



Spatial correlation in weather forecast accuracy: a functional time series approach

Phillip A. Jang¹  · David S. Matteson²

Received: 15 April 2019 / Accepted: 22 February 2023 / Published online: 14 March 2023

© The Author(s), under exclusive licence to Springer-Verlag GmbH Germany, part of Springer Nature 2023

Abstract

A functional time series approach is proposed for investigating spatial correlation in daily maximum temperature forecast errors for 111 cities spread across the U.S. The modelling of spatial correlation is most fruitful for longer forecast horizons, and becomes less relevant as the forecast horizon shrinks towards zero. For 6-day-ahead forecasts, the functional approach uncovers interpretable regional spatial effects, and captures the higher variance observed in inland cities versus coastal cities, as well as the higher variance observed in mountain and midwest states. The functional approach also naturally handles missing data through modelling a continuum, and can be implemented efficiently by exploiting the sparsity induced by a B-spline basis. The temporal dependence in the data is modeled through temporal dependence in functional basis coefficients. Independent first order autoregressions with generalized autoregressive conditional heteroskedasticity [AR(1)+GARCH(1,1)] and Student-t innovations work well to capture the persistence of basis coefficients over time and the seasonal heteroskedasticity reflecting higher variance in winter. Through exploiting autocorrelation in the basis coefficients, the functional time series approach also yields a method for improving weather forecasts and uncertainty quantification. The resulting method corrects for bias in the weather forecasts, while reducing the error variance.

Keywords Functional data analysis · Spatiotemporal · Random effects · GARCH · Weather forecast data · Data Expo 2018

✉ Phillip A. Jang
phillipjang23@hotmail.com

¹ Center for Applied Mathematics, Cornell University, 657 Frank H.T. Rhodes Hall, Ithaca, NY 14853, USA

² Department of Statistics and Data Science, Cornell University, 1196 Comstock Hall, Ithaca, NY 14853, USA

1 Introduction

A *functional time series* is a time-indexed sequence of stochastic processes $\{f_i(\tau)\}_{i=1}^{\infty}$ where each $f_i(\cdot)$ is a random function on the domain \mathcal{T} . By unifying functional data analysis with time series analysis, it presents an approach to modelling randomness on curves, surfaces, and other phenomena varying over a spatial continuum where these functional data are observed regularly over time and exhibit serial dependence. Ramsay and Silverman (2005) and Tsay (2010) provide background for functional data analysis and time series analysis respectively, while (Hörmann and Kokoszka 2012) provides background for functional time series.

Hyndman and Shahid Ullah (2007) and Hyndman and Booth (2008) propose a forecasting approach for functional time series and apply it to demographic data. Hyndman and Shang (2009) proposes a weighted functional approach which assigns more weight to recent observations and yields an improvement in forecast accuracy. Further developments in functional forecasting with applications to demographic data are found in Shang et al. (2011), Hyndman et al. (2013), and Dokumentov et al. (2018). Aue et al. (2017) develops the theory for functional GARCH models. Kowal et al. (2017) introduces a Bayesian framework for functional time series, and Kowal et al. (2019) develops functional autoregression for sparsely and irregularly sampled data.

The most common examples of functional data are samples of 1-dimensional curves, and a contribution of this paper is the analysis of a model for 2-dimensional surface data. In this application, the spatial domain \mathcal{T} is a rectangle in \mathbb{R}^2 containing the range of longitudes and latitudes covering the lower 48 states. The data set is taken from the 2018 American Statistical Association Data Expo,¹ consisting of daily maximum temperature forecasts from the National Weather Service for 111 cities spread across the US (excluding Alaska and Hawaii) over the period from July 2014 to September 2017. The locations of the cities are illustrated in Fig. 1. Forecasts range from same-day to six-days-ahead and are compared to actual temperature recorded at city airports.

A functional time series approach is applied to investigate and extract the structure of the spatial correlation in forecast errors. Through modelling the entire continuum instead of individual points, the proposed functional data approach also naturally handles missing data. This is a vital benefit, as data records are frequently incomplete and forecasts are not always available at every location. This paper also extends the methodology of Hyndman and Shahid Ullah (2007) with the incorporation of a heteroskedastic time series model to account for the higher unpredictability of weather in winters compared to summers.

Through exploiting both the spatial and temporal dependence in weather forecast errors, one can estimate the next day's expected forecast error and correct the forecast accordingly.

¹ <http://www2.amstat.org/sections/graphics/datasets/DataExpo2018.zip>.

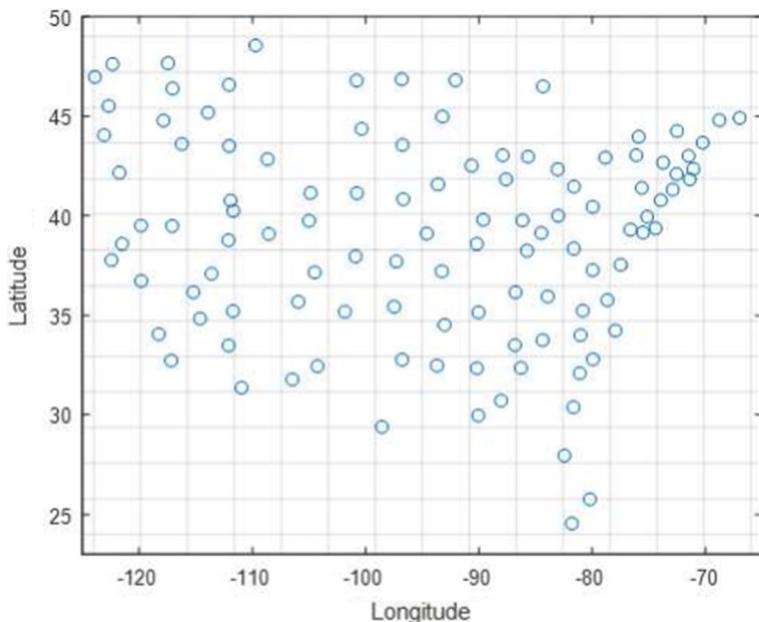


Fig. 1 Locations of cities included in weather forecast data shown in blue. The grey grid lines indicate the locations of the knots chosen for the B-spline bases

In this paper, a four-step procedure is used to facilitate this estimation. Section 2.1 describes the initial pre-processing step, where we first construct continuous surface data from the discrete observations using manually selected B-splines. The resolution of the chosen B-splines are set so that cities are able to be grouped into regions of interest. Once these surfaces are built, the second step is to reduce dimension to a smaller spatial basis which describes the most important modes of variation, and this is described in Sect. 2.2. After the spatial basis is chosen, the next step is to model the temporal dependence in the random coefficients which scale the basis functions, and this is described in Sect. 2.3. The time series model facilitates the final step, the prediction of forecast errors, and this is described in Sect. 3.1.

2 Methodology

We first fix a forecast horizon h of interest, in this case between 0 to 6 days ahead. For the given forecast horizon, let $Y_t(\tau)$ be the forecast error of maximum temperature on day t for the city located at spatial coordinates $\tau \equiv (\text{longitude}, \text{latitude})$. Spatial correlation is captured through the following spatio-temporal random effect model:

$$Y_t(\tau) = \mu(\tau) + \sum_{k=1}^K \beta_{kt} \varphi_k(\tau) + \varepsilon_t(\tau). \quad (1)$$

Here, $\mu(\tau)$ represents the mean forecast error and is assumed fixed over time. Estimation of the mean function is described in Sect. 2.1. The spatial basis functions $\varphi_1(\tau), \dots, \varphi_K(\tau)$ describe the main modes of variation in the forecast errors, capturing spatial dependence across different regions of the U.S. The construction of the basis functions is described in Sects. 2.1 and 2.2. The random coefficients β_{kt} capture temporal correlation. They are modelled using independent AR(1)+GARCH(1,1) processes, as described in Sect. 2.3. Lastly, $\varepsilon_t(\tau)$ is a white noise process independent of the random coefficients, and is assumed i.i.d. $N(0, \sigma^2)$ for all t and τ . It is used to capture, including any measurement errors, the remaining variation not explained by the spatio-temporal random effect.

Writing $\Phi(\tau) = [\varphi_1(\tau) \cdots \varphi_K(\tau)]^T$ and $\beta_t = [\beta_{1t} \cdots \beta_{Kt}]^T$, this model implies the following spatio-temporal covariance function:

$$\text{Cov}(Y_t(\tau), Y_{t'}(\tau')) = \Phi(\tau)^T \mathbb{E}[\beta_t \beta_{t'}^T] \Phi(\tau') + \sigma^2 \mathbb{1}_{\{t=t', \tau=\tau'\}}, \quad (2)$$

where the indicator function $\mathbb{1}_{\{t=t', \tau=\tau'\}}$ is 1 when both $t = t'$ and $\tau = \tau'$, and 0 otherwise.

2.1 Pre-processing with spline smoothing

A two-step procedure is employed to construct the spatial basis functions φ_k . As an initial pre-processing step, we first need to construct continuous surface data from the discretely observed values using manually selected B-splines, where the knot sequence is at the right level of tightness to capture regional groupings of the cities. Once these surfaces are built, the singular value decomposition (SVD) is used in Sect. 2.2 to reduce dimension and identify the φ_k which describe the most important modes of variation.

For simplicity, we consider the spatial domain as a subset of \mathbb{R}^2 . Specifically, we define a 2-D cubic B-spline basis over the rectangle $[-124, -66] \times [24, 49]$, which contains the range of longitudes and latitudes covering the lower 48 states. The 2-D splines are built from the tensor product of 1-D cubic B-splines on longitude and latitude individually. Refer to de Boor (2001) for further background on B-splines. The MATLAB package ‘bspline’ by Hunyadi (2020) is used to implement the splines.

For this particular dataset, knot sequences with 13 equally spaced interior knots were able to capture interesting regional groupings of the cities, and the knots are visualized against the cities by the grey grid lines in Fig. 1. This results in 17 cubic B-splines in each dimension, and 289 2-D splines in the resulting tensor product, denoted by $S_1(\tau), \dots, S_{289}(\tau)$, with

$$\mathcal{K}^{\text{Lon}} = [-124, -124, -124, -124, -119.86, -115.71, \dots, -66, -66, -66, -66],$$

$$\mathcal{K}^{\text{Lat}} = [24, 24, 24, 24, 25.79, 27.57, \dots, 49, 49, 49, 49], \text{ and}$$

$$\begin{bmatrix} S_1(lon, lat) \\ \vdots \\ S_{289}(lon, lat) \end{bmatrix} = \begin{bmatrix} B_{1,K^{\text{Lon}}}(lon) \\ \vdots \\ B_{17,K^{\text{Lon}}}(lon) \end{bmatrix} \otimes \begin{bmatrix} B_{1,K^{\text{Lat}}}(lat) \\ \vdots \\ B_{17,K^{\text{Lat}}}(lat) \end{bmatrix}.$$

For each day $t = 1, \dots, m$, we fit splines such that the coefficients $\hat{c}_{t,1}, \dots, \hat{c}_{t,289}$ solve the following optimization problem:

$$\min_{c_{t,1}, \dots, c_{t,289}} \sum_{i=1}^{n_t} \left[Y_t(\tau_{t,i}) - \sum_{j=1}^{289} c_{t,j} S_j(\tau_{t,i}) \right]^2, \quad (3)$$

where $\tau_{t,1}, \dots, \tau_{t,n_t}$ are the observation locations available on day t . This handles (moderate amounts of) missing data naturally since missing observations $Y_t(\tau)$ are simply omitted from the objective function. Also, since B-splines have compact support, the resulting system is sparse and can be solved efficiently.

As the number of coefficients exceeds the number of observation locations, a truncated singular value decomposition is used to solve for the coefficients. The truncated SVD acts as a form of regularization similar to ridge regression, and further details can be found in Hansen (1987). In Fig. 1, there is a lack of observations in regions of the rectangle outside the borders of the United States. The truncated SVD yields stable results by imposing zero coefficients for the splines that lie in these data-free regions.

We denote the resulting coefficient matrix by $C = [\hat{c}_{t,j}]$. We use the column means $\bar{C} = [\bar{c}_1, \dots, \bar{c}_{289}]$ to estimate the mean function $\mu(\tau)$ as:

$$\hat{\mu}(\tau) = \bar{Y} + \sum_{j=1}^{289} \bar{c}_j S_j(\tau) \quad (4)$$

where $\bar{Y} = (\sum_{t=1}^m n_t)^{-1} \sum_{t=1}^m \sum_{i=1}^{n_t} Y_t(\tau_{t,i})$ is the sample mean of all observed forecast errors.

2.2 Constructing spatial basis functions φ_k

As each spline function captures variation specific to a sub-rectangle of the domain, the amount of information in each region of the domain is encoded through the coefficients of its corresponding splines. Through studying the covariance of the B-spline coefficients, we are able to piece together variation across multiple sub-rectangles to study the important regional effects in the weather forecast errors. A principal component analysis yields a ranked order of groupings through the eigen-decomposition of the covariance matrix, or equivalently, the singular value decomposition of the mean-centered data matrix. The eigenvectors of the covariance matrix identify weightings of the spline coefficients which contribute the most variation, with said variation quantified by the size of the eigenvalue.

After mean-centering the columns of C , the singular value decomposition $C - \bar{C} = U\Sigma V^T$ provides principal component loadings as the columns of V , allowing for dimension reduction that is mean-square optimal in the coefficient space

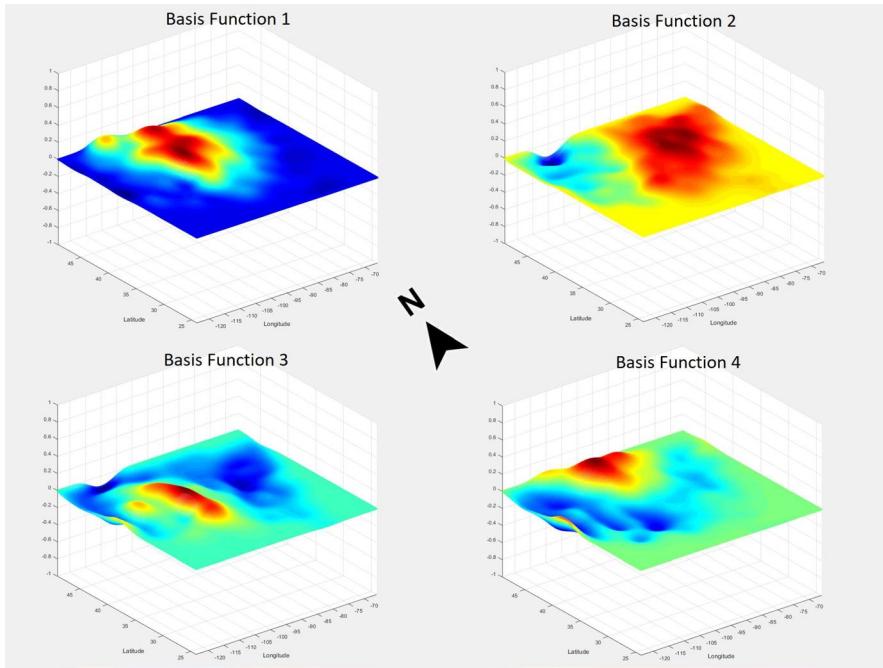


Fig. 2 First 4 basis functions for 6-day-ahead forecasts

$\text{Col}(C - \bar{C})$. Refer to Jolliffe (2002) for further background on principal component analysis and singular value decomposition. Assuming the SVD is written in descending order of singular values, the spatial basis functions $\varphi_k(\tau)$ are built using the first K columns of V , representing the K most important principal components:

$$\varphi_k(\tau) = \sum_{j=1}^{289} V_{jk} S_j(\tau) \text{ for } k = 1, \dots, K \ll 289. \quad (5)$$

The first four basis functions are shown in Fig. 2 for the 6-day-ahead forecast horizon. The first basis function represents an inland versus coastal effect, as cities further inland had greater variance in their forecast errors compared to coastal cities. The second basis function represents an east versus west effect, with the opposing signs of the regions allowing for a differentiation between the regions. The third and fourth basis functions represent mountain state and midwest state effects respectively, as these regions have the most unpredictable weather.

This structure of spatial correlation is most prevalent in 6-day-ahead forecasts, but vanishes as the forecast horizon shrinks to zero. For example, Fig. 3 shows the first four basis functions for same-day forecasts, and these basis functions lack any coherent spatial structure.

For the 6-day-ahead forecasts, a cutoff of $K = 20$ basis functions is used. Each of the 20 selected basis functions ranged from 9 to 1.4% of explained variance, while the basis functions after the 20th each contribute less than 1.4% of explained

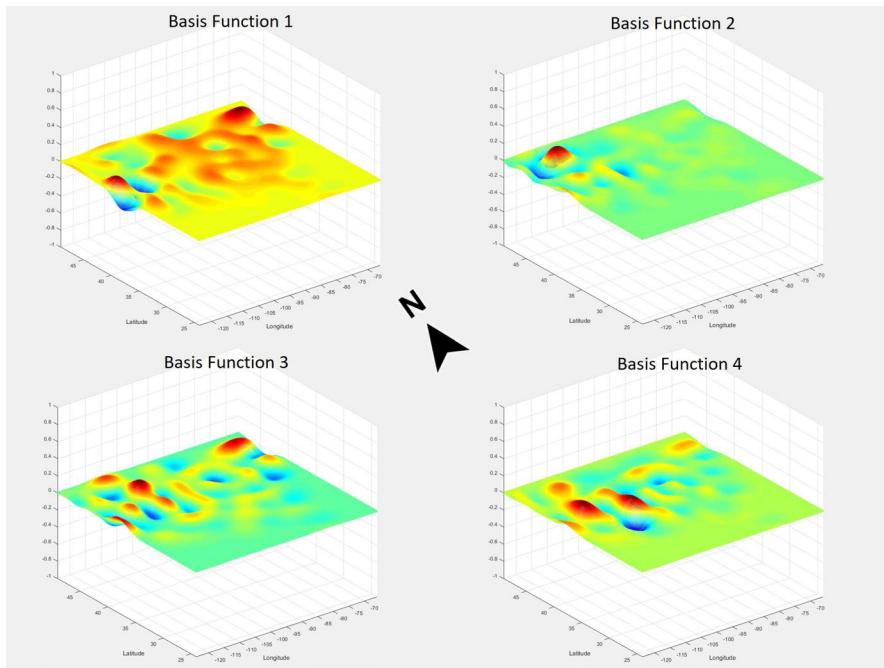


Fig. 3 First 4 basis functions for same-day (0-day-ahead) forecasts

variance and revealed no interesting spatial structure. The results of Table 2 were stable above $K = 20$ basis functions, justified by Fig. 7.

2.3 Modelling random coefficients

Once the spatial basis functions are defined as above, coefficients $\beta_{1t}, \dots, \beta_{Kt}$ for each day t are estimated on this reduced K -dimensional basis as solutions to the following optimization problem:

$$\min_{\beta_{1t}, \dots, \beta_{Kt}} \sum_{i=1}^{n_t} \left[Y_t(\tau_{t,i}) - \hat{\mu}(\tau_{t,i}) - \sum_{k=1}^K \beta_{kt} \varphi_k(\tau_{t,i}) \right]^2. \quad (6)$$

Temporal dependence in forecast errors is then modelled through these coefficients. Empirically, an AR(1)+GARCH(1,1) model with Student-t innovations was found to provide a good description of the observed coefficients. Specifically,

$$\beta_{kt} = \psi_k \beta_{k,t-1} + u_{kt}, \quad u_{kt} | \mathcal{F}_{t-1} \sim t_{\nu_k}(0, \eta_{kt}^2) \quad (7)$$

$$\eta_{kt}^2 = \omega_k + \alpha_k u_{k,t-1}^2 + \gamma_k \eta_{k,t-1}^2 \quad (8)$$

where \mathcal{F}_{t-1} is the information set up to time $t-1$ (the σ -field generated by $u_{k,t-1}, u_{k,t-2}, \dots$ for all k). The u_{1t}, \dots, u_{Kt} are assumed conditionally independent given \mathcal{F}_{t-1} . The proposed model implies the following conditional mean and covariance functions:

$$\mathbb{E}[Y_t(\tau)|\mathcal{F}_{t-1}] = \mu(\tau) + \sum_{k=1}^K \psi_k \beta_{k,t-1} \varphi_k(\tau) \quad (9)$$

$$\begin{aligned} \text{Cov}(Y_t(\tau), Y_t(\tau')|\mathcal{F}_{t-1}) &= \sum_{k=1}^K \frac{\nu_k}{\nu_k - 2} \left[\omega_k + \alpha_k u_{k,t-1}^2 + \gamma_k \eta_{k,t-1}^2 \right] \varphi_k(\tau) \varphi_k(\tau') \\ &\quad + \sigma^2 \mathbb{1}_{\{\tau=\tau'\}}. \end{aligned} \quad (10)$$

The resulting innovations for the first basis function u_{1t} are shown in Fig. 4 on the top left. Notably, the innovations exhibit a seasonal heteroskedasticity with winter weather being the most unpredictable. The GARCH process characterizes the heteroskedasticity well, as the standardized innovations u_{1t}/η_{1t} exhibit approximately constant variance, and the squared standardized innovations show no significant autocorrelation.

For the 6-day-ahead weather forecasts, the AR+GARCH parameter estimates for the first four basis functions are shown in Table 1. All fitted models were stationary with similar amounts of autocorrelation based on the similar values of ψ_k . Furthermore, all exhibit a high persistence in variance, indicated by the large values of γ_k . The earlier basis functions had conditional distributions with heavier tails, indicated by lower values of ν_k .

3 Empirical performance

For each pair of cities (i, j) , sample spatial correlation for the 6-day-ahead forecasts is computed before and after accounting for the spatial basis functions. More specifically, given a pair of cities located at τ_i and τ_j , the top of Fig. 5 shows

$$\rho_{i,j}^{\text{before}} = \frac{\sum_t [Y_t(\tau_i) - \bar{Y}_i(\tau_i)][Y_t(\tau_j) - \bar{Y}_j(\tau_j)]}{\sqrt{\sum_t [Y_t(\tau_i) - \bar{Y}_i(\tau_i)]^2 \sum_t [Y_t(\tau_j) - \bar{Y}_j(\tau_j)]^2}}, \quad (11)$$

and the bottom of Fig. 5 shows

$$\rho_{i,j}^{\text{after}} = \frac{\sum_t [\hat{\varepsilon}_i(\tau_i) - \bar{\hat{\varepsilon}}_i(\tau_i)][\hat{\varepsilon}_i(\tau_j) - \bar{\hat{\varepsilon}}_i(\tau_j)]}{\sqrt{\sum_t [\hat{\varepsilon}_i(\tau_i) - \bar{\hat{\varepsilon}}_i(\tau_i)]^2 \sum_t [\hat{\varepsilon}_i(\tau_j) - \bar{\hat{\varepsilon}}_i(\tau_j)]^2}}, \quad (12)$$

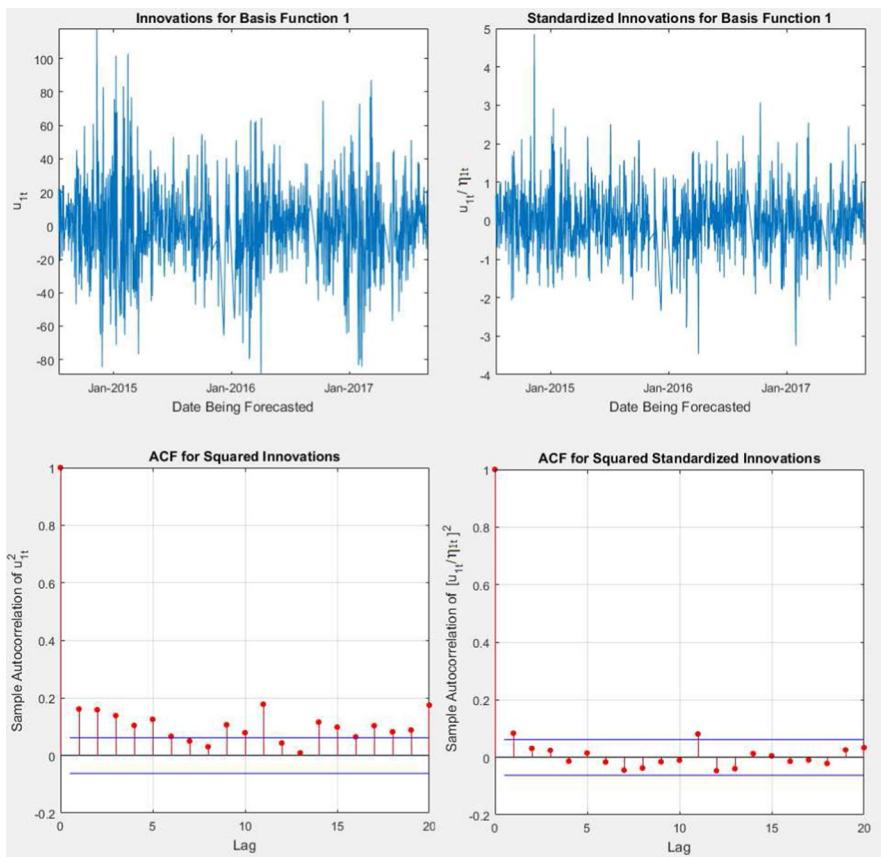


Fig. 4 Estimated innovations u_{1t} (top left) and standardized innovations u_{1t}/η_{1t} (top right) for basis function 1 at forecast horizon 6-days-ahead, and their respective sample autocorrelation functions (bottom, left and right, respectively). Approximate 95% pointwise confidence intervals are drawn as horizontal lines in the autocorrelation plots

for the residuals $\hat{\varepsilon}_t(\tau) = Y_t(\tau) - \hat{\mu}(\tau) - \sum_{k=1}^K \beta_{kt} \varphi_k(\tau)$. Above, the sums are over the days t with no missing observations, and $\bar{Y}_t(\tau_i)$ and $\bar{\varepsilon}_t(\tau_i)$ are sample averages of $Y_t(\tau_i)$ and $\hat{\varepsilon}_t(\tau_i)$ respectively over such t .

The cities are numbered from 1 to 111 (shown on the x - and y -axes), and cities are ordered from east to west, resulting in a concentration of high correlation along the main diagonal in the first figure. After accounting for $K = 20$ spatial basis functions, the second figure indicates the lack of spatial correlation in the residuals and provides evidence that the proposed model provides an adequate approximation of the observed spatial correlation in forecast errors.

The correlograms before and after accounting for the basis functions are shown in Fig. 6, visualizing the sample spatial correlations against city distance. The cutoff of $K = 20$ is justified in Fig. 7, where the squared Frobenius norm of the residual spatial correlation matrix ($\sum_{i,j} (\rho_{ij}^{\text{after}})^2$) is computed for different values of

Table 1 AR(1)+GARCH(1,1) parameter estimates for first four basis functions for 6-day-ahead forecasts

Basis Function 1	ψ_1	ω_1	α_1	γ_1	ν_1
Estimate	0.65	13.50	0.09	0.89	8.33
Approx. Std. Error	0.02	6.91	0.02	0.02	2.05
t-ratio	26.91	1.95	4.25	37.14	4.06
p-value	<0.0001	0.0507	<0.0001	<0.0001	<0.0001
Basis Function 2	ψ_2	ω_2	α_2	γ_2	ν_2
Estimate	0.57	18.67	0.05	0.92	10.90
Approx. Std. Error	0.03	11.43	0.02	0.03	3.95
t-ratio	21.08	1.63	3.00	34.18	2.76
p-value	<0.0001	0.1024	0.0027	<0.0001	0.0057
Basis Function 3	ψ_3	ω_3	α_3	γ_3	ν_3
Estimate	0.53	7.11	0.03	0.96	13.31
Approx. Std. Error	0.03	6.46	0.01	0.02	4.87
t-ratio	19.95	1.10	2.32	47.14	2.73
p-value	<0.0001	0.2709	0.0203	<0.0001	0.0063
Basis Function 4	ψ_4	ω_4	α_4	γ_4	ν_4
Estimate	0.57	13.01	0.06	0.91	14.75
Approx. Std. Error	0.03	7.56	0.02	0.03	6.97
t-ratio	21.21	1.72	3.50	31.04	2.11
p-value	<0.0001	0.0853	0.0005	<0.0001	0.0344

K . The Frobenius norm stops significantly decreasing after $K = 20$ indicating that additional basis functions do not explain any more spatial correlation, and capture individual city variation rather than regional variation.

3.1 Predicting forecast errors to improve forecast accuracy

Through exploiting autocorrelation in the coefficients, the AR+GARCH parameter estimates can be used to predict the next day's basis coefficients, as

$$\beta_{k,t+1} | \mathcal{F}_t \sim t_{\nu_k}(\psi_k \beta_{kt}, \eta_{k,t+1}^2), \text{ and} \quad (13)$$

$$\eta_{k,t+1}^2 = \omega_k + \alpha_k u_{kt}^2 + \gamma_k \eta_{kt}^2. \quad (14)$$

Substituting the parameter estimates $\hat{\psi}_k$, $\hat{\omega}_k$, $\hat{\alpha}_k$, $\hat{\gamma}_k$, and $\hat{\nu}_k$ in place of the true parameters yields an approximate distribution which can be used for prediction and uncertainty quantification.

By setting the predicted coefficient to $\hat{\beta}_{k,t+1} = \hat{\psi}_k \beta_{kt}$, this can then be used to predict next day's weather forecast errors by setting

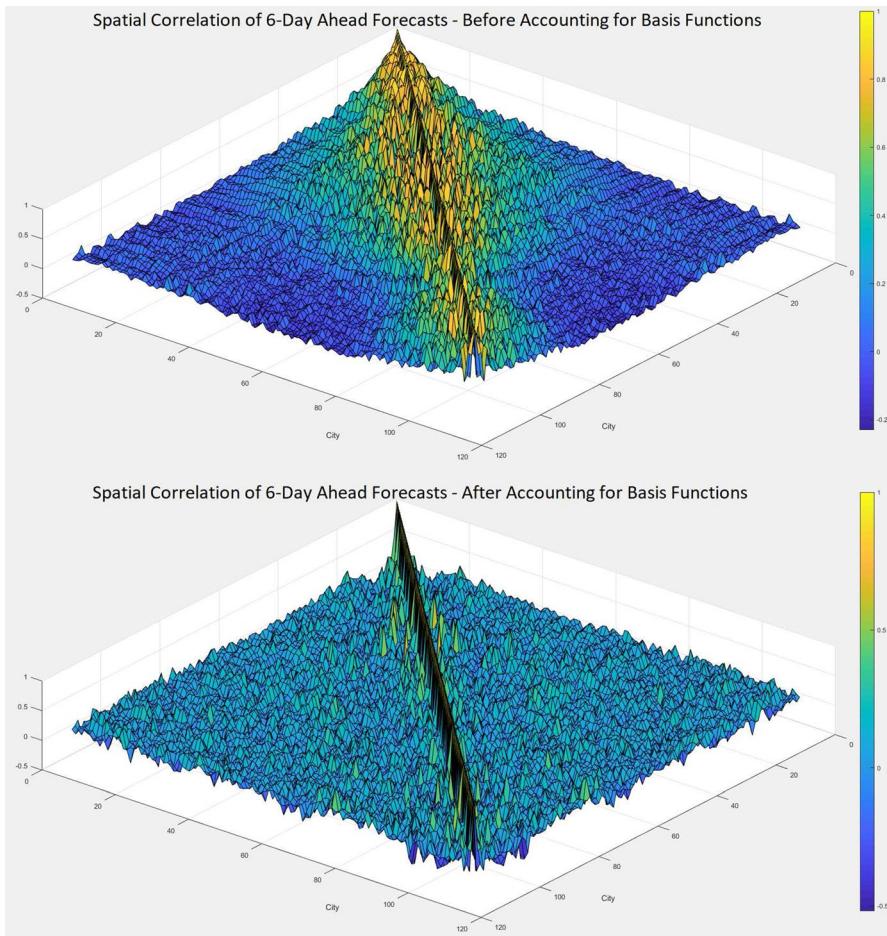


Fig. 5 Spatial correlation in 6-day-ahead forecast errors for before (top) and after (bottom) accounting for basis functions

$$\hat{Y}_{t+1}(\tau) = \hat{\mu}(\tau) + \sum_{k=1}^K \hat{\beta}_{k,t+1} \varphi_k(\tau). \quad (15)$$

These predicted errors $\hat{Y}_{t+1}(\tau)$ can then be used to adjust the next day's weather forecast (of the same horizon). More explicitly, for forecast horizon h and location τ , if $F_{t,t+h}(\tau)$ is the forecast on day t of the maximum temperature on day $t+h$ and $A_{t+h}(\tau)$ is the actual temperature on day $t+h$ so that $Y_t(\tau) = F_{t,t+h}(\tau) - A_{t+h}(\tau)$, we can define an adjusted forecast for day $t+1$ as

$$F_{t+1,t+1+h}^{adj}(\tau) = F_{t+1,t+1+h}(\tau) - \hat{Y}_{t+1}(\tau). \quad (16)$$

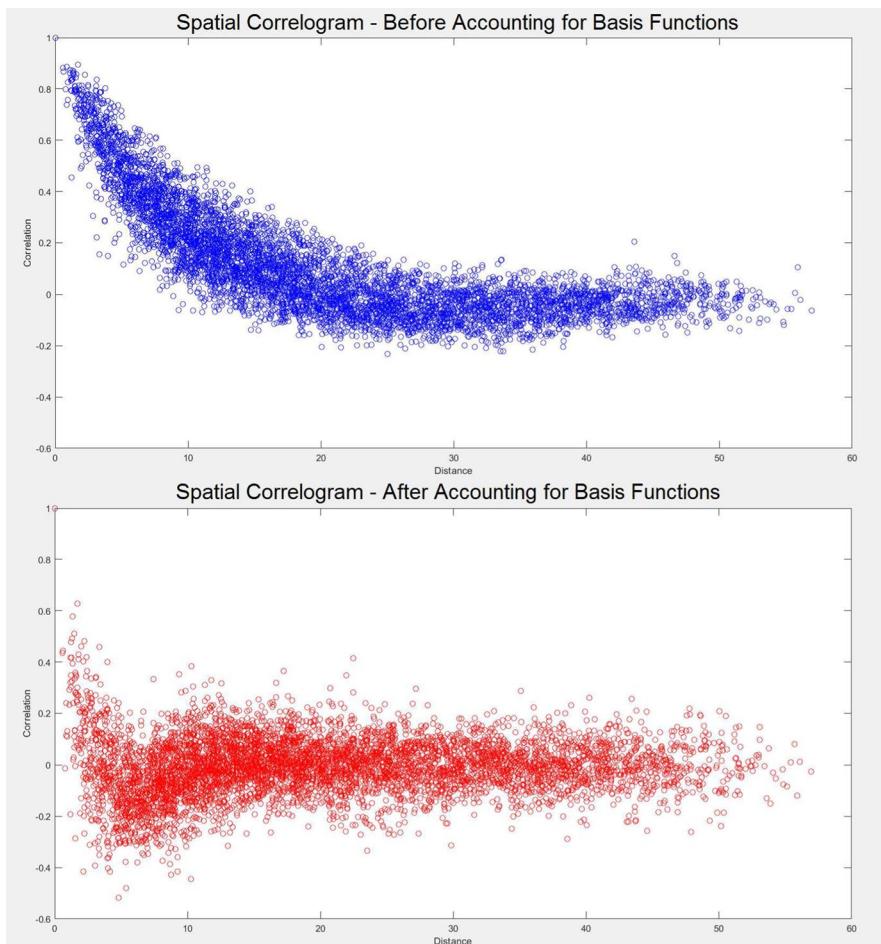


Fig. 6 Correlogram of 6-day-ahead forecast errors for before (top) and after (bottom) accounting for basis functions

The improvement on weather forecasts can then be assessed by comparing the adjusted errors

$$Z_{t+1}(\tau) = F_{t+1,t+1+h}^{adj}(\tau) - A_{t+1+h}(\tau) \quad (17)$$

to the unadjusted errors $Y_{t+1}(\tau)$.

The mean and variance of Y and Z across all t and τ is shown in Table 2.

The change in mean error indicates that the proposed method eliminates bias in the weather forecasts, while the reduction of standard deviation by 12% indicates a significant average improvement in forecast accuracy. The shift in error distribution is illustrated in Fig. 8, with the adjusted distribution centered at 0 and with smaller variance.

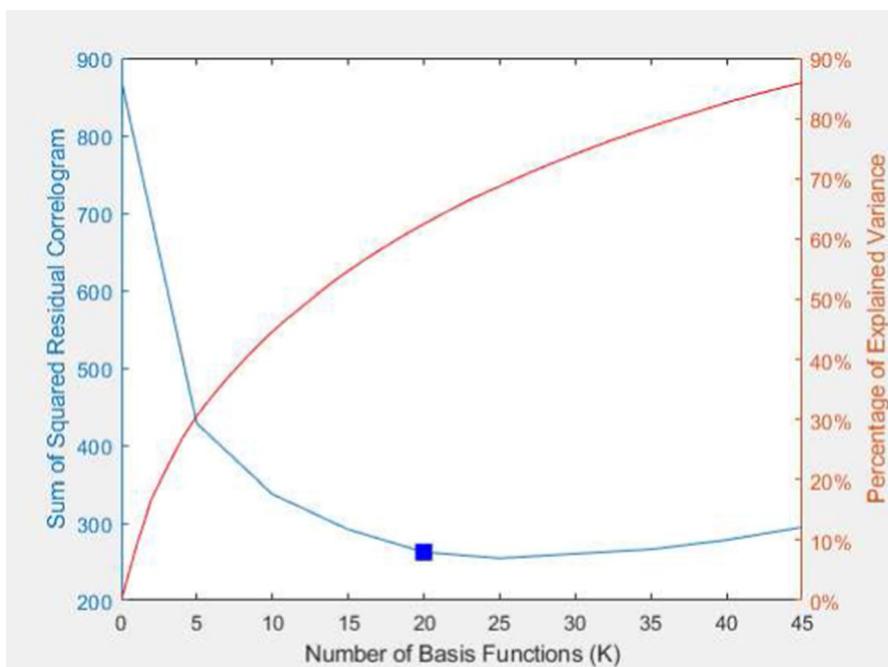


Fig. 7 Sum of squared residual correlogram $\sum_{i,j}(\rho_{ij}^{\text{after}})^2$ (blue) for a changing number of basis functions K . After $K = 20$ basis functions, the Frobenius norm of the residual spatial correlation matrix stops decreasing, indicating the remaining basis functions capture local variation rather than regional variation

4 Code

For reproducing our analysis and figures, MATLAB code is available at: <https://github.com/pjang23/dataexpo2018-functionaltimeseries>.

To accommodate those without the MATLAB Econometrics Toolbox™, the MATLAB package ‘ARMAX-GARCH-K-SK’ by Gabrielsen (2021) is included as an alternative for fitting the AR+GARCH models.

5 Conclusions

We have introduced a functional time series approach to investigating spatial correlation in weather forecast accuracy. The modelling of spatial correlation is most fruitful for the longer forecast horizons, and becomes less relevant as the forecast

Table 2 Comparison of 6-day ahead weather forecast errors before and after adjustment using $K = 20$ basis functions

	$Y_t(\tau)$	$Z_t(\tau)$
Mean	-1.17	0.00
Standard deviation	6.61	5.82

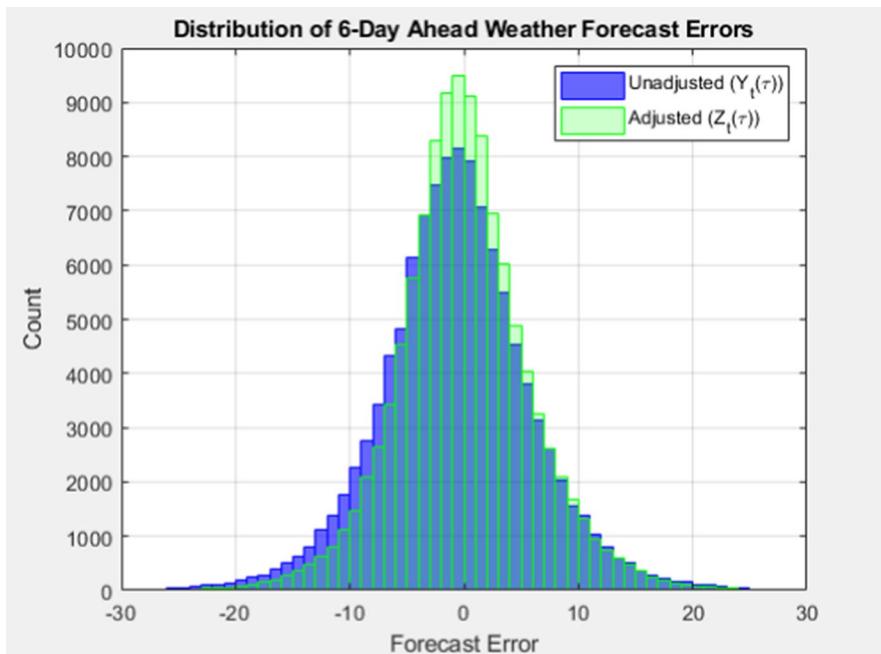


Fig. 8 Comparison of distributions of 6-day ahead weather forecast errors before and after adjustment using $K = 20$ basis functions

horizon shrinks towards zero. For 6-day-ahead weather forecasts, the functional approach uncovers interpretable regional spatial effects, and captures the higher variance observed in inland cities versus coastal cities, as well as the higher variance observed in mountain and midwest states. The functional approach also naturally handles missing data and can be implemented efficiently by exploiting the sparsity induced by using a B-spline basis.

Independent first order autoregressions with generalized autoregressive conditional heteroskedasticity [AR(1)+GARCH(1,1)] and Student-t innovations worked well to capture the persistence of coefficients over time and the seasonal heteroskedasticity reflecting higher variance in winter. Autocorrelation in the basis coefficients can further be exploited to improve weather forecasts, especially at longer horizons, and the resulting approach eliminates bias while reducing error variance.

Acknowledgements This work is supported in part by the Natural Sciences and Engineering Research Council of Canada (PGS-D 502888), the National Science Foundation (DMS-1455172), a Xerox PARC

Faculty Research Award, the United States Agency for International Development (USAID), and Cornell University Atkinson Center for a Sustainable Future (AVF-2017).

References

Aue A, Horváth L, Pellatt DF (2017) Functional generalized autoregressive conditional heteroskedasticity. *J Time Ser Anal* 38(1):3–21

de Boor C (2001) A practical guide to splines. Applied Mathematical Sciences, Springer, New York

Dokumentov A, Hyndman R, Tickle L (2018) Bivariate smoothing of mortality surfaces with cohort and period ridges. *Stat* 7(1):1–13. <https://doi.org/10.1002/sta4.199>

Gabrielsen A (2021) ARMAX-GARCH-K-SK toolbox (estimation, forecasting, simulation and Value-at-Risk applications). <https://www.mathworks.com/matlabcentral/fileexchange/32882-armax-garch-k-sk-toolbox-estimation-forecasting-simulation-and-value-at-risk-applications>, MATLAB Central File Exchange. Retrieved July 12, 2021

Hansen PC (1987) The truncated SVD as a method for regularization. *BIT Numer Math* 27(4):534–553. <https://doi.org/10.1007/bf01937276>

Hörmann S, Kokoszka P (2012) Functional time series. In: Time series analysis: methods and applications, handbook of statistics, vol 30, Elsevier, pp 157–186

Hunyadi L (2020) B-splines. <https://www.mathworks.com/matlabcentral/fileexchange/27374-b-splines>

Hyndman R, Booth H (2008) Stochastic population forecasts using functional data models for mortality, fertility and migration. *Int J Forecast* 24(3):323–342

Hyndman R, Shahid Ullah M (2007) Robust forecasting of mortality and fertility rates: a functional data approach. *Comput Stat Data Anal* 51(10):4942–4956

Hyndman R, Shang H (2009) Forecasting functional time series. *J Korean Stat Soc* 38(3):199–211. <https://doi.org/10.1016/j.jkss.2009.06.002>

Hyndman RJ, Booth H, Yasmeen F (2013) Coherent mortality forecasting: the product-ratio method with functional time series models. *Demography* 50(1):261–283

Jolliffe I (2002) Principal component analysis. Springer Series in Statistics. Springer, USA

Kowal DR, Matteson DS, Ruppert D (2017) A bayesian multivariate functional dynamic linear model. *J Am Stat Assoc* 112(518):733–744

Kowal DR, Matteson DS, Ruppert D (2019) Functional autoregression for sparsely sampled data. *J Bus Econ Stat* 37(1):97–109

Ramsay J, Silverman BW (2005) Functional data analysis. Springer Series in Statistics, Springer, USA

Shang HL, Booth H, Hyndman R (2011) Point and interval forecasts of mortality rates and life expectancy: a comparison of ten principal component methods. *Demogr Res* 25:173–214. <https://doi.org/10.4054/DemRes.2011.25.5>

Tsay RS (2010) Analysis of financial time series. John Wiley & Sons, UK

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.