

Water Resources Research

RESEARCH ARTICLE

10.1029/2017WR022473

Key Points:

- Uncertainty in model parameters propagates to weather simulations
- Spatial processes on model parameters enable gridded simulation
- BayGEN is coupled with crop simulation models for agricultural decision support

Correspondence to:

A. Verdin,
averdin@umn.edu

Citation:

Verdin, A., Rajagopalan, B., Kleiber, W., Podestá, G., & Bert, F. (2019). BayGEN: A Bayesian space-time stochastic weather generator. *Water Resources Research*, 55, 2900–2915. <https://doi.org/10.1029/2017WR022473>

Received 22 DEC 2017

Accepted 14 MAR 2019

Accepted article online 21 MAR 2019

Published online 8 APR 2019

BayGEN: A Bayesian Space-Time Stochastic Weather Generator

Andrew Verdin¹ , Balaji Rajagopalan^{1,2} , William Kleiber³, Guillermo Podestá⁴, and Federico Bert⁵

¹Minnesota Population Center, University of Minnesota, Minneapolis, MN, USA, ²Cooperative Institute for Research in Environmental Sciences, University of Colorado, Boulder, CO, USA, ³Department of Applied Mathematics, University of Colorado, Boulder, CO, USA, ⁴Rosenstiel School of Marine and Atmospheric Sciences, University of Miami, Miami, FL, USA, ⁵Asociación Argentina de Consorcios Regionales de Experimentación Agrícola, Ciudad de Buenos Aires, Argentina

Abstract We present a Bayesian hierarchical space-time stochastic weather generator (BayGEN) to generate daily precipitation and minimum and maximum temperatures. BayGEN employs a hierarchical framework with data, process, and parameter layers. In the data layer, precipitation occurrence at each site is modeled using probit regression using a spatially distributed latent Gaussian process; precipitation amounts are modeled as gamma random variables; and minimum and maximum temperatures are modeled as realizations from Gaussian processes. The latent Gaussian process that drives the precipitation occurrence process is modeled in the process layer. In the parameter layer, the model parameters of the data and process layers are modeled as spatially distributed Gaussian processes, consequently enabling the simulation of daily weather at arbitrary (unobserved) locations or on a regular grid. All model parameters are endowed with weakly informative prior distributions. The No-U Turn sampler, an adaptive form of Hamiltonian Monte Carlo, is used to maximize the model likelihood function and obtain posterior samples of each parameter. Posterior samples of the model parameters propagate uncertainty to the weather simulations, an important feature that makes BayGEN unique compared to traditional weather generators. We demonstrate the utility of BayGEN with application to daily weather generation in a basin of the Argentine Pampas. Furthermore, we evaluate the implications of crop yield by driving a crop simulation model with weather simulations from BayGEN and an equivalent non-Bayesian weather generator.

1. Introduction

Time series of weather variables can assist in agricultural and water resources planning and decision-making as input to crop simulation, hydrologic, and other process-based models (Berger, 2000; Berger et al., 2006; Bert et al., 2006, 2007, 2014; Ferreyra et al., 2001; Freeman et al., 2009; Happe et al., 2008; Schreinemachers & Berger, 2011). However, the historical record can be limited in that it covers too short of a period or it can be incomplete or inaccurate due to missing data and human error. Many process-based models require complete daily sequences of input data (i.e., missing data are not allowed), so the problem of “filling in the blanks” arises. If the problem requires the use of distributed hydrologic models, the user is often required to input daily weather sequences on a regular grid, which is difficult, if not impossible, to collect for even moderately sized basins. Alternatively, spatial interpolation based on available observations is required, introducing yet another possible source of error. Stochastic weather generators have long been used as a way to produce daily weather sequences for use in such applications, as they provide a means to simulate a plausible range of climatic scenarios while representing the true underlying variability of the natural processes being modeled. For agricultural planning on seasonal to multidecadal scales, weather generators can be used to simulate daily input data consistent with forecasted or projected scenarios (Apipattanas et al., 2010; Verdin et al., 2018), which can drive crop simulation models to assess how such scenarios may impact crop yields. A more complete history on stochastic weather generators can be found in Wilks and Wilby (1999) and Verdin, Rajagopalan, Kleiber, and Katz (2015); a brief summary is provided below.

Long held in regard as the original stochastic weather generator, the methodology of Richardson (1981) has paved the way for other approaches to producing synthetic daily weather sequences. These traditional single-site weather generators model precipitation occurrence as a chain-dependent process (Katz, 1977); thus, they are able to capture the length and frequency of wet and dry spells. Precipitation amounts are

modeled using various probability distributions (e.g., gamma and lognormal). Minimum and maximum temperatures are modeled using autoregressive time series models. There have been several multisite extensions to traditional weather generators (Baigorria & Jones, 2010; Khalili et al., 2009; Qian et al., 2002; Wilks, 1998; Wilks & Wilby, 1999). However, parametric weather generators can quickly become unwieldy when applied for multisite simulation, as a large number of parameters is needed to capture the spatial correlation and statistics (Mehrotra et al., 2006). Wilks (2008, 2009) showed that using spatial process models on weather generator parameters eases the simulation of gridded daily weather that adapts to local climates.

Nonparametric weather generators using k -nearest neighbor time series bootstrapping, proposed for single-site (Rajagopalan & Lall, 1999) and subsequently extended to multisite (Apipattanavis et al., 2010; Beersma & Buishand, 2003; Buishand & Brandsma, 2001; Sharif & Burn, 2007; Yates et al., 2003), are robust in their ability to capture nonlinearities in both space and time and can easily capture the spatial correlation and statistics for multisite simulation. However, they have two main limitations—(i) due to bootstrap resampling of the historical record, values outside the historical observations are not possible, and (ii) they cannot easily simulate sequences at unobserved locations.

Generalized linear models (GLMs) can greatly reduce the effort of modeling nonnormal variables through a suite of link functions and can more easily simulate at arbitrary locations when coupled with spatial models (Kleiber et al., 2012, 2013; Verdin, Rajagopalan, Kleiber, & Katz, 2015). Early use of GLMs for weather generation was by Stern and Coe (1984), with subsequent work by Yang et al. (2005) and Chandler (2005). GLMs can incorporate any number of covariates to facilitate the simulation of weather sequences under various conditions, such as wet or dry years. Such covariates have included large-scale climate drivers (i.e., El Niño Southern Oscillation and North Atlantic Oscillation) and total seasonal precipitation and mean temperature trajectories, to name a few (Asong et al., 2016; Furrer & Katz, 2008; Hauser & Demirov, 2013; Verdin et al., 2018).

GLM-based space-time stochastic weather generators have become more popular recently due to their ability to generate weather sequences at any desired location (e.g., grids with different spatial resolutions) and conditioned on large-scale climate drivers. However, they do not incorporate model parameter uncertainty in their simulations. Specifically, point estimates of model parameters based on maximum likelihood are fixed for all ensemble members, and the variability between ensemble members result solely from the residual terms. Lack of treatment of parameter uncertainty is a major shortcoming, as it may reduce the variability of the simulated weather sequences, which, from a planning perspective, can lead to overconfident estimates of system variables by not representing the full range of plausible future scenarios. Bayesian methods are known for their ability to capture uncertainty by providing a full distribution of model parameters. A Bayesian stochastic weather generator would be advantageous in that uncertainty in model parameters would propagate to the simulated weather sequences and, consequently, to process-based model output for natural resources decision support.

Bayesian methods are widely used to quantify uncertainty in a variety of models and applications. They are computationally intensive, but with increasing computation power, Bayesian methods are becoming more attractive. One of the more common Bayesian methods is Bayesian model averaging, first introduced by Leamer (1978) and used in a variety of disciplines including hydrology, climate science, ecology, economics, medicine, and politics. The applications include combining information from multiple sources (e.g., multiple model outputs), making skillful predictions, calibrating forecasts, risk assessment, and quantifying model uncertainty (Duan et al., 2007; Montgomery & Nyhan, 2010; Raftery et al., 2005; Slougher et al., 2007; Trujillo-Barreto et al., 2004; Viallefont et al., 2001; Volinsky et al., 1997; Vrugt & Robinson, 2007; Wintle et al., 2003; Wright, 2008).

Bayesian hierarchical models have been developed for a variety of applications—coupling with stochastic weather generators for enhanced decision making (Hashmi et al., 2009; Pezzulli et al., 2006), spatial and temporal modeling of precipitation extremes (Cooley et al., 2007; Cooley & Sain, 2010; Reich & Shaby, 2012), daily precipitation modeling (Lima & Lall, 2009; Smith & Robinson, 1997), analysis of precipitation and temperature trends (Tebaldi et al., 2004; Tebaldi & Sanso, 2009), and Bayesian kriging (Aelion et al., 2009; Cui et al., 1995; Handcock & Stein, 1993; Jin et al., 2014; Omre, 1987; Sahu & Mardia, 2005; Verdin, Rajagopalan, Kleiber, & Funk, 2015). A spatially consistent Bayesian precipitation state generator was developed by Cano et al. (2004), though a full Bayesian weather generator was yet to be developed.

GLM-based space-time weather generators offer a versatile framework for simulating a suite of variables of various types—continuous, discrete, binary, categorical, etc. However, the simulations do not incorporate the uncertainties in the model parameters, which is a notable drawback. Motivated by the need to fully quantify and incorporate parameter uncertainties in daily weather simulation, we developed a GLM-based Bayesian hierarchical space-time stochastic weather generator, hereafter referred to as BayGEN. This is a novel contribution in that, to our knowledge, this is one of the first attempts at a fully Bayesian weather generator. There are Bayesian generators of precipitation and streamflow (see references above), but none that simultaneously simulate a suite of weather variables in space and time. BayGEN builds on a GLM-based space-time weather generator (Verdin, Rajagopalan, Kleiber, & Katz, 2015) by embedding it in a Bayesian hierarchical framework, thus enabling posterior distributions of the model parameters. Sampling from these posterior distributions to generate ensembles of weather sequences robustly propagates model parameter uncertainty to the weather simulations. Furthermore, with a Bayesian hierarchical framework, expert knowledge about the weather system can be easily incorporated in the model structure and in the prior distributions.

The BayGEN model is presented in section 2, followed by a description of the study region and data in section 3. We discuss the results in section 4, and section 5 concludes this manuscript with a summary of the research and future work.

2. BayGEN

2.1. Model Definition

There are three conceptual layers in the hierarchy of the BayGEN model. The first layer (the data layer) explicitly models the data, in this case the observed precipitation occurrence and amount, and maximum and minimum temperatures at a network of weather stations. The second layer (the process layer) models the latent Gaussian process (i.e., linear regression on an unobservable process) that drives the precipitation occurrence model (see section 2.2 for details). The third layer (the parameter layer) models the spatial dependence of the statistical parameters in the models defined in the data and process layers and includes the prior distributions.

As mentioned, the BayGEN model structure is motivated by the GLM framework of Verdin, Rajagopalan, Kleiber, and Katz (2015) (hereafter GLMGEM). In the GLMGEM framework, at each location, (i) precipitation occurrence is modeled using probit regression; (ii) precipitation amounts are modeled using gamma distributions; and (iii) maximum and minimum temperatures are modeled as separate Gaussian processes. The statistical coefficients within these GLMs are modeled as spatial Gaussian processes to enable simulation of daily weather at arbitrary (unobserved) locations or on a regular grid in such a way that the model adapts to local climatological features. Much of the BayGEN model is similar to that of Verdin, Rajagopalan, Kleiber, and Katz (2015), with the introduction of the Bayesian framework to account for parametric uncertainty representing the main novelty of this section.

Let $O(\mathbf{s}, t)$, $A(\mathbf{s}, t)$, $Z_X(\mathbf{s}, t)$, and $Z_N(\mathbf{s}, t)$ denote precipitation occurrence, amount, maximum temperature, and minimum temperature, respectively, at spatial location $\mathbf{s} \in \mathbb{R}^2$ and day $t \in \mathbb{Z}$. In the following sections, we outline the data, process, and parameter layers of the hierarchical model for each of these processes.

To ease the exposition, an overview of the statistical model is included in equations (1)–(9), followed by detailed descriptions of the model in the following subsections. Equation (1) uses the indicator symbol \mathbb{I} to refer to the relationship with latent Gaussian process, $W_O(\mathbf{s}, t)$. Equation (2) uses “Gamma” to refer to a gamma distribution with spatially varying shape parameters, $\alpha_A(\mathbf{s})$, and spatially and temporally varying scale parameters, $\alpha_A(\mathbf{s})/\mu_A(\mathbf{s}, t)$. Equation (6) defines the denominator of the gamma scale parameters as a regression on some covariates (see section 2.3 for details). Equations (3) and (4) use “GP” to refer to Gaussian processes with means, $\mu_i(\mathbf{s}, t)$, $i = X, N$, and covariance functions, $C_i(t)$, $i = X, N$. Equation (8) uses “GP” to refer to the Gaussian processes for the GLM model parameters with spatially varying means, $\mu_{\beta_j}(\mathbf{s})$, $j = O, A, N, X$, and covariance functions, C_{β_j} , $j = O, A, N, X$. The spatially varying means for these Gaussian processes are defined by first-order polynomials fitted on the point estimates of $\beta_j(\mathbf{s})$, obtained for every sample in the posterior, using the spatial coordinates as predictors. Similarly, equation (9) refers to a Gaussian process for the gamma shape parameter with spatially varying mean, $\mu_{\alpha_A}(\mathbf{s})$, and covariance

function, C_{α_A} . The spatially varying mean is defined by a first-order polynomial fitted on the point estimates, $\alpha_A(\mathbf{s})$, using the spatial coordinates as predictors.

Data layer:

$$O(\mathbf{s}, t) = \mathbb{1}_{\{W_O(\mathbf{s}, t) > 0\}} \quad (1)$$

$$A(\mathbf{s}, t) \sim \text{Gamma}\left(\alpha_A(\mathbf{s}), \frac{\alpha_A(\mathbf{s})}{\mu_A(\mathbf{s}, t)}\right) \quad (2)$$

$$Z_X(\mathbf{s}, t) \sim \text{GP}(\mu_X(\mathbf{s}, t), C_X(t)) \quad (3)$$

$$Z_N(\mathbf{s}, t) \sim \text{GP}(\mu_N(\mathbf{s}, t), C_N(t)) \quad (4)$$

Process layer:

$$W_O(\mathbf{s}, t) \sim N(\mu_O(\mathbf{s}, t), \sigma_O^2) \quad (5)$$

Parameter layer:

$$\mu_A(\mathbf{s}, t) = \exp\left(X_A(\mathbf{s}, t)' \beta_A(\mathbf{s})\right) \quad (6)$$

$$\mu_i(\mathbf{s}, t) = X_i(\mathbf{s}, t)' \beta_i(\mathbf{s}), i = O, N, X \quad (7)$$

$$\beta_j(\mathbf{s}) \sim \text{GP}(\mu_{\beta_j}(\mathbf{s}), C_{\beta_j}), j = O, A, N, X \quad (8)$$

$$\alpha_A(\mathbf{s}) \sim \text{GP}(\mu_{\alpha_A}(\mathbf{s}), C_{\alpha_A}) \quad (9)$$

2.2. Precipitation Occurrence

Precipitation occurrence is defined as any daily amounts exceeding 0.1 mm and is modeled as a binary space-time process that is locally equivalent to a probit regression. Our data layer specification is $O(\mathbf{s}, t) = \mathbb{1}_{\{W_O(\mathbf{s}, t) > 0\}}$, where $\mathbb{1}_S$ denotes an indicator function that is unity if the set condition S holds and is zero otherwise. The process layer specifies $W_O(\mathbf{s}, t)$ as a latent Gaussian process with mean function $\mu_O(\mathbf{s}, t)$ that is independent in time but correlated in space with the associated covariance $C_O(\mathbf{h})$, $\mathbf{h} \in \mathbb{R}^2$. The mean function is assumed to be a regression on covariates 1, $O(\mathbf{s}, t - 1)$, $\cos(2\pi t/365)$, and $\sin(2\pi t/365)$ with associated regression coefficients β_j , $\beta_j(\mathbf{s}), j = 1, \dots, 4$, that vary over space. Assuming β_j , $\beta_j(\mathbf{s})$ are independent Gaussian processes completes the third level of hierarchy. The spatially varying means of these Gaussian processes are defined by first-order polynomials fitted on the point estimates, using the spatial coordinates as predictors, each with isotropic exponential covariance functions defining the second-order structure. The partial sills and range parameters are random variables, and the nuggets are fixed at 0.001. We fix a finitely small nugget value as a computational convenience.

For our data application, multivariate probit regression would necessitate the inclusion of approximately 500,000 additional parameters, which would make the model unwieldy, overparameterized, and unlikely to run in reasonable time. To this end, after the model is fitted, probit regression residuals are calculated for each day and each location from the posterior distributions, and the empirical spatial correlation matrices are calculated. From these empirical correlation matrices, partial sill and effective range parameters are estimated with a fixed nugget of 0.001, which enables simulation at arbitrary locations using the theoretical exponential covariance function. The optimization is achieved by minimizing the squared difference between the empirical correlations and a theoretical exponential variogram model. There are distinct correlation matrices—and thus distinct partial sills and effective ranges—for each calendar month to account for seasonality. In total, there are 60 parameters in the precipitation occurrence model.

2.3. Precipitation Amount

Precipitation amounts are locally modeled as gamma random variables with spatially varying shape parameter, $\alpha_A(\mathbf{s})$, to account for the climatological gradients, and spatially and temporally varying scale

parameter, $\alpha_A(\mathbf{s})/\mu_A(\mathbf{s}, t)$, to account for seasonality. The scale parameter is a function of the shape parameter and a spatiotemporal mean function, $\mu_A(\mathbf{s}, t)$. Local estimates (that is, at a single observation location) of shape and scale parameters are obtained by maximum likelihood, viewing each day as independent realizations of precipitation amounts—thus the likelihood is simply a product of gamma probability density functions. Given these local estimates, spatial dependence is maintained by modeling $\alpha_A(\mathbf{s})$ as a Gaussian process, with spatially varying mean and isotropic exponential covariance function, $C_{\alpha_A}(\mathbf{h})$ for spatial lag $\mathbf{h} \in \mathbb{R}^2$. The spatially varying mean is defined by a first-order polynomial fitted on the point estimates, $\alpha_A(\mathbf{s})$, using the spatial coordinates as predictors. Partial sill and range parameters are treated as random variables, and the nugget effect is fixed at 0.001. We model $\log(\mu_A(\mathbf{s}, t))$ as a Gaussian process whose mean is a regression on covariates 1, $\cos(2\pi t/365)$, and $MT(t)$, where $MT(t)$ denotes monthly average precipitation. The regression coefficients $\beta_{j, A}(\mathbf{s}), j = 1, \dots, 3$ are modeled as independent Gaussian processes with spatially varying means defined by first-order polynomials fitted on the point estimates, using the spatial coordinates as predictors, each with isotropic exponential covariance functions. Again, the partial sill and range parameters are random variables, while the nugget effects are fixed at 0.001. In total there are 60 parameters in the precipitation amounts model.

Although locally $A(\mathbf{s}, t)$ is modeled as a gamma random variable, we require spatially correlated fields of precipitation amounts. This is attained via a spatially varying Gaussian copula function (Chilés & Delfiner, 1999), that is, $W(\mathbf{s}, t) = \Phi^{-1}(G_{s, t}(A(\mathbf{s}, t)))$ where $G_{s, t}$ is the local gamma cumulative distribution function (cdf) and Φ^{-1} is a standard normal quantile function. Setting $W(\mathbf{s}, t) = W_O(\mathbf{s}, t)$ maintains the physical relationship between precipitation amounts and occurrence in that there is spatial continuity between areas with zero precipitation bordering positive precipitation events.

2.4. Maximum and Minimum Temperatures

As in Richardson (1981), we condition the maximum and minimum temperature models on precipitation occurrence. The maximum temperature process, $Z_X(\mathbf{s}, t)$, is modeled as a Gaussian process whose mean is a regression on covariates 1, $Z_X(\mathbf{s}, t - 1)$, $\cos(2\pi t/365)$, $\sin(2\pi t/365)$, and $O(\mathbf{s}, t)$. Z_X has isotropic exponential spatial covariance functions $C_{m, X}(\mathbf{h}), \mathbf{h} \in \mathbb{R}^2, m = 1, \dots, 12$ that are separately modeled for each month of the year to allow for climatologically varying spatial dependence structure. The partial sills and range parameters are treated as random variables, and the nuggets are fixed at 0.001. Similarly, minimum temperature, $Z_N(\mathbf{s}, t)$, is also a Gaussian process whose mean is a regression on covariates 1, $Z_N(\mathbf{s}, t - 1)$, $\cos(2\pi t/365)$, $\sin(2\pi t/365)$, and $O(\mathbf{s}, t)$, with monthly varying isotropic exponential spatial covariance functions $C_{m, N}(\mathbf{h}) \in \mathbb{R}^2, m = 1, 2, \dots, 12$.

Note that both maximum and minimum temperatures on day t depend on the current day's precipitation occurrence. Maximum temperature on day t depends on the previous day's maximum temperature; minimum temperature on day t depends on the previous day's minimum temperature. This mean function dependency serves to temporally correlate the maximum and minimum temperature processes without having to specify a complete space-time covariance structure.

Each temperature process' mean function is a regression; the regression coefficients are denoted by $\beta_{j, X}(\mathbf{s})$ and $\beta_{j, N}(\mathbf{s})$ for $j = 1, \dots, 5$ for maximum and minimum temperature, respectively. These coefficients are assumed to vary in space to capture local climatological influences of the covariates and are modeled as Gaussian processes with means defined as first-order polynomials fitted on the point estimates, using the spatial coordinates as predictors, each with isotropic exponential covariance functions. The partial sills and range parameters are random variables, while the nuggets are fixed at 0.001. In total there are 112 parameters in both maximum temperature and minimum temperature models, totaling 224 parameters.

2.5. Likelihood and Priors

Let $\theta \in \mathbb{R}^P, P = 112 + 112 + 60 + 60 = 334$ be the vector of model parameters and let $\mathbf{Y} \in \mathbb{R}^N, N = 19,358 \times 13 \times 4 = 1,006,616$ be the observational data vector of all occurrences, amounts, maximum temperatures, and minimum temperatures. We seek the posterior distribution $p(\theta | \mathbf{Y}) \propto f(\mathbf{Y} | \theta)\pi(\theta)$, where f is the data likelihood and π is the joint prior distribution. The likelihood function $f(\mathbf{Y} | \theta)$ factors into products of marginal likelihoods for each weather variable. The likelihood for the precipitation occurrence process is simply a product of Bernoulli probability mass functions; these probabilities are modeled via probit

regressions. The likelihood for the precipitation amount process is a product of gamma probability density functions. The likelihoods for the maximum and minimum temperature processes are independently multivariate normal. The likelihoods for the exponential covariance function parameters are independently univariate normal.

We use weakly informative independent normal priors, thus $\pi(\theta)$ factors as follows: the prior distributions for the GLM regression parameters are set to be independent normal distributions centered on the respective MLE with standard deviations set to 100. Exponential covariance parameters, specifically the partial sill and range parameters, are given independent normal prior distributions centered on the least squares optimization estimates. The optimization is achieved by minimizing the squared difference between the empirical covariances and theoretical exponential variogram models. The standard deviations for the normal priors on the partial sill parameters are defined as the empirical standard deviation of the monthly optimized estimates of partial sill (i.e., the standard deviation of the 12 values used to center the normal distributions). The standard deviation for the normal priors on the range parameters are defined by the standard deviation of all pairwise distances (i.e., the standard deviation of the distance matrix). All nugget parameters are fixed at 0.001. Nonnegative parameters (i.e., partial sill and range parameters) are given truncated normal priors with lower bounds of 0. The range parameters are assigned an upper bound equal to the greatest distance between all observation sites multiplied by 1,000 (i.e., the maximum value in the distance matrix multiplied by 1,000). Truncated normal priors are defined by specifying bounds on the parameters themselves, regardless of the specified priors, which is an option in the *rstan* library.

2.6. Implementation

We use the R (R Core Team, 2014) library *rstan* (Carpenter, 2015; Stan Development Team, 2015a, 2015b) for model development. This software employs the No-U Turn sampler (NUTS), an adaptive form of Hamiltonian Monte Carlo (HMC) sampling (see Hoffman & Gelman, 2014). Other existing software typically uses Markov chain Monte Carlo (MCMC) sampling, which, while effective, can take a long time to converge with its random walk behavior. HMC avoids the random walk behavior of Markov chain Monte Carlo by taking steps based on first-order gradient information. The efficacy of HMC is highly sensitive to user-defined parameters, particularly step size and the number of steps per iteration. NUTS was developed to adaptively and intrinsically set these key HMC parameters by (i) building a set of likely candidate points spanning the target distribution until it “U-turns” or doubles back and retraces its steps, eliminating the need for the user to set a discrete number of steps per iteration, and (ii) implementing a primal-dual averaging (Nesterov, 2009) method for adapting the step size parameter on the fly, eliminating the need for the user to set a discrete step size (Hoffman & Gelman, 2014). The partial derivatives required by the NUTS algorithm are computed via automatic differentiation variational inference (see Kucukelbir et al., 2017).

We have chosen the *rstan* library because it provides an R interface to the Stan modeling language (Carpenter, 2015; Stan Development Team, 2015a, 2015b), which has an active users group mailing list for real-time troubleshooting.

3. Study Region and Data

Application of BayGEN is focused on the Salado A subbasin (hereafter SAB) of the Argentine Pampas (see Figure 1). The climatology of the region surrounding the SAB exhibits longitudinal gradients, with higher temperatures and lower rainfall in the west and lower temperatures and higher rainfall in the east. The SAB is unique in that it is located where the climatological gradients are the greatest. Additionally, the SAB is one of the most agriculturally productive regions of the Pampas. For a detailed description of the SAB and its hydroclimatology, see Verdin et al. (2018).

Daily time series for three weather variables—precipitation, maximum temperature, and minimum temperature—are available for a network of 17 weather stations located in and around the SAB for the period from January 1961 to December 2013 (53 years). Because four stations have many missing data (>25%), they were removed from the network. The remaining 13 stations were used in the analyses (see Figure 1). The data were collected and organized by the Servicio Meteorológico Nacional (National Meteorological Service) of Argentina.

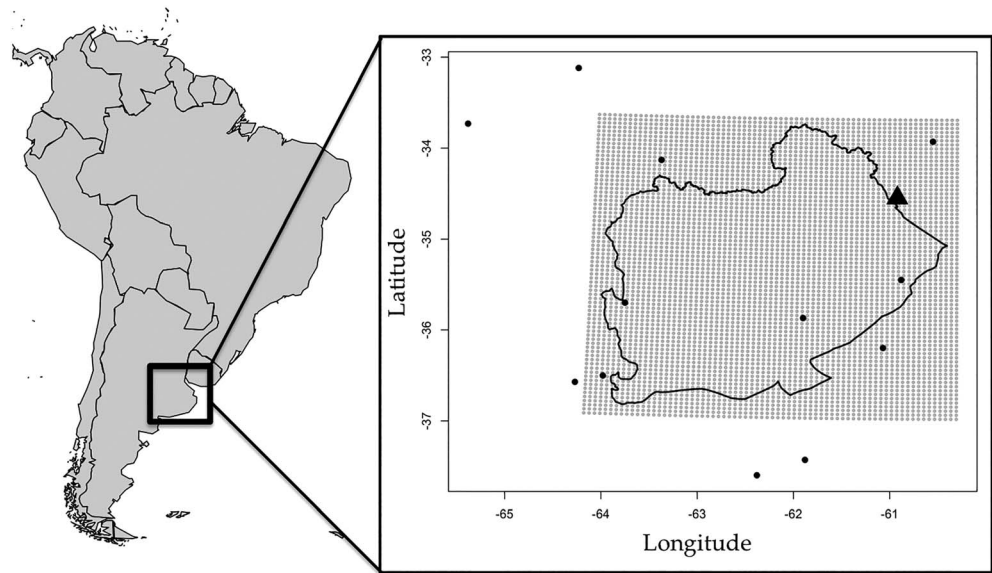


Figure 1. Salado A subbasin location relative to South America; station locations shown as black dots; Junin shown as large black triangle; grid cell locations shown as gray dots; and Salado A sub-basin is outlined.

4. Results

4.1. Model Fit

We employed BayGEN in eight parallel chains (i.e., independent sampling routines) of 1,000 iterations, the first 500 of which are discarded as “burn-in” (i.e., warmup), which results in 4,000 posterior samples of θ . Each parallel chain is initialized with a random sample from a uniform distribution, which ensures that the sampler converges to the appropriate limiting distribution. An example of the sampler failing to converge to the appropriate limiting distribution would be if one or more of the eight parallel chains converges on different posterior modes. Additionally, it is advantageous to run multiple chains because it is easier to diagnose convergence. Each chain uses NUTS to generate a sequence of plausible values for all parameters in the model. The sampled values from each chain are combined into a single distribution for each model parameter, which produces the full posterior distributions for each parameter in the model. From the posterior distributions, we obtain a full representation of the range of parametric uncertainty. In *rstan*, the data are first transformed to the unconstrained scale, and the initialization values are random samples from a uniform distribution spanning $[-2, +2]$. Numerical convergence is monitored by a weighted combination of the between- and within-sequence variances, as follows:

$$\hat{R} = \sqrt{\frac{\text{var}^+(\theta|\mathbf{Y})}{W}}, \quad (10)$$

where

$$\text{var}^+(\theta|\mathbf{Y}) = \frac{n-1}{n}W + \frac{1}{n}B. \quad (11)$$

In equations (10) and (11), θ is a vector of the parameters of interest, \mathbf{Y} is the response data, W is the within sequence variance, B is the between-sequence variance, and n is the length of the simulated sequences after burn-in. \hat{R} is the factor by which the scale of the current distribution for θ might be reduced if the simulations were continued in the limit $n \rightarrow \infty$. A value of $\hat{R} < 1.1$ implies convergence (Gelman et al., 2004). Every parameter in this model has a value of $\hat{R} < 1.1$ (most are $\hat{R} = 1.0$, which is optimal); thus, all parameters have converged within the allotted burn-in period. Visual inspection of traceplots (i.e., the evolution of the sampled values of θ) was also used to assess convergence (figures not shown), confirming what is reported by \hat{R} .

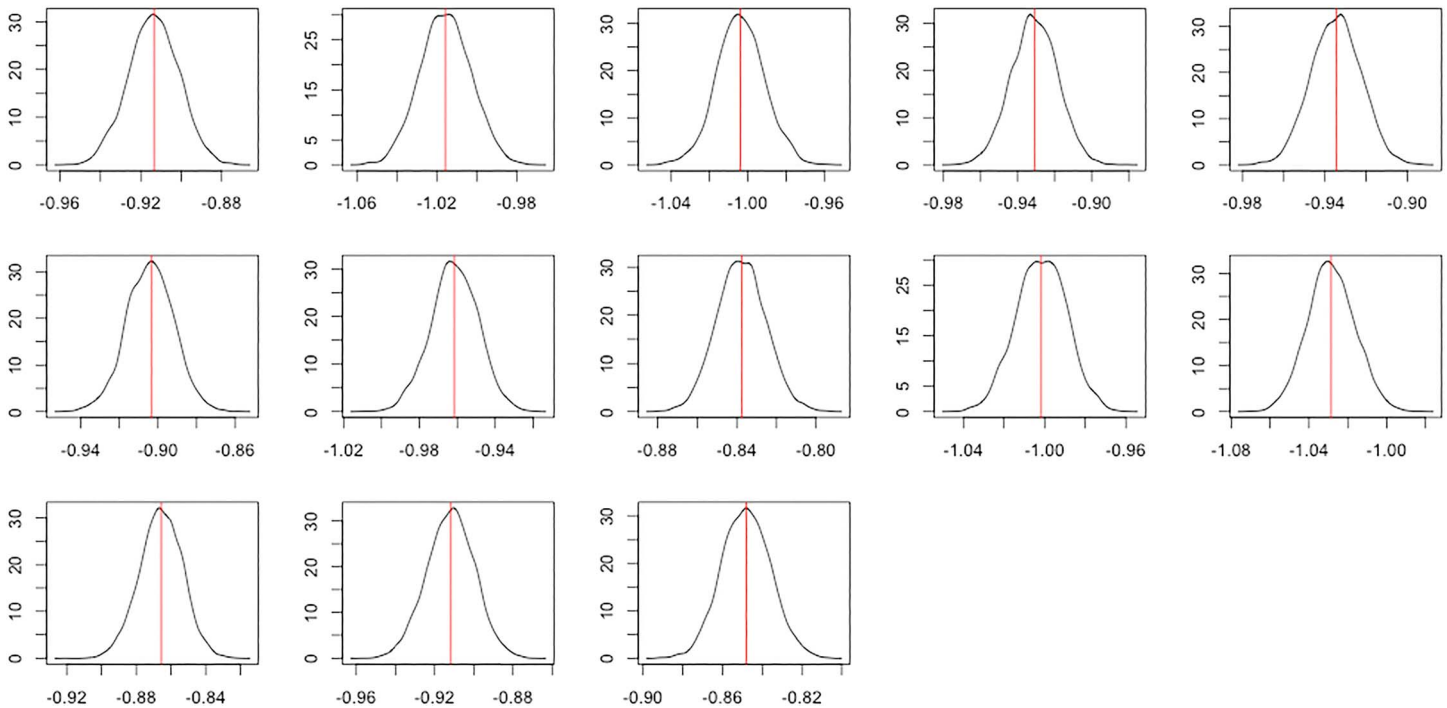


Figure 2. Posterior distributions and maximum likelihood estimates of the intercept terms for the precipitation occurrence model at each of the 13 locations shown in Figure 1.

4.2. Posterior Distributions

The posterior distributions of the model parameters were compared to their respective MLE values for consistency. Figure 2 shows the posterior density of the intercept terms of the precipitation occurrence latent Gaussian processes and the MLE at the 13 station locations. The MLE are very close to the mode of the distributions. Furthermore, the distributions are wide, which is indicative of the uncertainty not quantified by using the point estimates of MLE. Although it is possible to obtain the unconditional distribution for the point MLE from asymptotic arguments, it is not clear that such asymptotics are appropriate for these data. Hence, the uncertainties in the model parameters are not reflected in non-Bayesian GLM-based stochastic weather generation. Posterior distributions of other model parameters exhibit similar relationships with their respective MLE (figures not shown).

4.3. Multisite Weather Simulation

We fitted the BayGEN and generated daily weather ensembles at the network of 13 stations. Each simulated ensemble is 53 years long (the same length as the observational record). As there are 4,000 samples in θ , which form the posterior distribution of the parameters, we generated an ensemble of 4,000 daily weather series, each using a unique parameter set from θ . For each parameter set, the daily precipitation occurrence, amounts, and maximum and minimum temperatures are generated following the equations and methods described in section 2. A suite of monthly and spatial statistics was computed and compared with the observed values.

4.3.1. Model Validation

We computed basic statistics for each simulation, resulting in a distribution of basic statistics. These basic statistics were computed for each month—monthly mean and standard deviations of daily precipitation and minimum and maximum temperatures. Spatial correlations between the simulations across the locations were computed and compared to the historic values. In order to investigate the ability of BayGEN in capturing extremes, we show quantile-quantile (*Q-Q*) plots of observed and simulated daily minimum and maximum temperatures, and domain maximum precipitation.

4.3.2. Basic Statistics

Figure 3 shows boxplots of the means and standard deviations of daily rainfall, maximum temperature, and minimum temperature for each month at Junín (see Figure 1 for the location of this station), with the

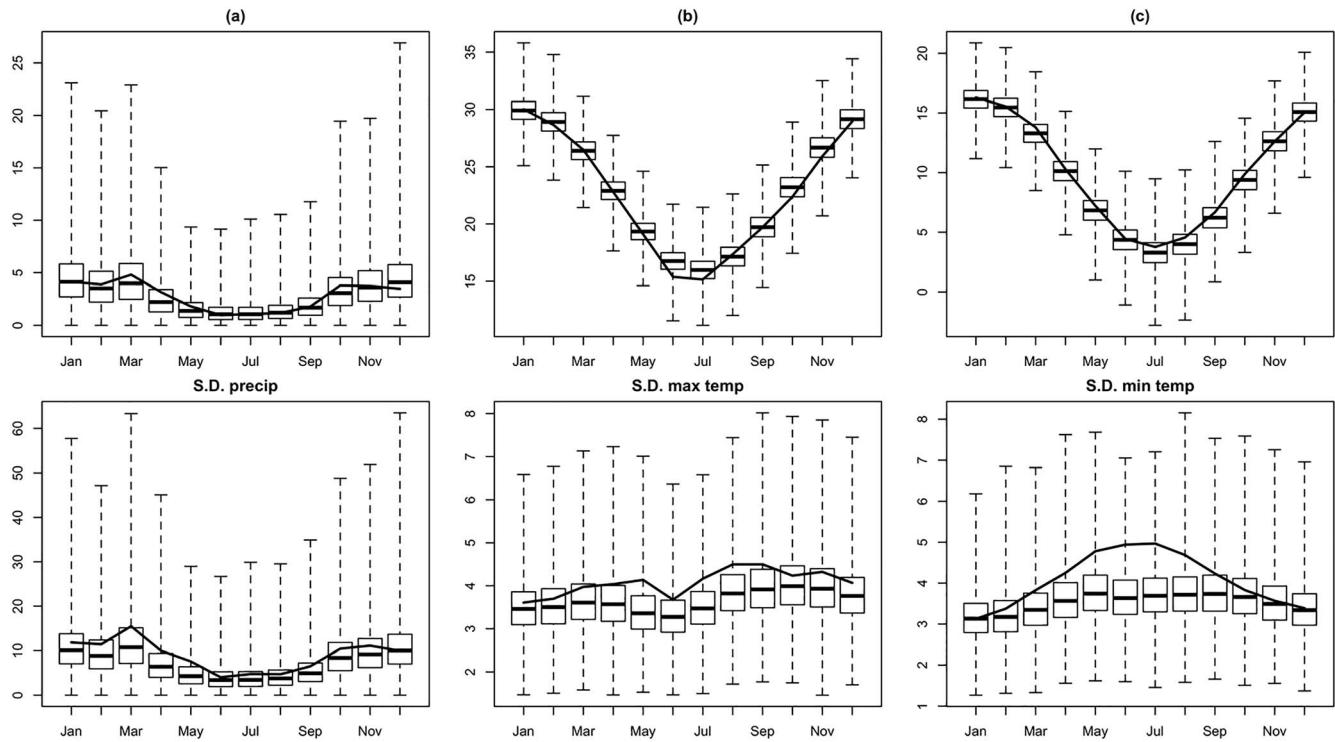


Figure 3. The 4,000 simulations at Junin: (a–c) Climatological mean of daily precipitation, maximum temperature, and minimum temperature; (d–f) monthly standard deviation of daily precipitation, maximum temperature, and minimum temperature. BayGEN shown as boxplots; observed statistics shown as lines.

corresponding observed statistics as solid lines. It can be seen that the historic means of all the variables for all the months (Figures 3a, 3b, and 3c) are well captured by the simulations, as the historic values fall within the interquartile range of the boxplots. The historic monthly standard deviations of precipitation (Figure 3d) are well captured by the BayGEN simulations, while the simulations of maximum and minimum temperatures (Figures 3e and 3f) show smaller standard deviation compared to observations; more so with minimum temperatures.

Figure 4 shows the pairwise correlations of monthly mean precipitation, maximum temperature, and minimum temperature from the simulations along with the observed values, with a 1:1 line for reference. The

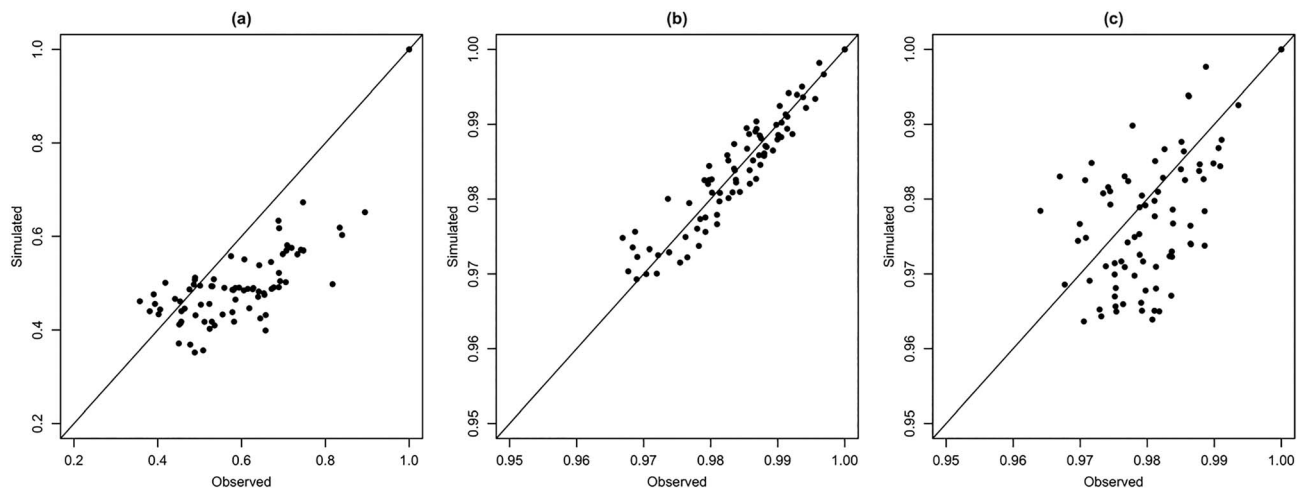


Figure 4. Spatial pairwise correlations of (a) monthly mean precipitation, (b) monthly mean maximum temperature, and (c) monthly mean minimum temperature. Observed correlations are on the x axis, and posterior median correlations are on the y axis.

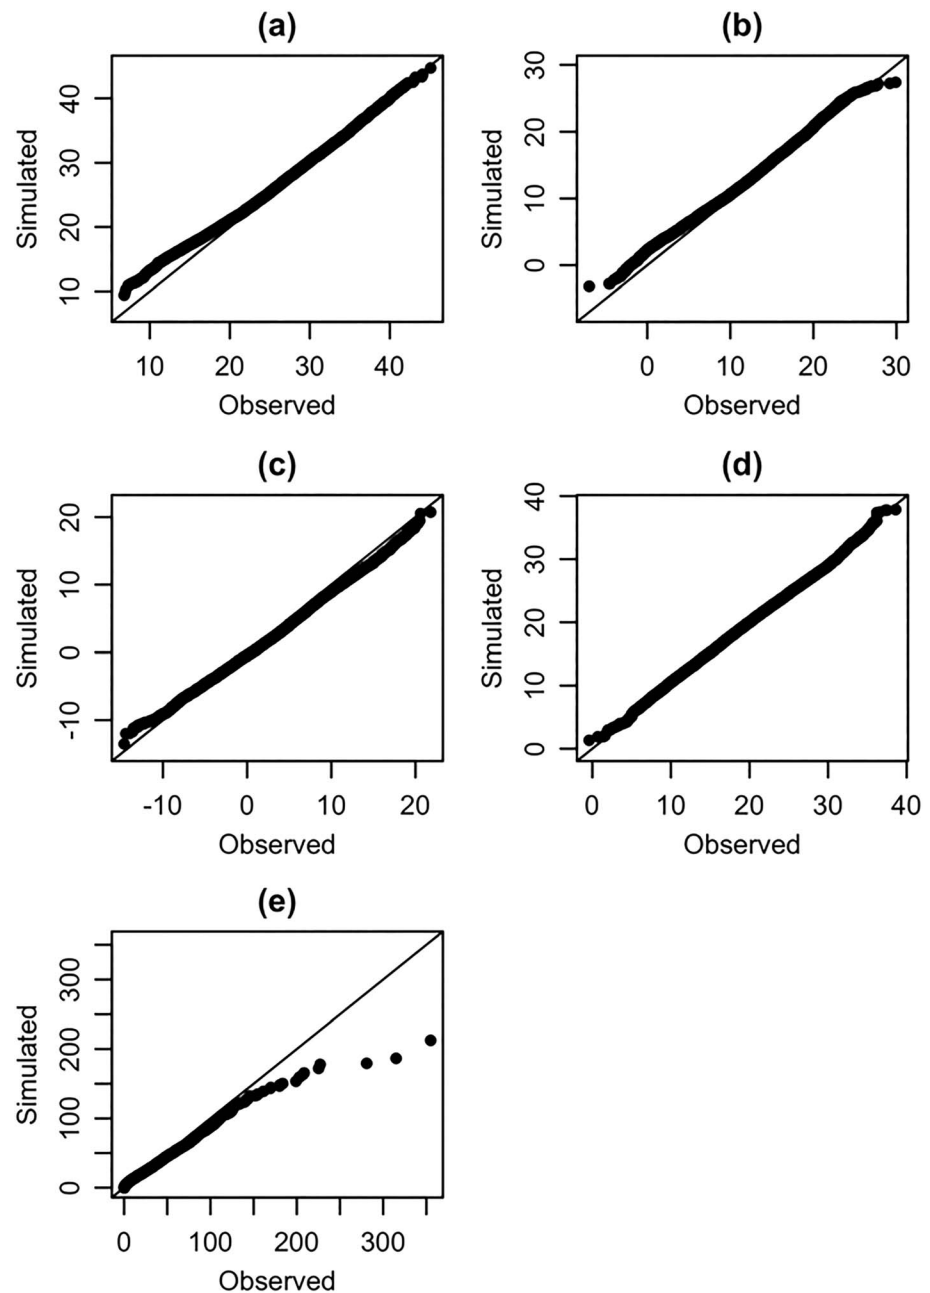


Figure 5. (a–d) Q–Q plots of BayGEN daily domain temperature extrema computed as maximum or minimum of daily values at all the locations: (a) domain maxima of maximum temperatures, (b) domain minima of maximum temperatures, (c) domain minima of minimum temperatures, and (d) domain maxima of minimum temperatures. Units are degrees Celsius. (e) Q–Q plot of BayGEN daily domain maxima precipitation computed as maximum of daily precipitation at all locations. Units are millimeters.

spatial correlations of the variables across the stations are very well captured by the simulations (Figures 4a, 4b, and 4c) with a slight undersimulation of higher correlations in precipitation (Figure 4a).

4.3.3. Extremes

For crop and hydrologic modeling, it is important to reproduce precipitation and temperature extremes. We assess the capability of BayGEN to produce extremes by Q–Q plots of domain extremes. Figure 5 shows Q–Q plots of the observed versus simulated domain daily extrema of temperature and precipitation (i.e., domain maxima and minima of maximum temperature and minimum temperature and domain maxima of daily

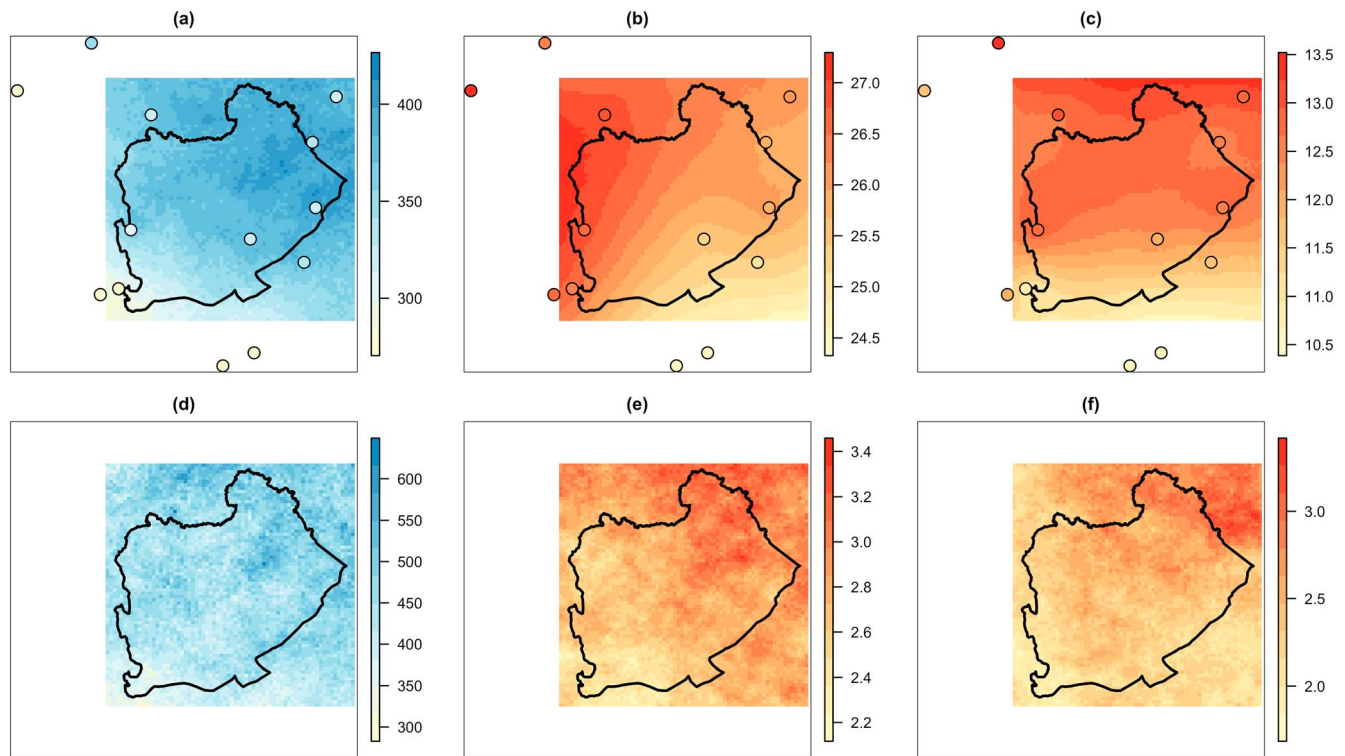


Figure 6. BayGEN: (a–c) Ensemble mean total precipitation, mean maximum temperature, and mean minimum temperature, for the OND season; observed OND values are shown as colored points; (d–f) corresponding 95% ensemble spread. Units are millimeters and degrees Celsius, respectively.

precipitation). Daily extrema across the 13 stations (i.e., domain) for a randomly selected simulation are paired with those from the corresponding historic values in computing the *Q-Q* plot. The quantiles of the simulated and observed daily extreme temperatures lie on the 1:1 line, except for slight deviations in the lower values of maxima of maximum temperatures (Figure 4a) and the higher values of minima of maximum temperatures (Figure 4b). The precipitation extremes are undersimulated for higher values, which is not surprising, as precipitation extremes in space and time are difficult to capture. Furthermore, considering that the model structure does not explicitly model extremes, the performance of the simulations in capturing the extrema is quite good.

4.4. Gridded Weather Simulation

The BayGEN framework is structured to enable weather simulation at arbitrary locations, such as on a regular grid, via spatial process models. First, model parameters are simulated at desired locations using conditional simulation informed by samples from the posterior distribution of the spatial process parameters. Conditional simulation is used instead of other spatial estimation techniques (such as interpolation) to preserve the uncertainty defined by the posterior distributions of spatial process parameters. Gaussian random fields are used to produce spatially correlated residuals, akin to the multisite simulation. For consistency between precipitation occurrence and amounts, the Gaussian random field that is used to simulate precipitation occurrence is transformed to gamma space using a Gaussian copula function (Chilés & Delfiner, 1999). We show the model’s ability in simulating spatial weather scenarios by displaying the ensemble mean and spread from the simulations for the wet season, October–December.

Because gridded weather simulation is computationally intensive, we randomly sample 100 posterior samples from the 4,000 available in the posterior distributions, resulting in 100 gridded weather scenarios. Gridded weather was simulated for the austral spring season (October–December; OND). For visualization, we aggregated daily weather to the seasonal scale (seasonal total precipitation and seasonal average temperatures). The ensemble mean and the 95% ensemble spread (i.e., the 97.5th percentile minus the 2.5th percentile) are calculated and shown in Figure 6. It can be seen that there is coherent spatial structure in the

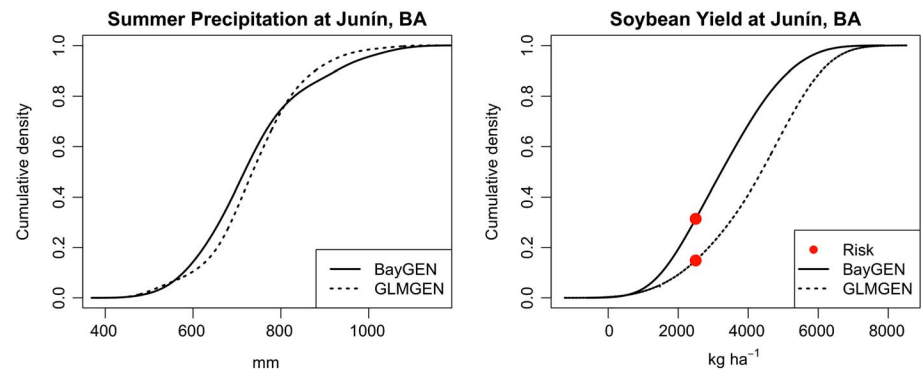


Figure 7. Cumulative density functions of summer growing season (left) total precipitation and (right) soybean yields at Junin, based on 100 simulations of daily weather for a 2-year period. Simulations from BayGEN are shown as a solid line and GLMGEN as a dashed line. Break-even yields are shown as dots on the distributions of soybean yield; these values can be used as indicators of production risk.

simulated fields, and the gradients in both precipitation and temperatures are consistent with observations (colored points in Figure 6). The 95% ensemble spread indicates significant uncertainty in precipitation amounts ($\sim 300\text{--}600$ mm), maximum temperature ($\sim 3.0\text{--}4.5$ °C), and minimum temperatures ($\sim 2.5\text{--}4.0$ °C).

4.5. Coupling With DSSAT

The utility of stochastic weather generators for decision-making is assessed, typically, by coupling them with process-based models (e.g., crop yield, hydrologic, and water resources) to generate ensembles of decision variables that will be of use in risk-based decision-making. For the sake of completeness, we offer a limited demonstration of this by integrating the weather ensembles from BayGEN with the Decision Support System for Agrotechnology Transfer (DSSAT; Jones et al., 2003), a widely used modeling system for simulating yields of many crops. DSSAT can be used to assist in the analysis of complex alternative decisions in agricultural management and adaptation. For a selected crop, the agronomic management (e.g., genotype, planting date, and fertilization rate) and soil type (values for many parameters that describe the soil of the agricultural plot of interest) are prescribed. Daily weather (i.e., precipitation, minimum and maximum temperatures, and solar radiation) are required inputs to DSSAT; optional inputs include dew point and wind speed. Outputs include crop yield (at the seasonal scale), total biomass, flowering and harvest date, soil water content for each layer, and root depth, among many other variables. DSSAT simulation models have been calibrated and tested in the Pampas for soybean, maize, and wheat (Apipattanavis et al., 2010; Bert et al., 2006, 2007; Ferreyra et al., 2001; Mercau et al., 2007; Meira et al., 1999; Podestá et al., 2009), which are commonly grown crops. Currently, DSSAT simulates crop development and growth for one point (i.e., plot); however, there are at present some tools that facilitate running a gridded version of DSSAT (Elliott et al., 2014; McNider et al., 2011).

BayGEN models precipitation and temperature but in its current form does not simulate additional weather variables (e.g., solar radiation, wind speed, and reference evapotranspiration). However, BayGEN can easily be modified to model any relevant additional variables. Here we estimated solar radiation after the simulation process using a modified Bristow-Campbell method (Bristow & Campbell, 1984). Bristow and Campbell (1984) suggest using the mean of minimum temperature for days t and $t + 1$ to help reduce the effect of large-scale hot or cold air masses moving through the domain. The modification to the Bristow-Campbell method that we used involves using minimum temperature for day t directly. We had these modifications implemented from prior research (e.g., Bert et al., 2006), thus readily available for this demonstration.

We simulated daily weather sequences for 2 years using 100 randomly selected parameter samples from the BayGEN posterior distributions. Using a calibrated DSSAT soybean simulation model, 100 soybean yields were computed for the austral spring-summer growing season (October–March) at Junin, an agriculturally productive region in the northeastern corner of the SAB. We also simulated 100 soybean yields using weather simulations from GLMGEN for comparison. Figure 7 shows the cdf of the 100 soybean yields from BayGEN and GLMGEN simulations. The moderate slope in the cdf of soybean yields from BayGEN indicates

Table 1

Climatological Monthly Means and Standard Deviations From the Observational Record (Denoted as Obs) for Precipitation, Maximum Temperature, and Minimum Temperature

	Precipitation			Maximum temperature			Minimum temperature		
	Obs	BayGEN	GLMGEN	Obs	BayGEN	GLMGEN	Obs	BayGEN	GLMGEN
Jan mean	4.1	4.1 (9.1)	4.0 (8.9)	30.0	29.9 (4.5)	30.4 (5.3)	16.3	16.1 (4.2)	15.9 (4.3)
Jan std. dev.	11.8	10.1 (20.4)	10.0 (20.2)	3.6	3.5 (2.2)	3.5 (2.5)	3.1	3.1 (2.1)	3.1 (2.1)
Feb mean	3.9	3.5 (8.5)	3.4 (8.4)	28.7	28.9 (4.7)	29.5 (5.5)	15.5	15.5 (4.4)	15.2 (4.3)
Feb std. dev.	11.4	8.8 (19.5)	8.7 (19.3)	3.7	3.5 (2.4)	3.5 (2.6)	3.4	3.2 (2.2)	3.1 (2.2)
Mar mean	4.8	4.0 (9.9)	3.9 (9.7)	26.5	26.4 (4.5)	26.9 (5.2)	13.4	13.3 (4.3)	13.0 (4.4)
Mar std. dev.	15.5	10.7 (24.2)	10.7 (23.9)	4.0	3.6 (2.4)	3.6 (2.7)	3.8	3.4 (2.3)	3.3 (2.3)
Apr mean	3.2	2.2 (6.2)	2.2 (6.1)	22.8	22.9 (4.5)	23.3 (5.1)	10.3	10.1 (4.7)	9.7 (4.8)
Apr std. dev.	10.0	6.4 (16.2)	6.4 (16.1)	4.0	3.6 (2.4)	3.5 (2.7)	4.3	3.6 (2.5)	3.6 (2.6)
May mean	1.8	1.4 (4.1)	1.4 (4.1)	19.1	19.3 (4.1)	19.7 (4.9)	7.2	6.8 (4.8)	6.3 (5.1)
May std. dev.	7.6	4.3 (11.6)	4.2 (11.5)	4.1	3.4 (2.3)	3.3 (2.5)	4.8	3.7 (2.5)	3.9 (2.7)
Jun mean	1.0	1.1 (3.5)	1.0 (3.4)	15.4	16.8 (4.2)	17.1 (4.8)	4.4	4.4 (4.8)	3.8 (5.2)
Jun std. dev.	4.0	3.4 (10.0)	3.4 (9.9)	3.7	3.3 (2.2)	3.1 (2.3)	4.9	3.6 (2.5)	3.8 (2.6)
Jul mean	1.1	1.1 (3.4)	1.0 (3.4)	15.1	16.0 (4.4)	16.3 (5.0)	3.8	3.3 (4.9)	2.8 (5.2)
Jul std. dev.	4.8	3.4 (9.9)	3.4 (9.9)	4.2	3.5 (2.2)	3.3 (2.4)	5.0	3.7 (2.4)	3.8 (2.6)
Aug mean	1.2	1.2 (3.7)	1.2 (3.6)	17.4	17.1 (4.7)	17.3 (5.2)	4.6	4.0 (4.9)	3.6 (5.2)
Aug std. dev.	4.7	3.8 (10.3)	3.7 (10.2)	4.5	3.8 (2.5)	3.5 (2.5)	4.7	3.7 (2.5)	3.9 (2.6)
Sep mean	1.8	1.7 (4.8)	1.6 (4.6)	19.7	19.7 (4.9)	20.0 (5.5)	6.7	6.2 (4.9)	6.0 (5.1)
Sep std. dev.	6.5	4.9 (12.5)	4.8 (12.5)	4.5	3.9 (2.6)	3.7 (2.8)	4.2	3.7 (2.5)	3.7 (2.6)
Oct mean	3.8	3.1 (7.7)	2.9 (7.5)	22.4	23.2 (4.9)	23.5 (5.5)	9.9	9.4 (4.7)	9.2 (4.7)
Oct std. dev.	10.4	8.4 (19.1)	8.2 (18.8)	4.2	4.0 (2.6)	3.8 (2.8)	3.8	3.7 (2.5)	3.5 (2.5)
Nov mean	3.7	3.6 (8.5)	3.5 (8.3)	25.9	26.7 (5.0)	27.1 (5.7)	12.6	12.6 (4.6)	12.5 (4.7)
Nov std. dev.	11.1	9.2 (19.7)	9.0 (19.5)	4.3	3.9 (2.6)	3.8 (2.9)	3.6	3.5 (2.4)	3.4 (2.5)
Dec mean	3.5	4.1 (9.0)	4.0 (8.8)	29.0	29.2 (4.8)	29.7 (5.7)	15.1	15.1 (4.4)	14.9 (4.6)
Dec std. dev.	10.0	10.0 (20.3)	9.9 (20.2)	4.1	3.8 (2.4)	3.8 (2.7)	3.4	3.3 (2.3)	3.4 (2.3)

Note. BayGEN ensemble medians are presented with 95% ensemble spread (97.5th percentile minus 2.5th percentile) included in parentheses; GLMGEN medians are presented with 95% ensemble spread included in parentheses.

higher risk of lower crop yields relative to GLMGEN. We select a heuristic “break-even” production quantity (~2,300 kg/ha) based on expert input to enable comparison between BayGEN and GLMGEN. The risks of not meeting break-even soybean production are 31% and 16% for BayGEN and GLMGEN, respectively. The higher risk can be attributed to the robust quantification of uncertainties in model parameters from BayGEN as reflected in higher standard deviation of simulations, especially precipitation during the wet season (Table 1). The decisions resulting from simulations exhibiting a higher risk estimate will likely be conservative, thus minimizing losses to farmers.

5. Summary and Discussion

We developed a novel GLM-based Bayesian hierarchical space-time stochastic weather generator, BayGEN. To our knowledge, this is one of the first attempts at a fully Bayesian weather generator. In this, the GLM-based space-time weather generator of Verdin, Rajagopalan, Kleiber, and Katz (2015) is embedded in a Bayesian hierarchical framework to obtain the posterior distributions of the model parameters. Sampling from these posterior distributions provides ensembles of weather sequences, thereby robustly reflecting the uncertainty in the model parameters.

We demonstrated the utility of BayGEN with application to daily weather generation in the Salado A subbasin of the Argentine Pampas. Daily weather simulations, each 53 years long (the same length as the observational record), were generated for each of the 4,000 posterior parameter samples, which produced a rich variety of weather sequences that incorporate parametric uncertainty. The BayGEN daily weather ensemble captured the basic statistics and domain extremes of the observations well, which is noteworthy as extremes were not explicitly modeled.

As is apparent in Table 1, BayGEN generates precipitation with greater variability than GLMGEN; however, minimum and maximum temperatures generated by BayGEN exhibit slightly less variability than those

from GLMGEN. The fact that temperature variability from BayGEN and GLMGEN is approximately equivalent is indicative of the fact that the main drivers of variability in the temperature processes are the random noise processes. It would follow that the noise processes overpower the parametric uncertainty for temperature simulations. However, the greater variability of daily precipitation in the BayGEN daily weather simulations was magnified when the weather simulations were used to drive a DSSAT soybean crop simulation model. The greater variability in DSSAT yields can be attributed to the greater variability in seasonal precipitation totals, which is the main driver of crop development in the DSSAT model. For comparison, weather simulations from GLMGEN were also used to drive the same DSSAT model. Break-even production risks from the BayGEN runs (31%) were nearly double that of the GLMGEN runs (16%), which is likely due to the increase in variability between the simulated weather series.

Extending BayGEN to include seasonal covariates to enable conditional weather simulation, informed with seasonal forecasts and multidecadal projections similar to the work of Verdin et al. (2018), would be beneficial for short- and long-term planning of agriculture and other resources. Additionally, teleconnections between the El Niño Southern Oscillation and the climate of a region may also be included to enable seasonal forecasting.

While effective in quantifying parameter uncertainty, BayGEN comes with a high computational cost, especially when applied to large domains or dense grids. Even in this application, spatial correlation of latent residuals in the precipitation occurrence probit regression is modeled during postprocessing using empirical correlation matrices for each month. Future work will investigate new and efficient ways to incorporate this step into the BayGEN workflow.

Acknowledgments

This research was funded by the National Science Foundation's Dynamics of Coupled Natural-Human Systems program (award 1211613). The authors would like to thank the Servicio Meteorológico Nacional (SMN) for providing the weather data. These data were extracted from the SMN database (<https://www.smn.gob.ar>), where it is operationally archived by an institution that has that mandate. The data cannot be uploaded to a repository by the authors, though readers may request the data through SMN. This work utilized the Janus supercomputer, which is supported by the National Science Foundation (award number CNS-0821794) and the University of Colorado Boulder. The Janus supercomputer is a joint effort of the University of Colorado Boulder, the University of Colorado Denver, and the National Center for Atmospheric Research. We are grateful to three anonymous reviewers for their insightful comments, which significantly improved the manuscript. Our special thanks to Jasper Vrugt for his detailed comments, which helped to enhance the overall presentation of the methodology and results in the manuscript.

References

- Aelion, C., Davis, H., Liu, Y., Lawson, A., & McDermott, S. (2009). Validation of Bayesian kriging of arsenic, chromium, lead, and mercury surface soil concentrations based on internode sampling. *Environmental Science & Technology*, *43*(12), 4432–4438. <https://doi.org/10.1021/es803322w>
- Apipattanavis, S., Bert, F., Podestá, G., & Rajagopalan, B. (2010). Linking weather generators and crop models for assessment of climate forecast outcomes. *Agricultural and Forest Meteorology*, *150*(2), 166–174. <https://doi.org/10.1016/j.agrformet.2009.09.012>
- Asong, Z. E., Khaliq, M. N., & Wheeler, H. S. (2016). Multisite multivariate modeling of daily precipitation and temperature in the Canadian prairie provinces using generalized linear models. *Climate Dynamics*, *47*(9–10), 2901–2921. <https://doi.org/10.1007/s00382-016-3004-z>
- Baigorria, G. A., & Jones, J. W. (2010). GiST: A stochastic model for generating spatially and temporally correlated daily rainfall data. *Journal of Climate*, *23*(22), 5990–6008. <https://doi.org/10.1175/2010JCLI3537.1>
- Beersma, J., & Buishand, T. (2003). Multi-site simulation of daily precipitation and temperature conditional on the atmospheric circulation. *Climate Research*, *25*, 121–133. <https://doi.org/10.3354/cr025121>
- Berger, T. (2000). Agent-based spatial models applied to agriculture: A simulation tool for technology diffusion, resource use changes, and policy analysis. *Agricultural Economics*, *25*(2–3), 245–260.
- Berger, T., Schreinemachers, P., & Woelcke, J. (2006). Multi-agent simulation for the targeting of development policies in less-favored areas. *Agricultural Systems*, *88*(1), 28–43. <https://doi.org/10.1016/j.agsy.2005.06.002>
- Bert, F. E., Laciána, C. E., Podestá, G. P., Satorre, E. H., & Menéndez, A. N. (2007). Sensitivity of CERES-Maize simulated yields to uncertainty in soil properties and daily solar radiation. *Agricultural Systems*, *94*(2), 141–150. <https://doi.org/10.1016/j.agsy.2006.08.003>
- Bert, F. E., Rovere, S. L., Macal, C. M., North, M. J., & Podestá, G. P. (2014). Lessons from a comprehensive validation of an agent based-model: The experience of the Pampas Model of Argentinean agricultural systems. *Ecological Modelling*, *273*, 284–298. <https://doi.org/10.1016/j.ecolmodel.2013.11.024>
- Bert, F. E., Satorre, E. H., Toranzo, F. R., & Podestá, G. P. (2006). Climatic information and decision-making in maize crop production systems of the Argentinean Pampas. *Agricultural Systems*, *88*(2–3), 180–204. <https://doi.org/10.1016/j.agsy.2005.03.007>
- Bristow, K. L., & Campbell, G. S. (1984). On the relationship between incoming solar radiation and daily maximum and minimum temperature. *Agricultural and Forest Meteorology*, *31*(2), 159–166. [https://doi.org/10.1016/0168-1923\(84\)90017-0](https://doi.org/10.1016/0168-1923(84)90017-0)
- Buishand, T., & Brandsma, T. (2001). Multisite simulation of daily precipitation and temperature in the Rhine Basin by nearest-neighbor resampling. *Water Resources Research*, *37*(11), 2761–2776. <https://doi.org/10.1029/2001WR000291>
- Cano, R., Sordo, C., & Gutiérrez, J. M. (2004). Applications of Bayesian networks in meteorology. In J. A. Gámez, S. Moral, & A. Salmerón (Eds.), *Advances in Bayesian networks, studies in fuzziness and Soft Computing* (Vol. 146, pp. 309–328). Berlin Heidelberg: Springer.
- Carpenter, B. (2015). *Stan: A probabilistic programming language*. Journal of Statistical Software.
- Chandler, R. E. (2005). On the use of generalized linear models for interpreting climate variability. *Environmetrics*, *16*(7), 699–715. <https://doi.org/10.1002/env.731>
- Chilés, J. P., & Delfiner, P. (1999). *Geostatistics: Modeling Spatial Uncertainty*. New York: Wiley. <https://doi.org/10.1002/9780470316993>
- Cooley, D., Nychka, D., & Naveau, P. (2007). Bayesian spatial modeling of extreme precipitation return levels. *Journal of the American Statistical Association*, *102*(479), 824–840. <https://doi.org/10.1198/016214506000000780>
- Cooley, D., & Sain, S. (2010). Spatial hierarchical modeling of precipitation extremes from a regional climate model. *Journal of Agricultural, Biological, and Environmental Statistics*, *15*(3), 381–402. <https://doi.org/10.1007/s13253-010-0023-9>
- Cui, H., Stein, A., & Myers, D. (1995). Extension of spatial information, Bayesian kriging and updating of prior variogram parameters. *Environmetrics*, *6*, 373–384.
- Duan, Q., Ajami, N. K., Gao, X., & Sorooshian, S. (2007). Multi-model ensemble hydrologic prediction using Bayesian model averaging. *Advances in Water Resources*, *30*(5), 1371–1386. <https://doi.org/10.1016/j.advwatres.2006.11.014>

- Elliott, J., Kelly, D., Chryssanthacopoulos, J., Glotter, M., Jhunjhunwala, K., Best, N., et al. (2014). The parallel system for integrating impact models and sectors (psims). *Environmental Modelling & Software*, *62*, 509–516. <https://doi.org/10.1016/j.envsoft.2014.04.008>
- Ferreira, R. A., Podestá, G. P., Messina, C. D., Letson, D., Dardanelli, J., Guevara, E., & Meira, S. (2001). A linked-modeling framework to estimate maize production risk associated with ENSO-related climate variability in Argentina. *Agricultural and Forest Meteorology*, *107*(3), 177–192. [https://doi.org/10.1016/S0168-1923\(00\)00240-9](https://doi.org/10.1016/S0168-1923(00)00240-9)
- Freeman, T., Nolan, J., & Schoney, R. (2009). An agent-based simulation model of structural change in Canadian prairie agriculture, 1960–2000. *Canadian Journal of Agricultural Economics-Revue Canadienne D Agroeconomie*, *57*(4), 537–554. <https://doi.org/10.1111/j.1744-7976.2009.01169.x>
- Furrer, E. M., & Katz, R. W. (2008). Improving the simulation of extreme precipitation events by stochastic weather generators. *Water Resources Research*, *44*, W12439. <https://doi.org/10.1029/2008WR007316>
- Gelman, A., Carlin, J. B., Stern, H. S., & Rubin, D. B. (2004). *Bayesian data analysis*. Boca Raton: Chapman & Hall/CRC.
- Handcock, M., & Stein, M. (1993). A Bayesian analysis of kriging. *Technometrics*, *35*(4), 403–410. <https://doi.org/10.1080/00401706.1993.10485354>
- Happe, K., Balmann, A., Kellermann, K., & Sahrbacher, C. (2008). Does structure matter? The impact of switching the agricultural policy regime on farm structures. *Journal of Economic Behavior and Organization*, *67*(2), 431–444. <https://doi.org/10.1016/j.jebo.2006.10.009>
- Hashmi, M., Shamseldin, A., & Melville, B. (2009). Statistical downscaling of precipitation: State-of-the-art and application of Bayesian multi-model approach for uncertainty assessment. *Hydrology and Earth System Sciences Discussions*, *6*(5), 6535–6579. <https://doi.org/10.5194/hessd-6-6535-2009>
- Hauser, T., & Demirov, E. (2013). Development of a stochastic weather generator for the sub-polar North Atlantic. *Stochastic Environmental Research and Risk Assessment*, *27*(7), 1533–1551. <https://doi.org/10.1007/s00477-013-0688-z>
- Hoffman, M., & Gelman, A. (2014). The No-U-Turn sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research*, *15*, 1593–1623.
- Jin, B., Wu, Y., Miao, B., Wang, X., & Guo, P. (2014). Bayesian spatiotemporal modeling for blending in situ observations with satellite precipitation estimates. *Journal of Geophysical Research: Atmospheres*, *119*, 1806–1819. <https://doi.org/10.1002/2013JD019648>
- Jones, J., Hoogenboom, G., Porter, C., Boote, K., Batchelor, W., Hunt, L., et al. (2003). The DSSAT cropping system model. *European Journal of Agronomy*, *18*(3–4), 235–265. [https://doi.org/10.1016/S1161-0301\(02\)00107-7](https://doi.org/10.1016/S1161-0301(02)00107-7)
- Katz, R. W. (1977). Precipitation as a chain-dependent process. *Journal of Applied Meteorology*, *16*(7), 671–676. [https://doi.org/10.1175/1520-0450\(1977\)016<0671:PAACDP>2.0.CO;2](https://doi.org/10.1175/1520-0450(1977)016<0671:PAACDP>2.0.CO;2)
- Khalili, M., Brissette, F., & Leconte, R. (2009). Stochastic multi-site generation of daily weather data. *Stochastic Environmental Research and Risk Assessment*, *23*(6), 837–849. <https://doi.org/10.1007/s00477-008-0275-x>
- Kleiber, W., Katz, R. W., & Rajagopalan, B. (2012). Daily spatiotemporal precipitation simulation using latent and transformed Gaussian processes. *Water Resources Research*, *48*, W01523. <https://doi.org/10.1029/2011WR011105>
- Kleiber, W., Katz, R. W., & Rajagopalan, B. (2013). Daily minimum and maximum temperature simulation over complex terrain. *Ann. Appl. Stat.*, *7*(1), 588–612. <https://doi.org/10.1214/12-AOAS602>
- Kucukelbir, A., Tran, D., Ranganath, R., Gelman, A., & Blei, D. M. (2017). Automatic differentiation variational inference. *The Journal of Machine Learning Research*, *18*(1), 430–474.
- Leamer, E. E. (1978). *Specification searches*. New York: Wiley.
- Lima, C. H. R., & Lall, U. (2009). Hierarchical Bayesian modeling of multisite daily rainfall occurrence: Rainy season onset, peak, and end. *Water Resources Research*, *45*, W07422. <https://doi.org/10.1029/2008WR007485>
- McNider, R. T., Christy, J. R., Moss, D., Doty, K., & Handyside, C. (2011). A real-time gridded crop model for assessing spatial drought stress on crops in the southeastern United States. *Journal of Applied Meteorology and Climatology*, *50*(7), 1459–1475. <https://doi.org/10.1175/2011JAMC2476.1>
- Mehrotra, R., Srikanthan, R., & Sharma, A. (2006). A comparison of three stochastic multi-site precipitation occurrence generators. *Journal of Hydrology*, *331*(1–2), 280–292. <https://doi.org/10.1016/j.jhydrol.2006.05.016>
- Meira, S., Gaigori, H., Guevara, E., & Maturano, M. (1999). Calibration of soybean cultivars for the SOYGRO model in two environments of Argentina, in Global Soy Forum (pp. 4–7).
- Mercau, J., Dardanelli, J., Collino, D., Andriani, J., Irigoyen, A., & Satorre, E. (2007). Predicting on-farm soybean yields in the pampas using CROPGRO-soybean. *Field Crops Research*, *100*(2–3), 200–209. <https://doi.org/10.1016/j.fcr.2006.07.006>
- Montgomery, J. M., & Nyhan, B. (2010). Bayesian model averaging: Theoretical developments and practical applications. *Political Analysis*, *18*(02), 245–270. <https://doi.org/10.1093/pan/mpq001>
- Nesterov, Y. (2009). Primal-dual subgradient methods for convex problems. *Mathematical Programming*, *120*(1), 221–259. <https://doi.org/10.1007/s10107-007-0149-x>
- Omre, H. (1987). Bayesian kriging—Merging observations and qualified guesses in kriging. *Mathematical Geology*, *19*(1), 25–39. <https://doi.org/10.1007/BF01275432>
- Pezzulli, S., Frederic, P., Majithia, S., Sabbagh, S., Black, E., Sutton, R., & Stephenson, D. (2006). The seasonal forecast of electricity demand: A hierarchical Bayesian model with climatological weather generator. *Applied Stochastic Models in Business and Industry*, *22*(2), 113–125. <https://doi.org/10.1002/asmb.622>
- Podestá, G., Bert, F., Rajagopalan, B., Apipattanas, S., Laciana, C., Weber, E., et al. (2009). Decadal climate variability in the Argentine Pampas: Regional impacts of plausible climate scenarios on agricultural systems. *Climate Research*, *40*, 199–210. <https://doi.org/10.3354/cr00807>
- Qian, B., Corte-Real, J., & Xu, H. (2002). Multisite stochastic weather models for impact studies. *International Journal of Climatology*, *22*(11), 1377–1397. <https://doi.org/10.1002/joc.808>
- R Core Team (2014). *R: A language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing.
- Raftery, A., Gneiting, T., Balabdaoui, F., & Polakowski, M. (2005). Using Bayesian model averaging to calibrate forecast ensembles. *American Meteorological Society*, *133*(5), 1155–1174.
- Rajagopalan, B., & Lall, U. (1999). A *k*-nearest-neighbor simulator for daily precipitation. *Water Resources Research*, *35*(10), 3089–3101. <https://doi.org/10.1029/1999WR900028>
- Reich, B., & Shaby, B. (2012). A hierarchical max-stable spatial model for extreme precipitation. *Annals of Applied Statistics*, *6*(4), 1430–1451. <https://doi.org/10.1214/12-AOAS591>
- Richardson, C. W. (1981). Stochastic simulation of daily precipitation, temperature, and solar radiation. *Water Resources Research*, *17*(1), 182–190. <https://doi.org/10.1029/WR017001p00182>

- Sahu, S., & Mardia, K. (2005). A Bayesian kriged Kalman model for short-term forecasting of air pollution levels. *Journal of the Royal Statistical Society: Series C: Applied Statistics*, *54*(1), 223–244. <https://doi.org/10.1111/j.1467-9876.2005.00480.x>
- Schreinemachers, P., & Berger, T. (2011). An agent-based simulation model of human-environment interactions in agricultural systems. *Environmental Modelling and Software*, *26*(7), 845–859. <https://doi.org/10.1016/j.envsoft.2011.02.004>
- Sharif, M., & Burn, D. H. (2007). Improved k -nearest neighbor weather generating model. *Journal of Hydrologic Engineering*, *12*(1), 42–51. [https://doi.org/10.1061/\(ASCE\)1084-0699\(2007\)12:1\(42\)](https://doi.org/10.1061/(ASCE)1084-0699(2007)12:1(42))
- Sloughter, J., Raftery, A., Gneiting, T., & Fraley, C. (2007). Probabilistic quantitative precipitation forecasting using Bayesian model averaging. *Monthly Weather Review*, *135*(9), 3209–3220. <https://doi.org/10.1175/MWR3441.1>
- Smith, R., & Robinson, P. (1997). A Bayesian approach to the modeling of spatial-temporal precipitation data. In C. Gatsonis, J. S. Hodges, R. E. Kass, R. McCulloch, P. Rossi, & N. D. Singpurwalla (Eds.), *Case studies in Bayesian statistics, Lecture Notes in Statistics* (Vol. 121, pp. 237–269). New York: Springer.
- Stan Development Team (2015a). Stan modeling language user's guide and reference manual, version 2.8.0. Retrieved from <http://mc-stan.org>
- Stan Development Team (2015b). Stan: A C++ library for probability and sampling, version 2.8.0. Retrieved from <http://mc-stan.org>
- Stern, R. D., & Coe, R. (1984). A model fitting analysis of daily rainfall data. *Journal of the Royal Statistical Society*, *147*(1), 1–34. <https://doi.org/10.2307/2981736>
- Tebaldi, C., Mearns, L., Nychka, D., & Smith, R. (2004). Regional probabilities of precipitation change: A Bayesian analysis of multimodel simulations. *Geophysical Research Letters*, *31*, L24213. <https://doi.org/10.1029/2004GL021276>
- Tebaldi, C., & Sanso, B. (2009). Joint projections of temperature and precipitation change from multiple climate models: A hierarchical Bayesian approach. *Journal of the Royal Statistical Society, Series A*, *172*(1), 83–106.
- Trujillo-Barreto, N. J., Aubert-Vázquez, E., & Valdés-Sossa, P. A. (2004). Bayesian model averaging in eeg/meg imaging. *NeuroImage*, *21*(4), 1300–1319. <https://doi.org/10.1016/j.neuroimage.2003.11.008>
- Verdin, A., Rajagopalan, B., Kleiber, W., & Funk, C. (2015). A Bayesian kriging approach for blending satellite and ground precipitation observations. *Water Resources Research*, *6*, 446–414. [https://doi.org/10.1016/0022-1694\(68\)90080-2](https://doi.org/10.1016/0022-1694(68)90080-2)
- Verdin, A., Rajagopalan, B., Kleiber, W., & Katz, R. W. (2015). Coupled stochastic weather generation using spatial and generalized linear models. *Stochastic Environmental Research and Risk Assessment*, *29*(2), 347–356. <https://doi.org/10.1007/s00477-014-0911-6>
- Verdin, A., Rajagopalan, B., Kleiber, W., Podestà, G., & Bert, F. (2018). A conditional stochastic weather generator for seasonal to multi-decadal simulations. *Journal of Hydrology*, *556*(1), 835–846. <https://doi.org/10.1016/j.jhydrol.2015.12.036>
- Viallefont, V., Raftery, A. E., & Richardson, S. (2001). Variable selection and Bayesian model averaging in case-control studies. *Statistics in Medicine*, *20*(21), 3215–3230. <https://doi.org/10.1002/sim.976>
- Volinsky, C. T., Madigan, D., Raftery, A. E., & Kronmal, R. A. (1997). Bayesian model averaging in proportional hazard models: Assessing the risk of a stroke. *Journal of the Royal Statistical Society: Series C: Applied Statistics*, *46*(4), 433–448. <https://doi.org/10.1111/1467-9876.00082>
- Vrugt, J. A., & Robinson, B. A. (2007). Treatment of uncertainty using ensemble methods: Comparison of sequential data assimilation and Bayesian model averaging. *Water Resources Research*, *43*, W01411. <https://doi.org/10.1029/2005WR004838>
- Wilks, D. S. (1998). Multisite generalization of a daily stochastic precipitation generation model. *Journal of Hydrology*, *210*(1–4), 178–191. [https://doi.org/10.1016/S0022-1694\(98\)00186-3](https://doi.org/10.1016/S0022-1694(98)00186-3)
- Wilks, D. S. (2008). High-resolution spatial interpolation of weather generator parameters using local weighted regressions. *Agricultural and Forest Meteorology*, *148*(1), 111–120. <https://doi.org/10.1016/j.agrformet.2007.09.005>
- Wilks, D. S. (2009). A gridded multisite weather generator and synchronization to observed weather data. *Water Resources Research*, *45*, W10419. <https://doi.org/10.1029/2009WR007902>
- Wilks, D. S., & Wilby, R. L. (1999). The weather generation game: A review of stochastic weather models. *Progress in Physical Geography*, *23*(3), 329–357. <https://doi.org/10.1177/030913339902300302>
- Wintle, B. A., McCarthy, M. A., Volinsky, C. T., & Kavanagh, R. P. (2003). The use of Bayesian model averaging to better represent uncertainty in ecological models. *Conservation Biology*, *17*(6), 1579–1590. <https://doi.org/10.1111/j.1523-1739.2003.00614.x>
- Wright, J. H. (2008). Bayesian model averaging and exchange rate forecasts. *Journal of Econometrics*, *146*(2), 329–341. <https://doi.org/10.1016/j.jeconom.2008.08.012>
- Yang, C., Chandler, R. E., Isham, V. S., & Wheeler, H. S. (2005). Spatial-temporal rainfall simulation using generalized linear models. *Water Resources Research*, *41*, W11415. <https://doi.org/10.1029/2004WR003739>
- Yates, D., Gangopadhyay, S., Rajagopalan, B., & Strzepek, K. (2003). A technique for generating regional climate scenarios using a nearest-neighbor algorithm. *Water Resources Research*, *39*(7), 1199. <https://doi.org/10.1029/2002WR001769>