Spatial autoregressive fractionally integrated moving average model

Philipp Otto
University of Glasgow, United Kingdom

Philipp Sibbertsen
Leibniz University Hannover, Germany

September 13, 2023

Abstract

In this paper, we introduce the concept of fractional integration for spatial autoregressive models. We show that the range of the dependence can be spatially extended or diminished by introducing a further fractional integration parameter to spatial autoregressive moving average models (SARMA). This new model is called the spatial autoregressive fractionally integrated moving average model, briefly sp-ARFIMA. We show the relation to time-series ARFIMA models and also to (higher-order) spatial autoregressive models. Moreover, an estimation procedure based on the maximum-likelihood principle is introduced and analysed in a series of simulation studies. Eventually, the use of the model is illustrated by an empirical example of atmospheric fine particles, so-called aerosol optical thickness, which is important in weather, climate and environmental science.

Keywords: Spatial ARFIMA, spatial fractional integration, long-range dependence, aerosol optical depth.

1 Introduction

Long memory of time series is a well-studied problem in statistics (see, e.g., Beran 2017 for an overview). A process is called to have long memory if the temporal autocorrelation is rather slowly decreasing, e.g. compared to autoregressive processes. For instance, consider a fractional Gaussian noise with $H = d + 0.5$, which coincides with an ARFIMA(0,d,0) process

$$(1 - B)^d Y_t = \varepsilon_t \,,$$

where B denotes the backshift operator. This process has temporal long memory. For finite samples $Y_1, \ldots, Y_T$, the model can be rewritten in a vector notation as follows

$$(\mathbf{I} - \mathbf{B})^d \boldsymbol{Y} = \boldsymbol{\varepsilon}$$

with $\boldsymbol{Y} = (Y_t)_{t=1,\ldots,T}$, $\boldsymbol{\varepsilon} = (\varepsilon_t)_{t=1,\ldots,T}$, $\mathbf{I}$ being the identity matrix and

$$\mathbf{B} = \begin{pmatrix} 0 & \cdots & 0 & 0 \\ 1 & \cdots & 0 & 0 \\ \vdots & \ddots & \vdots & \vdots \\ 0 & \cdots & 1 & 0 \end{pmatrix}.$$

Apparently, the process is a random walk if $d = 1$. Moreover, it is important to note that $\mathbf{B}$ is a lower triangular matrix. This ensures that there is some lead-lag relation (i.e., there are future and past values) and that the process is well-defined (i.e., $(\mathbf{I} - \mathbf{B})$ is non-singular).

Now, consider a spatial setting with n locations $\boldsymbol{s}_1, \ldots, \boldsymbol{s}_n$ instead of time points $1, \ldots, T$. These locations are supposed to lie in a q-dimensional space $D \subset \mathbb{R}^q$. Let $\boldsymbol{Y} = (Y(\boldsymbol{s}_i))_{i=1,\ldots,N}$. In this case, there is no clear lead-lag relationship between the observations. Thus, the observation at one specific location $\boldsymbol{s}$ influences all adjacent regions, but the adjacent ones usually also influence the observation in $\boldsymbol{s}$. There are no "future" and "past" observations anymore and, therefore, $\mathbf{B}$ is not necessarily a triangular matrix (this would only be the case for directional spatial processes, see, e.g., Basak et al. 2018; Merk and Otto 2021). Thus, further assumptions are needed such that the process is well-defined. However, in general, we define a spatial autoregressive fractionally integrated process analogously by

$$(\mathbf{I} - \mathbf{B})^d \boldsymbol{Y} = \boldsymbol{\varepsilon}.$$

In spatial settings, the fractional difference operator $(\mathbf{I} - \mathbf{B})^d$ serves to control both the spatial autocorrelation and the fractional differencing. In this regard, time-series ARFIMA processes and the spatial autoregressive fractionally integrated are slightly different. Moreover, for $d = 1$, the model

$$(\mathbf{I} - \mathbf{B})\boldsymbol{Y} = \boldsymbol{\varepsilon}$$

coincides with the commonly known spatial autoregressive model, where $\mathbf{B}$ determines the spatial dependence structure. Usually, $\mathbf{B}$ is chosen as $\rho \mathbf{W}$ with known, prespecified weighting matrices $\tilde{\mathbf{W}}_1, \ldots, \tilde{\mathbf{W}}_k$ and unknown scalar parameters $\rho_1, \ldots, \rho_k$, which has to be estimated (see, e.g., Elhorst et al., 2012).

In this paper, we extend this important class of models to spatial autoregressive fractionally integrated moving average models (spARFIMA). For this reason, we introduce a parameter d, which controls the range and the strength of the spatial dependence. Unlike the spatial autoregressive parameters controlling the degree of spatial dependence on all neighbours, the parameter d influences the shape of the spatial autocorrelation function. That is, this parameter allows to increase the range of the spatial dependence while the process is still stationary. However, we always have to assume that $\mathbf{I} - \mathbf{B}$ is non-singular, restricting the strength of the spatial dependence and leading to a stationary process. Thus, the interpretation of d differs from the time series case. Nevertheless, such a process can be considered to be long-range dependent in the q-dimensional space.

Previous approaches of long-range/memory dependence models for spatial models have mostly focussed on geostatistical settings. In contrast to the spatial econometrics framework, where the spatial dependence is modelled via a suitable spatial weights matrix, which defines the extent of the correlation to all adjacent regions, geostatistical approaches capture the spatial dependence by properly choosing the covariance matrix of a multivariate process. The entries of this covariance matrix usually follow a certain parametric covariance function $C : \mathbb{R}^q \to \mathbb{R}^+$ depending on the difference between two locations $\boldsymbol{s}_i - \boldsymbol{s}_j$. In particular, two-dimensional spatial lattice data has been considered, where the spatial dependence is separable (e.g., Robinson and Sanz 2006). That is, the spatial dependence is fully symmetric in both ways for each direction, meaning longitudinal and latitudinal directions. Hence, two separate backward-shift operators can be applied for each index. They are also called double-geometric processes (cf. Leonenko and Taufer 2013; Martin 1979). Boissy et al. (2005) introduce a fractionally integrated spatial model by considering two d parameters, one for each backshift operator. Thus, this process has a symmetric, long-range dependence in each direction and directly extends the long-memory idea in time series analysis to spatial settings (two-dimensional separable and symmetric settings). Further, Shitan (2008); Ghodsi and Shitan (2009) discussed this model. While Boissy et al. (2005); Robinson and Sanz (2006) focus on Whittle-type estimations of the long-range parameter, Beran et al. (2009) introduced a least-squares estimator. Moreover, a central limit theorem for processes having such kind of spatial dependence has been introduced by Lahiri et al. (2016), applicable even for higher-order and irregular lattices. In contrast to these geostatistical approaches, we focus on so-called spatial econometrics models, which account for spatial autoregressive dependence via weighting matrices.

3

The remainder of the paper is structured as follows. In the following Section 2, the new spARFIMA process is introduced. We present conditions for the existence and stationarity of such a process, and we also point out the differences between time-series ARFIMA processes and geostatistical long-memory processes that assumed separable spatial correlation. For this new spatial model, a quasi-maximum likelihood estimator is derived in Section 3. Furthermore, we carried out an extensive simulation study to show the performance of this QML estimator. The results are presented in Section 4. Eventually, the model is applied to a real-world example important in environmental science in Section 5. More precisely, we analyse raster data on aerosol optical depth with different resolutions. Section 6 concludes the paper.

2   Spatial autoregressive fractionally integrated model

Let $\{Y(\boldsymbol{s}) : \boldsymbol{s} \in D\}$ be a univariate process in the spatial domain D. For instance, D could be the two-dimensional space of integers, i.e., $D \subset \mathbb{Z}^2$, this would cover classical image processes, such as satellite or microscopic images. In spatial statistics, one would commonly refer to this case as a two-dimensional regular lattice process. In econometrics, however, we are often faced to irregular spatial lattice data, like in the case of polygon data (e.g., county-level data). Thus, we generally assume that D is a subset with a positive volume of the q-dimensional real space $\mathbb{R}^q$. That is, contrary to Robinson (2020), we do not restrict ourselves on the case that the process is regularly spaced in two dimensions (i.e., two-dimensional lattice) or that the spatial correlation structure should be symmetric and separable. In our case, the process is observed at a set of n locations, $\{\boldsymbol{s}_1, \ldots, \boldsymbol{s}_n\}$. It is worth noting that this definition also includes spatiotemporal processes if one of the q dimensions is the time axis. For a convenient notation, let $\boldsymbol{Y} = (Y(\boldsymbol{s}_i))_{i=1,\ldots,n}$ be a random vector of all locations and $\boldsymbol{y} = (y(\boldsymbol{s}_i))_{i=1,\ldots,n}$ its observation. In spatial econometrics, it is common to assume that the spatial dependence structure is described by a spatial weights matrix $\mathbf{B} = (b_{ij})_{i,j=1,\ldots,n}$. The diagonal elements of $\mathbf{B}$ are assumed to be zero to prevent self-influences, i.e., $Y(\boldsymbol{s}_i)$ is influenced by itself. In network modelling, this is also known as self-loops.

We define a spatial autoregressive fractionally integrated moving average (spARFIMA) process as follows

$$(\mathbf{I} - \mathbf{B}_1)^d \boldsymbol{Y} = \boldsymbol{\alpha} + (\mathbf{I} - \mathbf{B}_2)\boldsymbol{\varepsilon} \tag{1}$$

with $\boldsymbol{\varepsilon}$ being a vector of independent and identically distributed random variables. The site-specific intercept $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_n)'$ can also be easily extended to linear regression model $\mathbf{X}\boldsymbol{\beta}$. However, we initially focus on the general setting, namely having a site-specific intercept $\boldsymbol{\alpha}$ and general weights matrices $\mathbf{B}_1$ and $\mathbf{B}_2$ for the autoregressive and moving average term, respectively.

In practice, the intercept is often replaced by a constant vector $\boldsymbol{\alpha} = \alpha \mathbf{1}$, and the weighting matrices will be replaced by certain parametric models. In the general case, $\mathbf{B}_1$ and $\mathbf{B}_2$ would consist of $n(n-1)$ unknown parameters, while there are only n observations. Classical choices of such models are, for instance,

$$\mathbf{B}_1 = \rho \mathbf{W}_1 \quad \text{and} \quad \mathbf{B}_2 = \lambda \mathbf{W}_2 \tag{2}$$

with known, pre-specified matrices $\mathbf{W}_1$ and $\mathbf{W}_2$, which describe the structure of the spatial dependence, e.g., they could be first-order contiguity, k-nearest neighbours, or inverse-distance matrices. Moreover, higher-order dependencies can be modelled by a linear combination

$$\mathbf{B}_1 = \sum_{i=1}^{k} \rho_i \mathbf{W}_{i,1} \,,$$

where $\mathbf{W}_{i,1}$ is a contiguity matrix having positive weights for neighbours of spatial lag-order i only. The order of the spatial autoregression would be k in this case. However, more commonly, first-order spatial autoregressive models are considered, and higher-order dependencies are directly included in the spatial weighting matrix. Some recent approaches also considered estimating $\mathbf{B}$ directly by assuming a certain degree of sparsity (e.g. Otto and Steinert 2018; Lam et al. 2013; Lam and Souza 2016). Similarly, higher-order spatial lags can be included in the moving average term, but this is only rarely found in practical applications.

The following theorem shows that the process is well-defined under common conditions for spatial autoregressive models. That is, for any positive d there exists a one-to-one mapping between $\boldsymbol{Y}$ and $\boldsymbol{\varepsilon}$, i.e., $\boldsymbol{Y} = \xi^{-1}(\boldsymbol{\varepsilon}) \; \boldsymbol{\varepsilon} = \xi(\boldsymbol{Y})$.

Theorem 1. *If all diagonal entries of $\mathbf{B}_1$ and $\mathbf{B}_2$ are zero, $||\mathbf{B}_1|| < 1$, $||\mathbf{B}_2|| < 1$, and $d > 0$, the process given by* (1) *is well-defined and there exists one and only one real-valued sequence $Y(\boldsymbol{s}_1), \ldots, Y(\boldsymbol{s}_n)$ that corresponds to $\varepsilon(\boldsymbol{s}_1), \ldots, \varepsilon(\boldsymbol{s}_n)$. Such a process is called a spatial autoregressive fractionally integrated moving average (spARFIMA) process.*

Proof. The process is well-defined and real-valued if and only if $(\mathbf{I} - \mathbf{B}_1)^d$ is non-singular. Applying a binomial expansion, we get that

$$(\mathbf{I} - \mathbf{B}_1)^d = \sum_{k=0}^{\infty} \binom{d}{k} (-1)^k \mathbf{B}_1^k.$$

Because $||\mathbf{B}_1|| < 1$, the series $\mathbf{B}_1^k$ converges for $k \to \infty$ and $(\mathbf{I} - \mathbf{B}_1)$ is invertible. Then, $\mathbf{A} = (\mathbf{I} - \mathbf{B}_1)^{-1}$ and $\mathbf{Y} = \mathbf{A}^d(\boldsymbol{\alpha} + (\mathbf{I} - \mathbf{B}_2)\boldsymbol{\varepsilon})$. Moreover, if $||\mathbf{B}_2|| < 1$ $(\mathbf{I} - \mathbf{B}_2)$ is non singular as well and there is one-to-one mapping from $\boldsymbol{\varepsilon}$ to $\mathbf{Y}$. $\square$

This result makes use of the fact that

$$\left((\mathbf{I} - \mathbf{B})^d\right)^{-1} = \left((\mathbf{I} - \mathbf{B})^{-1}\right)^d. \tag{3}$$

Thus, there is a close relation to spatial autoregressive models, and many results about the existence of spatial autoregressive models also hold for the fractionally integrated model. To be precise, if the spatial autoregressive process for $d = 1$ is well-defined, also the fractionally integrated version exists. For instance, for the common specification with $\mathbf{B} = \rho\mathbf{W}$, all results about the range of the unknown parameter ρ are valid. This also includes higher-order models, as demonstrated by Elhorst et al. (2012). However, it is important to note that d should be too large; otherwise, the inverse in (3) gets unreasonably large, and its values are almost identical. From a practical perspective, this means that the process is not causal; that is, the observations cannot be determined by all other observations because the range of the spatial dependence exceeds the spatial domain. Thus, the process tends to have either extremely large or small values. This depends on the spatial setting, i.e., the number of locations, neighbourhood structure, etc.

Like for spatial autoregressive models, we also observe locally varying mean levels and heteroscedastic variances. However, the long-range dependence parameter d only affects the global spill-over effects, i.e., those associated with the autoregressive term. Whereas the moving average term only locally affects the first and second-lag neighbours – via $\mathbf{B}_2$ and $\mathbf{B}_2\mathbf{B}_2'$ – the autoregressive has global spill-over effects, which are diminished or strengthened by the parameter d. The mean vector and covariance matrix of a spARFIMA process is given in the following proposition.

Proposition 1. *Suppose that ε are identically and independently distributed random errors with mean zero, variance σ_ε^2 and finite fourth moments. Moreover, assume that all diagonal entries*

of **B** *are zero,* $||\mathbf{B}|| < 1$, *and* $d > 0$. *Then, the spatial autoregressive fractionally integrated process given by* (1) *has mean*

$$E(\boldsymbol{Y}) = (\mathbf{I} - \mathbf{B}_1)^{-d}\boldsymbol{\alpha} \tag{4}$$

and covariance matrix

$$Cov(\boldsymbol{Y}) = \sigma_\varepsilon^2(\mathbf{I} - \mathbf{B}_1)^{-d}(\mathbf{I} + \mathbf{B}_2 + \mathbf{B}_2' + \mathbf{B}_2\mathbf{B}_2')(\mathbf{I} - \mathbf{B}_1')^{-d}. \tag{5}$$

Proof. The process can easily be written in a matrix notation as

$$\boldsymbol{Y} = (\mathbf{I} - \mathbf{B}_1)^{-d}\left[\boldsymbol{\alpha} + (\mathbf{I} - \mathbf{B}_2)\boldsymbol{\varepsilon}\right].$$

The result can be obtained by straightforward calculations. $\qquad\square$

Below, because of their close similarity, we briefly discuss the relation to higher-order SAR models. Such higher-order models typically include multiple spatially lagged variables. For instance, a second-order spatial autoregressive model results by

$$\boldsymbol{Y} = \mathbf{B}_{1,1}\boldsymbol{Y} + \mathbf{B}_{1,2}\boldsymbol{Y} + \boldsymbol{\varepsilon} = (\mathbf{I} - \mathbf{B}_{1,1} - \mathbf{B}_{1,2})^{-1}\boldsymbol{\varepsilon}. \tag{6}$$

In contrast, the polynomial expansion would lead to the following model

$$\boldsymbol{Y} = (\mathbf{I} - \mathbf{B}_{1,1})(\mathbf{I} - \mathbf{B}_{1,2})\boldsymbol{Y} + \boldsymbol{\varepsilon}. \tag{7}$$

If $\mathbf{B}_1$ and $\mathbf{B}_2$ takes the easiest parametric form as defined by (2), then the parameter space of ρ_1 and ρ_2 for the process being stationary is much easier to obtain for (7) than for (6), as already pointed out by Elhorst et al. (2012).

2.1   Illustration of Interaction between ρ and d

Eventually, we illustrate the interaction between the spatial autoregressive dependence implied by the weight matrices and the range parameter d using some numerical examples. For simplicity, we only focus on the spatial autoregressive fractionally integrated process without a moving average component (i.e., $\mathbf{B}_2 = \mathbf{0}$). In contrast to time-series or directional spatial models, we allowed $\mathbf{B}_1$ to be non-triangular. Thus, we have to assume that $\mathbf{I} - \mathbf{B}_1$ is invertible. This, in turn, limits the overall spatial dependency to a certain extent so that the interpretation of d is different compared to the time-series case. That is, there is certain interaction between $\mathbf{B}_1$ and d, which we will describe below in more detail for the classic case with $\mathbf{B}_1 = \rho\mathbf{W}_1$.

In Figures 1 and 2, we have illustrated the influence of the central location s on its neighbours of a 20×20 spatial lattice for different values of ρ and d, where $\mathbf{W}_1$ is row-standardised Queen's contiguity matrix in all cases. We particularly focus on processes having a strong spatial autocorrelation, namely $\rho \in \{0.85, 0.9\}$. The dependence of a standard spatial autoregressive model (i.e., $d = 1$) is depicted by the red curves in Figure 1. Obviously, increasing values of ρ (solid vs dashed curves) lead to increased spatial dependence. Since $\mathbf{I} - \rho\mathbf{W}_1$ must be invertible, the parameter ρ must be smaller than one that again limits the spatial autocorrelation. That is, the fractional integration parameter d allows to increase or diminish the spatial autocorrelation, which can be seen by the blue and black curves for $d = 1.5$ and $d = 0.5$, respectively. The intensity of the spatial dependence implied by the blue curves can only be achieved by choosing ρ very close to one for a spatial autoregressive process. However, such a model is close to the ill-defined case.

Now, one might think that by choosing ρ appropriately, one could also achieve the spatial dependence of any other d. However, this is not the case, as we illustrate in Figure 2. Here, we consider a spatial autoregressive model with $\rho = 0.85$ and computed the distance of the closest models with $d = 1.5$ and $d = 2$. That is, we selected ρ such that the squared distances between the curves is minimised – this leads to $\rho = 0.702$ and $\rho = 0.588$ for $d = 1.5$ and $d = 2$, respectively. Obviously, these curves differ in a way that a larger value of d increases the size of the spatial dependence to the higher-order neighbours (i.e., the ones with a distance of at least $\sqrt{5} \approx 2.23$), while dependence to directly adjacent regions is decreased. Hence, the shape implied by different values of fractional integration parameter is different compared to the classical autoregressive process.

3 QML estimation

If the spatial dependence structures $\mathbf{B}_1$ and $\mathbf{B}_2$ are unknown, i.e., it is not known in advance which observations influence each other, each individual link is not generally identifiable. This is a well-known result in spatial econometrics initially pointed out by Manski (1993) (see also Gibbons and Overman 2012). For the spatial long-range dependence model, these results hold equivalently. Generally, suppose that there are two different spatial weight matrices $\mathbf{B}_1 \neq \mathbf{B}_1^*$
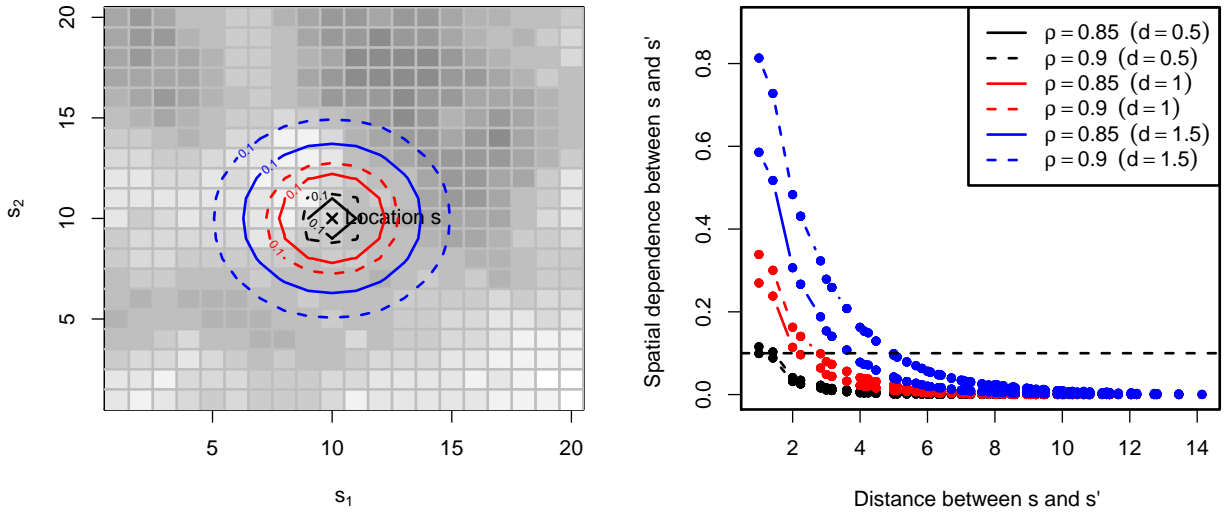
8

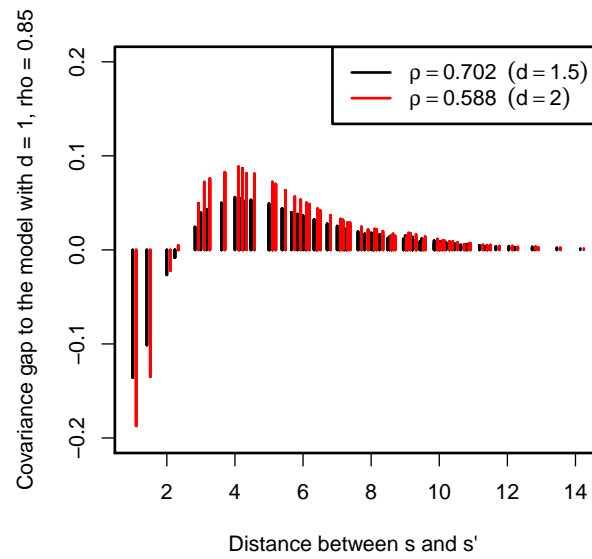Figure 1: Numerical examples: interaction between ρ and d.

Figure 2: Numerical examples: interaction between ρ and d.

9

Table 1: Overview of nested models and special cases.

Weighting matrix	Spatial dimension q	Resulting model
Triangular matrix	1	Time-series ARFIMA(0, d, 1) model
		Note: Adding a further weighting term $(\mathbf{I}-\eta\mathbf{W})$ with $|\eta| < 1$ leads to an ARIMA(1, d, 1) process
	> 1	Causal/directional spatial ARFIMA process
		Note: With $\mathbf{B}_1 = \rho\mathbf{W}$ and $\rho = 1$ a non-stationary spatial random walk is obtained if $d = 1$
Non-triangular matrix	1	Non-causal time-series model
	> 1	Spatial ARFIMA process
		Note: $(\mathbf{I} - \mathbf{B}_1)$ must be invertible (usually, $\mathbf{B}_1 = \rho\mathbf{W}$ with a known, standardised matrix $\mathbf{W}$ and $|\rho| < 1$) to obtain a stationary and well-defined spatial model (i.e., $(\mathbf{I} - \mathbf{B}_1)^d$ serves to control both the fractional differencing and the autoregression)

and long-range dependence parameters $d \neq d^*$. The model is observationally equivalent if

$$(\mathbf{I} - \mathbf{B}_1)^d \boldsymbol{u} = (\mathbf{I} - \mathbf{B}_1^*)^d \boldsymbol{u} \,, \text{ that is,} \tag{8}$$

$$(\mathbf{I} - \mathbf{B}_1)^d = (\mathbf{I} - \mathbf{B}_1^*)^d \,. \tag{9}$$

Here, $\boldsymbol{u}$ denotes the mean and moving average component $\boldsymbol{\alpha} + (\mathbf{I} - \mathbf{B}_2)\boldsymbol{\varepsilon}$. Thus, if $\mathbf{B}_1$ is identifiable, i.e., $\mathbf{B}_1 = \mathbf{B}_1^*$, and $\mathbf{B}_1$ is not equal to a zero matrix, then d is uniquely identifiable. That means that d can only be identified for spatially correlated processes. For the identifiability of $\mathbf{B}_1$ all results that hold for spatial autoregressive models can be applied (see, e.g., Manski 1993). Thus, we follow the common parametric setting described above. That is, suppose that $\boldsymbol{\alpha} = \alpha \mathbf{1}$, $\mathbf{B}_1 = \rho \mathbf{W}_1$, and $\mathbf{B}_2 = \lambda \mathbf{W}_2$.

Let $\boldsymbol{\varepsilon}$ be a vector of independent and identically distributed random variables with the density $f_{\boldsymbol{\varepsilon}}$. Then, the joint likelihood is given by

$$f_{\boldsymbol{Y}}(\boldsymbol{y}) = \left| (\mathbf{I} - \lambda \mathbf{W}_2)^{-1}(\mathbf{I} - \rho \mathbf{W}_1)^d \right| f_{\boldsymbol{\varepsilon}}(\xi(\boldsymbol{y})) \,, \tag{10}$$

where $\boldsymbol{y}$ is the vector of observations. With $f_{\boldsymbol{\varepsilon}}$ being the density of a normal distribution with mean zero and covariance matrix $\sigma_\varepsilon^2 \mathbf{I}$, the logarithmic likelihood function is obtained as

$$\mathcal{L}(\boldsymbol{\vartheta}|\boldsymbol{y}) = -\frac{N}{2}\log(2\pi) - \frac{N}{2}\log(\sigma_\varepsilon^2) - \log|\mathbf{I} - \lambda \mathbf{W}_2| + d\log|\mathbf{I} - \rho \mathbf{W}_1| - \frac{1}{2\sigma_\varepsilon^2}\xi(\boldsymbol{y})'\xi(\boldsymbol{y}) \,. \tag{11}$$

The QML estimator of the parameters $\boldsymbol{\vartheta} = (\alpha, \rho, \lambda)'$ is then given by

$$\hat{\boldsymbol{\vartheta}} = \underset{\boldsymbol{\vartheta} \in \Theta}{\arg\max}\, \mathcal{L}(\boldsymbol{\vartheta}|\boldsymbol{y}) \,. \tag{12}$$

The parameter space Θ depends on the choice of the weight matrices $\mathbf{W}_1$ and $\mathbf{W}_2$, such that the assumptions of Theorem 1 are fulfilled. The main drawback of the QML approach is the scalability to large data sets because it involves the computation of the determinants of the Jacobian, i.e., $|\mathbf{I} - \rho \mathbf{W}_1|$ and $|\mathbf{I} - \lambda \mathbf{W}_2|$. To avoid repeatedly computing the determinant, we suggest following the approach by Ord (1975) for both determinants, that is,

$$\log|\mathbf{I} - a\mathbf{W}| = \sum_{i=1}^{n} \log(1 - a\lambda_W) \,,$$

where λ_W are the eigenvalues of $\mathbf{W}$, which have to be computed only once. This is the main bottleneck of the QML approach regarding scalability. An alternative method is the generalised method of moments, for instance (see Doğan and Taşpınar (2013) for SARMA models).

4   Simulation Studies

We conducted various simulation studies to analyse the algorithm's performance and scalability. For all of them, we considered the classical parametric setup defined above ($\boldsymbol{\alpha} = \alpha\mathbf{1}$, $\mathbf{B}_1 = \rho\mathbf{W}_1$, and $\mathbf{B}_2 = \lambda\mathbf{W}_2$). The $n \times n$ spatial weight matrix $\mathbf{W}_1 = \mathbf{W}_2 = \mathbf{W}$ is a first-order Queen's contiguity matrix, i.e., all surrounding first-lag neighbours are equally affected. This leads to an isotopic setting. Moreover, the locations are assumed to be on a two-dimensional square grid $D = \{\boldsymbol{s} \in \mathbb{Z}^2 : (0,0)' \leq \boldsymbol{s} \leq (\delta, \delta)'\}$. We simulated the process for increasing dimensions of the field $\delta \in \{15, 20, 25\}$ leading to increasing sample sizes of $n \in \{15^2, 20^2, 25^5\}$.

Firstly, we focus on the fractional integration parameter d and purely autoregressive dependencies. That is, we set λ equal to zero. Secondly, we simulated a spARFIMA process with $\lambda = 0.5$. The range parameter was between 0.5 and 2, namely $d \in \{0.8, 1, 1.5\}$. For $d = 1$, the classical spatial autoregressive (with/without a moving average term) is obtained, while for $d = 0.5$ the spatial autoregressive effect is diminished, leading to locally constraint spillovers, and for $d > 1$, the range of the spillover effects is increased compared the SAR case. We considered a medium and large spatial autoregressive dependence, namely $\rho = 0.5$ and $\rho = 0.9$. The results of the simulations experiments in terms of the root mean square errors (RMSE) and the average bias of the estimates can be found in Table 2 and 3 for the setting without and with moving average dependencies, respectively. As expected, the MAE decreases with the increasing size of the spatial fields, while the average bias fluctuates around zero for all cases. Moreover, if the magnitude of the spatial dependence is increasing, the estimates of both ρ and d are getting more precise. This can be seen by the decreasing RMSEs.

We also computed the average time needed to estimate the parameters using a standard `R` implementation for all simulations. The eigenvalues of the weight matrix were computed using the `eigen` function in `R`, and the optimisation of (12) was done numerically using the algorithm implemented in `solnp()` (see Ghalanos and Theussl 2012). The computation time is shown in Figure 3 for both simulation studies, i.e., with $\lambda = 0$ and $\lambda = 0.5$.

5   Real-world illustrative example

To illustrate the usefulness of the range parameter d in practice, we will examine a real-world example below. For this reason, we consider a specific set-up, namely the identical data set,

Table 2: Performance of the QML estimators for $(\rho, \sigma_\varepsilon^2, d)'$ and $\lambda = 0$: root mean square errors (RMSE) and average bias in parentheses, 100 replications.

	Estimation of ρ		Estimation of $\sigma_\varepsilon^2 = 1$		Estimation of d	
	$\rho = 0.5$	$\rho = 0.9$	$\rho = 0.5$	$\rho = 0.9$	$\rho = 0.5$	$\rho = 0.9$
$d = 0.8$						
$n = 15^2$	0.310 (0.069)	0.143 (-0.040)	0.098 (-0.018)	0.099 (-0.004)	0.831 (0.203)	0.387 (0.129)
$n = 20^2$	0.297 (0.029)	0.139 (-0.056)	0.069 (-0.007)	0.071 (0.006)	0.806 (0.220)	0.354 (0.105)
$n = 25^2$	0.290 (0.038)	0.106 (-0.038)	0.062 (0.000)	0.064 (0.011)	0.771 (0.269)	0.235 (0.080)
$d = 1.0$						
$n = 15^2$	0.280 (0.076)	0.119 (-0.027)	0.098 (-0.020)	0.099 (-0.005)	0.732 (0.099)	0.298 (0.133)
$n = 20^2$	0.258 (0.037)	0.102 (-0.029)	0.069 (-0.009)	0.071 (0.005)	0.705 (0.151)	0.289 (0.104)
$n = 25^2$	0.252 (0.048)	0.076 (-0.025)	0.062 (-0.001)	0.063 (0.010)	0.652 (0.131)	0.198 (0.060)
$d = 1.5$						
$n = 15^2$	0.216 (0.086)	0.058 (-0.012)	0.100 (-0.024)	0.098 (-0.007)	0.560 (-0.121)	0.238 (0.042)
$n = 20^2$	0.188 (0.057)	0.055 (-0.018)	0.069 (-0.012)	0.070 (0.003)	0.534 (-0.044)	0.212 (0.065)
$n = 25^2$	0.186 (0.062)	0.045 (-0.012)	0.063 (-0.005)	0.063 (0.008)	0.519 (-0.069)	0.170 (0.040)

Table 3: Performance of the QML estimators for $(\rho, \sigma_\varepsilon^2, d)'$ and $\lambda = 0.5$: root mean square errors (RMSE) and average bias in parentheses, 100 replications.

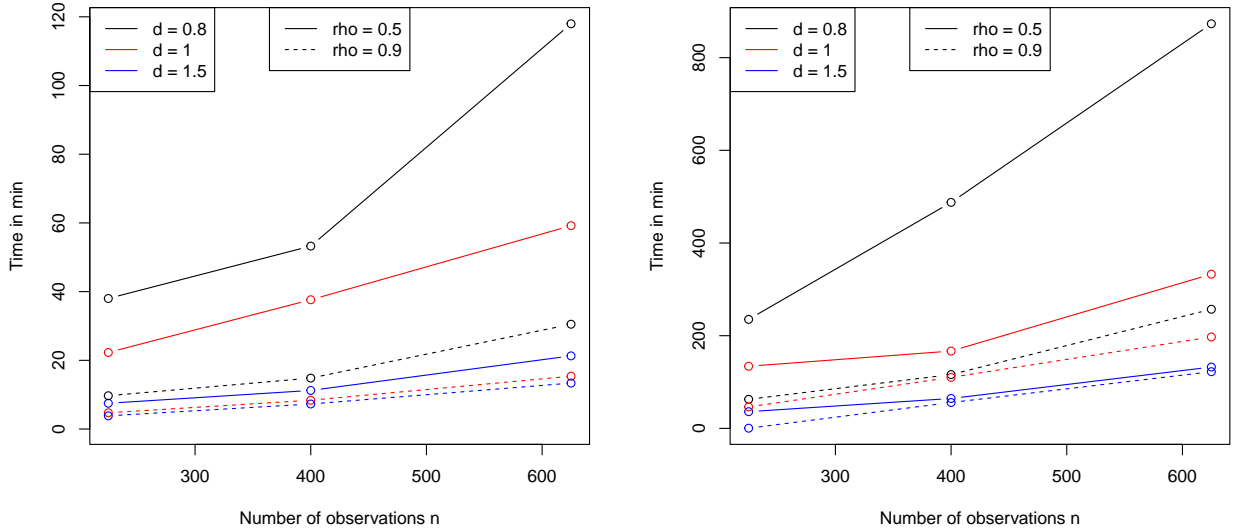|	Estimation of ρ		Estimation of $\sigma_\varepsilon^2 = 1$		Estimation of d		Estimation of $\lambda = 0.5$	
	$\rho = 0.5$	$\rho = 0.9$	$\rho = 0.5$	$\rho = 0.9$	$\rho = 0.5$	$\rho = 0.9$	$\rho = 0.5$	$\rho = 0.9$
$d = 0.8$								
$n = 15^2$	0.351 (0.126)	0.166 (-0.03)	0.333 (0.083)	0.287 (-0.064)	0.099 (-0.011)	0.099 (0.001)	0.663 (0.042)	0.325 (0.0
$n = 20^2$	0.346 (0.187)	0.155 (-0.03)	0.305 (0.104)	0.274 (-0.051)	0.070 (0.001)	0.071 (0.011)	0.647 (0.119)	0.308 (0.0
$n = 25^2$	0.322 (0.151)	0.082 (-0.00)	0.288 (0.082)	0.255 (-0.044)	0.064 (0.006)	0.065 (0.015)	0.615 (0.038)	0.271 (-0.0
$d = 1.0$								
$n = 15^2$	0.384 (0.142)	0.104 (-0.010)	0.350 (0.002)	0.320 (-0.137)	0.099 (-0.016)	0.098 (-0.001)	0.643 (-0.036)	0.324 (-0.1
$n = 20^2$	0.365 (0.124)	0.101 (-0.008)	0.328 (0.018)	0.307 (-0.172)	0.070 (-0.003)	0.071 (0.011)	0.637 (0.028)	0.311 (-0.1
$n = 25^2$	0.355 (0.133)	0.065 (0.003)	0.340 (0.023)	0.262 (-0.143)	0.063 (0.001)	0.064 (0.014)	0.608 (0.004)	0.284 (-0.1
$d = 1.5$								
$n = 15^2$	0.321 (0.017)	0.063 (0.002)	0.395 (-0.210)	0.321 (-0.168)	0.101 (-0.023)	0.098 (-0.001)	0.610 (-0.246)	0.344 (-0.1
$n = 20^2$	0.320 (-0.032)	0.050 (0.004)	0.393 (-0.253)	0.312 (-0.195)	0.069 (-0.012)	0.070 (0.009)	0.597 (-0.172)	0.324 (-0.1
$n = 25^2$	0.313 (-0.026)	0.040 (0.008)	0.387 (-0.243)	0.286 (-0.174)	0.062 (-0.005)	0.064 (0.014)	0.595 (-0.198)	0.312 (-0.1

Figure 3: Average computation time for the estimation of the parameters (numerical maximisation) for the case autoregressive model, $\lambda = 0$ (left), and the autoregressive moving average model, $\lambda = 0.5$ (right)

but in three different resolutions. At the same time, we apply classical weighting matrices (i.e., Queen's contiguity matrices) that weigh all neighbouring grid cells equally. Since the data set does not change and the dependence structure (in a geographical sense) thus remains the same, the shape of the spatial autocorrelation function changes for higher resolutions (because the neighbouring raster cells are geographically closer). In the lowest resolution, the distance between grid cells is greater in a geographical sense, and spatial dependence is, therefore, faster, declining to zero (in terms of the number of spatial lags). Hence, the parameter d provides additional flexibility for the shape and the range of the spatial dependence.

More precisely, we consider raster data on the aerosol optical depth obtained from NEO, NASA Earth Observations, measured by NASA's Moderate Resolution Imaging Spectroradiometer (MODIS). The aerosol optical thickness measures the concentration of solid and liquid particles in the atmosphere, so-called aerosols. This aerosol concentration plays an important role in weather, climate, air quality, and thus human's health (cf. Kumar et al. 2007; Wang and Christopher 2003; Gupta et al. 2013; Van Donkelaar et al. 2010). Moreover, these aerosols are one of the greatest sources of uncertainty in climate modelling.

A climactic active and interesting area is the Northern Atlantic Ocean over the equator.

15

Thus, we considered this area, N 0°-25°, E -45°-20°, which also covers the area of the most Northern Atlantic hurricanes, in different resolutions of 0.5×0.5, 1×1, and 2×2 degrees. This leads to quadratic lattices of sizes 12×12, 25×25, and 50×50 for the highest, medium, and lowest resolution, respectively. Hence, the sample size increases from $n = 144$, $n = 625$, and $n = 2,500$ observations. It is worth noting that this implies a $2,500$-dimensional weighting matrix for the computationally largest problem. Because the focus is on the range of the spatial dependence, we standardised each data set in advance.

The full data set is shown along with the subset of the three considered resolutions in Figure 4. In addition to showing the data, we also provide the estimated spatial autocorrelation functions based on Moran's I in the bottom row of this figure. From these spatial ACFs, one could see that the lower the resolution, the faster the spatial autocorrelation is decaying – because the directly neighbouring pixels for the lowest resolution already cover a larger geographical distance than the directly adjacent pixels for the highest resolution. Thus, the fractional integration parameter d can provide further flexibility for the model, especially for the larger ranges in higher resolutions. Moreover, one could see that the clusters appear more pronounced with rather sharp edges for the images with a higher resolution compared to the third case with a low resolution. In Table 4, we report the resulting estimated parameter along with their estimated standard errors of a spatial ARFIMA model for all three resolutions. The standard errors are obtained from the Hessian of log-likelihood as Cramer-Rao bounds. Because the moving average component seems to be irrelevant (non-significant and leading to lower AIC/BIC), all models have been estimated for $\lambda = 0$. As a benchmark model, we also report the results of a classical spatial autoregressive model (i.e., $d = 1$). For the sake of completeness, we also report the results of a SARMA model. At this point, it is worth noting that one could also test for the difference of the parameter d to 1.

Looking at the information criteria reported in Table 4, we see that the fractional integration of the spatial autoregressive is particularly useful for medium and high resolutions. While we are getting good model fits for a SAR process in the case of the lowest resolution, both the AIC and BIC criteria are smaller for the spARFIMA process in the two other cases. Moreover, we see that the autoregressive parameters are larger while the parameter d is smaller compared to the low-resolution case. That is, there is a strong spatial dependence on the directly adjacent pixels, which decays fast with the spatial distance. This leads to more pronounced and sharp clusters compared to the low-resolution case, where the clusters rather fade out across space

16

Table 4: Estimated parameters of an spatial ARFIMA process (with $\lambda = 0$) and classical SAR and SARMA models as benchmark.

	Resolution	spARFIMA		SAR model		SARMA model	
		Estimate	Standard error	Estimate	Standard error	Estimate	Standard error
d	low	1.3178	0.4029				
	medium	0.7656	0.0610	$(d=1)$		$(d=1)$	
	high	0.6927	0.0271				
ρ	low	0.8576	0.1228	0.9440	0.0275	0.9440	0.0373
	medium	0.9911	0.0090	0.9466	0.0126	0.9719	0.0119
	high	0.9967	0.0025	0.9435	0.0066	0.9758	0.0054
λ	low					0.0000	0.2102
	medium	$(\lambda=0)$		$(\lambda=0)$		0.2526	0.0913
	high					0.3275	0.0435
σ_ε^2	low	0.1654	0.0216	0.1637	0.0200	0.1637	0.0203
	medium	0.1338	0.0077	0.1347	0.0085	0.1306	0.0076
	high	0.1370	0.0040	0.1378	0.0040	0.1328	0.0039
AIC	low	189.1391		188.3013*		190.3013	
	medium	660.91*		667.1188		662.5683	
	high	2620.552*		2677.81		2634.109	
BIC	low	198.0485		194.2409*		199.2107	
	medium	674.2232*		675.9943		675.8815	
	high	2638.024*		2689.462		2651.581	
Residuals' standard deviation	low	0.4076		0.4053		0.4053	
	medium	0.3661		0.3673		0.3617	
	high	0.3702		0.3713		0.3644	
Moran's I of the residuals (p-value)	low	0.0003 (0.4346)		0.0331 (0.1822)		0.0331 (0.1822)	
	medium	0.0025 (0.4207)		-0.0386 (0.9645)		0.0006 (0.4567)	
	high	0.0053 (0.2867)		-0.0509 (1.0000)		0.0008 (0.4535)	

(because of the averaging of the grid cells). To a limited degree, the moving average residuals could also capture this behaviour. Thus, the SARMA model shows a better fit compared to the SAR model for medium and high resolution.

6 Discussion and Conclusions

Motivated by time-series fractionally integrated autoregressive models, we have introduced the concept of fractional integration for spatial autoregressive processes. More precisely, we developed a spatial autoregressive fractionally integrated moving average model (spatial ARFIMA) that is suitable for data observed in multidimensional space. Moreover, we do not restrict the process to regularly spaced grid data so that the process can be applied to irregular polygon data, as it is often the case in economics, but also to regular grids, like image, geostatistical,
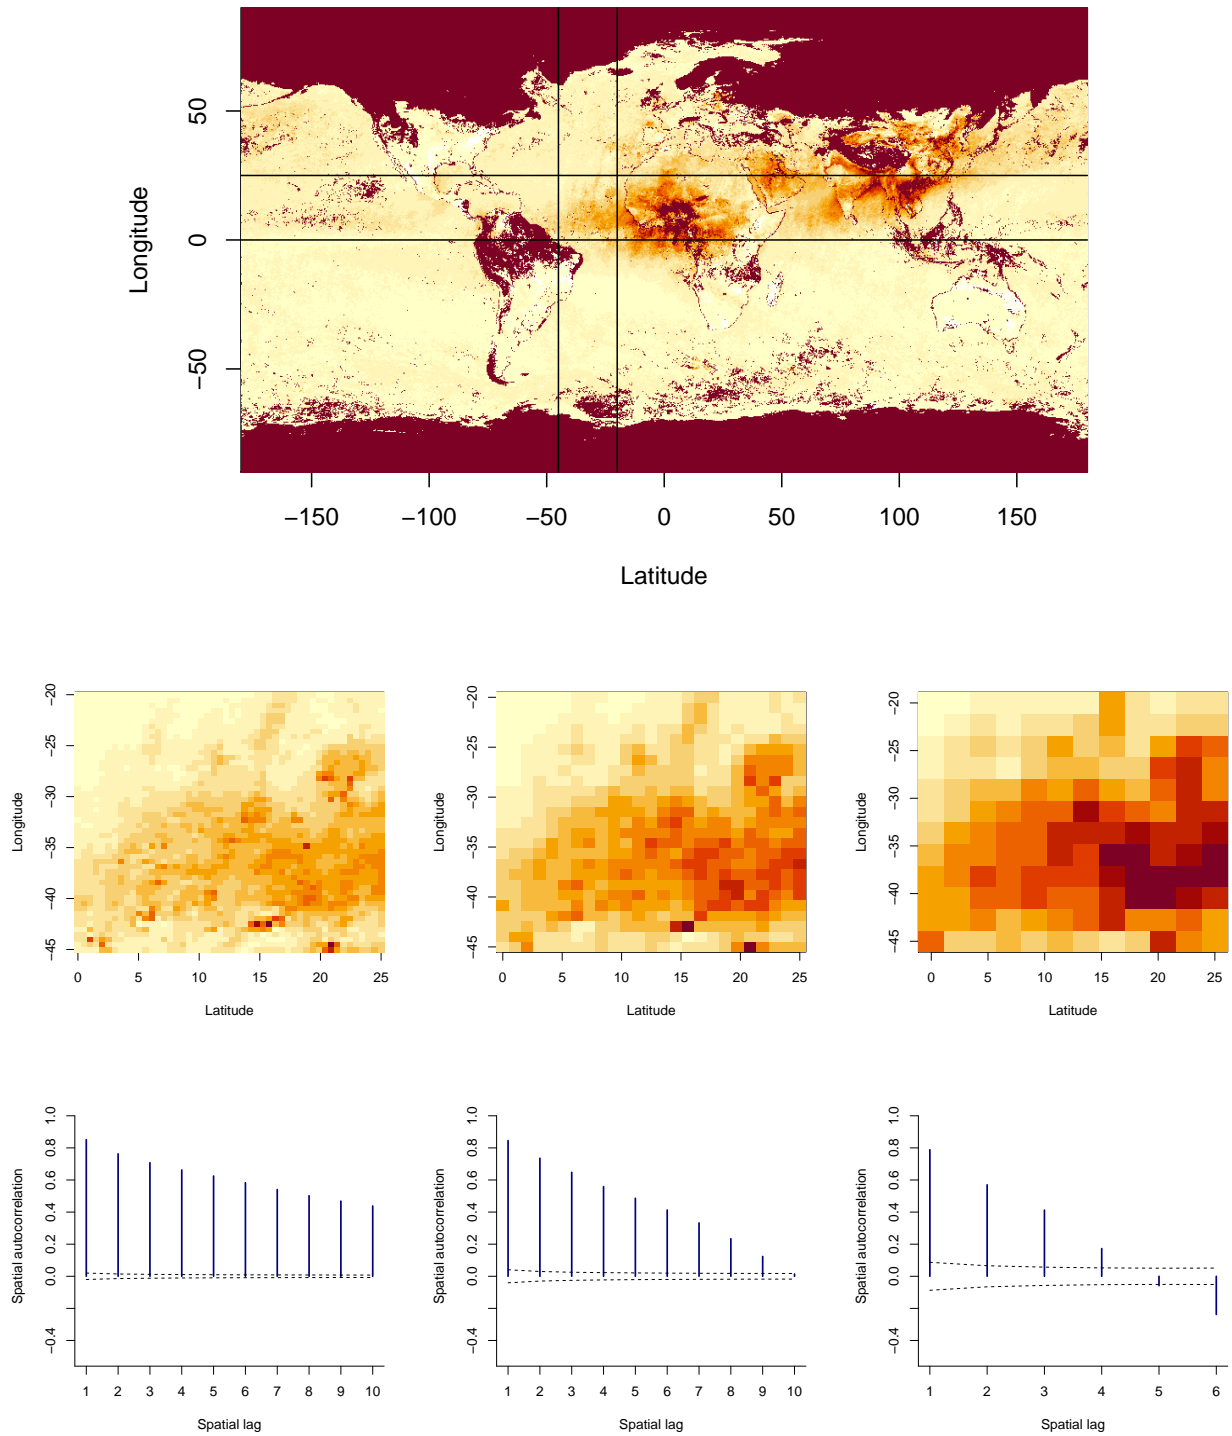
Figure 4: Optical aerosol depth (top: global data, middle: high, medium, and low resolution from left to right, bottom: spatial ACF)

or raster data. The latter examples are often present in environmental studies.

In contrast to time-series ARFIMA processes, fractional integration is directly included in the spatial autoregressive term. Alternatively, two different spatial weight matrices could be considered – one for the fractional integration and one for the autoregressive dependence. In spatial settings, however, the choice of the weight matrix is complicated, and often it has a prespecified structure, so it is preferable to combine these two effects into one term.

This new spatial ARFIMA model is closely related to SAR models, so many results can be directly applied, e.g., on the identification or estimation. This paper considers the frequently applied QML approach to estimate the parameters. We paid particular attention to the scalability of this approach. Furthermore, we analysed the performance and the computation time in a series of Monte-Carlo simulation studies.

Finally, the model has been applied to real data – aerosol optical depth. We focussed on the interaction between the fractional integration and the spatial autoregressive parameter because the same data was analysed in different resolutions. We found a pronounced spatial dependence for all resolutions. The fractional integration parameter was particularly useful for images in higher resolutions.

References

Basak, G. K., Bhattacharjee, A., and Das, S. (2018). Causal ordering and inference on acyclic networks. *Empirical Economics*, 55(1):213–232.

Beran, J. (2017). *Statistics for long-memory processes.* Routledge.

Beran, J., Ghosh, S., and Schell, D. (2009). On least squares estimation for long-memory lattice processes. *Journal of Multivariate Analysis*, 100(10):2178–2194.

Boissy, Y., Bhattacharyya, B., Li, X., Richardson, G., et al. (2005). Parameter estimates for fractional autoregressive spatial processes. *The Annals of Statistics*, 33(6):2553–2567.

Doğan, O. and Taşpınar, S. (2013). Gmm estimation of spatial autoregressive models with moving average disturbances. *Regional Science and Urban Economics*, 43(6):903–926.

Elhorst, J. P., Lacombe, D. J., and Piras, G. (2012). On model specification and parameter

space definitions in higher order spatial econometric models. *Regional Science and Urban Economics*, 42(1-2):211–220.

Ghalanos, A. and Theussl, S. (2012). *Rsolnp: General Non-linear Optimization Using Augmented Lagrange Multiplier Method.* R package version 1.14.

Ghodsi, A. and Shitan, M. (2009). Estimation of the memory parameters of the fractionally integrated separable spatial autoregressive (FISSAR (1, 1)) model: A simulation study. *Communications in Statistics-Simulation and Computation*, 38(6):1256–1268.

Gibbons, S. and Overman, H. G. (2012). Mostly pointless spatial econometrics? *Journal of Regional Science*, 52(2):172–191.

Gupta, P., Khan, M. N., da Silva, A., and Patadia, F. (2013). Modis aerosol optical depth observations over urban areas in pakistan: quantity and quality of the data for air quality monitoring. *Atmospheric pollution research*, 4(1):43–52.

Kumar, N., Chu, A., and Foster, A. (2007). An empirical relationship between pm2. 5 and aerosol optical depth in delhi metropolitan. *Atmospheric Environment*, 41(21):4492–4503.

Lahiri, S., Robinson, P. M., et al. (2016). Central limit theorems for long range dependent spatial linear processes. *Bernoulli*, 22(1):345–375.

Lam, C. and Souza, P. C. (2016). Detection and estimation of block structure in spatial weight matrix. *Econometric Reviews*, 35(8-10):1347–1376.

Lam, C., Souza, P. C., et al. (2013). Regularization for spatial panel time series using the adaptive lasso. *Journal of Regional Science*.

Leonenko, N. and Taufer, E. (2013). Disaggregation of spatial autoregressive processes. *Spatial Statistics*, 3:1–20.

Manski, C. F. (1993). Identification of endogenous social effects: The reflection problem. *The review of economic studies*, 60(3):531–542.

Martin, R. (1979). A subclass of lattice processes applied to a problem in planar sampling. *Biometrika*, 66(2):209–217.

Merk, M. S. and Otto, P. (2021). Directional spatial autoregressive dependence in the conditional first-and second-order moments. *Spatial Statistics*, 41.

Ord, K. (1975). Estimation methods for models of spatial interaction. *Journal of the American Statistical Association*, 70(349):120–126.

Otto, P. and Steinert, R. (2018). Estimation of the Spatial Weighting Matrix for Spatiotemporal Data under the Presence of Structural Breaks. Technical report.

Robinson, P. M. (2020). Spatial long memory. *Japanese Journal of Statistics and Data Science*, 3(1):243–256.

Robinson, P. M. and Sanz, J. V. (2006). Modified whittle estimation of multilateral models on a lattice. *Journal of Multivariate Analysis*, 97(5):1090–1120.

Shitan, M. (2008). Fractionally integrated separable spatial autoregressive (FISSAR) model and some of its properties. *Communications in Statistics—Theory and Methods*, 37(8):1266–1273.

Van Donkelaar, A., Martin, R. V., Brauer, M., Kahn, R., Levy, R., Verduzco, C., and Villeneuve, P. J. (2010). Global estimates of ambient fine particulate matter concentrations from satellite-based aerosol optical depth: development and application. *Environmental health perspectives*, 118(6):847–855.

Wang, J. and Christopher, S. A. (2003). Intercomparison between satellite-derived aerosol optical thickness and pm2. 5 mass: Implications for air quality studies. *Geophysical research letters*, 30(21).