

A FAST TUNABLE BLURRING ALGORITHM FOR SCATTERED DATA ^{*}

GREGOR ROBINSON & IAN GROOMS [†]

Abstract. A blurring algorithm with linear time complexity can reduce the small-scale content of data observed at scattered locations in a spatially extended domain of arbitrary dimension. The method works by forming a Gaussian interpolant of the input data, and then convolving the interpolant with a multiresolution Gaussian approximation of the Green's function to a differential operator whose spectrum can be tuned for problem-specific considerations. Like conventional blurring algorithms that the new algorithm generalizes to data measured at locations other than a uniform grid, applications include deblurring and separation of spatial scales. An example illustrates a possible application toward enabling importance sampling approaches to data assimilation of geophysical observations, which are often scattered over a spatial domain, since blurring observations can make particle filters more effective at state estimation of large scales. Another example, motivated by data analysis of dynamics like ocean eddies that have strong separation of spatial scales, uses the algorithm to decompose scattered oceanographic float measurements into large-scale and small-scale components.

Key words. signal processing, data assimilation, multiresolution algorithms

AMS subject classifications. 65D15, 68U10, 65D10, 62M20, 93E11, 60G35

1. Introduction. We present a linear-time blurring algorithm that works for data measured at scattered points in \mathbb{R}^d , for arbitrary dimension $d > 0$, and that permits choice in shaping its response to different length scales. This algorithm is equivalent to solving a positive definite self-adjoint elliptic partial differential equation.

By blurring, we mean the process of attenuating small-scale content of a dataset. Doing so offers a way to denoise and simplify spatial data whose large features are of primary relevance. Blurring is also a mechanism for isolating small-scale content by simply subtracting the blurred version from the original. We show an example of each of these use cases.

Applications of blurring in scientific computing generally benefit from having control over the blurring spectrum, i.e. the factor by which it attenuates inputs of different spatial scales. This concept is made more rigorous in [2](#). Regularly-spaced data on a periodic domain would enable straightforward application of Fourier methods to implement a blurring algorithm that obeys a desired spectrum. But applications in geophysics, and remote sensing in general, often involve measurements made at irregularly scattered locations in a spatially-extended domain. It is therefore desirable not to require a regular grid.

The paper is organized as follows. Our blurring method is described in [section 2](#); an illustrative synthetic example is shown in [section 3](#); an example of separating large and small scales of scattered oceanographic data is presented in [section 4](#); how this blur can be applied in its original motivating context of meteorological data assimilation is described in [section 5](#) along with a connection to generalized Gaussian random fields that led to our algorithm's discovery; another example using real meteorological data is in [section 6](#) to show the blur has the desired effect on particle filtering; algorithmic complexity and generalizations toward practical application of our method

^{*}Submitted to the editors June 16, 2019.

Funding: This work was funded by the National Science Foundation under award no. 1821074.

[†]Department of Applied Mathematics, University of Colorado Boulder, Boulder, CO (gor.robinson@colorado.edu, ian.grooms@colorado.edu).

are discussed in [section 7](#); and conclusions follow in [section 8](#).

2. Method. Let $\mathbf{z} \in \mathbb{R}^{N_z}$ be a vector of data at locations $\mathbf{q}_i \in \mathbb{R}^d$ for each $i \in (1, \dots, N_z)$, where N_z is the number of observations. The proposed blur \mathbf{S} works by solving for a discrete approximation of $\mathcal{D}^{-1}\zeta$, where \mathcal{D} is an elliptic differential operator described in the next paragraph and $\zeta : \mathbb{R}^d \rightarrow \mathbb{R}$ is a continuous-domain interpolant of the data \mathbf{z} expressed as a sum of Gaussians. This obtains by approximating the Green's function g of \mathcal{D} as a multiresolution sum of Gaussians, computing the convolution of that approximation with ζ , and evaluating the result at the locations $\{\mathbf{q}_i\}$.

Define \mathcal{D} to be the fractional bound-state Helmholtz operator:

$$(2.1) \quad \mathcal{D} = (1 - \ell^2 \Delta)^\beta,$$

where Δ is the formal Laplacian operator, $\ell > 0$ is a tuning parameter with dimensions of length, and $\beta > 0$ is a dimensionless tuning parameter that controls the rate of growth of eigenvalues. Eigenfunctions of \mathcal{D} are Fourier modes of wavenumber k and corresponding eigenvalues $(1 + \ell^2 |k|^2)^\beta$. The characteristic scale of this operator is $\ell/(2\pi\sqrt{2^{1/\beta} - 1})$, in the sense that eigenfunctions with length scales longer than this have corresponding eigenvectors close to 1.

A drawback of this approach is that convolution with g attenuates all but constant functions on \mathbb{R}^d , so even a constant data vector \mathbf{z} will be attenuated to some degree. We will discuss a way to mitigate this effect in [section 3](#).

To represent the data in a continuous form that allows convolution with g , we choose radial basis function (RBF) interpolation [\[4\]](#). RBF interpolation of the observations requires us to choose a kernel $\psi : \mathbb{R}^+ \rightarrow \mathbb{R}$ that is used as the *radial basis* in which the interpolant will represent the data. The interpolant takes the form

$$(2.2) \quad \zeta(\cdot) = \sum_{j=1}^{N_z} b_j \psi(\|\cdot - \mathbf{q}_j\|),$$

where (b_j) are interpolation weights such that

$$(2.3) \quad \zeta(\mathbf{q}_i) = \sum_{j=1}^{N_z} b_j \psi(\|\mathbf{q}_i - \mathbf{q}_j\|) = z_i.$$

In matrix form, this linear system becomes

$$(2.4) \quad \mathbf{B}\mathbf{b} = \mathbf{z}.$$

We take the RBF kernel to be $\psi(\|\cdot\|) = \phi(\cdot; 0, \xi \mathbf{I})$, where $\phi(\cdot; \mu, \Sigma)$ is the density of a d -variate Gaussian random variable with mean μ and covariance Σ

$$(2.5) \quad \phi(\cdot; \mu, \Sigma) = (2\pi \det \Sigma)^{-d/2} \exp((\cdot - \mu)^T \Sigma^{-1} (\cdot - \mu)).$$

This notation is used as a convenient description of Gaussian functions even though we will not use them to describe any random variables.

One may worry that this method is hardly fast, despite using a fast PDE solver, since a naive approach to solving for \mathbf{b} requires $\mathcal{O}(N_y^3)$ operations. Computational complexity of our algorithm, including faster alternatives to solving for \mathbf{b} , is discussed in [section 7](#).

The multiresolution Gaussian approximation of the Green's function begins by writing the Fourier transform of the inverse of (2.1) as an inverse-power function $\hat{g}(t) = t^{-\beta}$ where $t = 1 + \ell^2|k|^2$. Exponential approximations of inverse power functions like this are studied in [3, 12]. The approach therein is to write a finite trapezoid-type discretization of an integral representation of \hat{g} . With the change of variables introduced by McLean [12], the integral representation to be discretized is

$$(2.6) \quad \frac{1}{t^\beta} = \frac{1}{\Gamma(\beta)} \int_{-\infty}^{\infty} \exp(-\varphi(x, t)) (1 + e^{-x}) dx,$$

where

$$(2.7) \quad \varphi(x, t) = t \exp(x - e^{-x}) - \beta(x - e^{-x}).$$

The finite trapezoid rule discretizes this into

$$(2.8) \quad \frac{1}{t^\beta} \approx \frac{1}{\Gamma(\beta)} \sum_{n=-M_-}^{M_+} v_n e^{-a_n t},$$

where $\Gamma(\beta)$ is the Gamma function and

$$(2.9) \quad a_n = \exp(nh - e^{-nh}),$$

$$(2.10) \quad v_n = h(1 + e^{-nh}) \exp(\beta(nh - e^{-nh})).$$

Ref. [12] Lemma 4 shows that the total required number of terms $M_- + M_+ + 1$ scales as $(\ln E)^2$ to achieve uniform relative error bounded by $E > 0$ in the limit $E \downarrow 0$.

The approximation in (2.8) can now be rewritten in terms of normalized multivariate isotropic Gaussian functions of k . Given weights v_n and exponential rates a_n from the exponential approximation above, we can derive the multiplicative factors required of this equivalent formulation:

$$(2.11) \quad v_n e^{-a_n(1+\ell^2 k^2)} = v_n e^{-a_n} e^{-a_n \ell^2 k^2}$$

$$(2.12) \quad = v_n e^{-a_n} (2\pi/2\ell^2 a_n)^{d/2} \left((2\pi/2\ell^2 a_n)^{-d/2} e^{-2\ell a_n^2 k^2/2} \right)$$

$$(2.13) \quad = v_n e^{-a_n} (\pi/\ell^2 a_n)^{d/2} \phi(k; 0, 1/2\ell^2 a_n).$$

The second line obtains from simultaneously multiplying and dividing by the constant required to normalize the Gaussian term in large parentheses, which is written in that manner to ease visual comparison to the standard form of an isotropic d -variate Gaussian probability density function of mean 0 and variance $1/2\ell^2 a_n$. Combining (2.8)-(2.13) yields

$$(2.14) \quad \frac{1}{(1 + \ell^2|k|^2)^\beta} \approx \frac{1}{\Gamma(\beta)} \sum_{n=-M_-}^{M_+} v_n e^{-a_n} (\pi/\ell^2 a_n)^{d/2} \phi(k; 0, 1/2\ell^2 a_n).$$

A plot of the relative error committed by this approximation is shown in Figure 1.

Taking the inverse Fourier transform and combining terms finally yields the desired approximation of the Green's function in physical space in terms of normalized




FIG. 1. *Relative error of a Gaussian approximation in the form (2.13) of the Fourier-space Green's function $(1 - \ell^2 k^2)^{-1/\beta}$ with parameters $\ell = 1.0$, $\beta = 0.5$, $h = 0.2$, $M = 32$, and $N = 28$.*

Gaussians:

$$(2.15) \quad g(\cdot) \approx \sum_{n=-M_-}^{M_+} c_n \phi(\cdot; 0, \rho_n),$$

$$(2.16) \quad \rho_n = 2\ell^2 a_n,$$

$$(2.17) \quad c_n = \frac{v_n (\pi/\ell^2 a_n)^{d/2-1}}{\Gamma(\beta) e^{a_n}}.$$

With approximations of the data and the Green's function now constructed in terms of d -variate isotropic normalized Gaussians, convolution of which is trivial, applying a discrete version of the integral operator that inverts \mathcal{D} is just:

$$(2.18) \quad \mathcal{D}^{-1} y \approx \left(\sum_{i=-M_-}^{M_+} c_i \phi(\cdot; 0, \rho_i I) \right) * \left(\sum_{j=1}^{N_y} b_j \phi(\cdot; \mu_j, \xi I) \right)$$

$$(2.19) \quad = \sum_{i=-M_-}^{M_+} \sum_{j=1}^{N_y} c_i b_j \phi(\cdot; 0, \rho_i I) * \phi(\cdot; \mu_j, \xi I)$$

$$(2.20) \quad = \sum_{j=1}^{N_y} b_j \left(\sum_{i=-M_-}^{M_+} c_i \phi(\cdot; \mu_j, (\rho_i + \xi_j) I) \right).$$

Evaluating this quantity at the observation locations yields the output of our blur.

The last line in the manipulation above shows that the continuous function we evaluate to arrive at outputs can be interpreted as an RBF interpolant of the blurred data in terms of a new blurred basis function $\tilde{\psi}$, given by

$$(2.21) \quad \tilde{\psi}(\cdot) = \sum_{i=-M_-}^{M_+} c_i \phi(\cdot; 0, (\rho_i + \xi_j)I).$$

The weights b_j of the blurred interpolant are identical to the weights of the input interpolant, which will be valuable later in this section.

We now describe our blurring algorithm more concretely. Algorithm 2.1 ties together pieces of the Green's function approximation specified in (2.6)-(2.17). Algorithm 2.2 combines RBF interpolation of the data with the output of Algorithm 2.1 with a convolution and evaluates the result at the data locations. These algorithms, used together, are a complete description of our blur.

Algorithm 2.1 Gaussian approximation of fractional bound-state Helmholtz kernel

Input

Blurring scale parameter $\ell > 0$.
 Blurring shape parameter $\beta > 0$.
 Integration mesh size $h > 0$.
 Number of negative integration steps $M_- \in \mathbb{N}$.
 Number of positive integration steps $M_+ \in \mathbb{N}$.
 Dimension of each measurement location $d \in \mathbb{N}$.

Output

Vector of positive weights $\mathbf{c} \in \mathbb{R}^{M_-+M_++1}$.

Vector of variances $\boldsymbol{\rho} \in \mathbb{R}^{M_-+M_++1}$.

function GAUSSIANBSH($\ell, \beta, h, M_-, M_+, d$)

for $n = -M_- \rightarrow M_+$ **do**

$\hat{a} \leftarrow \exp(nh - \exp(-nh))$.

$\hat{w} \leftarrow h(1 + \exp(-nh)) \exp(\beta(nh - \exp(-nh)))$.

$\hat{w}' \leftarrow \hat{w} \exp(-\hat{a}) (\pi/\ell^2 \hat{a})^{d/2} / \Gamma(\beta)$.

$\rho_n \leftarrow 2\ell^2 \hat{a}$.

$c_n \leftarrow \hat{w}' \rho / 2\pi$.

end for

return $\mathbf{c}, \boldsymbol{\rho}$.

end function

Algorithm 2.2 defines a linear operator \mathbf{S} on \mathbb{R}^{N_z} . We will prove that \mathbf{S} is positive definite in [Theorem 2.1](#), which will be useful in [section 5](#). The theorem is more general than the specific algorithm so far presented, which will set the stage for potential variants to be described in [section 7](#). To ease into the theorem, we will summarize the preceding development of \mathbf{S} and connect it to a briefer alternative formulation that is easier to treat analytically.

Algorithm 2.2 Blur

1: **Input**

2: Vector of measured data $\mathbf{z} \in \mathbb{R}^{N_z}$.

3: Array of data location vectors $\mathbf{q}_i \in \mathbb{R}^d$, $i \in (1, \dots, N_z)$.

4: RBF scale parameter $\xi > 0$.

5: Vector of positive weights $\boldsymbol{\omega} \in \mathbb{R}^{M_- + M_+ + 1}$.

6: Vector of variances $\boldsymbol{\rho} \in \mathbb{R}^{M_- + M_+ + 1}$.

7: **Output**

8: Array of blurred data $\tilde{\mathbf{z}}_i \in \mathbb{R}$, $i \in (1, \dots, N_z)$

9: **function** BLUR(\mathbf{z} , \mathbf{q} , ξ , $\boldsymbol{\omega}$, $\boldsymbol{\rho}$)

10: Let $\phi(\|\cdot\|; 0, \xi \mathbf{I})$ be a unit-mass Gaussian to use as the RBF kernel.

11: Generate RBF weight matrix \mathbf{B} with elements:

12: **for** $i = 1 \rightarrow N_z, j = 1 \rightarrow N_z$ **do**

13: $B_{ij} \leftarrow \phi(\|\mathbf{z}_i - \mathbf{q}_j\|; 0, \xi \mathbf{I})$.

14: **end for**

15: $\mathbf{b} \leftarrow \mathbf{B}^{-1} \mathbf{z}$.

16: **for** $i = 1 \rightarrow N_y$ **do**

17: $\tilde{z}_i \leftarrow \sum_{i',n} (b_{i'} \cdot \omega_n) \phi(\|y_i - \mathbf{q}'_i\|; 0, (\xi + \rho_n) I)$.

18: **end for**

19: **return** $\tilde{\mathbf{z}}$.

20: **end function**

We described a sequence of mappings between vector spaces, with blurring taking place most explicitly in the function space $L^2(\mathbb{R}^{N_y})$ of interpolants by way of the convolution $\psi_i \mapsto \mathcal{G}\psi_i$, where \mathcal{G} denotes an operator that performs convolution with the Gaussian approximation of g , and ψ_i is defined as the interpolation basis function $\psi(\|\cdot - \mathbf{q}_i\|)$ centered at location \mathbf{q}_i . Taken literally, that conceptual development prescribes the following composition of linear operations:

$$(2.22) \quad Y \xrightarrow{\mathbf{B}^{-1}} W \xrightarrow{\mathcal{F}} X \xrightarrow{\mathcal{G}} \tilde{X} \xrightarrow{\tilde{\mathcal{F}}^{-1}} \tilde{W} \xrightarrow{\tilde{\mathbf{B}}} \tilde{Y}.$$

Nodes in this diagram represents the various vector spaces found along the way of describing our blurring algorithm:

- Y is the space of input data,
- W is the space of interpolant weights in the basis $\{\psi_i\}$,
- $X = \text{span}\{\psi_i\} \subset L^2(\mathbb{R}^n)$ is the space of interpolants,
- $\tilde{X} = \text{span}\{\mathcal{G}\psi_i\} \subset L^2(\mathbb{R}^n)$ is the space of blurred interpolants,
- \tilde{W} is the space of blurred interpolant weights in the basis $\{\mathcal{G}\psi_i\}$, and
- \tilde{Y} is the space of blurred data.

Arrows in the diagram represent the action of the operators superscribed on them:

- \mathbf{B}^{-1} maps input data to RBF weights,
- \mathcal{F} maps RBF weights to interpolated functions,
- \mathcal{G} maps interpolated functions to blurred functions,
- $\tilde{\mathcal{F}}^{-1}$ maps blurred functions to weights in a blurred RBF basis, and
- $\tilde{\mathbf{B}}$ maps blurred weights to blurred data.

The complicated sequence of steps above can simplify greatly; observe in (2.20) that the weights in the blurred basis $\{\mathcal{G}\psi_i\}$ are always identical to the weights in the unblurred basis $\{\psi_i\}$. Therefore $\tilde{\mathcal{F}}^{-1}\mathcal{G}\mathcal{F} = \mathbf{I}$, leaving just $\mathbf{S} = \tilde{\mathbf{B}}\mathbf{B}^{-1}$ where $\tilde{\mathbf{B}}$ is the

RBF matrix in the blurred RBF basis:

$$(2.23) \quad \tilde{B}_{i,j} = \tilde{\psi}(\|\mathbf{q}_i - \mathbf{q}_j\|).$$

This alternative perspective demonstrates that \mathbf{S} is equivalent to finding weights \mathbf{b} for an RBF interpolant of the unblurred data using a basis $\{\psi_i\}$, and then evaluating an RBF interpolant of the *blurred* data using the same weights \mathbf{b} that now act as coefficients on a blurred basis $\{\mathcal{G}\psi_i\}$. The following theorem is stated in terms of this simplified perspective.

THEOREM 2.1. *Let $\psi : \mathbb{R}^+ \rightarrow \mathbb{R}$ be an interpolating radial basis function and let $g : \mathbb{R}^+ \rightarrow \mathbb{R}$ be a convolution kernel. Suppose ψ and g each have positive Fourier transforms, and define $\tilde{\psi} = g * \psi$. Then the matrices \mathbf{B} with entries $B_{ij} = \psi(\|\mathbf{q}_i - \mathbf{q}_j\|)$ and $\tilde{\mathbf{B}}$ with entries $\tilde{B}_{ij} = \tilde{\psi}(\|\mathbf{q}_i - \mathbf{q}_j\|)$ are symmetric positive definite, and the product $\mathbf{S} = \tilde{\mathbf{B}}\mathbf{B}^{-1}$ is positive definite.*

Proof. A standard theorem of RBF interpolation (e.g. Section 3 of [4]) states that \mathbf{B} is positive definite under the assumption that the Fourier transform of ψ is positive.

The Convolution Theorem guarantees that $\tilde{\psi} = g * \psi$ has a positive Fourier transform if g and ψ both have positive Fourier transforms, so $\tilde{\mathbf{B}}$ is positive definite for the same reason that \mathbf{B} is positive definite.

Observe that \mathbf{B} and $\tilde{\mathbf{B}}$ are symmetric by construction, and that \mathbf{B}^{-1} is symmetric positive definite since it is the inverse of a symmetric positive definite matrix. Theorem 7.6.3 in [7] states that the product of a positive definite matrix \mathbf{P} and a Hermitian matrix \mathbf{Q} is a matrix with the same number of negative, zero, and positive eigenvalues as \mathbf{Q} . It follows that the product of two Hermitian positive definite matrices is also positive definite. Therefore, since $\tilde{\mathbf{B}}$ and \mathbf{B}^{-1} are Hermitian positive definite matrices, $\tilde{\mathbf{B}}\mathbf{B}^{-1}$ is positive definite. \square

The coefficients in the multiresolution approximation (2.15) are all positive, and the Fourier transform of a positive Gaussian is also a positive Gaussian. Therefore \mathbf{S} , as defined by Algorithm 2.2 together with Algorithm 2.1, is positive definite as a corollary of Theorem 2.1.

3. Example 1: circular measurement locations embedded in a 2-plane.

We want to verify that the blur's effect resembles what we would expect of a discrete approximation to \mathcal{D}^{-1} . To that end, we blur equally-spaced data on a circle embedded in \mathbb{R}^2 to provide insight into the spectral properties of the blur in practice. This example will also describe heuristics in choosing parameters ℓ , β , and ξ . In the course of this example we will also demonstrate an undesirable phenomenon whereby \mathbf{S} attenuates even the largest scales, and suggest a workaround.

Locations were chosen to encircle the origin in \mathbb{R}^2 with $N_y = 100$ distinct locations separated by unit distance from nearest neighbors, i.e.

$$\mathbf{q}_i = \left| e^{2in\pi/100} - 1 \right|^{-1} \begin{bmatrix} \cos(2n\pi/100) \\ \sin(2n\pi/100) \end{bmatrix}.$$

The interpolation kernel was chosen to be the isotropic Gaussian PDF with standard deviation $\xi^{1/2} = 2.5$. The convolution kernel is the multiresolution Gaussian approximation (2.15) to the fractional bound-state Helmholtz kernel with $\ell = 1$ and $\beta = 1$, using approximation parameters $h = 0.2$, $M = 32$, and $N = 28$. These parameters yield an approximation of \hat{g} with $< 0.05\%$ relative error up to $k_{max} = 49$, the Nyquist number for the one-dimensional problem that this example simulates embedded in two dimensions. Recall that Figure 1 shows this relative error as a function of k .

The blur thus constructed defines a linear operator \mathbf{S} on \mathbb{R}^{N_y} . For the purpose of inspecting the effect of blurring at different scales we numerically constructed a matrix representation of \mathbf{S} (though for practical application of the blurring algorithm it is inadvisable to actually construct \mathbf{S}). Due to the rotational symmetry of observation locations, the matrices \mathbf{B}^{-1} and \mathbf{B} are circulant. The class of circulant matrices is stable under inversion, transposition, and matrix multiplication, so $\mathbf{S} = \tilde{\mathbf{B}}\mathbf{B}^{-1}$ is also circulant. Therefore eigenvectors of \mathbf{S} are discrete Fourier vectors. This connection provides a rationale for comparing eigenvalues of \mathbf{S} to the spectrum of \mathcal{D}^{-1} , whose eigenfunctions are Fourier modes. However, it is important to recognize that the comparison is imprecise because the interpolant (2.3) represents these discrete Fourier eigenvectors as functions that differ from Fourier modes on \mathbb{R}^2 .

Eigenvalues of \mathbf{S} are plotted in Figure 2 as circles, and some examples of eigenfunctions of \mathbf{S} are visualized beneath the plot for $k \in (1, 2, 25, 49)$. *Eigenfunctions of \mathbf{S}* are defined here as continuous interpolants of the matrix's eigenvectors found by the RBF interpolation scheme utilized in the blur. This figure also shows a solid trace labelled “Fourier” that plots $(1 + \ell^2 k^2)^{-\beta}$, which is the spectrum of \mathcal{D}^{-1} that corresponds to \mathbb{R}^2 Fourier modes. Since eigenvectors of \mathbf{S} do not correspond to \mathbb{R}^2 Fourier modes, this trace of the Fourier spectrum is only a rough comparison rather than an analytical prediction that we are trying to match. The observed spectrum of \mathbf{S} behaves as expected, with gradual blurring of small scale features.

Recall from section 2 that this method has a drawback of attenuating large scales. That behavior is evident in the eigenvalues plotted in Figure 2, which are all less than 1. Eigenvalues less than 1 correspond to attenuation, so the blur attenuates even the largest scale (eigenmode index 0). This over-attenuation occurs because convolution with g attenuates every Fourier eigenmode of \mathcal{D} except constant functions in \mathbb{R}^d . Since a finite Gaussian approximation can never fully describe a nonzero spatial constant, even the largest-scale function in the space of possible RBF interpolants will be attenuated by our blur. The largest-scale eigenfunctions in this example are thin in the direction transverse to the circle, causing those modes to be blurred more than continuous Fourier modes in \mathbb{R}^2 with the same wavenumber (i.e. Fourier modes with *planar* length scale equal to the *circumferential* length scale of \mathbf{S} eigenfunctions)

For a similar reason that large scales are attenuated too much, the blur does not suppress the smallest-scale eigenmodes as much as \mathcal{G} would suppress a true \mathbb{R}^2 Fourier mode of the same wavenumber. This is because the RBF interpolants of the most highly-oscillatory eigenvectors in this example have more large-scale content than \mathbb{R}^2 Fourier modes with the same length scale.

Over-attenuation can be mitigated, so that the largest scales are closer to unity, by rescaling the operator by replacing $\mathbf{S} \mapsto \mathbf{S}/\|\mathbf{S}\mathbf{1}\|$, where $\mathbf{1}$ is a unit-norm vector with all entries identical. The eigenvectors of \mathbf{S} are usually *not* discrete Fourier vectors like they are in this symmetric example, so the largest-scale eigenvector is not necessarily $\mathbf{1}$. Therefore this mitigation technique is only a heuristic, which derives from the idea that an input with identical entries contains little small-scale information.

Choosing the RBF standard deviation parameter $\xi^{1/2}$ is not to be taken lightly. We recommend choosing it to be roughly on the order of the nearest-neighbor distance between measurements. A value too small prevents the RBF interpolation step from resolving gradual transitions from location to location, causing the interpolant to appear as a rugged set of “spikes” that are overly suppressed by the convolution step on account of their inappropriately small scale. Choosing an interpolation kernel that is too large, however, can cause numerical problems related to ill-conditioning of the linear system we must solve to arrive at RBF coefficients. Choosing $\xi^{1/2}$ to be as




FIG. 2. Top: points indicate eigenvalues of the matrix \mathbf{S} defined using interpolation by Gaussian radial basis functions of standard deviation $\xi^{1/2} = 2.5$, followed by convolution with a Gaussian approximation of the Green's function for the bound-state fractional Helmholtz kernel of $\mathcal{D} = (1 - \Delta)^\beta$ with $\ell = 1.0$ and $\beta = 1$, acting on 100 equally spaced points around the origin in \mathbb{R}^2 with unit nearest-neighbor distance. The solid trace shows the spectrum $(1 + k^2)^{-2\beta}$ of \mathcal{D}^{-1} , eigenfunctions of which are Fourier modes; this serves to highlight the similarity between the spectra of \mathbf{S} and of \mathcal{D}^{-1} , but is not an analytical solution to match since interpolating the eigenfunctions of \mathbf{S} are not Fourier modes in the plane. Bottom: some example eigenfunctions, defined as interpolants given by (2.3) of the eigenvectors of \mathbf{S} , for $k \in \{1, 2, 25, 49\}$. Duplicate eigenpairs that arise due to symmetry are suppressed in this figure.

large as possible, balanced against insurmountable instability due to ill-conditioning, is considered a best practice in RBF literature [4].

4. Example 2: scale separation of ocean temperature measurements.

The application of our algorithm is demonstrated with data from Argo, an international instrumentation project to measure the world's oceans [16]. Argo uses an array of approximately 3600 profiling floats. Every 10 days, these devices sink to a depth of 2000 meters to capture a depth-wise profile of measurements as they float back to the surface.

We apply our blurring algorithm to Argo sea surface temperature measurements to decompose the data into large-scale and small-scale components. Physical oceanographers are interested in how ocean heat content varies across spatial and temporal scales, and the Argo data set is an invaluable tool [23]. Given the irregularity of the




FIG. 3. *Sea surface temperature measurements made by floats in the Argo oceanographic instrumentation project. The horizontal and vertical coordinates describe longitudinal and latitudinal displacement, in kilometers, from a center of 50° N 35° W. Contours and colors indicate interpolated temperature differences in degrees Celsius, as estimated by the original interpolant of the data used in the blurring algorithm with $\xi^{1/2} = 175\text{km}$, $\ell = 70\text{km}$, and $\beta = 8$ (top); by the blurred interpolant representing large-scale features (center), and the small-scale interpolant defined as the difference of the blurred interpolant subtracted from the original interpolant (bottom). Although the linear fit is an important part of both the original and large scales, it is not added back for the purpose of this visualization in order to show other large-scale features that would otherwise be difficult to see.*

Argo data in space and time, it is usually interpolated to a regular grid before use [15]. Our blurring algorithm enables extraction of the large-scale (blurred) part of the data without the need to first interpolate onto a grid. The data used in this example, accessed on February 13, 2020, were taken from the region 40°N–60°N and 20°–50°W over the time period February 1, 2020 through February 10, 2020.

A small number of data points were removed to maintain a minimum nearest-neighbor separation of 50 kilometers. A least-squares linear fit is subtracted from the remaining data to arrive at deviations. These deviations are blurred with section 2 using RBF length scale $\xi^{1/2} = 175$ kilometers, $\ell = 70$ kilometers, and exponent $\beta = 8$. Subtracting the linear fit to obtain deviations avoids rapid variation of the interpolant that could otherwise happen near the boundaries due to the basis functions' rapid decay.

The original field of sea surface temperature deviations, as interpolated by RBF stage of our algorithm, is depicted in the top panel of Figure 3. The center panel shows the large-scale component, which is the result of blurring the temperature deviations. The small scale component, plotted in the bottom panel, is the difference of the unblurred and blurred temperature deviations.

5. Application to particle filtering. The blur described above was originally motivated by an effort to mitigate the dimensional curse of particle filtering, which we will refer to as sequential importance sampling with resampling (SIR) hereinafter.¹ This section will describe how the blur was derived from the motivating considerations about SIR, in terms of an observation error covariance matrix.

SIR begins by running an ensemble of forecasts; each ensemble member is a *particle*. A set of observations of the true system is then taken; the observations are corrupted by instrument errors, whose distribution defines a likelihood. This likelihood is used to compute weights for each particle so that the weighted ensemble approximates the Bayesian posterior.

SIR is susceptible to a phenomenon called *collapse*, characterized by essentially all the ensemble weight accumulating on a single ensemble member that is closest to the observations, causing the filter to catastrophically underestimate posterior dispersion. The number of ensemble members required to avoid SIR collapse depends on system covariance and observation error covariance, scaling exponentially in an effective system dimension.

Specific estimates of ensemble size required to avoid collapse, provided in [21, 22], suggest one can reduce the required ensemble size by increasing eigenvalues of the observation error covariance.² Doing so carefully can also improve uncertainty quantification for a fixed number of ensemble members. Ref. [14] suggests inflating the observation error variance at small scales, letting variance grow in wavenumber, since small scales have very limited predictability in geophysical flows [10, 17, 8].

To be more precise, consider an observing system

$$(5.1) \quad y(\mathbf{q}) = \mathcal{H}\{x\}(\mathbf{q}) + r^{1/2}(\mathbf{q}) \epsilon(\mathbf{q}),$$

where $y(\mathbf{q}) \in \mathbb{R}$ is the observation at location $\mathbf{q} \in \mathbb{R}^d$, \mathcal{H} is a function-valued observation operator acting on x , which describes the scalar system state as a function of

¹Sequential importance sampling (SIS) is also known as *particle filtering*. In this paper we specialize on SIS with resampling (SIR), the most famous variety of particle filter, though the dimensional curse and presumably our applications also extend to other particle filter varieties.

²Decreasing eigenvalues of the system covariance has the same effect, for the same reasons. However, it is harder to justify changes to a dynamical model, and to implement those changes, than to modify the observation model as we propose.

location, $r^{1/2}(\mathbf{q})$ can be imagined as the standard deviation of the observation error at \mathbf{q} , and $r^{1/2}(\mathbf{q}) \epsilon(\mathbf{q})$ is the random observation error.

It is natural to think of ϵ as a random field. But letting the spectrum grow in wavenumber precludes pointwise definition of ϵ , with probability 1, so it is not a random field in the traditional sense. The idea of imposing a correlation structure with a growing spectrum can instead be understood in the framework of generalized random fields. In this case ϵ can be treated as a random process with realizations taking the form of tempered distributions, i.e. elements of the topological dual to a Schwartz space of rapidly decaying functions on \mathbb{R}^d .

Since realizations of ϵ with a growing spectrum cannot be described pointwise, instead interpret the spatially parametrized terms in (5.1) as averages with respect to a Schwartz function ν that is closely concentrated near \mathbf{q} . For example,

$$(5.2) \quad \epsilon(\mathbf{q}) \equiv \int_{\mathbb{R}^d} \epsilon \nu \, dx / \int_{\mathbb{R}^d} \nu \, dx.$$

Narrowing our attention within the scope of generalized random fields, let ϵ be a mean-zero stationary Gaussian generalized random field (GGRF). Then the vector $(\epsilon(\mathbf{q}_1), \dots, \epsilon(\mathbf{q}_{N_y}))$ is a multivariate normal random variable with zero mean and a covariance matrix \mathbf{C} with entries C_{ij} that depend only on $\|\mathbf{q}_i - \mathbf{q}_j\|$. Hence the vector of observations $\mathbf{y} \equiv (y(\mathbf{q}_1), \dots, y(\mathbf{q}_{N_y}))$ conditioned on x is a multivariate normal random variable with mean $H(\mathbf{x})$ and covariance

$$\mathbf{R} = \mathbf{R}_0^{1/2} \mathbf{C} \mathbf{R}_0^{1/2},$$

where $\mathbf{R}_0^{1/2}$ is a diagonal matrix of the discrete observation standard deviations $r^{1/2}(\mathbf{q}_i)$ that can be treated as instrument errors and $H(\cdot) : \mathbb{R}^{N_y} \rightarrow \mathbb{R}^{N_y}$ is an observation operator acting on the discrete vector \mathbf{x} that characterizes the underlying system state. The discrete observing system can be summarized in the form

$$(5.3) \quad \mathbf{y} = \mathbf{H}\mathbf{x} + \mathbf{R}_0^{1/2} \epsilon \in \mathbb{R}^{N_y}.$$

In the spirit of [9, 18], a connection between elliptic stochastic partial differential equations and random fields enables us to make use of fast algorithms for PDEs in the context of solving for the likelihood under GGRF models, rather than naively developing a dense approximation of \mathbf{C} and then solving the associated linear system. In seeking to build a GGRF error model with a likelihood that is cheap to evaluate in high dimensions, it was this connection that led to the development of the more generally-applicable blurring algorithm presented in this paper.

We specifically treat the continuous field of observation error as $\epsilon = \mathcal{D}\mathcal{W}$, where $\mathcal{D} = (1 - \ell^2 \Delta)^\beta$ acts on a spatial white noise \mathcal{W} with mean zero and unit pointwise variance. As in (2.1), Δ is the formal Laplacian operator, $\ell > 0$ is a tuning parameter with dimensions of length, and $\beta > 0$ is a dimensionless tuning parameter that controls the rate of growth of eigenvalues. Recall that eigenfunctions of \mathcal{D} are Fourier modes of wavenumber k and corresponding eigenvalues $(1 + \ell^2 |k|^2)^\beta$. Recall also that the characteristic scale of this operator is $\ell / (2\pi\sqrt{2^{1/\beta} - 1})$, in the sense that eigenfunctions with length scales longer than this have corresponding eigenvalues close to 1. Modeling the observation error ϵ in this manner therefore ascribes a variance to large scales that is commensurate with instrument error, but it also progressively and unboundedly inflates variance for small scales at a rate controlled by β . The GGRF description of observation error is thus a kind of surrogate model for the assumption

of uncorrelated observations at large scales, but with inflated variance at small scales that are of lesser concern in geophysical forecasting.

We will use the fact that preferentially inflating observation variance at small scales is equivalent to treating blurred innovations as uncorrelated.³ To see this equivalence, observe how the correlation matrix features in the Gaussian likelihood, the logarithm of which is proportional to

$$(5.4) \quad (\mathbf{y} - \mathbf{Hx})^T \mathbf{R}_0^{-1/2} \mathbf{C}^{-1} \mathbf{R}_0^{-1/2} (\mathbf{y} - \mathbf{Hx}).$$

Then consider preprocessing the “standardized innovations” $\mathbf{R}_0^{-1/2}(\mathbf{y} - \mathbf{Hx})$ with a linear operation

$$(5.5) \quad \mathbf{R}_0^{-1/2}(\mathbf{y} - \mathbf{Hx}) \longmapsto \mathbf{S} \mathbf{R}_0^{-1/2}(\mathbf{y} - \mathbf{Hx}).$$

If these blurred observations are now assimilated under the assumption that the errors in the blurred field are standard normal, then the log-likelihood is proportional to

$$(5.6) \quad (\mathbf{y} - \mathbf{Hx})^T \mathbf{R}_0^{-1/2} \mathbf{S}^T \mathbf{S} \mathbf{R}_0^{-1/2} (\mathbf{y} - \mathbf{Hx}).$$

Any valid covariance matrix must be symmetric and positive definite. If \mathbf{S} is a positive definite blurring operator — i.e. a positive definite operator with a decaying spectrum toward small scales — then $\mathbf{C} = (\mathbf{S}^T \mathbf{S})^{-1}$ is a symmetric positive definite operator with a spectrum that grows toward small scales. In light of these considerations, take \mathbf{S} to be the algorithm presented in section 2. Although Theorem 2.1 shows that \mathbf{S} is positive definite, it is typically *not* symmetric, and the symmetric part of \mathbf{S} is not necessarily positive definite. For this reason we must treat $\mathbf{R}_0^{1/2}(\mathbf{S}^T \mathbf{S})^{-1} \mathbf{R}_0^{1/2}$ as a covariance matrix, rather than $\mathbf{R}_0^{1/2} \mathbf{S}^{-1} \mathbf{R}_0^{1/2}$.

In summary, we propose blurring standardized innovations with \mathbf{S} in such a way that $(\mathbf{S}^T \mathbf{S})^{-1}$ is a covariance for a discretized GGRF.

6. Example 3: SIR with radiosonde data. To demonstrate the behavior of our blurring algorithm on scattered data and its impact on SIR weights, we make use of data from the U.S. National Center for Atmospheric Research (NCAR) Convection Allowing Ensemble [19, 20]. The NCAR ensemble produced real-time 48 hour forecasts over the conterminous United States (CONUS) from April 7, 2015 to December 30, 2017. The ensemble forecasting system consisted of two components: an 80 member ensemble assimilation system operating at 15 km resolution and a 10 member ensemble forecast system operating at 3 km resolution. We make use of the 80 member ensemble data. The assimilation system used the Advanced Research version of the Weather Research and Forecasting (WRF) model; observations were assimilated in a 6 hour cycle via the Ensemble Adjustment Kalman Filter [2] implemented in the Data Assimilation Research Testbed software suite [1]. Every assimilation cycle processed between 66,000 and 70,000 observations from a variety of sources including radiosondes, aircraft measurements, satellite wind measurements, and Global Positioning System radio occultation data, among others. Further details are provided in [19].

To verify that our blurring algorithm performs as expected on scattered data, we apply it to radiosonde temperature measurements at a single pressure level. Every 12 hours, i.e. every other assimilation window, there are between 90 and 97 radiosonde

³The *innovation* of an ensemble member is the difference between observation and forecast.




FIG. 4. Left: Temperature fields interpolated from radiosonde data measured on May 15, 2017 at a pressure level of 70kPa. Right: the result of applying our blur with parameters $\xi^{1/2} = 5^\circ$, $\beta = 1$, and $\ell = 4^\circ$. The shape of the North America is underlaid to give a sense of scale, and small circles indicate measurement locations.

measurements scattered across North and Central America and the Caribbean available at various pressure levels. An example of the locations of these observations at a pressure level of 70 kPa on May 15, 2017 is shown in Fig. 4. The left panel shows the locations of the measurements along with an interpolated temperature field obtained using Gaussian RBFs with standard deviation $\xi^{1/2} = 5^\circ$. The right panel shows the result of applying our blurring algorithm with blurring exponent $\beta = 1/2$ and blurring length scale $\ell = 4^\circ$. Figure 4 provides visual evidence that our algorithm indeed blurs scattered data. Interpolating the raw data would leave strange regions of approximately zero Kelvins in the interpolants depicted. So for the purpose of visualization, we subtract the mean before applying the blur, and then add the mean back to the blurred data.

We next verify that the algorithm has a controllable degree of blurring with the desired effect on SIR weights, viz. that the effective sample size increases as the blurring length scale ℓ increases. To that end we use the 80 member ensemble forecast for temperature at the locations of the radiosonde temperature observations, assume that the forecast weights are all equal to 1/80, and update the weights based on mismatch to the observations using the standard SIR update formula.

To be precise, let \mathbf{y} be the vector of radiosonde temperature observations at a given time, let $\mathbf{Hx}^{(i)}$ be the vector of forecast temperatures at the same time and locations for ensemble member i , and let $\mathbf{R}_0^{1/2}$ be a diagonal matrix whose diagonal contains the standard deviations of the observation errors. The un-normalized weight for the i^{th} ensemble member is

$$(6.1) \quad \bar{w}_i = \exp \left\{ -\frac{1}{2\sigma} \left(\mathbf{y} - \mathbf{Hx}^{(i)} \right)^T \mathbf{R}_0^{-1/2} \mathbf{S}^T \mathbf{S} \mathbf{R}_0^{-1/2} \left(\mathbf{y} - \mathbf{Hx}^{(i)} \right) \right\}$$

where \mathbf{S} is the matrix corresponding to the blurring operator and $\sigma = \|\mathbf{S}\mathbf{1}\|^2$ is a rescaling factor with $\mathbf{1}$ a unit vector of identical entries. The normalized weights are

$$(6.2) \quad w_i = \frac{\tilde{w}_i}{\sum_{j=1}^N \tilde{w}_j}$$




FIG. 5. *Distributions of effective sample sizes, for different values of blurring length scale ℓ , observed for the posterior ensembles on radiosonde temperature. The posterior ensembles are obtained by performing an importance sampling update on the WRF forecast as a prior ensemble, using actual radiosonde temperature measurements at an atmospheric pressure level of 50 kPa. The distributions in this plot depict ESS values for SIR weights computed twice daily over the entire month of May 2017.*

and the effective sample size (ESS) is

$$(6.3) \quad \text{ESS} = \frac{1}{\sum_{i=1}^N w_i^2}.$$

Figure 5 shows violin plots of the ESS computed twice daily over the entire month of May 2017 using radiosonde temperature data at a pressure level of 50 kPa. The standard particle filter (i.e. without blurring, or $\ell = 0$) exhibits very poor performance with ESS only rarely rising much beyond 1.

As the blurring length scale ℓ increases from 0, the distribution of ESS also increases. With $\ell = 4^\circ$, the median ESS is approximately 3 and ESS occasionally rises beyond 5. We emphasize that these values of ESS are still quite small for an 80 member ensemble, but the goal here has not been to demonstrate the performance of the particle filter per se, but of the blurring algorithm for scattered data. That said, even a tiny increase of the ESS beyond its minimum possible value of 1 is promising because it offers hope that uncertainty quantification will improve faster upon increasing the ensemble size.

7. Computational considerations. The first step of our blur involves solving a dense linear system (2.3) for interpolation weights \mathbf{b} . A naive approach to doing so would require $\mathcal{O}(N_y^2)$ storage and $\mathcal{O}(N_y^3)$ operations. This would be highly undesirable, especially if the interpolation matrix \mathbf{B} changes between assimilation cycles due to changing observation locations. Happily, there exist algorithms to solve a Gaussian RBF problem much faster. These notably include PetRBF, which is based on GMRES iteration with a restricted additive Schwarz method preconditioner. PetRBF requires $\mathcal{O}(N_y)$ storage and $\mathcal{O}(N_y)$ operations in arbitrary dimension d , and is scalable to many cores as implemented in PETSc [25].

Another potential bottleneck is evaluating the sum of $N_y M$ Gaussians at N_y target locations, with $M = M_- + M_+ + 1$ terms in the approximation of g . A direct

approach to evaluating this sum, as is written in Algorithm 2.2, has time complexity $\mathcal{O}(N_y^2 M)$. This Gaussian sum approximation can be reduced to $\mathcal{O}(N_y + N_y M)$ with the Fast Gauss Transform (FGT) [6]. The FGT has exponential time complexity in location dimensionality d , so it often runs slower than direct evaluation when d is greater than 2. The Improved Fast Gauss Transform (IFGT) exists to eliminate that exponential scaling that hinders the original FGT for large d [24]. The IGBT can be challenging to use in practice, but there exist approaches to assist in automatic tuning such as [13]. Finally, ASKIT offers a new approach to the kernel summation problem that extends to non-Gaussian kernels and which may be less fragile in high dimensions [11].

Our algorithm uses an interpolation basis function whose width remains fixed throughout the domain. This may be problematic when the density of observation locations is highly heterogeneous, since an RBF standard deviation ξ large enough to resolve smooth features in a sparsely-sampled region may be large enough that it causes numerical problems in densely-sampled regions. Those numerical problems may arise from ill-conditioning of \mathbf{B} or from insufficient data locality expected of some divide-and-conquer solvers like PetRBF. There is some extant literature on the use of nonuniform RBF width parameters to address this situation [5], so that the size of the basis function can adapt to the density of observations. Using adaptive width involves interpolation with a basis $\{\psi_i\}$ that is allowed to vary with i . Adaptive width can be incorporated into our blur, just by modifying (2.3) and (2.18) to let ξ vary with i . Using nonuniform width parameters no longer comes with guaranteed nonsingularity, but [5] suggests that singularity is more of an exception than a rule.

Unfortunately, many fast solvers for the RBF problem are incompatible with basis functions that vary by location. One possibility to reduce the cost of solving for interpolation weights in this case is to choose compactly-supported basis functions ψ_i so that \mathbf{B} is sparse. We are unaware of any compactly-supported radial basis functions with positive Fourier transforms that are simple to convolve with a Gaussian, particularly for arbitrary d . But performing the interpolation in terms of compactly-supported bases ψ_i can be made compatible with the rest of our blurring method, simply by approximating each ψ_i with a sum of Gaussians. The resulting Gaussian approximation of the data will not be an interpolant, but careful construction can make it accurate. Therefore [Theorem 2.1](#) does not apply, but we can still expect this substitution to yield a good approximation of the convolution acting on the original interpolant.

It is similarly possible to choose a different convolution kernel g to approximate with a sum of Gaussians. This idea can be used to implement a blur of the form presented here with a wider variety of characteristics, such as a non-monotonic response in length scale. If the Gaussian approximation kernel possesses a positive Fourier transform, and the RBF interpolation employs a uniform basis function, then [Theorem 2.1](#) still applies to guarantee that $\mathbf{S}^T \mathbf{S}$ is a valid covariance matrix in its application shown in [section 5](#).

To reduce the M prefactor in the convolution step, we can apply a reduction algorithm based on Prony's method with the suboptimal approximation (2.8) as a starting point [3]. Doing so yields an optimal multiresolution approximation of the integral kernel for given uniform relative error bounds, which may require substantially fewer terms to attain the same relative accuracy. This reduction method may be particularly helpful for different forms of \mathcal{D} (ergo g) that do not yield such a rapidly-convergent approximation as (2.15).

8. Conclusions. We have described a method to blur data measured at N_y locations that are arbitrarily scattered in \mathbb{R}^d , for arbitrary d , by applying a discrete approximation to the integral equation that inverts the fractional bound state Helmholtz operator $(1 - \ell^2 \Delta)^\beta$. The degree of attenuation for different length scales can be tuned by adjusting the parameters $\ell > 0$ and $\beta > 0$; large scales are attenuated little, but length scales shorter than $\ell/(2\pi\sqrt{2^{1/\beta} - 1})$ are rapidly suppressed with a rate determined by β .

The discrete approximation results from a multiresolution Gaussian approximation to the differential operator's Green's function. This readily permits convolution with a sum of Gaussians that approximate the data; we take the sum of Gaussians approximation to be a radial basis function (RBF) interpolant with a Gaussian kernel. The blur is shown to be a positive definite linear operator on \mathbb{R}^{N_y} in a more general context where the interpolation basis and the convolution kernel have positive Fourier transforms.

Spectral properties of our blur are examined with an example shown in [section 3](#). This example provides evidence that the algorithm operates as expected: attenuation gradually increases in wavenumber, roughly approximating the differential operator's inverse spectrum, with a caveat that our blur attenuates even the largest scale. In order to preserve large scales, we propose dividing \mathbf{S} by $\|\mathbf{S}\mathbf{1}\|$, where $\mathbf{1}$ is a unit vector with all entries identical.

[Section 4](#) shows an example application of our blurring algorithm on sea surface temperature data measured by oceanographic floats in the Argo project. In contrast to the typical approach of decomposing these data into large-scale and small-scale components, our method circumvents the need for interpolating onto a regular grid.

Our blur is developed with application to Sequential Importance Sampling with Resampling (SIR) particle filters in mind. [Section 5](#) describes how blurring observations with \mathbf{S} before assimilating them as if they have uncorrelated errors is equivalent to assuming that the observation errors have covariance $(\mathbf{S}^T \mathbf{S})^{-1}$, which gives observation errors the correlation structure of a stationary generalized Gaussian random field. Relative to an uncorrelated model, an observation error model of this type decreases the number of ensemble members required to achieve good uncertainty quantification from SIR for spatially-extended dynamical systems [\[14\]](#).

[Section 6](#) demonstrates that this blur has the desired effect of helping balance SIR weights in an example with real meteorological data, which improves uncertainty quantification by reducing the tendency of SIR to produce underdispersed posterior distributions in high dimensions. This example is chosen to be provocative of potential future applications to geophysical fluid dynamics, but it is worth characterizing traits of applications that would be more appropriate. The extratropical temperature field in [section 6](#) probably features little dynamical nonlinearity at large scales, and its measurements are linear and Gaussian, so this corpus of data is an excellent candidate for assimilation with any one of the many variants of the Ensemble Kalman Filter. A more appropriate application of blurred-observation SIR would feature substantially non-Gaussian behavior at large scales. That can arise due to nonlinear dynamics of large scales or due to large dispersion of a non-negative state variable relative to its mean, or due to a nonlinear observation operator inducing a non-Gaussian posterior distribution. Moist convective systems, for example, have nonlinear dynamics and substantially skewed sign-definite variables. Examples of nonlinear observation operators that could be similar motivation for SIR include satellite radiance and precipitation measurements. Any of these features could provide motivation for accepting the computational challenge of SIR in exchange for provable convergence to the

smoothed-observation surrogate model with its controllable bias. However, it remains to be seen whether a smoothed-observation SIR filter can beat methods in the EnKF family when applied to real atmospheric problems.

A naive implementation of Algorithm 2.1 requires $\mathcal{O}(N_y^2)$ memory and $\mathcal{O}(N_y^3)$ operations to solve for interpolant weights. However section 7 describes how specialized kernel matrix solvers and fast kernel summation methods can reduce the asymptotic complexity of our algorithm to $\mathcal{O}(N_y)$.

Acknowledgments. We are grateful to Jeff Anderson and Glen Romine for their help in accessing and using the NCAR ensemble data. We are grateful to Gregory Beylkin for pointing us towards a multiresolution Gaussian approximation of the blurring kernel. The Argo data used here (<http://doi.org/10.17882/42182#70590>) were collected and made freely available by the International Argo Program and the national programs that contribute to it (<http://www.argo.ucsd.edu>, <http://argo.jcommops.org>). The Argo Program is part of the Global Ocean Observing System.

REFERENCES

- [1] J. ANDERSON, T. HOAR, K. RAEDER, H. LIU, N. COLLINS, R. TORN, AND A. AVELLANO, *The data assimilation research testbed: A community facility*, B Am Metorol Soc, 90 (2009), pp. 1283–1296.
- [2] J. L. ANDERSON, *A local least squares framework for ensemble filtering*, Mon. Wea. Rev., 131 (2003), pp. 634–642.
- [3] G. BEYLKIN AND L. MONZÓN, *Approximation by exponential sums revisited*, Appl Comput Harmon A, 28 (2010), pp. 131–149.
- [4] B. FORNBERG AND N. FLYER, *A primer on radial basis functions with applications to the geosciences*, SIAM, 2015.
- [5] B. FORNBERG AND J. ZUEV, *The Runge phenomenon and spatially variable shape parameters in RBF interpolation*, Computers & Mathematics with Applications, 54 (2007), pp. 379–398.
- [6] L. GREENGARD AND J. STRAIN, *The fast gauss transform*, SIAM Journal on Scientific and Statistical Computing, 12 (1991), pp. 79–94.
- [7] R. A. HORN AND C. R. JOHNSON, *Matrix analysis*, Cambridge Univ. Press, Cambridge, 23. print ed., 1990.
- [8] F. JUDT, *Insights into atmospheric predictability through global convection-permitting model simulations*, J. Atmos. Sci., 75 (2018), pp. 1477–1497.
- [9] F. LINDGREN, H. RUE, AND J. LINDSTRÖM, *An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach*, J Roy Stat Soc B, 73 (2011), pp. 423–498.
- [10] E. N. LORENZ, *The predictability of a flow which possesses many scales of motion*, Tellus, 21 (1969), pp. 289–307.
- [11] W. B. MARCH, B. XIAO, AND G. BIROS, *Askit: Approximate skeletonization kernel-independent treecode in high dimensions*, SIAM Journal on Scientific Computing, 37 (2015), pp. A1089–A1110.
- [12] W. MCLEAN, *Exponential sum approximations for $t - \beta$* , in *Contemporary Computational Mathematics-A Celebration of the 80th Birthday of Ian Sloan*, Springer, 2018, pp. 911–930.
- [13] V. I. MORARIU, B. V. SRINIVASAN, V. C. RAYKAR, R. DURAISWAMI, AND L. S. DAVIS, *Automatic online tuning for fast gaussian summation*, in *Advances in neural information processing systems*, 2009, pp. 1113–1120.
- [14] G. ROBINSON, I. GROOMS, AND W. KLEIBER, *Improving particle filter performance by smoothing observations*, Monthly Weather Review, 146 (2018), pp. 2433–2446.
- [15] D. ROEMMICH AND J. GILSON, *The 2004–2008 mean and annual cycle of temperature, salinity, and steric height in the global ocean from the argo program*, Progress in oceanography, 82 (2009), pp. 81–100.
- [16] D. ROEMMICH, G. C. JOHNSON, S. RISER, R. DAVIS, J. GILSON, W. B. OWENS, S. L. GARZOLI, C. SCHMID, AND M. IGNASZEWSKI, *The argo program: Observing the global ocean with profiling floats*, Oceanography, 22 (2009), pp. 34–43.
- [17] R. ROTUNNO AND C. SNYDER, *A generalization of lorenz's model for the predictability of flows with many scales of motion*, J. Atmos. Sci., 65 (2008), pp. 1063–1076.

- [18] H. RUE AND L. HELD, *Gaussian Markov random fields: theory and applications*, CRC press, 2005.
- [19] C. S. SCHWARTZ, G. S. ROMINE, R. A. SOBASH, K. R. FOSSELL, AND M. L. WEISMAN, *NCars experimental real-time convection-allowing ensemble prediction system*, Weather Forecast., 30 (2015), pp. 1645–1654.
- [20] C. S. SCHWARTZ, G. S. ROMINE, R. A. SOBASH, K. R. FOSSELL, AND M. L. WEISMAN, *NCARS real-time convection-allowing ensemble project*, B Am Metorol Soc, 100 (2019), pp. 321–343.
- [21] C. SNYDER, T. BENGTSSON, P. BICKEL, AND J. ANDERSON, *Obstacles to high-dimensional particle filtering*, Mon. Wea. Rev., 136 (2008), pp. 4629–4640.
- [22] C. SNYDER, T. BENGTSSON, AND M. MORZFIELD, *Performance bounds for particle filters using the optimal proposal*, Mon. Wea. Rev., 143 (2015), pp. 4750–4761.
- [23] C. WORTHAM AND C. WUNSCH, *A multidimensional spectral description of ocean variability*, Journal of physical oceanography, 44 (2014), pp. 944–966.
- [24] C. YANG, R. DURAISWAMI, N. A. GUMEROV, AND L. DAVIS, *Improved fast gauss transform and efficient kernel density estimation*, IEEE, 2003, p. 464.
- [25] R. YOKOTA, L. A. BARBA, AND M. G. KNEPLEY, *Petrbf4 parallel o(n) algorithm for radial basis function interpolation with gaussians*, Computer Methods in Applied Mechanics and Engineering, 199 (2010), pp. 1793–1804.