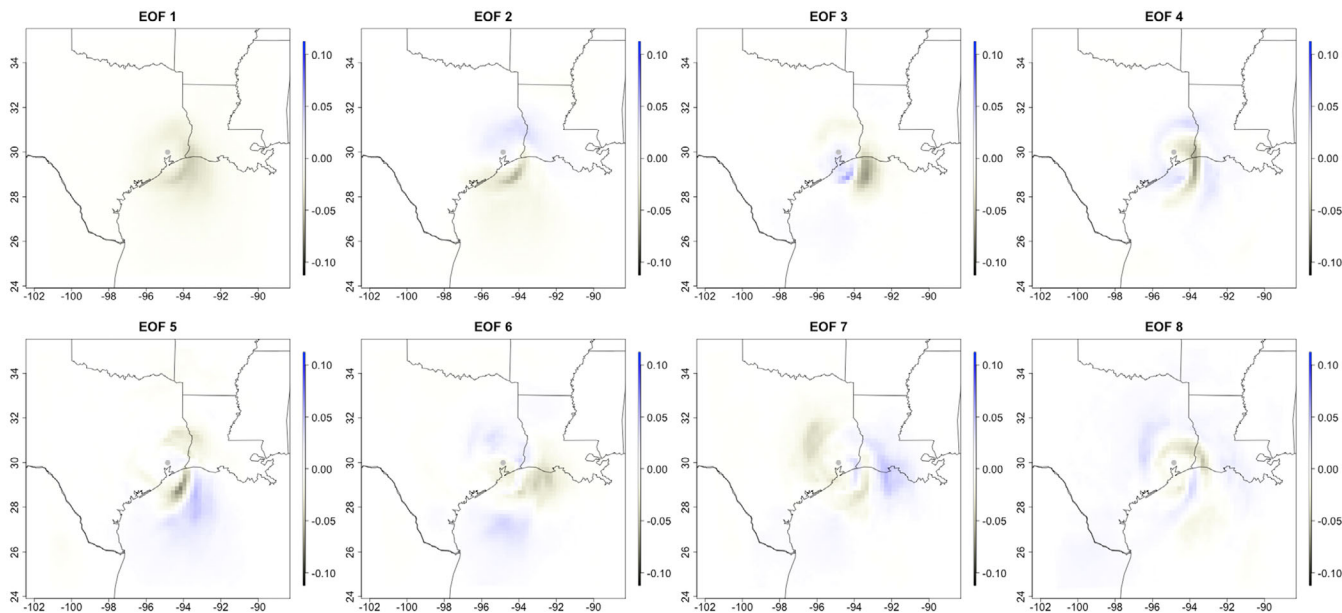


FIGURE 2 First eight empirical orthogonal functions the storm center indicated by a grey dot.

TABLE 1 Variable importance measured as total reduction in variance for the first seven principal components' random forest models; values are scaled by 1×10^{-5} .

	c_1	c_2	c_3	c_4	c_5	c_6	c_7
Radius of maximal winds	1.07	0.65	0.48	0.56	0.34	0.19	0.18
Pressure deficit at center	9.97	4.61	1.06	0.70	0.89	0.64	0.61
U component of direction	1.57	1.06	0.78	0.47	0.41	0.29	0.24
V component of direction	3.82	0.84	0.98	0.69	0.45	0.51	0.51
U component of center	2.88	1.91	0.74	0.44	0.32	0.38	0.35
V component of center	1.86	1.10	0.61	0.96	0.40	0.32	0.29
Distance to coast	1.44	1.04	0.45	0.64	0.26	0.37	0.33

few predictors are those that are available in the widely-used IBTrACS data base or in other simulated products. With only these features, the model generates plausible precipitation patterns.

To test the ability of the model to generalize beyond the few storms we have available, we consider a cross-validation experiment in which the estimation routine of Section 3.2 is implemented on six of seven storms, the last of which is held out for comparison purposes. The estimation process is repeated for each possible hold-out storm, and in each case we generate an ensemble of 100 space-time realizations from the statistical model which we compare against the held-out WRF data.

4.2.1 | Space-time structure

To begin this section, Figure 3 shows realizations from the cross-validation experiment for a handful of instances. The three rows correspond to different TCs, the left column contains the held-out WRF simulation for a particular hour and the right three columns contain randomly chosen simulations from the simulated ensemble. We can see, for example in the second two instances, apparent asymmetry in the WRF fields that are replicated by the statistical model. In the first row, the first and second realization of the storm show heavier precipitation patterns, but that are more circular. The third realization shows stronger asymmetry with less precipitation.

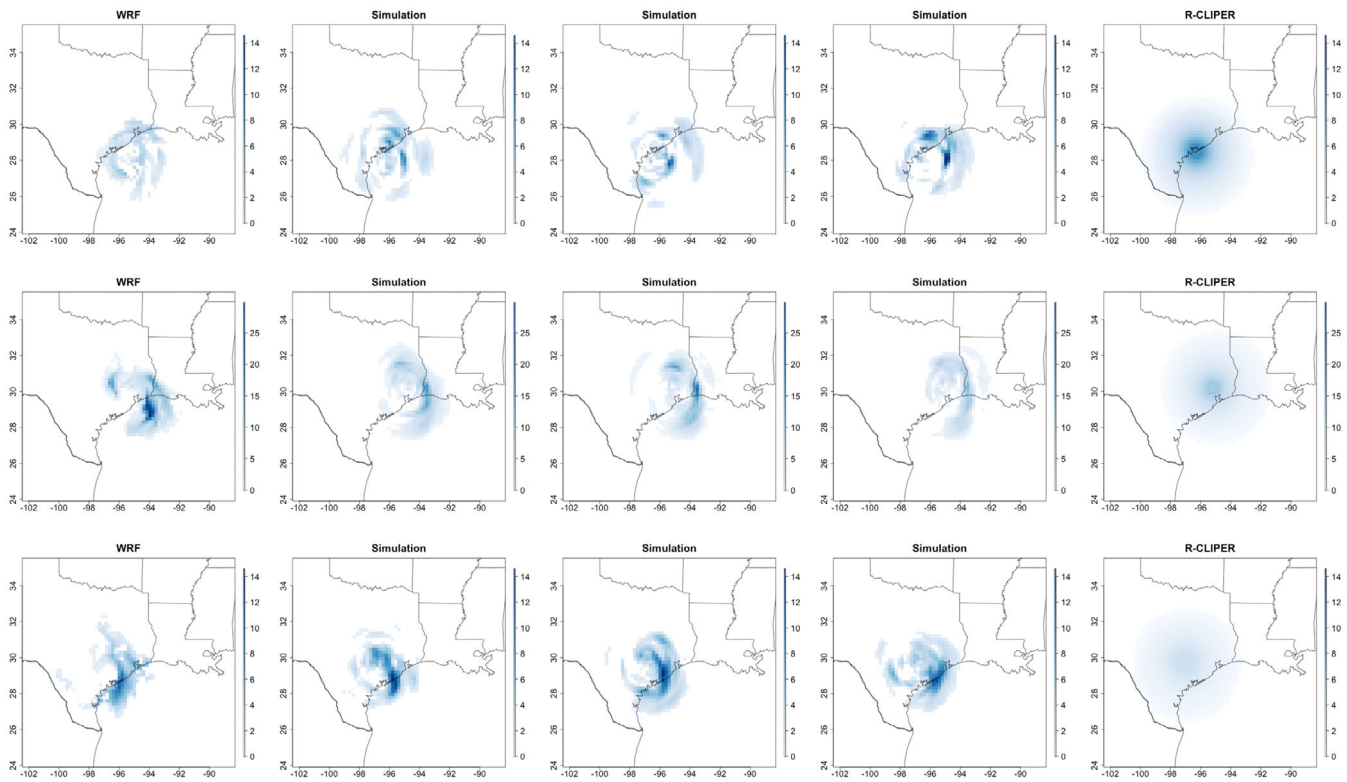


FIGURE 3 Cross-validation realizations for three snapshots. The three rows represent hurricanes Claudette on July 15, 2003 at 4:00 p.m., Ike on September 13, 2008 at 12:00 p.m. and Bill on June 17, 2015 at 4:00 a.m., respectively. The first column is WRF output for the corresponding time, the middle three columns show cross-validation realizations from the statistical model while the last column shows the R-CLIPER prediction. Units are mm and all plots are on the same color scale.

The asymmetrical patterns returned by the cross-validation realizations are consistent with observed patterns of tropical cyclone precipitation, particularly near landfall, where precipitation is enhanced to the right of the storm motion (here, to the east side of the storm). In addition, localized enhancements are seen in the portion of these rain bands immediately adjacent to the coast, particularly in the first and second row, which is also consistent with physical expectations (Rogers et al., 2009). In the third row, this is a particularly difficult time to replicate, as the storm has moved well inland and is losing its structure as a tropical cyclone. Factors outside of the immediate environment of the tropical cyclone (including pre-existing atmospheric patterns over the continent) are more strongly influencing the structure of the decaying storm and increasing the variability in precipitation patterns away from the coast, and the statistical model struggles somewhat to capture this.

Although capturing the instantaneous spatial patterns of precipitation is important, it is also important to replicate temporally-aggregated behavior such as the total rainfall over the duration of the cyclone for flood risk modeling exercises. Figure 4 shows pointwise-integrated precipitation over the study domain for one example event. The left panel represents integrated WRF and the second panel is one particular realization's integration; the next two panels show pointwise 5% and 95% quantiles of integrated precipitation over all realizations. First, the simulated totals appropriately capture the spatial distribution of total precipitation that cut in a diagonal pattern across the study domain. The statistical model exhibits some smoothing as compared to the actual WRF data, but capturing such high frequency behavior with our limited features is likely an unrealistic goal.

A third assessment of the model's ability to replicate space-time structure is in the temporal behavior of domain-wide integrated (i.e., total hourly) precipitation. Figure 5 shows functional boxplots (Sun & Genton, 2011) for integrated precipitation at each hour for each of the seven cross-validation storms. The red line indicates the held-out WRF statistic, while the grey band shows interquartile range of the WRF ensemble. There is some additional high-frequency variation in the statistical model's integrated precipitation, where the WRF time series exhibits smoother variation of total precipitation than the statistical model. This could potentially be resolved by putting moderately-strong prior distributions on the autoregressive aspects of the model, encouraging greater temporal regularity, but it is not clear that the additional




FIGURE 4 Integrated pointwise precipitation for the event beginning August 14, 2007 for WRF and a realization from the statistical model (left two panels), as well as pointwise 5%, 50% and 95% quantiles over an ensemble of 100 realizations of the statistical model. Units are mm.

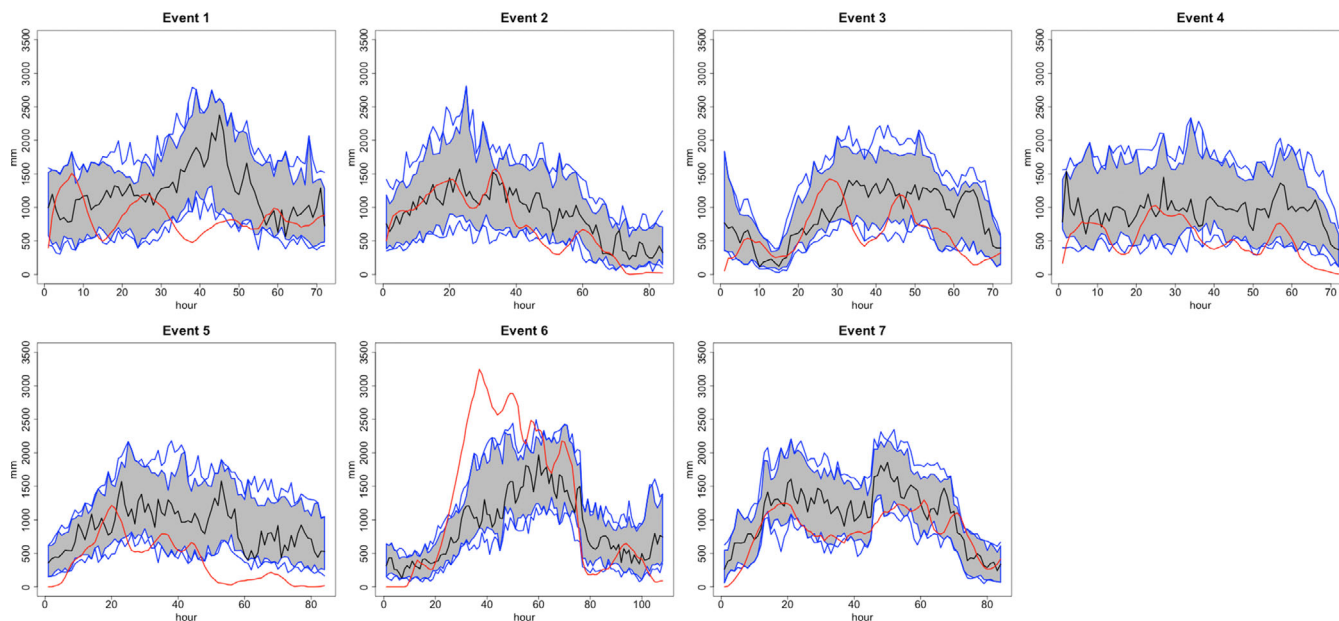


FIGURE 5 Functional boxplots of integrated precipitation over the study domain for each storm event based on an ensemble of 100 realizations of the statistical model. The black line is the functional boxplot median while the red line is the integrated precipitation of WRF. Units are mm.

computational cost of such an implementation would be prohibitive. Generally the model realizations capture the verification time series within the 95% functional boxplot confidence intervals, and also tend to follow the increasing and decreasing trends of precipitation. The sixth storm is an unusual or outlying case where substantially more precipitation fell over the domain (around hours 30–60) than in any of the other study storms; the statistical model captures the ramp-up period of increasing rainfall, but whose median does not adequately represent such intensity. This failure to capture the extreme integrated rainfall is due to the fact that the training storms (the other six) had more similar statistical characteristics with much less instantaneous precipitation falling on average.

4.2.2 | Local verification and case studies

We now turn to local verification and case study examples of the statistical model. We consider five case study locations shown in Figure 6 on and near the Texas coastline which represents the primary area of interest for the original experimental design. The study locations are Port Arthur, Houston, Corpus Christi, TX and two pixels covering Galveston, TX: one pixel that is primarily over the land, the second is primarily over the ocean/coast adjacent to Galveston. The motivation for choosing the two Galveston locations is that precipitation patterns can change substantially with the introduction of land.

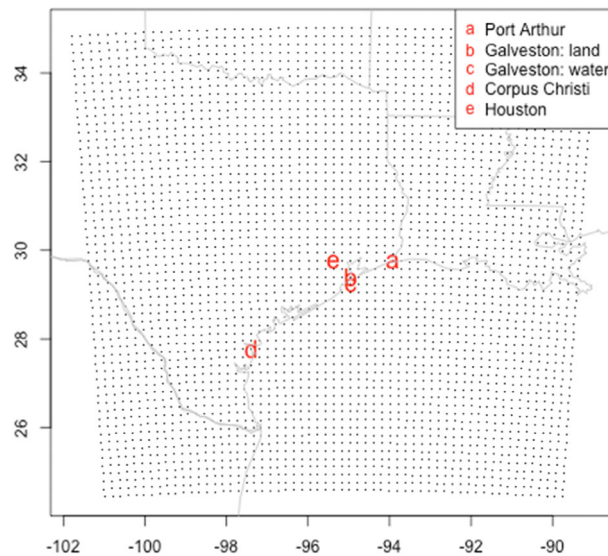


FIGURE 6 Locations of five case study locations on and near the Texas coastline; dots represent the computational WRF grid.

We consider distributional aspects of hourly precipitation at each of the five case study locations in Figure 6 over all seven cross-validation events. Figure 7 shows Q-Q plots for each storm and location and each of the 100 realizations of the model, as well as values from R-CLIPER. In particular, these Q-Q plots represent hourly precipitation data at the particular pixel of interest. These plots provide much information about the ability of the statistical model to appropriately localize precipitation. For example, in approximately six of the 35 Q-Q plots, the statistical model undersimulates precipitation at the extremes. Specifically, the sixth storm (column six) has high precipitation at Port Arthur and both Galveston pixels that are not adequately captured by the statistical model; as with Figure 5 this is likely due to the fact that this storm was statistically unusual in having substantially higher precipitation compared to the other events. Generally speaking the Q-Q plots typically contain the identity line in approximate 95% confidence bands, even at more extreme precipitation amounts, although there are instances where the statistical model oversimulates precipitation (e.g., the first event at Houston, TX). Moreover, it is clear from these plots that R-CLIPER typically severely oversimulates, or undersimulates, at the point-level, whereas in our simulations the identity line is more often contained within the central bulk of our simulations, see, for example, the Port Arthur Q-Q plots. Capturing such local distributions is a high demand on the model given the highly heterogeneous behavior of local precipitation, including its discrete-continuous nature. We next consider validating the occurrence of precipitation.

The TC generator provides a separate probability distribution at each spatial location and hour; it is appropriate and useful to quantify the quality of these distributions with respect to the held-out WRF data using verification tools from probabilistic forecasting (Gneiting, Balabdaoui, & Raftery, 2007). In particular, we present two statistics, the first validating the model's ability to realistically provide probability of precipitation, and the second the full local simulation distribution. A Brier score is a proper scoring rule used for predicting dichotomous events: for a set of forecast probabilities $f_i \in [0, 1]$ and observations $o_i \in \{0, 1\}$, the Brier score $i = 1, \dots, n$ is defined as

$$\frac{1}{n} \sum_{i=1}^n (o_i - f_i)^2.$$

Brier scores vary between zero and one, and are inversely oriented in that lower scores indicate better probability forecasts. We approximate the model's probability of precipitation by the empirical rate that the 100 ensemble members are nonzero at any given spatial location and hour. Figure 8 shows a histogram of Brier scores for presence of positive precipitation over all hours from all events with almost all scores falling below 0.2. The average Brier score over all data is 0.05, indicating high quality model probabilities of zero or positive precipitation that generally correctly identify locations of precipitation occurrence.

As a final verification statistic, we consider the full local predictive distributions of precipitation. In theme with validating the probability of precipitation, at each spatial location and time point there is a predictive distribution of precipitation

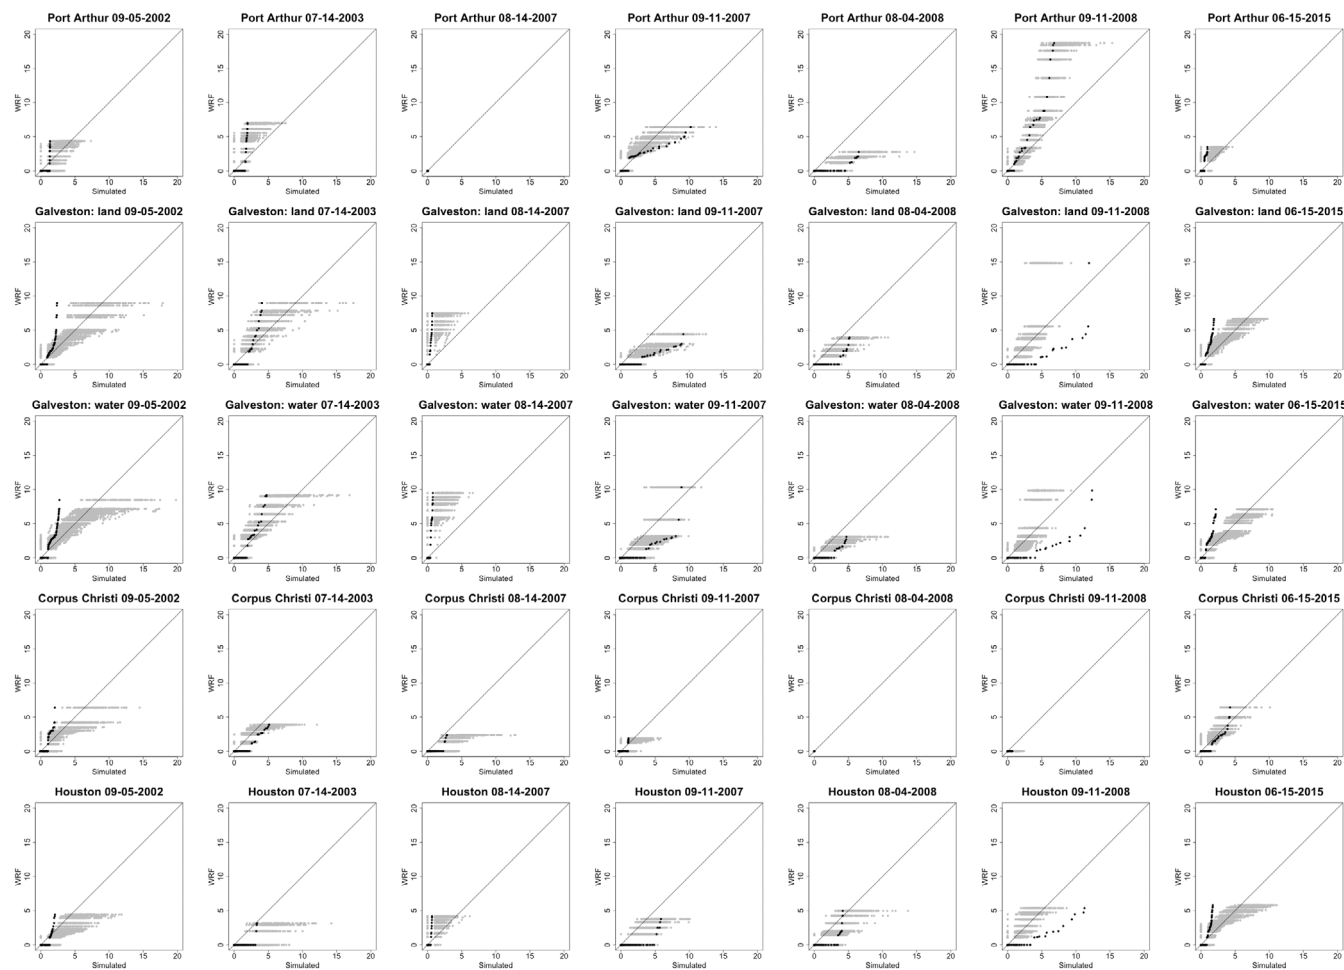


FIGURE 7 Q-Q plots of hourly precipitation at the five case study sites for all seven events. Each row is a case study location and each column is a separate event. Grey dots represent values different simulations from the proposed model, while black dots represent R-CLIPER quantiles. Units are mm.

sampled by the 100 realizations from the statistical model. Such predictive distributions can be validated using verification rank histograms (Gneiting, Balabdaoui, & Raftery, 2007). Rank histograms represent counts of the rank of the WRF precipitation (at a given pixel and time) within the 100 simulated values from the statistical model. If the model is well-calibrated, then these ranks should approximately follow a uniform distribution.

Figure 9 shows such rank histograms for each separate verification storm with a horizontal line indicating perfect calibration. Events with ties between the WRF data and simulated data are randomly disaggregated. Generally the histograms are flat and indicate high-quality calibration without much substantial under- or over-dispersion. Evidence shows in almost all cases of slight undersimulation at the most extreme events. That is, the presence of the peak at the highest rank indicates that there are an exaggerated number of cases where the WRF precipitation is higher than all 100 ensemble members from the model. These are not necessarily cases of extreme precipitation, although they are contributing factors. Future research will focus on incorporating extreme value distributions into the space-time dynamical framework proposed in this article.

5 | DISCUSSION

In this article we introduce an ambitious approach to simulating quantitative precipitation fields for tropical cyclones given a handful of features that are typically available in historical or simulated tropical cyclone track databases. The method relies on modeling in Lagrangian coordinates, using the estimated storm center as the coordinate system origin.




FIGURE 8 Histogram of Brier scores validating probability of precipitation over all hours and all TC events.

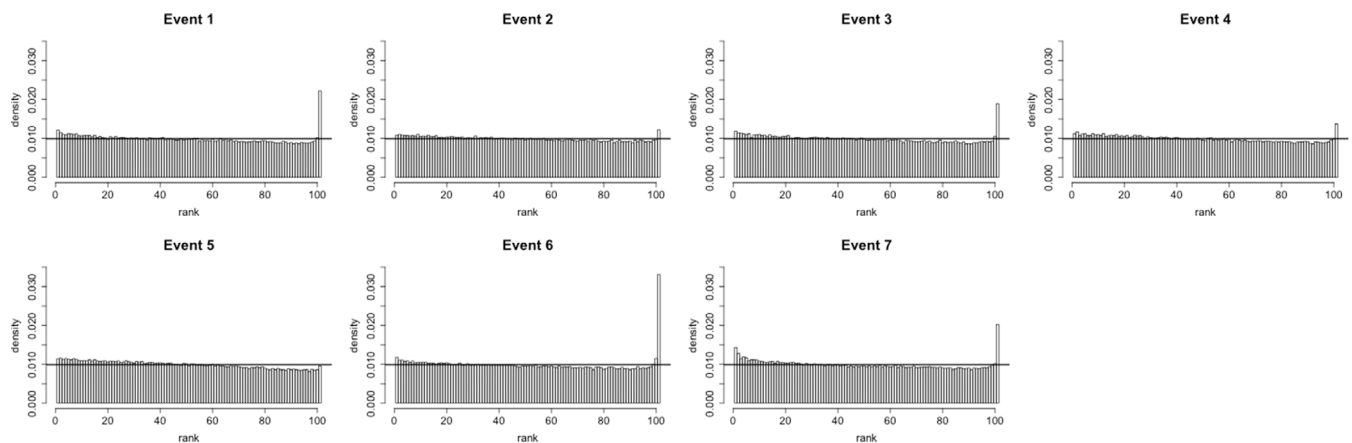


FIGURE 9 Verification rank histograms for each tropical cyclone event in the cross-validation experiment.

Basis expansion models are used both within the mean and residual processes, exploiting links to standard kernel convolution approaches for nonstationary processes in the spatial statistical literature. The limited information in the tracks is aided by the use of training data based on observed TC tracks and high-resolution WRF model output, and the results here represent a positive first step in developing a statistical approach for simulating precipitation fields for assessing flooding due to TCs.

Estimation of the model is closely tied to model development, allowing for assumption checking and model building at each stage. A natural criticism of our approach is such a stepwise estimation procedure which is not expected to be statistically efficient. Taking a more principled likelihood-based approach would be desirable for a fully statistical solution, but in our experiments the current approach aligns well with the goal of this article: to produce realistic space-time correlated fields of tropical cyclone precipitation. It is not clear that the added effort to more fully incorporate parametric uncertainty would have any substantive benefit, and would certainly come at the cost of significantly greater effort.

While our interest focuses on TC precipitation characteristics, coastal flooding is driven by both precipitation and storm surge. A future research direction is to fold in a wind field model, such as Reich and Fuentes (2007), to estimate the joint hazard posed by these two aspects of TCs. Another relevant direction would be to experiment with the proposed approach in other geographical domains, as tropical cyclone climatological characteristics differ over distinct bodies of water.

ORCID

William Kleiber  <https://orcid.org/0000-0003-0411-9108>

REFERENCES

- Alegria, A., & Porcu, E. (2017). The dimple problem related to space-time modeling under the Lagrangian framework. *Journal of Multivariate Analysis*, *162*, 110–121.
- Bandyopadhyay, S., & Lahiri, S. N. (2010). Asymptotic properties of discrete Fourier transforms for spatial data. *Sankhya*, *71*, 221–259.
- Breiman, L. (2001). Random forests. *Machine Learning*, *45*, 5–32.
- Chavas, D., Reed, R. K. A., & Knaff, J. A. (2017). Physical understanding of the tropical cyclone wind-pressure relationship. *Nature Communications*, *8*, 1–11.
- Cressie, N., & Johannesson, G. (2008). Fixed rank kriging for very large spatial data sets. *Journal of the Royal Statistical Society, Series B*, *70*, 209–226.
- Davis, C., Wang, W., Chen, S. S., Chen, Y., Corbosiero, K., DeMaria, M., Dudhia, J., Holland, G., Klemp, J., Michalakes, J., & Reeves, H. (2008). Prediction of land falling hurricanes with the advanced hurricane WRF mode. *Monthly Weather Review*, *136*, 1990–2005.
- Davis, C. A., Wang, W., Cavallo, S., Done, J. M., Dudhia, J., Fredrick, S. M., Michalakes, J., Caldwell, G., Engel, T., & Torn, R. (2010). High-resolution hurricane forecasts. *Computing in Science & Engineering*, *13*, 22–30.
- Dee, D. P., Uppala, S. M., Simmons, A. J., Berrisford, P., Poli, P., Kobayashi, S., Andrae, U., Balmaseda, M. A., Balsamo, G., Bauer, P., Bechtold, P., Beljaars, A. C. M., van de Berg, L., Bidlot, J., Bormann, N., Delsol, C., Dragani, R., Fuentes, M., Geer, A. J., ... Vitart, F. (2011). The ERA-Interim reanalysis: Configuration and performance of the data assimilation system. *Quarterly Journal of the Royal Meteorological Society*, *137*, 553–597.
- Elsberry, R. L. (2002). Predicting hurricane landfall precipitation: Optimistic and pessimistic views from the symposium on precipitation extremes. *Bulletin of the American Meteorological Society*, *83*, 1333–1339.
- Emanuel, K. (2006). Climate and tropical cyclone activity: A new model downscaling approach. *Journal of Climate*, *19*, 4797–4802.
- Fuentes, M., & Smith, R. L. (2001). *A new class of nonstationary spatial models* [Technical report]. North Carolina State University, Department of Statistics.
- Gneiting, T. (2013). Strictly and non-strictly positive definite functions on spheres. *Bernoulli*, *19*, 1327–1349.
- Gneiting, T., Balabdaoui, F., & Raftery, A. E. (2007). Probabilistic forecasts, calibration and sharpness. *Journal of the Royal Statistical Society, Series B*, *69*, 243–268.
- Gneiting, T., Genton, M. G., & Guttorp, P. (2007). *Geostatistical space-time models, stationarity, separability and full symmetry*. In B. Finkenstadt, L. Held, & V. Isham (Eds.), *Statistical methods for spatio-temporal systems* (pp. 151–175). Chapman & Hall/CRC Press.
- Hacker, J. P., Exby, J., Gill, D., Jimenez, I., Maltzahn, C., See, T., Mullendore, G., & Fossell, K. (2017). A containerized mesoscale model and analysis toolkit to accelerate classroom learning, collaborative research, and uncertainty quantification. *Bulletin of the American Meteorological Society*, *98*, 1129–1138.
- Hall, T. M., & Jewson, S. (2007). Statistical modeling of North Atlantic tropical cyclone tracks. *Tellus*, *59A*, 486–498.
- Hastie, T., Tibshirani, R., & Friedman, J. (2009). *The elements of statistical learning: Data mining, inference, and prediction*. Springer Science & Business Media.
- Higdon, D. (1998). A process-convolution approach to modelling temperatures in the North Atlantic Ocean. *Environmental and Ecological Statistics*, *5*, 173–190.
- Hong, S., Dudhia, J., & Chen, S. (2004). A revised approach to ice microphysical processes for the bulk parameterization of clouds and precipitation. *Monthly Weather Review*, *132*, 103–120.
- Hong, S., Noh, Y., & Dudhia, J. (2006). A new vertical diffusion package with an explicit treatment of entrainment processes. *Monthly Weather Review*, *134*, 2318–2341.
- Iacono, M. J., Delamere, J. S., Mlawer, E. J., Shepard, M. W., Clough, S. A., & Collins, W. D. (2008). Radiative forcing by long-lived greenhouse gases: Calculations with the AER radiative transfer models. *Journal of Geophysical Research*, *113*, D13103.
- Jona-Lasinio, G., Gelfand, A., & Jona-Lasinio, M. (2012). Spatial analysis of wave direction data using wrapped Gaussian processes. *Annals of Applied Statistics*, *6*, 1478–1498.
- Kain, J. S. (2004). The Kain-Fritsch convective parametrization: An update. *Journal of Applied Meteorology*, *43*, 170–181.
- Kleiber, W., Katz, R. W., & Rajagopalan, B. (2012). Daily spatiotemporal precipitation simulation using latent and transformed Gaussian processes. *Water Resources Research*, *48*.
- Knapp, K., Kruk, M. C., Levinson, D. H., Diamond, H. J., & Neumann, C. J. (2010). The international best track archive for climate stewardship (IBTrACS): Unifying tropical cyclone data. *Bulletin of the American Meteorological Society*, *91*, 363–376.

- Knutson, T., Camargo, S., Chan, J. C. L., Emanuel, K., Ho, C. H., Kossin, J., Mohapatra, M., Satoh, M., Sugi, M., Walsh, K., & Wu, L. (2020). Tropical cyclones and climate change assessment: Part II. Projected response to anthropogenic warming. *Bulletin of the American Meteorological Society*, E303–E322.
- Lai, Y., Li, J., Gu, X., Liu, C., & Chen, Y. (2021). Global compound floods from precipitation and storm surge: Hazards and the roles of cyclones. *Journal of Climate*, 34, 8319–8339.
- Langousis, A., & Veneziano, D. (2009). Theoretical models of rainfall in tropical cyclones for the assessment of long-term risk. *Journal of Geophysical Research*, 114, D02106.
- Liaw, A., & Wiener, M. (2002). Classification and regression by randomForest. *R News*, 2, 18–22.
- Lin, Y., & Mitchell, K. E. (2005). *The NCEP stage II/IV hourly precipitation analyses: Development and applications*. Proceedings of the 19th Conference Hydrology, American Meteorological Society (Vol. 1.2). San Diego, CA: American Meteorological Society.
- Lindgren, F., Rue, H., & Lindström, J. (2011). An explicit link between Gaussian fields and Gaussian Markov random fields: The stochastic partial differential equation approach. *Journal of the Royal Statistical Society, Series B*, 73, 423–498.
- Lonfat, M., Marks, F. D., & Chen, S. S. (2004). Precipitation distribution in tropical cyclones using the tropical rainfall measuring mission (TRMM) microwave imager. A global perspective. *Monthly Weather Review*, 132, 1645–1660.
- Lonfat, M., Rogers, M. R., Marchok, T., & Marks, F. (2007). A parametric model for predicting hurricane rainfall. *Monthly Weather Review*, 135, 3086–3097.
- Lu, P., Lin, N., Emanuel, K., Chavas, D., & Smith, J. (2018). Assessing hurricane rainfall mechanisms using a physics-based model: Hurricanes Isabel (2003) and Irene (2011). *Journal of the Atmospheric Sciences*, 75, 2337–2358.
- Mardia, K. V., & Jupp, P. E. (2009). *Directional statistics* (Vol. 494). John Wiley & Sons.
- Modlin, D., Fuentes, M., & Reich, B. (2012). Circular conditional autoregressive modeling of vector fields. *Environmetrics*, 23, 46–53.
- Nakamura, J., Lall, U., Kushnir, Y., & Rajagopalan, B. (2015). HITS: Hurricane intensity and track simulator with North Atlantic Ocean implications for risk assessment. *Journal of Applied Meteorology and Climatology*, 54, 1620–1636.
- Naveau, P., Huser, R., Ribereau, P., & Hannart, A. (2016). Modeling jointly low, moderate, and heavy rainfall intensities without a threshold selection. *Water Resources Research*, 52, 2735–2769.
- Nelson, B. R., Prat, O. P., Seo, D. J., & Habib, E. (2016). Assessment and implications of NCEP stage IV quantitative precipitation estimates for product intercomparisons. *Weather and Forecasting*, 31, 371–394.
- Nychka, D., Bandyopadhyay, S., Hammerling, D., Lindgren, F., & Sain, S. (2015). A multi-resolution Gaussian process model for the analysis of large spatial datasets. *Journal of Computational and Graphical Statistics*, 24, 579–599.
- Omranian, E., Sharif, H. O., & Tavakoly, A. A. (2018). How well can global precipitation measurement (GPM) capture hurricanes? Case study: Hurricane Harvey. *Remote Sensing*, 10, 1150.
- Padonou, E., & Roustant, O. (2016). Polar Gaussian processes and experimental designs in circular domains. *SIAM/ASA Journal on Uncertainty Quantification*, 4, 1014–1033.
- Pollock, M. D., O'Donnell, G., Quinn, P., Dutton, M., Black, A., Wilkinson, M. E., Colli, M., Stagnaro, M., Lanza, L. G., Lewis, E., Kilsby, C. G., & O'Connell, P. E. (2018). Quantifying and mitigating wind-induced undercatch in rainfall measurements. *Water Resources Research*, 54, 3863–3875.
- Porcu, E., Furrer, R., & Nychka, D. (2021). 30 years of space–Time covariance functions. *Wiley Interdisciplinary Reviews: Computational Statistics*, 13, e1512.
- Rasmussen, C. E., & Williams, C. K. I. (2006). *Gaussian processes for machine learning*. MIT Press.
- Reich, B. J., & Fuentes, M. (2007). A multivariate semiparametric Bayesian spatial modeling framework for hurricane surface wind fields. *Annals of Applied Statistics*, 1, 249–264.
- Rekabdarkolaee, H. M., Krut, C., Fuentes, M., & Reich, B. J. (2019). A Bayesian multivariate functional model with spatially varying coefficient approach for modeling hurricane track data. *Spatial Statistics*, 29, 351–365.
- Rodriguez-Iturbe, I., Cox, D. R., & Isham, V. (1987). Some models for rainfall based on stochastic point processes. *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences*, 410, 269–288.
- Rogers, R., Marks, F., & Marchok, T. (2009). *Tropical cyclone rainfall*. In M. G. Anderson & J. J. McDonnell (Eds.), *Encyclopedia of hydrological sciences*. Wiley.
- Skamarock, W. C., Klemp, J. B., Dudhia, J., Gill, D. O., Liu, Z., Berner, J., Wang, W., Powers, J. G., Duda, M. G., Barker, D. M., & Huang, X. Y. (2019). *A description of the advanced research WRF version 4* [Technical report NCAR/TN-556+STR]. National Center for Atmospheric Research.
- Stein, M. L. (1999). *Interpolation of spatial data: Some theory for kriging*. Springer-Verlag.
- Stern, R. D., & Coe, R. (1984). A model fitting analysis of daily rainfall data. *Journal of the Royal Statistical Society, Series A (General)*, 147, 1–34.
- Sun, Y., & Genton, M. G. (2011). Functional boxplots. *Journal of Computational and Graphical Statistics*, 20, 316–334.
- Ueno, M. (2007). Observational analysis and numerical evaluation of the effects of vertical wind shear on the rainfall asymmetry in the typhoon inner core. *Journal of the Meteorological Society of Japan*, 85, 115–136.
- Villarini, G., Mandapaka, P. V., Krajewski, W. F., & Moore, R. J. (2008). Rainfall and sampling uncertainties: A rain gauge perspective. *Journal of Geophysical Research*, 113, D11102.
- Villarini, G., & Smith, J. A. (2010). Flood peak distributions for the eastern United States. *Water Resources Research*, 46, W06504.
- von Storch, H., & Zwiers, F. W. (2002). *Statistical analysis in climate research*. Cambridge University Press.

- Wahl, T., Jain, S., Bender, J., Meyers, S. D., & Luther, M. E. (2015). Increasing risk of compound flooding from storm surge and rainfall for major US cities. *Nature Climate Change*, 5, 1093–1097.
- Wikle, C. K. (2010). *Low rank representations for spatial processes*. In A. E. Gelfand, P. Diggle, P. Guttorp, & M. Fuentes (Eds.), *Handbook of spatial statistics* (pp. 107–118). Chapman & Hall/CRC Press.
- Zhu, L., Quiring, S. M., & Emanuel, K. A. (2013). Estimating tropical cyclone precipitation risk in Texas. *Geophysical Research Letters*, 40, 6225–6230.

How to cite this article: Kleiber, W., Sain, S., Madaus, L., & Harr, P. (2023). Stochastic tropical cyclone precipitation field generation. *Environmetrics*, 34(1), e2766. <https://doi.org/10.1002/env.2766>

APPENDIX A. WRF DETAILS

Here we describe the configuration of the WRF model and the extraction of predictors from the WRF model output.

A.1 WRF model configuration

Version 4.0.1 of the Advanced Research WRF (ARW; Skamarock et al., 2019) model is used to make these simulations. The WRF model has a number of settings that can be configured to change the exact formulation of the model; many of these settings are configured to simulate particular types of weather events. Several WRF model simulations with different configuration options were compared to hourly Automated Surface Observing System (ASOS) observations at major airports within the WRF domain, and the configuration that best matched the available in-situ observations was chosen for these simulations. The configuration options are outlined in Table A1.

A.2 Extraction of WRF predictors

The WRF output contains U and V direction wind speeds (in m/s) at multiple geopotential heights; we use speeds at 850 mb to derive both the storm center and radius of maximal winds. First, winds are spatially smoothed using a double exponential kernel which reduces aberrant artifacts in the derivation of predictors. We calculate the pointwise curl of U/V wind speeds using second-order discretized finite difference derivatives at each time step. At each hour, the storm center is identified as the grid cell with maximal curl. We restrict attention to storm centers that are within two degrees L_∞ distance of the associated IBTrACS estimated storm centers to ensure that instantaneous aberrant storm center values are not used. Finally, the storm centers, which are registered to the WRF output grid, are then independently smoothed in both the U and V directions with a cubic smoothing spline.

Radius of maximal winds is calculated as the great circle distance to the point of maximal wind speed within 400 km of the estimated storm center at each hour. These radii are then smoothed over time with a cubic smoothing spline to reduce artifacts of the WRF gridding. Storm direction is calculated at hour t as the difference of the estimated storm center at time t minus that at time $t - 1$; such an estimate was found to be quantitatively similar to higher order approximations of direction. Central pressure deficit is quantified as 1013 minus pressure at the estimated storm center.

TABLE A1 WRF model configuration, see text for details.

Configuration	Selection	Citation
Grid spacing	20 km	
Initial and boundary conditions	ERA-Interim	Dee et al. (2011)
Microphysics	WSM5	Hong et al. (2004)
Cumulus	Kain–Fritsch	Kain (2004)
Planetary boundary layer	YSU	Hong et al. (2006)
Radiation	RRTMG	Iacono et al. (2008)

APPENDIX B. MODEL DETAILS

This appendix contains details of the model fit, including comparisons against simpler alternatives at various stages of the model.

B.1 Polar coordinate smoothing

Figure B1 shows two example snapshots of raw precipitation data from WRF (without the spatial tapering), the same data in polar coordinates and finally the kriged data in polar coordinate space. The critical point is that the two datasets in the first column show rather distinct spatial structure of precipitation, but when viewed in polar coordinates (middle column), the structure becomes more similar. However, the two sets of data are at irregularly-spaced locations in polar coordinates, while the kriged data (third column) are mapped to a regular grid which is then comparable across all storms and hours, and still maintains the important spatial structure seen in geographical coordinates.

B.2 Specification of transformation function

We next consider the choice of F in (1), instead using a standard normal cdf. Note that a standard normal cdf is the approach used in Kleiber et al. (2012); however, in this application the process is highly nonstationary, and it is not a priori clear that the more complicated version used in (1) is necessary. Figure B2 shows verification rank histograms for our case study locations when using a normal cdf transformation—there is clear and strong evidence of bias and over-simulation.

B.3 Specification of taper function

We apply a taper function to each WRF and simulated precipitation field at each time point. The function is

$$g(r; RMAX) = \begin{cases} 1, & r \leq RMAX, \\ \exp\left(-\left(\frac{r-RMAX}{60}\right)^2\right), & r > RMAX, \end{cases}$$

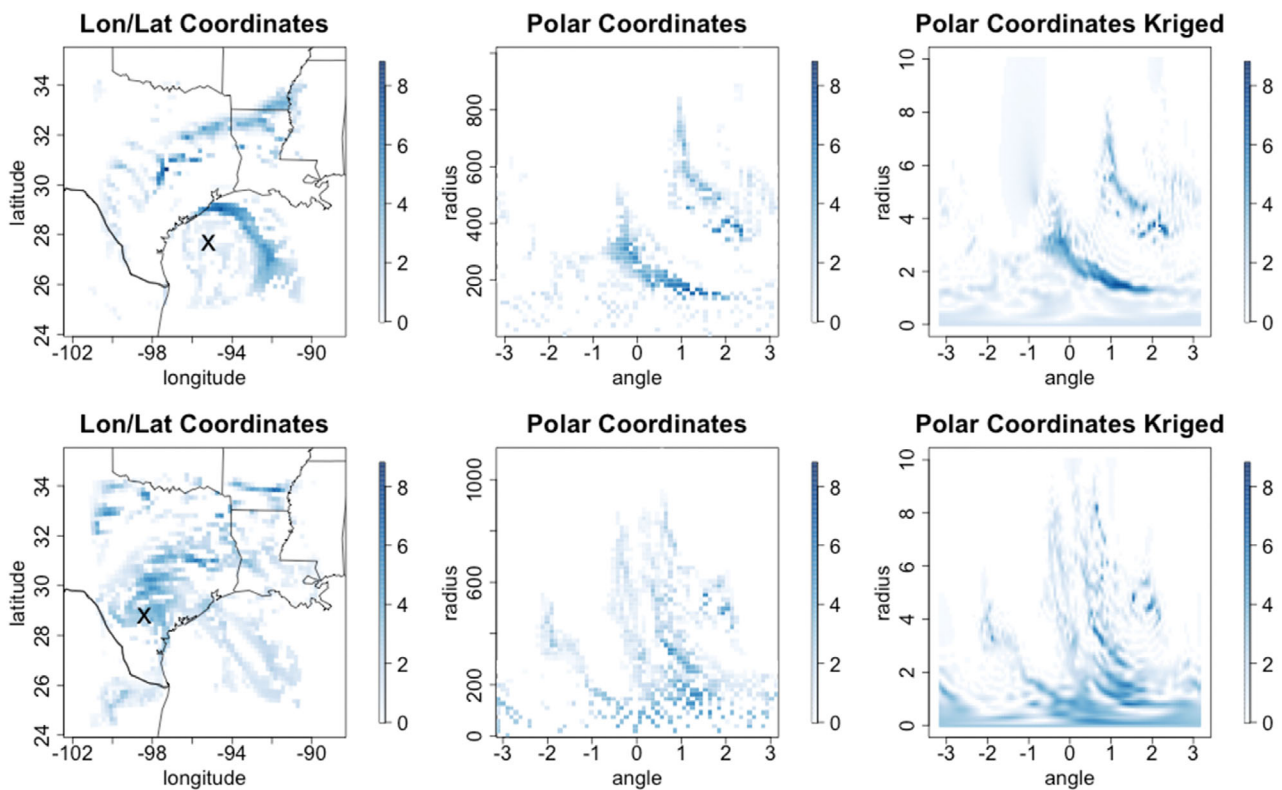


FIGURE B1 The first column contains two snapshots of for TC Fay on September 6, 2002 at 11 p.m. and September 8, 2002 at midnight. The second column are the same data, in polar coordinates, while the third contains the kriged data that are regular in polar coordinates.

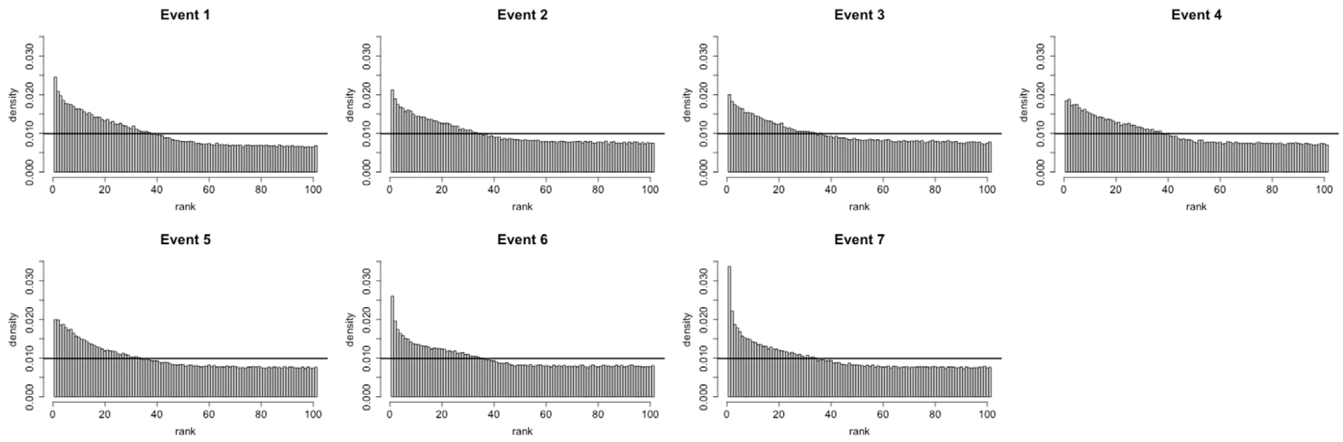


FIGURE B2 Verification rank histograms for each tropical cyclone event in the cross-validation experiment using F as a standard normal cdf in (1).

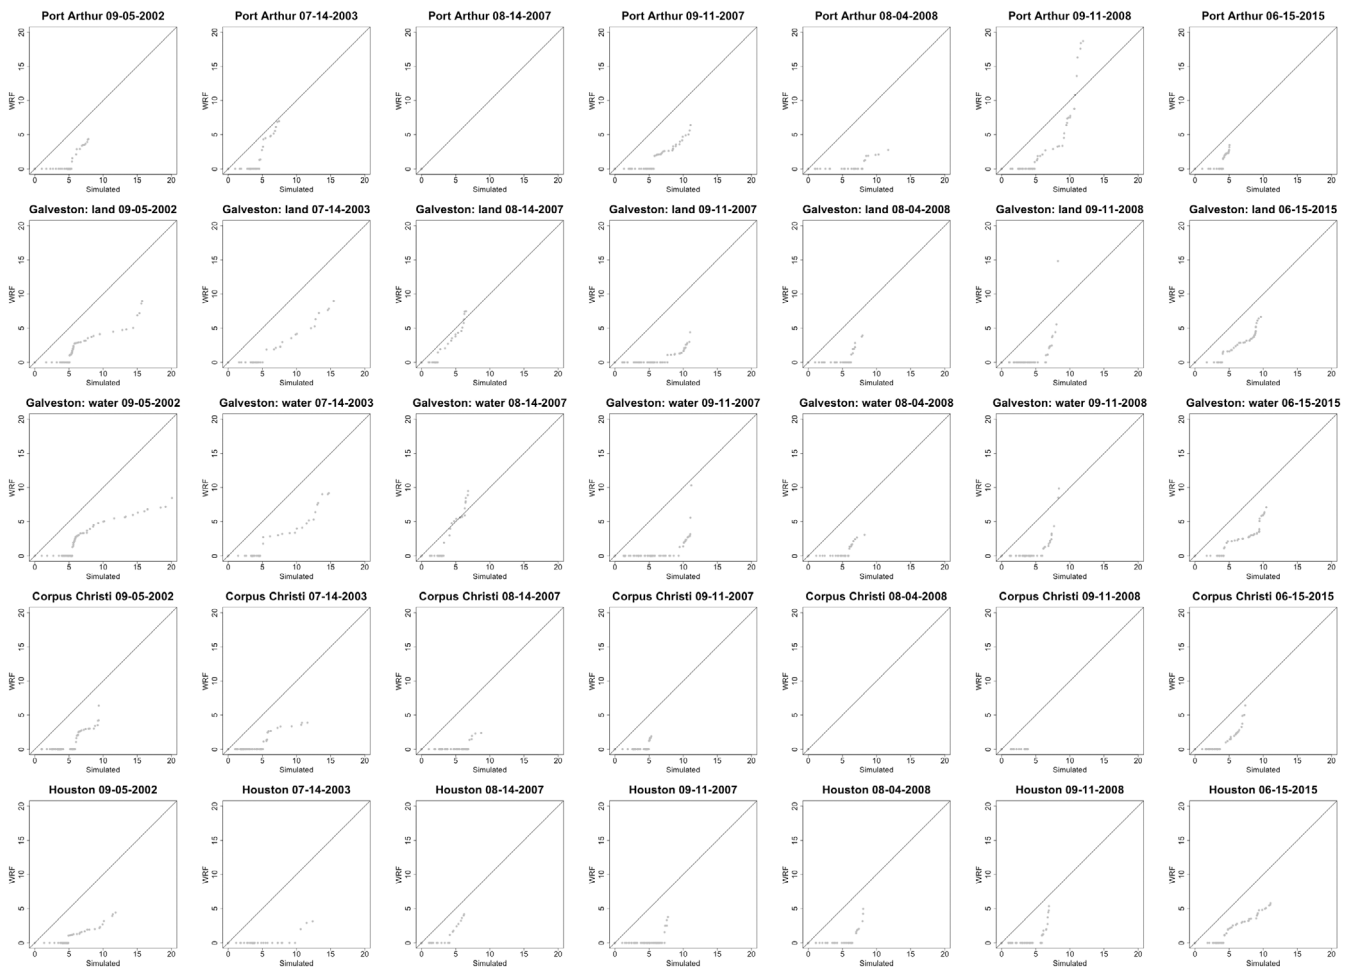


FIGURE B3 Q-Q plots for the fixed mean trend model.

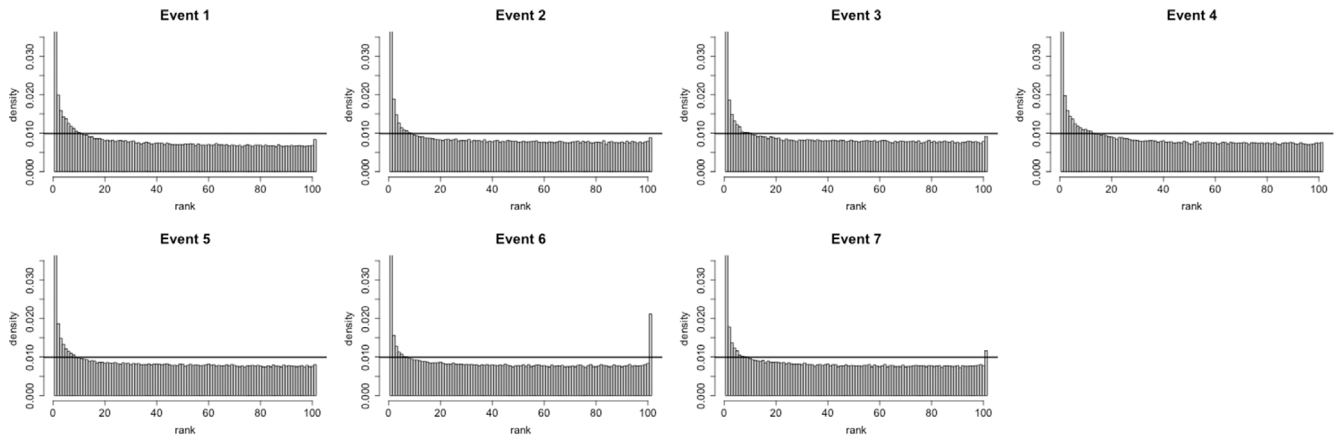


FIGURE B4 Verification rank histograms for the stochastic mean-only model.

TABLE B1 Percent of variance explained for the first seven principal component models using a multiple linear regression model for $\mu_\ell(\cdot)$ versus the random forest model in the text.

Model	μ_1	μ_2	μ_3	μ_4	μ_5	μ_6	μ_7
Linear regression	0.69	0.42	0.05	0.16	0.09	0.12	0.09
Random forest	0.97	0.94	0.86	0.83	0.81	0.86	0.86

where r is a distance of a pixel in kilometers from the storm center at a particular time point, and $RMAX$ is the radius of maximal winds at the corresponding time point. This taper function maintains the exact simulated precipitation fields below the radius of maximal winds, and applies a squared exponential taper beyond this radius with a range of 60 km, which was seen to work well in visual assessments.

B.4 Model without residual process

The model (2) includes both a stochastic mean trend, as well as stochastic residual component. We performed the cross-validation experiment under two separate models that remove the residual process, so that

$$Y(r, \theta, t) = \sum_{\ell=1}^L c_\ell(t) \varphi_\ell(r, \theta).$$

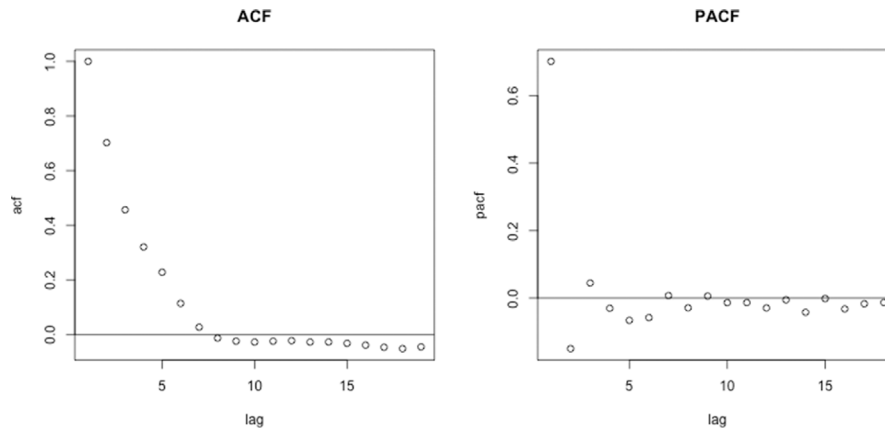


FIGURE B5 Averaged autocorrelation (ACF) and partial autocorrelation (PACF) functions for the $d_0(r, t), \{d_{1,m}(r, t), d_{2,m}(r, t)\}_m$ coefficients.

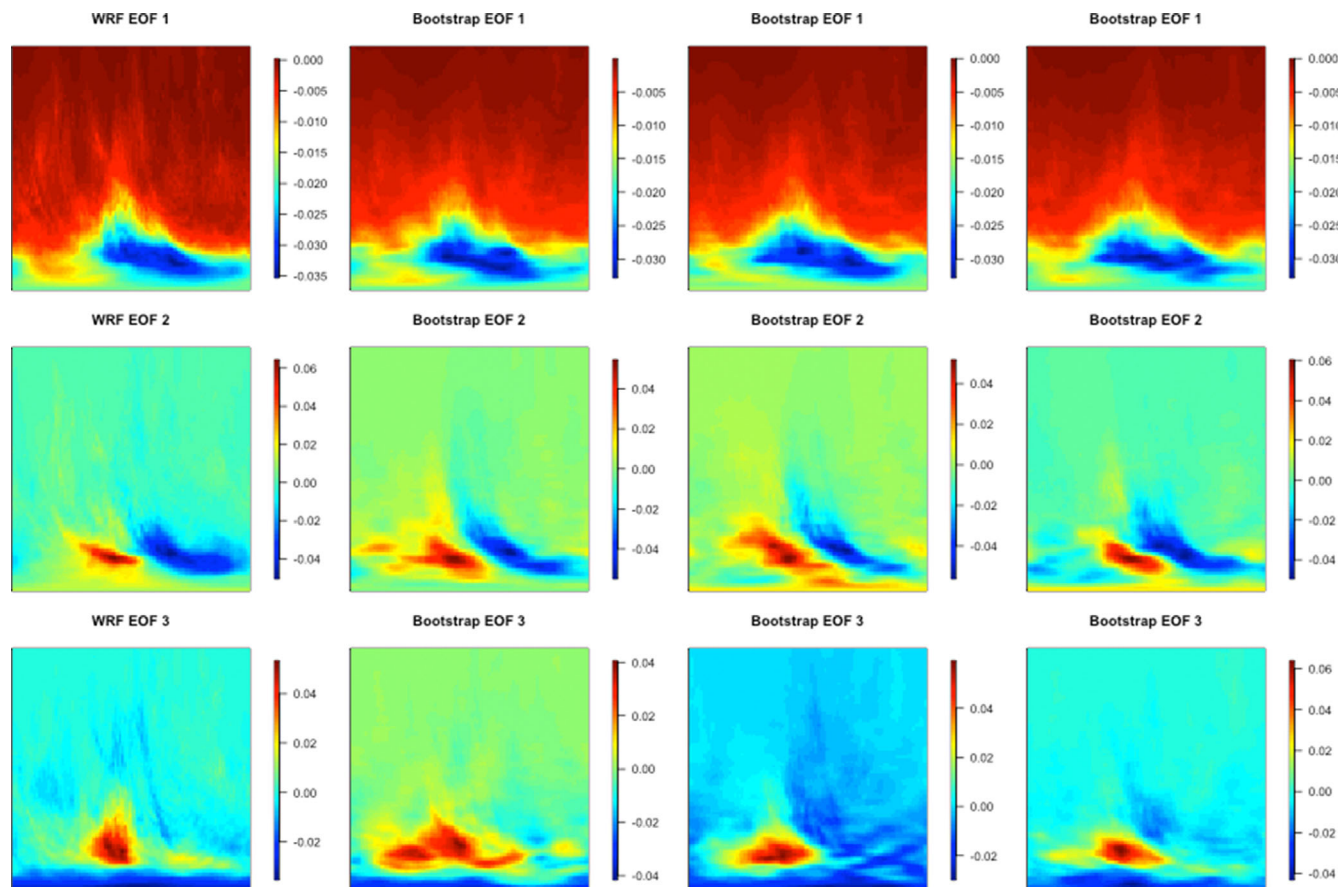


FIGURE B6 Parametric bootstrap estimates of the first three EOFs of (2) in the cross-validation study with storm Ike leftout.

In the first experiment, we allow $c_\ell(t)$ to be random as in the manuscript; in the second experiment, we use fixed predicted values according to the fitted random forests, that is we update (3) to $c_\ell(t) = \mu_\ell(t)$ without the random time series process. Figures B3 and B4 show the local Q-Q plot studies for the fixed, non-random mean approach, and the stochastic mean-only models, respectively. With the fixed-mean model, there is no quantification of uncertainty, and clearly some bias is introduced through the model's nonlinear transformations. With the stochastic mean-only model, there is evidence of clear underdispersion in the simulated precipitation distributions, suggesting that the residual component of (2) is a necessary piece of the model.

B.5 Models for basis function coefficients

The choice to model $\mu_\ell(t)$ in (3) as a random forest is due to the apparent nontrivial and nonlinear relationships between the features and the principal components. The random forest model for $\mu_\ell(t)$ explains a substantial amount of variability in the coefficients $c_\ell(t)$ as compared to a multiple linear regression model—Table B1 compares the r-squared values for each model fit for the first seven principal components. There is, essentially, little-to-no predictive ability when using a linear model, whereas the nonlinear relationships identified by the random forests explain more than 80% of variability for these coefficients. Of course, it is unsurprising that a nonparametric model will outperform a straightforward linear model, but in this case it seems critical to use the nonlinear model given its vastly superior explanatory ability.

B.6 Residual process autocorrelation

Figure B5 shows lag-wise averaged empirical autocorrelation and partial autocorrelation functions for the $d_0(r, t)$, $\{d_{1,m}(r, t), d_{2,m}(r, t)\}_m$ coefficients, averaged over m and r . In our model we assume an autoregressive structure of order one, which seems reasonable given the approximate exponential decay of the autocorrelation and negligible values of the partial autocorrelation function of after order one.

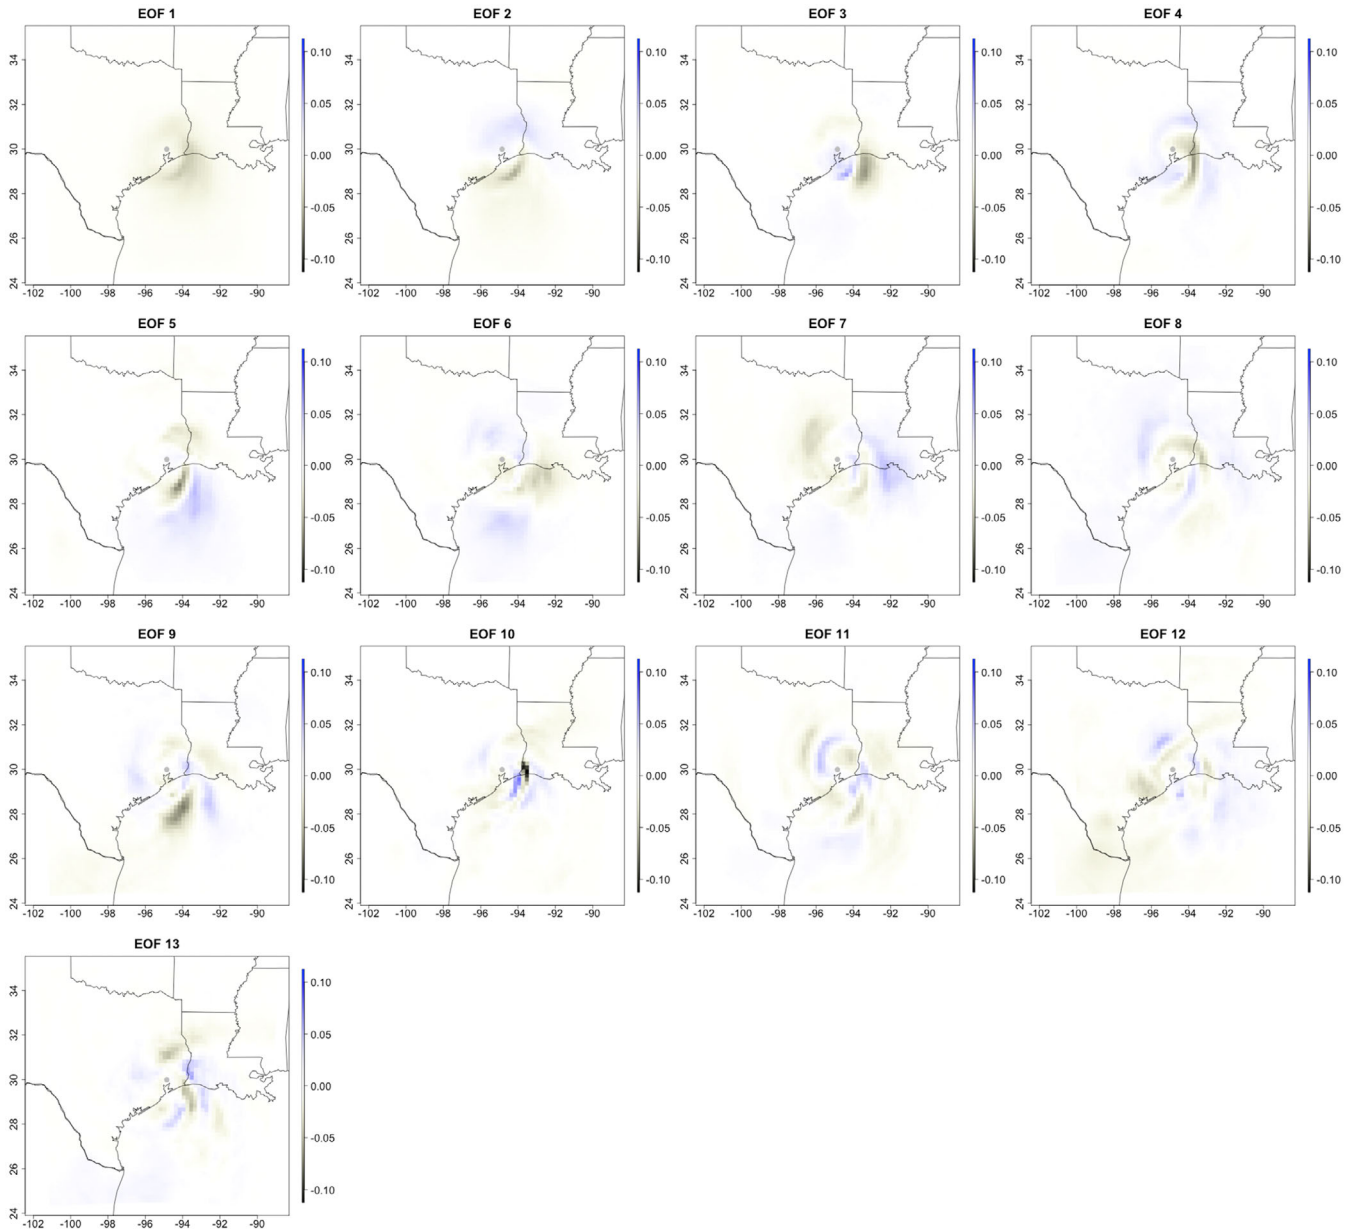


FIGURE B7 All empirical orthogonal functions the storm center indicated by a grey dot.

TABLE B2 Variable importance measured as total reduction in variance for all principal components' random forest models; values are scaled by 1×10^{-5} .

	c_1	c_2	c_3	c_4	c_5	c_6	c_7	c_8	c_9	c_{10}	c_{11}	c_{12}	c_{13}
Radius of maximal winds	1.07	0.65	0.48	0.56	0.34	0.19	0.18	0.29	0.21	0.16	0.14	0.16	0.11
Pressure deficit at center	9.97	4.61	1.06	0.70	0.89	0.64	0.61	0.33	0.35	0.31	0.22	0.16	0.14
U component of direction	1.57	1.06	0.78	0.47	0.41	0.29	0.24	0.34	0.29	0.20	0.14	0.16	0.14
V component of direction	3.82	0.84	0.98	0.69	0.45	0.51	0.51	0.39	0.22	0.22	0.31	0.22	0.18
U component of center	2.88	1.91	0.74	0.44	0.32	0.38	0.35	0.27	0.31	0.23	0.15	0.17	0.15
V component of center	1.86	1.10	0.61	0.96	0.40	0.32	0.29	0.24	0.22	0.17	0.18	0.13	0.18
Distance to coast	1.44	1.04	0.45	0.64	0.26	0.37	0.33	0.23	0.19	0.19	0.14	0.15	0.15

B.7 EOF estimation

It is not a priori clear that our method of estimating the EOFs in (2) will work, as EOFs are not invariant under nonlinear transformations. Figure B6 shows parametric bootstrap samples of the first three EOFs using simulated data from the model. Note these were estimated in the cross-validation study with TC Ike held out, to illustrate that even with a slightly reduced sample size, our estimates of these functions are reasonable. The exact spatial patterns in higher-order EOFs do not always match the bootstrap estimates, but the amount of variability explained is still consistent with our estimation approach.

As a final addition, we include the full set of 13 EOFs used in the application, as well as their coefficient's random forest variable importance, seen in Figure B7 and Table B2.