

Regularization of the ensemble Kalman filter using a non-parametric, non-stationary spatial model

Michael Tsyrlunikov^{a,*}, Arseniy Sotskiy^{a,b}

^a*HydroMeteorological Research Center of Russia, Russian Federation*

^b*National Research University Higher School of Economics, Russian Federation*

Abstract

The sample covariance matrix of a random vector is a good estimate of the true covariance matrix if the sample size is much larger than the length of the vector. In high-dimensional problems, this condition is never met. As a result, in high dimensions the Ensemble Kalman Filter's (EnKF) ensemble does not contain enough information to specify the prior covariance matrix accurately. This necessitates the need for regularization of the analysis (observation update) problem. We propose a regularization technique based on a new spatial model. The model is a constrained version of the general Gaussian process convolution model. The constraints include local stationarity and smoothness of local spectra. We regularize EnKF by postulating that its prior covariances obey the spatial model. Placing a hyperprior distribution on the model parameters and using the likelihood of the prior ensemble data allows for an optimized use of both the ensemble and the hyperprior. A linear version of the respective estimator is shown to be consistent. A more accurate nonlinear neural-Bayes implementation of the estimator is developed. In simulation experiments, the new technique led to substantially better EnKF performance than several existing techniques.

Keywords: Process convolution model, Spatial non-stationarity, Data assimilation, Ensemble Kalman Filter, Neural Bayes estimator, Gaussian process

1. Introduction

Modern data assimilation in geosciences and related areas relies on the ensemble approach, in which the probability distribution of the unknown system state is represented by a finite sample (ensemble) of pseudo-random realizations (ensemble members). This allows for realistic, spatially varying prior covariances in the analysis (the observation update step of the data assimilation cycle). In practical applications, the most widely-used approach is the Ensemble Kalman Filter (EnKF) and its numerous modifications and

*Corresponding author

Email address: michael.tsyrlunikov@gmail.com (Michael Tsyrlunikov)

This is the authors' final version of the paper published in *Spatial Statistics*, 2024, v.64, 100870. DOI: <https://doi.org/10.1016/j.spasta.2024.100870>

©2024. This manuscript version is made available under the CC-BY-NC-ND 4.0 license <https://creativecommons.org/licenses/by-nc-nd/4.0/>

extensions². The principal problem of the ensemble approach is that maintaining a large ensemble in the course of data assimilation (sequential filtering) is computationally expensive. For real-world systems modeled with high-resolution models with billions of degrees of freedom, this means that only relatively small ensembles are affordable. As a result, such ensembles can provide the analysis with only scarce information on the prior distribution. The sample covariance matrix is a poor estimate of the true covariance matrix if the sample size is much lower than the dimensionality of the state space. This requires a kind of *regularization*, that is, the introduction of additional information on the prior covariances.

1.1. Related work

The following EnKF regularization techniques have been proposed so far.

1. The most popular approach is covariance localization (tapering), (e.g. Furrer and Bengtsson, 2007, Houtekamer and Zhang, 2016), which reduces spurious long-distance correlations through element-wise multiplication of the sample covariance matrix by an ad-hoc positive definite localization matrix. This technique efficiently removes a lot of noise in the sample covariance matrix. However, it cannot cope with the noise at small distances (where the noise is the strongest). And it inevitably leads to an underestimation of the assumed length scale of the forecast-error field and thus to imbalances between different fields (Lorenc, 2003).
2. Blending prior sample covariances with static (time-mean) covariances is widely used in meteorological ensemble-variational schemes (Buehner et al., 2013, Lorenc et al., 2014). In statistical literature, similar techniques are known as shrinkage estimators (Ledoit and Wolf, 2004). Sample covariances are noisy but they do contain the useful flow-dependent ‘signal’. Static covariances are noise-free but can be irrelevant under some flow conditions and observation coverage. A convex combination of the two covariance matrices proved to be useful (see the above references) but it is not selective: the noise in the sample covariances is reduced to the same extent as the flow-dependent non-stationary signal.
3. Another approach is spatial averaging of sample covariances (that is, blending with neighboring in space covariances), e.g., Berre and Desroziers (2010). The technique reduces noise in sample covariances due to an increase in the effective ensemble size but it does so at the expense of somewhat distorting (smoothing) the covariances. Optimal linear filtering of sample covariances (Ménétrier et al., 2015) further develops this idea.
4. A similar technique is temporal averaging of prior sample covariances (i.e., blending with recent past covariances). Berre et al. (2015) and Bonavita et al. (2016) used

²See section 2 for a brief introduction and Appendix A for a more technical description of sequential ensemble filtering in the linear setting (which is sufficient for our purposes in this study).

prior ensembles from several previous assimilation cycles to increase the ensemble size. Lorenc (2017) used time-shifted ensemble perturbations for the same purpose. Tsyrlunikov and Rakitko (2017) arrived at this technique by assuming that the true covariance matrix is an unknown random matrix and introducing a secondary filter in which the covariances are updated.

5. Ueno and Tsuchiya (2009) proposed to regularize the sample covariance matrix by imposing a sparse structure in the inverse covariance (precision) matrix. A similar approach was taken by Hou et al. (2021). Gryvill and Tjelmeland (2023) imposed a Markov random field structure to obtain a sparse prior precision matrix. Boyles and Katzfuss (2021) postulated sparsity in inverse Cholesky factors of the prior covariance matrix.
6. One more option is to adopt a parametric model for the prior covariances and estimate parameters of the model from the forecast ensemble. Skauvold and Eidsvik (2019) used maximum likelihood estimates of a few model parameters in stationary and non-stationary settings.
7. Finally, imposing a structure in wavelet space can be used to regularize the analysis. Postulating the simplest structure, the diagonal wavelet covariance matrix, proved to be effective in data assimilation applications (Fisher, 2003, Pannekoucke et al., 2007), see also Kasanický et al. (2015). This indicates that the regularizing effect of that approach can outweigh methodological errors related to the neglect of off-diagonal covariances in wavelet space.

1.2. Motivation and contributions

In this paper we consider systems whose state vector represents a *spatial field*, e.g., atmosphere, ocean, and the like. Taking advantage of this interpretation of the state vector, we propose to regularize the analysis problem by imposing a *spatial model*. In this regard, our approach is similar to the two latter techniques (items 6 and 7 above). In contrast to parametric covariance models, we aim at a non-parametric formulation to obtain a more flexible model. By non-parametric, we mean defined by parameters whose number is not fixed (as with a parametric model) but grows with the increasing spatial resolution of the model. Compared to the wavelet-diagonal models, which provide regularizing information by imposing zero cross-correlations between wavelet coefficients, our model does so by imposing several constraints on the general process convolution model. The two principal constraints are local stationarity (under which a spatially non-stationary process can be approximated by a stationary process locally in space) and smoothness of what we call local spectra.

We assume that the prior distribution of the state at each filtering time step is multivariate Gaussian with the mean equal to the ensemble mean (or, for a linear forecast model, to the deterministic forecast, see Remark in Appendix A.4). The prior covariances are postulated to be governed by the spatial model introduced in this paper. We place a

hyperprior distribution on the model parameters and use the likelihood of the ensemble data to online estimate the spatial model. The estimated model is then employed to compute a square root of the prior covariance matrix, which is finally used in the analysis.

The reason why we propose to rely on a Bayesian estimate of a spatial model in the EnKF analysis is that a more principled (than the above ad-hoc techniques) way of covariance regularization can lead to a better estimate of the true prior covariance matrix and thus better performance of the filter. It is also worth noting that the introduction of the hyperprior distribution implies an *enhancement* of EnKF that goes beyond blending of sample covariances with time-mean covariances (mentioned in item 2 above).

Our key contributions are:

- A constrained non-parametric, non-stationary process convolution model on the sphere. The concept of local stationarity of the resulting spatial process defined by conditioning on the random convolution kernel.
- An estimator of the spatial model suitable for high-dimensional applications. A proof that a linear version of the estimator is consistent. A nonlinear, neural Bayes realization of the estimator.
- Implementation of the model and its estimator in EnKF, leading to a new Locally Stationary Ensemble Filter (LSEF).
- Numerical experiments with simulated data with the LSEF analysis on the sphere and the LSEF on the circle.

The rest of the paper is organized as follows. Section 2 gives a brief overview of sequential ensemble filtering. The theoretical part of the paper consists of two big sections, 3 and 4, supported by several Appendices. In section 3 we introduce a non-parametric locally stationary Gaussian process convolution model on the sphere (and on the circle). In section 4 we propose an estimator of the model from an ensemble of random field realizations. Section 5 describes how the spatial model is used to regularize the analysis step of the EnKF. The experimental part of the paper is presented in sections 6 and 7.

Reproducible Python and R code for numerical experiments presented in this paper can be found at <https://github.com/Arseniy100/LSEF>.

2. Background on sequential ensemble filtering

The goal of sequential filtering is the estimation of the current state of an evolving *system*, having (i) a forecast model capable of predicting the future state of the system given its current state, (ii) a set of observations on the system, and (iii) an observation model that relates the observations to the system state (the truth).

In the time-discrete filtering, observations are available at discrete time instants. The estimates of the (ground) truth are sought at the same time instants. A sequential filter consists of alternating short-term forecast (time update) and analysis (observation update) steps. In the time-discrete *linear* filtering, which is sufficient to test the covariance

regularization technique we propose in this research, the evolution of the truth \mathbf{x}_t between consecutive time instants t is assumed to be governed by the stochastic equation

$$\mathbf{x}_t = \mathbf{F}_t \mathbf{x}_{t-1} + \boldsymbol{\eta}_t, \quad (1)$$

where \mathbf{F}_t is the known forecast-model operator (a matrix) and $\boldsymbol{\eta}_t$ is the unknown random model error. The assimilated observations are assumed to be related to the truth as

$$\mathbf{x}_t^{\text{obs}} = \mathbf{H}_t \mathbf{x}_t + \boldsymbol{\varepsilon}_t, \quad (2)$$

where \mathbf{H}_t is the known observation-model operator and $\boldsymbol{\varepsilon}_t$ is the unknown random observation error.

If the errors $\boldsymbol{\eta}_t$ and $\boldsymbol{\varepsilon}_t$ are all mutually uncorrelated and have known first and second moments, then the mean-square optimal estimate of \mathbf{x}_t given $\mathbf{x}_t^{\text{obs}}, \mathbf{x}_{t-1}^{\text{obs}}, \dots$ is given by the Kalman filter (e.g. Jazwinski, 1970). In high dimensions, however, the Kalman filter is prohibitively expensive. Its Monte-Carlo version known as the Ensemble Kalman filter (EnKF) is an affordable approximate filter, (e.g. Asch et al., 2016).

EnKF propagates in time a sample (ensemble) of ‘possible truths’ consistent with observations. A *forecast (prior) ensemble* is computed in EnKF at every assimilation cycle t before assimilation of current observations, and an *analysis (posterior) ensemble* is computed after current observations are assimilated. The key assumption in EnKF is that the prior ensemble is a sample from the true prior distribution. EnKF computes a *regularized sample covariance matrix* $\widehat{\mathbf{B}}_t$ from the prior ensemble and substitutes it for the *true prior covariance matrix* \mathbf{B}_t in the Kalman-filter analysis equation. It is the prior sample covariance matrix that we aim to regularize in this research using a spatial model — for systems whose state is a spatial field.

Having a deterministic forecast (the mean of the prior ensemble or the so-called control forecast), we deal in what follows with the *forecast-error* random field, that is, the deviation of the forecast from the truth. Online information on its current (i.e., at the time instant t) probability distribution is given by *ensemble perturbations*, the differences of the prior ensemble members and the deterministic forecast.

For more details on approximate sequential filtering, see Appendix A.

3. Spatial model

We employ a non-parametric spatial model for the forecast-error random field to regularize the data assimilation problem. Following Higdon et al. (1999), we rely on the Gaussian *process convolution* model. In contrast to most applications, which postulate a parametric model for the spatial convolution kernel (e.g. Lemos and Sansó, 2009, Sham Bhat et al., 2012, Heaton et al., 2014, Li and Zhu, 2020), we choose a *non-parametric* approach to allow for variable shapes of spatial covariances. Non-parametric convolution kernels were considered by Tobar et al. (2015) and Bruinsma et al. (2022) in the stationary

time series context. Here we propose a non-parametric non-stationary spatial convolution model on compact manifolds such as the sphere and the circle.

The default domain is the two-dimensional unit sphere \mathbb{S}^2 (the domain of interest in geosciences). We also formulate and test the model on the unit circle \mathbb{S}^1 (used as a lower-dimensional counterpart of the sphere). We refer to \mathbb{S}^2 and \mathbb{S}^1 just as the sphere and the circle in the sequel. We denote points on the sphere as x, y, \dots and area elements simply by dx, dy, \dots . On the sphere, we use the terms stationarity and isotropy interchangeably. The terms field and process are also used interchangeably. We confine ourselves to band-limited processes because in the context of sequential filtering, the spatial field to be estimated inevitably has limited spatial resolution. We also assume that probability distributions we deal with in this study have densities with respect to the Lebesgue measure.

The rest of this section is structured as follows. We start by defining our Gaussian process convolution model, including the introduction of so-called local spectra. Then, we realize that to be identifiable from the prior ensemble, the model needs to be constrained. We propose the following four constraints: (i) local isotropy of the convolution kernel, (ii) positive definiteness of the non-stationary convolution kernel as a function of distance, (iii) local stationarity, (iv) smoothness of the local spectra.

3.1. General process convolution model

Let $\xi(x)$ be a general real-valued space-continuous zero-mean linear Gaussian process:

$$\xi(x) = \int_D w(x, y) \alpha(y) dy \equiv \int_D w(x, y) Z(dy). \quad (3)$$

Here D is the domain of interest, α is the white Gaussian noise, Z is the Gaussian orthogonal stochastic measure such that $\mathbb{E} Z(dx) = 0$, $\mathbb{E} Z^2(dx) = |dx|$, $\mathbb{E} Z(dx)Z(dy) = 0$ whenever $dx \cap dy = \emptyset$, \mathbb{E} is expectation, and $w(x, y)$ is a non-random real function called the convolution kernel.

The (non-stationary in general) covariance function of the process $\xi(x)$ defined by Eq. (3) is $B(x, x') = \mathbb{E} \xi(x)\xi(x') = \int_D w(x, y)w(x', y) dy$.

3.2. Constraining the model: strategy

In the formulation Eq. (3), the model is not identifiable, that is, the function $w(x, y)$ is not unique given the process covariance function $B(x, x')$. Indeed, consider an isotropic convolution kernel $\psi(\rho(y, y'))$, where $\rho(y, y')$ stands for the great-circle distance between the points y and y' . Let the Fourier-Legendre transform of $\psi(\rho)$ be equal to one in modulus (see Eq. (6) for the definition of the Fourier-Legendre transform we use). Then, it is straightforward to see that the covariance function remains unchanged if $w(x, y)$ is convolved with ψ^3 . (The convolution with such a kernel is analogous to the multiplication by an orthogonal matrix in the space-discrete case.)

³By the convolution of $w(\cdot, \cdot)$ with $\psi(\rho(\cdot, \cdot))$ we mean $w_*(x, y) = \int w(x, z)\psi(\rho(z, y)) dz$.

We intend to introduce constraints on the convolution kernel to make the model identifiable and reduce the effective number of degrees of freedom in the model, thus facilitating its estimation from an ensemble of random field realizations.

3.3. Locally isotropic kernel

With a fixed x , we require that $w(x, y)$ depends on y only through the great-circle distance $\rho(x, y)$ between the two points. This is the **first constraint** we impose on the convolution kernel. It implies that there is a real-valued function $u(x, \rho(x, y))$ such that

$$w(x, y) = u(x, \rho(x, y)). \quad (4)$$

Substituting Eq. (4) into Eq. (3) yields

$$\xi(x) = \int_D u(x, \rho(x, y)) \alpha(y) dy. \quad (5)$$

In the remainder of this subsection, we develop spectral representations of the process defined by Eq. (5) and its covariance function. We perform the spectral (Fourier-Legendre) expansion of $u(x, \rho)$ with x being fixed:

$$u(x, \rho) = \sum_{\ell=0}^{\ell_{\max}} \frac{2\ell+1}{4\pi} \sigma_{\ell}(x) P_{\ell}(\cos \rho), \quad (6)$$

where ℓ_{\max} is the bandwidth, $\sigma_{\ell}(x)$ are real-valued functions (we call them *spectral functions*) and P_{ℓ} are the Legendre polynomials. We substitute $\rho = \rho(x, y)$ into Eq. (6) and apply the addition theorem of spherical harmonics,

$$\sum_{m=-\ell}^{\ell} Y_{\ell m}(x) Y_{\ell m}^*(y) = \frac{2\ell+1}{4\pi} P_{\ell}(\cos \rho(x, y)), \quad (7)$$

where $Y_{\ell m}$ is the spherical harmonic, ℓ its degree (sometimes called the total wavenumber), m its order (the meridional wavenumber), and the asterisk denotes complex conjugation.

As a result, Eq. (6) can be rewritten as

$$u(x, \rho(x, y)) = \sum_{\ell=0}^{\ell_{\max}} \sigma_{\ell}(x) \sum_{m=-\ell}^{\ell} Y_{\ell m}(x) Y_{\ell m}^*(y). \quad (8)$$

Then, we make use of the spectral expansion of the white noise,

$$\alpha(y) = \sum_{\ell=0}^{\ell_{\max}} \sum_{m=-\ell}^{\ell} \tilde{\alpha}_{\ell m} Y_{\ell m}(y), \quad (9)$$

where $\tilde{\alpha}_{\ell m}$ are mutually uncorrelated complex-valued random variables with $\mathbb{E} \tilde{\alpha}_{\ell m} = 0$ and $\text{Var} \tilde{\alpha}_{\ell m} = 1$. More specifically, $\tilde{\alpha}_{\ell 0}$ are real-valued and all the other $\tilde{\alpha}_{\ell m}$ are complex circularly symmetric random variables (e.g. Gallager, 2013, section 3.7) such that $\tilde{\alpha}_{\ell, -m} = \tilde{\alpha}_{\ell m}^*$.

Substituting Eqs. (8) and (9) into (5) and utilizing orthonormality of spherical harmonics, we obtain the basic spectral representation of the random process that satisfies Eq. (5):

$$\xi(x) = \sum_{\ell=0}^{\ell_{\max}} \sum_{m=-\ell}^{\ell} \sigma_{\ell}(x) \tilde{\alpha}_{\ell m} Y_{\ell m}(x). \quad (10)$$

The spatial covariances $B(x, x') = \mathbb{E} \xi(x) \xi(x')$ can readily be obtained from Eq. (10) by taking into account that all $\tilde{\alpha}_{\ell m}$ are mutually uncorrelated and again applying the addition theorem of spherical harmonics. Equivalently, we may use Eq. (5) for this purpose so that

$$B(x, x') = \int_D u(x, \rho(x, y)) u(x', \rho(x', y)) dy = \sum_{\ell=0}^{\ell_{\max}} \frac{2\ell+1}{4\pi} \sigma_{\ell}(x) \sigma_{\ell}(x') P_{\ell}(\cos \rho(x, x')). \quad (11)$$

If $u(x, \rho) = u(\rho)$ or, equivalently, if all spectral functions $\sigma_{\ell}(x)$ do not depend on x , then $\xi(x)$ becomes stationary, in which case the process variance $\text{Var} \xi(x)$ can be obtained by summing $\frac{2\ell+1}{4\pi} \sigma_{\ell}^2$ known as the variance spectrum (energy per degree ℓ). Besides, $f_{\ell} = \sigma_{\ell}^2$ is often called the *modal* spectrum (energy per ‘mode’, i.e., per pair (ℓ, m)). In the non-stationary case, in which σ_{ℓ} depends on x , we call $f_{\ell}(x) = \sigma_{\ell}^2(x)$ the *local spectrum*. Equation (11) shows that the non-stationary process variance can be computed in the same way as the stationary one: $\text{Var} \xi(x) = \sum_{\ell} \frac{2\ell+1}{4\pi} \sigma_{\ell}^2(x) \equiv \sum_{\ell} \frac{2\ell+1}{4\pi} f_{\ell}(x)$.

Non-negative definiteness (i.e., validity) of the model’s non-stationary covariance function $B(x, x')$ is established in Appendix B.

The formulation of the spatial model Eq. (10) on the circle is outlined in Appendix C.

3.4. Positive definite convolution kernel

In Appendix D, we consider the process ξ_{statio} defined by the process convolution model with the *stationary* kernel $u(\rho)$, see Eq. (D.3). As discussed above, the kernel is not unique given the covariance function of the process. In an attempt to select a unique kernel, we impose a computationally motivated *spatial localization* requirement, seeking $u(\rho)$ that has the smallest spatial scale. We show that there are three solutions to this optimization problem, all of them having the same shape of the kernel. One of these three equivalent solutions is *positive definite*, which we select. So, if a stationary process with a given covariance function is defined by the convolution $\xi = u * \alpha$ (where α is the white noise), then the most spatially localized kernel u can be taken to be a positive definite function of distance.

Motivated by this result, we postulate that in the non-stationary case, for any x , the kernel $u(x, \rho)$ is a *positive definite function* of distance ρ . As a consequence, $\sigma_{\ell}(x) > 0$ both in the spherical and in the circular case. This constitutes our **second constraint** imposed on the general process convolution model.

3.5. Uniqueness of the convolution kernel

In Appendix E we prove that under the local isotropy constraint (section 3.3), the positive-definite-kernel constraint (section 3.4), and a technical assumption, the kernel

$u(x, \rho)$ is uniquely determined by the model's non-stationary covariance function $B(x, x')$.

Aiming to further reduce the effective number of degrees of freedom, we introduce two additional regularizing constraints (presented in the following two subsections). Both are smoothness constraints, one imposed in physical space and the other in spectral space.

3.6. Local stationarity

The notion of local stationarity has been defined differently by different authors (most often for processes on the real line). The general idea is that a locally stationary process can be approximated by a stationary process locally, i.e., in the vicinity of any point in time (Mallat et al., 1998).

Starting from Dahlhaus (1997), the common approach to studying local stationarity is the use of the ‘infill’ asymptotics, in which the non-stationarity of the process $\xi(t)$ is determined by a function, $a(\tau)$, defined in rescaled time $\tau \in [0, 1]$. Then, $\xi(t)$ is considered on $t = T\tau \in [0, T]$ as $T \rightarrow \infty$. The prototypical example is the autoregressive process $\xi_{t+1} = a(t/T)\xi_t + \eta_t$, where $t = 0, 1, \dots, T$ and $\eta_t \sim N(0, 1)$. The mean time-scale of this process does not change as $T \rightarrow \infty$, whereas the non-stationarity pattern is infinitely stretched. As a result, for any point τ_0 in rescaled time, ξ_t can be approximated by a stationary process in the vicinity of the point $t = T\tau_0$, with the approximation error vanishing as $T \rightarrow \infty$.

On compact manifolds like the circle or the sphere, this rescaling cannot be used because of their compactness. Therefore, we propose a different approach to local stationarity with our process convolution model.

Recall that the non-stationary random process in question is defined by Eq. (10), where the non-stationarity comes from the dependencies of σ_ℓ on the location x . We assume that these are caused by $\sigma_\ell(x)$ *being random processes themselves*. More specifically, we postulate that $\sigma_\ell(x)$ are, jointly, a vector-valued mean-square differentiable *stationary random process*. Then, we obtain spatial non-stationarity by considering the conditional distribution of the process ξ given all $\sigma_\ell(\cdot)$. We denote this conditioning by $\xi | \sigma$. In other words, we model spatial non-stationarity by assuming that $\boldsymbol{\sigma}(x) = \{\sigma_\ell(x)\}_{\ell=0}^{\ell_{\max}}$ is a random process and examine $\xi(x) = \sum \sigma_\ell(x) \tilde{\alpha}_{\ell m} Y_{\ell m}(x)$ with a random realization of $\boldsymbol{\sigma}(x)$ being fixed. Equivalently, conditioning on the convolution kernel $u(x, \rho)$ (viewed as a Hilbert space-valued stationary random process of x) can be considered.

We define the length scale Λ_ℓ of the stationary process $\sigma_\ell(x)$ as the micro-scale from the equation

$$\mathbb{E} \left| \frac{\partial \sigma_\ell}{\partial x} \right|^2 = \frac{\text{Var } \sigma_\ell}{\Lambda_\ell^2} \quad (12)$$

(where $\partial/\partial x$ is the spatial gradient operator) and the overall *non-stationarity length* scale Λ from

$$\sum_{\ell=-\ell_{\max}}^{\ell_{\max}} \frac{2\ell + 1}{4\pi} \frac{\text{Var } \sigma_\ell}{\Lambda_\ell^2} = \frac{\mathbb{E} \xi^2}{\Lambda^2}. \quad (13)$$

It is straightforward to prove that the process $\xi(x)$ is indeed locally stationary in the sense that for any point x_0 , it can be approximated by a stationary process $\zeta(x; x_0)$, with the mean-square approximation error vanishing as $\Lambda \rightarrow \infty$. To do so, we define $\zeta(x; x_0) = \sum_{\ell} \sigma_{\ell}(x_0) \sum_m \tilde{\alpha}_{\ell m} Y_{\ell m}(x)$, that is, with the argument of σ_{ℓ} fixed at x_0 . Then, the mean approximation error variance is

$$\begin{aligned} \mathbb{E} \mathbb{E}[(\xi(x) - \zeta(x; x_0))^2 | \sigma] &= \sum_{\ell=-\ell_{\max}}^{\ell_{\max}} \frac{2\ell+1}{4\pi} \mathbb{E}[\sigma_{\ell}(x) - \sigma_{\ell}(x_0)]^2 \leq \\ &\rho^2(x, x_0) \sum_{\ell=-\ell_{\max}}^{\ell_{\max}} \frac{2\ell+1}{4\pi} \mathbb{E} \left| \frac{\partial \sigma_{\ell}}{\partial x} \right|^2 = \frac{\rho^2(x, x_0)}{\Lambda^2} \mathbb{E} \xi^2. \end{aligned} \quad (14)$$

Here the last equality is due to Eqs. (12) and (13).

Theoretically, Eq. (14) shows that in the locally stationary limit, as $\Lambda \rightarrow \infty$ and $\mathbb{E} \xi^2 = \text{const}$, the approximation error vanishes globally, that is, for all points x on the sphere. Practically, we wish a locally stationary process to have correlations that, approximately, depend only on the distance ρ for distances *within a process length scale*, L (whose definition is application dependent). Then, Eq. (14) implies that in applications, we may assume local stationarity if $L \ll \Lambda$.

Local stationarity is the **third constraint** we introduce. Note that the slow variation of the kernel $u(x, \rho)$ with location x (compared with its variation as a function of distance ρ), actually, justifies the term ‘local spectrum’ (section 3.3) introduced first in the time series context by Priestley (1965), who called it evolutionary spectrum.

3.7. Smoothness of local spectra

Studies of real-world spatio-temporal processes showed that spatial (and temporal) spectra in turbulent flows are quite smooth, exhibiting typically a power-law behavior at large wavenumbers (e.g. Gage and Nastrom, 1986, Trenberth and Solomon, 1993). For this reason and to regularize the spatial model by further reducing its effective number of degrees of freedom, we postulate that the spectral functions $\sigma_{\ell}(x)$ are *smooth functions of the wavenumber ℓ* — this is our **fourth constraint**. As for a measure of smoothness of the local spectrum, in section 4.4 we propose one definition aimed to facilitate the proof of consistency of the estimator, and in section 4.5 we impose another, more practically oriented smoothness constraint by placing a prior distribution over the functions $\sigma_{\ell}(x)$.

Note that if $\sigma_{\ell}(x)$ is smooth as a function of ℓ , then its inverse spectral transform, that is, the convolution kernel $u(x, \rho)$, rapidly decreases with distance ρ . The rapidly decreasing convolution kernel prohibits spurious (and damaging in the analysis) long-distance correlations, which are present in sample (ensemble) covariances.

3.8. Summary of constraints

- (1) The convolution kernel has the locally isotropic form $w(x, y) = u(x, \rho(x, y))$.
- (2) The convolution kernel $u(x, \rho)$ is a positive definite function of distance ρ .

- (3) The process $\xi(x)$ is locally stationary.
- (4) The spectral functions $\sigma_\ell(x)$ are smooth functions of the wavenumber ℓ .

Below we refer to the above spatial model that obeys these four constraints as the Locally Stationary Convolution Model.

3.9. Spatial discretization

When used in the analysis to specify the prior covariance matrix (see section 5), the model needs to be discretized, that is, defined on a spatial grid. This is done by approximating the convolution integral in Eq. (5) as a sum. In this proof-of-concept study, we restricted the band-limited field $\xi(\cdot)$ to the regular latitude-longitude grid described in section 6.3. In a real-world high-dimensional application, the key point will be to design a numerical scheme in which, for each grid point x , the sum runs over a relatively small number of neighboring grid points. This will make the resulting matrix operator (denoted in section 5 as $\widehat{\mathbf{W}}$) *sparse*. Developing such an algorithm will be done in a follow-up study using a multi-scale/multigrid approach. A non-ensemble version of the multi-scale technique (with static spatial covariances) has been successfully tested in an operational meteorological data assimilation system (Tsyrlunikov et al., 2019). In that application, we employed a non-regular, quasi-uniform grid on the sphere to maintain sparsity of the $\widehat{\mathbf{W}}$ matrix in the polar regions, where the latitude-longitude grid is dense.

4. Estimation of the spatial model from the ensemble

In this section, we propose an estimator of the Locally Stationary Convolution Model from an ensemble (sample) of K independent identically distributed random field realizations (ensemble perturbations), $\{\xi^{(k)}(x)\}_{k=1}^K$. The goal is to estimate the convolution kernel $u(x, \rho(x, \cdot))$, where x runs over points of the spatial grid. In this study, we confine ourselves to *point estimates* of $u(x, \rho(x, y))$. We recall that for any x , the convolution kernel $u(x, \rho(x, y))$ is the inverse Fourier-Legendre transform of the spectral functions $\sigma_\ell(x) = \sqrt{f_\ell(x)}$, see Eq. (6), so it suffices to estimate $\sigma_\ell(x)$. We seek an estimator that is applicable to high-dimensional problems with 10^9 degrees of freedom or more.

The outline of this section is as follows. We start by motivating, in section 4.1, the application of a *multi-scale bandpass filter* to ensemble perturbations with the intention to extract aggregated local spectra. Then we explain how the (computationally intractable) likelihood of the raw data $\{\xi^{(k)}(x_i)\}$ (with $i = 1, 2, \dots, n_x$ and $k = 1, 2, \dots, K$) can be approximated by a much simpler *composite likelihood* (Lindsay, 1988, Reid, 2013), see section 4.2. After that, we propose a two-stage estimator, section 4.3. At the first stage, aggregated local spectra are computed as sample variances of the bandpass filtered fields. At the second stage, the local spectra are disaggregated from these sample variances for each grid point independently. We explore two approaches to the disaggregation. The first one, based on a linear estimating equation, allows us to prove *consistency* of the

resulting model estimator, section 4.4. The second one relies on a *neural network* to estimate the local spectra, section 4.5. Finally, the spatially varying convolution kernel is computed from the estimated local spectra (by substituting them into Eq. (6)) and used in the analysis. The estimator’s workflow is summarized in section 4.6.

Our approach can be regarded as a modification of the technique by Priestley (1965), who, for each frequency ω_0 , applied a narrow bandpass filter centered at ω_0 to a non-stationary time series and used the filter’s mean-squared output as the estimate of his evolutionary spectrum at the same frequency. Another relevant reference is Marinucci and Peccati (2011, chapter 11), who applied a wavelet filter to extract the variance spectrum of an isotropic random field on the sphere.

4.1. Multiple spatial bandpass filters

In the *stationary* case, it is straightforward to derive from Eq. (D.1) that the likelihood of the data, i.e., the ensemble of random-field realizations $\{\xi^{(k)}(x)\}_{k=1}^K$, given the spectrum $\{f_\ell\}_{\ell=0}^{\ell_{\max}}$, depends on the data through the set of spectral-space single-wavenumber sample variances $\hat{v}_\ell = \frac{1}{K} \sum_{k=1}^K \sum_{m=-\ell}^{\ell} |\tilde{\xi}_{\ell m}^{(k)}|^2$. (Note that with $K = 1$, the set of \hat{v}_ℓ is the spherical analog of what is known in the time series literature as the *periodogram*.) Therefore, the vector of the single-wavenumber sample variances \hat{v}_ℓ (with $\ell = 0, 1, \dots, \ell_{\max}$) is a sufficient statistic for the spectrum f_ℓ . This implies that no information on the true spectrum is lost if we switch from the raw ensemble to the set of \hat{v}_ℓ .

In the locally stationary case, however, relying on the single-wavenumber sample variances is not a good idea because isotropic spatial filters that isolate individual spectral components have non-local impulse response functions. Indeed, the filter’s response function is the backward Fourier-Legendre transform of its spectral transfer function $H(\ell)$. If $H(\ell)$ is equal to one at the single wavenumber ℓ_0 and zero otherwise, then $h(\rho(x, y)) = \frac{2\ell_0+1}{4\pi} P_{\ell_0}(\cos(\rho(x, y)))$ (see Eq. (6)). Due to the non-locality of $P_{\ell_0}(\cos(\rho))$, the filtered signal, which is the convolution of the input signal with $h(\rho(x, y))$, will contain significant contributions from the whole global domain. This is not acceptable in the non-stationary case, in which we seek primarily local contributions.

Therefore, instead of $\ell_{\max} + 1$ filters that isolate individual spectral components, we introduce a smaller number, J , of *bandpass filters* \mathcal{H}_j ($j = 1, 2, \dots, J$) that have broader spectral transfer functions $H_j(\ell)$. Broader transfer functions imply spatially localized impulse response functions, $h_j(\rho)$, as required with a locally stationary process. Different $H_j(\ell)$ represent partly overlapping spectral *bands* that, jointly, cover the whole spectral range. Importantly, the filtered signals $\varphi_j(x) = \mathcal{H}_j \xi(x)$ evaluated at a single grid point x contain spatial information about the signal in question $\xi(x)$ in a vicinity of x . So, we can infer local spatial statistics of $\xi(\cdot)$ from the set of $\varphi_j(x)$ evaluated at a single point x .

On the other hand, the likelihood of $\xi(x)$ is computationally intractable because it depends on a huge number of parameters, that is, the spectral functions $\sigma_\ell(x)$ evaluated at all spatial grid points x . In the estimator, we propose to replace the raw data $\xi(x)$ with the set of filtered fields $\varphi_j(x)$ and, crucially, switch from the likelihood of the whole spatial

fields to the likelihood of $\boldsymbol{\varphi}(x) = \{\varphi_j\}_{j=1}^J(x)$ evaluated at each grid point x individually. As we show below in section 4.2, the likelihood of $\boldsymbol{\varphi}(x)$ depends, approximately, on $\sigma_\ell(x)$ evaluated at the same point. This reduces the global estimation problem to a series of local problems of much smaller dimensionality.

We note that broadening the filters' transfer functions reduces the spectral resolution of the estimator. But, according to our fourth constraint (section 3.8), the spatial spectra are smooth, so high spectral resolution is not needed. For more details on the bandpass filters, see sections 4.2, 6, and 7.

4.2. Approximate likelihood of the filtered signals

Consider the application of the filter \mathcal{H}_j to $\xi(x)$. As shown in Appendix F, if $L_{\mathcal{H}_j} \ll \Lambda$, (where $L_{\mathcal{H}_j}$ is the length scale of the j th bandpass filter's impulse response function and Λ the process's non-stationarity length scale introduced in section 3.6), then the spectral-space filtering can be approximately performed as if the random field $\xi(x)$ were stationary:

$$\varphi_j(x) = \sum_{\ell=0}^{\ell_{\max}} \sum_{m=-\ell}^{\ell} H_j(\ell) \sigma_\ell(x) \tilde{\alpha}_{\ell m} Y_{\ell m}(x) + \delta\varphi_j(x), \quad (15)$$

where $\delta\varphi_j(x)$ is the discrepancy due to the non-stationarity, which has mean zero and variance $\mathbb{E}|\delta\varphi_j|^2 = O(L_{\mathcal{H}_j}^2/\Lambda^2)$ as $L_{\mathcal{H}_j}/\Lambda \rightarrow 0$, see Eq. (F.9). The key point here is that, up to the small discrepancy, the distribution of $\varphi_j(x)$ is determined by the spectral functions $\sigma_\ell(x)$ evaluated at the same spatial point x only.

Specifically, given $\boldsymbol{\sigma}(x) = \{\sigma_\ell(x)\}_{\ell=0}^{\ell_{\max}}$, the random vector $\boldsymbol{\varphi}(x) = \{\varphi_j(x)\}_{j=1}^J$ is multivariate Gaussian with mean zero and covariances $\mathbb{E}\varphi_j(x)\varphi_{j'}(x) = \sum_{\ell=0}^{\ell_{\max}} \frac{2\ell+1}{4\pi} H_j(\ell)H_{j'}(\ell) \sigma_\ell^2(x) + O(L_{\mathcal{H}_j}/\Lambda) + O(L_{\mathcal{H}_{j'}/\Lambda})$. If the small $O(\cdot)$ contributions are neglected, this fully determines its likelihood $p(\boldsymbol{\varphi}(x) | \boldsymbol{\sigma}(x))$ for each grid point x .

For computational reasons, instead of relying on the exact likelihood of $\boldsymbol{\varphi}(x)$, we neglect cross-correlations between different $\varphi_j(x)$ at the same grid point x . Besides, considering the likelihood grid point by grid point, we, effectively, neglect cross-correlations between different grid points as well. As a result, we end up with the approximate Gaussian *composite likelihood* defined for each spatial grid point separately. For the whole sample of K ensemble members the composite log-likelihood reads (up to a constant)

$$c\ell[\boldsymbol{\sigma}(x)] = -\frac{K}{2} \sum_{j=1}^J \left(\ln v_j(x) + \frac{d_j(x)}{v_j(x)} \right), \quad (16)$$

where $v_j(x) = \mathbf{Var}(\varphi_j(x) | \boldsymbol{\sigma}(x))$ are the *band variances* and $d_j(x)$ are their sample counterparts, the *sample band variances*. From Eq. (15),

$$v_j(x) = \underline{v}_j(x) + \delta_j^{\text{nsa}}(x), \quad (17)$$

where

$$\underline{v}_j(x) = \sum_{\ell=0}^{\ell_{\max}} \frac{2\ell+1}{4\pi} H_j^2(\ell) \sigma_\ell^2(x) \quad (18)$$

is the approximate band variance and $\delta_j^{\text{nsta}}(x) = O(L_{\mathcal{H}_j}/\Lambda)$ is the approximation error.

According to Eqs. (17) and (18), the band variances $v_j(x)$ are, approximately, the aggregated local spectra $f_\ell(x) = \sigma_\ell^2(x)$. Jointly, $\{v_j(x)\}_{j=1}^J$ bear all the information on the (smooth) local spectrum we wish to infer from the sample (up to the small discrepancy δ_j^{nsta}). On the other hand, the sample band variances $d_j(x)$ summarize information on the local spectrum contained in the raw sample $\{\xi^{(k)}(x)\}_{k=1}^K$ (up to the sampling error: $d_j(x) = v_j + \delta_j^{\text{samp}}$). Note that $v_j(x)$ and $d_j(x)$ are the locally stationary counterparts of the stationary spectrum f_ℓ and periodogram \widehat{v}_ℓ , respectively, see section 4.1.

4.3. A two-stage estimator

Equations (16)–(18) suggest the following two-stage estimator applied point-wise, for each spatial grid point x independently:

1. Compute the sample band variances $d_j(x)$:
 - (a) Bandpass filter K ensemble perturbations, $\xi^{(k)}(\cdot)$, $k = 1, 2, \dots, K$. The result is the JK filtered fields $\varphi_j^{(k)}(\cdot)$, $j = 1, 2, \dots, J$.
 - (b) At each spatial grid point x , compute sample variances of $\varphi_j^{(k)}(x)$, that is, the sample band variances $d_j(x)$.
2. Recover $\boldsymbol{\sigma}(x)$ from the set of sample band variances $\mathbf{d}(x) = \{d_j(x)\}_{j=1}^J$. We consider two ways of doing so.

- The first approach relies on linearity of Eq. (18) with respect to the local spectrum $f_\ell(x) = \sigma_\ell^2(x)$ and neglect of $\delta_j^{\text{nsta}}(x)$. We rewrite Eqs. (17) and (18) as

$$\mathbf{v}(x) \approx \boldsymbol{\Omega} \mathbf{f}(x), \quad (19)$$

where $\mathbf{v}(x) = \{v_j(x)\}_{j=1}^J$, $\mathbf{f}(x) = \{f_\ell(x)\}_{\ell=0}^{\ell_{\max}}$, and $\boldsymbol{\Omega}$ is the J by $\ell_{\max} + 1$ matrix with the known entries $(\boldsymbol{\Omega})_{j\ell} = \frac{2\ell+1}{4\pi} H_j^2(\ell)$.

Equations (16) and (19) imply that the composite likelihood $cl[\boldsymbol{\sigma}(x)]$ is approximately maximized at any vector $\mathbf{f}(x)$ that satisfies

$$\mathbf{d}(x) = \boldsymbol{\Omega} \mathbf{f}(x). \quad (20)$$

This estimating equation is underdetermined and thus needs to be regularized. In section 4.4, we do so by imposing a deterministic smoothness constraint and show that the error of the resulting solution vanishes in the large sample limit so that the whole estimator is consistent.

- The second approach is to rely on the available unapproximated nonlinear data generation (or ‘forward’) model $\boldsymbol{\omega}$ such that

$$\mathbf{d}(x) = \boldsymbol{\omega}(\boldsymbol{\sigma}(\cdot), x). \quad (21)$$

In section 4.5, we build an optimized neural-network based mapping $\mathbf{d}(x) \mapsto \boldsymbol{\sigma}(x)$ that approximates the classical Bayes estimator.

In the next two subsections, we elaborate on the two implementations of stage 2, one based on Eq. (20) and the other on Eq. (21).

4.4. Consistency of the linear estimator

To solve the underdetermined Eq. (20) we employ the deterministic regularization technique known as ‘spectral cut-off’, (e.g. Cavalier, 2011). Specifically, we impose a smoothness constraint that requires the local spectrum to be band-limited as a function of log-wavenumber, see below.

To formulate the smoothness constraint, we switch from wavenumbers ℓ to log-wavenumbers $\vartheta(\ell) = \pi \log(\ell + 1) / \log(\ell_{\max} + 1)$. We set $\vartheta(-\ell) = \vartheta(\ell)$ so that ϑ ranges from $-\pi$ to π . Then, we let $F(\vartheta(\ell)) = f_\ell$. We also assume that the spatial resolution is large enough and the spectra are smooth enough to treat $F(\vartheta)$ as a function of continuous argument on the unit circle.

Definition 4.1. We say that the spectrum f_ℓ is *smooth with the bandwidth M* if $F(\vartheta)$ is a *band-limited function*:

$$F(\vartheta) = \sum_{m=-M}^M \tilde{F}_m e^{im\vartheta}, \quad (22)$$

where \tilde{F}_m is the the Fourier transform of $F(\vartheta)$ and M is a positive integer.

The smoothness constraint in this section is the assumption that the spectrum is *smooth with the bandwidth M* (the lower M , the smoother $F(\cdot)$). Since $F(\vartheta)$ is an even real-valued function, it follows from Eq. (22) that the spectrum becomes parameterized with the vector, \mathbf{p} , of $M + 1$ real-valued parameters: $\mathbf{f} = \mathbf{\Psi}\mathbf{p}$ with the known $(\ell_{\max} + 1) \times (M + 1)$ matrix $\mathbf{\Psi}$. Therefore, we can select $J = M + 1$ bandpass filters so that the number of unknowns becomes equal to the number of data (the sample band variances \mathbf{d}). Then, Eq. (20) reduces to a system of linear algebraic equations:

$$\mathbf{d} = \mathbf{\Omega}\mathbf{\Psi}\mathbf{p}. \quad (23)$$

The solution to this equation is subject to both sampling error $\boldsymbol{\delta}^{\text{samp}}(x) = \{\delta_j^{\text{samp}}(x)\}_{j=1}^J$ and non-stationarity error $\boldsymbol{\delta}^{\text{nsta}}(x) = \{\delta_j^{\text{nsta}}(x)\}_{j=1}^J$. In Appendix G we prove that their impact on the estimated local spectrum results in the estimation error that can be made arbitrarily small by increasing the sample size and appropriately increasing spectral widths of the filters. More specifically, we prove the following

Proposition 4.1. *Let the local spectrum be smooth in terms of Definition 4.1 with the bandwidth M . Then, for every sample size K , there exist $M + 1$ bandpass filters such that the solution to Eq. (23) yields vanishing mean absolute error in the local spectrum as $K \rightarrow \infty$. The spectral widths of the filters (in terms of log-wavenumbers) grow with the increasing sample size K as K^γ for some $\gamma > 0$.*

Due to Markov’s inequality, Proposition 4.1 entails that the estimator is consistent.

Two points are worth noting. First, while proving consistency we found that the widths of the bandpass filters’ transfer functions are to be selected depending on the ensemble size: the bigger the ensemble, the broader bandpass filters’ spectral transfer functions can and should be taken. Broader transfer functions imply more localized filters’ impulse response functions and therefore less non-stationarity error δ_j^{nsta} in $d_j(x)$. This allows the estimator to extract increasingly detailed spatially non-stationary structures, resulting in more accurately estimated local spectra. Second, this conclusion has important consequences not just for the theoretical proof of consistency but also for a more practical implementation of stage 2 of the estimator (presented in the next subsection), where such a conclusion could not be derived analytically.

4.5. Neural Bayes estimator of local spectrum from sample band variances

We recall that the second implementation of stage 2 of the estimator, see section 4.3, abandons the linear approximation made in the previous subsection and relies on Eq. (21). This makes the ‘forward model’ more accurate and reduces the noise level — but at the expense of rendering the problem nonlinear and high-dimensional. Indeed, the exact expression for the filtered process $\varphi_j(x)$ is

$$\varphi_j(x) = \int_D h_j(\rho(x, y)) \xi(y) dy = \sum_{\ell=0}^{\ell_{\max}} \sum_{m=-\ell}^{\ell} \tilde{\alpha}_{\ell m} \int_D h_j(\rho(x, y)) \sigma_{\ell}(y) Y_{\ell m}(y) dy, \quad (24)$$

where Eq. (10) was used. From Eq. (24), it follows that the respective band variance $v_j(x) = \mathbb{E}(|\varphi_j(x)|^2 | \boldsymbol{\sigma}(\cdot))$ depends not only on $\sigma_{\ell}^2(x)$ (i.e., evaluated at the same point x at which φ_j is evaluated), but also on the products $\sigma_{\ell}(y) \sigma_{\ell}(y')$ for all pairs of spatial points (y, y') . This justifies writing the data model as $\mathbf{d}(x) = \boldsymbol{\omega}(\boldsymbol{\sigma}(\cdot), x)$, as in Eq. (21). In high dimensions, this also implies that the likelihood of $\mathbf{d}(x)$ is numerically intractable. Therefore we build an estimator of $\boldsymbol{\sigma}(x)$ that does not require evaluations of the likelihood function.

We start with postulating that the estimator is of the form

$$\hat{\boldsymbol{\sigma}}(x) = \boldsymbol{\Phi}(\mathbf{d}(x), \boldsymbol{\psi}), \quad (25)$$

where $\boldsymbol{\Phi} = (\Phi_0, \Phi_1, \dots, \Phi_{\ell_{\max}})$ is a known vector function of its arguments and $\boldsymbol{\psi}$ is the unknown vector of the estimator’s parameters.

We place a prior distribution, $p(\boldsymbol{\sigma}(\cdot))$, over the unknown spectral functions $\sigma_{\ell}(\cdot)$. Note that $p(\boldsymbol{\sigma}(\cdot))$ is a hyperprior in the context of ensemble filtering. The hyperprior $p(\boldsymbol{\sigma}(\cdot))$ encodes the four constraints of our spatial model, see section 3.8.

To optimize the estimator, that is, to find the best $\boldsymbol{\psi}$, we introduce a *loss function* chosen to be quadratic:

$$\mathcal{L}[\hat{\boldsymbol{\sigma}}(\cdot), \boldsymbol{\sigma}(\cdot)] = \frac{1}{|D|} \int_D \sum_{\ell=0}^{\ell_{\max}} \frac{2\ell + 1}{4\pi} (\hat{\sigma}_{\ell}(x) - \sigma_{\ell}(x))^2 dx \quad (26)$$

(where $|D|$ is the area of the analysis domain D). Aiming at the Bayes point estimator (e.g. Lehmann and Casella, 2006, Chapter 4), we minimize the Bayes risk, that is, the expectation of the loss function with respect to the joint distribution of the parameter $\boldsymbol{\sigma}(\cdot)$ and the data $\mathbf{d}(\cdot)$:

$$\hat{\boldsymbol{\psi}} = \underset{\boldsymbol{\psi}}{\operatorname{argmin}} \mathbb{E} \mathcal{L}[\boldsymbol{\Phi}(\mathbf{d}(\cdot), \boldsymbol{\psi}), \boldsymbol{\sigma}(\cdot)], \quad (27)$$

getting the mean-square optimal estimator's parameters $\boldsymbol{\psi}$ and thus the optimal estimator defined by Eq. (25).

As the likelihood function is numerically intractable, we rely on a Monte-Carlo approximation to solve the optimization problem Eq. (27). In doing so, we assume that, first, there is a sample $\{\boldsymbol{\sigma}^i(\cdot)\}_{i=1}^I$ (I is its size) drawn from the hyperprior distribution $p(\boldsymbol{\sigma}(\cdot))$. Second, we exploit the fact that the data model $\boldsymbol{\omega}(\boldsymbol{\sigma}(\cdot), x)$ is *generative*: given $\boldsymbol{\sigma}^i(\cdot)$, it can simulate the data $\mathbf{d}^i(x_r)$, that is, J sample band variances at R randomly and uniformly selected grid points x_r , where $r = 1, 2, \dots, R$. Since, by construction, $\boldsymbol{\sigma}^i(\cdot) \sim p(\boldsymbol{\sigma}(\cdot))$ and $\mathbf{d}^i(\cdot) \sim p(\mathbf{d}(\cdot) | \boldsymbol{\sigma}^i(\cdot))$, we can approximate the expectation in Eq. (27) and the integral in Eq. (26) using Monte-Carlo:

$$\frac{1}{n_{\text{train}}} \sum_{i=1}^I \sum_{r=1}^R \sum_{\ell=0}^{\ell_{\text{max}}} \frac{2\ell+1}{4\pi} (\Phi_{\ell}(\mathbf{d}^i(x_r), \boldsymbol{\psi}) - \sigma_{\ell}^i(x_r))^2, \quad (28)$$

where $n_{\text{train}} = IR$ is the size of the training sample.

In this study, we choose a neural network as the parametric estimator $\boldsymbol{\Phi}(\mathbf{d}(x), \boldsymbol{\psi})$ with $\boldsymbol{\psi}$ being the collection of the net's weights and biases. If (i) the training sample $\{\boldsymbol{\sigma}^i(\cdot)\}_{i=1}^I$ is large and representative of the hyperprior $p(\boldsymbol{\sigma})$, (ii) the neural network is expressive enough, and (iii) the network optimizer is capable of finding a good proxy to the global minimum in Eq. (28), then the resulting estimator is a good approximation to the Bayes estimator. In this way, we obtain a *neural Bayes estimator* (Sainsbury-Dale et al., 2024). Technical details on the neural net are given in Appendix I.

The above neural-Bayes based disaggregation of the local spectrum from the set of sample band variances makes the technique applicable to high-dimensional problems. This is because (i) the estimator is local in space, that is, for each spatial point x , it maps $\mathbf{d}(x)$ to $\hat{\boldsymbol{\sigma}}(x)$, (ii) the resulting inference is *amortized*: no retraining is needed when new data arrive (unless there are changes in the hyperprior $p(\boldsymbol{\sigma}(\cdot))$ or in the ensemble size, or in the bandpass filters), (iii) the computations can be performed at different grid points in parallel.

It is worth reiterating that the estimator $\hat{\boldsymbol{\sigma}}(x) = \boldsymbol{\Phi}(\mathbf{d}(x), \boldsymbol{\psi})$ is trained only once, that is, its parameters $\boldsymbol{\psi}$ are optimized during the training and then kept fixed. The application of the estimator to fresh data $\mathbf{d}(x)$ requires no additional optimization. The neural network just transforms its input $\mathbf{d}(x)$ to the estimate $\hat{\boldsymbol{\sigma}}(x)$.

At the final stage of the estimator, we substitute $\hat{\boldsymbol{\sigma}}_{\ell}(x)$ into Eq. (6), getting the estimated spatially varying convolution kernel $\hat{u}(x, \rho)$, which, in turn, is used to compute

the $\widehat{\mathbf{W}}$ matrix, the key ingredient of the analysis algorithm, see section 5. The resulting spatial covariances, being determined by both the hyperprior (‘climatological’) distribution of the local spectra and the current ensemble, are non-stationary in time as well as in space. The uncertainty of the estimate is not taken into account in this study.

4.6. Summary of the estimator’s workflow

The practical estimation procedure consists of the following steps.

- Ensemble perturbations are bandpass filtered.
- Sample variances of the filtered fields are computed at each grid point.
- Local spectra are inferred from the sample band variances using a neural network grid point by grid point.
- The spatially varying convolution kernel is computed from the estimated local spectra.

5. LSEF: A Locally Stationary Ensemble Filter

The approach taken in this study is to use a spatial random field model to regularize the prior covariance matrix in sequential ensemble filtering. We call the resulting regularized EnKF the Locally Stationary Ensemble Filter (LSEF). Its forecast step coincides with that of the stochastic EnKF. The only difference between the LSEF’s and the traditional-EnKF’s analyses is how the proxy $\widehat{\mathbf{B}}$ to the true prior covariance matrix is defined. In this way, we can isolate the impact of our spatial-model-based regularization approach.

Testing and validation of our EnKF regularization technique can well be performed in the *linear* setting, avoiding unnecessary complications related to nonlinearities of observation and forecast models. For this reason, in the sequel we confine ourselves to the linear filtering setup. Appendix A reviews optimal and sub-optimal, ensemble and non-ensemble, time-discrete linear filters that are used in the numerical experiments presented below.

In the LSEF analysis scheme, we postulate that the forecast error obeys the Locally Stationary Convolution Model. The model, or more specifically the spatial convolution kernel $u(x, \rho(x, y))$, is estimated from the forecast ensemble and a hyperprior distribution as detailed in section 4. With the estimated convolution kernel $\widehat{u}(x, \rho(x, y))$ in hand, we build the space-discrete approximation of Eq. (5):

$$\boldsymbol{\xi} = \widehat{\mathbf{W}}\boldsymbol{\alpha}, \quad (29)$$

where $\boldsymbol{\xi}$ is the vector that represents the random field $\xi(x)$ evaluated on the spatial grid, $\boldsymbol{\alpha} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and $\widehat{\mathbf{W}}$ is a matrix with the entries

$$\widehat{w}_{ii'} = \widehat{u}(x_i, \rho(x_i, x_{i'})) \sqrt{\Delta x_{i'}}, \quad (30)$$

where $\Delta x_{i'}$ is the area of i' th grid cell. Note that Eqs. (29) and (30) approximate the integral in Eq. (5) using the simplest rectangle rule, more accurate approximations will be considered elsewhere. Equation (29) entails that

$$\widehat{\mathbf{B}} = \widehat{\mathbf{W}}\widehat{\mathbf{W}}^\top. \quad (31)$$

We note that the representation of the forecast-error covariance matrix $\widehat{\mathbf{B}}$ in the ‘square-root’ form, Eq. (31), is common in data assimilation practice because it provides efficient preconditioning of the analysis equations (e.g. Asch et al., 2016). The need for preconditioning follows from Eq. (A.4) because the matrix $\mathbf{B}^{-1} + \mathbf{H}^\top \mathbf{R}^{-1} \mathbf{H}$ to be inverted is, normally, ill-conditioned. Substituting Eq. (31) into Eq. (A.4) and rearranging the terms yields

$$\widehat{\mathbf{K}} = \widehat{\mathbf{W}} \left(\mathbf{I} + \widehat{\mathbf{W}}^\top \mathbf{H}^\top \mathbf{R}^{-1} \mathbf{H} \widehat{\mathbf{W}} \right)^{-1} \widehat{\mathbf{W}}^\top \mathbf{H}^\top \mathbf{R}^{-1}. \quad (32)$$

Now the matrix to be inverted, $\mathbf{I} + \widehat{\mathbf{W}}^\top \mathbf{H}^\top \mathbf{R}^{-1} \mathbf{H} \widehat{\mathbf{W}}$, is well conditioned (as a sum of the identity matrix and a non-negative definite matrix). This is how the LSEF analysis is done, that is, without explicitly computing $\widehat{\mathbf{B}}$. The LSEF analysis assimilates all available observations at once, as this is done in variational data assimilation schemes (e.g. Asch et al., 2016).

The representation of the covariance matrix $\widehat{\mathbf{B}}$ as $\widehat{\mathbf{W}}\widehat{\mathbf{W}}^\top$ also guarantees positive definiteness of $\widehat{\mathbf{B}}$ for any $\widehat{\mathbf{W}}$. This property can be utilized to perform *thresholding* of the $\widehat{\mathbf{W}}$ matrix, i.e., nullifying its small in modulus entries to make it sparse and facilitate fast computations without jeopardizing positive definiteness of the implied covariance matrix.

6. Numerical experiments with static analyses

In this section, we experimentally examine the above LSEF *analysis* technique in the setting with known synthetic ‘truth’ (ground truth) that obeys the Locally Stationary Convolution Model on the sphere. Besides, the hyperprior distribution of the model parameters was known in these experiments as well, so the neural Bayes estimator could be trained using the perfect training sample. This setting allowed us to reveal the best possible performance our approach can provide.

We compared deterministic analyses produced by the following schemes:

- True-B: the optimal analysis (which has access to the true \mathbf{B} matrix).
- LSEF-B: the new analysis proposed in this paper (with the neural Bayes estimator).
- Mean-B: analysis with time-mean \mathbf{B} (known in geosciences as 3D-Var).
- EnKF-B: stochastic EnKF with tuned localization.
- Hybrid-B: analysis with the half-sum of the time-mean \mathbf{B} and EnKF’s \mathbf{B} (known as EnVar).

See section 5 for the description of the LSEF analysis and Appendix A.3 and Appendix A.4 for the description of the other (reference) analyses. The quality of the analysis ensemble was not tested, it will be assessed in sequential filtering experiments presented in section 7.

6.1. Experimental methodology

Pseudo-random realizations of the ‘true’ spectral functions $\sigma_\ell(x)$ (which, we recall, are non-negative square roots of the local spectra $f_\ell(x)$) and the respective ‘true’ background error field $\xi(x)$ were generated following the Parametric Locally Stationary Convolution Model of ‘truth’ described in Appendix H. To give the reader an impression of the ground truth fields, we show in Fig. 1 two pseudo-random realizations⁴ of ξ for two different values of the non-stationarity length parameter μ_{NSL} , which is defined in Appendix H.3 as the ratio of the non-stationarity length scale Λ_{NSL} and the median process length scale L_{med} . The left panel of Fig. 1 shows the field with the small $\mu_{\text{NSL}} = 1$, that is, with the strong




Figure 1: Two realizations of the random field produced by the parametric model of ‘truth’. Left: $\mu_{\text{NSL}} = 1$. Right: $\mu_{\text{NSL}} = 3$. x-axis: longitude index, y-axis: co-latitude index.

spatial non-stationarity. The left panel of Fig. 1 displays a realization of the random field with the default (moderate) non-stationarity, $\mu_{\text{NSL}} = 3$. In both panels of Fig. 1, one can see a substantial degree of spatial non-stationarity, both in the field’s magnitude and length scale. With $\mu_{\text{NSL}} = 1$, Fig. 1(left), areas where the local length scale of ξ is small (seen as patchy spots) and where the local length scale is large (smooth spots) are both seen to be smaller in size compared with those in the right panel. This reflects a smaller non-stationarity length scale with $\mu_{\text{NSL}} = 1$ than with $\mu_{\text{NSL}} = 3$, as expected.

The true background-error vector $\boldsymbol{\xi}$ was computed by evaluating $\xi(\cdot) | \sigma$ on the analysis grid (defined below in section 6.3). We assumed (without loss of generality) that $\mathbf{x}^f = 0$ so that the simulated ground truth was $\mathbf{x} = \mathbf{x}^f - \boldsymbol{\xi} = -\boldsymbol{\xi}$ and the ensemble members were $\mathbf{x}^{\text{fe}(k)} = \mathbf{x}^f - \boldsymbol{\xi}^{(k)} = -\boldsymbol{\xi}^{(k)}$ (see Appendix A for notation).

Having computed the true $\sigma_\ell(x)$, we used Eq. (6) to compute the true kernel $u(x, \rho)$. After that, we built the true \mathbf{W} matrix following Eq. (30) and applied it to independent realizations of the white noise $\boldsymbol{\alpha}$ to generate the background-error vector $\boldsymbol{\xi}$ and the

⁴We used SHTools (Wieczorek and Meschede, 2018) to generate pseudo-random fields on the sphere.

ensemble perturbations $\boldsymbol{\xi}^{(k)}$, see Eq. (29). The true covariance matrix (used by the optimal True-B analysis) was $\mathbf{B} = \mathbf{W}\mathbf{W}^\top$.

With the truth $\mathbf{x} = -\boldsymbol{\xi}$ in hand, we generated observations \mathbf{x}^{obs} following Eq. (2), assimilated them using the five analysis techniques listed above, and compared the resulting deterministic analyses with the truth.

6.2. Details of the analysis schemes

The four sub-optimal analysis schemes differed only in the way the prior covariance matrix was specified. In the EnKF analysis, the localization matrix was specified using the popular Gaspari-Cohn correlation function (Gaspari and Cohn, 1999), whose length scale was tuned to obtain the best performance for each configuration of the model of truth and for each ensemble size separately.

In the Mean-B analysis, the static prior covariance matrix $\mathbf{B}_{\text{static}} = \mathbf{B}_{\text{mean}}$ was computed for each configuration of the model of truth as follows. We generated $n_{\text{repl}} = 33$ replicates of the true spectral functions $\sigma_\ell(x)$. Given each replicate, $\sigma_\ell^{(i)}(x)$, we generated a sample of $n_{\text{sample}} = 10$ realizations of $\xi | \sigma$: $\xi^{(i,k)}(\cdot)$, where $i = 1, 2, \dots, n_{\text{repl}}$ and $k = 1, 2, \dots, n_{\text{sample}}$. Then, we expanded each realization in spherical harmonics, getting the spectral coefficients $\tilde{\xi}_{\ell m}^{(i,k)}$. For each ℓ , we summed $|\tilde{\xi}_{\ell m}^{(i,k)}|^2$ over m in the range from $-\ell$ to ℓ , divided by 4π , and averaged over i and k , obtaining a sample estimate of the *mean variance spectrum* of ξ . From the mean variance spectrum, we computed the mean (stationary in space) covariance function and used it to build \mathbf{B}_{mean} .

In the Hybrid-B analysis, a convex combination of the static and localized sample covariance matrices was used: $\hat{\mathbf{B}}_{\text{hybr}} = (1 - w_e)\mathbf{B}_{\text{mean}} + w_e\hat{\mathbf{B}}_{\text{EnKF}}$, where $w_e = 0.5$ was found to be nearly optimal and used in all experiments.

In LSEF-B, the bandpass filters' spectral transfer functions were specified as

$$H_j(\ell) = \exp\left(-\left|\frac{\ell - \ell_j^c}{\Delta_{\mathcal{H}_j}}\right|^{q_{\text{shape}}}\right), \quad (33)$$

where ℓ_j^c is the central wavenumber of the j th filter, $\Delta_{\mathcal{H}_j}$ is the half-width parameter, and q_{shape} is the shape parameter. In the experiments described below, we took $q_{\text{shape}} = 3$ ($q_{\text{shape}} = 2$ also worked well). Both ℓ_j^c and $\Delta_{\mathcal{H}_j}$ were taken exponentially growing with the band index j . Figure 2 shows the spectral transfer functions and the respective impulse response functions for six bandpass filters used in the experiments presented in this section. The spatial filtering was performed in spectral space.

6.3. Experimental setup

The default setup of the parametric model of truth is specified in Appendix H.4.

The maximal wavenumber was $\ell_{\text{max}} = 50$. The analysis grid was of the regular latitude-longitude type with $\ell_{\text{max}} + 1$ latitudes (so that the mesh size in the latitudinal direction was $\Delta x = \pi/\ell_{\text{max}}$) and $2\ell_{\text{max}}$ longitudes. Note that according to the spherical sampling theorem by Driscoll and Healy (1994, Theorem 3), a band-limited (with the bandwidth ℓ_{max}) function defined on the sphere can be losslessly represented on the regular




Figure 2: Spectral transfer functions $H_j(\ell)$ (left) and their respective normalized impulse response functions $h_j(\rho)$ (right) for six bandpass filters.

latitude-longitude grid that has $n_{\text{lat}} = 2\ell_{\text{max}}$ latitudes and $n_{\text{lon}} = 2\ell_{\text{max}}$ longitudes. With our decaying spectra, reducing n_{lat} to $\ell_{\text{max}} + 1$ led to negligible errors.

Observations were located at randomly and independently sampled grid points selected with probabilities proportional to the spherical surfaces of the respective grid cells. The number of observations was $n_{\text{obs}} = \frac{1}{2}n_{\text{x}}$ (recall that n_{x} stands for the total number of grid points, that is, the length of the state vector \mathbf{x}). The observation-error variance was equal to the median background-error variance (we tried other choices and obtained qualitatively very similar results, not shown).

The ensemble size was between 5 and 80 with the default 20.

6.4. Training sample

Training the neural network involved in the estimation of the spatial model was done as follows. The Parametric Locally Stationary Convolution Model, see Appendix H, was used to generate $n_{\text{repl}} = 33$ replicates of the true spectral functions $\sigma_{\ell}(x)$. For each replicate of $\sigma^{(i)}(\cdot)$ (with $i = 1, 2, \dots, n_{\text{repl}}$), we generated an ensemble of K pseudo-random realizations of $\xi | \sigma$, from which we computed vectors of sample band variances $\mathbf{d}^{(i)}(\cdot)$ at all analysis grid points following Eq. (21). As a result, with n_{x} grid points, we had the training sample of $n_{\text{train}} = n_{\text{repl}} \cdot n_{\text{x}}$ pairs of vectors $\mathbf{d}^{(i)}(x)$ (inputs of the neural network) and $\sigma^{(i)}(x)$ (targets to be predicted by the neural network). The neural network was trained only once, for the default setup.

6.5. Results

The accuracy of each deterministic analysis was measured using the root-mean-square error (RMSE) with respect to the known truth. The RMSE was computed by averaging over the spatial grid and 100 independent analyses. In this computation, each point of the spatial grid had a weight proportional to the surface area of the respective grid cell. In the figures below, we show the analysis performance score

$$\frac{\text{RMSE} - \text{RMSE}_{\text{TrueB}}}{\text{RMSE}_{\text{TrueB}}}, \quad (34)$$

where $\text{RMSE}_{\text{TrueB}}$ is the RMSE of the optimal analysis. We also show 90% bootstrap confidence intervals (shaded). In the bootstrap, we resampled the (independent by construction) spatial-grid-averaged mean-square errors.

We compared the five deterministic analyses (listed at the beginning of section 6) in the default experimental setup and examined effects of varying the ensemble size K and the hyperparameters \varkappa (controls the strength of the non-stationarity, Appendix H.2) and μ_{NSL} (controls the length scale of the non-stationarity, Appendix H.3).

Figure 3 shows dependencies of the RMSEs on the ensemble size. It is seen that all ensemble-based analyses improved with the increasing K , as it should be. The advantage of the LSEF-B analysis over the competitors was very significant.




Figure 3: Analysis performance scores (the lower the better, =0 best possible), 90% bootstrap confidence bands (shaded). Dependence on ensemble size.

Figure 4 shows how the five analyses performed with the varying non-stationarity strength parameter \varkappa (defined in Appendix H.2). The advantage of LSEF-B over EnKF-B was big and nearly uniform for all $\varkappa > 1$ on the plot. It was Mean-B that suffered/benefited most from strong/weak non-stationarity, as expected. In the stationary regime (i.e., with $\varkappa = 1$), Mean-B was almost as accurate as True-B because \mathbf{B}_{mean} was a very accurate estimate of the constant (in this regime) \mathbf{B}_{true} . For this reason, Mean-B could not be rivaled by any ensemble-based analysis under stationarity.

Finally in this section, we study the impact of the non-stationarity length scale parameter μ_{NSL} (defined in Appendix H.3, see also section 6.1) on the performance of LSEF. Figure 5 displays RMSEs of the five analyses as functions of μ_{NSL} . Again, LSEF-B was significantly better than the other sub-optimal analyses, and that was the case even beyond the local stationarity assumption (section 3.6), which requires that




Figure 4: Analysis performance scores (the lower the better, =0 best possible), 90% bootstrap confidence bands (shaded). Dependence on the non-stationarity strength parameter \varkappa .

the non-stationarity length scale is much larger than the typical length scale of the field itself. Indeed, LSEF-B was substantially better than the competing techniques even when $\mu_{\text{NSL}} = 1$ (see Fig. 5), that is, when the non-stationarity length scale Λ_{NSL} was *equal* to the median length scale L_{med} (see Appendix H.3).

It is important to notice that in LSEF-B, there was no need to retrain the neural network for different setups. This is in contrast to EnKF-B, where we had to tune the localization length scale for each triple $K, \varkappa, \mu_{\text{NSL}}$. Similarly, \mathbf{B}_{mean} was computed for use in Mean-B and Hybrid-B for every pair $\varkappa, \mu_{\text{NSL}}$. This suggests that the neural net-based LSEF-B was more robust than the competing analyses.

The neural network was trained on the sample of the same size as we used to calculate \mathbf{B}_{mean} , yet it allowed LSEF-B to exhibit much better results than Hybrid-B. This suggests that the neural Bayes estimator is capable of *selective* blending of sample covariances with time-mean covariances. This is in contrast to Hybrid-B, which just mixes the two kinds of covariances with, effectively, equal weights for all spatial scales.

It is worth noting that the neural-Bayes-based LSEF-B analysis significantly outperformed an LSEF-B analysis based on the linear model estimator presented in section 4.4 (not shown). This was the expected outcome because the linear estimator (for which consistency is proven) is less accurate due to the linearity approximation.

6.6. Effects of covariance regularization on spatial covariances

We reproduced the above experimental methodology (the model of truth and the five analyses) on the circle and obtained very similar results (not shown). The code was written by the other author in a different programming language (R vs. Python).




Figure 5: Analysis performance scores (the lower the better, =0 best possible), 90% bootstrap confidence bands (shaded). Dependence on the non-stationarity length parameter μ_{NSL} .

Here we take advantage of the simpler circular domain to demonstrate how the regularization procedure proposed in this paper impacts spatial covariances. The spatial grid had 120 points and the ensemble size was 10.

Figure 6 shows typical examples of how the true covariances (available in the setting described above in this section) were estimated from the ensemble by (i) the localized sample covariances (used in the baseline stochastic EnKF method) and (ii) the model covariances generated by the estimated Locally Stationary Convolution Model. The left panel of Fig. 6 shows the variances and the right panel the correlations⁵.

Figure 6(left) shows that the sample (EnKF’s) variances are much noisier than the model variances. This illustrates the fact that the EnKF’s covariance localization cannot improve sample variances, which are quite uncertain for small ensemble sizes. Less noise in the model (LSEF) variances is due to the regularizing effect of the constrained spatial model. In particular, the local stationarity constraint implicitly allows for the use of neighboring covariances in the estimation. The specific mechanism for that is provided by the LSEF’s bandpass filters.

The right panel of Fig. 6 shows that, similar to the left panel, the model correlations contain much less noise than their localized sample counterparts. Another feature seen in this figure is that, in contrast to model correlations, localized sample correlations

⁵The realizations shown in Fig. 6 were typical in the following sense. Their single-realization mean absolute errors with respect to the truth were close (up to 10 percent difference) to the large-sample (of size 300) counterparts. For correlations, the mean absolute errors were computed by averaging over distances up to 15 mesh sizes.




Figure 6: Comparison of typical true, localized sample (EnKF), and model (LSEF) covariances on the circle. Left: variances at all grid points. Right: correlations with respect to a single grid point.

suffer from a negative bias in their length scale (see section 1.1, item 1). Furthermore, Fig. 6(right) illustrates that fact that the LSEF analysis produces almost noise-free large-distance background-error correlations without any ad-hoc localization (tapering) device. This emphasizes the usefulness of the smooth-local-spectrum constraint (item (4) in section 3.8) and the ability of the neural Bayes estimator to ensure this constraint is met. One can also notice in Fig. 6(right) that the regularizing constraint (1) (section 3.8), which requires the kernel to be locally isotropic, does not lead to isotropic correlations (note the asymmetry in both the true and model correlations).

Averaged over 300 realizations of the truth and the ensemble, the mean absolute error in the model (LSEF) variances was about 1.5 times less than the respective error in the sample variances. The mean absolute error in the model (LSEF) spatial correlations was about *half* of the error in the localized sample correlations. So, both LSEF’s variances and correlations were substantially more accurate than their localized sample counterparts. This explains the significant advantage of the LSEF analysis over the EnKF analysis in the experiments presented above in this section.

6.7. Conclusions on the experiments with static analyses

When the truth was generated by a (parametric) Locally Stationary Convolution Model, the proposed analysis technique performed very well compared with the competing techniques. The question remains how the proposed approach performs in a more realistic situation when background errors do not obey the Locally Stationary Convolution Model and the true local spectra are not just unavailable but they do not exist. To answer this question we turn to sequential data assimilation.

7. Numerical experiments with LSEF

Here we compare LSEF with the Kalman filter and the three sub-optimal filters whose analyses were tested in section 6. In the sequential filtering experiments reported in this section, the spatial domain was the circle.

7.1. Model of truth and forecast model

To test LSEF we took the Doubly Stochastic Advection-Diffusion-Decay Model (Tsyrlunikov and Rakitko, 2019):

$$\mathbf{x}_t = \mathbf{F}(\boldsymbol{\theta}_t)\mathbf{x}_{t-1} + \boldsymbol{\Sigma}_t\boldsymbol{\eta}_t, \quad (35)$$

where \mathbf{x}_t is the spatial field defined on a regular grid on the unit circle (a vector), t is the time index, $\mathbf{F}(\boldsymbol{\theta}_t)$ is the forecast model operator (a matrix), $\boldsymbol{\theta}_t$ and $\boldsymbol{\Sigma}_t$ are the spatial ‘parameter’ fields (vectors) defined below in Eqs. (36) and (37), $\boldsymbol{\eta}_t$ is the standard spatio-temporal white noise, and multiplication in the second term on the right-hand side is component-wise. Double stochasticity means that both the forcing $\boldsymbol{\eta}_t$ and the parameters $\boldsymbol{\theta}_t, \boldsymbol{\Sigma}_t$ are random. $\boldsymbol{\theta}_t, \boldsymbol{\Sigma}_t$ are postulated to be stationary (in space and time) random fields satisfying the same model Eq. (35) but with constant non-random coefficients $\theta_\theta, \theta_\Sigma, \Sigma_\theta, \Sigma_\Sigma$:

$$\boldsymbol{\theta}_t = \mathbf{F}(\theta_\theta)\boldsymbol{\theta}_{t-1} + \Sigma_\theta\boldsymbol{\eta}'_t, \quad (36)$$

$$\boldsymbol{\Sigma}_t = \mathbf{F}(\theta_\Sigma)\boldsymbol{\Sigma}_{t-1} + \Sigma_\Sigma\boldsymbol{\eta}'_t, \quad (37)$$

where $\boldsymbol{\eta}'_t, \boldsymbol{\eta}'_t$ are independent standard spatio-temporal white noises. Variability in $\boldsymbol{\theta}_t, \boldsymbol{\Sigma}_t$ leads to non-stationarity (in space and time) of the conditional distribution $\mathbf{x} | (\boldsymbol{\theta}, \boldsymbol{\Sigma})$, i.e., when $\boldsymbol{\theta}_t, \boldsymbol{\Sigma}_t$ are simulated and then fixed in the simulation of \mathbf{x}_t .

The reason for using this model was its ability to generate non-stationary random fields \mathbf{x}_t without sacrificing linearity. Linearity of the forecast model allowed us to compare the filters with the unbeatable benchmark, the exact Kalman filter. This is normally not possible with nonlinear forecast models.

In the filters, the forecast model operator was $\mathbf{F}(\boldsymbol{\theta}_t)$ and the forcing was zero. First, we generated the parameter fields $\boldsymbol{\theta}_t, \boldsymbol{\Sigma}_t$ and computed the model error covariance matrices $\mathbf{Q}_t = \boldsymbol{\Sigma}_t\boldsymbol{\Sigma}_t^\top$ (made available to the Kalman filter only). Then, $\boldsymbol{\theta}_t, \boldsymbol{\Sigma}_t$ were fixed and we computed the truth \mathbf{x}_t (as a solution to Eq. (35)) and observations $\mathbf{x}_t^{\text{obs}}$ (as the noise-contaminated truth at some grid points and every analysis time instant t). After that, having fixed the truth and the observations, we ran the filters.

As a result, we were able to set up and examine the Kalman Filter (KF), LSEF, and the three competing ensemble filters: EnKF, the Mean-B filter, and the Hybrid-B filter.

7.2. Filtering setup

The Doubly Stochastic Advection-Diffusion-Decay Model was defined on a 120-point regular grid on the circle (so that $\ell_{\text{max}} = 60$). Otherwise, its setup coincided with that reported in Tsyrlunikov and Rakitko (2019).

The ensemble size was between 5 and 160 with the default 10. The number of bandpass filters and the setup of the neural network were the same as in the above experiments on the sphere. 5000 consecutive assimilation cycles were used to validate the filters. \mathbf{B}_{mean} was computed by averaging Kalman filter’s \mathbf{B} over 100,000 cycles and, additionally, by averaging in space. The latter was done by averaging entries of \mathbf{B} over its main diagonal

(getting the mean variance) and each cyclic super-diagonal (getting the whole set of mean covariances). ‘Synthetic’ observations were located at every 10th model grid point. The observation error variance was specified to roughly match the mean Kalman-filter forecast error variance (this resembles the relationship between observation and forecast errors in meteorological data assimilation, our application area).

In the EnKF and Hybrid-B filters, multiplicative covariance inflation (e.g. Houtekamer and Zhang, 2016) was used and tuned to optimize their performance. Ensemble perturbations were multiplied by a factor of 1.02. There was no covariance inflation in LSEF (as we found that it deteriorated its performance, not shown).

The performance of the LSEF analysis reported in section 6 was so impressive that we did no tuning there (other than choosing the neural network’s hyperparameters). We recall that the reason for that was the setting in which (i) the prior uncertainty was, by construction, governed by the Locally Stationary Convolution Model and (2) the hyperprior of the model parameters was also available to train the neural network. Here, as neither of these two conditions was met, some tuning was done. First, we found it beneficial to increase the number of bandpass filters from 6 to 8. Second, we specified $q_{\text{shape}} = 2$ and the same width $\Delta_{\mathcal{H}_j} = \Delta_{\mathcal{H}}$ of the spectral transfer function (see Eq. (33)) for all bandpass filters. Third, we retrained the neural network for each ensemble size. Fourth, we increased $\Delta_{\mathcal{H}}$ for larger ensemble sizes: $\Delta_{\mathcal{H}} = 5$ for $n_e \leq 40$ and $\Delta_{\mathcal{H}} = 10$ for $n_e > 40$ (following the finding in Appendix G that it is worth broadening the spectral transfer functions of the bandpass filters with the growing ensemble size, see also section 4.4). Fifth, we specified the (hyper)prior sample of the spectral functions $\sigma_\ell(x)$ in a different, ad-hoc way compared to section 6, see the next subsection.

7.3. Training sample

In the course of sequential filtering with LSEF, the forecast error field does not have to obey the Locally Stationary Convolution Model, its imposition in the LSEF analysis is an approximation. Therefore we had no access to true local spectra here (as they do not exist in this setting) and thus could not immediately apply the technique described in section 6.4 to train the neural network.

To build the training sample, we initially tried to manually set up a circular analogue of the Parametric Locally Stationary Convolution Model (described in Appendix H) in such a way that the spatial non-stationarity of the pseudo-random fields generated by it resembled those of the forecast-error fields we obtained in the course of Kalman filtering. In that way, we were able to generate the training sample following section 6.4. However, the filtering results were mediocre (not shown). Another option we tested was to fit, having the EnKF forecast ensemble at every assimilation cycle, the Locally Stationary Convolution Model. The fitted model provided us with the estimated local spectra, which were then used in the training sample as the ‘truth’. The respective sample band variances were computed from a new ensemble generated by the fitted model. This approach did yield an improvement, but the following technique was simpler and somewhat more

successful, so it was used in the experiments presented below.

During an EnKF run, at each assimilation cycle $t = 1, 2, \dots, 10^5$, we performed *spatial averaging* of the forecast-ensemble sample covariance matrices \mathbf{S}_t (in the same way the spatial averaging was done in the computation of \mathbf{B}_{mean} , see section 7.2). The resulting spatially averaged covariance functions $b_t(x)$ (where $x = 0, 1, \dots, n_x/2$ labels the distance measured in mesh sizes) were stored. Then, offline, for each covariance function $b_t(x)$ we computed its spectrum (Fourier transform) $\mathbf{f}_t = (f_t(0), f_t(1), \dots, f_t(\ell_{\text{max}}))$. Knowing \mathbf{f}_t , we computed, for each t , an ensemble of pseudo-random realizations of the respective (stationary) random process, subjected them to the spatial bandpass filters (section 4.1), and finally computed the vectors \mathbf{d}_t of sample band variances (section 4.3). The resulting training sample consisted of 100,000 pairs $(\mathbf{d}_t, \sqrt{\mathbf{f}_t})$. The above spatial averaging of the EnKF sample covariances reduced the diversity in the training sample. To compensate for that effect, we set up the Doubly Stochastic Advection-Diffusion-Decay Model in the strongly non-stationary regime (Tsyrlunikov and Rakitko, 2019, Table 3) when running EnKF to generate the training sample.

7.4. Results

The performance of each of the five filters was measured using the RMSE (with respect to the truth) of their control forecasts \mathbf{x}_t^f . The RMSE was computed by averaging over the spatial grid, over 5000 assimilation cycles, and over 10 replicates of the parameter fields $\boldsymbol{\theta}_t, \boldsymbol{\Sigma}_t$ of the model of truth (outlined in section 7.1). As in section 6.5, we show here the analysis performance score, Eq. (34), where $\text{RMSE}_{\text{TrueB}}$ is the RMSE of Kalman filter forecasts.

We compared the five filters for different ensemble sizes and in four regimes of non-stationarity of the true spatio-temporal field. The regimes are labeled by numbers from 0 (stationarity) to 3 (strong non-stationarity). For a more detailed description of the non-stationarity regimes, see Tsyrlunikov and Rakitko (2019). The default non-stationarity regime was 2.

Figure 7 shows that LSEF was significantly better than the competing EnKF, Mean-B, and Hybrid-B filters. Comparing Fig. 7 with Fig. 3 demonstrates that the great advantage of LSEF observed when the truth was generated by the Parametric Locally Stationary Convolution Model (Fig. 3) was more modest in a more realistic situation here (Fig. 7). Nevertheless, the LSEF outperformed the three competing approximate filters even for very large ensemble sizes. Note that for a system with 120 degrees of freedom, the ensemble of size 160 is indeed very large, a situation we never have in high-dimensional systems. For $K \leq 40$, the superiority of LSEF was substantial.

Figure 8 shows how the five filters performed under different non-stationarity regimes. The superiority of LSEF was nearly uniform except for the stationary regime 0. Like in section 6.5 (see Fig. 5 at $\mu_{\text{NSL}} = 1$), the advantage of LSEF remained substantial even under strong non-stationarity (regime 3 on the x-axis in Fig. 8).

It is also worth mentioning that with each ensemble-based filter examined in our




Figure 7: Filters performance scores (the lower the better, =0 best possible), 90% bootstrap confidence bands (shaded). Dependence on ensemble size.

experiments, the spread (standard deviation) of the forecast ensemble did match the forecast RMSE pretty well. The difference was within 3% of the RMSE.

8. Discussion and conclusions

This research aims to enhance the ensemble Kalman filter by introducing a spatial model at the analysis (observation update) step. The idea is to abandon the use of sample covariances, which are notoriously noisy in high-dimensional problems. Instead, we propose to build a non-parametric, non-stationary spatial model, constrain it to reduce its number of degrees of freedom, and estimate it directly from the prior ensemble. The estimated model is then used to specify prior covariances in the analysis.

The model is formulated (on the sphere and the circle, but can be extended to other domains) for the underlying forecast-error random field rather than for its spatial covariances, which guarantees that the implied covariance matrices are positive definite. The model is non-stationary in space, which can be essential for real-world (e.g., geoscientific) applications, and non-parametric, which has the potential of producing more diverse shapes of spatial correlations compared to those produced by parametric models. The spatial model is formulated as a Gaussian process convolution model. The spatially varying convolution kernel $u(x, \rho)$ is *constrained* to reduce the effective number of degrees of freedom and regularize the estimation task.

One fundamental regularizing constraint is local stationarity, hence the name, the




Figure 8: Filters performance scores (the lower the better, =0 best possible), 90% bootstrap confidence bands (shaded). Dependence on the non-stationarity regime (0 is stationary, 1 is weakly non-stationary, 2 is default, 3 is highly non-stationary).

Locally Stationary Convolution Model. Local stationarity means that the random field (process) in question can be approximated by a stationary process locally in space. Local stationarity takes place whenever the typical length scale of the process is significantly less than a *non-stationarity length scale*, the scale on which the local structure of the process varies in space. In other words, whenever the convolution kernel $u(x, \rho)$ changes slowly with location (x) compared with its change as a function of distance (ρ). The stationary process-convolution model is a special case of our non-stationary model.

Another fundamental constraint imposed on the spatial model is formulated in spectral space for *spectral functions* $\sigma_\ell(x)$ defined to be the spectral transform of $u(x, \rho)$ with respect to ρ . We require $\sigma_\ell(x)$ (or the *local spectra* $f_\ell(x) = \sigma_\ell^2(x)$) to be *smooth* functions of the wavenumber ℓ . This constraint means that there is no need for high spectral resolution in estimating the spectral functions $\sigma_\ell(x)$. The relatively low spectral resolution allows for a higher resolution in physical space, thus permitting to capture spatial non-stationarity in the model estimator. Besides, smooth local spectra imply rapidly decaying spatial correlations. As a result, spurious long distance correlations are suppressed (without additional ad-hoc devices like localization or tapering), which is important in data assimilation applications.

The constraints are part of a (hyper)prior model for the model parameters. The list of constraints can be modified or extended. For example, we may know the typical degree of *roughness/smoothness* of the analyzed field (determined by the behavior of spatial

correlations at small distances). With our spatial model, this knowledge can be encoded in a constraint on the rate of decay of the spectral functions $\sigma_\ell(x)$ at large wavenumbers.

In general, using a spatial model to regularize the analysis step in a high-dimensional ensemble filter is a way to incorporate problem-specific spatial information into the analysis. This information may be available or can be acquired by a dedicated effort.

Aiming at high-dimensional applications, we built a computationally efficient model estimator. Relying on the local stationarity and smooth-local-spectra constraints, the estimator employs spatial bandpass filters to extract aggregated local spectra from the input ensemble of random field realizations. The aggregated local spectra are then disaggregated (deconvolved) from sample variances of the filtered fields. With a linear disaggregation technique, we prove that the bandpass filters can be specified in such a way that the resulting model estimator is consistent. A more accurate nonlinear disaggregation method based on the use of a neural network is developed.

Due to the regularizing constraints, the estimator produces much less noise in the resulting spatial covariances compared with the traditional localized (tapered) sample covariances. However, it does so at the expense of introducing a systematic distortion in the covariances/spectra—because the imposed constraints need not be respected in real-world spatial processes. The main hypothesis underlying this research is that there are filtering problems in which the reduction of noise in spatial covariances (i.e., the reduced random uncertainty) has a greater positive impact on the filter’s accuracy than the negative effect of the induced systematic uncertainty in the filter’s prior distribution. Our numerical experiments suggest that this hypothesis is plausible.

We postulated that in the course of sequential ensemble filtering, the prior uncertainty is governed by the Locally Stationary Convolution Model estimated online from the prior ensemble and a hyperprior distribution. This allowed us to optimize the use of both the ensemble and the forecast-error ‘climatology’. As a result, we obtained a new sequential filter termed the Locally Stationary convolutional Ensemble Filter (LSEF). Testing the filter with a non-stationary spatio-temporal model of truth yielded encouraging results. It is worth noting that ad-hoc regularization devices like localization and covariance inflation, which are, normally, used in EnKF and hybrid filters were absent in LSEF. This can be important in practice since ad-hoc devices require extensive (and expensive) tuning. Another interesting result is that the advantage of LSEF was observed not only under weak non-stationarity, as intended, but in strongly non-stationary regimes as well.

Our spatial-model-based covariance regularization technique was tested with linear observation models on the sphere and linear forecast and observation models on the circle. We believe that linearity of these tests is not a significant restriction, expecting that the technique will be useful whenever an EnKF-like filter is useful, that is, with non-linear models as well.

Using a neural Bayes estimator to recover local spectra from the outputs of the spatial bandpass filters proved to be successful even in the situation when true hyperprior distribution was *not* available to train the neural net (in sequential filtering). In case

when the true hyperprior was available (in our experiments with static analyses on the sphere), the advantage of the neural-Bayes-estimator based analysis was dramatic. Our results also suggest that placing a hyperprior on the spatial model’s parameters leads to better extraction of ‘climatological’ information from the training sample than the mixing of ensemble prior covariances with the time-mean prior covariance matrix (the popular technique reviewed in section 1.1, see item 2). One reason for this is that the hyperprior distribution is more informative than just the mean covariance matrix. Another reason is that the expressivity of a neural network is much greater than that of a linear combination with scalar weights.

Applying the developed regularization technique to real-world high-dimensional problems will require a multigrid/multi-scale version of the algorithm. Work on it is underway. The key idea is to approximate a given convolution operator using sparse operators defined on a hierarchy of spatial grids. The computational cost will be drastically reduced by representing large-distance spatial interactions on coarse grids.

Adapting our model to fields with strong non-stationarities that occur at fixed known locations (like ocean coastlines) can also be subject of future research. Further extensions of the spatial model may include directional anisotropy, and, possibly, non-Gaussianity. In applications to non-global domains, spectral bandpass filtering can be replaced with spatial wavelet filtering.

Last but not least, a promising avenue for further research is to extend the approach proposed in this paper to deal with imperfect prior ensembles. This can be especially important if the ensemble is generated by a fast data-driven forecast model. Even though such an ensemble can be large, it can suffer from probability biases. The role of the spatial model will then be to correct for those biases.

Acknowledgments

We would like to thank Dmitry Gayfulin, who contributed to the development of the regularization technique at an early stage of this study. Constructive criticism of two anonymous reviewers of the first (rejected) version of the manuscript is highly appreciated. Valuable comments of two anonymous reviewers and the editor, Alfred Stein, of the second version of the manuscript allowed us to substantially improve and clarify the manuscript. The study has been conducted under Task 1.1.1 of the Scientific Research and Technology Plan of the Russian Federal Service for Hydrometeorology and Environmental Monitoring.

Appendix A. Approximate sequential filtering

We review the Kalman filter and sub-optimal non-ensemble and ensemble filters that are used in our numerical experiments. The sub-optimal (approximate) filters are reviewed in the linear setting sufficient for our purposes in this paper. For more details and a broader perspective, see, e.g., Asch et al. (2016), Houtekamer and Zhang (2016).

Appendix A.1. Linear filtering setup

Here, we add a few notational and technical details to what is presented in section 2 above.

By n_x we denote the dimensionality of the system, i.e., the length of its state vector \mathbf{x}_t . By n_t^{obs} we denote the number of observations available at the time instant t . The model-error covariance matrix is $\mathbb{E}\boldsymbol{\eta}_t\boldsymbol{\eta}_t^\top = \mathbf{Q}_t$. The observation-error covariance matrix is $\mathbb{E}\boldsymbol{\varepsilon}_t\boldsymbol{\varepsilon}_t^\top = \mathbf{R}_t$.

The time-discrete linear filtering aims at finding *mean-square optimal* estimates and their uncertainties. A mean-square optimal estimate is a linear combination of all available data (observations) such that the estimation error is orthogonal (with respect to the covariance inner product) to the available data.

Appendix A.2. Forecast step

In the optimal (Kalman) filter, assume we have the *optimal analysis* vector \mathbf{x}_{t-1}^a , that is, the point estimate of \mathbf{x}_{t-1} whose error is orthogonal to the observations $\mathbf{x}_{t-1}^{\text{obs}}, \mathbf{x}_{t-2}^{\text{obs}}, \dots$, along with the analysis error covariance matrix \mathbf{A}_{t-1} . Then, the *optimal forecast* vector (the point estimate of \mathbf{x}_t from the same observations) is simply the deterministic (control) forecast

$$\mathbf{x}_t^f = \mathbf{F}_t\mathbf{x}_{t-1}^a \tag{A.1}$$

provided that the model error $\boldsymbol{\eta}_t$ is uncorrelated with the observation error $\boldsymbol{\varepsilon}_{t'}$ for any $t' < t$. It is straightforward to write down the error covariance matrix of the deterministic forecast: $\mathbf{B}_t = \mathbf{F}_t\mathbf{A}_{t-1}\mathbf{F}_t^\top + \mathbf{Q}_t$. So, the Kalman filter propagates in time the optimal point estimate of the truth and its error covariance matrix.

In a sub-optimal (ensemble or non-ensemble) filter, the forecast Eq. (A.1) is no longer optimal but it is unbiased if so is \mathbf{x}_{t-1}^a (subtract Eq. (1) from Eq. (A.1)). For this reason, we use Eq. (A.1) to compute the control forecast in the experiments presented in this paper (see also Remark in Appendix A.4).

Instead of propagating the error covariance matrix, an ensemble filter propagates an *ensemble* that represents the uncertainty in \mathbf{x}_t^f :

$$\mathbf{x}_t^{\text{fe}(k)} = \mathbf{F}_t\mathbf{x}_{t-1}^{\text{ae}(k)} + \boldsymbol{\eta}_t^{(k)}, \tag{A.2}$$

where $k = 1, 2, \dots, K$ labels the ensemble member, K is the ensemble size, the analysis-ensemble members $\mathbf{x}_{t-1}^{\text{ae}(k)}$ are defined below in Appendix A.4, and the simulated model error $\boldsymbol{\eta}_t^{(k)}$ is drawn from a probability distribution that, ideally (and in the experiments in this study), coincides with the distribution of the true model error $\boldsymbol{\eta}_t$.

Appendix A.3. Non-ensemble analysis

In a non-ensemble analysis scheme, we are given, first, the forecast (analysis background) \mathbf{x}^f , see Eq. (A.1). (Till the end of this subsection, the time subscripts are omitted since all vectors correspond to the same time instant.) Second, we have current

observations \mathbf{x}^{obs} . Third, a static proxy $\mathbf{B}_{\text{static}}$ to the true background-error covariance matrix \mathbf{B} is available. The point estimate of the truth (often called the deterministic or control analysis) \mathbf{x}^{a} is sought as a linear combination of \mathbf{x}^{f} and \mathbf{x}^{obs} .

If (i) the true background-error covariance matrix \mathbf{B} is known, (ii) \mathbf{x}^{f} is an optimal estimate of the truth by itself (meaning that no affine transformation of the forecast is better in the mean-square sense), and (iii) $\boldsymbol{\varepsilon}$ is probabilistically independent of the forecast error $\boldsymbol{\xi} = \mathbf{x}^{\text{f}} - \mathbf{x}$ (this is the case if $\boldsymbol{\varepsilon}$ is independent of all past observation and model errors), then the *optimal* (Kalman filter) analysis scheme can be employed:

$$\mathbf{x}^{\text{a}} = \mathbf{x}^{\text{f}} + \mathbf{K}(\mathbf{x}^{\text{obs}} - \mathbf{H}\mathbf{x}^{\text{f}}), \quad (\text{A.3})$$

where

$$\mathbf{K} = \mathbf{B}\mathbf{H}^{\text{T}}(\mathbf{H}\mathbf{B}\mathbf{H}^{\text{T}} + \mathbf{R})^{-1} \equiv (\mathbf{B}^{-1} + \mathbf{H}^{\text{T}}\mathbf{R}^{-1}\mathbf{H})^{-1}\mathbf{H}^{\text{T}}\mathbf{R}^{-1} \quad (\text{A.4})$$

is the (Kalman) gain matrix. The optimal analysis error covariance matrix is $\mathbf{A} = (\mathbf{B}^{-1} + \mathbf{H}^{\text{T}}\mathbf{R}^{-1}\mathbf{H})^{-1} \equiv (\mathbf{I} - \mathbf{K}\mathbf{H})\mathbf{B}$.

In a *sub-optimal* non-ensemble analysis scheme, only a proxy $\mathbf{B}_{\text{static}}$ to the true background-error covariance matrix is available. In this situation, the common practice is to plug in $\mathbf{B}_{\text{static}}$ instead of \mathbf{B} into Eq. (A.4) and use the resulting gain matrix to compute the analysis vector using Eq. (A.3). This approach is taken in so-called variational data assimilation (e.g. Asch et al., 2016).

Appendix A.4. Ensemble analysis

The ensemble analysis scheme has on input the control forecast \mathbf{x}^{f} and the forecast ensemble $\{\mathbf{x}^{\text{fe}(k)}\}_{k=1}^K$, see Eq. (A.2). The static covariance matrix $\mathbf{B}_{\text{static}}$ can also be available. The goal is to compute the deterministic analysis vector \mathbf{x}^{a} and the *analysis ensemble* $\{\mathbf{x}^{\text{ae}(k)}\}_{k=1}^K$ aimed to characterize the error in \mathbf{x}^{a} .

Remark. With nonlinear forecast models, the deterministic forecast \mathbf{x}^{f} is normally replaced with the ensemble mean forecast $\overline{\mathbf{x}^{\text{fe}}}$. However, with the linear forecast model, linear observations, and zero-mean forecast and observation errors, the *deterministic forecast* is unbiased (see below in this subsection and Appendix A.2). Therefore it is a better proxy to the truth than the ensemble mean. Indeed, $\overline{\mathbf{x}^{\text{fe}}} - \mathbf{x} \equiv (\overline{\mathbf{x}^{\text{fe}}} - \mathbf{x}^{\text{f}}) + (\mathbf{x}^{\text{f}} - \mathbf{x}) = -\frac{1}{K} \sum \boldsymbol{\xi}^{(k)} + \boldsymbol{\xi}$, where $\boldsymbol{\xi}^{(k)} = \mathbf{x}^{\text{fe}} - \mathbf{x}^{\text{fe}(k)}$ are the pseudo-random *ensemble perturbations* and $\boldsymbol{\xi}$ is the error in the control forecast. If $\mathbb{E}\boldsymbol{\xi} = 0$ and $\boldsymbol{\xi}^{(k)}$ are independent of $\boldsymbol{\xi}$, then the mean square error of the ensemble mean is seen to be greater than that of the control forecast. For this reason, we used \mathbf{x}^{f} rather than $\overline{\mathbf{x}^{\text{fe}}}$ to compute the deterministic analysis in our experiments.

The ensemble analysis computes the gain matrix $\widehat{\mathbf{K}}$ using Eq. (A.4) in which \mathbf{B} is replaced with its estimate, $\widehat{\mathbf{B}}$, computed from the ensemble $\mathbf{x}^{\text{fe}(k)}$ and, possibly, from the static covariances $\mathbf{B}_{\text{static}}$. Recent-past estimated covariances can also be utilized, see references in section 1.1, item 4. The Ensemble Kalman Filter (EnKF) explicitly or implicitly relies on the sample covariance matrix \mathbf{S} to specify the prior covariances. In the so-called stochastic variant of EnKF (which is used in our experiments), $\widehat{\mathbf{B}}$ equals

\mathbf{B}_{EnKF} , which is \mathbf{S} element-wise multiplied by a so-called localization matrix. Hybrid ensemble-variational data assimilation schemes use a convex combination of $\mathbf{B}_{\text{static}}$ and \mathbf{B}_{EnKF} to specify $\widehat{\mathbf{B}}$.

The resulting gain matrix $\widehat{\mathbf{K}}$ is used to compute the deterministic analysis \mathbf{x}^a following Eq. (A.3). Since $\widehat{\mathbf{K}}$ is sub-optimal, \mathbf{x}^a is not optimal either, but it is unbiased if so is the control forecast. Indeed, writing $\widehat{\mathbf{K}} = \widehat{\mathbf{K}}(\boldsymbol{\xi}^{(1)}, \boldsymbol{\xi}^{(2)}, \dots, \boldsymbol{\xi}^{(K)})$ to indicate its dependence on the forecast ensemble perturbations $\boldsymbol{\xi}^{(k)}$, substituting $\widehat{\mathbf{K}}$ into Eq. (A.3), and using Eq. (2), we obtain the analysis error

$$\mathbf{x}^a - \mathbf{x} = \boldsymbol{\xi} + \widehat{\mathbf{K}}(\boldsymbol{\xi}^{(1)}, \boldsymbol{\xi}^{(2)}, \dots, \boldsymbol{\xi}^{(K)}) \cdot (\boldsymbol{\varepsilon} - \mathbf{H}\boldsymbol{\xi}), \quad (\text{A.5})$$

which has zero expectation because $\boldsymbol{\xi}$ and $\boldsymbol{\varepsilon}$ are zero-mean random vectors independent of $\boldsymbol{\xi}^{(k)}$.

As for the analysis ensemble, we confine our attention to the so-called perturbed-observations approach to generate its members:

$$\mathbf{x}^{\text{ae}(k)} = \mathbf{x}^{\text{fe}(k)} + \widehat{\mathbf{K}} \cdot (\mathbf{x}^{\text{obs}} + \boldsymbol{\varepsilon}^{(k)} - \mathbf{H}\mathbf{x}^{\text{fe}(k)}), \quad (\text{A.6})$$

where $\boldsymbol{\varepsilon}^{(k)}$ are independent pseudo-random draws from the true observation-error distribution.

Appendix B. Non-negative definiteness of the non-stationary model

By definition, a real-valued symmetric kernel, $f(x, x')$, is non-negative definite (and thus is a valid covariance function of a random field) if for any set of spatial points x_1, x_2, \dots, x_N and any real numbers c_1, c_2, \dots, c_N , it holds that

$$\sum_{p=1}^N \sum_{q=1}^N c_p c_q f(x_p, x_q) \geq 0. \quad (\text{B.1})$$

The spatial covariance function of our non-stationary spatial convolution model is given by Eq. (11). Applying the addition theorem of spherical harmonics, we rewrite Eq. (11) as

$$B(x, x') = \sum_{\ell=0}^{\ell_{\max}} \sigma_{\ell}(x) \sigma_{\ell}(x') \sum_{m=-\ell}^{\ell} Y_{\ell m}(x) Y_{\ell m}^*(x'), \quad (\text{B.2})$$

where $\sigma_{\ell}(\cdot)$ are real-valued. Non-negative definiteness of $B(x, x')$ immediately follows from

$$\sum_{p=1}^N \sum_{q=1}^N c_p c_q B(x_p, x_q) = \sum_{\ell=0}^{\ell_{\max}} \sum_{m=-\ell}^{\ell} \left| \sum_{p=1}^N c_p \sigma_{\ell}(x_p) Y_{\ell m}(x_p) \right|^2 \geq 0. \quad (\text{B.3})$$

Appendix C. The spatial model on the circle

On \mathbb{S}^1 , the model Eq. (5) specializes to

$$\xi(x) = \int_{\mathbb{S}^1} u(x, y - x) \alpha(y) dy, \quad (\text{C.1})$$

where the real-valued convolution kernel $u(x, \rho)$ is an even function of the second argument. Let us employ the truncated Fourier series expansion of the kernel in the form

$$u(x, \rho) = \frac{1}{\sqrt{2\pi}} \sum_{\ell=-\ell_{\max}}^{\ell_{\max}} \sigma_{\ell}(x) e^{i\ell\rho}. \quad (\text{C.2})$$

Here $\sigma_{\ell}(x)$ are real-valued functions such that $\sigma_{-\ell}(x) = \sigma_{\ell}(x)$. With the positive-definite convolution kernel (section 3.4), $\sigma_{\ell}(x) > 0$. The band-limited Gaussian white noise α in Eq. (C.1) has the spectral expansion

$$\alpha(y) = \frac{1}{\sqrt{2\pi}} \sum_{\ell=-\ell_{\max}}^{\ell_{\max}} \tilde{\alpha}_{\ell} e^{i\ell y}, \quad (\text{C.3})$$

$\tilde{\alpha}_{\ell}$ are mutually uncorrelated zero-mean and unit-variance random variables. $\tilde{\alpha}_0$ is real-valued and the others are complex-valued circularly symmetric random variables.

Substituting Eqs. (C.2) and (C.3) into Eq. (C.1) yields the spectral version of the spatial model on the circle:

$$\xi(x) = \sum_{\ell=-\ell_{\max}}^{\ell_{\max}} \sigma_{\ell}(x) \tilde{\alpha}_{\ell} e^{i\ell x}, \quad (\text{C.4})$$

for which the local spectrum is defined as $f_{\ell}(x) = \sigma_{\ell}^2(x)$. From Eq. (C.4) it follows that the respective covariance function on \mathbb{S}^1 is

$$B(x, x') = \mathbb{E} \xi(x) \xi(x') = \sum_{\ell=-\ell_{\max}}^{\ell_{\max}} \sigma_{\ell}(x) \sigma_{\ell}(x') e^{i\ell(x'-x)}. \quad (\text{C.5})$$

Appendix D. The most localized kernel in the stationary case

Here we consider the stationary process convolution model. We show that its convolution kernel is not unique given the model's covariance function (or, equivalently, the spectrum). Uniqueness can be achieved by requiring that the kernel be most spatially localized.

Let us define a stationary (isotropic) process on the sphere through its spectral representation:

$$\xi_{\text{statio}}(x) = \sum_{\ell=0}^{\ell_{\max}} \sum_{m=-\ell}^{\ell} \tilde{\xi}_{\ell m} Y_{\ell m}(x) \equiv \sum_{\ell=0}^{\ell_{\max}} \sum_{m=-\ell}^{\ell} \sigma_{\ell} \tilde{\alpha}_{\ell m} Y_{\ell m}(x), \quad (\text{D.1})$$

where it is straightforward to see that $\mathbb{E} \tilde{\xi}_{\ell m} \tilde{\xi}_{\ell' m'}^* = \sigma_{\ell}^2 \delta_{\ell}^{\ell'} \delta_m^{m'}$. The (real-valued) spectrum of the process is $f_{\ell} = \sigma_{\ell}^2$. The process $\xi_{\text{statio}}(x)$ can be modeled as the convolution of the kernel

$$u(\rho) = \sum_{\ell=0}^{\ell_{\max}} \frac{2\ell+1}{4\pi} \sigma_{\ell} P_{\ell}(\cos \rho) \quad (\text{D.2})$$

with the white noise α (see Eq. (9)):

$$\xi_{\text{statio}}(x) = \int_{\mathbb{S}^2} u(\rho(x, y)) \alpha(y) dy. \quad (\text{D.3})$$

The kernel $u(\rho)$ is unique given the set of σ_ℓ but it is not unique given the spectrum $f_\ell = \sigma_\ell^2$ of the process $\xi_{\text{statio}}(x)$ — just because σ_ℓ is defined by f_ℓ up to sign. To isolate a unique kernel, we require it to be *most spatially localized* in the sense that it has a minimal width. We define the width as a kind of macro-scale from

$$R_u^2 = \frac{\int_{\mathbb{S}^2} u^2(\rho(x, y)) dy}{(\max |u(\rho)|)^2}. \quad (\text{D.4})$$

Here, the numerator is equal to $\sum \frac{2\ell+1}{4\pi} f_\ell$ and thus is fixed given the spectrum f_ℓ . So, R_u is minimized when $\max |u(\rho)|$ is maximal among all kernels with the fixed $|\sigma_\ell|$. That is, we seek to maximize $|u(\rho)|$ over both ρ and the signs of σ_ℓ . From Eq. (D.2), we have

$$|u(\rho)| = \left| \sum_{\ell=0}^{\ell_{\max}} \frac{2\ell+1}{4\pi} \sigma_\ell P_\ell(\cos \rho) \right| \leq \sum_{\ell=0}^{\ell_{\max}} \frac{2\ell+1}{4\pi} |\sigma_\ell| \quad (\text{D.5})$$

because $|P_\ell(t)| \leq 1$ for $-1 \leq t \leq 1$ (Szegő, 1939, section 7.21). Since $|P_\ell(t)| = 1$ if and only if $t = \pm 1$, the upper bound (i.e., the maximum of $|u(\rho)|$) indicated in Eq. (D.5) is reached if all $|P_\ell(\cos \rho)| = 1$ (which implies that $\rho = 0$ or $\rho = \pi$) and all products $\sigma_\ell P_\ell(\cos \rho)$ have the same sign. This happens in three cases so there are three solutions to the optimization problem $R_u \rightarrow \min$.

The first solution is $\sigma_\ell > 0$ for all ℓ . The corresponding kernel $u_1(\rho)$ is positive definite and its modulus is maximized at $\rho = 0$.

The second solution is $\sigma_\ell \leq 0$ for all ℓ . The corresponding kernel $u_2(\rho)$ is non-positive definite and its modulus is maximized at $\rho = 0$ so that $u_2(\rho) = -u_1(\rho)$.

The third solution is $\sigma_\ell = (-1)^\ell |\sigma_\ell|$. The corresponding kernel $u_3(\rho)$ is non-definite, its modulus is maximized at $\rho = \pi$, and $u_3(\rho) = u_1(\pi - \rho)$ (this follows from the identity $P_\ell(-t) = (-1)^\ell P_\ell(t)$).

As the above three kernels that minimize the length scale R_u have the same shape, we select the first solution: the *positive definite kernel* $u(\rho)$.

On the circle, similar arguments lead to virtually the same conclusion: the most localized kernel is a positive definite function $u(x)$ or its negated/translated version $\pm u(x-h)$ (for any $h \in \mathbb{S}^1$, the proof is omitted).

Appendix E. Identifiability of the model

Here we examine the process convolution model with the locally isotropic real-valued convolution kernel $u(x, \rho(x, y))$. The kernel $u(x, \rho)$ is a positive definite function of ρ for any x (equivalently, the spectral functions $\sigma_\ell(x)$ are positive for any x). We prove that if the non-stationary covariance function $B(x, x')$ is produced by this model, then the model

is unique, i.e., $\sigma_\ell(x)$ are uniquely determined by $B(x, x')$. For simplicity, we consider the circular case.

Let us start with expanding $\sigma_\ell(x)$ into the truncated Fourier series

$$\sigma_\ell(x) = \sum_{q=-Q}^Q \tilde{\sigma}_{\ell q} e^{iqx}, \quad (\text{E.1})$$

where Q is the half-bandwidth of the processes $\sigma_\ell(x)$ as functions of location x . Note that since $\sigma_{-\ell}(x) = \sigma_\ell(x)$ (see Appendix C), we have $\tilde{\sigma}_{-\ell, q} = \tilde{\sigma}_{\ell q}$. Since $\sigma_\ell(x)$ is real-valued, we have $\tilde{\sigma}_{\ell, -q} = \tilde{\sigma}_{\ell q}^*$.

We assume that the bandwidth $[-Q, Q]$ is tight in the sense that $\tilde{\sigma}_{\ell q} \neq 0$ for all ℓ .

Let us substitute the expansion Eq. (E.1) into Eq. (C.4) and change the summation variable q to $n = \ell + q$ ⁶:

$$\xi(x) = \sum_{\ell=-L}^L \tilde{\alpha}_\ell e^{i\ell x} \sum_{q=-Q}^Q \tilde{\sigma}_{\ell q} e^{iqx} = \sum_{n=-N}^N \tilde{\xi}_n e^{inx}, \quad (\text{E.2})$$

where $N = L + Q$ and

$$\tilde{\xi}_n = \sum_{\substack{\ell=n-Q \\ |\ell| \leq L}}^{n+Q} \tilde{\alpha}_\ell \tilde{\sigma}_{\ell, n-\ell}. \quad (\text{E.3})$$

The summation area in Eq. (E.3) for all possible n and \bar{n} is shown in Fig. E.9 for $L = 2$ and $Q = 1$.

Since $\xi(x)$ is uniquely represented by the set of its spectral coefficients $\{\tilde{\xi}_n\}_{n=-N}^N$, the covariance function $B(x, x')$ is uniquely represented by the set of the spectral covariances $\tilde{B}_{n\bar{n}} = \mathbb{E} \tilde{\xi}_n \tilde{\xi}_{\bar{n}}^*$.

Our goal is to prove that all $\tilde{\sigma}_{\ell q}$ are uniquely determined by the set of $\tilde{B}_{n\bar{n}}$ (and thus by the covariance function $B(x, x')$).

We start with $n = N \equiv L + Q$, for which the sum in Eq. (E.3) reduces to the single term, $\tilde{\xi}_N = \tilde{\alpha}_L \tilde{\sigma}_{LQ}$ (the upper right corner of the summation area in Fig. E.9). This means that the covariance $\tilde{B}_{Nn} = \mathbb{E} \tilde{\xi}_N \tilde{\xi}_n^*$ for any $n = L + q \in [L - Q, L + Q]$ will also contain just one term,

$$\tilde{B}_{Nn} \equiv \tilde{B}_{L+Q, L+q} = \mathbb{E} \tilde{\xi}_{L+Q} \tilde{\xi}_{L+q}^* = \tilde{\sigma}_{LQ} \tilde{\sigma}_{Lq}^*, \quad (\text{E.4})$$

so that we can uniquely restore $\tilde{\sigma}_{Lq}$ for all q . To realize this, take $q = -Q$ in Eq. (E.4) (the lower right corner of the summation area in Fig. E.9). We obtain $\tilde{B}_{N, L-Q} = \tilde{\sigma}_{LQ} \tilde{\sigma}_{L, -Q}^*$. Recalling that $\tilde{\sigma}_{\ell, -q} = \tilde{\sigma}_{\ell q}^*$, we have $\tilde{B}_{N, L-Q} = \tilde{\sigma}_{LQ}^2$, which implies that $\tilde{\sigma}_{LQ}$ is uniquely determined by $\tilde{B}_{N, L-Q}$, up to sign. By the tight bandwidth assumption, $\tilde{\sigma}_{LQ} \neq 0$, therefore, we recover $\tilde{\sigma}_{Lq}$ for all $|q| < Q$ from Eq. (E.4) (the right edge of the summation parallelogram in Fig. E.9), again, up to sign. The ambiguity in the sign disappears when

⁶In this Appendix (in contrast to the rest of the paper), we denote ℓ_{\max} by L to simplify the notation.




Figure E.9: Circular case. Summation area for $L = 2$ and $Q = 1$

it comes to $\tilde{\sigma}_{L0}$, which is positive because $\sigma_\ell(x) > 0$ due to the positive-definiteness constraint (see item (2) in section 3.8 and Appendix C). Thus, $\tilde{\sigma}_{Lq}$ are uniquely recovered for all q from the model covariances.

Then, we consider $\tilde{\xi}_{N-1}$ and realize that it, again, has just one term, $\tilde{\alpha}_{L-1} \tilde{\sigma}_{L-1, L-1+Q}$, besides the term that contains $\tilde{\sigma}_{L, L+Q-1}$, which has already been recovered. This allows us to repeat the above process and recover $\tilde{\sigma}_{L-1, L-1+q}$ for all $q \in [-Q, Q]$. And so on, we uniquely recover all non-zero $\tilde{\sigma}_{\ell q}$, and therefore all spectral functions $\sigma_\ell(x)$, from the set of spectral covariances $\tilde{B}_{n\bar{n}}$. This completes the proof that the non-stationary spatial convolution model on the circle is unique provided that the convolution kernel $u(x, \rho)$ is real-valued, is an even function of ρ , and the half-bandwidth Q of $\sigma_\ell(x)$ is such that $\tilde{\sigma}_{\ell Q} \neq 0$ for any ℓ . On \mathbb{S}^2 , the same reasoning is valid (not shown).

Appendix F. Bandpass filtering of a locally stationary signal

Here we consider spatial bandpass filtering of a random process governed by the Locally Stationary Convolution Model. We consider the filtered process in the locally stationary regime with large Λ or, formally, in the limit $\Lambda \rightarrow \infty$, with the mean process spectrum $\mathbb{E} \sigma_\ell^2$ (see section 3.6) being kept constant. We examine the approximation errors involved in Eqs. (15) and Eq. (17) and show that these errors are small whenever Λ is much greater than the length scale of the filter's impulse response function (formally, they vanish in the above limit). For simplicity of presentation, we consider the circular case.

Let the bandpass filter \mathcal{H} have the spectral transfer function

$$H(\ell) = \kappa \left(\frac{\ell - \ell^c}{T} \right), \quad (\text{F.1})$$

where $\kappa(\cdot)$ is a smooth real-valued bell-shaped function of continuous argument such that $\kappa(0) = \max |\kappa(\cdot)| = 1$ and $\max |\kappa'(\cdot)| = 1$, ℓ^c is the band's central wavenumber, and T is the half-bandwidth (*cf.* Eq. (33)). Let $L_{\mathcal{H}} = 1/T$ denote the length scale of the filter's impulse response function. Apply the filter \mathcal{H} to the process

$$\xi(x) = \sum_{\ell=-\ell_{\max}}^{\ell_{\max}} \sigma_{\ell}(x) \tilde{\alpha}_{\ell} e^{i\ell x} \quad (\text{F.2})$$

(we reproduce here Eq. (C.4) for the reader's convenience). In spectral space, the action of the filter on the signal $\xi(x) = \sum \tilde{\xi}_n e^{inx}$ amounts to the multiplication of the spectral coefficients $\tilde{\xi}_n$ by $H(n)$. The spectral-space representation of $\xi(x)$ is given by Eq. (E.2) so the filtered process $\varphi = \mathcal{H}\xi$ reads

$$\varphi(x) = \sum_{n=-\ell_{\max}-Q}^{\ell_{\max}+Q} H(n) \tilde{\xi}_n e^{inx} = \sum_{\ell=-\ell_{\max}}^{\ell_{\max}} \tilde{\alpha}_{\ell} e^{i\ell x} \sum_{q=-Q}^Q H(\ell+q) \tilde{\sigma}_{\ell q} e^{iqx}, \quad (\text{F.3})$$

where, we recall, $\tilde{\sigma}_{\ell q}$ are the spectral coefficients of the expansion $\sigma_{\ell}(x) = \sum_q \tilde{\sigma}_{\ell q} e^{iqx}$. We suppose that the spatial filter in question is selected such that $L_{\mathcal{H}} \ll \Lambda$. Here Λ is the non-stationarity length scale defined on the circle from

$$\frac{\mathbb{E} \xi^2}{\Lambda^2} = \sum_{\ell=-\ell_{\max}}^{\ell_{\max}} \mathbb{E} \left(\frac{\partial \sigma_{\ell}}{\partial x} \right)^2 \quad (\text{F.4})$$

(*cf.* Eqs. (12) and (13)). Heuristically, $H(\ell+q)$ in Eq. (F.3) is then much flatter than $\tilde{\sigma}_{\ell q}$ (both being considered as functions of q). As a result, within the width of $\tilde{\sigma}_{\ell q}$ (regarded as a function of q with the fixed ℓ), it holds that $H(\ell+q) \approx H(\ell)$. This suggests the approximation, $\varphi(x)$, that results from replacing, in Eq. (F.3), $H(\ell+q)$ with $H(\ell)$:

$$\varphi(x) \approx \varphi(x) = \sum_{\ell=-\ell_{\max}}^{\ell_{\max}} \sigma_{\ell}(x) \tilde{\alpha}_{\ell} H(\ell) e^{i\ell x}. \quad (\text{F.5})$$

Its spherical counterpart is given in Eq. (15). Now we study the approximation error

$$\delta\varphi(x) = \varphi(x) - \varphi(x) = \sum_{\ell=-\ell_{\max}}^{\ell_{\max}} \tilde{\alpha}_{\ell} e^{i\ell x} \sum_{q=-Q}^Q [H(\ell) - H(\ell+q)] \tilde{\sigma}_{\ell q} e^{iqx}. \quad (\text{F.6})$$

It has mean zero and variance

$$\mathbb{E} [|\delta\varphi(x)|^2 | \sigma] = \sum_{\ell=-\ell_{\max}}^{\ell_{\max}} \sum_{q=-Q}^Q \sum_{q'=-Q}^Q [H(\ell) - H(\ell+q)] [H(\ell) - H(\ell+q')] \tilde{\sigma}_{\ell q} \tilde{\sigma}_{\ell q'}^* e^{i(q-q')x}. \quad (\text{F.7})$$

We take here expectation over the randomness of the spectral functions $\sigma_{\ell}(x)$. Since the vector-valued random process $\boldsymbol{\sigma}(x) = \{\sigma_{\ell}(x)\}_{\ell=-\ell_{\max}}^{\ell_{\max}}$ is postulated to be stationary, see section 3.6, it holds that $\mathbb{E} \tilde{\sigma}_{\ell q} \tilde{\sigma}_{\ell q'}^* = 0$ for $q \neq q'$. As a result, Eq. (F.7) simplifies to

$$\mathbb{E} |\delta\varphi|^2 = \mathbb{E} \mathbb{E} [|\delta\varphi(x)|^2 | \sigma] = \sum_{\ell=-\ell_{\max}}^{\ell_{\max}} \sum_{q=-Q}^Q [H(\ell) - H(\ell+q)]^2 \mathbb{E} |\tilde{\sigma}_{\ell q}|^2. \quad (\text{F.8})$$

Now, we note that for each ℓ , there exists $\bar{\ell}$ such that $H(\ell + q) - H(\ell) = H'(\bar{\ell})q = \kappa'(\cdot)q/T \equiv \kappa'(\cdot)qL_{\mathcal{H}}$. Hence

$$\mathbb{E} |\delta\varphi|^2 = L_{\mathcal{H}}^2 \sum_{\ell} [\kappa'(\cdot)]^2 \sum_q q^2 \mathbb{E} |\tilde{\sigma}_{\ell q}|^2 \leq L_{\mathcal{H}}^2 \sum_{\ell} \mathbb{E} \left(\frac{\partial \sigma_{\ell}}{\partial x} \right)^2 = \left(\frac{L_{\mathcal{H}}}{\Lambda} \right)^2 \mathbb{E} \xi^2, \quad (\text{F.9})$$

where Eq. (F.4) as well as the assumption $|\kappa'(\cdot)| \leq 1$ were used.

Note that $\mathbb{E} |\varrho(x)|^2 = \mathbb{E} \mathbb{E} [|\varrho(x)|^2 | \sigma] = \sum H^2(\ell) \mathbb{E} |\sigma_{\ell}|^2$ does not change neither with Λ nor with $L_{\mathcal{H}}$ as long as the mean process spectrum $\mathbb{E} \sigma_{\ell}^2$ is not changed, as we assume. Therefore the relative approximation error variance tends to zero as $L_{\mathcal{H}}/\Lambda \rightarrow 0$ at the same rate $(L_{\mathcal{H}}/\Lambda)^2$ as the absolute error variance.

Finally, we assess the mean absolute approximation error in Eq. (17), $\mathbb{E} |\delta^{\text{nsta}}(x)|$, where $\delta^{\text{nsta}}(x) = \underline{v}(x) - v(x)$, $v(x) = \mathbb{E}(|\varphi(x)|^2 | \sigma)$, and $\underline{v}(x) = \mathbb{E}(|\varrho(x)|^2 | \sigma)$. We have $|\delta^{\text{nsta}}| = |\mathbb{E}(|\varrho|^2 - |\varphi|^2 | \sigma)| \leq \mathbb{E}(|\varrho|^2 - |\varphi|^2 | \sigma) = \mathbb{E}(|\varrho| - |\varphi|(|\varrho| + |\varphi|) | \sigma)$. Further, the inequality $||a| - |b|| \leq |a - b|$ entails that $|\delta^{\text{nsta}}| \leq \mathbb{E}(|\delta\varphi|(|\varrho| + |\varphi|) | \sigma)$. Next, taking expectation with respect to the randomness of the spectral functions σ , we obtain $\mathbb{E} |\delta^{\text{nsta}}| \leq \mathbb{E}(|\delta\varphi|(|\varrho| + |\varphi|))$. Here, the Cauchy-Schwarz inequality yields $\mathbb{E} |\delta^{\text{nsta}}| \leq \sqrt{\mathbb{E} |\delta\varphi|^2} \sqrt{\mathbb{E} (|\varrho| + |\varphi|)^2}$. Finally, assuming that $\mathbb{E} (|\varrho| + |\varphi|)^2$ is bounded above, we conclude from Eq. (F.9) that the mean absolute non-stationarity error in $\underline{v}(x)$, that is, the error in the approximate Eq. (17), is

$$\mathbb{E} |\delta^{\text{nsta}}| = O \left(\frac{L_{\mathcal{H}}}{\Lambda} \right) \quad (\text{F.10})$$

as $L_{\mathcal{H}}/\Lambda \rightarrow 0$.

Appendix G. Proof of Proposition 4.1: Consistency of the estimator

Here, we examine the accuracy of the linear estimator described in section 4.4. We consider the large sample limit, $K \rightarrow \infty$. The degree of the spatial non-stationarity is kept fixed in this limit. We show that the bandpass filters can be selected in such a way that the estimation error vanishes, implying consistency of the estimator.

We start by recalling that according to Eqs. (17) and (18), the j th band variance can be represented as

$$v_j = \underline{v}_j + \delta_j^{\text{nsta}}, \quad (\text{G.1})$$

where \underline{v}_j is linearly related to the local spectrum and δ_j^{nsta} is the error of the linear approximation (here and in the sequel, dependencies on the spatial point x are dropped).

Then, following section 4.4, we replace the sum in Eq. (18) with the integral:

$$\underline{v}_j = \int_{\mathbb{S}^1} \omega_j(\vartheta; T) F(\vartheta) d\vartheta, \quad (\text{G.2})$$

where ϑ is the log-wavenumber, $F(\vartheta)$ represents the local spectrum in the sense that $f_{\ell} = \sigma_{\ell}^2 = F(\vartheta(\ell))$, $\omega_j(\vartheta; T)$ stands for $\frac{2\ell+1}{4\pi} H_j^2(\ell)$, and T is a real number that controls

the widths of the functions $\omega_j(\cdot; T)$ (the way it does so is specified below). The discrepancy δ_j^{nsta} is studied in Appendix F, where it is generically denoted as δ^{nsta} .

We specify the transfer functions $\omega_j(\vartheta, T)$ of the bandpass filters in Eq. (G.2) as follows. First, we let the number of the filters, J , be equal to $M + 1$, where M is the bandwidth of the local spectrum, see Definition 4.1. Second, we define the function $\Omega : \mathbb{R} \rightarrow \mathbb{R}$ as a positive even function that is monotonic on \mathbb{R}_+ , vanishes at infinity, and satisfies for some $\beta > 0$

$$\Omega(t) = 1 - |t|^\beta + o(|t|^\beta) \quad (\text{G.3})$$

as $t \rightarrow 0$. Third, we define

$$\omega(\vartheta; T) = \Omega(\vartheta/T) \quad (\text{G.4})$$

and let

$$\omega_j(\vartheta; T) = \omega(\vartheta_j - \vartheta; T), \quad (\text{G.5})$$

where $\vartheta_j = \frac{\pi}{M}(j - 1)$, $j = 1, 2, \dots, J$, determines the j th passband's central log-wavenumber. We note that by definition, see Eqs. (G.4) and (G.5), T determines the length scale of the resulting filters' impulse response functions: $L_{\mathcal{H}_j} = O(T^{-1})$ as $T \rightarrow \infty$. From Eq. (F.10), we obtain

$$\mathbb{E} |\delta_j^{\text{nsta}}| = O(T^{-1}). \quad (\text{G.6})$$

The data we have are the sample band variances d_j , which estimate the band variances v_j up to the sampling error δ_j^{samp} :

$$d_j = v_j + \delta_j^{\text{samp}}. \quad (\text{G.7})$$

The magnitude of the sampling error can be assessed as

$$\mathbb{E} |\delta_j^{\text{samp}}| \leq \sqrt{\mathbb{E}(\delta_j^{\text{samp}})^2} \sim \sqrt{\frac{2}{K}} \mathbb{E} v_j = O(K^{-1/2}) \quad (\text{G.8})$$

as $K \rightarrow \infty$. From Eqs. (G.1), (G.2), (G.5), and (G.7), we have

$$\int_{\mathbb{S}^1} \omega(\vartheta_j - \vartheta, T) F(\vartheta) d\vartheta = d_j - \delta_j^{\text{samp}} - \delta_j^{\text{nsta}}. \quad (\text{G.9})$$

This is a linear inverse problem to be solved with respect to $F(\cdot)$ from the data $\{d_j\}_{j=1}^J$. The left-hand side of Eq. (G.9) is the *convolution* $\omega(\cdot; T) * F(\cdot)$ evaluated at ϑ_j . This convolution can readily be computed in Fourier space, see Eq. (22), where Eq. (G.9) becomes

$$\widetilde{F}_m \widetilde{\omega}_m(T) \propto \widetilde{d}_m - \widetilde{\delta}_m^{\text{samp}} - \widetilde{\delta}_m^{\text{nsta}}, \quad (\text{G.10})$$

where the variables with the tilde are Fourier transforms of the respective variables without the tilde and $0 \leq m \leq M$. From Eq. (G.10), we can write down the solution in Fourier space, $\widetilde{\widetilde{F}}_m \propto \widetilde{d}_m / \widetilde{\omega}_m(T)$, and its error (up to a constant factor):

$$|\delta \widetilde{\widetilde{F}}_m| \leq \frac{|\widetilde{\delta}_m^{\text{samp}}| + |\widetilde{\delta}_m^{\text{nsta}}|}{\widetilde{\omega}_m(T)}. \quad (\text{G.11})$$

From Eqs. (G.3) and (G.4), it follows that $1/\tilde{\omega}_m(T) = O(T^\beta)$ where, we recall, $\beta > 0$. We substitute this latter asymptotic relation in Eq. (G.11) and take expectation of the resulting equation with respect to the distribution of the random spectral functions $\sigma_\ell(\cdot) = \sqrt{f_\ell(\cdot)}$ (see section 3.6). We examine the asymptotic behavior of the mean absolute error $\mathbb{E}|\delta\hat{F}(\vartheta)|$ as $K \rightarrow \infty$ and, importantly, $T = K^\gamma$ (where $\gamma > 0$), with $M = \text{const.}$ From Eq. (G.11), we obtain, again up to a constant factor,

$$\mathbb{E}|\delta\hat{F}(\vartheta)| \leq \mathbb{E} \sum_{m=-M+1}^M |\delta\tilde{F}_m| \leq O(T^\beta K^{-1/2}) + O(T^{\beta-1}). \quad (\text{G.12})$$

Here the first inequality is because $\delta\hat{F}(\vartheta) = \sum_{m=-M}^M \delta\tilde{F}_m e^{im\vartheta}$ and therefore $|\delta\hat{F}(\vartheta)| \leq \sum_{m=-M}^M |\delta\tilde{F}_m|$. The second inequality is due to Eqs. (G.11), (G.6), (G.8), and the assumption that M is finite and fixed.

From Eq. (G.12), taking $0 < \beta < 1$ and $0 < \gamma < 1/(2\beta)$, we have $\mathbb{E}|\delta\hat{F}(\vartheta)| \rightarrow 0$ as $K \rightarrow \infty$. Recalling that $f_\ell = F(\vartheta(\ell))$, we conclude that the local spectrum $f_\ell(x)$ can be estimated from the set of sample band variances $d_j(x)$, and the mean absolute estimation error vanishes as the ensemble size goes to infinity.

It is worth reiterating that this result is obtained in the limit $K \rightarrow \infty$ while the degree of non-stationarity is held fixed.

Appendix H. The Parametric Locally Stationary Convolution Model

Here we introduce a parametric spatial model used in section 6 as a hyperprior model to train the neural estimator and as a model of ‘truth’ to test the LSEF analysis on the sphere. The model generates (in a hierarchical manner) realizations of the spatially variable random spectral functions $\sigma_\ell(x)$ and realizations of the respective non-stationary random field $\xi(x) | \sigma$.

Appendix H.1. General design

At the first (i.e., the lowest) level in the hierarchy is the random field $\xi(x)$ defined conditionally on the spectral functions $\sigma_\ell(x)$ according to Eq. (10).

At the second level are the random spectral functions $\sigma_\ell(x)$, which are postulated to obey the following parametric model:

$$\sigma_\ell(x) = \sqrt{\frac{c(x)}{1 + [\lambda(x)\ell]^{\gamma(x)}}}. \quad (\text{H.1})$$

Here $c(x)$ is the normalizing variable that ensures that $\text{Var} \xi(x) = \sum_\ell \frac{2\ell+1}{4\pi} \sigma_\ell^2(x) = s^2(x)$ and $s(x)$, $\lambda(x)$, $\gamma(x)$ are random fields, which we call *parameter fields*.

Appendix H.2. Parameter fields

The random field $s(x)$ equals the spatially variable standard deviation of the random field in question. $\lambda(x)$ is the local length scale parameter. $\gamma(x)$ is the local shape

parameter. These three fields are set to be non-Gaussian *stationary* random fields computed as follows:

$$\begin{aligned} s(x) &= s_{\text{add}} + s_{\text{mult}} \cdot g(\ln \varkappa_s \cdot \chi_s(x)), \\ \lambda(x) &= \lambda_{\text{add}} + \lambda_{\text{mult}} \cdot g(\ln \varkappa_\lambda \cdot \chi_\lambda(x)), \\ \gamma(x) &= \gamma_{\text{add}} + \gamma_{\text{mult}} \cdot g(\ln \varkappa_\gamma \cdot \chi_\gamma(x)). \end{aligned} \tag{H.2}$$

Here g is the nonlinear transformation function defined in the next paragraph and $\chi_s, \chi_\lambda, \chi_\gamma$ are the three *pre-transform* fields (introduced in the next subsection). The pre-transform fields are independent stationary zero-mean Gaussian fields. The hyperparameters with the subscript _{add} prohibit unrealistically small values in the respective parameter fields and, together with the hyperparameters with the subscript _{mult}, specify the median values of the respective parameter fields. The hyperparameters $\varkappa_\bullet > 0$ affect the magnitude of deviations of the parameter fields from their median values.

The transformation function g in Eq. (H.2) is taken, following Tsyrlunikov and Rakitko (2019), to be the scaled and shifted logistic function:

$$g(z) = \frac{1 + e^b}{1 + e^{b-z}}, \tag{H.3}$$

where $b = 1$ is the constant. The function $g(z)$ behaves like the ordinary exponential function everywhere except for $z \gg b$, where the exponential growth is tempered and $g(z) \approx 1 + e^b$. The reason to replace the more common e^z with $g(z)$ is the desire to avoid too large values in the parameter fields, which can give rise to unrealistically large spikes in ξ .

Due to the nonlinearity of the transformation function g , the above parameter fields $s(x), \lambda(x), \gamma(x)$ are non-Gaussian. Their point-wise distributions are known as logit-normal or logit-Gaussian. If some $\varkappa_\bullet = 1$, the respective parameter field is constant in space. The higher \varkappa_\bullet , the more variable in space becomes the respective parameter. In the numerical experiments described below, all \varkappa_\bullet are the same, equal to \varkappa . We refer to \varkappa as the *non-stationarity strength* hyperparameter because it specifies the degree of spatial deviations of the fields $s(x), \lambda(x)$, and $\gamma(x)$ (according to Eq. (H.2)) and thus determines the magnitude of spatial variability in the spectral functions $\sigma_\ell(x)$, see Eq. (H.1), which cause non-stationarity.

Appendix H.3. Pre-transform fields

The *pre-transform* random fields $\chi_s(x), \chi_\lambda(x), \chi_\gamma(x)$ are mutually independent zero-mean and unity-variance *stationary* Gaussian processes whose common spatial spectrum is

$$f_\ell^x \propto \frac{1}{1 + (\Lambda_{\text{NSL}} \cdot \ell)^\Gamma}, \tag{H.4}$$

where $\Gamma = \gamma_{\text{add}} + \gamma_{\text{mult}}$,

$$\Lambda_{\text{NSL}} = L_{\text{med}} \cdot \mu_{\text{NSL}} \tag{H.5}$$

is the non-stationarity length scale,

$$L_{\text{med}} = \lambda_{\text{add}} + \lambda_{\text{mult}} \tag{H.6}$$

is the median process length scale, and μ_{NSL} is the *non-stationarity length* scale hyperparameter.

Back substitution of the pre-transform random fields $\chi_S, \chi_\lambda, \chi_\gamma$ into Eq. (H.2) followed by plugging in the resulting parameter fields $s(x), \lambda(x),$ and $\gamma(x)$ into Eq. (H.1) (after $c(x)$ is adjusted point-wise as explained in Appendix H.1) yields the non-stationary random spectral functions $\sigma_\ell(x)$. Given $\sigma_\ell(\cdot)$, the random field $\xi(\cdot)$ is computed using Eq. (10).

Appendix H.4. Hyperparameters

The default hyperparameters of the model were the following: $\varkappa = 2, \mu_{\text{NSL}} = 3, s_{\text{add}} + s_{\text{mult}} = 1, s_{\text{add}} = (s_{\text{add}} + s_{\text{mult}})/10, \lambda_{\text{add}} + \lambda_{\text{mult}} = \Delta x \cdot 3, \lambda_{\text{add}} = \Delta x/3, \gamma_{\text{add}} + \gamma_{\text{mult}} = 4,$ and $\gamma_{\text{add}} = 1$.

Appendix H.5. Local stationarity

The (random) local spectra f_ℓ defined above in this Appendix give rise to a locally stationary random field $\xi(x)$ if $\varkappa \approx 1$ (low non-stationarity strength) or $\mu_{\text{NSL}} \gg 1$ (large non-stationarity length), *cf.* section 3.6.

Indeed, if $\varkappa \approx 1$, the parameter fields are nearly constants in space, see Eq. (H.2). Therefore, the spatial variability in the local spectrum $f_\ell(x)$ is small, implying local stationarity given that $\mu_{\text{NSL}} \geq 1$.

If $\mu_{\text{NSL}} \gg 1$, the non-stationarity length Λ_{NSL} (Eq. (H.5)) is much larger than the median local length scale L_{med} , implying local stationarity according to section 3.6.

Appendix I. Details on the neural network

We used PyTorch (Paszke et al., 2019) in our experiments on the sphere, section 6, and R Torch (Keydana, 2023) on the circle, section 7. In the standard feed-forward network, the size of the input layer was J (the number of bandpass filters). The output layer contained $\ell_{\text{max}} + 1$ neurons (the size of the local spectrum). With $J = 5\text{--}7$ and $\ell_{\text{max}} = 50\text{--}60$, we found that the optimal number of hidden fully connected layers was two and the optimal number of neurons in each of the hidden layers was 120.

In the hidden layers, the activation function was ReLU (Goodfellow et al., 2016). We tried a few other activation functions (leaky ReLU, tanh, and sigmoid) and found little difference in the performance of the estimator (not shown). To make the magnitudes of the inputs less different, we performed the square-root transformation of band variances before feeding them to the network (this appeared to be slightly beneficial).

In the training process, we used the Adam optimizer (a stochastic gradient-based numerical optimization algorithm (e.g. Kingma and Ba, 2014, Goodfellow et al., 2016)). The network’s hyperparameters were taken according to recommendations in Kingma and

Ba (2014). The size of the minibatch (a random subsample without replacement of the training sample, used to compute an estimate of the current gradient of the loss function with respect to the network weights) was tuned to be 2500. The number of iterations over the whole training sample (epochs) was 30 on the sphere and 200 on the circle.

References

- M. Asch, M. Bocquet, and M. Nodet. *Data assimilation: methods, algorithms, and applications*. SIAM, 2016.
- L. Berre and G. Desroziers. Filtering of background error variances and correlations by local spatial averaging: A review. *Mon. Weather Rev.*, 138(10):3693–3720, 2010.
- L. Berre, H. Varella, and G. Desroziers. Modelling of flow-dependent ensemble-based background-error correlations using a wavelet formulation in 4D-Var at Météo-France. *Quart. J. Roy. Meteorol. Soc.*, 141(692):2803–2812, 2015.
- M. Bonavita, E. Hólm, L. Isaksen, and M. Fisher. The evolution of the ECMWF hybrid data assimilation system. *Quart. J. Roy. Meteorol. Soc.*, 142(694):287–303, 2016.
- W. Boyles and M. Katzfuss. Ensemble Kalman filter updates based on regularized sparse inverse Cholesky factors. *Mon. Weather Rev.*, 149(7):2231–2238, 2021.
- W. P. Bruinsma, M. Tegnér, and R. E. Turner. Modelling non-smooth signals with complex spectral structure. In *International Conference on Artificial Intelligence and Statistics*, pages 5166–5195. PMLR, 2022.
- M. Buehner, J. Morneau, and C. Charette. Four-dimensional ensemble-variational data assimilation for global deterministic weather prediction. *Nonlin. Process. Geophys.*, 20(5):669–682, 2013.
- L. Cavalier. Inverse problems in statistics. In *Inverse Problems and High-Dimensional Estimation: Stats in the Château Summer School, August 31-September 4, 2009*, pages 3–96. Springer, 2011.
- R. Dahlhaus. Fitting time series models to nonstationary processes. *Ann. Stat.*, 25(1):1–37, 1997.
- J. R. Driscoll and D. M. Healy. Computing Fourier transforms and convolutions on the 2-sphere. *Advances in Applied Mathematics*, 15(2):202–250, 1994.
- M. Fisher. Background error covariance modelling. *Proc. ECMWF Semin. on recent developments in data assimilation for atmosphere and ocean, 8-12 September 2003*, pages 45–64, 2003.
- R. Furrer and T. Bengtsson. Estimation of high-dimensional prior and posterior covariance matrices in Kalman filter variants. *J. Multivar. Anal.*, 98(2):227–255, 2007.

- K. S. Gage and G. D. Nastrom. Theoretical interpretation of atmospheric wavenumber spectra of wind and temperature observed by commercial aircraft during GASP. *Journal of Atmospheric Sciences*, 43(7):729–740, 1986.
- R. G. Gallager. *Stochastic processes*. Cambridge University Press, 2013.
- G. Gaspari and S. E. Cohn. Construction of correlation functions in two and three dimensions. *Quart. J. Roy. Meteorol. Soc.*, 125(554):723–757, 1999.
- I. Goodfellow, Y. Bengio, and A. Courville. *Deep learning*. MIT press, 2016.
- H. Gryvill and H. Tjelmeland. A sparse matrix formulation of model-based ensemble Kalman filter. *Statistics and Computing*, 33(3):63, 2023.
- M. Heaton, M. Katzfuss, C. Berrett, and D. Nychka. Constructing valid spatial processes on the sphere using kernel convolutions. *Environmetrics*, 25(1):2–15, 2014.
- D. Higdon, J. Swall, and J. Kern. Non-stationary spatial modeling. *Bayes. Statist.*, 6(1):761–768, 1999.
- E. Hou, E. Lawrence, and A. O. Hero. Penalized ensemble Kalman filters for high dimensional non-linear systems. *PloS one*, 16(3):e0248046, 2021.
- P. L. Houtekamer and F. Zhang. Review of the ensemble Kalman filter for atmospheric data assimilation. *Mon. Weather Rev.*, 144(12):4489–4532, 2016.
- A. Jazwinski. *Stochastic processes and filtering theory*. Academic Press, 1970.
- I. Kasanický, J. Mandel, and M. Vejmelka. Spectral diagonal ensemble Kalman filters. *Nonlin. Proc. Geophys.*, 22(4):485–497, 2015.
- S. Keydana. *Deep Learning and Scientific Computing with R Torch*. CRC Press, 2023.
- D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- O. Ledoit and M. Wolf. A well-conditioned estimator for large-dimensional covariance matrices. *J. Multivar. Anal.*, 88(2):365–411, 2004.
- E. L. Lehmann and G. Casella. *Theory of point estimation*. Springer Science & Business Media, 2006.
- R. T. Lemos and B. Sansó. A spatio-temporal model for mean, anomaly, and trend fields of North Atlantic sea surface temperature. *Journal of the American Statistical Association*, 104(485):5–18, 2009.
- Y. Li and Z. Zhu. Spatio-temporal modeling of global ozone data using convolution. *Japanese Journal of Statistics and Data Science*, 3(1):153–166, 2020.

- B. Lindsay. Composite likelihood methods. *Contemporary Mathematics*, 80:221–239, 1988.
- A. C. Lorenc. The potential of the ensemble Kalman filter for NWP — A comparison with 4D-Var. *Quart. J. Roy. Meteorol. Soc.*, 129(595):3183–3203, 2003.
- A. C. Lorenc. Improving ensemble covariances in hybrid variational data assimilation without increasing ensemble size. *Quart. J. Roy. Meteorol. Soc.*, 143(703):1062–1072, 2017.
- A. C. Lorenc, N. E. Bowler, A. M. Clayton, S. R. Pring, and D. Fairbairn. Comparison of hybrid-4DVar and hybrid-4DVar data assimilation methods for global NWP. *Mon. Weather Rev.*, 143(2015):212–229, 2014.
- S. Mallat, G. Papanicolaou, and Z. Zhang. Adaptive covariance estimation of locally stationary processes. *Ann. Stat.*, 26(1):1–47, 1998.
- D. Marinucci and D. Peccati. *Random Fields on the Sphere*. Cambridge University Press, 2011.
- B. Ménétrier, T. Montmerle, Y. Michel, and L. Berre. Linear filtering of sample covariances for ensemble-based data assimilation. Part I: Optimality criteria and application to variance filtering and covariance localization. *Mon. Weather Rev.*, 143(5):1622–1643, 2015.
- O. Pannekoucke, L. Berre, and G. Desroziers. Filtering properties of wavelets for local background-error correlations. *Quart. J. Roy. Meteorol. Soc.*, 133(623):363–379, 2007.
- A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, pages 8024–8035. 2019.
- M. B. Priestley. Evolutionary spectra and non-stationary processes. *J. Roy. Statist. Soc. Series B (Methodological)*, 27(2):204–237, 1965.
- N. Reid. Aspects of likelihood inference. *Bernoulli*, 19(4):1404–1418, 2013.
- M. Sainsbury-Dale, A. Zammit-Mangion, and R. Huser. Likelihood-free parameter estimation with neural Bayes estimators. *The American Statistician*, pages 1–14, 2024.
- K. Sham Bhat, M. Haran, R. Olson, and K. Keller. Inferring likelihoods and climate system characteristics from climate models and multiple tracers. *Environmetrics*, 23(4):345–362, 2012.

- J. Skauvold and J. Eidsvik. Parametric spatial covariance models in the ensemble Kalman filter. *Spatial Statistics*, 29:226–242, 2019.
- G. Szegő. *Orthogonal polynomials*. American Mathematical Soc., 1939.
- F. Tobar, T. D. Bui, and R. E. Turner. Learning stationary time series using Gaussian processes with nonparametric kernels. In *Advances in neural information processing systems 28*. 2015.
- K. E. Trenberth and A. Solomon. Implications of global atmospheric spatial spectra for processing and displaying data. *J. Clim.*, 6(3):531–545, 1993.
- M. Tsyrlunikov and A. Rakitko. A hierarchical Bayes ensemble Kalman filter. *Physica D*, 338:1–16, 2017.
- M. Tsyrlunikov and A. Rakitko. Impact of non-stationarity on hybrid ensemble filters: A study with a doubly stochastic advection-diffusion-decay model. *Quart. J. Roy. Meteorol. Soc.*, 145(722):2255–2271, 2019.
- M. Tsyrlunikov, P. Svirenko, D. Gayfulin, M. Gorbunov, and A. Uspensky. Development of the operational meteorological data assimilation system in the Hydrometcenter of Russia (in Russian). *Hydrometeorological Research and Forecasting*, 4(374):112–126, 2019. URL <https://method.meteorf.ru/publ/tr/tr374/07.pdf>.
- G. Ueno and T. Tsuchiya. Covariance regularization in inverse space. *Quart. J. Roy. Meteorol. Soc.*, 135(642):1133–1156, 2009.
- M. A. Wiczorek and M. Meschede. SHTools: Tools for working with spherical harmonics. *Geochemistry, Geophysics, Geosystems*, 19(8):2574–2592, 2018.