

A DYNAMIC NONSTATIONARY SPATIO-TEMPORAL MODEL FOR SHORT TERM PREDICTION OF PRECIPITATION

BY FABIO SIGRIST, HANS R. KÜNSCH AND WERNER A. STAHEL

Seminar for Statistics, ETH Zurich

Precipitation is a complex physical process that varies in space and time. Predictions and interpolations at unobserved times and/or locations help to solve important problems in many areas. In this paper, we present a hierarchical Bayesian model for spatio-temporal data and apply it to obtain short term predictions of rainfall. The model incorporates physical knowledge about the underlying processes that determine rainfall, such as advection, diffusion and convection. It is based on a temporal autoregressive convolution with spatially colored and temporally white innovations. By linking the advection parameter of the convolution kernel to an external wind vector, the model is temporally nonstationary. Further, it allows for nonseparable and anisotropic covariance structures. With the help of the Voronoi tessellation, we construct a natural parametrization, that is, space as well as time resolution consistent, for data lying on irregular grid points. In the application, the statistical model combines forecasts of three other meteorological variables obtained from a numerical weather prediction model with past precipitation observations. The model is then used to predict three-hourly precipitation over 24 hours. It performs better than a separable, stationary and isotropic version, and it performs comparably to a deterministic numerical weather prediction model for precipitation and has the advantage that it quantifies prediction uncertainty.

1. Introduction. Precipitation is a very complex phenomenon that varies in space and time, and there are many efforts to model it. Predictions and interpolations at unobserved times and/or locations obtained from such models help to solve important problems in areas such as agriculture, climate science, ecology and hydrology. Stochastic models have the great advantage of providing not only point estimates, but also quantitative measures of uncertainty. They can be used, for instance, as stochastic generators [Wilks (1998), Makhnin and McAllister (2009)] to provide realistic inputs to flooding, runoff and crop growth models. Moreover, they can be applied as components within general circulation models used in climate change studies [Fowler et al. (2005)], or for postprocessing precipitation forecasts [Sloughter et al. (2007)].

1.1. *Distributions for precipitation.* A characteristic feature of precipitation is that its distribution consists of a discrete component, indicating occurrence of

Received December 2011; revised March 2012.

Key words and phrases. Rainfall modeling, space–time model, hierarchical Bayesian model, Markov chain Monte Carlo (MCMC), censoring, Gaussian random field.

precipitation, and a continuous one, determining the amount of precipitation. As a consequence, there are two basic statistical modeling approaches. The continuous and the discrete part are either modeled separately [Coe and Stern (1982), Wilks (1999)] or together [Bell (1987), Wilks (1990), Bardossy and Plate (1992), Hutchinson (1995), Sansó and Guenni (2004)]. Typically, in the second approach, the distribution of the rainfall amounts and the probability of rainfall are determined together using what is called a censored distribution. Originally, this idea goes back to Tobin (1958) who analyzed household expenditure on durable goods. For modeling precipitation, Stidd (1973) took up this idea and modified it by including a power-transformation for the nonzero part so that the model can account for skewness.

1.2. *Correlations in space and time.* For modeling processes that involve dependence over space and time, there are two basic approaches [see, e.g., Cressie and Wikle (2011)]: one which models the space–time covariance structure without distinguishing between the time and space dimensions, and a dynamic one which takes the natural ordering in the time dimension into account.

The first approach usually follows the traditional geostatistical paradigm of assuming a parametric covariance function [for an introduction into geostatistics, see, e.g., Cressie (1993) or Gelfand et al. (2010)]. Several parametric families specifying explicitly the joint space–time covariance structure have been proposed [Jones and Zhang (1997), Cressie and Huang (1999), Gneiting (2002), Ma (2003), Stein (2005), Paciorek and Schervish (2006)]. Interpretability and, especially, computational complexity are challenges when working with parametric space–time covariance functions.

There is, however, a fundamental difference between the spatial and the temporal dimensions. Whereas there is an order in the time domain, there exists no obvious order for space. It is therefore natural to assume a dynamic temporal evolution combined with a spatially correlated error term [Sølna and Switzer (1996), Wikle and Cressie (1999), Huang and Hsu (2004), Xu, Wikle and Fox (2005), Gelfand, Banerjee and Gamerman (2005)]. As Wikle and Hooten (2010) state, the dynamic approach can be used to construct realistic space–time dependency structures based on physical knowledge. Further, the temporal Markovian structure offers computational benefits.

1.3. *Models for precipitation.* Isham and Cox (1994) state that there are three broad types of mathematical models of rainfall: deterministic meteorological models [Mason (1986)], intermediate stochastic models [Le Cam (1961), Cox and Isham (1988), Waymire, Gupta and Rodriguez-Iturbe (1984)], and empirical statistical models. Meteorological models represent as realistically as possible the physical processes involved. As noted by Kyriakidis and Journel (1999), deterministic models typically require a large number of input parameters that are difficult to determine, whereas stochastic models are usually based on a small number of

parameters. Nevertheless, statistical models can also incorporate knowledge about physical processes. Parametrizations can be chosen based on physical knowledge and covariates reflecting information about the physical processes can be included.

In the following, we briefly review statistical models for precipitation. For modeling daily precipitation at a single measuring site, Stern and Coe (1984) use a nonstationary second-order Markov chain to describe precipitation occurrence and a gamma distribution to describe rainfall amounts. Hughes and Guttorp (1994) and Hughes, Guttorp and Charles (1999) model precipitation occurrence using a nonhomogeneous hidden Markov model. With the help of an unobserved weather state they link large scale atmospheric circulation patterns with the local precipitation process. Bellone, Hughes and Guttorp (2000) and Charles, Bates and Hughes (1999) both extend this approach by also modeling precipitation amounts. The former propose to use gamma distributions, whereas the latter use empirical distribution functions. Ailliot, Thompson and Thomson (2009) present a hidden Markov model in combination with the transformed and censored Gaussian distribution approach used in Bardossy and Plate (1992). Also building on the same censoring idea, Sansó and Guenni (1999a) model precipitation occurrence and amount of precipitation using a transformed multivariate Gaussian model with a spatial correlation structure. Further works on statistical precipitation modeling include Sansó and Guenni (1999b, 2000), Brown et al. (2001), Stehlik and Bardossy (2002), Allcroft and Glasbey (2003), Slughter et al. (2007), Berrocal, Raftery and Gneiting (2008) and Fuentes, Reich and Lee (2008).

1.4. *Outline.* The model presented in the following is a hierarchical Bayesian model for spatio-temporal data. At the data stage, we opt for a modeling approach that determines the discrete and the continuous parts of the precipitation distribution together. This is done by assuming the existence of a latent Gaussian variable which can be interpreted as a precipitation potential. The mean of the Gaussian variable is related to covariates through a regression term. The advantages of this one-part modeling strategy are twofold: the model contains less parameters and it can deal with the so-called spatial (and temporal) intermittence effect [Bardossy and Plate (1992)] which suggests smooth transitions between wet and dry areas. This means that at the edge of a dry area the amount of rainfall should be low. Wilks (1998) notes that, indeed, lower rainfall intensity is observed when more neighboring stations are dry. This feature also reflects the idea that if a model determines a low probability of rainfall for a given situation, it should also give a small expected value for its amount conditional on this event, and vice versa. However, we note that there is no consensus in the literature whether the two parts of precipitation should be modeled together or separately.

At the process level, we use a dynamic model for accounting for spatio-temporal variation. The model explicitly incorporates knowledge about the underlying physical processes that determine rainfall, such as advection, diffusion and convection. Approximating an integrodifference equation, we obtain a temporally autoregressive convolution with spatially colored and temporally white innovations.

The model is nonstationary, anisotropic, and it allows for nonseparable covariance structures, that is, covariance structures where spatial and temporal variation interact. While our approach builds on existing models, it includes several novel features. With the help of the Voronoi tessellation, a natural parametrization for data lying on an irregular grid is obtained. The parametrization based on this tessellation is space as well as time resolution consistent, physically realistic and allows for modeling irregularly spaced data in a natural way. To our knowledge, the use of the Voronoi tessellation for spatio-temporal data on an irregular grid is new. By linking the advection parameter of the kernel to an external wind vector, the model is temporally nonstationary.

The model is applied to predict three-hourly precipitation. The prediction model is based on three forecasted meteorological variables obtained from an NWP model as well as past rainfall observations. We compare predictions from the statistical model with the precipitation forecasts obtained from the NWP.

The remainder is organized as follows. In Section 2 the model specifications are presented. In Section 3 it is shown how the model can be fitted to data using a Markov chain Monte Carlo (MCMC) algorithm and how predictions can be obtained. Next, in Section 4 the model is applied to obtain short term predictions of three-hourly rainfall. Conclusions are given in Section 5.

2. The model. It is assumed that the rainfall $Y_t(\mathbf{s})$ at time t on site $\mathbf{s} = (x, y)' \in \mathbb{R}^2$ depends on a latent normal variable $W_t(\mathbf{s})$ through

$$(1) \quad \begin{aligned} Y_t(\mathbf{s}) &= 0 && \text{if } W_t(\mathbf{s}) \leq 0 \\ &= W_t(\mathbf{s})^\lambda && \text{if } W_t(\mathbf{s}) > 0, \end{aligned}$$

where $\lambda > 0$. A power transformation is needed since precipitation amounts are more skewed than a truncated normal distribution and since the scatter of the precipitation amounts increases with the average amount. The latent variable $W_t(\mathbf{s})$ can be interpreted as a precipitation potential.

This latent variable $W_t(\mathbf{s})$ is modeled as a Gaussian process that is specified as

$$(2) \quad W_t(\mathbf{s}) = \mathbf{x}_t(\mathbf{s})^T \boldsymbol{\beta} + \xi_t(\mathbf{s}) + \nu_t(\mathbf{s}),$$

where $\boldsymbol{\beta} \in \mathbb{R}^k$, and $\nu_t(\mathbf{s}) \sim N(0, \tau^2)$, $\tau^2 > 0$, are i.i.d. The mean $\mathbf{x}_t(\mathbf{s})^T \boldsymbol{\beta}$ of $W_t(\mathbf{s})$ is assumed to depend linearly on regressors $\mathbf{x}_t(\mathbf{s}) \in \mathbb{R}^k$. For notational convenience, we split the terms specifying the covariance function into a structured part $\xi_t(\mathbf{s})$ and an unstructured “nugget” $\nu_t(\mathbf{s})$. The term $\xi_t(\mathbf{s})$ is a zero-mean Gaussian process that accounts for structured variation in time and space. It is specified below in Section 2.1. The nugget $\nu_t(\mathbf{s})$ models microscale variability and measurement errors. Since, typically, the resolution of the data does not allow for distinguishing between microscale variability and measurement errors, we model these two sources of variation together. Note that the covariates $\mathbf{x}_t(\mathbf{s})$ will usually be time and location dependent. In addition to weather characteristics, Fourier harmonics can be included to account for seasonality, and functions of coordinates can account for smooth effects in space.

2.1. *The convolution autoregressive model.* We follow the dynamic approach and define an explicit time evolution through the following integrodifference equation (IDE):

$$(3) \quad \xi_t(\mathbf{s}) = \int_{\mathbb{R}^2} h_{\boldsymbol{\vartheta}}(\mathbf{s} - \mathbf{s}') \xi_{t-1}(\mathbf{s}') d\mathbf{s}' + \varepsilon_t(\mathbf{s}), \quad \mathbf{s} \in \mathbb{R}^2,$$

where $\varepsilon_t(\mathbf{s})$ is a Gaussian innovation that is white in time and colored in space, and $h_{\boldsymbol{\vartheta}}$ is a Gaussian kernel,

$$(4) \quad h_{\boldsymbol{\vartheta}}(\mathbf{s} - \mathbf{s}') = \phi \exp(-(\mathbf{s} - \mathbf{s}' - \boldsymbol{\mu}_t)^T \boldsymbol{\Sigma}^{-1}(\mathbf{s} - \mathbf{s}' - \boldsymbol{\mu}_t)),$$

where the parameter vector $\boldsymbol{\vartheta}$ combines ϕ and the elements of $\boldsymbol{\mu}_t$ and $\boldsymbol{\Sigma}^{-1}$. Note that $\boldsymbol{\mu}_t$ shifts the kernel and $\boldsymbol{\Sigma}^{-1}$ determines the range and the degree of anisotropy. The parameter ϕ controls the amount of temporal correlation. More details on the interpretation of the model and specific choices of the parameters $\boldsymbol{\mu}_t$ and $\boldsymbol{\Sigma}$ are discussed below in Section 2.2. An illustration of this kernel can be found in the application in Section 4.2.

In the following, we assume that we have N measurement locations \mathbf{s}_i , $i = 1, \dots, N$, where measurements are made at times $t = 1, \dots, T$. Instead of working with a fine spatial grid with many missing observations, we formulate an approximate model for the values at the stations only, $\boldsymbol{\xi}_t = (\xi_t(\mathbf{s}_1), \dots, \xi_t(\mathbf{s}_N))'$. Discretizing the integral in (3), we obtain

$$(5) \quad \begin{aligned} \int_{\mathbb{R}^2} h_{\boldsymbol{\vartheta}}(\mathbf{s}_i - \mathbf{s}') \xi_{t-1}(\mathbf{s}') d\mathbf{s}' &\approx \int_A h_{\boldsymbol{\vartheta}}(\mathbf{s}_i - \mathbf{s}') \xi_{t-1}(\mathbf{s}') d\mathbf{s}' \\ &\approx \sum_{j=1}^N h_{\boldsymbol{\vartheta}}(\mathbf{s}_i - \mathbf{s}_j) \xi_{t-1}(\mathbf{s}_j) |A_j|. \end{aligned}$$

Here $A \subset \mathbb{R}^2$ is an area which contains the convex hull of all stations, the sets A_i , $i = 1, \dots, N$, form a tessellation of A with $\mathbf{s}_i \in A_i$ and $|A_j|$ denotes the area of cell A_j .

Our model can then be written as the vector autoregression

$$(6) \quad \boldsymbol{\xi}_t = \phi \mathbf{G}_t \boldsymbol{\xi}_{t-1} + \boldsymbol{\varepsilon}_t, \quad \mathbf{G}_t \in \mathbb{R}^{N \times N},$$

where

$$(7) \quad (\mathbf{G}_t)_{ij} = \exp(-(\mathbf{s}_i - \mathbf{s}'_j - \boldsymbol{\mu}_t)^T \boldsymbol{\Sigma}^{-1}(\mathbf{s}_i - \mathbf{s}'_j - \boldsymbol{\mu}_t)) \cdot |A_j|,$$

and where $\boldsymbol{\varepsilon}_t = (\varepsilon_t(\mathbf{s}_1), \dots, \varepsilon_t(\mathbf{s}_N))'$.

Note that this process does not exhibit explosive growth if the largest eigenvalue of $\phi \mathbf{G}_t$ is smaller than one. To ensure this, we check in our application that the largest eigenvalue is smaller than one for the parameters at the posterior modes.

If the \mathbf{s}_i 's form a regular grid, a tessellation is straightforward. Otherwise, we propose to use the Voronoi tessellation [Voronoi (1908)] which decomposes the

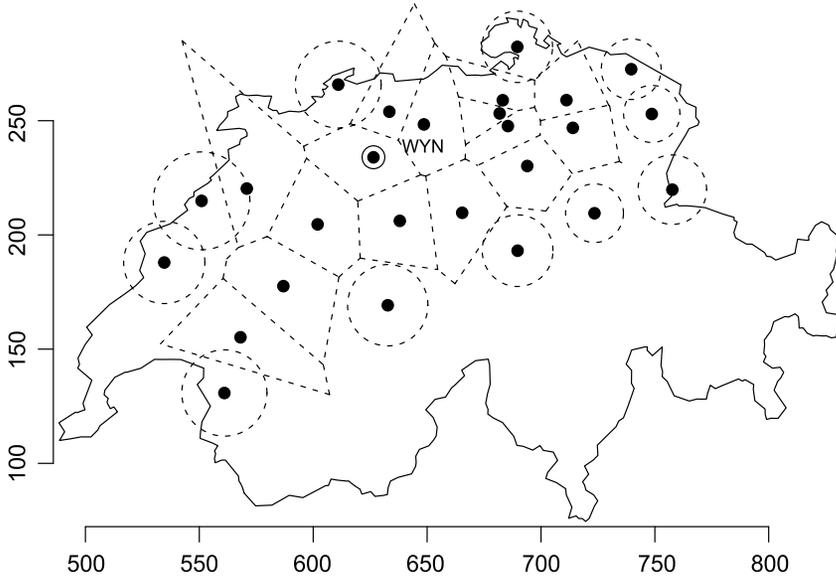


FIG. 1. Locations of stations. Both axes are in km using the Swiss coordinate system (CH1903). The lines illustrate the Voronoi tessellation. Cells with unbounded area have been replaced by circles whose area is determined as described in the text.

space. Specifically, each site s_i has a corresponding Voronoi cell consisting of all points closer to s_i than to any other site s_j , $j \neq i$ [see, e.g., Okabe et al. (2000) for more details]. Stations on the boundary of the convex hull have cells with infinite area. For these stations, we define $|A_i|$ as described in the following. We first calculate the Voronoi tessellation of \mathbb{R}^2 . We then replace unbounded cells by cells whose area is the average area of the neighboring bounded cells. In Figure 1, the Voronoi tessellation for the Swiss stations used in the application below is shown as an example. Concerning the stations on the boundary, the circles represent the surface area $|A_i|$.

As mentioned before, the ϵ_t 's are assumed to be independent over time and colored in space. More precisely, we assume a stationary, isotropic Gaussian random field

$$(8) \quad \epsilon_t \sim N(0, \sigma^2 \mathbf{V}_{\rho_0}), \quad \sigma^2 > 0,$$

with

$$(9) \quad (\mathbf{V}_{\rho_0})_{ij} = \exp(-d_{ij}/\rho_0), \quad \rho_0 > 0, \quad 1 \leq i, j, \leq N,$$

where d_{ij} denotes the Euclidean distance between two sites i and j . The exponential correlation function is used for computational convenience. In principle, it is possible to use other covariance functions, for instance, other members of the Matérn family.

The approximation in (5) assumes that h_{θ} is approximately constant in each cell. If some cells are considered to be too large for this approximation to be reasonable, additional points \mathbf{s}_j^* can be added for which all observations are missing. Since such additional points increase the computational load, some compromise has to be found between accuracy and computational feasibility.

2.2. *Interpretation and parametrization of the kernel function.* For the purpose of interpretation, we note that, in the limit when the temporal spacing goes to zero, the solution of the IDE (3) can also be written as the solution of the stochastic partial differential equation (SPDE) [see Brown et al. (2000)]

$$(10) \quad \frac{\partial}{\partial t} \xi_t(\mathbf{s}) = -\boldsymbol{\mu}_t \cdot \nabla \xi_t(\mathbf{s}) + \frac{1}{4} \nabla \cdot \boldsymbol{\Sigma} \nabla \xi_t(\mathbf{s}) - \eta \xi_t(\mathbf{s}) + B_t(\mathbf{s}),$$

where $\nabla = (\frac{\partial}{\partial x}, \frac{\partial}{\partial y})$ is the gradient operator and where $B_t(\mathbf{s})$ is temporally independent and spatially dependent. The terms have the following interpretations: $\boldsymbol{\mu}_t \cdot \nabla \xi_t(\mathbf{s})$ models advection, $\boldsymbol{\mu}_t$ being a drift or velocity vector. The second term is a diffusion term that can incorporate anisotropy, and $-\eta \xi_t(\mathbf{s})$ accounts for damping. The damping parameter η is related to ϕ and $\boldsymbol{\Sigma}$ through $\eta = -\log(\phi\pi|\boldsymbol{\Sigma}|^{1/2})$. $B_t(\mathbf{s})$ is a source-sink or stochastic forcing term that can be interpreted as modeling convective phenomena. This interpretation is based on the reasoning that typically convective precipitation cells emerge and cease on the domain of interest in contrast to larger scale advective precipitation that is being transported over the area.

We now turn to the discussion of the parameterization of $\boldsymbol{\mu}_t$ and $\boldsymbol{\Sigma}$. In our application, we have information about wind. It is assumed that the drift term $\boldsymbol{\mu}_t$ is proportional to this external wind vector. With $\boldsymbol{\mu}_t$ varying over time, the model is temporally nonstationary. It is also conceivable that in certain situations $\boldsymbol{\Sigma}$ or η may vary over time and/or space, thus obtaining different forms of nonstationarity. Concerning $\boldsymbol{\Sigma}$, it is thought that potential anisotropy is related to topography. Denoting by \mathbf{w}_t the wind vector at time t , we assume

$$(11) \quad \begin{aligned} \boldsymbol{\mu}_t &= u \cdot \mathbf{w}_t \quad \text{and} \\ \boldsymbol{\Sigma}^{-1} &= \frac{1}{\rho_1^2} \begin{pmatrix} \cos \alpha & \sin \alpha \\ -c \cdot \sin \alpha & c \cdot \cos \alpha \end{pmatrix}^T \begin{pmatrix} \cos \alpha & \sin \alpha \\ -c \cdot \sin \alpha & c \cdot \cos \alpha \end{pmatrix}, \end{aligned}$$

where $u \in \mathbb{R}$, $c > 0$, and $\alpha \in [0, \pi/2]$. We use a wind vector which is averaged over the entire area, but the wind could also change locally. The motivation for writing $\boldsymbol{\Sigma}$ in the given form comes from considering a coordinate transformation

$$(12) \quad \begin{pmatrix} x' \\ y' \end{pmatrix} = \begin{pmatrix} \cos \alpha & \sin \alpha \\ -c \cdot \sin \alpha & c \cdot \cos \alpha \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix},$$

where the parameter α is the angle of rotation, and c determines the degree of anisotropy, $c = 1$ corresponding to the isotropic case. ρ_1 is a range parameter that

determines the degree of interaction between spatial and temporal correlation. See Section 4.2 for an illustration of a kernel with the above parametrization.

The resulting model is nonstationary and incorporates anisotropy. Finally, we note that there are various other possible choices of parametrizations. For instance, a relatively simple model can be obtained by assuming

$$(13) \quad \boldsymbol{\mu}_t = \mathbf{0} \quad \text{and} \quad \boldsymbol{\Sigma}^{-1} = \frac{1}{\rho_1^2} \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix},$$

that is, no drift and an isotropic diffusion term. There is still spatio-temporal interaction, though, which implies that the model is not separable in the sense that (16) does not hold. We can simplify further and take not only $\boldsymbol{\mu}_t = \mathbf{0}$, but also $\boldsymbol{\Sigma} = \mathbf{0}$, leading to \mathbf{G}_t being the identity matrix

$$(14) \quad \boldsymbol{\xi}_t = \phi \boldsymbol{\xi}_{t-1} + \boldsymbol{\varepsilon}_t.$$

This means that each point at time $t - 1$ only has an influence on itself at time t , that is, there is no spatio-temporal interaction and the model is separable.

2.3. Discussion of the model.

Propagator matrix \mathbf{G}_t . Using a parametrized propagator matrix \mathbf{G}_t in (6) has the obvious advantage that less parameters are needed than in the general case, in which each entry in the matrix has to be estimated, resulting in N^2 parameters. Moreover, in contrast to the general case, the parametric approach allows for making predictions at sites where no measurements are available, which is often of interest in applications.

Space resolution consistency. At first sight, it might be tempting to use a simpler parametrization of \mathbf{G}_t not based on a convolution but of the form

$$(15) \quad (\mathbf{G}_t)_{ij} = \exp(-(d_{ij}/\rho_1)^2).$$

However, such a model has the following important drawback. Assume, for instance, that a station i is surrounded by two neighboring sites j and k . Say that both stations j and k lie at the same distance from i but in different directions. Consequently, j and k at time $t - 1$ exercise the same influence on i at time t . If one adds an additional station l very close to k , the joint influence of k and l at time $t - 1$ on site i at time t would then approximately be twice as big as the one of site j . This means that the distribution of the process at point i depends on the number and the location of stations in the neighborhood at which it has been observed. The convolution model, on the other hand, does not exhibit this drawback. Furthermore, the convolution model has the advantage that it is “space resolution consistent,” that is, it retains approximately its temporal Markovian structure if one, or several, sites are removed from the domain. This does not hold true for the simpler vector autoregressive model as specified in (15).

Space–time covariance structure. In the following, let us turn to the spatio-temporal dependence structure of the latent process $\boldsymbol{\xi}_t$. A random field $\xi_t(\mathbf{s})$,

$(\mathbf{s}, t) \in \mathbb{R}^2 \times \mathbb{R}$ is said to have a *separable* covariance structure [Gneiting, Genton and Guttorp (2007)] if there exist purely spatial and purely temporal covariance functions C_S and C_T , respectively, such that

$$(16) \quad \text{cov}(\xi_{t_1}(\mathbf{s}_1), \xi_{t_2}(\mathbf{s}_2)) = C_S(\mathbf{s}_1, \mathbf{s}_2) \cdot C_T(t_1, t_2).$$

The convolution based approach allows for nonseparable covariance structures, whereas the separable autoregressive model in (14) has a separable covariance structure.

Extremal events. For the data model as specified in equation (1), Hernández, Guenni and Sansó (2009) showed that the distribution of the maxima is a Gumbel. If the focus lies on extremal events, other distributions, which have Fréchet maxima, can be used, for instance, a t -distribution. The t -distribution is particularly attractive since it is a scale mixture of normal distributions. To be more specific, if S_t has a χ_{df}^2 distribution, then $\mathbf{W}_t = \mathbf{x}_t^T \boldsymbol{\beta} + (\boldsymbol{\xi}_t + \mathbf{v}_t) / \sqrt{S_t/df}$ has a multivariate t -distribution. This means that the fitting algorithm introduced below can be extended to the t -distribution case by introducing an additional latent variable S_t .

3. Fitting and prediction. Fitting is done using a Markov chain Monte Carlo method (MCMC), the Metropolis–Hastings algorithm [Metropolis et al. (1953), Hastings (1970)]. Concerning most parameters, it will be shown that the full conditionals are known distributions. Therefore, Gibbs sampling [Gelfand and Smith (1990)] can be used in these cases.

For convenience and later use, we combine the parameters characterizing the model into a vector $\boldsymbol{\theta} = (\lambda, \boldsymbol{\beta}', \tau^2, \sigma^2, \rho_0, \boldsymbol{\vartheta}')$ and call them *primary parameters*. Our goal is to simulate from the joint posterior distribution of these parameters and the latent variables $\boldsymbol{\xi} = (\boldsymbol{\xi}_1, \dots, \boldsymbol{\xi}_T), \boldsymbol{\xi}_0$, and $\mathbf{W} = (\mathbf{W}_1, \dots, \mathbf{W}_T)$. We note that those $W_t(\mathbf{s}_i)$ that correspond to observed values above zero are known. In that case the full conditional distribution consists of a Dirac distribution at $Y_t(\mathbf{s}_i)^{1/\lambda}$. For handling the censored values and for allowing for missing values, we adopt a data augmentation approach [Smith and Roberts (1993)] as specified below in equation (19). See Section 3.1 for more details.

Assuming prior independence among the primary parameters, the prior distributions are specified as

$$(17) \quad \begin{aligned} &P[\lambda, \boldsymbol{\beta}, \tau^2, \sigma^2, \rho_0, \phi, u, \rho_1, \alpha, c, \boldsymbol{\xi}_0] \\ &\propto \frac{1}{\tau^2} \frac{1}{\sigma^2} P[\rho_0] P[\rho_1] P[u] P[c] P[\alpha] P[\boldsymbol{\xi}_0 | \sigma^2, \rho_0] \end{aligned}$$

with $\boldsymbol{\xi}_0$ having a normal prior $P[\boldsymbol{\xi}_0 | \sigma^2, \rho_0] = N(0, \sigma^2 \mathbf{V}_{\rho_0})$. Further, ρ_0 and ρ_1 have gamma priors with mean μ_ρ and variance σ_ρ^2 . For c , we assume a gamma prior with mean 1 and variance 1, α has a uniform prior on $[0, \pi/2]$, and u has a normal prior with mean 0 and variance 10^4 . Further, we assume locally uniform priors on $\log(\tau^2)$ and $\log(\sigma^2)$ as well as for ϕ, λ and $\boldsymbol{\beta}$.

In our application, we choose to use informative priors for ρ_0 and ρ_1 . It is known that in model-based geostatistics difficulties can arise when estimating the variance and scale parameters of the exponential covariogram [see, e.g., [Warnes and Ripley \(1987\)](#), [Mardia and Watkins \(1989\)](#), [Diggle, Tawn and Moyeed \(1998\)](#)]. For the geostatistical covariance model, [Zhang \(2004\)](#) shows that the product of the two parameters can be estimated consistently, and [Stein \(1990\)](#) shows that it is the product of the two parameters that matters more than the individual parameters for spatial interpolation. Further, [Berger, De Oliveira and Sansó \(2001\)](#) show that, at least in the simplest setting, the posterior of the range parameters is improper for most noninformative priors. Given these considerations, we think that using informative priors for the two range parameters ρ_0 and ρ_1 is appropriate. In our example, we chose priors with mean $\mu_\rho = 100$ and variance $\sigma_\rho^2 = 10$. We have tried different informative priors. The less informative they are, the worse are the mixing properties of the MCMC algorithm. In line with the results of [Stein \(1990\)](#) and [Zhang \(2004\)](#), we have made the experience that different choices of priors on these range parameters do not have a strong impact on the predictive performance of the model.

The posterior distribution is then proportional to

$$\begin{aligned}
 & \left(\frac{1}{\sigma^2}\right)^{N(T+1)/2+1} \left(\frac{1}{\tau^2}\right)^{NT/2+1} |\mathbf{V}_{\rho_0}|^{-(T+1)/2} \prod_{Y_t(\mathbf{s}_i) > 0} \frac{Y_t(\mathbf{s}_i)^{1/\lambda-1}}{\lambda} \\
 & \times \exp\left(-\frac{1}{2} \sum_{t=1}^T \frac{1}{\tau^2} \|\mathbf{W}_t - \mathbf{x}_t^T \boldsymbol{\beta} - \boldsymbol{\xi}_t\|^2\right. \\
 (18) \quad & \left. + \frac{1}{\sigma^2} (\boldsymbol{\xi}_t - \phi \mathbf{G}_t \boldsymbol{\xi}_{t-1})' \mathbf{V}_{\rho_0}^{-1} (\boldsymbol{\xi}_t - \phi \mathbf{G}_t \boldsymbol{\xi}_{t-1})\right) \\
 & \times \exp\left(-\frac{1}{2} \frac{1}{\sigma^2} \boldsymbol{\xi}_0' \mathbf{V}_{\rho_0}^{-1} \boldsymbol{\xi}_0\right) \cdot P[\rho_0] \cdot P[\boldsymbol{\theta}] \cdot \mathbf{1}_{\{W_t(\mathbf{s}_i) \leq 0 \ \forall i, t: Y_t(\mathbf{s}_i) = 0\}}.
 \end{aligned}$$

The product in the first line is the Jacobian for the power transformation in (1). Note that missing observations do not cause any problem. If $Y_t(\mathbf{s}_i)$ is missing, there is no respective term in the product nor a corresponding condition for the indicator function.

3.1. *Full conditional distributions.* In the following, we derive full conditional distributions for the individual parameters.

It is readily seen that the full conditional of $\boldsymbol{\beta}$ is a multivariate normal distribution, and the full conditional distribution of ϕ is a normal distribution as well. The full conditionals of both σ^2 and τ^2 are inverse gamma distributions.

For obtaining the full conditionals of \mathbf{W}_t , we partition its components according to whether $Y_t(\mathbf{s}_i)$ is above zero, equal to zero, or missing. Denote by $i_t^{[+]}$ those

indices for which $Y_t(\mathbf{s}_i) > 0$, by $i_t^{[0]}$ those with $Y_t(\mathbf{s}_i) = 0$, and by $i_t^{[m]}$ the missing ones. The vector \mathbf{W}_t can then be partitioned into $\mathbf{W}_t^{[+]}$, $\mathbf{W}_t^{[0]}$, and $\mathbf{W}_t^{[m]}$ accordingly. We remark that $\mathbf{W}_t^{[0]}$ and $\mathbf{W}_t^{[m]}$ are latent variables, whereas $\mathbf{W}_t^{[+]}$ corresponds to transformed observed values. In addition, $\mathbf{W}_t^{[0]}$ has the restriction that all its values must be smaller than zero, $\mathbf{W}_t^{[0]} \leq \mathbf{0}$. For facilitating understanding, we note that $W_t(\mathbf{s}_i)$ can be written as

$$\begin{aligned}
 (19) \quad W_t(\mathbf{s}_i) &= W_t^{[+]}(\mathbf{s}_i) = Y_t(\mathbf{s}_i)^{1/\lambda} && \text{if } Y_t(\mathbf{s}_i) > 0 \\
 &= W_t^{[0]}(\mathbf{s}_i) && \text{if } Y_t(\mathbf{s}_i) = 0 \\
 &= W_t^{[m]}(\mathbf{s}_i) && \text{if } Y_t(\mathbf{s}_i) \text{ is missing.}
 \end{aligned}$$

The full conditional of $\mathbf{W}_t^{[m]}$ is then a multivariate normal distribution with mean and covariance

$$(20) \quad \mu_{\mathbf{W}_t^{[m]}} = (\mathbf{x}_t^T \boldsymbol{\beta} + \boldsymbol{\xi}_t)^{[m]} \quad \text{and} \quad \boldsymbol{\Sigma}_{\mathbf{W}_t^{[m]}} = \tau^2 \cdot \mathbf{I}.$$

Similarly, the full conditional distribution of $\mathbf{W}_t^{[0]}$ is a truncated multivariate normal distribution with mean and covariance

$$(21) \quad \mu_{\mathbf{W}_t^{[0]}} = (\mathbf{x}_t^T \boldsymbol{\beta} + \boldsymbol{\xi}_t)^{[0]} \quad \text{and} \quad \boldsymbol{\Sigma}_{\mathbf{W}_t^{[0]}} = \tau^2 \cdot \mathbf{I}.$$

As mentioned before, the full conditional of $\mathbf{W}_t^{[+]}$ is a Dirac distribution with point mass at $(\mathbf{Y}_t^{[+]})^{1/\lambda}$.

Concerning the latent variables $(\boldsymbol{\xi}_0, \boldsymbol{\xi}_1, \dots, \boldsymbol{\xi}_T)$, we note that conditional on $\boldsymbol{\theta}$, $(\boldsymbol{\xi}_t, \mathbf{W}_t)$ is a linear Gaussian state space model. Therefore, a sample from the joint full conditional of $(\boldsymbol{\xi}_0, \boldsymbol{\xi}_1, \dots, \boldsymbol{\xi}_T)$ can be obtained using the forward filtering backward sampling (FFBS) algorithm proposed by Carter and Kohn (1994) and Frühwirth-Schnatter (1994). The forward filtering step corresponds to the Kalman filter [see, e.g., West and Harrison (1997) and Künsch (2001)].

Alternatively, one can also use single t updates. The full conditional of one $\boldsymbol{\xi}_t$, $0 \leq t \leq T$, is a normal distribution $N(\boldsymbol{\mu}_{\boldsymbol{\xi}_t}, \boldsymbol{\Sigma}_{\boldsymbol{\xi}_t})$. In the case of the separable model, the mean $\boldsymbol{\mu}_{\boldsymbol{\xi}_t}$ depends on $\boldsymbol{\xi}_{t-1}$ and $\boldsymbol{\xi}_{t+1}$, whereas the covariance matrix $\boldsymbol{\Sigma}_{\boldsymbol{\xi}_t}$ does not depend on t . This is convenient for simulation since its Cholesky decomposition has to be calculated only once in each update cycle. In contrast, in the sampling step of the FFBS algorithm, one has to calculate a Cholesky decomposition for each t . The advantage that the FFBS algorithm mixes better than the single t update algorithm per update cycle is outweighed by the fact that an update cycle of the single t update algorithm is a lot faster than one of the FFBS algorithm. Thus, more effective samples can be obtained with the single t update algorithm per time. In the case of the nonstationary anisotropic drift model, though, $\boldsymbol{\Sigma}_{\boldsymbol{\xi}_t}$ in the single t update algorithm is not constant over time. Thus, a Cholesky decomposition needs to be computed for each t anyway, meaning that the FFBS algorithm is preferable.

In summary, we made the experience that it is recommendable to use single t updates for temporally stationary models where the covariance Σ_{ξ_t} of the full conditional of one ξ_t is constant over time. If Σ_{ξ_t} changes over time, we recommend using the FFBS algorithm.

For the remaining parameters, that is, ρ_0 , ϑ (excluding ϕ) and λ , there is no apparent distribution family from which one can simulate. Therefore, Metropolis steps will be used. We note that the full conditional distribution of λ is proportional to

$$(22) \quad \prod_{Y_t(\mathbf{s}_i) > 0} \left(\frac{Y_t(\mathbf{s}_i)^{1/\lambda - 1}}{\lambda} \right) \exp\left(-\frac{1}{2} \sum_{Y_t(\mathbf{s}_i) > 0} \frac{1}{\tau^2} \|Y_t(\mathbf{s}_i)^{1/\lambda} - \mathbf{x}_t^T \boldsymbol{\beta} - \xi_t\|^2 \right).$$

The parameter λ is sampled on the log-scale. This means that we first transform it to the log scale. Then a proposal is obtained by sampling from a normal distribution with the mean equal to the last value of the parameter. Thereafter, this proposal is accepted with a probability that is given by the usual Metropolis–Hasting algorithm [see, e.g., Chib and Greenberg (1995)].

Finally, ρ_0 and ϑ (excluding ϕ) are sampled together. The full conditional is proportional to

$$(23) \quad \exp\left(-\frac{1}{2\sigma^2} \left(\sum_{t=1}^T (\xi_t - \phi \mathbf{G}_t \xi_{t-1})' \mathbf{V}_{\rho_0}^{-1} (\xi_t - \phi \mathbf{G}_t \xi_{t-1}) + \xi_0' \mathbf{V}_{\rho_0}^{-1} \xi_0 \right) \right) \times |\mathbf{V}_{\rho_0}|^{-(T+1)/2}.$$

3.2. *Prediction.* We consider predictions at new locations and/or times as well as predictions of areal averages. It turns out that in the case of areal averages, the Voronoi tessellation is again useful.

One way to obtain predictions is to augment the data \mathbf{Y}_{obs} with missing values at the locations or times where predictions are made. When doing so, the MCMC algorithm implicitly draws from the corresponding predictive distribution. See the previous Section 3.1 on how to handle missing values.

If one does not specify the points in space and time where predictions are to be made prior to model fitting, the predictive distribution of a new set of observations $\mathbf{Y}^* = (Y_{t_1}^*(\mathbf{s}_1^*), \dots, Y_{t_k}^*(\mathbf{s}_k^*))'$ is calculated as

$$(24) \quad \begin{aligned} P[\mathbf{Y}^* | \mathbf{Y}_{\text{obs}}] &= \int P[\mathbf{Y}^* | \xi^*, \boldsymbol{\theta}] P[\xi^* | \xi, \boldsymbol{\theta}] P[\xi, \boldsymbol{\theta} | \mathbf{Y}_{\text{obs}}] d\xi^* d\xi d\boldsymbol{\theta} \\ &\approx \frac{1}{m} \sum_{i=1}^m \int P[\mathbf{Y}^* | \xi^*, \boldsymbol{\theta}^{(i)}] P[\xi^* | \xi^{(i)}, \boldsymbol{\theta}^{(i)}] d\xi^* \\ &\approx \frac{1}{m} \sum_{i=1}^m P[\mathbf{Y}^* | \xi^{*(i)}, \boldsymbol{\theta}^{(i)}], \end{aligned}$$

where \mathbf{Y}_{obs} denotes the observed data, $\boldsymbol{\xi}$ and $\boldsymbol{\xi}^*$ the latent Gaussian process at the observed and predicted sites, respectively, and $\boldsymbol{\theta}$ all the remaining parameters. Samples $\boldsymbol{\theta}^{(i)}$ and $\boldsymbol{\xi}^{(i)}$, $i = 1, \dots, m$, from their posterior distribution are obtained by the MCMC algorithm, and $\boldsymbol{\xi}^{*(i)}$ is sampled from $P[\boldsymbol{\xi}^* | \boldsymbol{\xi}^{(i)}, \boldsymbol{\theta}^{(i)}]$.

When $\boldsymbol{\xi}^*$ is modeled at the same sites as $\boldsymbol{\xi}$ but at different time points, the distribution $P[\boldsymbol{\xi}^* | \boldsymbol{\xi}^{(i)}, \boldsymbol{\theta}^{(i)}]$ is Gaussian and readily obtained using (6).

In the case when predictions are made at unobserved sites $s \in S$ and time t , $P[\boldsymbol{\xi}_t^* | \boldsymbol{\xi}, \boldsymbol{\theta}]$ can be calculated as described in the following. First, because of the temporal Markov property, $P[\boldsymbol{\xi}_t^* | \boldsymbol{\xi}, \boldsymbol{\theta}]$ is equal to $P[\boldsymbol{\xi}_t^* | \boldsymbol{\xi}_{t-1}, \boldsymbol{\xi}_t, \boldsymbol{\xi}_{t+1}, \boldsymbol{\theta}]$. This density is then obtained by considering the augmented model

$$(25) \quad \begin{pmatrix} \boldsymbol{\xi}_t \\ \boldsymbol{\xi}_t^* \end{pmatrix} = \phi \begin{pmatrix} \mathbf{G}_t \\ \mathbf{G}_t^* \end{pmatrix} \boldsymbol{\xi}_{t-1} + \begin{pmatrix} \boldsymbol{\varepsilon}_t \\ \boldsymbol{\varepsilon}_t^* \end{pmatrix},$$

$$\boldsymbol{\xi}_{t+1} = \phi(\mathbf{H}_{t+1} \quad \mathbf{H}_{t+1}^*) \begin{pmatrix} \boldsymbol{\xi}_t \\ \boldsymbol{\xi}_t^* \end{pmatrix} + \boldsymbol{\varepsilon}_{t+1},$$

where \mathbf{G}_t^* is defined analogously to (7), \mathbf{H}_{t+1} and \mathbf{H}_{t+1}^* are obtained from the same approximations as in (5), and the covariances of $\boldsymbol{\varepsilon}_t$ and $\boldsymbol{\varepsilon}_t^*$ are as in (9). By (25), the conditional distribution of $\boldsymbol{\xi}_t, \boldsymbol{\xi}_t^*, \boldsymbol{\xi}_{t+1}$ given $\boldsymbol{\xi}_{t-1}$ is normal. Therefore, also the conditional distribution of $\boldsymbol{\xi}_t^*$ given $\boldsymbol{\xi}_{t-1}, \boldsymbol{\xi}_t, \boldsymbol{\xi}_{t+1}$ is Gaussian. Its mean and covariance can be computed by noting that

$$(26) \quad \begin{aligned} P[\boldsymbol{\xi}_t^* | \boldsymbol{\xi}_{t-1}, \boldsymbol{\xi}_t, \boldsymbol{\xi}_{t+1}, \boldsymbol{\theta}] &\propto P[\boldsymbol{\xi}_{t+1} | \boldsymbol{\xi}_t^*, \boldsymbol{\xi}_t, \boldsymbol{\theta}] P[\boldsymbol{\xi}_t^* | \boldsymbol{\xi}_{t-1}, \boldsymbol{\xi}_t, \boldsymbol{\theta}] \\ &\propto P[\boldsymbol{\xi}_{t+1} | \boldsymbol{\xi}_t^*, \boldsymbol{\xi}_t, \boldsymbol{\theta}] P[\boldsymbol{\xi}_t, \boldsymbol{\xi}_t^* | \boldsymbol{\xi}_{t-1}, \boldsymbol{\theta}] \end{aligned}$$

and then completing the square in the exponent of the last expression.

In many cases, for instance, when the focus lies on flooding, areal averages

$$(27) \quad \bar{Y}_t^{(A^*)} = \frac{1}{|A^*|} \int_{A^*} Y_t(\mathbf{s}) \, ds$$

of precipitation are of interest. If $Y_t(\mathbf{s})$ is observed on an irregular grid, one could first define a regular grid, then interpolate the nonobserved grid points, and approximate the integral in (27) by a Riemann sum. However, since the regular grid usually becomes very large, this is computationally expensive. Instead, we propose to use the Voronoi tessellation once again to approximate the integral

$$(28) \quad \bar{Y}_t^{(A^*)} = \frac{1}{|A^*|} \int_{A^*} Y_t(\mathbf{s}) \, ds \approx \frac{1}{|A^*|} \sum_{j=1}^N Y_t(\mathbf{s}_j) |A_j \cap A^*|.$$

Thereby, an adequate weight $|A_j \cap A^*|$ is given to each station. Samples from the predictive distribution of $\bar{Y}_t^{(A^*)}$ can be obtained by simulating $Y_t^{(i)}(\mathbf{s}_j)$ from their predictive distribution and inserting them in (28).

We note that the areal prediction becomes deterministic if all $Y_t(\mathbf{s}_j)$ consist of observed values. This means that uncertainty about values of $Y_t(\mathbf{s})$ at locations

where no observations are made is implicitly ignored with the above approximation. This can be amended for by first making predictions at a few sites where no observations were made. Inserting additional unobserved sites can also be useful in other cases. For instance, if A^* cuts off a substantial part of any A_j , that is, $A_j \cap A^*$ is much smaller than A_j but not empty, the areal prediction might be improved by replacing $Y_t(s_j)$ by the prediction of Y_t at the center of gravity of $A_j \cap A^*$, or if the area A^* is small and contains only a few stations, improved predictions of the areal average can be obtained by making predictions at a few additional points inside the area.

4. Application to short term prediction of precipitation. We apply the model to obtain short term forecasts of precipitation. Such forecasts are important, for instance, for agriculture and flooding. The traditional way for obtaining precipitation forecasts is the use of numerical weather prediction (NWP) models. NWP models solve complex, nonlinear equations emulating the dynamics of the atmosphere. Typically, NWP models require a lot of computational resources to run. Fitting our statistical model using the MCMC algorithm presented above is also computationally intensive. However, once the statistical model is fitted and assuming that the posterior of the primary parameters does not change (see Section 4.3 for more details), predictions are computationally a lot cheaper. Furthermore, the statistical model can be used in situations where there are no NWP models available or to obtain predictions at different temporal resolutions than the one at which the NWP model operates.

4.1. *The data.* The data consists of three-hourly precipitation amounts collected by 26 stations around the Swiss Plateau from the beginning of December 2008 to the end of March 2009, making a total of 968 time periods. The data were provided by MeteoSwiss. We use the first three months, consisting of 720 time periods, for fitting the model. The remaining month March, consisting of 248 time periods, is set aside for model evaluation. The locations of these stations are shown in Figure 1. In Figure 2, a time series plot of the observed precipitation at one station (corresponding to the station with the acronym WYN in Figure 1) and of the weighted areal average is shown. Concerning the latter, we take the weighted average over the entire spatial domain. Figure 3 shows the spatial distribution of the precipitation accumulated over time.

The covariates consist of the x- and y-coordinates (km), altitude (m), temperature ($^{\circ}\text{C}$), dew point ($^{\circ}\text{C}$) and specific humidity (%). Specific humidity is the ratio of water vapor to dry air in a particular mass. It is expected to be positively related to precipitation. The dew point is the temperature to which a given parcel of humid air must be cooled, at constant barometric pressure, for water vapor to condense into water. Thus, the lower the dew point, the lower is the chance for precipitation. However, specific humidity and dew point are considerably negatively

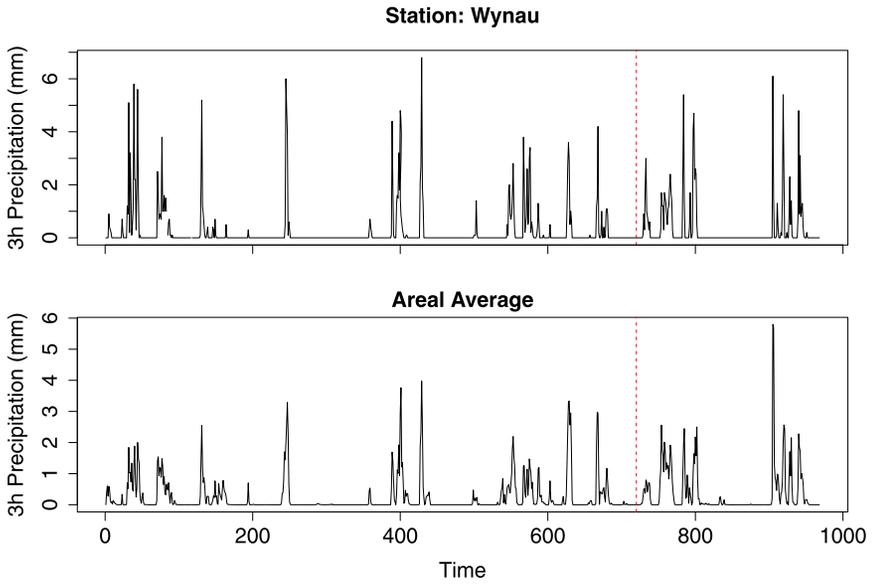


FIG. 2. Precipitation versus time. The lines are observed precipitation of one station (corresponding to the station with the acronym WYN in Figure 1) and of the areal average. The time axis is in 3 h steps starting at December 1, 2008. The dotted vertical line separates the training and test data.

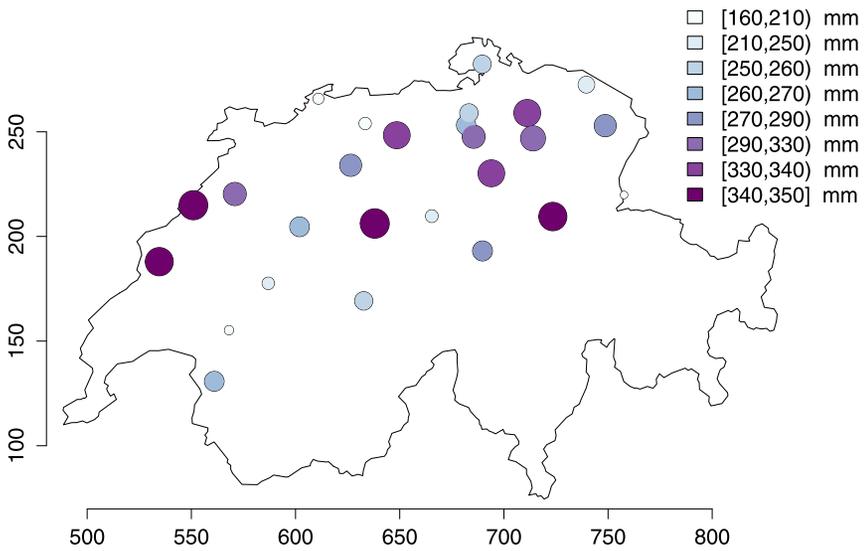


FIG. 3. Illustration of the spatial distribution of precipitation. The circles display the cumulative rainfall amounts over time at the stations. The larger the circle and the darker the color, the higher is the cumulative precipitation amount. Both axes are in km.

correlated. This makes it unclear, a priori, what their joint relation to precipitation is like. Temperature, dew point and specific humidity are predicted variables obtained from an NWP model called COSMO-2. From the same model, we also obtain wind predictions (speed is in m/s). Predictions of the statistical model are evaluated by comparing them to precipitation forecasts from the same NWP. Having a high resolution with a grid spacing of 2.2 km, the NWP model is able to resolve convective dynamics. The NWP model produces predictions once a day for 24 hours ahead starting at 0:00UTC. After assimilation and computation, forecasts are available at around 1:30UTC. For all meteorological variables, we use values at approximately 1000 m above ground. This is the height where we think these variables to be most influential for precipitation. All covariates are centered and standardized to unit variance. Centering covariates around their means is used in order to avoid correlations of the regression coefficients with the intercept and to reduce posterior correlations.

4.2. *Fitting and results.* In the following, the nonstationary anisotropic model incorporating the wind as an external drift term (see Section 2) is fitted. In addition, we also fit a separable model. We simulate from the posterior distributions of these models as outlined in Section 3.

After the burn-in period consisting of 5000 draws, 195,000 samples from the Markov chain were used to characterize posterior distributions. Convergence was monitored by inspecting trace plots.

In Table 1 we show posterior modes as well as 95% credible intervals for the different parameters of the nonstationary anisotropic drift model. The coefficients of the geographic coordinates are not significant. Specific humidity has a large positive coefficient. As expected, higher humidity implies more rainfall. The dew point is also positively related to precipitation. Higher temperatures, on the other hand, seem to imply less precipitation.

For interpreting the fitted parameters governing the convolution kernel (ρ_1 , c , α and u), we illustrate in Figure 4 the convolution kernel over the region where the stations lie. The parameters ρ_1 , c , α and u are taken at their posterior mode. The plot is interpreted as follows. The height of the kernel is the level of influence that $\xi_{t-1}(s')$ at location s' has on $\xi_t(s)$ at location s as a function of $s' - s$. In other words, the colors represent the lag-1 influence of the other stations on the station Wynau which is used as origin in the plot. The white arrow represents the drift vector $\mu_t = u \cdot \mathbf{w}_t$ at time $t = 429$, \mathbf{w}_t being the wind vector. Note that this transport vector changes over time, thus causing temporal nonstationarity. The time $t = 429$ illustrates a meteorological situation with the typically predominant southwestern wind direction.

With c and α being approximately 4 and 0.7, we observe anisotropy along the south-east north-west direction. This corresponds to the topography of the region, as the area containing a majority of the stations lies between two mountain ranges:

TABLE 1
Posterior modes and 95% credible intervals for the nonstationary, anisotropic model with an external drift

	Mode	2.5%	97.5%
Intercept	-1.05	-1.21	-0.929
X	-0.0473	-0.133	0.0541
Y	-0.0108	-0.0846	0.0531
Z	0.00347	-0.0169	0.0247
Temp	-0.717	-0.856	-0.583
Dew point	0.406	0.187	0.601
Spec hum	1.14	0.949	1.33
λ	1.58	1.54	1.62
τ^2	0.0685	0.0451	0.0943
σ^2	1.04	0.953	1.17
ρ_0	92	86.4	97.9
ϕ	0.000159	0.000147	0.00017
ρ_1	93.6	88.1	99.4
c	4.1	3.61	4.63
α	0.704	0.658	0.777
u	0.879	0.645	1.1

the Jura to the north-west and the Alps to the south-east. Correlations are expected to be higher along the flat part between these two mountain ranges.

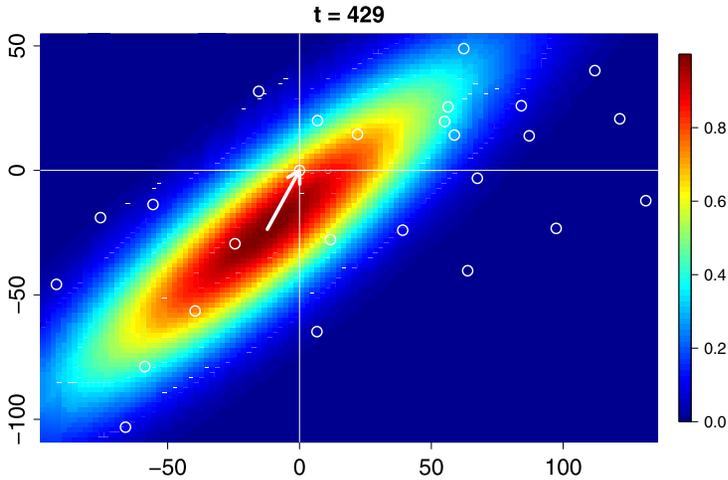


FIG. 4. *Illustration of the convolution kernel at time $t = 429$. The colors indicate the lag-1 influence of the other stations on the station Wynau. The white arrow represents the drift caused by a south-west wind at this time point. The dots represent the observation stations. The axes are in km.*

Furthermore, the plot shows how the external drift shifts the convolution kernel. Apparently, the southwestern neighbor (Bern) has the highest influence on Wynau in this situation, with wind coming from the southwest. Gneiting et al. (2006) observe a similar phenomenon in wind speed data over the U.S. Pacific Northwest where there is also a predominant wind direction causing asymmetric cross-correlations.

4.3. *Short term prediction of precipitation.* In the following, we apply the fitted models to produce short term predictions of precipitation. As mentioned before, we have fitted the model to the first 720 time periods from December 2008 to February 2009. From this we obtain posterior distributions for the primary parameters. Predictions for the time periods in March that were set aside are obtained as described in the following.

Ideally, one would run the full MCMC algorithm at each time point, including all data up to the point, and obtain predictive distributions from this. However, since this is rather time consuming, we make the following approximation. We assume that the posterior distribution of the primary parameters given $\mathbf{Y}_{1:t} = \{\mathbf{Y}_1, \dots, \mathbf{Y}_t\}$ is the same for all $t \geq 720$. That is, we neglect the additional information that the observations in March give about the primary parameters. In practice, this means that posterior distributions of the primary parameters are calculated only once, namely, on the data set from December 2008 to February 2009.

For each time $t \geq 720$, we make up to 8 steps ahead forecasts. That is, we sample from the predictive distribution of \mathbf{Y}_{t+k}^* , $k = 1, \dots, 8$, given $\mathbf{Y}_{1:t} = \{\mathbf{Y}_1, \dots, \mathbf{Y}_t\}$ and given the posterior of the primary parameters based on the data from December 2008 to February 2009. Since the NWP produces forecasts for the three meteorological covariates once a day, for each prediction time $t + k$, the forecasts made at 0:00UTC of the same day are used. Sampling from the predictive distribution consists of imputing the augmented data \mathbf{W} and sampling from the latent process ξ . These two steps are done as described in Section 3. To generate one sample from the predictive distribution takes around 3.5 seconds on an AMD Athlon(tm) 64 X2 Dual Core Processor 5600+ with a 2900 MHz CPU clock rate. We use 200 samples to characterize each predictive distribution.

The assumption that the posterior of the primary parameters does not change may be questionable over longer time periods and when one moves away from the time period from which data is used to obtain the posterior distribution. But since all our data lies in the winter season, we think that this assumption is reasonable. If longer time periods are considered, one could use sliding training windows or model the primary parameters as evolving dynamically over time. One can also investigate how the predictive performance deteriorates with increasing lags between predictions and last time point from which data is used to fit the model.

In addition to the separable model and the nonstationary anisotropic drift model, we fit a model with no autoregressive term, that is, with $\phi = 0$. Further, to assess

how much information stems from the three meteorological covariates (temperature, dew point and specific humidity) and how much from the dynamic spatio-temporal model, we also fit the nonstationary anisotropic drift model without including these covariates. For each model, we calculate pointwise predictions for the individual stations and also predictions for the areal average. The latter are obtained using the Voronoi tessellation as described in Section 3.2.

In order to assess the performance of the probabilistic predictions, we use the continuous ranked probability score (CRPS) [Matheson and Winkler (1976)]. The CRPS is a strictly proper scoring rule [Gneiting and Raftery (2007)] that assigns a numerical value to probabilistic forecasts and assesses calibration and sharpness simultaneously [Gneiting, Balabdaoui and Raftery (2007)]. It is defined as

$$(29) \quad \text{CRPS}(F, y) = \int_{-\infty}^{\infty} (F(x) - \mathbf{1}_{\{y \leq x\}})^2 dx,$$

where F is the predictive cumulative distribution function, y is the observed realization, and $\mathbf{1}$ is an indicator function. It can be equivalently calculated as

$$(30) \quad \text{CRPS}(F, y) = E_F |Y - y| - \frac{1}{2} E_F |Y - Y'|,$$

where Y and Y' are independent random variables with distribution F . If a sample $Y^{(1)}, \dots, Y^{(m)}$ from F is available, it can be approximated by

$$(31) \quad \frac{1}{m} \sum_{i=1}^m |Y^{(i)} - y| - \frac{1}{2m^2} \sum_{i,j=1}^m |Y^{(i)} - Y^{(j)}|.$$

In Figure 5 the average CRPS of the pointwise predictions and the areal predictions are plotted versus lead times. In the left plot, the mean is taken over all stations and time periods, whereas the areal version is an average over all time periods. Predictions \mathbf{Y}_{t+k}^* , $k = 1, \dots, 8$, for the next 8 time steps are made at each time point t . We recall that the NWP model produces predictions for 8 consecutive periods once a day at midnight. For simplicity, potential diurnal variation in the accuracy of the predicted covariates is ignored.

We see that the nonstationary anisotropic drift model (“ConvAR”) has clearly the best performance among the three models. In particular, the nonseparable convolution based model performs better than the simpler separable spatio-temporal model (“SAR”). Not surprisingly, the model without temporal dependency (“NoAR”) performs worse than the other two models. Comparing the “ConvAR” model, the nonstationary convolution model without covariates (“ConvAR No Cov”), and the “NoAR” model, we see that the main source of predictive performance at small lead times is not the covariates but the dynamic spatio-temporal model. In the areal case, the nonstationary convolution model without covariates even outperforms the simple autoregressive model including covariates at small lead times. With increasing lead time, the meteorological covariates contribute more to the predictive performance and the dynamic spatio-temporal model becomes less important.

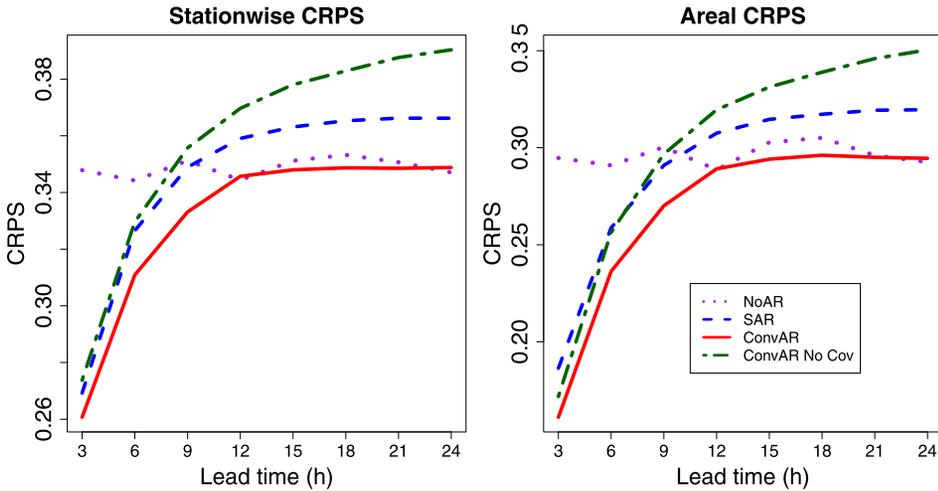


FIG. 5. Comparison of statistical models. The continuous ranked probability score (CRPS) of forecasts versus number of consecutive time periods for which predictions are made is shown. On the left are CRPSs of station specific forecasts and on the right are CRPSs of areal forecasts. “NoAR” denotes the model without an autoregressive term, “SAR” the one with a separable covariance structure, and “ConvAR” the convolution based nonstationary anisotropic drift model. All three models include the covariates described in Section 4.1. A convolution based model without including covariates (“ConvAR No Cov”) is also fitted. The unit of the CRPS is mm.

We also compare the performance of the predictions from the nonstationary anisotropic drift model with predictions obtained from the NWP model. Since the NWP model produces deterministic forecasts, we use the mean absolute error (MAE). In order to make the comparison fair, we first reduce the statistical distributional forecast to a point forecast by taking the median [see Gneiting (2011) on why this is a reasonable choice]. As mentioned, the NWP model produces predictions once a day starting at 0:00UTC. Predictions are then made for eight consecutive time periods corresponding to 24 h ahead. This means that the time of day also corresponds to the lead time. This is in contrast to the above comparison of the different statistical models where 8 step ahead predictions were made at all time periods.

In Figure 6 the mean absolute error (MAE) of forecasts versus lead time, or time of day is shown. In addition, in Table 2 we report MAEs averaged over all lead times. Note that there is one particular day (March 24) when heavy rainfall occurred shortly after 0:00UTC. We report results including (thin lines) and excluding (bold lines) this day.

Table 2 shows that overall the statistical model outperforms the NWP on a stationwise base. When considering the areal average, the two models perform similarly. Depending on whether March 24 is included or not, the NWP or the statistical model has a slightly lower average MAE.

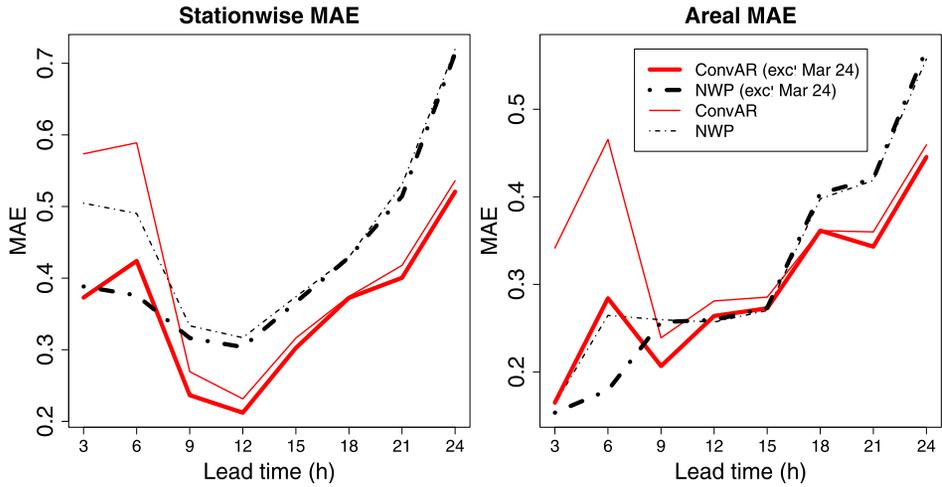


FIG. 6. Comparison of statistical and NWP model. The mean absolute error (MAE) of forecasts versus lead time is shown. Lead time also corresponds to the time of day. The left panel shows MAEs of station specific forecasts averaged over time and the stations, and the right panel shows MAEs of areal forecasts averaged over time. “ConvAR” denotes the convolution based nonstationary anisotropic drift model and “NWP” the NWP model. The bold lines show the results when excluding March 24, 2009. The unit of the MAE is mm.

Furthermore, Figure 6 shows that March 24 considerably affects the performance of the one- and two-step ahead predictions of the statistical model as well as the stationwise performance of the NWP model. When excluding this day, the corresponding MAEs are considerably lower. This shows a typical behavior of our model and statistical models in general: they perform well when, at the time of prediction, the major phenomena (advective fronts) are already observable. In this case, the spatio-temporal statistical model can extrapolate the space–time dynamics of the rainfall process into the future.

Earlier studies have shown that nowcasting methods, including statistical approaches, perform usually better at short lead times (up to one day), while NWP have higher predictive skills at medium ranges [see Kober et al. (2012) or Little,

TABLE 2

Comparison of statistical and NWP model. The mean absolute error (MAE) averaged over all days and lead times is reported. “ConvAR” denotes the convolution based nonstationary anisotropic drift model and “NWP” the NWP model. The unit of the MAE is mm

	ConvAR	NWP	Areal ConvAR	Areal NWP
March 2009	0.41	0.46	0.35	0.32
Excluding March 24	0.36	0.43	0.29	0.31

McSharry and Taylor (2009)]. Our results are in line with these findings in the sense that all lead times used in our application are still in the range of what is considered “short” lead times. However, our model is not just based on past precipitation observations but also on other predicted meteorological variables.

5. Conclusions. A hierarchical Bayesian spatio-temporal model is presented. Incorporating physical knowledge, the dynamic model is nonstationary, anisotropic, and allows for nonseparable covariance structures. It incorporates a drift term that depends on a wind vector. At the data stage, the model determines the probability of rainfall and the rainfall amount distribution together. The model is fitted using Markov chain Monte Carlo (MCMC) methods and applied to obtain short term precipitation forecasts. It performs better than a separable, stationary and isotropic model, and it performs comparably to a deterministic numerical weather prediction model and has the advantage that it quantifies prediction uncertainty.

Even though we have applied the model to prediction of precipitation, it can also be used to predict or interpolate other meteorological quantities of interest.

Future research could focus on adapting the model so that it can be applied to spatially highly resolved data. Using Markov random fields [Rue and Held (2005), Lindgren, Rue and Lindström (2011)] for the innovation process $\boldsymbol{\varepsilon}_t$ might be a potential direction. Alternatively, a dimension reduction approach could be examined; cf. Banerjee et al. (2008). For instance, Sigrist, Künsch and Stahel (2012) approximate an advection-diffusion SPDE to cope with large data sets. Further, the model can be extended by additionally relaxing some assumptions. For instance, the parameters σ^2 , ϕ , ρ_0 , ρ_1 and λ were assumed to be constant over time. Assuming periodicity, Fourier harmonics could be used to model parameters that vary seasonally during the year. Alternatively, the parameters could evolve dynamically over time according to an equation of the form $\vartheta_t = \vartheta_{t-1} + N(0, \sigma_\vartheta^2)$.

Acknowledgments. We thank Vanessa Stauch from MeteoSwiss for providing parts of the data and for interesting discussions. We also would like to thank the Editor and three anonymous referees for their insightful comments and suggestions.

REFERENCES

- AILLIOT, P., THOMPSON, C. and THOMSON, P. (2009). Space–time modelling of precipitation by using a hidden Markov model and censored Gaussian distributions. *J. R. Stat. Soc. Ser. C. Appl. Stat.* **58** 405–426. [MR2750013](#)
- ALLCROFT, D. J. and GLASBEY, C. A. (2003). A latent Gaussian Markov random-field model for spatiotemporal rainfall disaggregation. *J. R. Stat. Soc. Ser. C. Appl. Stat.* **52** 487–498. [MR2012972](#)
- BANERJEE, S., GELFAND, A. E., FINLEY, A. O. and SANG, H. (2008). Gaussian predictive process models for large spatial data sets. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **70** 825–848. [MR2523906](#)
- BARDOSSY, A. and PLATE, E. J. (1992). Space–time model for daily rainfall using atmospheric circulation patterns. *Water Resources Research* **28** 1247–1259.

- BELL, T. (1987). A space–time stochastic model of rainfall for satellite remote-sensing studies. *Journal of Geophysical Research* **92** 9631–9643.
- BELLONE, E., HUGHES, J. P. and GUTTORP, P. (2000). A hidden Markov model for downscaling synoptic atmospheric patterns to precipitation amounts. *Climate Research* **15** 1–12.
- BERGER, J. O., DE OLIVEIRA, V. and SANSÓ, B. (2001). Objective Bayesian analysis of spatially correlated data. *J. Amer. Statist. Assoc.* **96** 1361–1374. [MR1946582](#)
- BERROCAL, V. J., RAFTERY, A. E. and GNEITING, T. (2008). Probabilistic quantitative precipitation field forecasting using a two-stage spatial model. *Ann. Appl. Stat.* **2** 1170–1193. [MR2655654](#)
- BROWN, P. E., KÅRESEN, K. F., ROBERTS, G. O. and TONELLATO, S. (2000). Blur-generated non-separable space–time models. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **62** 847–860. [MR1796297](#)
- BROWN, P. E., DIGGLE, P. J., LORD, M. E. and YOUNG, P. C. (2001). Space–time calibration of radar rainfall data. *J. R. Stat. Soc. Ser. C. Appl. Stat.* **50** 221–241. [MR1833274](#)
- CARTER, C. K. and KOHN, R. (1994). On Gibbs sampling for state space models. *Biometrika* **81** 541–553. [MR1311096](#)
- CHARLES, S., BATES, B. and HUGHES, J. (1999). A spatiotemporal model for downscaling precipitation occurrence and amounts. *Journal of Geophysical Research* **104** 31657–31669.
- CHIB, S. and GREENBERG, E. (1995). Understanding the Metropolis–Hastings algorithm. *Amer. Statist.* **49** 327–335.
- COE, R. and STERN, R. (1982). Fitting models to daily rainfall data. *Journal of Applied Meteorology* **21** 1024–1031.
- COX, D. R. and ISHAM, V. (1988). A simple spatial–temporal model of rainfall. *Proc. R. Soc. Lond. Ser. A Math. Phys. Eng. Sci.* **415** 317–328. [MR0932924](#)
- CRESSIE, N. A. C. (1993). *Statistics for Spatial Data*, 2nd ed. Wiley, New York.
- CRESSIE, N. and HUANG, H.-C. (1999). Classes of nonseparable, spatio-temporal stationary covariance functions. *J. Amer. Statist. Assoc.* **94** 1330–1340. [MR1731494](#)
- CRESSIE, N. and WIKLE, C. K. (2011). *Statistics for Spatio-Temporal Data*. Wiley, Hoboken, NJ. [MR2848400](#)
- DIGGLE, P. J., TAWN, J. A. and MOYEED, R. A. (1998). Model-based geostatistics. *J. R. Stat. Soc. Ser. C. Appl. Stat.* **47** 299–350. With discussion and a reply by the authors. [MR1626544](#)
- FOWLER, H. J., KILSBY, C. G., O’CONNELL, P. E. and BURTON, A. (2005). A weather-type conditioned multi-site stochastic rainfall model for the generation of scenarios of climatic variability and change. *Journal of Hydrology* **308** 50–66.
- FRÜHWIRTH-SCHNATTER, S. (1994). Data augmentation and dynamic linear models. *J. Time Series Anal.* **15** 183–202. [MR1263889](#)
- FUENTES, M., REICH, B. and LEE, G. (2008). Spatial–temporal mesoscale modeling of rainfall intensity using gage and radar data. *Ann. Appl. Stat.* **2** 1148–1169. [MR2655653](#)
- GELFAND, A. E., BANERJEE, S. and GAMERMAN, D. (2005). Spatial process modelling for univariate and multivariate dynamic spatial data. *Environmetrics* **16** 465–479. [MR2147537](#)
- GELFAND, A. E. and SMITH, A. F. M. (1990). Sampling-based approaches to calculating marginal densities. *J. Amer. Statist. Assoc.* **85** 398–409. [MR1141740](#)
- GELFAND, A. E., DIGGLE, P. J., FUENTES, M. and GUTTORP, P., eds. (2010). *Handbook of Spatial Statistics*. CRC Press, Boca Raton, FL. [MR2761512](#)
- GNEITING, T. (2002). Nonseparable, stationary covariance functions for space–time data. *J. Amer. Statist. Assoc.* **97** 590–600. [MR1941475](#)
- GNEITING, T. (2011). Making and evaluating point forecasts. *J. Amer. Statist. Assoc.* **106** 746–762. [MR2847988](#)
- GNEITING, T., BALABDAOUI, F. and RAFTERY, A. E. (2007). Probabilistic forecasts, calibration and sharpness. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **69** 243–268. [MR2325275](#)
- GNEITING, T., GENTON, M. G. and GUTTORP, P. (2007). Geostatistical space–time models, stationarity, separability and full symmetry. In *Statistical Methods for Spatio-Temporal Systems*

- (B. Finkenstädt, L. Held and V. Isham, eds.). *Monographs on Statistics and Applied Probability* **107** 151–175. Chapman & Hall/CRC, Boca Raton.
- GNEITING, T. and RAFTERY, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. *J. Amer. Statist. Assoc.* **102** 359–378. [MR2345548](#)
- GNEITING, T., LARSON, K., WESTRICK, K., GENTON, M. G. and ALDRICH, E. (2006). Calibrated probabilistic forecasting at the stateline wind energy center: The regime-switching space–time method. *J. Amer. Statist. Assoc.* **101** 968–979. [MR2324108](#)
- HASTINGS, W. K. (1970). Monte Carlo sampling methods using Markov chains and their applications. *Biometrika* **57** 97–109.
- HERNÁNDEZ, A., GUENNI, L. and SANSÓ, B. (2009). Extreme limit distribution of truncated models for daily rainfall. *Environmetrics* **20** 962–980. [MR2838498](#)
- HUANG, H.-C. and HSU, N.-J. (2004). Modeling transport effects on ground-level ozone using a non-stationary space–time model. *Environmetrics* **15** 251–268.
- HUGHES, J. and GUTTORP, P. (1994). A class of stochastic models for relating synoptic atmospheric patterns to regional hydrologic phenomena. *Water Resources Research* **30** 1535–1546.
- HUGHES, J., GUTTORP, P. and CHARLES, S. P. (1999). A non-homogeneous hidden Markov model for precipitation occurrence. *J. R. Stat. Soc. Ser. C. Appl. Stat.* **48** 15–30.
- HUTCHINSON, M. (1995). Stochastic space–time weather models from ground-based data. *Agricultural and Forest Meteorology* **73** 237–264.
- ISHAM, V. and COX, D. R. (1994). Stochastic models of precipitation. In *Statistics for the Environment, Vol. 2* (V. Barnett and K. F. Turkman, eds.). Wiley, Chichester.
- JONES, R. H. and ZHANG, Y. (1997). Models for continuous stationary space–time processes. In *Statistical Methods for Spatio-Temporal Systems* (T. G. Gregoire, D. R. Brillinger, P. J. Diggle, E. Russek-Cohen, W. G. Warren and R. D. Wolfinger, eds.). *Lecture Notes in Statistics* **122** 289–298. Springer, New York.
- KOBER, K., CRAIG, G. C., KEIL, C. and DÖRNBRACK, A. (2012). Blending a probabilistic now-casting method with a high-resolution numerical weather prediction ensemble for convective precipitation forecasts. *Quarterly Journal of the Royal Meteorological Society* **138** 755–768.
- KÜNSCH, H. R. (2001). State space and hidden Markov models. In *Complex Stochastic Systems (Eindhoven, 1999)*. *Monographs on Statistics and Applied Probability* **87** 109–173. Chapman & Hall/CRC, Boca Raton, FL. [MR1893412](#)
- KYRIAKIDIS, P. C. and JOURNAL, A. G. (1999). Geostatistical space–time models: A review. *Math. Geol.* **31** 651–684. [MR1694654](#)
- LE CAM, L. (1961). A stochastic description of precipitation. In *Proc. 4th Berkeley Sympos. Math. Statist. and Prob., Vol. III* 165–186. Univ. California Press, Berkeley, CA. [MR0135598](#)
- LINDGREN, F., RUE, H. and LINDSTRÖM, J. (2011). An explicit link between Gaussian fields and Gaussian Markov random fields: The stochastic partial differential equation approach. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **73** 423–498. With discussion and a reply by the authors. [MR2853727](#)
- LITTLE, M. A., MCSHARRY, P. E. and TAYLOR, J. W. (2009). Generalized linear models for site-specific density forecasting of U.K. daily rainfall. *Monthly Weather Review* **137** 1029–1045.
- MA, C. (2003). Families of spatio-temporal stationary covariance models. *J. Statist. Plann. Inference* **116** 489–501. [MR2000096](#)
- MAKHININ, O. V. and MCALLISTER, D. L. (2009). Stochastic precipitation generation based on a multivariate autoregression model. *Journal of Hydrometeorology* **10** 1397–1413.
- MARDIA, K. V. and WATKINS, A. J. (1989). On multimodality of the likelihood in the spatial linear model. *Biometrika* **76** 289–295. [MR1016019](#)
- MASON, J. (1986). Numerical weather prediction. *Proc. R. Soc. Lond. Ser. A Math. Phys. Eng. Sci.* **407** 51–60.
- MATHESON, J. E. and WINKLER, R. L. (1976). Scoring rules for continuous probability distributions. *Manag. Sci.* **22** 1087–1096.

- METROPOLIS, N., ROSENBLUTH, A. W., ROSENBLUTH, M. N., TELLER, A. H. and TELLER, E. (1953). Equation of state calculations by fast computing machines. *J. Chem. Phys.* **21** 1087–1092.
- OKABE, A., BOOTS, B., SUGIHARA, K. and CHIU, S. N. (2000). *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams*, 2nd ed. Wiley, Chichester. [MR1770006](#)
- PACIOREK, C. J. and SCHERVISH, M. J. (2006). Spatial modelling using a new class of nonstationary covariance functions. *Environmetrics* **17** 483–506. [MR2240939](#)
- RUE, H. and HELD, L. (2005). *Gaussian Markov Random Fields: Theory and Applications. Monographs on Statistics and Applied Probability* **104**. Chapman & Hall/CRC, Boca Raton, FL. [MR2130347](#)
- SANSÓ, B. and GUENNI, L. (1999a). Venezuelan rainfall data analysed by using a Bayesian space-time model. *J. R. Stat. Soc. Ser. C. Appl. Stat.* **48** 345–362.
- SANSÓ, B. and GUENNI, L. (1999b). A stochastic model for tropical rainfall at a single location. *Journal of Hydrology* **214** 64–73.
- SANSÓ, B. and GUENNI, L. (2000). A nonstationary multisite model for rainfall. *J. Amer. Statist. Assoc.* **95** 1089–1100. [MR1821717](#)
- SANSÓ, B. and GUENNI, L. (2004). A Bayesian approach to compare observed rainfall data to deterministic simulations. *Environmetrics* **15** 597–612.
- SIGRIST, F., KÜNSCH, H. R. and STAHEL, W. A. (2012). An SPDE based spatio-temporal model for large data sets with an application to postprocessing precipitation forecasts. Preprint. Available at <http://arxiv.org/abs/1204.6118>.
- SLOUGHTER, J. M., RAFTERY, A. E., GNEITING, T. and FRALEY, C. (2007). Probabilistic quantitative precipitation forecasting using Bayesian model averaging. *Monthly Weather Review* **135** 3209–3220.
- SMITH, A. F. M. and ROBERTS, G. O. (1993). Bayesian computation via the Gibbs sampler and related Markov chain Monte Carlo methods. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **55** 3–23. [MR1210421](#)
- SØLNA, K. and SWITZER, P. (1996). Time trend estimation for a geographic region. *J. Amer. Statist. Assoc.* **91** 577–589. [MR1395727](#)
- STEHLIK, J. and BARDOSSY, A. (2002). Multivariate stochastic downscaling model for generating daily precipitation series based on atmospheric circulation. *Journal of Hydrology* **256** 120–141.
- STEIN, M. (1990). Uniform asymptotic optimality of linear predictions of a random field using an incorrect second-order structure. *Ann. Statist.* **18** 850–872. [MR1056340](#)
- STEIN, M. L. (2005). Space-time covariance functions. *J. Amer. Statist. Assoc.* **100** 310–321. [MR2156840](#)
- STERN, R. D. and COE, R. (1984). A model fitting analysis of daily rainfall data. *J. Roy. Statist. Soc. Ser. A* **147** 1–34.
- STIDD, C. K. (1973). Estimating the precipitation climate. *Water Resources Research* **9** 1235–1241.
- TOBIN, J. (1958). Estimation of relationships for limited dependent variables. *Econometrica* **26** 24–36. [MR0090462](#)
- VORONOI, G. (1908). Nouvelles applications des paramètres continus à la théorie des formes quadratiques. Deuxième mémoire. Recherches sur les paralléloèdres primitifs. *Journal Für die Reine und Angewandte Mathematik (Crelles Journal)* **1908** 198–287.
- WARNES, J. J. and RIPLEY, B. D. (1987). Problems with likelihood estimation of covariance functions of spatial Gaussian processes. *Biometrika* **74** 640–642. [MR0909370](#)
- WAYMIRE, E., GUPTA, V. K. and RODRIGUEZ-ITURBE, I. (1984). A spectral theory of rainfall intensity at the meso- β scale. *Water Resources Research* **20** 1453–1465.
- WEST, M. and HARRISON, J. (1997). *Bayesian Forecasting and Dynamic Models*, 2nd ed. Springer, New York. [MR1482232](#)
- WIKLE, C. K. and CRESSIE, N. (1999). A dimension-reduced approach to space-time Kalman filtering. *Biometrika* **86** 815–829. [MR1741979](#)

- WIKLE, C. K. and HOOTEN, M. B. (2010). A general science-based framework for dynamical spatio-temporal models. *TEST* **19** 417–451. [MR2745992](#)
- WILKS, D. (1990). Maximum likelihood estimation for the gamma distribution using data containing zeros. *Journal of Climate* **3** 1495–1501.
- WILKS, D. (1998). Multisite generalization of a daily stochastic precipitation generation model. *Journal of Hydrology* **210** 178–191.
- WILKS, D. (1999). Multisite downscaling of daily precipitation with a stochastic weather generator. *Climate Research* **11** 125–136.
- XU, K., WIKLE, C. K. and FOX, N. I. (2005). A kernel-based spatio-temporal dynamical model for nowcasting weather radar reflectivities. *J. Amer. Statist. Assoc.* **100** 1133–1144. [MR2236929](#)
- ZHANG, H. (2004). Inconsistent estimation and asymptotically equal interpolations in model-based geostatistics. *J. Amer. Statist. Assoc.* **99** 250–261. [MR2054303](#)

SEMINAR FOR STATISTICS
DEPARTMENT OF MATHEMATICS
ETH ZURICH
SWITZERLAND
E-MAIL: sigrist@stat.math.ethz.ch
kuensch@stat.math.ethz.ch
stahel@stat.math.ethz.ch