Geophysical Journal International

Geophys. J. Int. (2026) **00,** 1–38
Advance Access publication 2026 October 9
Research Paper

https://doi.org/10.1093/gji/ggag425

Irregularly and incompletely sampled random fields in the Earth sciences: analysis and synthesis of parametrized covariance models

Olivia L. Walbert [1], Frederik J. Simons [1,2], Arthur P. Guillaumin[3] and Sofia C. Olhede [4]

[1]*Department of Geosciences, Princeton University, Princeton, NJ 08544, USA. E-mail: olwalbert@princeton.edu*
[2]*Program in Applied & Computational Mathematics, Princeton University, Princeton, NJ 08544, USA*
[3]*School of Mathematical Sciences, Queen Mary University of London, E1 4NS London, UK*
[4]*Institute of Mathematics, Ecole Polytechnique Fédérale de Lausanne, 1015 Lausanne, Switzerland*

Accepted 2026 September 30. Received 2026 September 11; in original form 2026 April 1

SUMMARY

Geophysical data observed in the Earth, planetary and lunar sciences are spatially limited to sampling plans designed to span a feature of interest whose resolution and regularity in discretization are bound by site, observational and instrumental constraints. Common wisdom tells us that the more observations we make, the less variable our estimates of statistical properties of measured physical quantities will be. However, given the limited spatial extent and possibly irregular sampling of any data set, how confident can we be in our estimates of the parameters of statistical processes underlying such fields? We study how sampling geometry contributes to uncertainty in modelling spatial geophysical observations as sampled random fields characterized by stationary, isotropic, parametric covariance functions. We incorporate the signature of discrete spatial sampling patterns into an asymptotically unbiased spectral maximum-likelihood estimation method along with analytical uncertainty calculation. We illustrate the broad applicability of our modelling through synthetic and real data examples with sampling patterns that include irregularly bounded contiguous regions of interest, structured sweeps of instrumental measurements and missing observations dispersed across the domain of a field, which spur behaviours in the estimator. We find through asymptotic studies that allocating samples following a growing-domain strategy rather than a densifying, infill scheme best reduces estimator bias and (co)variance, whether the field has been sampled regularly or not. As our modelling assumptions, too, shape how (well) an observed random field can be characterized, we study the effect of covariance parameters assumed *a priori*. We demonstrate the desirable behaviour of the general Matérn class and show how to interrogate goodness-of-fit criteria to detect departures from the null hypothesis of Gaussianity, stationarity and isotropy.

Key words: Fourier analysis; Spatial analysis; Statistical methods.

1 INTRODUCTION

How does sampling (discretization over limited domains) affect what can be learned about the statistical properties of a spatial data set? Geophysical data are collected following spatial patterns designed to best survey a region of interest within practical experimental constraints, such as sampling cost, site accessibility and anticipated measurement quality. Data sets in practice are always finite and discrete, often departing from regular (rectangular or circular), evenly spaced, full grids. Properly accounting for the finiteness and discrete nature of spatial data when attempting to characterize the second-order, that is, correlation, structure, of the underlying random field, is a must.

Writing as F.J. Simons *et al.* (2026), we argued for the merits of modelling in the spectral domain, showing how to recover the parameters of the spectral density of locally stationary, isotropic, processes. Working entirely in the space domain, J.A. Goff & T.H. Jordan (1988) developed a procedure for the parametrized estimation of the spatial covariance of geophysical fields. Both the spatial covariance and the spectral density are 'covariance functions', Fourier pairs that carry the same information and are governed by the same parameters, but estimation procedures anchored in the spatial domain contend with different issues than approaches rooted in the spectral domain. For the former, based on spatial likelihood maximization or least-squares fitting of sample spatial covariances, the size of matrix inversion invariably becomes a computational bottleneck, scaling as $\mathcal{O}(K^3)$ with the size of the data K. With the latter, often based on maximizing a spectral pseudo-likelihood, unbiased parameter recovery and incorporating irregular sampling patterns were typically seen as challenging, especially in higher dimensions. The attractive $\mathcal{O}(K \log K)$ scaling of Fourier-based algorithms, however, stimulated our work on addressing both of these problems. Adopting the debiased spatial Whittle maximum-likelihood

© The Author(s) 2026. Published by Oxford University Press on behalf of The Royal Astronomical Society. All rights reserved. For commercial re-use, please contact reprints@oup.com for reprints and translation rights for reprints. All other permissions can be obtained through our RightsLink service via the Permissions link on the article page on our site-for further information please contact journals.permissions@oup.com

estimation strategy of A.P. Guillaumin *et al.* (2022), we first showed, for geospatial data with structure across multiple physical scales, that aliasing, finite-field and edge effects can be treated, resulting in the unbiased estimation of isotropic Matérn covariance function parameters. The Matérn covariance is a class of covariance functions that possess finite variance, non-negative correlation and monotonic decay.

Initially, we limited our treatment and analysis to completely and regularly sampled rectangular domains, see F.J. Simons *et al.* (2026). Our present work broadens the applicability of (univariate) spectral estimation to the more challenging settings familiar to the geophysical community, namely: Realistic data sampling patterns defined by irregular boundaries, structured tracks and general 'missingness' that may arise as voids, sparsity or random deletions. We provide geometric and asymptotic insight into the effects of sampling on the spatial statistics inferred through the spectral estimation process, and shed light on a number of other issues, such as how and when (not) to fix parameters in the inversion, how to recognize the effect of departures from isotropy and stationarity and how to design geophysical sampling experiments.

Here, in modelling geophysical data as spatial random fields, we continue to work under the hypothesis that the data observed are a realization of a locally stationary, isotropic, Gaussian process whose statistical structure is estimable through three parameters characterizing features related to the size, scale and shape of the field through its covariance structure. In the case of the Matérn density, those are variance σ^2, smoothness ν and range ρ. We emphasize the connection between the general Matérn covariance and its special cases (P. Guttorp & T. Gneiting 2006), which are frequently assumed within the geophysical community as descriptions of the process that underlies an observed field, either from rules-of-thumb or through domain-specific analysis and precedent (e.g. J. Przybilla *et al.* 2009; A.P. Valentine & D.R. Davies 2020; U. Beyaztas *et al.* 2021; G. Tunnicliffe Wilson *et al.* 2022; T. Schwaiger *et al.* 2023; Y. Zhang *et al.* 2023). We illustrate the relationship between the general and special cases to bring awareness to the implications of selecting, for example, a specific smoothness, in terms of the bias it imparts when chosen incorrectly.

We show how parameter uncertainty can be analytically predicted as a function of the sampling pattern, allowing us to study the asymptotic behaviour of the estimator and its covariance. We explore these attributes of our estimator in the context of fixed and growing domains with regular and irregular sampling, with the intent of informing the experimental design of sampling patterns. We demonstrate how structure within the spectral-domain residuals can be informative of when, how and at what wavelength data defy the assumptions of the null hypothesis, and we quantify model goodness-of-fit through a number of carefully crafted diagnostics. We demonstrate the broad applicability of our methodology through several synthetic simulation studies with realistic sampling strategies common to the geosciences, and through several actual case studies that use data sets from multiple subdisciplines. We provide all developed computational tools openly as a complete software suite for simulation, estimation and diagnostics of spatial data sets across scientific domains.

To lead with a worked example while relegating all the details of the theory, implementation and interpretational framework to the remainder of this paper, we show in Fig. 1 how modelling the statistical properties of composite data sets presents the geoscientist with the need for acknowledging how the (supposedly 'noiseless', correctly 'positioned') 'data', themselves are (geophysical) 'models'. In the case of sea level (e.g. D.C. Slobbe *et al.* 2012), or seafloor bathymetry (e.g. W.H.F. Smith & D.T. Sandwell 1997; D.T. Sandwell *et al.* 2022), for example, spatially distinct domains may incorporate observational strategies with very different bandwidths and resolutions. When these are spatially merged and presented as a *data product* (e.g. GEBCO Bathymetric Compilation Group 2024), acknowledgment must be made of their disjoint spatial footprints prior to analysis. In the case of parametrized covariance estimation, by working directly in the spatial domain, J.A. Goff & T.H. Jordan (1988) naturally achieve spatial data selectivity in the inversion, but it is not *a priori* clear how modelling carried out in the spectral domain can be made to be similarly aware of data provenance. And yet this is among the merits of the debiased spatial Whittle likelihood (A.P. Guillaumin *et al.* 2022): That it can leverage the scaling and speed of the Fast Fourier Transform in the spectral domain while running on incompletely observed spatial grids. Fig. 1 shows a portion of northeast Atlantic bathymetry, identifying the scope of shipboard (Fig. 1a), satellite-augmented (Fig. 1b) and merged values (Fig. 1c). Along the structured sampling tracks of Fig. 1(a) and within their complement (Fig. 1b), examples of the bounded contiguous and dispersed sampling patterns that we study throughout this work meld. Our modelling is able to recover the underlying Matérn parameters from those (dis)joint subdomains, and simulate synthetic realizations on the corresponding subdomain (Figs 1d–f), or anywhere else. The simulated realizations preserve the (mean and deterministic trends) first-order (obtained via least-squares first-order polynomial fitting in the space-domain) and (two-point correlation) second-order statistics (estimated parametrically from the data using the spectral Whittle likelihood in a manner that is unbiased by the sampling process). Fig. 1(i) puts forward that distinct domains can be compared, but conclusions based on their unseparated, merged mixture are meaningless. Parameter inference and uncertainty estimation lead to models that are 'right' (different, testable) when they are individually informed from the correct spatial support ('`dir`' and '`ind`'), while mixing and mingling of data sources ('`mer`') leads to models, incorrectly posited as stationary, that are (demonstrably, verifiably) 'wrong'. For comparison, Fig. 1(g) shows space-based naive estimators for the variance and range of the three data domains, with their uncertainties. The sampling variance was calculated on a $6\times$ decimated observed grid, and the range lengths estimated empirically from the drop-off to $1/3$ in the autocorrelation (Fig. 1h), ignoring the effect of smoothness in that measure. Naive estimators for smoothness are not readily available. We will return to the directly measured data set (Fig. 1a) as one of our real-world data examples later in this paper.

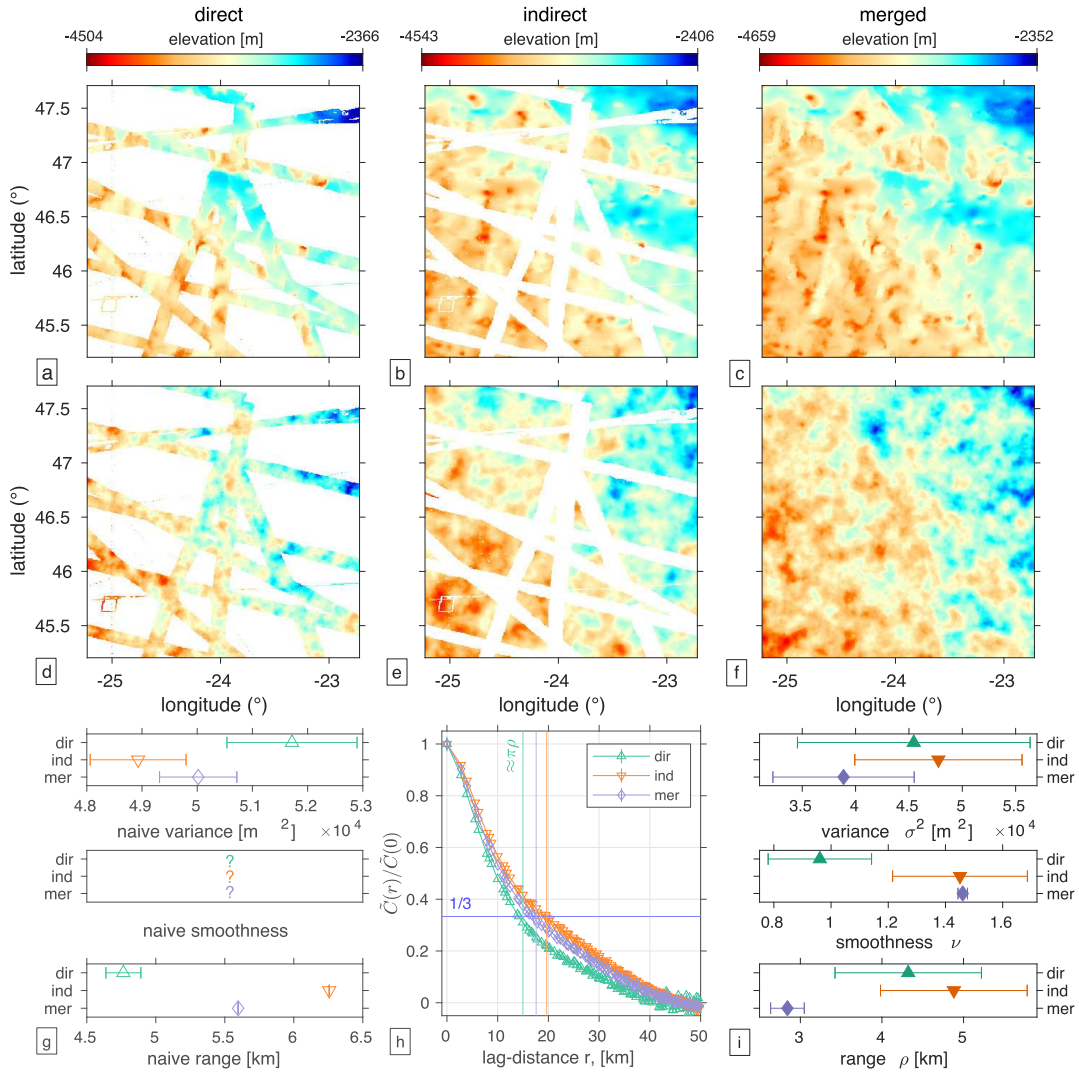


Figure 1. Analysing geophysical random fields on domains with disjoint support. The debiased Whittle maximum-likelihood analysis developed in this paper is able to estimate statistical parameters of random field covariance structure, suitable for simulation, out-of-sample extension, evaluation, interpretation and uncertainty appraisal. Bathymetric data from the northeast Atlantic (a–c). Synthetic realizations (d–f) that share the first- and second-order statistics of the superjacent data within their window of observation. In triptych (i), debiased Whittle likelihood Matérn model parameter estimates (for variance σ^2, smoothness ν and range ρ) and their one standard deviation uncertainties for the direct (`dir`, green upward triangles), indirect (`ind`, orange downward triangles) and merged (`mer`, purple diamonds) data. The parameter estimates for the merged, non-stationary, data set are not simple averages of those of the underlying component processes, and their uncertainties are irreconcilable due to model misspecification. In triptych (g), naive estimators (unfilled symbols) and their uncertainties: sample variance of the 6× decimated data set, and correlation length obtained from the sample autocovariance (h).

2 THE COMPLETE METHODOLOGY: MATÉRN MEETS WHITTLE

Our Whittle maximum-likelihood method for the spectral-domain covariance estimation of data drawn from a stationary, isotropic Matérn class (P. Whittle 1954; B. Matérn 1960) grew out of developments in a (geo)physical context made by F.J. Simons & S.C. Olhede (2013), and was put on firm statistical footing by A.M. Sykulski *et al.* (2019) for time-series and A.P. Guillaumin *et al.* (2022) for spatial data in any number of dimensions d.

F.J. Simons *et al.* (2026) fully established the debiased spatial Whittle maximum-likelihood method for univariate analysis as used here, but we did not then study its behaviour under realistic sampling scenarios, which is the subject of this paper. To make the present contribution self-contained, we review the essential elements of the methodology in this section. We identify special cases of the Matérn covariance, introduce a new measure for understanding spatial correlation, discuss the simulation of spatial realizations and the unbiased recovery of parameters via the Whittle likelihood and we review its quality assessment and uncertainty quantification.

We conclude this section with practical take-aways for data preparation and pre-processing and a reader's guide to the rest of the paper.

2.1 Stationary, isotropic parametric covariance models of spatial random fields

The Matérn class of covariance functions offers flexibility and generality for stationary processes (M.L. Stein 1999). Adopting the parametrization of M.S. Handcock & J.R. Wallis (1994), the space-domain representation of the isotropic Matérn covariance for spatial separation lag distances r is

$$C_{\boldsymbol{\theta}}(r) = \sigma^2 \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\frac{2\nu^{\frac{1}{2}}}{\pi\rho} r \right)^{\nu} K_{\nu} \left(\frac{2\nu^{\frac{1}{2}}}{\pi\rho} r \right) \qquad \text{for} \qquad r \geq 0\,,\ \sigma^2 > 0\,,\ \nu > 0\,,\ \text{and}\ \rho > 0. \tag{1}$$

The three Matérn parameters $\boldsymbol{\theta} = [\sigma^2,\ \nu,\ \rho]^{\mathrm{T}}$ flexibly characterize the covariance function, encoding the amplitude of the field in the variance, σ^2, its mean-squared differentiability (R.J. Adler 1981; G. Christakos 1992) through the smoothness, ν and its lateral correlation scale through the range, ρ. Two special functions appear in eq. (1): the Gamma function $\Gamma(z)$ defined in eq. (6.1.1) of M. Abramowitz and I.A. Stegun (1965), and the modified Bessel function of the second kind of order ν, $K_{\nu}(z)$, see eq. (9.6.1) of M. Abramowitz & I.A. Stegun (1965).

In the stationary, isotropic case, the Matérn spatial covariance possesses a Wiener–Khintchine Fourier-pair relationship with the Matérn spectral density, itself the covariance of orthogonal spectral increments (H. Cramér 1942), specified in d dimensions for wavenumbers k as

$$\mathcal{S}^d_{\boldsymbol{\theta}}(k) = \sigma^2 \frac{\Gamma(\nu + d/2)}{\Gamma(\nu)} \frac{1}{\pi^{d/2}} \left(\frac{4\nu}{\pi^2\rho^2} \right)^{\nu} \left(\frac{4\nu}{\pi^2\rho^2} + k^2 \right)^{-\nu-d/2} \qquad \text{for} \qquad k \in \mathbb{R}^+. \tag{2}$$

All that follows will be in terms of 2-D random fields with generality to higher dimensions straightforward. Evaluating eq. (2) for the 2-D case, we drop the dimensionality superscript and express the Matérn spectral density as

$$\mathcal{S}_{\boldsymbol{\theta}}(k) = \sigma^2 \frac{\pi\rho^2}{4} \left(\frac{4\nu}{\pi^2\rho^2} \right)^{\nu+1} \left(\frac{4\nu}{\pi^2\rho^2} + k^2 \right)^{-\nu-1} \qquad \text{and} \qquad k \in \mathbb{R}^+. \tag{3}$$

The three-parameter form of the Matérn class lends to its generality, with several, well-known two-parameter special cases being simplifications of eqs (1) and (2), where the value of the smoothness parameter, ν, is fixed, and the two parameters that remain free characterize the shape and scale of the covariance. Often, the special value of ν reduces the $K_{\nu}(z)$ term to a more wieldy function, such as the product of an exponential and a polynomial term when ν is half integer. While many of these two-parameter covariance functions are widely adopted throughout the geosciences, the Matérn class that they generalize from is less commonly recognized as such. We present the analytical expressions for several special models in Tables A1 and A2, as direct simplifications of the parametrization of the two- and d-D Matérn forms in eqs (1)–(3).

In Fig. 2, we display random fields generated for several special cases with increasing smoothness values ν and a common variance σ^2 and range ρ, on a fully observed, evenly spaced, rectangular sampling grid. The distinction between these covariance models is evident: Fields characterized by a larger ν display more gentle, undulous transitions in amplitude between spatially sampled values. To make this figure, as the only example in this paper where we wish to preserve the dominant spatial patterns, we scaled a single set of complex normal Gaussian variates with spectral densities of varying smoothness. The omission of wavenumber correlation that that entails causes wrap-around periodicity and loss of stationarity. Hence, for simulations everywhere else, we will use circulant embedding, see Section 2.2.

As Fig. 3 reveals, in their spectral representations, higher-smoothness processes show reduced high-frequency content, and their correlation functions take off from the origin at gentler angles. F.J. Simons *et al.* (2026) quantified the first notion by quoting the cumulative power of the isotropic Matérn form integrated from 0 to positive k. They defined λ_{α}, their eq. (18), as the wavenumber (practically, limited by the Nyquist frequency) at which the power reaches $100 \times \alpha$ per cent of the total, which we plotted in the top row of Fig. 3. Here, as to the second notion, analogously, in the space domain, we find that the cumulative spatial covariance as a function of lag distance r' (in practice, limited by the diagonal length of the domain) is given by

$$\int_0^{r'} r\, C_{\boldsymbol{\theta}}(r)\, \mathrm{d}r = \frac{1}{2}\sigma^2(\pi\rho)^2 \left[1 - \frac{2^{-\nu}}{\Gamma(\nu+1)} \left(\frac{2\nu^{\frac{1}{2}}}{\pi\rho} r' \right)^{\nu+1} K_{\nu+1} \left(\frac{2\nu^{\frac{1}{2}}}{\pi\rho} r' \right) \right]. \tag{4}$$

We obtained the above result via integration by parts (see also eq. 7.14.1.3 of A. Erdélyi 1953). The total covariance accumulated over all lag distances can be expressed (using eq. 6.561.16 of I.S. Gradshteyn & I.M. Ryzhik 2000) as

$$\int_0^{\infty} r\, C_{\boldsymbol{\theta}}(r)\, \mathrm{d}r = \frac{1}{2}\sigma^2(\pi\rho)^2 = \frac{\mathcal{S}_{\boldsymbol{\theta}}(0)}{2\pi}. \tag{5}$$

To determine the lag distance r_{α} that is associated with α per cent of the total cumulative covariance, we implement a bisection root-finding strategy to evaluate eqs (4) and (5) numerically, as we cannot isolate r_{α} from the argument of the Bessel function analytically, per

$$\int_0^{r_{\alpha}} r\, C_{\boldsymbol{\theta}}(r)\, \mathrm{d}r = \alpha \int_0^{\infty} r\, C_{\boldsymbol{\theta}}(r)\, \mathrm{d}r. \tag{6}$$

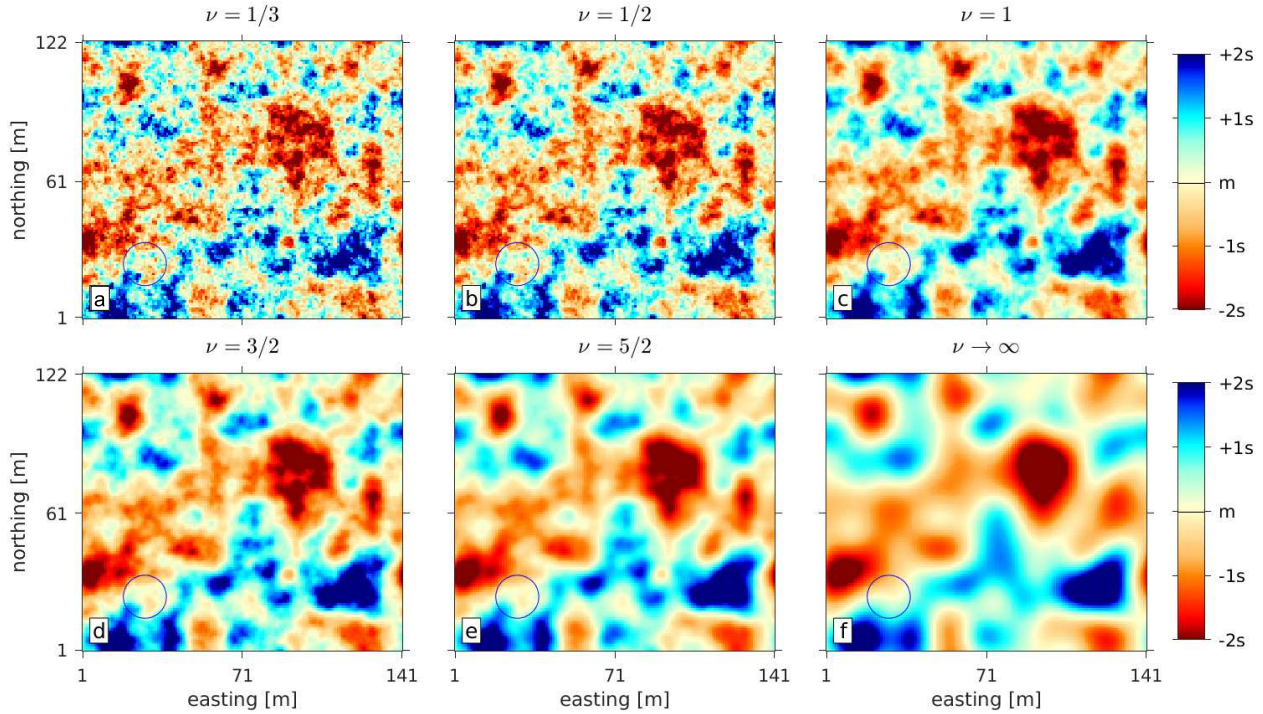


Figure 2. Spatial random fields simulated from special cases of the Matérn class, for a common variance ($\sigma^2 = 1\,\mathrm{m}^2$) and range ($\rho = 5\,\mathrm{m}$) over a rectangular, regular, grid, 122 m × 141 m, with even 1 m spacing, for different smoothness parameters ν. (a–c) Smoothness parameter $\nu = 1/3$, 1/2, and 1; (d–f) $\nu = 3/2$, 5/2, and ∞ (see Table A1 for their analytic forms). Blue circles mark the approximate 1/3 decorrelation length of the field by their radii $\pi\rho$, all of which are equal here due to the common range ρ of the six models. Colourmaps range between ±2 standard deviations s of the fields with means m.

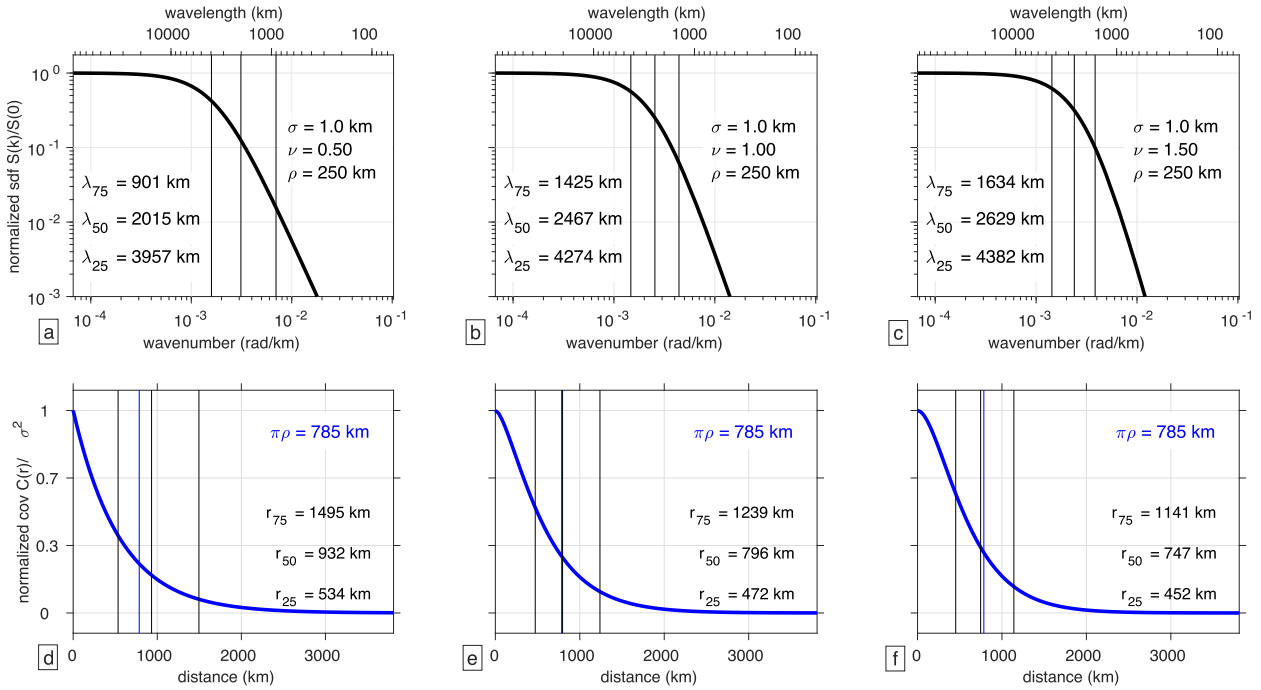


Figure 3. Normalized, radial spectral densities $\mathcal{S}_\theta(k)/\mathcal{S}_\theta(0)$ (a–c, black) as a function of wavenumber k, and normalized spatial covariances $\mathcal{C}_\theta(r)/\sigma^2$ (d–f, blue) as a function of lag distance r. We show three special cases of smoothness of the Matérn covariance class (see Tables A1 and A2), including $\nu = 1/2$ (a, d), $\nu = 1$ (b, e) and $\nu = 3/2$ (c, f), for a common variance $\sigma^2 = 1\,\mathrm{km}^2$ and range $\rho = 250\,\mathrm{km}$ over a regular rectangular 3840 km × 3840 km grid with even 30 km spacing. (a–c) Vertical black lines at λ_α correspond to percentages of cumulative spectral variance, λ_{25}, λ_{50}, λ_{75}, appearing left to right on the top wavelength axis. (d–f) Vertical black lines at r_α correspond to percentages of cumulative spatial variance (eq. 6), r_{25}, r_{50}, r_{75}, appearing left to right on the bottom distance axis, and vertical blue lines indicate a distance of $\pi\rho = 785\,\mathrm{km}$, where the correlation reaches about one-third, depending on ν.

In Fig. 3, the bottom row is annotated at lag distances r_α that correspond to 25, 50 and 75 per cent of the total spatial covariance by thin, black vertical lines. Likewise, using a thin, blue vertical line, we mark the lag distance $\pi\rho$, commonly quoted as the point at which the correlation decays to 1/3. We find that the quality of that approximation, as to whether $\pi\rho \approx r_{33}$, is strongly dependent on the smoothness ν.

2.2 Simulation of spatial random fields from parametric covariance models

We generate synthetic spatial random fields $\mathcal{H}(\mathbf{x})$ as a stationary Gaussian process through the circulant embedding (C.R. Dietrich & G.N. Newsam 1997; G. Chan & A.T.A. Wood 1999; D.P. Kroese & Z.I. Botev 2015) of the spatial covariance model (eq. 1, Table A1) on the sampling grid (for an embedding that is in practice $2^d\times$ as large). For a 2-D, $\mathrm{N_y}\times\mathrm{N_x}$, observation window discretized by a spacing of $\Delta y\times\Delta x$, we define $w(\mathbf{x})$ as the data taper that indicates the presence or absence of data observed on the discretized grid, that is, $w(\mathbf{x})\in\{0,1\}^{\mathrm{N_y}\times\mathrm{N_x}}$, or as a smoothing function $w(\mathbf{x})\in[0,1]^{\mathrm{N_y}\times\mathrm{N_x}}$ used as non-binary weights. Hence, we consider irregularities and incompleteness in the sampling of $\mathcal{H}(\mathbf{x})$ as modifications to an otherwise unitary $w(\mathbf{x})$ that spans an encompassing regular grid. We define $K=\sum w(\mathbf{x})\leq \mathrm{N_yN_x}$.

The circulant embedding method is advantageous in that it produces, on average, a zero-mean random field without the wrap-around effects nor omission of wavenumber correlations that can burden spectral-domain simulation techniques (e.g. via Cholesky decomposition of spectral densities, C.R. Dietrich & G.N. Newsam 1997). Successful embedding requires that the circulant matrix of the spatial covariance is positive semi-definite, a condition met in practice by grids that are sufficiently large relative to the correlation length. Independent realizations generated by the circulant embedding method make for sampled, spatial random fields that possess the requisite second-order statistics.

2.3 Maximum-likelihood estimation of parametric covariance models from spatial random fields

We estimate the parametric covariance of an observed field, $\mathcal{H}(\mathbf{x})$, as the parameters $\hat{\boldsymbol{\theta}}$ that maximize the blurred, debiased Whittle likelihood

$$\bar{\mathcal{L}}(\boldsymbol{\theta})=-\frac{1}{K}\sum_{\mathbf{k}}\left[\ln\bar{\mathcal{S}}_{\boldsymbol{\theta}}(\mathbf{k})+\frac{|H(\mathbf{k})|^2}{\bar{\mathcal{S}}_{\boldsymbol{\theta}}(\mathbf{k})}\right]\quad\text{where}\quad H(\mathbf{k})\equiv\frac{1}{2\pi}\left(\frac{\Delta y\Delta x}{\mathrm{N_yN_x}}\right)^{\frac{1}{2}}\sum_{\mathbf{x}}w(\mathbf{x})\mathcal{H}(\mathbf{x})e^{-i\mathbf{k}\cdot\mathbf{x}},\tag{7}$$

following eq. (5) of A.P. Guillaumin *et al.* (2022) and eqs (48) and (49) of F.J. Simons *et al.* (2026). All summations are limited to the spatial grid $\mathbf{x}$ and the Nyquist grid $\mathbf{k}$. By 'modifying' the data $\mathcal{H}(\mathbf{x})$ by the spatial observation window $w(\mathbf{x})$, sampling bias for regular and irregular data can be corrected for in the periodogram as a non-parametric estimator (e.g. R. Dahlhaus & H. Künsch 1987; F.J. Simons *et al.* 2000; S. Deb *et al.* 2017), but in the likelihood, that term, $|H(\mathbf{k})|^2$, must not be directly compared to the theoretical parametrized spectral density $\mathcal{S}_{\boldsymbol{\theta}}(\mathbf{k})$ unless it is smoothed, for example, via convolution with the spectral window $|w(\mathbf{k})|^2$, turning it into the 'blurred' spectral density denoted with the overbar as $\bar{\mathcal{S}}_{\boldsymbol{\theta}}(\mathbf{k})$. The blurring operation earns eq. (7) the designation *debiased*. Compared to space-based estimation methods which require the construction and *inversion* of the empirical spatial covariance matrix (J.A. Goff & T.H. Jordan 1988), after Fast Fourier Transformation to the wave vector domain, the simple spectral *division* in eq. (7) is efficiently evaluated.

Calculating the debiased Whittle likelihood function requires the window-modified empirical periodogram $|H(\mathbf{k})|^2$, formed from the Fourier-transformed data vector $H(\mathbf{k})$. In eq. (7), the periodogram is equivalent *in expectation* to the blurred spectral density $\bar{\mathcal{S}}_{\boldsymbol{\theta}}(\mathbf{k})$ by which it is divided, as shown by Lemma 1 of A.P. Guillaumin *et al.* (2022) and eqs (28)–(30) of F.J. Simons *et al.* (2026). It is calculated exactly as

$$\bar{\mathcal{S}}_{\boldsymbol{\theta}}(\mathbf{k})=\frac{1}{(2\pi)^2}\left(\frac{\Delta y\Delta x}{\mathrm{N_yN_x}}\right)\sum_{\mathbf{y}}W(\mathbf{y})\,\mathcal{C}_{\boldsymbol{\theta}}(y)e^{-i\mathbf{k}\cdot\mathbf{y}}\quad\text{where}\quad W(\mathbf{y})=\sum_{\mathbf{x}}w(\mathbf{x})w(\mathbf{x}+\mathbf{y}),\tag{8}$$

that is, we form $\bar{\mathcal{S}}_{\boldsymbol{\theta}}(\mathbf{k})$ as the Fourier transform of the isotropic spatial covariance $\mathcal{C}_{\boldsymbol{\theta}}(y)$ modified by the correlation of the sampling window $W(\mathbf{y})$ computed over the relevant separation grid with lag distances $\boldsymbol{y}$. The blurred spectral density $\bar{\mathcal{S}}_{\boldsymbol{\theta}}(\mathbf{k})$ can be computed using a multidimensional Fast Fourier Transform for sampling patterns that are modulations of regular grids (A.P. Guillaumin *et al.* 2022, their Lemma 2). The explicit dependence of eqs (7) and (8) on spatial wave vector $\mathbf{k}$ rather than isotropic wavenumber $k=||\mathbf{k}||$ is intentional, as even for an isotropic model, the operation of blurring is not radially symmetric. This is evident as the autocorrelation of the window $W(\mathbf{y})$ is clearly not isotropic, even for a full rectangular grid. Compared to eq. (8), in the original formulation of F.J. Simons & S.C. Olhede (2013), the blurring of the spectral density by the sampling window $w(\mathbf{x})$ was inexact, through numerically discretized convolution with the *spectral* window $|w(\mathbf{k})|^2$ (the Fejér kernel in the unitary case), for which they used a slightly different notation.

In Fig. 4, we illustrate the equivalence of $\langle|H(\mathbf{k})|^2\rangle$ to $\bar{\mathcal{S}}_{\boldsymbol{\theta}}(\mathbf{k})$, along with several of the quantities that appear in the debiased Whittle likelihood (eq. 7), which are significant to the spectral-domain method. One realization of a spatial field $\mathcal{H}(\mathbf{x})$ simulated from a known $\boldsymbol{\theta}_0$ and sampled on a rectangular observation grid is shown in Fig. 4(a). Missingness in its sampling pattern is shown as white in the colourmap. In the data taper $w(\mathbf{x})$ shown in Fig. 4(d) the light colour indicates observations and dark colour indicates missingness.

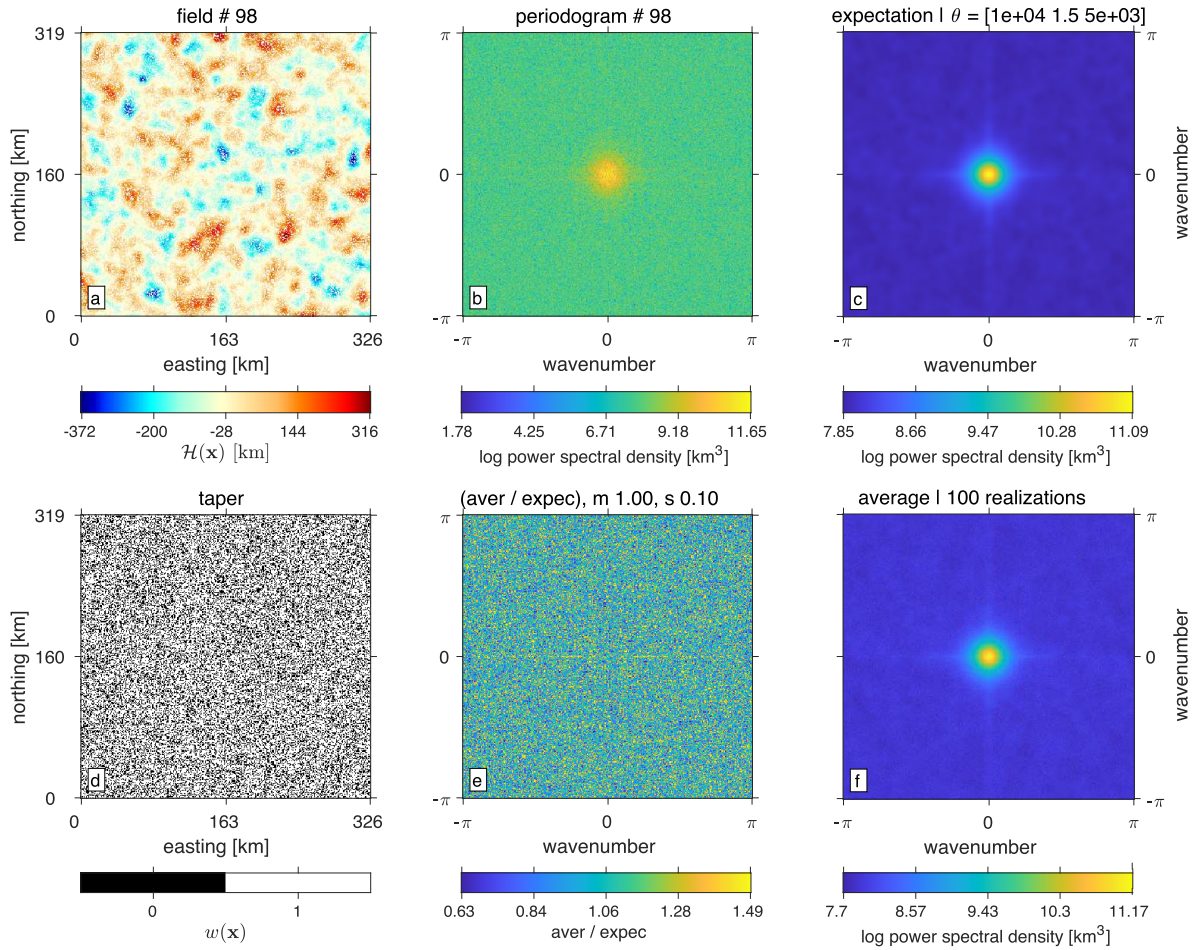


Figure 4. Maximum-likelihood estimation (Section 2.3) on a randomly subsampled (66.7 per cent observed) grid: the effect of spectral blurring. (a) A single realization of a spatial random field $\mathcal{H}(\mathbf{x})$ simulated through the circulant embedding of the Matérn covariance with $\boldsymbol{\theta} = [10\,\mathrm{km}^2\ 1.5\ 5\,\mathrm{km}]$ on a 319×326 grid with $1\,\mathrm{km} \times 1\,\mathrm{km}$ spacing. (b) The empirical periodogram of the field $|H(\mathbf{k})|^2$ in the Nyquist wavenumber plane $[-\pi\ \pi]\,\mathrm{km}^{-1}$. (c) The blurred spectral density $\bar{\mathcal{S}}_{\boldsymbol{\theta}}(\mathbf{k})$. (d) Spatial taper $w(\mathbf{x})$ indicating the presence or absence of data in white and black, respectively. (e) The ratio of the average empirical periodogram over 100 realizations to the blurred spectral density. (f) The average empirical periodogram. The empirical periodogram approaches the blurred spectral density in expectation. Boundary effects and sampling irregularities spoil the isotropy of both, while keeping their ratio unstructured and uninformative.

This sampling pattern was created through the uniformly random deletion of one-third of an otherwise evenly sampled grid. The empirical periodogram $|H(\mathbf{k})|^2$ of the spatial realization is shown in Fig. 4(b). The blurred spectral density $\bar{\mathcal{S}}_{\boldsymbol{\theta}}(\mathbf{k})$ evaluated at $\hat{\boldsymbol{\theta}}$ is shown in Fig. 4(c). In contrast to the regularly sampled rectangular fields shown by F.J. Simons *et al.* (2026), the blurred spectral density displayed here demonstrates the loss of not only isotropy but mirror symmetry, too, which arises from the sampling pattern, in the blurred model. From many independent experimental realizations, we calculate the average periodogram $\overline{|H(\mathbf{k})|^2}$, shown in Fig. 4(f), and compare this in ratio to the blurred spectral density in Fig. 4(e), which displays very low perturbations that lack structure. We interpret this as evidence that, in expectation, the periodogram approaches the modelled blurred spectral density.

Maximizing the debiased Whittle likelihood (eq. 7) yields an asymptotically unbiased, random estimator $\hat{\boldsymbol{\theta}}$, such that its expectation is the true underlying process model parameter set, $\langle\hat{\boldsymbol{\theta}}\rangle = \boldsymbol{\theta}_0$. Defining the blurred *score* vector $\bar{\boldsymbol{\gamma}}(\boldsymbol{\theta})$ as the gradient of the blurred likelihood function $\bar{\mathcal{L}}(\hat{\boldsymbol{\theta}})$, evaluation at the local optimum $\hat{\boldsymbol{\theta}}$ will in expectation yield the zero-vector

$$\nabla_{\boldsymbol{\theta}}\bar{\mathcal{L}}(\hat{\boldsymbol{\theta}}) = \bar{\boldsymbol{\gamma}}(\hat{\boldsymbol{\theta}}) = \mathbf{0}. \tag{9}$$

We find that various iterative numerical methods are suitable for the optimization of eq. (7), including gradient-based (both constrained and unconstrained) methods and non-gradient, simplex-based methods. In practice, we choose to implement an unconstrained quasi-Newton algorithm for ease of comparison between numerically calculated gradients and Hessians of the likelihood function, evaluated at the numerical estimate $\hat{\boldsymbol{\theta}}$, with our evaluations of the analytical partial derivatives of eq. (7). Our initialization of the numerical optimizer randomly perturbs guesses of the Matérn parameters that we propose as follows: We choose an initial σ_0^2 as the variance of the spatial field as observed within the sampling window $w(\mathbf{x})$, we assign a guess for ν_0 at a value we take to be intermediate to the typical range found in practice ($\nu_0 = 2$), and we arbitrarily calculate an initial ρ_0 relative to the grid spacing and the area of the grid as $(1/20\pi)\sqrt{\Delta y \Delta x \mathrm{N_y N_x}}$. While practitioners may provide initialization values for the parameters, or fix any of their values as

hyperparameters, all of the initialized guesses we use in practice follow this automated strategy that is independent of prior knowledge of the truth.

Upon determining an estimator $\hat{\boldsymbol{\theta}}$ for the random field $\mathcal{H}(\mathbf{x})$ within the observation window $w(\mathbf{x})$, we must assess our confidence in the estimator and how well the model, with all its inherent assumptions (i.e. local stationarity, Gaussianity, isotropy and a local non-varying mean), complies with the observed data.

2.4 Residual diagnostics and their interpretation

In the spectral domain, our spectral residuals are the ratio of the periodogram to the blurred spectral density parametrized by our estimator,

$$X_{\boldsymbol{\theta}}(\mathbf{k}) = \frac{|H(\mathbf{k})|^2}{\bar{S}_{\boldsymbol{\theta}}(\mathbf{k})}. \tag{10}$$

In expectation, the periodogram $|H(\mathbf{k})|^2$ approximates the blurred spectral density $\bar{S}_{\boldsymbol{\theta}}(\mathbf{k})$, and for a (Gaussian) process their ratio scales (asymptotically) like a chi-squared distribution with two degrees of freedom, except at 0 and the Nyquist wavenumbers (F.J. Simons *et al.* 2026, their section 4.6), hence,

$$X_{\boldsymbol{\theta}}(\mathbf{k}) \sim \frac{\chi^2_2}{2} \qquad \text{and} \qquad s^2_X = \frac{1}{K}\sum_{\mathbf{k}}[X_{\boldsymbol{\theta}}(\mathbf{k}) - 1]^2 \xrightarrow{\text{law}} \mathcal{N}(1, 8/K). \tag{11}$$

The convergence of s^2_X for data $H(\mathbf{k})$ that abide by model assumptions forms a basis for quantitative assessments. The appropriateness of the covariance model as stationary, isotropic and parametric, and the adequacy of the sample size K, are addressed by posing the Matérn estimator as the null hypothesis to be evaluated based on the test statistic in eq. (11).

While we report s^2_X and the results of the test as such in our data examples, these values are to be used for initial guidance only, as the relations in eq. (11) are asymptotic, and sample sizes are often too small. A much more valuable qualitative assessment tool is the visual inspection of the spectral residuals $X_{\boldsymbol{\theta}}(\mathbf{k})$ for unmodelled structure.

In the spatial domain, we employ simulation as a secondary qualitative tool that allows us to compare by eye the appropriateness of a synthetic spatial field that possesses the first- and second-order structure estimated from the observed spatial field (e.g. Fig. 1). This is useful for real data examples where we appraise the average variability of the data, rather than unique or deterministic features.

2.5 Analytical prediction of estimator covariance

Obtaining confidence intervals on our estimated parameters requires their covariance for the given sampling pattern $w(\mathbf{x})$. This is possible analytically and exactly without direct reference to any particular data realization, which provides us the opportunity to investigate the asymptotic behaviour of our estimator under various designed sampling plans. Our exact and analytical calculations are consistent with copious empirical estimates of model covariance from simulation-derived estimator ensembles across the wide range of sampling patterns we have generated. Here, we summarize the parameter covariance calculation method that we find most intuitive and which emphasizes the role of the sampling window $w(\mathbf{x})$ and its correlation $W(\mathbf{x})$.

Following A.P. Guillaumin *et al.* (2022) and F.J. Simons *et al.* (2026), the parameter covariance for the estimated model is determined from $\mathrm{cov}\{\tilde{\boldsymbol{\gamma}}(\boldsymbol{\theta})\}$, the covariance of the blurred score (the gradient of the log likelihood) and $\bar{\boldsymbol{\mathcal{F}}}^{-1}$, the blurred inverse Fisher information matrix (the inverse of the curvature of the log likelihood), both evaluated in practice at the estimator,

$$\mathrm{cov}(\hat{\boldsymbol{\theta}}) \approx \bar{\boldsymbol{\mathcal{F}}}^{-1}(\boldsymbol{\theta}_0)\,\mathrm{cov}\{\tilde{\boldsymbol{\gamma}}(\boldsymbol{\theta}_0)\}\,\bar{\boldsymbol{\mathcal{F}}}^{-1}(\boldsymbol{\theta}_0). \tag{12}$$

Asymptotic expressions often further approximate $\mathrm{cov}(\hat{\boldsymbol{\theta}}) \approx \bar{\boldsymbol{\mathcal{F}}}^{-1}$, but we have rarely found ourselves in the regime where that holds comfortably, see Section 8. The blurred score vector, which we minimize per eq. (9), is

$$\tilde{\gamma}_\theta(\boldsymbol{\theta}) = -\frac{1}{K}\sum_{\mathbf{k}} \bar{m}_\theta(\mathbf{k})\left[1 - \frac{|H(\mathbf{k})|^2}{\bar{S}_{\boldsymbol{\theta}}(\mathbf{k})}\right], \tag{13}$$

which depends on the sampling pattern $w(\mathbf{x})$ through the exactly blurred spectral density (eq. 8) and window-modified periodogram (eq. 7), as well as the partial derivatives of the logarithmic blurred spectral density presented as

$$\bar{m}_\theta(\mathbf{k}) = \bar{S}^{-1}_{\boldsymbol{\theta}}(\mathbf{k})\frac{\partial \bar{S}_{\boldsymbol{\theta}}(\mathbf{k})}{\partial\theta} = \frac{\bar{S}^{-1}_{\boldsymbol{\theta}}(\mathbf{k})}{(2\pi)^2}\left(\frac{\Delta y \Delta x}{\mathrm{N_y N_x}}\right)\sum_{\mathbf{y}} W(\mathbf{y})\,\frac{\partial \mathcal{C}_{\boldsymbol{\theta}}(y)}{\partial\theta}e^{-i\mathbf{k}\cdot\mathbf{y}}. \tag{14}$$

All of the partial derivatives $\partial_\theta \mathcal{C}_{\boldsymbol{\theta}}$ have closed-form analytical expressions. The components of the Hessian, given by

$$F_{\theta\theta'}(\boldsymbol{\theta}) = -\frac{1}{K}\sum_{\mathbf{k}}\left[\frac{\partial \bar{m}_{\theta'}(\mathbf{k})}{\partial\theta} + \left\{\bar{m}_\theta(\mathbf{k})\bar{m}_{\theta'}(\mathbf{k}) - \frac{\partial \bar{m}_{\theta'}(\mathbf{k})}{\partial\theta}\right\}\frac{|H(\mathbf{k})|^2}{\bar{S}_{\boldsymbol{\theta}}(\mathbf{k})}\right], \tag{15}$$

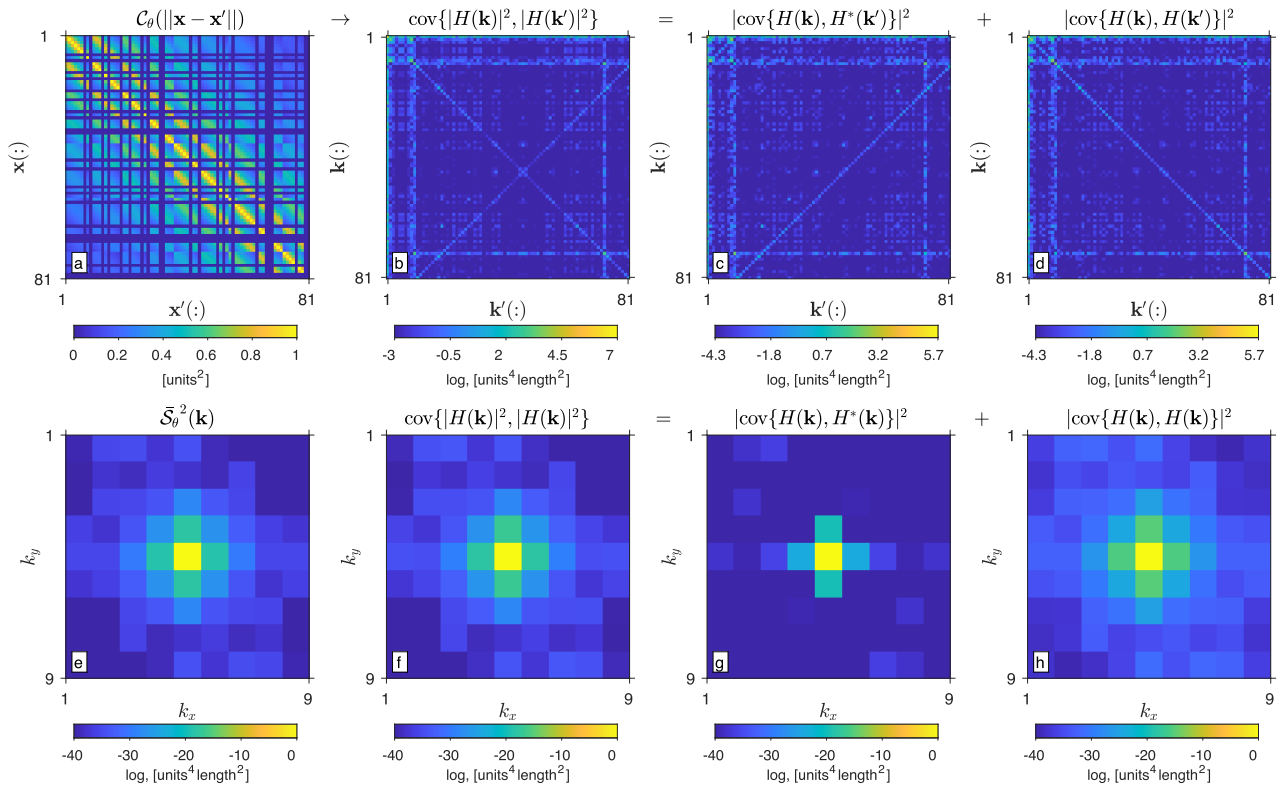


Figure 5. Maximum-likelihood estimation (Section 2.3) on a randomly subsampled (66.7 per cent observed) grid: the effect of wavenumber correlation. The covariance and pseudo-covariance terms of the periodogram covariance (eq. 18) for $\boldsymbol{\theta} = [1\,\mathrm{unit}^2\ 0.5\ 2\,\mathrm{length}]$. The product of the observed grid taper with the spatial autocovariance over all lag interactions (a) reveals a tartan pattern of negligibly contributing lag pairs that manifests in the spectral-domain terms that include all wavenumber interactions (b–d) as elongated hatches of diffuse power in contrast to the fully observed grid, concentric about the intersection of the diagonals of the covariance and pseudo-covariance. When $\mathbf{k} = \mathbf{k}'$ (e–h), the random deletions of the observation grid display an oblique symmetry in the periodogram variance.

and its negative expectation, the Fisher information matrix, both depend on the product of eq. (14) evaluated for parameter interactions,

$$\bar{\mathcal{F}}_{\theta\theta'}(\boldsymbol{\theta}) = \frac{1}{K}\sum_{\mathbf{k}} \bar{m}_\theta(\mathbf{k})\ \bar{m}_{\theta'}\ (\mathbf{k}). \tag{16}$$

The covariance of the blurred score with full wave vector correlations,

$$\mathrm{cov}\big\{\bar{\gamma}_\theta, \bar{\gamma}_{\theta'}\big\} = \frac{1}{K^2}\sum_{\mathbf{k}}\sum_{\mathbf{k}'} \bar{m}_\theta(\mathbf{k})\,\frac{\mathrm{cov}\{|H(\mathbf{k})|^2, |H(\mathbf{k}')|^2\}}{\bar{\mathcal{S}}_{\boldsymbol{\theta}}(\mathbf{k})\bar{\mathcal{S}}_{\boldsymbol{\theta}'}(\mathbf{k}')}\ \bar{m}_{\theta'}\ (\mathbf{k}'), \tag{17}$$

depends on values calculated from eqs (8) and (14), as well as the covariance of the window-modified periodogram and with correlated wave vectors $\mathrm{cov}\{|H(\mathbf{k})|^2, |H(\mathbf{k}')|^2\}$. We calculate this term by applying the L. Isserlis (1918) theorem for Gaussian processes (D.B. Percival & A.T. Walden 1993; A.T. Walden *et al.* 1994; M.L. Stein 1995; F.J. Simons *et al.* 2026), which we have assumed holds for $\mathcal{H}(\mathbf{x})$, so that

$$\begin{aligned}
&\mathrm{cov}\big\{|H(\mathbf{k})|^2, |H(\mathbf{k}')|^2\big\} = |\mathrm{cov}\{H(\mathbf{k}), H^*\ (\mathbf{k}')\}|^2 + |\mathrm{cov}\{H(\mathbf{k}), H\ (\mathbf{k}')\}|^2 && (18)\\
&= |\langle H(\mathbf{k})H(\mathbf{k}')\rangle|^2 + |\langle H(\mathbf{k})H^*\ (\mathbf{k}')\rangle|^2 && (19)\\
&= \big|\mathsf{Q}\ \langle \mathcal{H}(\mathbf{x})\mathcal{H}(\mathbf{x}')\rangle\ \mathsf{Q}'\big|^2 + \big|\mathsf{Q}\ \langle \mathcal{H}(\mathbf{x})\mathcal{H}(\mathbf{x}')\rangle\ \mathsf{Q}^{*\prime}\big|^2 && (20)\\
&= \big|\mathsf{Q}\ \tilde{C}_{\boldsymbol{\theta}}\ \mathsf{Q}'\big|^2 + \big|\mathsf{Q}\ \tilde{C}_{\boldsymbol{\theta}}\ \mathsf{Q}^{*\prime}\big|^2, && (21)
\end{aligned}$$

where Q is the discrete Fourier transform matrix, and $'$ and $^{*\prime}$ denote the transpose and conjugate transpose. In our case of non-unitary tapers, the elements of the spatial covariance in eq. (21) are weighted by the outer product of the sampling window $w(\mathbf{x})$ with itself, denoted $\tilde{C}_{\boldsymbol{\theta}}$.

Implementing eq. (21) has proven to be computationally tractable for problems under $K = 10\,000$ samples on a multi-CPU desktop computer (equipped with Intel Xeon w7-3455 processor with 24 2.5–4.8 GHz cores with 64 GB memory) as well as on a contemporary MacBook Pro. To handle larger sample sizes, as in Fig. 1, we switch to a slower method that accesses the wave vector correlation matrix diagonal per diagonal as described by F.J. Simons *et al.* (2026), with identical results.

Fig. 5 illustrates the calculation of the periodogram covariance (eq. 18) through the discrete Fourier transform matrix implementation (eq. 21) for a small grid that is only 2/3 observed due to random deletions in its sampling pattern. Fig. 5, which, when compared

to fig. 6 in F.J. Simons *et al.* (2026), demonstrates how, for an irregularly sampled $w(\mathbf{x})$, asymmetric anisotropy is imparted in the periodogram's covariance (and thus the blurred spectral density). The top row (Figs 5a–d) visualizes each of the terms that appear in eqs (18)–(21). Fig. 5(b) in the top row labelled $\mathrm{cov}\{|H(\mathbf{k})|^2, |H(\mathbf{k}')|^2\}$ includes all wavenumber correlation terms. Fig. 5(f) in the bottom row labelled $\mathrm{cov}\{|H(\mathbf{k})|^2, |H(\mathbf{k})|^2\}$ only shows the diagonal ($\mathbf{k} = \mathbf{k}'$) terms, that is, the periodogram variance, where the diagonals of the second through fourth columns have been wrapped. The nature of the anisotropy depends on the sampling pattern, with the amplitude at $\mathbf{k}, \mathbf{k}'$ further depending on $\boldsymbol{\theta}$. We find a near equivalence of $\mathrm{cov}\{|H(\mathbf{k})|^2, |H(\mathbf{k})|^2\}$ (Fig. 5f) with $\bar{\mathcal{S}}^2_{\boldsymbol{\theta}}(\mathbf{k})$ (Fig. 5e).

Our ability to analytically calculate periodogram covariance (eq. 18), the covariance of the blurred score (eq. 17) and ultimately estimator covariance (eq. 12) is a boon of our methodology. As the covariance between the model parameter estimates only depends on the parameter values and the geometry of the experiment, without reference to a particular data realization, we are able to calculate confidence intervals for single data realizations without the need for additional simulation, and, in addition, we can employ eq. (12) for designing sampling plans that minimize estimation errors or other criteria, which is of great applicability across the sciences.

2.6 How to implement our method—and how to read this paper

Section 2.1 stated the Matérn model and pointed to its two-parameter special cases. Section 2.2 defined the observation window $w(\mathbf{x})$, and formalized the approach to creating synthetics that fully conform to all model assumptions, allowing for simulations and scenario testing. In Sections 2.3 (formulating the Whittle likelihood and maximizing it to obtain the Matérn model parameters), 2.4 (calculating residuals and interrogating them for model departures) and 2.5 (quantifying parameter uncertainties and inspecting their correlations to enable model comparisons), we discussed the three key steps involved in obtaining an estimate for the covariance parameters of a random field $\mathcal{H}(\mathbf{x})$, as well as tools for their assessment. Section 3 provides the full numerical validation of the robustness of our strategy for irregularly bounded and incompletely sampled fields with a series of visual indicators and a variety of metrics. When compared with their equivalents for the completely sampled, rectangular case considered by F.J. Simons *et al.* (2026), they corroborate that Whittle likelihood analysis of Matérn covariance models need not suffer sampling.

To give the reader a clear picture of how to implement our methodology (and for our computer code, see Section 11) for their own science, we now outline exploratory data analysis and pre-processing steps to help geospatial data sets align with our modelling assumptions.

A basic requirement is that finite-*variance* field data should be collected at an interval and over an extent that are broadly commensurate with domain-specific initial notions of *smoothness* (grid discretization needs to be attuned to the visible variability) and *correlation lengths* (grid size and spacing need to capture characteristic scales of the process) over physical domains that are not patently amalgamated (suggesting lack of *stationarity*) or showing preferred directionality (hinting at loss of *isotropy*). Next, observed data, histograms and summary statistics can be checked for obvious departures from *Gaussianity*. Non-negligible excess kurtosis might be treated by outlier removal, and areas marred by artefacts or revealing a likely non-random nature should be excluded. The spatial window $w(\mathbf{x})$ must capture data availability; it can be crafted to isolate mapped regions of specific interest or provenance, or to select for good behaviour: It should be designed as a mask, taper or weight with these purposes in mind. Lastly, zero-mean randomness is not guaranteed by simple mean removal: Guarding against locally varying means can be accomplished by detrending. In our modelling, we consider polynomials up to order two as deterministic, and we least-squares fit and remove them prior to analysis. Matérn covariance modelling using our procedure may be iterative: Inspecting residuals for unmodelled structure could motivate a reassessment of the data taken to be $\mathcal{H}(\mathbf{x})$.

Not all readers will be interested in reading beyond the numerical validation of statistical and computational efficiency presented in Section 3. Those focused on how to conduct Matérn covariance parameter estimation may wish to skip to Section 11 for our code availability.

Readers invested in utilizing specific models (e.g. exponential, squared-exponential) or deciding why they should use the Matérn class are advised to read Section 4, which deals with model specification by parameter fixing, for example, to its special cases. Section 5 builds the bridge between model selection and sampling patterns, and how these choices or externalities interact. Realistic geoscience sampling scenarios are brought in using synthetics and simulations. Section 6 dwells in more detail on the motivating example presented in our first real-data example of Fig. 1, furthering the discussion on the implications of analysing non-stationary fields, with or without knowledge of the geographical boundary separating them.

Section 7 takes us through a variety of real-data test cases that complement the gallery of geoscientific examples that we presented as F.J. Simons *et al.* (2026). They fit in this repertoire while being incompletely and irregularly sampled because the debiased Whittle likelihood effectively erases sampling-induced bias, while remaining computationally efficient in its spectral-domain formulation. These examples were not selected to prove that the world is governed only by Gaussian processes belonging to the Matérn covariance class; they were chosen to convince the reader that *if* they are, we can find their parameters while being cognizant of the sampling function, we can find and interpret their uncertainties, and we can synthesize simulated observations. If they are not, we can assess residuals for model departure: Such diagnostics help alert us to any and all of the aspects that our model cannot capture, helping identify alternatives and carrying information about the physical process under study (e.g. the directionality of seafloor fractures). Not all data sets will be Gaussian, stationary, isotropic and Matérn, but our analysis under those assumptions as a null hypothesis will not

be unduly influenced by how specifically we collect or mask (e.g. to remove clouds in remote sensing) the data. Thus, in Section 7, we demonstrate the spatial isolation of a homogeneous background from deterministic features, and how through our model evaluation tools, we can interpret the random portion of the data set to be isotropically Matérn. Additionally, we show how to extract disjoint processes from within the same data set in order to show data departures from a model that may well turn out to reveal a feature of physical interest (be it the phenomenon or how it is measured) as they can not be readily attributed to artefacts and biases arising from spatial sampling patterns. Finally, we present a directly observed, presumably single, statistical process, but under a sampling pattern so intricate that we must make sure any inference based on it is not unduly influenced by it. While the sampling pattern explains much of the variability within the periodogram, fine-scale, directional structure apparent within the model residuals shows that some regions of seafloor are not isotropic.

The penultimate section is for the statistically minded: Section 8 helps understand large-sample asymptotic behaviour of the estimators. Having established that taking sampling pattern 'shape' into account matters, sampling 'size' will matter, too. Increasing the *number* of samples taken within a specific 'template' influences the behaviour of the estimators. Growing the domain, reducing data missingness, infilling within a fixed domain: These are all strategies with consequences for parameter recovery for which we developed the machinery to study. Extensions of our methodology from two to d dimensions are mathematically intrinsically contained in our method and understanding them statistically will benefit from this section. Our root-mean-squared error decay with sample size K as $1/\sqrt{K}$ is favourable compared to methods, where parameter estimation bias is expected to dominate over variance with growing dimension (R. Dahlhaus & H. Künsch 1987).

The last section is for aficionados and sampling strategy enthusiasts devoted to experimental design. Section 9 prepares the reader for the most practical question of all: Where should they be adding more data to better understand their model, their science, their field of study? Designing experiments is possible within the context of this paper because we are able to calculate parameter uncertainties via dataless parameter covariance predictions conditioned on the sampling pattern.

3 UNBIASED PARAMETER ESTIMATES FROM INCOMPLETELY SAMPLED DATA

We illustrated the power of the debiased Whittle likelihood in F.J. Simons *et al.* (2026) for completely sampled, evenly spaced, rectangular grids. We recall their and our ability to incorporate exactly, and fast, the effects of sampling and bounding, into the likelihood (eq. 7). In their fig. 2, they showed that the blurred spectral density (eq. 8) matches the average periodogram of the data $\overline{|H(\mathbf{k})|^2}$, and in their fig. 6, they illustrated the geometrical structure of the covariance of the periodogram, $\mathrm{cov}\{|H(\mathbf{k})|^2, |H(\mathbf{k}')|^2\}$, which controls the parameter covariance (eq. 12). In this paper, we have so far, in our Fig. 4, documented the effect on the expected periodogram of missingness through random grid subsampling, and we have, in our Fig. 5, shown the corresponding structure and geometry of the periodogram covariance. In this section, we show how the debiased Whittle likelihood continues to deliver asymptotically unbiased, Gaussian estimates of Matérn parameters for those kinds of sampling scenarios, including how well our uncertainty quantification continues to capture the observed behaviour. We conclude this section by covering the behaviour of the moments of the periodogram for contiguously missing patches—and their complements.

3.1 Randomly subsampled regular grids

We demonstrate the robustness of our estimation strategy and quantitatively evaluate the quality of our estimated models through numerical experiments with stationary Matérn random fields. Figs 6–9 report on an ensemble of simulations and their maximum-likelihood estimates over many realizations from the same model $\boldsymbol{\theta}_0$ for sampling operators $w(\mathbf{x})$ characterized as 66.7 per cent observed rectangular grids. The $w(\mathbf{x})$ we implement for these first experiments all exhibit uniformly random missingness so that we can build intuition without the confounding factors of grid geometry, size or sampling pattern. Evaluating statistics from this ensemble and its realizations provides us with information about how well we can hope to know our parametrically modelled random fields, as a framework for interpretation when we are confronted with real data, when not all of our modelling assumptions will be met. We save the complications on estimators and their covariances related to varying grid geometry, growing sample size and (geographically) structured sampling patterns for later in this manuscript.

We visualize ensembles of parameter estimates for an experiment of 500 realizations of the same model and grid parameters as in Fig. 4. Fig. 6 shows the distribution of the marginals of the model estimates to be approximately normally distributed. Shown are grey kernel density estimates (kdes) with grey overlayed normal distribution fits to the sample mean and standard deviation in the top row, and as quantile–quantile (Q–Q) plot comparisons of the observed estimates to a theoretical normal distribution in the bottom row. From these diagnostics, we see agreement of the numerical estimates with theory in that the centrality of the distribution of estimates is aligned with $\boldsymbol{\theta}_0$ (i.e. $\hat{\boldsymbol{\theta}}$ is an unbiased estimator). The spread of the distribution of estimates is consistent with our analytical calculation of parameter covariance, shown as the thick black line. The analytically calculated parameter standard deviation follows eq. (12), evaluated at the mean of the estimator for this ensemble. The analytical calculation closely matches our numerical experiments for these irregularly sampled fields, as well as it did for the complete sampling scenarios presented by F.J. Simons *et al.* (2026), their

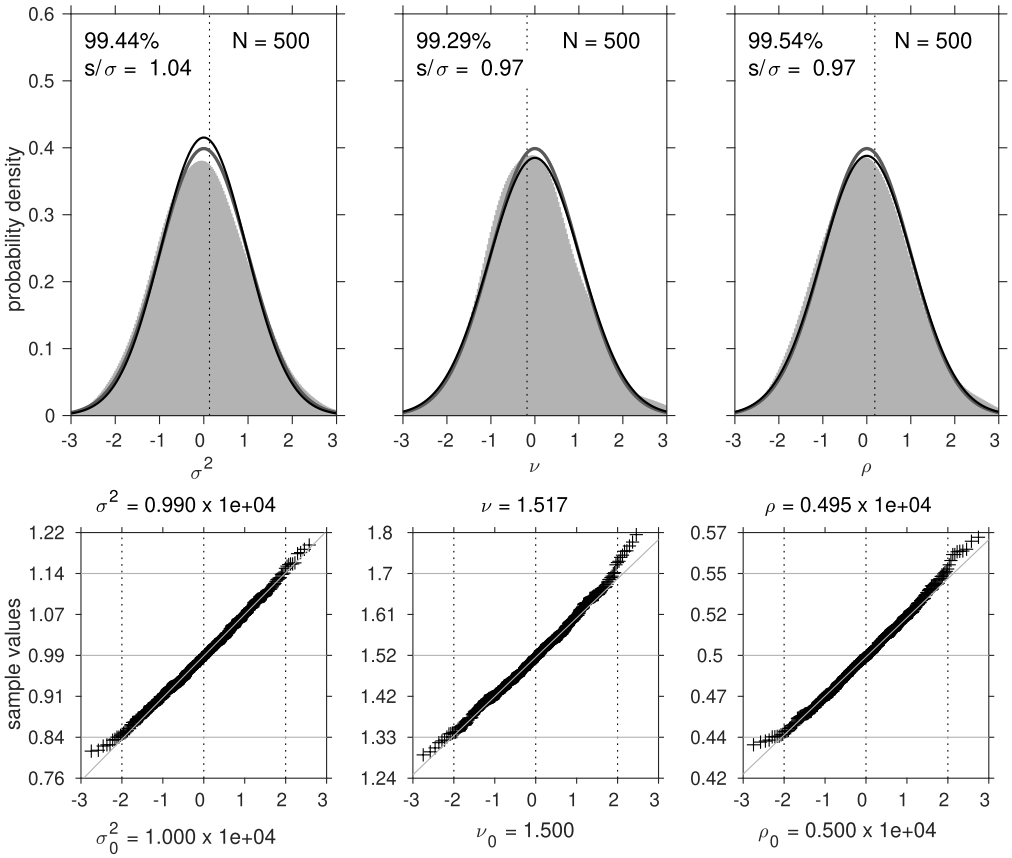


Figure 6. Maximum-likelihood estimation (Section 2.3) on a randomly subsampled (66.7 per cent observed) grid as shown in Fig. 4(a): ensemble statistics. Matérn parameter maximum-likelihood estimation statistics for 500 spatial-covariance embedding simulations carried out on a 319×326 grid with 1 km × 1 km spacing exhibiting uniform random deletions (see Fig. 4) for true Matérn parameter values of $\sigma_0^2 = 10$ km^2, $\nu_0 = 1.5$ and $\rho_0 = 5$ km, recovered via maximization of the exactly blurred, uncorrelated, likelihood (eq. 7). The estimates average to $\sigma^2 = 9.90 \pm 0.75$ km^2, $\nu = 1.52 \pm 0.09$, $\rho = 4.95 \pm 0.26$ km, quoting one observed standard deviation. The thick grey line is derived from the covariance calculated from the ensemble of simulation and estimation outcomes. The thick black line is based on the covariance exactly calculated from eq. (12) shown and calculated at the truth, shown by the dotted vertical line in the top panels. The quantile–quantile plots of the bottom row show the ensemble distribution between its $[-3, 3]$ empirical standard deviations (abscissas) in relation to the theoretical normal distribution parametrized by the sample statistics (ordinates).

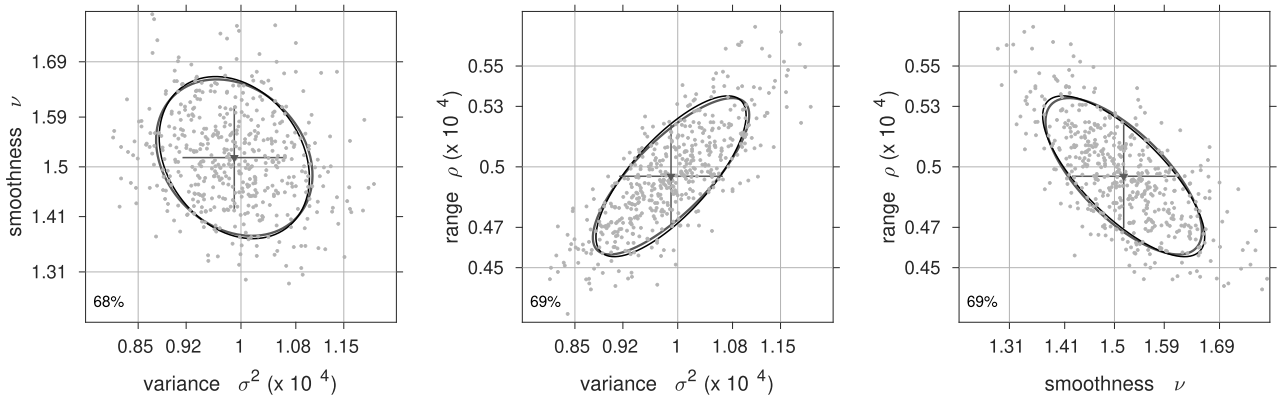


Figure 7. Maximum-likelihood estimation (Section 2.3) on a randomly subsampled (66.7 per cent observed) grid as shown in Fig. 4(a): parameter covariance. Maximum-likelihood estimation statistics for the ensemble of 500 simulation and recovery experiments reported in Fig. 6. The mean estimate is highlighted by the grey triangle and two observed standard deviations marked by grey lines. The heavy grey ellipse is the 68 per cent confidence region based on the ensemble. The closely matching heavy black ellipse is the predicted 68 per cent confidence region based on the covariance predicted from eq. (12). Values quoted in the bottom left are the percentages of estimates falling within the 68.27 per cent predicted confidence interval.

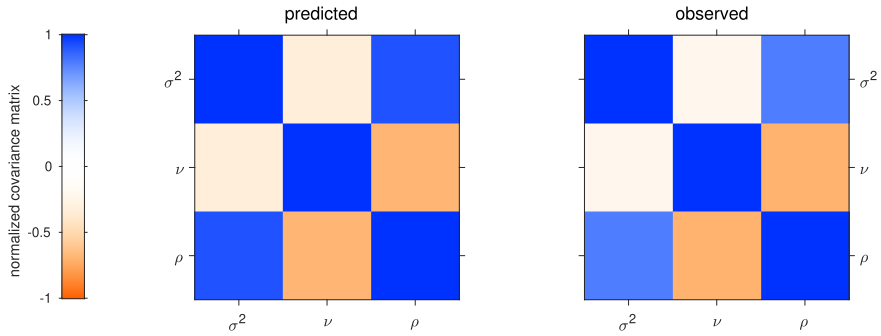


Figure 8. Maximum-likelihood estimation (Section 2.3) on a randomly subsampled (66.7 per cent observed) grid as shown in Fig. 4(a): ensemble correlations. Comparison of the covariance predicted via eq. (12) and the covariance observed on the set of experiments reported in Figs 6 and 7. Shown are the relevant correlation matrices between the estimators for the three Matérn parameters σ^2, ν and ρ, highlighting the relatively strong trade-off between ν and ρ.

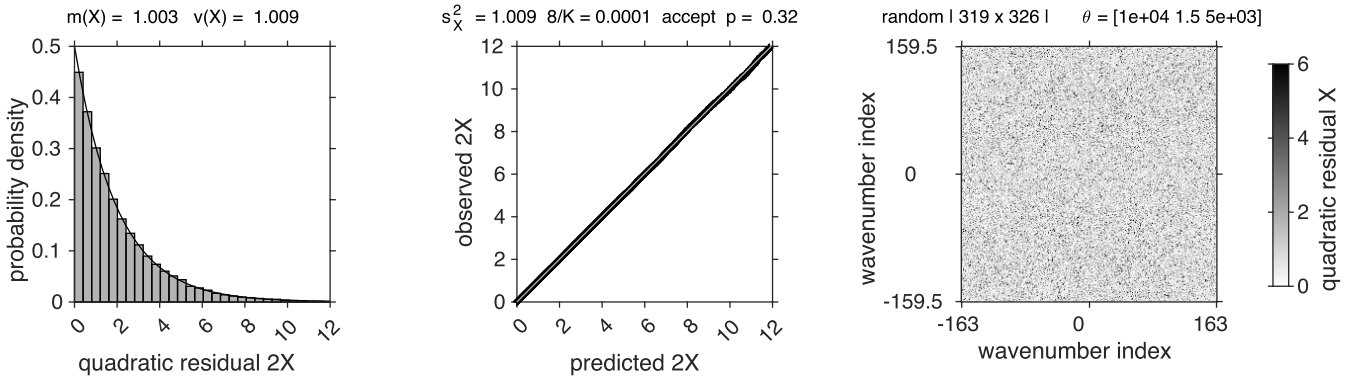


Figure 9. Maximum-likelihood estimation (Section 2.3) on a randomly subsampled (66.7 per cent observed) grid as shown in Fig. 4(a): residual statistics. We report on a single, randomly selected simulation carried out in the experiment shown in Fig. 4(a). As in fig. 11 of F.J. Simons *et al.* (2026), the distribution of the variable $X_{\boldsymbol{\theta}}(\mathbf{k})$ of eq. (10) is shown as a histogram across all wavenumbers with the theoretical distribution superposed, as a quantile–quantile plot for the distribution in question, and as a spectral-domain map. Plot labels refer to (*left*) the mean, $\mathsf{m}(\mathsf{X})$, and variance, $\mathsf{v}(\mathsf{X})$, of $X_{\boldsymbol{\theta}}(\mathbf{k})$, (centre) the test-statistic ($\mathsf{s}^2_{\mathsf{X}}$, eq. 11), variance of the convergence distribution given the number of samples K for the observed grid, whether the rigorous test was rejected, and with what significance (p-value, right), and the grid and model parametrization.

fig. 7. Optimizing the likelihood in eq. (7) correctly accounts for missingness; see Section 8 for a discussion of the asymptotics of this behaviour.

Fig. 7 shows our suite of estimates as parameter cross-plots. Each of the estimates is plotted in the three permutations of the 2-D parameter spaces. The 68 per cent confidence ellipses calculated are drawn from the observed samples, in grey, and from the analytically calculated covariance of eq. (12), in black. Here, for the 66.7 per cent observed field with random deletions, we see a subtle negative correlation between estimators of σ^2 and ν, a strong positive correlation between estimators of σ^2 and ρ and a strong negative correlation between those of ν and ρ. The empirical values are $\{\sigma^2, \nu\} = -0.2254$, $\{\sigma^2, \rho\} = 0.7830$ and $\{\nu, \rho\} = -0.7006$, respectively. The correlations theoretically calculated via eq. (12) are $\{\sigma^2, \nu\} = -0.2569$, $\{\sigma^2, \rho\} = 0.7637$ and $\{\nu, \rho\} = -0.7441$, which are good predictors for those observed. The results for these incompletely sampled fields remain in qualitative agreement with the complete sampling scenarios presented by F.J. Simons *et al.* (2026), their fig. 9. As we will see later, the strength of the parameter correlation is dependent on the sampling geometry, structure and size.

Fig. 8 gives a second means of assessing the correlation of the estimates. We render the normalized parameter covariance as a heat map. The correlation structure observed through experiment is consistent with our prediction, as it was for F.J. Simons *et al.* (2026) in their fig. 10.

Fig. 9 evaluates the model via a graphical representation of the distribution of the spectral residuals $X_{\boldsymbol{\theta}}(\mathbf{k})$ of eq. (10) for the same model, grid and taper as previously detailed in Fig. 4. The histogram and Q–Q plot follow the χ^2_2 theoretical distribution very closely, and the wave vector plot retains no structure. This behaviour mirrors the completely sampled case shown by F.J. Simons *et al.* (2026) in their fig. 11. In scenarios where we evaluate models for real data and the model is rejected, an inspection of such statistics and visualizations allows us to glean information regarding how the statistics of our data set depart from our model assumptions.

3.2 Irregularly bounded patches

Repeating the experiments as illustrated in Figs 6–9 for different non-random sampling scenarios, in which patches with particular boundary geometries are either present or absent, only corroborates our assertion that the debiased Whittle likelihood (eq. 7) produces unbiased Gaussian estimators with fully characterizable uncertainties with powerful diagnostics for evaluating model fitness. Hence, while we have validated this performance extensively, to save space, we are not showing the full results here. Our openly available code (see Section 11) allows the reader to explore all manner of sampling scenarios.

The behaviour of the estimates made via eq. (9) is controlled by the shape of the sampling kernels $w(\mathbf{x})$, their correlations $W(\mathbf{x})$ and the blurring effect of the spectral window $|w(\mathbf{k})|^2$, where $w(\mathbf{k})$ is the Fourier transform of $w(\mathbf{x})$, on the spectral densities $\tilde{\mathcal{S}}_{\boldsymbol{\theta}}(\mathbf{k})$ in eq. (8), which is equal to the expected periodogram of the data $\langle |H(\mathbf{k})|^2 \rangle$ as could be seen from the average $\overline{|H(\mathbf{k})|^2}$ shown in Fig. 4. For a sampling scenario where a New Jersey shaped domain is the region of interest or omission, Fig. 10 shows the analytically blurred spectral densities $\tilde{\mathcal{S}}_{\boldsymbol{\theta}}(\mathbf{k})$ and the average periodogram $\overline{|H(\mathbf{k})|^2}$ over 100 realizations for a Matérn model noted in the caption.

The parameter uncertainties and covariances obtained with eqs (12)–(21) are to be understood geometrically through the covariance of the periodogram $\langle |H(\mathbf{k})|^2, |H(\mathbf{k}')|^2 \rangle$ as shown in Fig. 5 for random sampling, and now illustrated by Fig. 10 for New Jersey and its complement. Sections 4–7 elaborate on this behaviour for cases that carry specific geophysical relevance. After discussing the statistics of asymptotics in Section 8, we take a final look in Section 9 at the effects of aliasing and blurring, sampling and bounding, across a variety of data observation scenarios that arise in the Earth sciences, specifically through the lens of the spectral window.

4 SIMPLIFYING THE MATÉRN DENSITY: FIXING PARAMETERS

Models frequently assumed within the geophysical community (e.g. A. Tarantola & A. Nercessian 1984; A. Frankel & R.W. Clayton 1986; J.A. Goff & T.H. Jordan 1988; G.R. North *et al.* 2011; J.B. Muir & Z.E. Ross 2023) are special cases of the general Matérn form. In Table A1, we provide analytic simplifications of the Matérn spatial covariance (eq. 1) and 2-D spectral density (eq. 3) by specifying the smoothness parameter ν to fixed values that correspond to covariance functions that are commonly selected within the physical sciences. Many of these are motivated by the mean-squared behaviour of the solutions of differential equations (e.g. T. de Karman & L. Howarth 1938; T. Bui-Thanh *et al.* 2014; A. Cárdenas-Avendaño *et al.* 2023). We provide these simplifications for the d-D spectral density (eq. 2) in Table A2.

In relating our forward and inverse modelling of the general three-parameter Matérn spatial and spectral covariance functions to that of a nested model where one or more parameters are fixed *a priori* (e.g. a two-parameter special case), we adjust the parametrization from

$$\boldsymbol{\theta} = [\sigma^2,\ \nu,\ \rho]^T \tag{22}$$

to the set of nested parameters that will be allowed to vary, for example,

$$\boldsymbol{\theta}^{101} = [\sigma^2,\ \cdot,\ \rho]^T, \tag{23}$$

where the superscript indicates the parameters actually inverted for (as 1), or held fixed (as 0).

Beyond taking a specialized covariance function to be representative of the statistics of a data set due to precedent or assumption, prior knowledge of its character may make working with a nested covariance model desirable. One can determine the potential consequences of adopting a special-case parametric covariance model rather than the general form through model selection tests (e.g. Q.H. Vuong 1989) or through simulations conducted to assess estimator ensemble bias and uncertainties. We include the option for estimating nested models with parameters specified as in eq. (23) within the implementation of our maximum-likelihood estimation strategy. Any (or all) parameters may be held fixed to a supplied initial value and the inversion calculations will only optimize the likelihood (eq. 7) by varying the free parameters. In the case of the smoothness parameter ν held fixed to a value that corresponds to one of the special cases (Tables A1 and A2), one may invert for the simplified analytic function rather than eq. (3). In such special nested model cases, we optimize only over variance σ^2 and range ρ (as in eq. 23), reducing the estimation problem to characterizing a covariance structure based upon the amplitude of the random field and a scale factor. Verification calculations of the gradient vector (eq. 13), Hessian (eq. 15) and Fisher (eq. 16) matrix to compare with the observed numerical behaviour and to evaluate eq. (12) at the estimate $\hat{\boldsymbol{\theta}}$, rely on first and second partial derivatives of the simplified nested models, for which we will provide analytic expressions in forthcoming work.

In Fig. 11, we study the influence of an increasingly large smoothness parameter ν on the reliability of our estimation strategy through two numerical experiments. In both experiments, we set a common variance $\sigma_0^2 = 1\,\mathrm{m}^2$ and range $\rho_0 = 3\,\mathrm{m}$, and select 20 logarithmically spaced values of $\nu_0 \in [1, 100]$. We simulate 48 realizations for each of these 20 target models on a common, fully sampled, evenly spaced $128\,\mathrm{m} \times 128\,\mathrm{m}$ grid. In the first experiment (Fig. 11, *left*), we estimate all three parameters of the Matérn model, and in the second (Fig. 11, *centre*), we fix the smoothness correctly to ν_0 and estimate the variance σ^2 and range ρ as two-parameter nested models. In the three-parameter inversion experiment, we observe that as ν_0 increases, all three parameters are well estimated until a point when they begin to accumulate bias which grows until it saturates. The one-standard error (1 se) empirical confidence interval for the variance parameter σ^2 excludes σ_0^2 when $\nu_0 > 4$, while the 1 se interval for the smoothness ν increases until

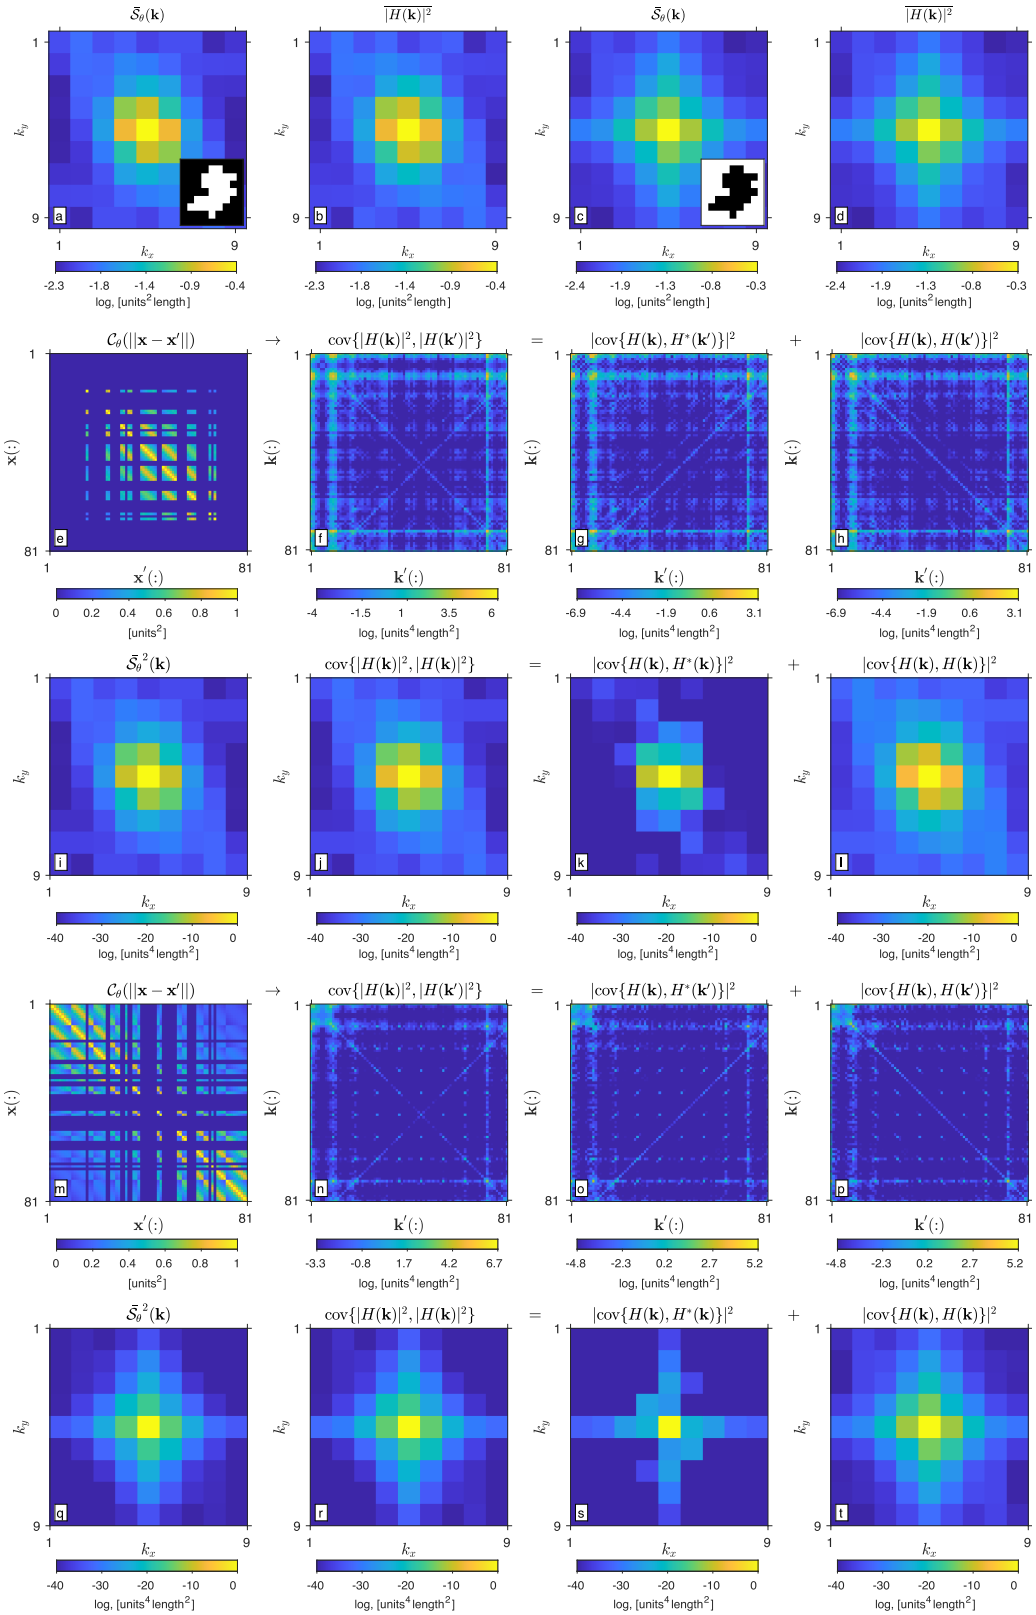


Figure 10. Expected and averaged periodograms with their computed covariances for contiguous, bounded fields. All panels and models are for a 9×9 perimeter with 1 m even spacing and $\boldsymbol{\theta} = [1\,\mathrm{m}^2\ 0.5\ 2\,\mathrm{m}]$. Top, (a–d) The expected periodogram $\bar{\mathcal{S}}_{\boldsymbol{\theta}}(\mathbf{k})$ and the average periodogram over 100 realizations $\overline{|H(\mathbf{k})|^2}$ for a field sampled within the interior of an irregular boundary comprising 68 per cent of the complete grid, and $\bar{\mathcal{S}}_{\boldsymbol{\theta}}(\mathbf{k})$ and $\langle|H(\mathbf{k})|^2\rangle$ for a field sampled from the complementary exterior of the irregular boundary as 32 per cent of the complete grid. The $w(\mathbf{x})$ for the interior and exterior sampling patterns are shown in the subset of the $\bar{\mathcal{S}}_{\boldsymbol{\theta}}(\mathbf{k})$ subfigures. The covariance and pseudo-covariance terms of the periodogram covariance (eq. 18) for the interior (e–l) and exterior (m–t) sampling patterns. Layouts of the periodogram covariance panel sets are as in Fig. 5.

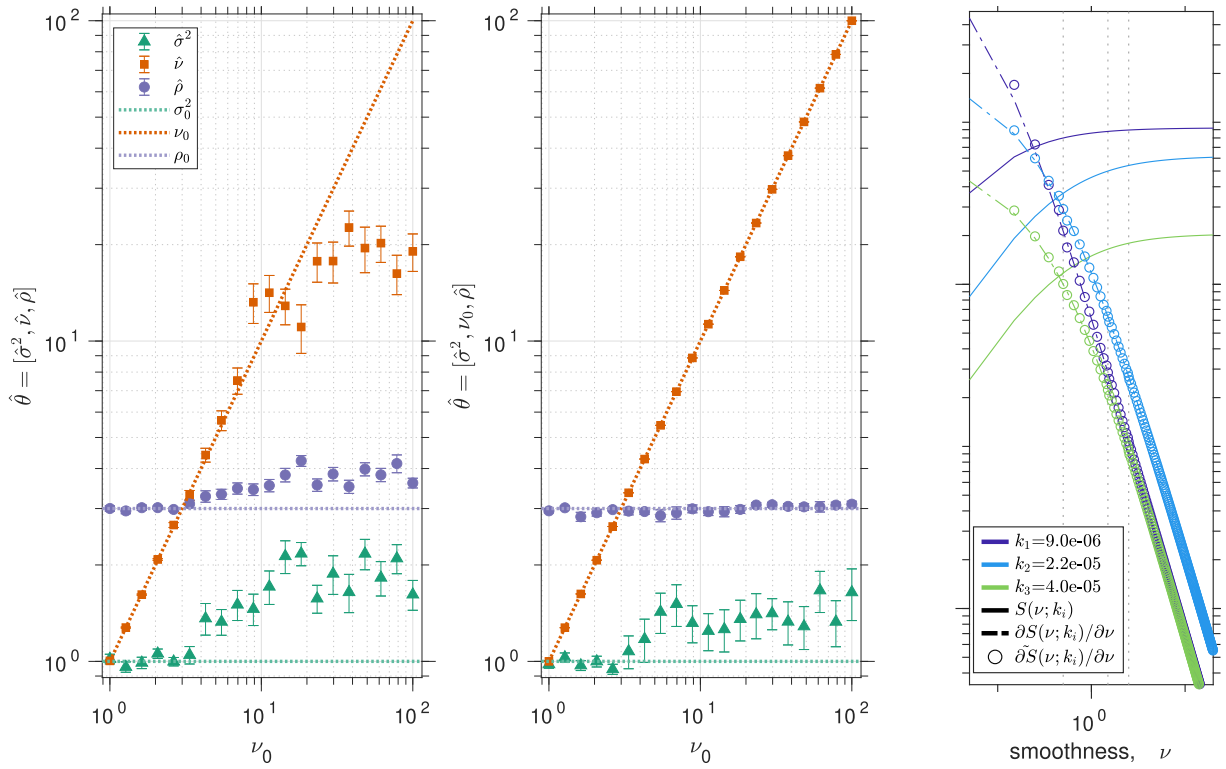


Figure 11. Dependence of Matérn parameter estimates on smoothness. (*Left*) Three-parameter estimates. (*Centre*) Two-parameter estimates with correctly fixed ν_0. The 128 m × 128 m sampling grid is completely, regularly and evenly spaced. Markers show mean estimates $\hat{\theta} = \{\hat{\sigma}^2, \hat{\nu}, \hat{\rho}\}$ for 48 realizations with their empirical standard errors. True values $\theta_0 = \{\sigma_0^2, \nu_0, \rho_0\}$ are displayed as coloured, dotted lines for each trial model where the variance $\sigma_0^2 = 1\,\mathrm{m}^2$ and range $\rho_0 = 3\,\mathrm{m}$ are held constant while the smoothness ν_0 increases from 1 to 100. (*Right*) At three selected wavenumbers k_i, $i = 1, 2, 3$, values of the normalized radial spectra (solid), and their derivatives with respect to the exact (dashed) and numerically approximated (open circles) smoothness are displayed.

about $\nu_0 > 9$, after which $\hat{\nu}$ saturates and hovers around 10–20. The 1 se empirical confidence interval for the range parameter ρ fails to contain ρ_0 when $\nu_0 > 4$. In the second experiment, the 1 se interval for the variance σ^2 exceeds σ_0^2 for $\nu_0 > 5$ and its bias plateaus at about half that of the three-parameter inversion. Not being estimated, ν is plotted along the 1:1 line. The confidence intervals of $\hat{\rho}$ overlap with ρ_0 for all trial models.

The difficulty in accurately estimating all three Matérn parameters for large ν should be expected due to problems with identifiability in 'sloppy' models (G.M. Monsalve-Bravo *et al.* 2022), a trait that can be observed, for example, in the parameter correlation structure of Figs 7 and 8. The strong trade-off between ν and ρ is apparent; when we know ν and fix it, ρ is more adeptly estimated and along with it, σ^2. The failure of $\hat{\nu}$ to approximate large ν_0 suggests there is an upper bound in being able to distinguish variation in ν for smooth fields, depending on sampling interval Δy, Δx. As a function of ν the Matérn spectral density (eq. 3) begins to plateau first for small wavenumbers, followed by intermediate and large wavenumbers. In Fig. 11 (*right*), we explore how the spectral density and its first derivative with respect to the smoothness varies with increasing smoothness for small, medium and large wavenumbers. In the presence of very smooth, large ν, fields, fixing the smoothness parameter should be considered. However, the squared-exponential ($\nu \to \infty$) model is still not an advisable choice due to its accelerated decay with wavenumber, which is numerically problematic without additional precautions taken in the analysis (e.g. including a nugget effect). Analytically, the limiting behaviour $\nu \to \infty$ of the spectral domain illuminates its asymptotically Gaussian nature, which depends on σ^2, ρ and wavenumber k, as in Table A1,

$$\lim_{\nu\to\infty} S(k) = \sigma^2 \left(\frac{\pi\rho^2}{4}\right) \exp\left(-\frac{\pi^2\rho^2 k^2}{4}\right) \qquad \text{for} \qquad k \in \mathbb{R}^+. \tag{24}$$

The experiments shown in Fig. 11 pertained to complete rectangular samplings. We next illustrate the nested special cases of the Matérn form for realistic synthetics comprising three geophysically motivated spatial sampling cases subject to irregular boundaries, structured tracks and random missingness. For each of those we will discuss both correctly and incorrectly selected nested models.

5 MODEL SELECTION AND SAMPLING PATTERNS

In this section, we create realistic synthetics to further elucidate two of the problems we introduced that are regularly encountered in modelling geophysical data: selecting appropriate covariance models (see Section 4) and handling discrete, missing data (see Section 3).

The flexible attributes of the Matérn density are the scale, shape and size parameters. The comparison of three-parameter and two-parameter inversions shown in Fig. 11 illustrated how in the case of large smoothness, the option to fix the smoothness parameter can be advantageous. When the fixed hyperparameter(s) are selected correctly, the behaviour of the estimators is approximately equivalent with that of the general Matérn model. If the value of the fixed parameter is chosen incorrectly, the model is effectively unable to characterize the others, and accuracy in the estimation of the parameters inverted for will suffer. The magnitude of this compensatory parameter bias depends on the spatial sampling of the data.

We now inspect the applicability of two-parameter special cases for three types of realistic sampling patterns: a geographic patch, sweeps of sampling tracks and uniform random missingness. We then address the issue of separability, both in space and in model parameters. We conduct numerical experiments for a variety of sampling patterns (contiguous, structured and dispersed) on the partial grid where we compare summary statistics of estimating the three Matérn parameters (eq. 22) versus two (eq. 23). Specifically, we estimate the σ^2 and ρ that capture the amplitude and length scale of the covariance structure while setting the smoothness ν as a hyperparameter.

Fig. 12 shows these cases where we recover approximately Gaussian distributed parameters centred at their true value for both three- and two-parameter inversions, but only for the latter if we have fixed ν to its true value. We demonstrate and discuss the repercussions of incorrectly selecting a two-parameter special case in terms of parameter bias and uncertainty for the variance σ^2 and range ρ estimates. We consider two overarching synthetic cases for this set of experiments where data availability comprises a third of the encompassing, regular, rectangular grid (Figs 12a–c), and its two-thirds complement (Figs 12d–f). For both cases, we show three sampling scenarios: a geographically motivated, bounded patch in the shape of New Jersey (left column, a, d), a sweep of the seafloor along bathymetric sounding tracks in the South Atlantic (GEBCO Bathymetric Compilation Group 2024, centre column, b, e) and another realization of uniformly random subsampling (right column, c, f).

For each of these sampling scenarios, Fig. 12 shows three subplots with the distributions of the estimators. Within the subplots, we display right-side-up, grey histograms for the three-parameter inversion. These are presented as they were in Fig. 6, as a kde of the estimator ensemble with curves displaying Gaussian distributions parametrized by the ensemble mean and the empirical (dark grey) and analytically calculated (black) standard deviations. The upside-down histograms are two cases of inverting for two parameters, one where the smoothness is held fixed $\nu = \nu_0 = 1$ at the truth (darker, cool colour), and the other held fixed at an incorrect value $\nu = 1/2$ (lighter, warm colour). As expected, the three-parameter inversion, regardless of whether the region of data availability is a New Jersey patch or an Atlantic hatch or a uniform stipple, yields unbiased (visibly centred at the truth $\boldsymbol{\theta}_0$) and (to the eye) Gaussian distributed ensembles, whose empirical parameter uncertainties are well predicted by the analytical calculation of eq. (12). Fig. 12 is our first demonstration of the effect of sampling patterns on the estimators, which so far we had only illustrated geometrically (in Figs 5 and 10) for the specific regions of interest. As the columns in the top and bottom panels each have the same number of data points, the visible distinctions in parameter deviations between the three sets of distributions can be attributed to the geometry of the acquired samples. These ensembles allow us to interpret the effects of boundary structure and regularity in the missingness. The numerical results of the experiments in Fig. 12 are summarized in Table A3, from which we will point out a few key observations.

The three-parameter inversions (upward-facing histograms in Fig. 12) result in unbiased estimators regardless of the sampling window $w(\mathbf{x})$. The spread of the ensembles, however, does show a dependence on $w(\mathbf{x})$. Beginning with the results from the top panel of Figs 12(a)–(c) for 33.4 per cent of the encompassing grid sampled, and progressing from left to right, we will report on the standard deviations of the three Matérn parameters, noting the empirical versus analytical values as ($\cdot$ per cent | $\cdot$ per cent). For the process variance σ^2, the relative standard deviations for the irregular boundary case (Fig. 12a) of the New Jersey patch (16.07 per cent | 17.09 per cent), the structured bathymetric tracks (Fig. 12b) from the Atlantic (8.91 per cent | 9.17 per cent), and the uniformly random subsampling (Fig. 12c) of the grid (6.67 per cent | 6.83 per cent) reveal a decrease in uncertainty with increased dispersion in sampling, which we interpret to be due to the subsampling effectively removing the confounding effects of spatial correlation. We initially observed this effect in the growing-domain experiments in F.J. Simons *et al.* (2026), their figs 4 and 5, and we will return to this in Section 8.1. In contrast, for the smoothness parameter ν, we see from the New Jersey (3.28 per cent | 4.06 per cent), Atlantic (13.03 per cent | 13.62 per cent) and uniform (11.10 per cent | 11.61 per cent) sampling experiments that more contiguous sampling patterns increase the precision on ν, revealing greater process discrimination with denser sampling. We will be studying this effect in Sections 8.2 and 8.3 through the perspective of infill asymptotics for regular and irregular allocations of additional samples. For the correlation length ρ, the New Jersey (9.27 per cent | 10.12 per cent), Atlantic (9.96 per cent | 10.53 per cent) and uniform (6.34 per cent | 7.37 per cent) samplings reduce the spread for the uniform random subsampling pattern only.

The two-parameter inversions (downward-facing, light coloured histograms in Fig. 12) display a pronounced bias in estimated variance σ^2 and range ρ when the smoothness ν is fixed to an incorrect value (in other words, when we chose the wrong model). For σ^2, the bias (sample mean of the estimators, $\overline{\sigma^2}$ versus truth σ_0^2) for the New Jersey (695 versus 1000 units^2), Atlantic (833 versus 1000 units^2) and uniform (972 versus 1000 units^2) sampling reveals a decrease in bias as the sampling becomes more random because the boundary terms no longer dominate the estimation procedure. Needless to say, the bias in ν is intentional because we (knowingly) fixed it to the incorrect value. For ρ, the bias for the New Jersey (58.1 versus 20.0 km), Atlantic (32.5 versus 20.0 km) and uniform (24.1 versus 20.0 km) sampling decreases as we further randomize sampling, which we interpret, too, as being due to reduced boundary effects. When the smoothness parameter is fixed correctly, however, the estimates of the two-parameter inversion are unbiased with

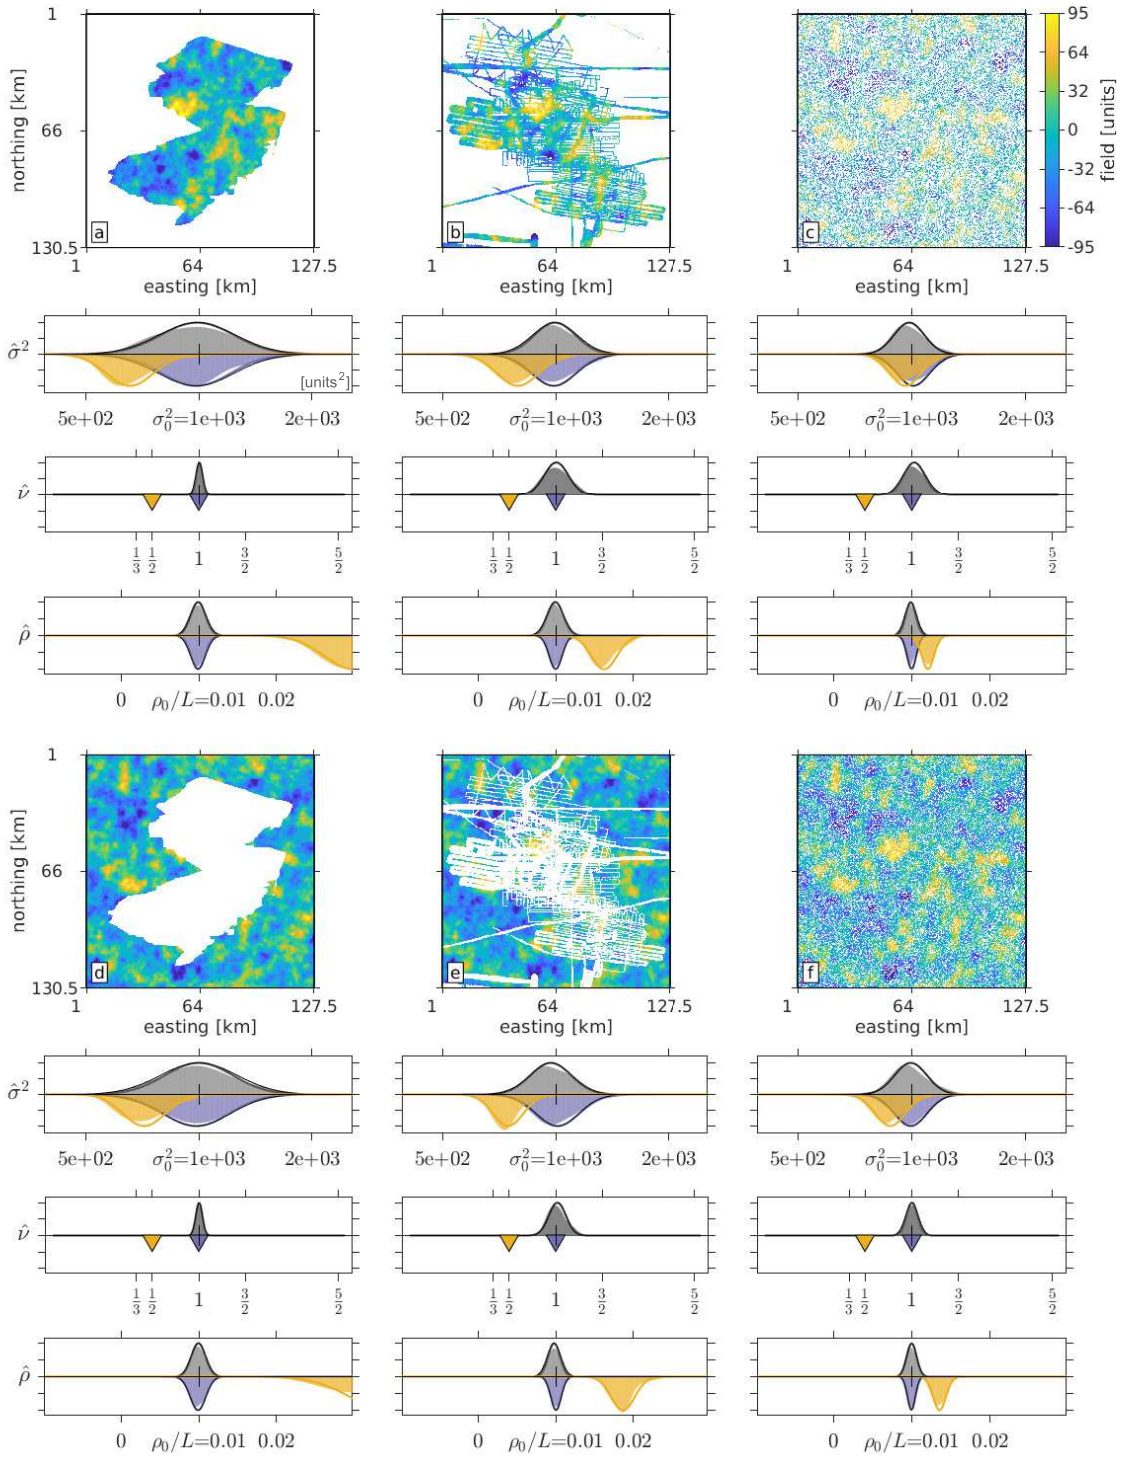


Figure 12. Maximum-likelihood estimation statistics for 100 spatial-covariance circulant embeddings for irregularly sampled fields (*top*: 33.6 per cent observed, *bottom*: complementary 66.4 per cent observed) on a 261×255 grid with spacings $\Delta y = \Delta x = 0.5$ km for true values (black vertical ticks) $\sigma_0^2 = 1000$ units^2, $\nu_0 = 1.0$ and $\rho_0 = 20$ km. Single realizations for each sampling pattern are shown, (a, d) irregular boundary, (b, e) structured and (c, f) random deletions. Parameter estimates for the general Matérn model are displayed in the top, up-right smoothed ensemble distributions (kdes; grey) with curves indicating the normal distribution fit for the sample mean and the ensemble observed (grey) and analytically calculated (black) standard deviations. Parameter estimates for two-parameter special case covariance models where the smoothness ν is correctly fixed to the true value (purple; 'Whittle', $\nu = 1$) and where it is incorrectly fixed (light yellow; 'Exponential', $\nu = 1/2$) are shown by the downward, mirrored kdes with normal distribution fits for the sample mean and ensemble observed (respective colours) and analytically calculated (black; only for the correctly fixed model) standard deviations in estimated variance σ^2 and range ρ. The fixed value associated with the special case ν is shown by the inverted colour-coded triangles. Tick labels in estimated smoothness ν represent additional popular special cases of the Matérn class mentioned in Table A1: the von Kármán case for $\nu = 1/3$, exponential $\nu = 1/2$, second-order autoregressive $\nu = 3/2$, Whittle $\nu = 1$ and third-order autoregressive $\nu = 5/2$. Estimated range ρ is normalized relative to the diagonal length of the grid, $L = [(N_y \Delta y)^2 + (N_x \Delta x)^2]^{1/2}$.

correct model specification, as they were for the three-parameter inversion. For σ^2, differences in parameter uncertainty are negligible between the three-parameter and correctly fixed, two-parameter model. For ρ, we observe a slight decrease in the relative standard deviation between the three-parameter versus the two-parameter model in the cases of the New Jersey (9.27 versus 8.60 per cent), Atlantic (9.96 versus 8.05 per cent) and uniform (6.34 versus 5.77 per cent) sampling patterns.

The complementary sampling scenarios in the bottom panel of Figs 12(d)–(f) substantiate all these behaviours, but we see the effects of having more data available. For example, comparing the right-side-up histograms for the three-parameter inversions from the top panel with the bottom panel (33 per cent grid observations versus 66 per cent), the parameter standard deviations obtained from the ensemble are similar or subtly reduced. The standard deviations of the variance σ^2 for the New Jersey (16.07 versus 14.91 per cent), the Atlantic (8.91 versus 8.88 per cent) and the uniform (6.67 versus 7.17 per cent) sampling scenarios remain relatively constant with increased data observations. For ν, in the case of the New Jersey (3.28 versus 3.46 per cent), Atlantic (13.03 versus 9.29 per cent) and uniform (11.10 versus 5.79 per cent) sampling, the empirical standard deviations remain similar in the bounded case and decrease in the distributed sampling cases for increased data size. For ρ, in the case of the New Jersey (9.27 versus 8.90 per cent), Atlantic (9.96 versus 7.52 per cent) and uniform (6.34 versus 6.17 per cent) sampling, as expected, the empirical standard deviations decrease with increased sample size. These observed changes in the spread of the ensemble between the three sets of models for increased data size are small. However, inspecting the upside-down histograms where the smoothness parameter ν has been fixed incorrectly, bias grows in the presence of more data. The relative bias in σ^2 (33 versus 66 per cent versus the truth) for the Atlantic (833 versus 790 versus 1000 units^2) and uniform (972 versus 906 versus 1000 units^2) sampling follows this behaviour, while in the New Jersey (696 versus 759 versus 1000 units^2) case, which has the greatest bias across the sampling scenarios, we see an improvement with sample size. The smoothness ν is again intentionally wrong. For ρ, the case of New Jersey (58.1 versus 68.9 versus 20.0 km), the Atlantic (32.5 versus 37.3 versus 20.0 km) and the uniform (24.1 versus 27.1 versus 20.0 km) sampling all show larger bias with sample size. We will further discuss the imprint of the sampling window on estimator behaviour in Section 9.

We recommend using the general three-parameter Matérn covariance in modelling (M.L. Stein 1999). There are no improvements to be made in terms of bias reduction when selecting the correct simplified form over the general, and parameter uncertainty decreases only subtly—for the range parameter ρ. In practice, the true value of ν can rarely, if ever, be assumed to be known. If one must, a preliminary estimation of the three parameters can be used to inform which two-parameter special case is most appropriate. Spatial sampling patterns should be considered in fixing the smoothness, as they tightly connect to estimator bias. Bias bred by an errant choice is reduced when samples are randomly dispersed and amplified when contiguous and dense.

As to the effects of having (for reasons pragmatic) or needing (for adjacent processes) a fixed (e.g. geographic) boundary that separates domains with potentially different Matérn parameters, additional considerations come into play, which we explore next.

6 SEPARABILITY OF SPATIALLY MERGED PROCESSES

We explore a series of numerical experiments to demonstrate the ramifications of analysing multiple Matérn processes that are spatially merged along an irregular boundary where we are either ignorant or aware of that boundary and analyse the observation as a single or multiple fields, respectively. In the case of the former, the random field departs from our model assumption of local stationarity.

Fig. 1 showed a first such example. In it, we showed 15-arcseconds resolution seafloor bathymetry comprised of direct soundings acquired by single and multibeam sonars along ship sampling tracks filling 38 per cent of the bounding grid with the complementary 62 per cent predicted from bathymetry data augmented by indirectly obtained sea surface satellite altimetry. We fit and removed a plane over the entire domain, split the data into its non-overlapping components, demeaned them and estimated the maximum-likelihood Matérn covariance parameters from the separate and merged data sets. The synthetic realizations simulated on the partial and complete grids were shown with the removed trends added back in. The parameter estimates for (σ^2, ν, ρ) are (38871, 1.46, 2.841) for the merged data set ('`mer`'), which is effectively non-stationary, is not a simple average of the estimates made on the constituent components (45 452, 0.96, 4.322) for '`dir`', and (47 764, 1.45, 4.881) for '`ind`'. The apparently concentrated uncertainty of the merged estimate ($\pm$6657, $\pm$0.02, $\pm$0.208) versus ($\pm$10 852, $\pm$0.18, $\pm$0.895) and ($\pm$7857, $\pm$0.24, $\pm$0.907), is illustrative of how estimable such a process is within its observation grid rather than how well it corresponds to the data, which, other than delivering the parameters, have no part in shaping the uncertainty.

For the case of bounded domains with an interior and exterior portion, the experiments in this section address three queries: to ascertain (1) how the debiased Whittle likelihood (eq. 7) accounts for spatial sampling patterns and produces an unbiased estimator $\langle\hat{\boldsymbol{\theta}}\rangle = \boldsymbol{\theta}_0$ whose covariance can be exactly predicted via eq. (12); (2) how the appearance of multiple processes within a single window of observation affects the estimator when their distinctions are not indicated by the taper $w(\mathbf{x})$ employed in the analysis; and (3) how the acknowledgment of the spatial extent of each of the merged processes by the sampling taper allows for the robust estimation of the processes individually.

In Fig. 13, we consider a rectangular grid with an interior region (the '*mask*') that comprises 40.17 per cent of the total sample size, and an exterior (the '*anti-mask*') possessing the complementary 59.83 per cent. The interior is separated from the exterior by an irregular, sock-like, closed boundary. Fig. 13 shows a suite of 10 experiments for two processes that we organized by permuting the parametric process (two sets of Matérn parameters, *I* and *II*), the extent of the process within the rectangular observation window and

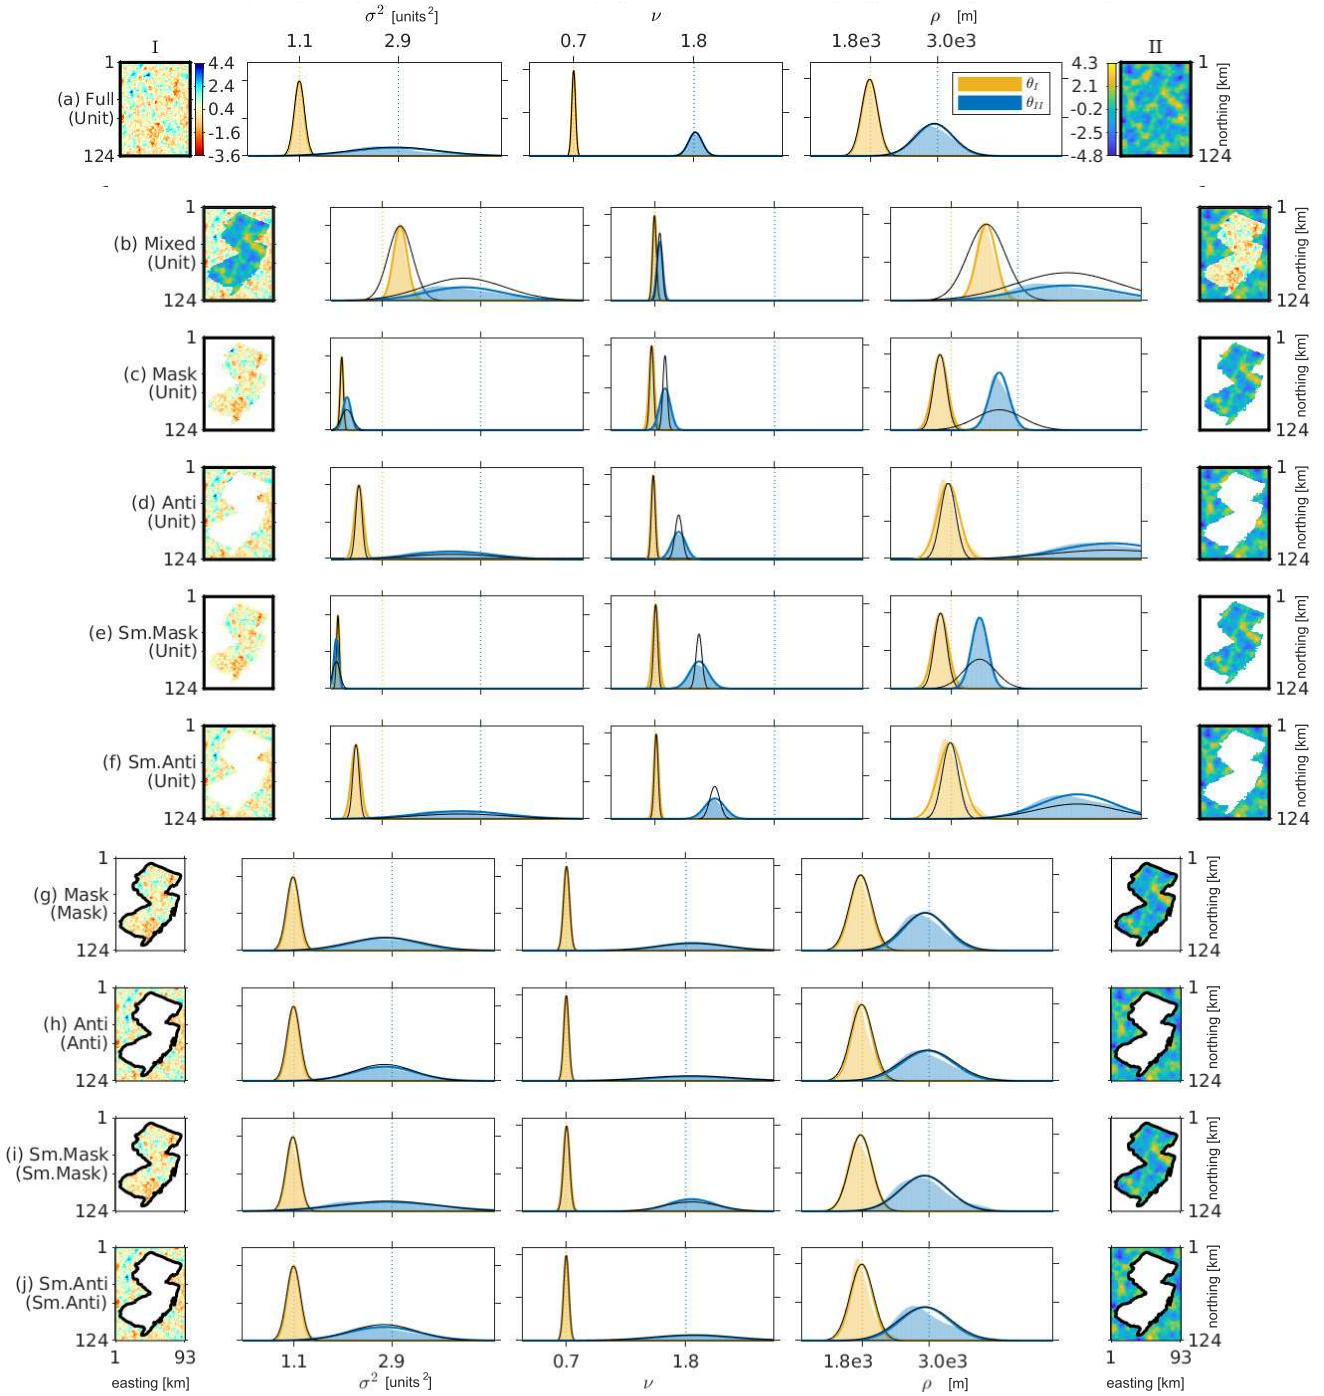


Figure 13. Recovery of spatially merged fields with different Matérn covariance structures: the role of the analysis window in the debiased Whittle likelihood. Synthetic experimental results conducted on a common rectangular grid, $N_y = 124$, $N_x = 93$ and $\Delta y = \Delta x = 1$ km, with different processes on an exterior (59.83 per cent of the total area) and interior (40.17 per cent) subdomain delineated by an irregular boundary. The far left and right columns display realizations for Matérn models $\boldsymbol{\theta}_I = [1.15\,\text{units}^2\ 0.65\ 1.8\text{km}]$ (light, yellow-toned colourmap) and $\boldsymbol{\theta}_{II} = [2.90\,\text{units}^2\ 1.75\ 3\text{km}]$ (deep, blue-toned colourmap), respectively, for each of the 10 experimental scenarios. In each row, the observation window and analysis taper are varied, which we denote along the rows with labelling and the display for a single realization of fields *I* (left) and *II* (right). Order of labelling: experiment (a–j), observation window of the field (*Full*, *Mixed* at the boundary, *Mask* of interior, *Anti*-mask of exterior, *Sm(ooth) Mask*, *Sm(ooth) Anti*-mask), taper used in analysis for the blurred likelihood (the $w(\mathbf{x})$ in eq. 7; *Unit*, *Mask*, *Anti*-mask, *Sm. Mask*, *Sm. Anti*-mask). Estimate ensembles are displayed as kernel density estimates for $\hat{\theta}_I$ (left, yellow) and $\hat{\theta}_{II}$ (right, blue), overlain with normal fits in matching colours, and with, in black, analytically predicted parameter standard deviations (eq. 12), evaluated at the sample means, for the analysis taper used in the experiment, which is plotted by the thick black line on the spatial renderings. The coloured empirical and black analytically predicted distributions are unbiased and agree when the correct taper is used in analysis, that is, for experiments (a) and (g–j), offset horizontally for distinction.

whether the sampling window applied in the analysis $w(\mathbf{x})$, smoothed or not, has acknowledged the region within it. The set covers the following 10 scenarios: (a) a single process observed and analysed on a rectangular grid; (b) two processes merged along an irregular interior boundary but analysed on the rectangular grid; a single process in the interior without (c) and with (e) smoothing along the boundary analysed on the rectangular grid; a single process in the exterior without (d) and with (f) smoothing along the boundary analysed on the rectangular grid. In all those cases (a–f), the analysis window $w(\mathbf{x})$ is a rectangle, rendered as a thick black line. The next four rows show a single process in the interior without (g) and with (i) smoothing along the boundary analysed in the interior only and in the exterior without (h) and with (j) smoothing along the boundary analysed in the exterior only. In those cases (g–j), the analysis window $w(\mathbf{x})$ is the boundary between the processes, rendered as a thick black line. If applied, the smoothing process used along the boundary is a cosine squared taper applied along each of the segments of the boundary, iteratively convolved to smooth the corners.

We simulate 200 realizations of all permutations, estimate the corresponding Matérn parameters and report on the empirical and analytical distribution (calculated from eq. 12 evaluated at the sample mean) of the ensemble in the inner panels of Fig. 13. In Table A4, we summarize, for each experiment and for both processes the taper used, the sample means and the standard deviations (relative, analytical and empirical) of the variance, smoothness and range of the process. In addition, Table A4 lists the correlations between the Matérn parameter estimates, in per cent, first as predicted and then as empirically observed.

In Fig. 13 and Table A4, we begin with the base case (a) of two rectangular, stationary sampled random fields. The first field, I, possesses the Matérn parametrization $\boldsymbol{\theta}_I = [1.15\,\text{units}^2 \quad 0.65 \quad 1800\,\text{km}]$, and the second, II, with $\boldsymbol{\theta}_{II} = [2.90\,\text{units}^2 \quad 1.75 \quad 3000\,\text{km}]$, is larger in magnitude for all three parameters. For both field I and II, we simulate realizations via the circulant embedding of the Matérn spatial covariance for $\boldsymbol{\theta}_I$ and $\boldsymbol{\theta}_{II}$, both on a common grid size of 124 km × 93 km with even 1 km spacing in each direction.

Case (a) in the top left of Fig. 13 displays a single realization for process I and in the top right for process II. For each simulation, we apply our maximum-likelihood estimation strategy to determine estimators $\hat{\boldsymbol{\theta}}_I$ and $\hat{\boldsymbol{\theta}}_{II}$, taking the sampling window for the analysis, $w(\mathbf{x})$ in eqs (7) and (8) to be a unit taper (i.e. fully observed). The ensembles of estimates for the three Matérn parameters are displayed in the centre columns as smooth distributions, with yellow kdes distinguishing the ensemble of $\hat{\boldsymbol{\theta}}_I$, yellow overlying curves marking the normal distribution fit to the $\hat{\boldsymbol{\theta}}_I$ sample mean and sample standard deviation, and black overlying curves showing the normal distribution fit to the sample mean and analytically calculated parameter standard deviation (eq. 12) evaluated at the ensemble sample mean for the unit taper sampling window on the rectangular grid. Similarly, the blue kdes, matching blue curves, and overlying black curves with the shared centres correspond to the distribution of the $\hat{\boldsymbol{\theta}}_{II}$ and its empirical and analytically calculated standard deviations. The true parameter values for $\boldsymbol{\theta}_I$ and $\boldsymbol{\theta}_{II}$ are labelled on the horizontal axis, displayed as dotted vertical lines in yellow and blue, respectively, with $\boldsymbol{\theta}_{II}$ falling to the right as its true parameter values were selected to always be greater than $\boldsymbol{\theta}_I$. The experimental outcomes of (a) include two ensembles centred on their truths with predictable parameter uncertainties that correspond to the observed ensembles. The spread of the ensembles scales with the magnitude of the parameter value: the larger the $\hat{\theta}_i$, the larger the $\text{var}\{\hat{\theta}_i\}$. We retain the setup of the fields and organization of results throughout the experiments that follow.

Case (b) considers the two processes occurring within the same unit window, either with process II (coded by the deeper colour, blue-toned colourmap within the realizations shown in Fig. 13) within the interior of process I (shown in the lighter, yellow-toned colourmap of Fig. 13) in the left hand-side 'I' column, or the opposite (process I within the interior of process II) in the right hand-side 'II' column. In this example, we are (knowingly incorrectly) attempting to estimate a stationary Matérn process from a non-stationary field. The estimate ensembles and normal approximations reveal estimates that are intermediate to I and II in comparison to case (a). Additionally, our analytical prediction of uncertainty evaluated at the ensemble sample mean for a unit taper window (black curve) reveals a larger anticipated spread than we observe from the ensemble.

Cases (c) through (f) consider a single process confined to the same bounded area (denoted *mask* in Fig. 13 and Table A4) or its complement (*antimask*) observed and analysed with the *unit* taper, with or without smoothing (*Sm.*) applied at the edges of the field. These experiments are deemed unsuccessful due to the bias imparted on the ensemble of estimates and the discrepancy between the observed and predicted parameter uncertainty.

Cases (g) through (j) are formatted as (c–f), however, they are analysed according to the proper observation window, either tapered by the *mask* or *antimask* with possible smoothing at the edges. In all of these experimental cases, the estimates are unbiased and the empirical and analytical parameter uncertainties are well matched, a testament to the debiased Whittle likelihood of eq. (7). Smoothing the outer edges of the field does not appear to improve the ensemble mean in these cases, though we do see a subtle increase in parameter uncertainty when smoothing is applied.

The outcome of our suite of experiments reveals what one might expect for our first stated query (1): Estimation of single processes whose sampling pattern is taken as the analysis tapering window ($w(\mathbf{x})$ in eq. 7; experiments a, g–j) results in an unbiased estimator with estimator uncertainty that is inversely proportional to the degrees of freedom. Deviations in $w(\mathbf{x})$ introduce increased bias and uncertainty, whose magnitudes depend on the field observation window and the magnitude of the process parameters. The 68.27 per cent confidence intervals for the variance σ^2 and smoothness ν do not overlap for I and II in the experiments conducted on the homogeneous subdomains, while the distributions of the range ρ parameters display overlap due in part to the increased parameter variance from the reduced degrees of freedom in the *mask* and *antimask* tapered fields. In the experiments for *(anti-)masked* single processes analysed with a *unit* taper (c–f), overlap is only a problem for σ^2 in (c) and (e) for the masked field with and without smoothing applied

to the boundary of the mask which is subject to large, negative bias. Bias is observed for all parameters in (c–f), with a greater effect displayed by the larger-magnitude parameter model II. Smoothing along the boundaries yields similar ensemble results to analysing fields with a hard indicator boundary.

The resolution to our query (2) is highlighted by the outcome of the mixed-process experiment (b), which is an example of multiple, spatially separable processes being observed and analysed within the same window. This is a realistic scenario for geophysical sampled random fields where multiple, spatially confined, physical processes have imprinted neighbouring regions. The results of this experiment for process II within the interior of process I in light yellow (left column of Fig. 13) and the inverse in blue (right), show ensemble estimates that are intermediate to model I and II in σ^2, estimates that favour the smaller ν of I and estimates that favour the larger value ρ of II, with overlap in the confidence intervals for both scenarios. In terms of model identifiability, A.P. Guillaumin *et al.* (2017, 2022) discuss how the interactions between the geometry of the spectral density and the geometry of the sampling process can be understood via what they term the *significant correlation contribution*.

What we wish to emphasize in our last query (3) for merged processes is that unbiased estimates with well predicted parameter covariances can be determined from a single data set if it is possible to spatially delineate the processes within the domain and analyse them separately for an observation window defined by the detected boundary (as in cases g through j). This strategy will allow practitioners to study the covariance structure of spatial geophysical data sets that develop under the action of multiple, spatially jigsawed phenomena, if their heterogeneity can be demarcated first. We utilize this strategy of demarcation for our real-data applications, where we bring in outside (e.g. geological) information. In the absence of such guiding information, our framework and analysis will help evaluate the success of alternative data-driven (e.g. segmentation, classification and clustering) approaches.

7 GEOSCIENCE APPLICATIONS: REAL DATA

Real, sampled spatial data are unlikely to be inherently and unquestionably Gaussian, stationary and isotropic. However, from our synthetic studies, we are heedful of the bias that observations of multiple merged, non-stationary, processes may impart on model estimates (Section 6) and of the bias and increased estimator uncertainty (namely of the smoothness parameter, ν) for irregularly, incompletely sampled fields (in comparison to those on a regular grid; Section 5). In our previous work (section 5 of F.J. Simons *et al.* 2026), we encountered the effect of departures from these assumptions on the residuals for covariance estimates of real-data examples sampled on a complete regular grid.

We apply our maximum-likelihood estimation strategy to four geophysical data examples that exhibit irregularity in their sampling to characterize their covariance structures with quantified uncertainties and a rigorous assessment of the model residuals. As in the synthetic example in Fig. 9 and the real-data examples of F.J. Simons *et al.* (2026), we estimate the covariance structure of a geophysical field and assess the model by a qualitative assessment of whether a synthetic realization simulated from the estimate corresponds to the behaviour of the original data and whether the model residuals $X_{\theta}(\mathbf{k})$ (eq. 10) retain any structures from the periodogram $|H(\mathbf{k})|^2$ (eq. 7) that have been left unresolved by the modelled blurred spectral density $\bar{S}_{\theta}(\mathbf{k})$ (eq. 8). Quantitatively, we evaluate the goodness-of-fit of our model and how well it scales with the theoretical distribution χ^2_2 (eq. 10) by way of a Q–Q plot and the test statistic s^2_X (eq. 11). Our analysis results for these four data applications are displayed in Figs 14–16 and Table 1. For the partially sampled data sets we consider here, we create an observation window that masks either indirect sources of data, secondary processes, or observed phenomena that we seek to omit, thus performing our analysis on the data type or feature of interest only, with exact blurring of the likelihood and covariance calculation. We pre-process the data by detrending up to a second-order polynomial fit when a locally varying mean is evident in the original data. A 5 per cent cosine window taper is applied to the encompassing rectangular perimeter of the observation window, only smoothing the outermost samples to reduce edge effects that appear in the spectral residuals and which obfuscate the remaining signal we wish to interpret.

In the first example (Fig. 14), we demonstrate a procedure for assessing spatial processes where an irregular boundary can be drawn to delineate features with geomorphic traits or histories that are distinct from their surroundings, for which we create a mask for the observation window of the process under analysis. Our data source is SLDEM2015, a lunar digital elevation model (DEM) developed by M.K. Barker *et al.* (2016) from observations gathered by the Lunar Orbiter Laser Altimeter onboard NASA's Lunar Reconnaissance Orbiter spacecraft that were corrected by co-registered stereo-derived DEMs from the SELENE Terrain Camera. We isolate a rectangular patch of lunar topography encompassing Taurus–Littrow Valley (centred at approximately 20°N, 31°E). One application of parametric covariance estimation is as a metric for surface roughness. For this data set, the historic landing site of Apollo 17, our analysis might provide a useful and interpretable surface roughness metric for comparative site planning in autonomous navigation (e.g. for landers or rovers; S. Thrun *et al.* 2006; M. Häselich *et al.* 2013; E. Pardo-Igúzquiza & P. Dowd 2022). To exclude deterministic geologic features from their homogeneous background, we create a mask that fits catalogued lunar impact craters with diameters greater than 1 km (S.J. Robbins 2018) as circles. The non-cratered lunar surface comprises 62.8 per cent of the bounding rectangle. The grid spacing is 0.24 km. Our analysis of the Matérn model for the non-cratered topography reveals a near χ^2_2 distribution of the spectral residuals $X_{\theta}(\mathbf{k})$, a strong fit to eq. (11), though not quite sufficient to fail to reject the hypothesis test. We see no structure within the spectral map of the residuals. In the visual comparison between the DEM and the simulation, subtle along-latitude streaking appears in the DEM, most likely arising from a data artefact, which our model for random processes is *not* meant to replicate.

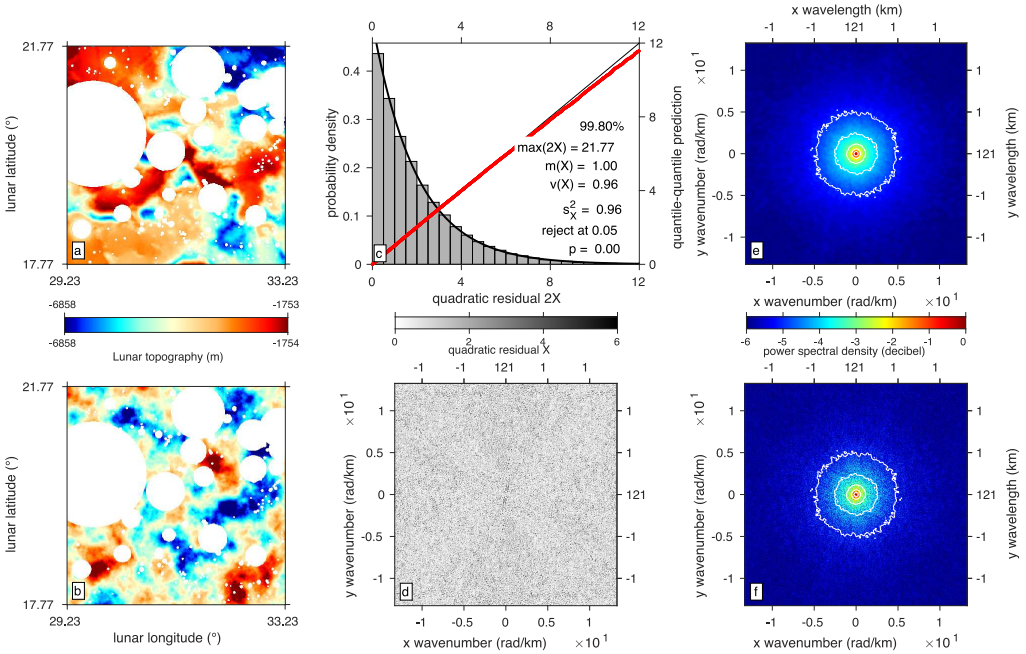


Figure 14. Maximum-likelihood analysis for the Matérn covariance structure of lunar topography (M.K. Barker *et al.* 2016) without craters (S.J. Robbins 2018) near Taurus–Littrow Valley. Shown are the observed data (a) and a synthetic generated through the circulant embedding of the Matérn spatial covariance parametrized by our estimate (b; Table 1), a histogram and a quantile–quantile plot of the quadratic residual $2X_{\boldsymbol{\theta}}$ (c), with $X_{\boldsymbol{\theta}}(\mathbf{k})$ being rendered and inspected for patterns in reciprocal-space (d; Section 2.4). Also shown are the expected (e) and observed (f) periodograms, $\mathcal{S}_{\boldsymbol{\theta}}(\mathbf{k})$ and $|H(\mathbf{k})|^2$, respectively, with contour lines at each integer decibel of the former overlain on the latter. This example demonstrates the analysis of a spatially contiguous process (a) that is geomorphically distinct from the union of its bounded impact craters. The synthetic simulation (b) is visually comparable to the random features within the observed data. The model for the non-cratered lunar surface fits very well (c), though not quite passing the rigorous statistical test (eq. 11), with no evident structure within the model residuals (d). The modelled spectral density (e) fits the data periodogram (f) very well.

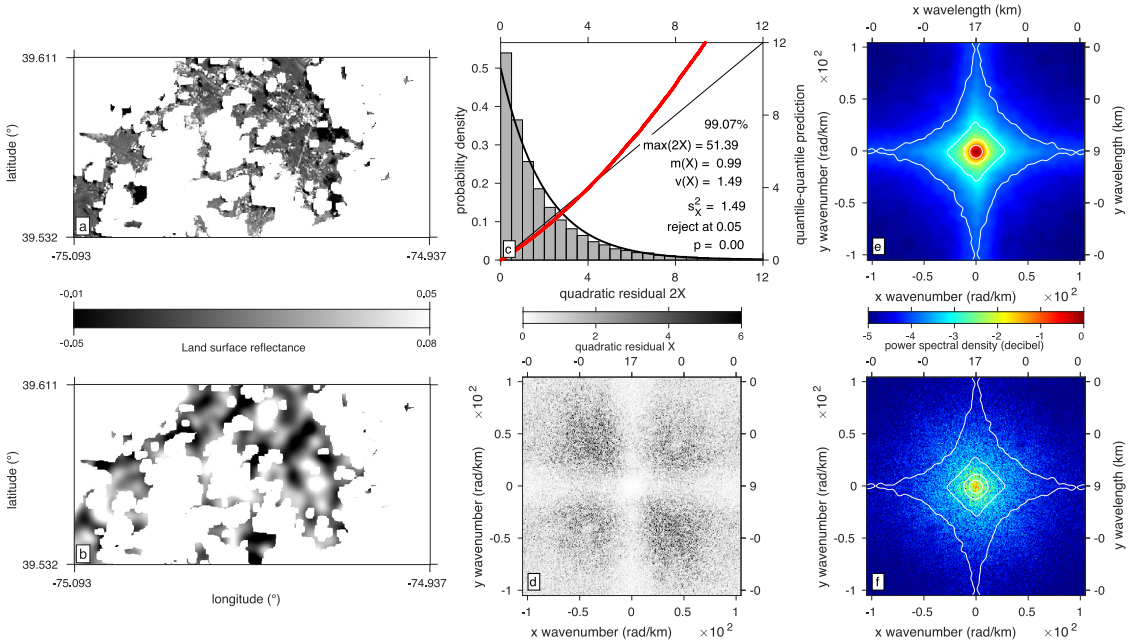


Figure 15. Maximum-likelihood analysis for the Matérn covariance structure of LANDSAT 8 Collection 2, Level-2 surface reflectance data for land observations. Red, green and blue spectral bands converted to a greyscale composite with pixels labelled as 'land' in the accompanying quality analysis file selected (a), which comprise 29.71 per cent of the encompassing rectangular grid; (diffuse) clouds, cloud shadows and minor water bodies make up the remaining masked pixels. The estimated covariance structure has a large smoothness (Table 1). The synthetic (b) shares some resemblance to the data. The $2X\sigma^2(\mathbf{k})$ model residuals (c) for the land surface reflectance data follow the χ^2_2 distribution well through the bulk of the distribution, with higher than expected small residuals and a long tail, as shown in the departures of the Q–Q plot. The residual plot in reciprocal-space (d) possesses a windowpane-like structure that our model is unable to capture. The modelled spectral density (e) captures the general structure of the noisy data periodogram (f). Layout and annotation are as in the caption of Fig. 14.

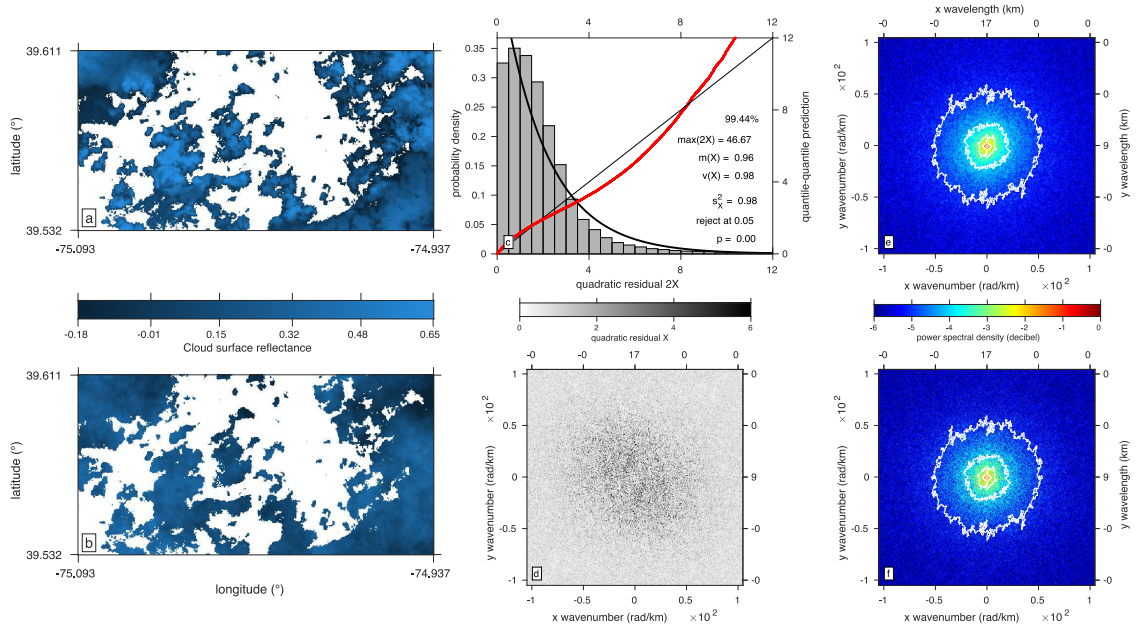


Figure 16. Maximum-likelihood analysis for the Matérn covariance structure of LANDSAT 8 Collection 2, Level-2 surface reflectance data for cloud observations (a), complementary to the data set in Fig. 15, comprising 54.23 per cent of the encompassing grid. The spatial synthetic (b) generated from the model is appropriate by eye. The model for the observed cloud data (Table 1) does not fit well in terms of the histogram of the residuals (c) as they do not conform to the theoretical χ^2_2 distribution of eq. (10) as would be anticipated for the ratio of uncorrelated, Gaussian random variables. The structure of the residuals in reciprocal-space (d) appears isotropic, albeit with higher than expected structure remaining for low wavenumbers. The modelled spectral density (e) fits the data periodogram (f). Layout and annotation are as in the caption of Fig. 14.

Table 1. Results from experiments with irregularly sampled geophysical data, sorted by degrees of freedom of the analysis for the tapered patch, $K = \sum w(\mathbf{x})$. The grid percentage observed is reported relative to the area of the encompassing regular grid as $K/(N_y N_x)$. The estimate for the variance σ^2 is quoted as a fraction of the sample variance of the observed data set, s^2. In the presence of significant correlation range, s^2 is small relative to σ^2. We present the estimate of the correlation range in units of field distance [km], see Fig. 3, and as a fraction of r, the enclosing grid diagonal, scaled by π. For all three parameters, the one-standard-deviation estimation uncertainty is listed to the nearest per cent of the parameter. The final three columns contain the correlation between estimates, in per cent.

Example	Fig.	K	Per cent obs	σ^2/s^2	±Per cent	ν	±Per cent	ρ	$\pi\rho/r$	±Per cent	$\{\sigma^2, \nu\}$	$\{\sigma^2, \rho\}$	$\{\nu, \rho\}$
Land surface reflectance	15	50 061	29.7	6.7290	12.3	30.53	533.7	0.18	0.0296	12.9	−29.5	45.5	−93.0
Cloud surface reflectance	16	91 371	54.2	2.1084	109.9	0.46	41.6	7.39	1.1969	190.3	−10.9	74.5	−72.6
Ocean bathymetry	17	137 251	38.0	2.3073	9.0	0.96	7.19	4.32	0.0110	7.9	21.1	49.0	−69.0
Lunar topography	14	165 255	62.8	1.0727	21.4	0.98	25.9	5.12	0.0936	38.8	1.9	32.7	−91.5

In our next two companion examples (Figs 15 and 16), we select a data set that exhibits missingness distributed throughout the observation window as pixel-labelled 'land' and 'cloud' features corresponding to surface reflectance data captured by LANDSAT 8 on 2025 July 20 provided as a Collection 2, Level-2 data product for the tile aligned on path 014, row 032 with 30 m resolution [Earth Resources Observation and Science (EROS) Center 2020]. From the red, green and blue surface reflectance bands, we convert the digital recorded values to physical measurements of surface reflectance ratios, masking pixels that exceed the specified valid ranges [Earth Resources Observation and Science (EROS) Center 2019] and converting the three bands into a greyscale composite for univariate analysis. We select a region within the tile between 39.53–39.61°N and 74.94–75.09°W, where 54 per cent of the pixels are labelled in the quality assessment band file as cloud coverage, 30 per cent as land and of the pixels we do not consider within the tile, 18 per cent are labelled as cloud shadows (possessing considerable overlap with other labels), 16 per cent dilated clouds and, negligibly, 0.02 per cent water bodies. We estimate the covariance structure for the land-tagged data (Fig. 15) and the cloud-tagged data (Fig. 16) as separate experiments. For the land surface reflectance data, the estimated smoothness is notably large, as are its parameter uncertainty bounds. Due to the (expected) strong, negative trade-off between the smoothness and range parameters in the Matérn model (see the correlations listed in Table 1) it is possible that, at this resolution, we are unable to estimate the covariance structure of this particular data set well as the true range parameter is smaller than the grid spacing in our observation window. We include this example to show how a 'too-smooth' physical data set may fall outside of our sampling grid constraints, which we will discuss in Section 8. The cloud data set fits its estimated Matérn model better, but not perfectly. We observe isotropic structure within the reciprocal-space plot of the residuals that is not captured by the model. The histogram of the residuals does not sufficiently conform to a χ^2_2 distribution.

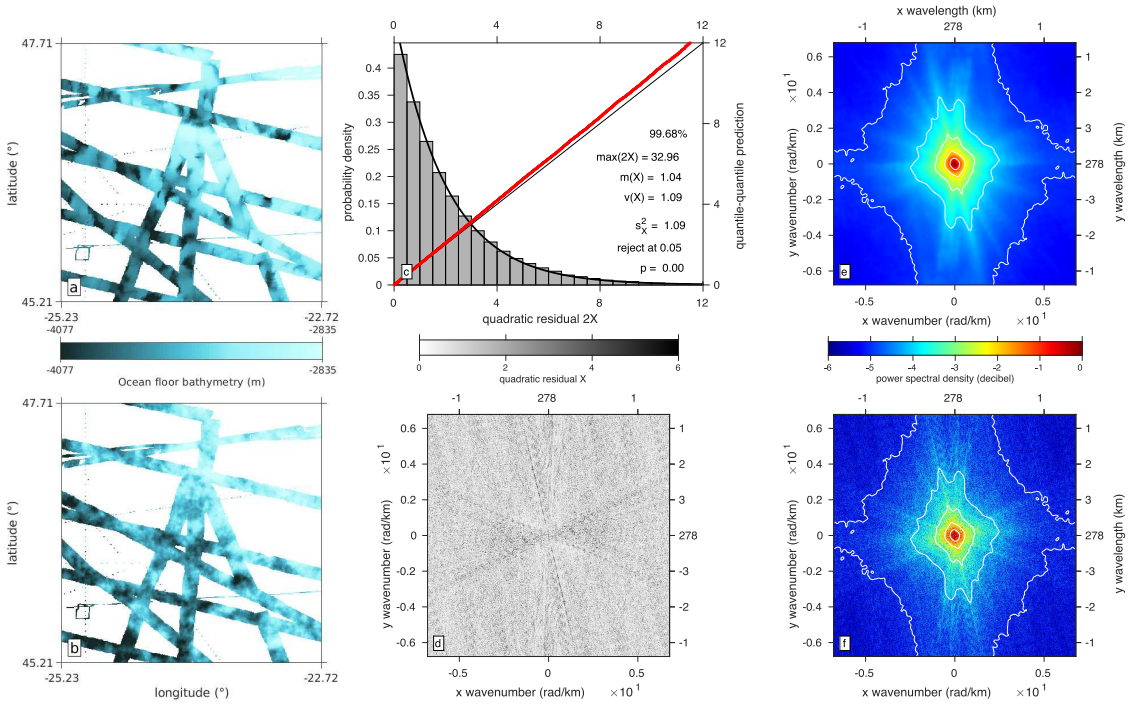


Figure 17. Maximum-likelihood analysis for the Matérn covariance structure of partially observed Atlantic seafloor bathymetry. This data set provides an example of structured geophysical data collected along sampling tracks by shipboard direct single-beam and multibeam measurements (a; GEBCO Bathymetric Compilation Group 2024). The synthetic (b) is very similar to the data. The spectral domain residual histogram and Q–Q plot (c) show the model to be a good fit, not quite passing the rigorous statistical test (eq. 11), with subtle, fine-scale radiating structure retained within the reciprocal-space residual plot (d), which we interpret as multidirectional geometric anisotropy within the data. The modelled spectral density (e) overlays the data periodogram (f) very well. Layout and annotation are as in the caption of Fig. 14.

This deviation of the data from their most likely estimated model and the residuals from their expectation could be indicative of non-Gaussianity.

Our final example (Fig. 17) is of the directly sampled, structured tracks with 0.46 km resolution of ocean floor bathymetry (GEBCO Bathymetric Compilation Group 2024) centred at 46.45°N, 23.97°W in the Atlantic from Fig. 1. The directly measured data are collected shipboard as predominantly multibeam (98.9 per cent) with few single beam (1.1 per cent) soundings that collectively make up 38.0 per cent of the enclosing rectangular window. As we stated in the bathymetric full patch example of F.J. Simons *et al.* (2026), we do not anticipate that ocean bathymetry in general will exhibit isotropic spatial behaviour due to the anisotropy of its plate tectonic formative processes. For this patch with only direct measurements considered, we do however find that the isotropic model fits very well. The residuals retain a minor, radial pattern that we take to represent multidirectional, geometric anisotropy, in contrast to the more diffuse directional anisotropy that we observed in fig. 16 of F.J. Simons *et al.* (2026).

8 ASYMPTOTICS OF SPATIAL SAMPLING

Our debiased Whittle maximum-likelihood estimation strategy provides an (asymptotically) unbiased estimator by incorporating the effects of finite data, possessing edges either hard or smooth in transition, regular or irregular in boundary and sampling patterns coarse or fine, discretized by spaces even or agape. The results from the experiments discussed in Section 5 indicated that sample size and spatial geometry bleed into estimator covariance, as seen both from empirical observation and through our analytical prediction (eq. 12). In the case of real data, as in Section 7, that are limited to a single spatial realization, understanding the (well predicted) distance at which a lone estimate may fall from the truth, and how increasing sample size according to particular modes of data collection can reduce this distance, is valuable. Here, we design synthetic experiments that demonstrate the asymptotic behaviour of the estimator bias and its covariance when sample size increases through three mechanisms: growing domain, infill (fixed domain) and reduced missingness.

In Fig. 18, we illustrate in a series of cartoons the allocation of samples in growing and fixed domains for fully observed, regular grids. Robust estimation of the underlying Matérn model from a realization requires that the sampling grid is crafted or observed such that it is sufficiently large in terms of how the correlation length $\pi\rho$ relates to the physical grid size and its sample spacing (Fig. 18, *left*). Growing-domain (H. Zhang 2004) asymptotics (Fig. 18, *centre*) are concerned with how increasing sample sizes by increasing the physical grid size with sample spacing (Δy, Δx) held constant affects the estimator. Infill (N. Cressie 1993; H. Zhang & D.L. Zimmerman 2005), or fixed-domain (M.L. Stein 1995, 1999), asymptotics (Fig. 18, *right*), on the other hand, explore increases in sample sizes

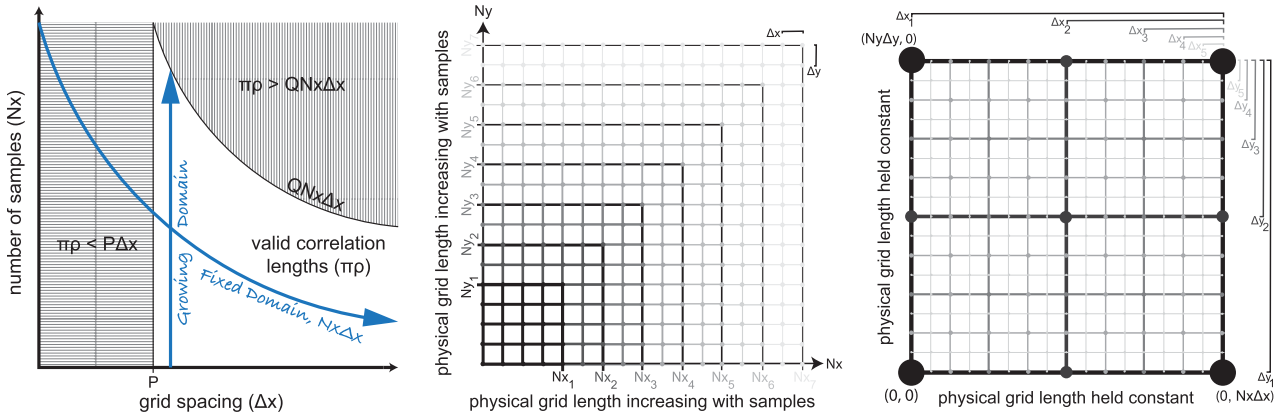


Figure 18. Cartoons illustrating sampling grid constraints (*left*, simplified to 1-D) and increasing sample size within a growing-domain (*centre*) and a fixed-domain (*right*) scheme. A Matérn model is estimable for a sufficiently large grid, where multiple observations of the correlation length ($\pi\rho$) may be made and the spacing is small in relation. Growing-domain asymptotics assume a fixed grid spacing (Δy, Δx), where samples are increased by extending the physical grid length ($N_y\Delta y \times N_x\Delta x$). In contrast, the fixed-domain or infill approach increases samples through densifying sampling on a fixed physical grid length by reducing the spacing between existing samples on a fixed physical grid size ($N_y\Delta y \times N_x\Delta x$).

through densifying the sampling grid by decreasing the spacing between samples such that the physical grid size ($N_y\Delta y \times N_x\Delta x$) is held constant. Mixed-domain asymptotics, as the name suggests, can capture the interplay between these end-member studies through sensitivity to local and large-scale features of estimated parametric covariance structures (T. Chu 2023). Another mechanism for increasing the sample size of a field is by replacing missing samples with new observations, which is effectively a variation on fixed-domain asymptotics for an irregular grid with an irregular filling strategy (S. Deb *et al.* 2017). In the growing-domain and infill asymptotic studies, we will consider observation windows of simulated fields that are full, regularly sampled rectangular grids and windows with 66.7 per cent of the grid observed, such that 33.3 per cent of the samples are missing at random. For the missingness-reduction study, we will explore broad proportions of random deletions and our results will emphasize the dependence of the estimator's behaviour on the true range parameter ρ_0 of the simulated model.

In each of the experiments that follow, the model is constant and only a single feature of the grid is adjusted in each trial. We therefore attribute any sensitivity in parameter bias or covariance between trials to the design of the sampling grids. The behaviour of $\bar{\mathcal{S}}_{\boldsymbol{\theta}}(\mathbf{k})$ and its partial derivatives $\partial\bar{\mathcal{S}}_{\boldsymbol{\theta}}(\mathbf{k})/\partial\theta$ and their interactions with the spectral window $|w(\mathbf{k})|^2$ as dependent on data size N_y, N_x, sampling size Δy, Δx, completeness and arrangement of the sampling, collectively control whether the chi-squaredness of the ordinates of the periodogram and the asymptotic normality of the estimators can be reached (S.C. Olhede *et al.* 2004; A.M. Sykulski *et al.* 2019; A.P. Guillaumin *et al.* 2022). Rather than discussing these effects theoretically, we opt for a suite of practical experiments conducted in regimes that geoscientists will recognize as relevant to their experience, if only to guard against the unrealistic optimism that large-sample behaviour will be un-intentionally met or realized. Section 9 will discuss means for arriving at desired regimes through the design of experiments.

8.1 Growing-domain asymptotics

In our growing-domain experiments, we perform four simulation trials on evenly spaced, square grids for which we generate 100 realizations parametrized by the Matérn model $\boldsymbol{\theta}_0 = [3\,\text{unit}^2\ 0.75\ 4\,\text{m}]$. Each trial corresponds to a sample size that increases by a power of two from 32×32 to 256×256, which, for a grid with 1 m spacing between samples, equates to a physical grid length of approximately $5\pi\rho$ to $20\pi\rho$. We perform this numerical experiment on the full, regularly sampled grid, and on a grid 66.7 per cent observed, with the complementary 33.3 per cent of samples randomly deleted over the grid.

In Fig. 19, we present the estimate ensembles yielded by the simulation trials by comparing their sample mean relative to the truth $\boldsymbol{\theta}_0$, and their sample parameter (co)variances to their analytical prediction (eq. 12). For both fully and partially observed grids, increasing sample size by increasing the size of the physical grid for a fixed spacing Δy, Δx reduces the relative estimator bias (Fig. 19, *first row*), parameter variance (Fig. 19, *second row*), parameter covariance (Fig. 19, *third row*) and the root-mean-squared error (Fig. 19, *fourth row*),

$$\text{rmse}(\hat{\theta}) = \sqrt{\text{var}(\hat{\theta}) + \text{bias}^2(\hat{\theta})} \tag{25}$$

for all parameter and cross-parameter terms. The rate of decay of each scenario is similar between the full and partial grid experiments.

The common approximation $\text{cov}(\hat{\boldsymbol{\theta}}) \approx \mathcal{F}^{-1}(\boldsymbol{\theta}_0)$ requires comparison to the full expression of eq. (12): How well does the inverse Fisher information matrix approximate the full theoretical calculation of parameter covariance for a Matérn model in terms of increasing window sizes? Inspired by the formulation of the Cramér–Rao (C.R. Rao 1945; H. Cramér 1946) lower bound efficiency, we

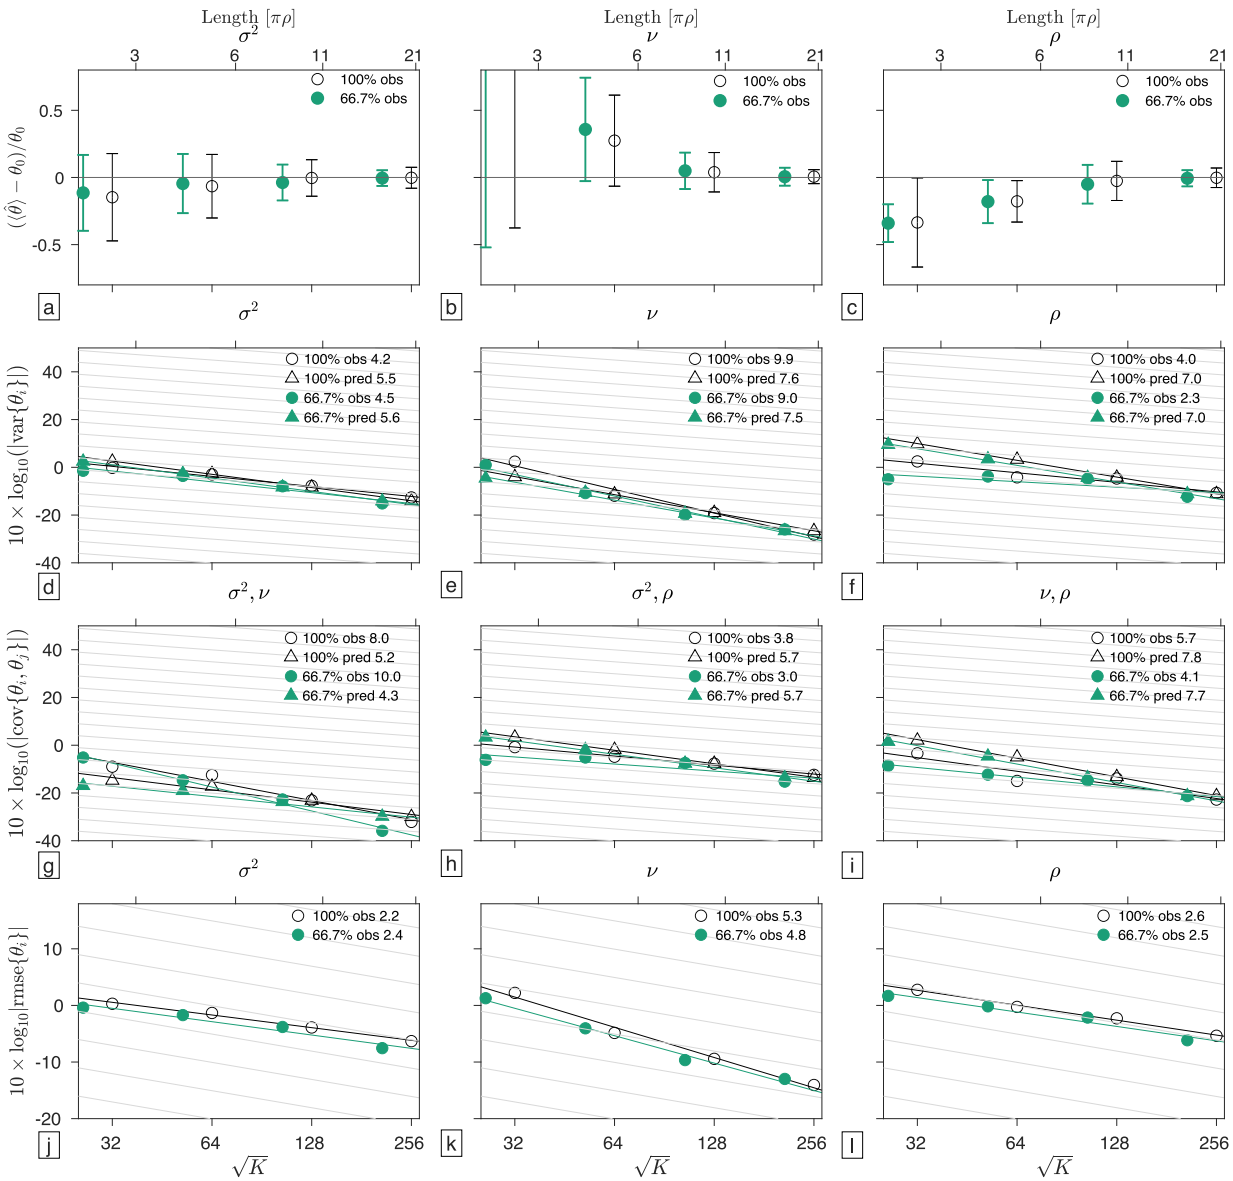


Figure 19. Growing-domain (i.e. constant $\Delta y = \Delta x$) asymptotic behaviour of observed and predicted estimator covariance for increasing sample size through physical grid length growth of a fully (*black*) and partially (66.7 per cent; *green*) observed random field. One hundred realizations are simulated for the Matérn model $\boldsymbol{\theta}_0 = [3\ \text{unit}^2\ 0.75\ 4\ \text{m}]$ with 1 m spacing in both spatial directions of each trial grid. Relative ensemble estimator bias $(\langle\hat{\theta}\rangle - \theta_0)/\theta_0$ (a, σ^2; b, ν; c, ρ) with empirical standard deviation error bars plotted against the square root of the sample size, $\sqrt{K}$. Estimator variance $\text{var}\{\theta_i\}$ (d, σ^2; e, ν; f, ρ) and absolute covariance $|\text{cov}\{\theta_i, \theta_j\}|$ (g, $\{\sigma^2, \nu\}$; h, $\{\sigma^2, \rho\}$; i, $\{\nu, \rho\}$) from the observed ensemble (*circles*) and predicted from eq. (12) (*triangles*) scaled as $10\log_{10}(\cdot)$. Root-mean-squared error $|\text{rmse}\{\theta_i\}|$ (j, σ^2; k, ν; l, ρ) calculated from the ensemble. Absolute best-fit linear slopes in the $10\log_{10} - \log_2$ space for the empirical and predicted (co)variances and rmse are reported in the subplot legends for each experiment; grey guide lines show $1/\sqrt{K}$ slopes. The number of approximate one-third decorrelation lengths ($\pi\rho$) for this grid are quoted on the upper horizontal axis.

define

$$e(\boldsymbol{\theta}_0) = \frac{\bar{\mathcal{F}}^{-1}(\boldsymbol{\theta}_0)}{\text{cov}\{\boldsymbol{\theta}_0\}}, \tag{26}$$

for each parameter (pair), that is, including the parameter interactions in the ratio of simulation scheme cross-terms $\bar{\mathcal{F}}^{-1}(\theta_i, \theta_j)$ and $\text{cov}\{\theta_i, \theta_j\}$.

In Fig. 20, we display these results within the context of a growing domain for the full, regularly sampled experimental grid and the Matérn parameters used in the previous experiment (Fig. 19). Both $\bar{\mathcal{F}}^{-1}(\boldsymbol{\theta}_0)$ and $\text{cov}\{\boldsymbol{\theta}_0\}$ decrease linearly in the $\log_{10} - \log_2$ space for all parameter interactions, with the exception of $\{\sigma^2,\ \nu\}$. With increasing sample size, $\bar{\mathcal{F}}^{-1}(\boldsymbol{\theta}_0)$ and $\text{cov}\{\boldsymbol{\theta}_0\}$ converge in $\{\nu,\ \nu\}$, remain constant in $\{\nu,\ \rho\}$, and diverge in all other parameter pairings. Repeating this experiment for additional models, we find that the blurred inverse Fisher information approximation inconsistently approaches the parameter uncertainty in only some of the Matérn parameters for some models with small smoothness ($\nu_0 < 3$) parameters observed for grid sizes with lengths of $\sqrt{K} \leq 2^{10.5}$. In comparison to the

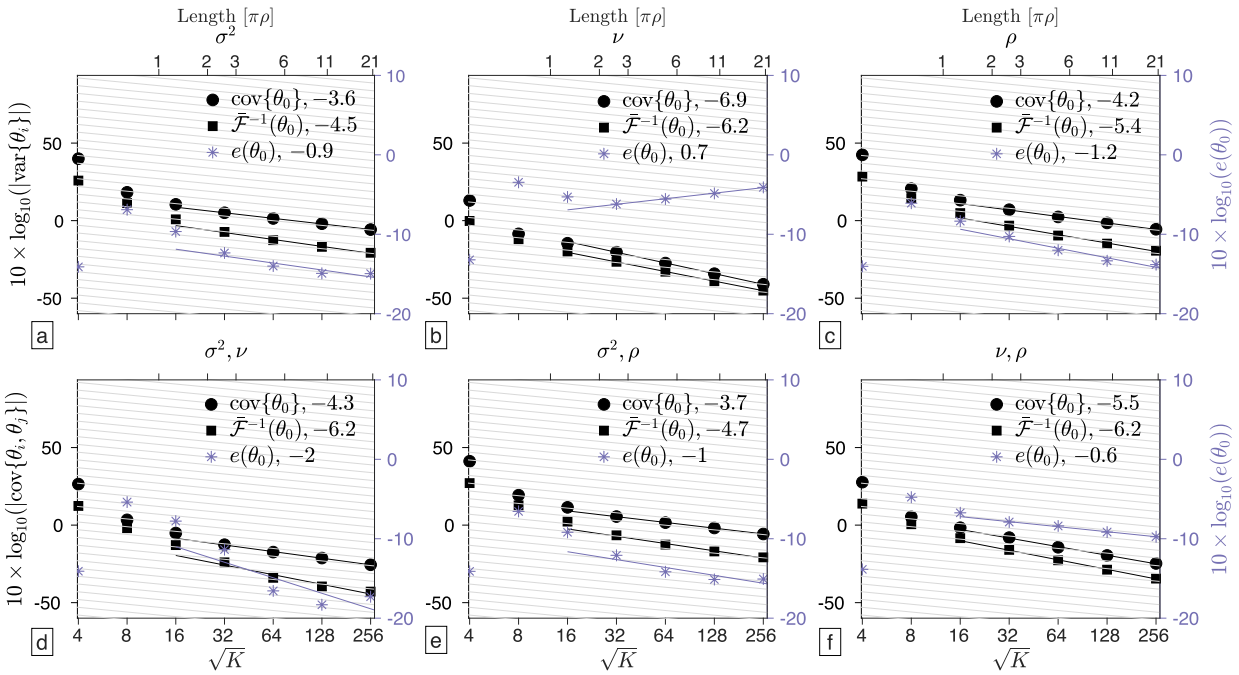


Figure 20. Growing-domain asymptotic behaviour of the elements of the absolute value of the blurred inverse Fisher information matrix (*black squares, left axis*), the absolute value of the analytically calculated estimator covariance (*black circles, left axis*) and their ratio $e(\boldsymbol{\theta}_0)$ (*purple stars, right axis*) as a function of the sample size via the grid-length increasing evenly in both spatial directions for a fully observed grid. Plot organization and annotation details are as in the second and third rows of Fig. 19, with horizontal axes in multiples of the correlation length $\pi\rho$ and in terms of number of samples $\sqrt{K}$.

analytically predicted parameter covariance calculated at the truth $\mathrm{cov}\{\boldsymbol{\theta}_0\}$, we find that $\bar{\mathcal{F}}^{-1}(\boldsymbol{\theta}_0)$ is an overly confident estimate of parameter uncertainty. The rate of growth of the ratio for many of these parameters, at least for the grid sizes we inspect, indicates that the blurred inverse Fisher information matrix does not consistently approach parameter uncertainty values and is thus not a sufficient approximation for sample window sizes that are realistic for much of the geophysical sciences.

8.2 Infill (fixed-domain) asymptotics

For our infill asymptotics experiments, we design square, evenly spaced grids with a common physical length dimension of $5\pi\rho$, that is, relative to the Matérn range parameter $\rho_0 = 16\,\mathrm{m}$, that we maintain by decreasing the sample spacing $\Delta y = \Delta x$ from 4 to 1 m for sample sizes that increase by powers of two from 64×64 to 256×256. We simulate 100 realizations for the Matérn model $\boldsymbol{\theta}_0 = [3\,\mathrm{unit}^2\ 0.5\ 16\,\mathrm{m}]$ on each trial grid for the case of full, regular sampling and the case of the partially (66.7 per cent) observed grid whose complement was randomly deleted. As we did in the growing-domain experiment (Section 8.1), we calculate the relative estimator bias of the ensembles, parameter variance and covariance from the ensemble and the analytic prediction (eq. 12), and the rmse (eq. 25) of the ensemble.

In Fig. 21, we show that densifying the full, regularly sampled grid only reduces the estimation variance of the smoothness parameter, denoted here as $\mathrm{var}\{\nu\}$, and the estimated parameter covariances that the smoothness pairs into, which we denote again here as $\mathrm{cov}\{\sigma^2, \nu\}$ and $\mathrm{cov}\{\nu, \rho\}$. Relative to the full, regularly sampled grid, the partially observed grid experiment reveals bias in the estimated smoothness ν and range ρ parameters that is roughly constant with increasing sample size, reduced parameter variance in the variance σ^2 parameter, larger parameter variance in the smoothness ν and some disagreement in magnitude between the ensemble and analytically calculated covariances. As increasing the number of samples within the interior of the fixed domain improves the resolution of the high-frequency content within the periodogram, to which the smoothness ν parameter is sensitive, the improvement in parameter (co)variance for the smoothness ν should be expected within the infill scenario (H. Zhang 2004).

8.3 Reducing missingness

We consider how reducing missingness in the sampling of an observed random field affects the estimator and the behaviour of its uncertainty. For increasing sample sizes, we replace irregularly omitted samples with data randomly distributed across the observation window. This experiment is akin to infill asymptotics in that we are (irregularly) densifying the sampling of an irregular grid as a function of sample size.

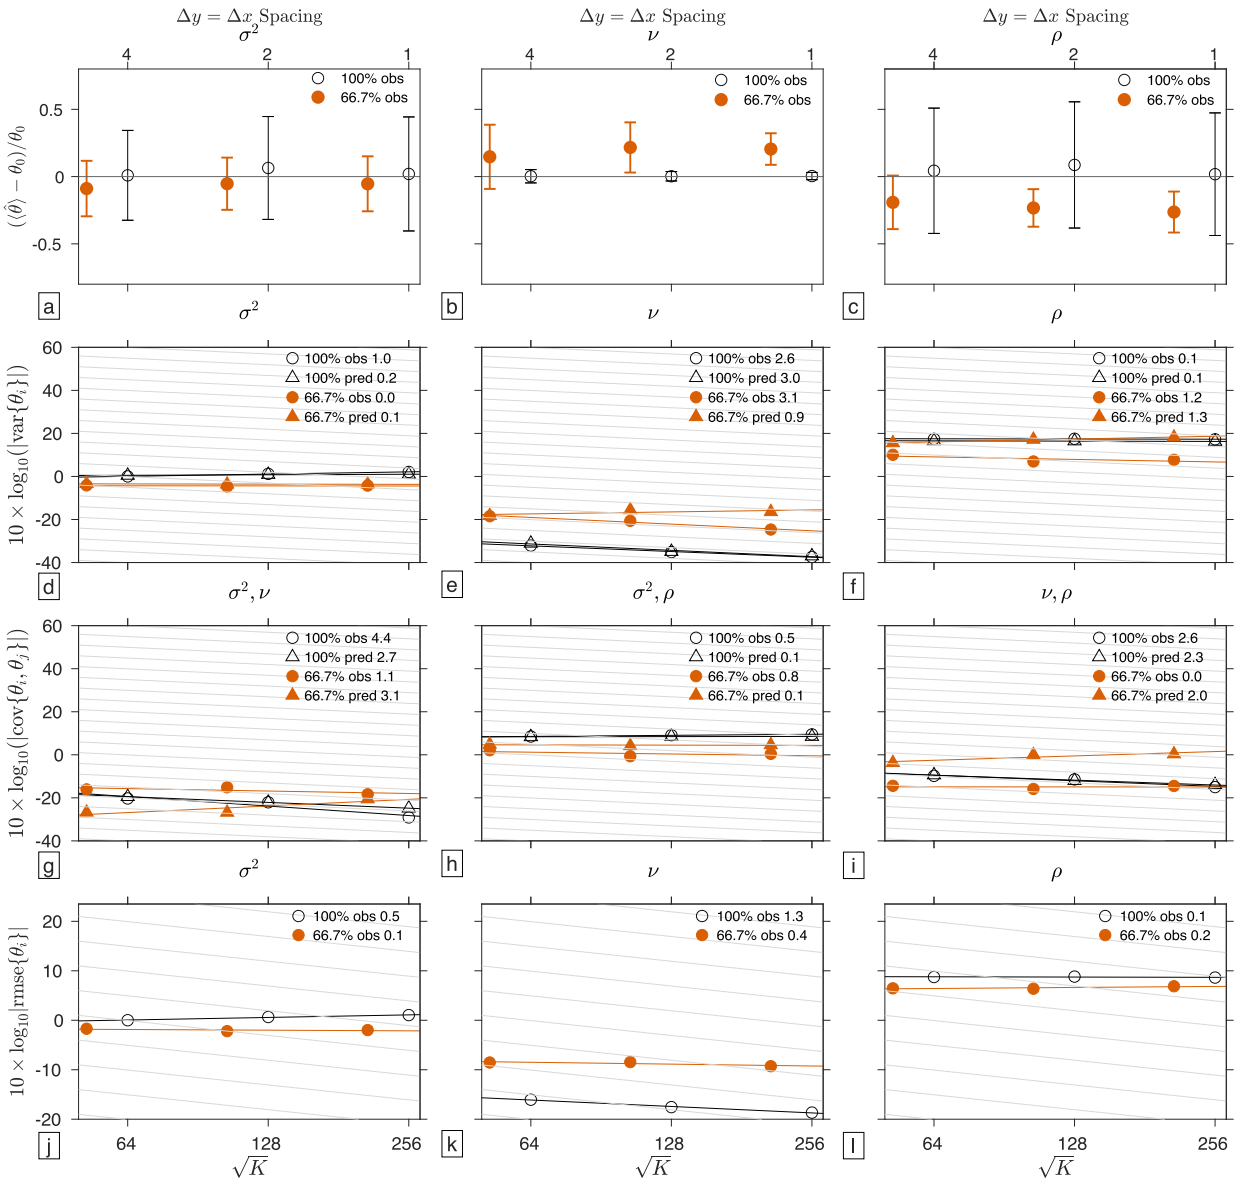


Figure 21. Fixed-domain (i.e. constant $N_y \Delta y = N_x \Delta x$) asymptotic behaviour of observed and predicted estimator covariance for increasing sample size through grid densification of a fully and partially (66.7 per cent) observed field for the Matérn model $\boldsymbol{\theta}_0 = [3\,\text{unit}^2\ 0.5\ 16\,\text{m}]$ with fixed physical grid-length of $5\pi\rho \times 5\pi\rho$ for 100 realizations of each trial grid. Plot organization and annotation are as in Fig. 19.

Fig. 22 displays the results of three experiments estimating the Matérn parameters for 100 realizations of a field with accumulating observations to simulate an increasing percentage of observation for three Matérn models with a common sampling grid and variance σ_0^2 and smoothness ν_0 parameters, and with three range parameters $\rho_0 = \{2, 4, 8\}$ that correspond to 1.6, 3.2 and 6.4 per cent of the diagonal length of the sampling grid. The relative estimator bias (Fig. 22, *first row*) reveals a positive dependence scaling with the magnitude of the range parameter ρ_0. For small enough ρ_0, the average parameter estimate is within 1 standard deviation (sd) of $\boldsymbol{\theta}_0$ for each trial grid size. However, for large ρ_0, the parameter bias exceeds 1 sd for ν and ρ. Parameter variance (Fig. 22, *second row*) increases at a slow linear rate for σ^2, decreases rapidly in ν and decreases slowly in ρ as the sample size and percent observation increase. Parameter covariance (Fig. 22, *third row*) increases rapidly in estimated $\{\sigma^2, \nu\}$ and $\{\nu, \rho\}$, but negligibly for $\{\sigma^2, \rho\}$. As in the infill experiment, the largest improvements are observed for the smoothness parameter, such as in the high rate of decay in its rmse (*fourth*) for all three ρ_0 models.

9 DESIGNING SPATIAL SAMPLING EXPERIMENTS

Sampling in the spatial domain shapes information in the spectral domain, which affects our statistical knowledge of the underlying process. As we illustrated in our realistic synthetic experiments in Figs 12 and 13, the geometry of the observation window contributes

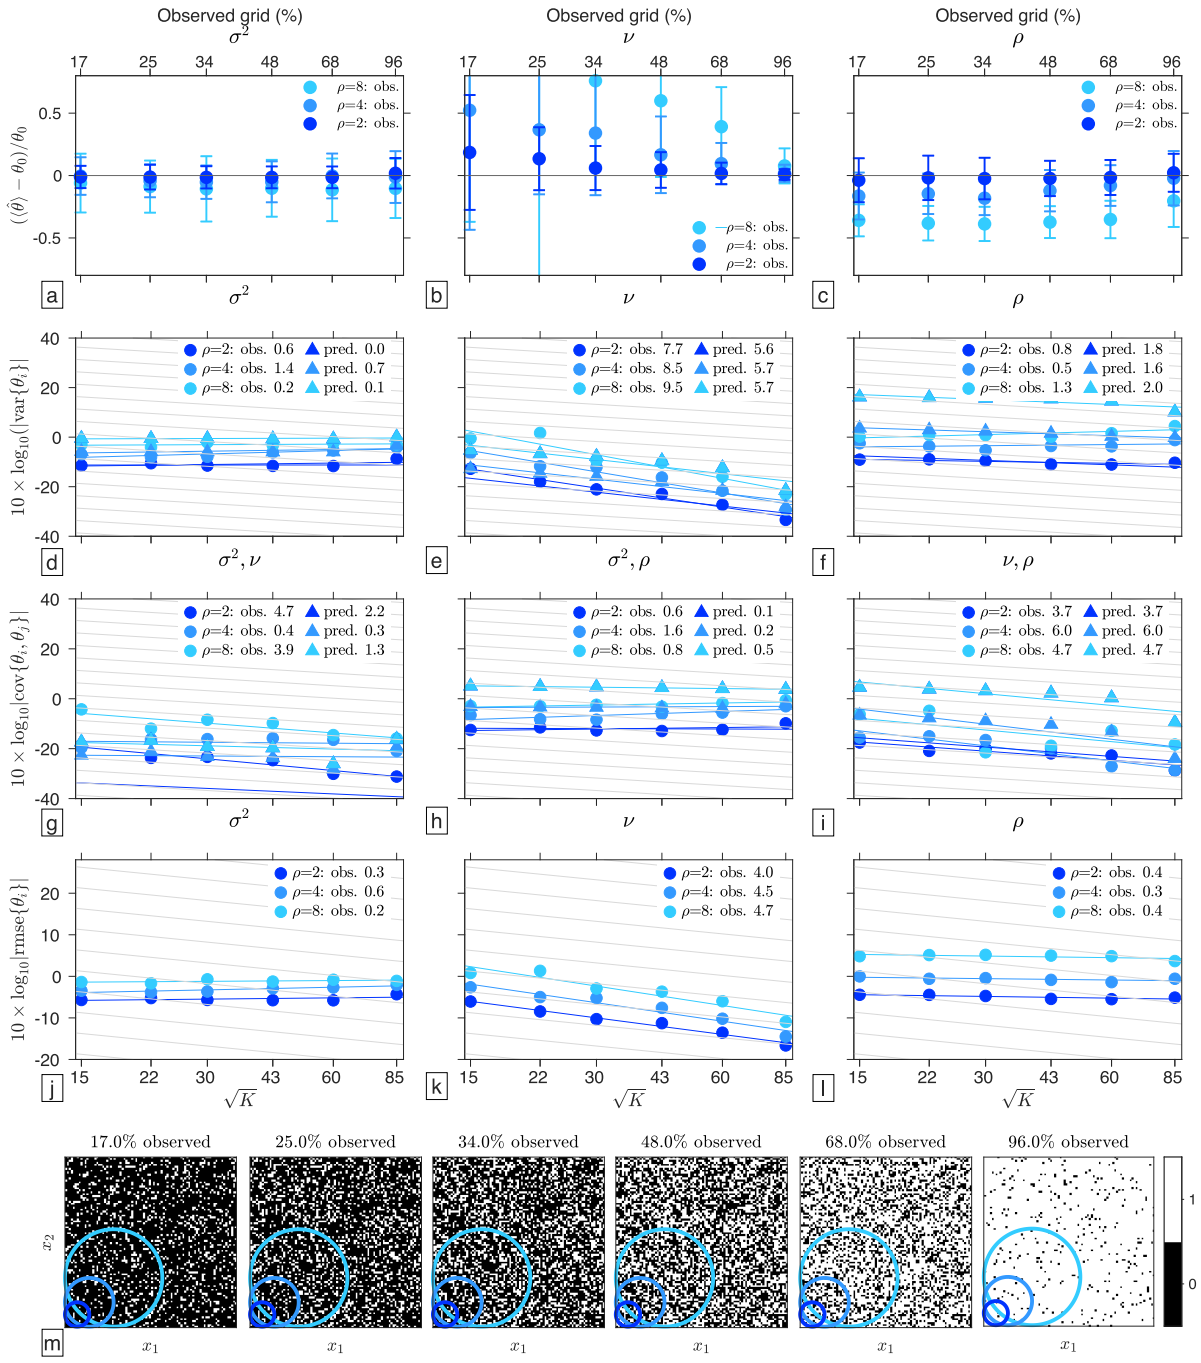


Figure 22. Estimator relative bias (*first row*), variance (*second*), covariance (*third*) and rmse (*fourth*) as a function of increasing observations within a fixed domain. Plot organization and annotation details for top four rows are as in Fig. 19. Experiments are performed for three Matérn models where $\sigma_0^2 = 3\,\text{unit}^2$ and $\nu_0 = 0.5$ for all, with values of $\rho_0 = 2\text{m}$ (light blue; 0.0160 of the grid diagonal), $\rho = 4\,\text{m}$ (medium blue, 0.0321 of the grid diagonal) and $\rho = 8\,\text{m}$ (dark blue; 0.0642 of the grid diagonal) on a 88×88 grid with even 1 m spacing. The spatial taper accumulates random samples after each trial of 100 realizations for the three models considered where the per cent observed is distributed from 17 to 96 per cent. Spatial sampling patterns for each per cent observation (omissions in black) are shown in the bottom row (*m*); circle radii correspond to the colour-coded value of $\pi\,\rho$ for visual context of the correlation lengths.

to parameter estimate covariance, in terms of both the spread of estimates and the strength of the parameter correlation. The spatial window $w(\mathbf{x})$ appears in the calculation of both the blurred inverse Fisher information matrix (eq. 16) and the covariance of the blurred score (eq. 17), which itself depends on the sampling-window informed covariance of the periodogram (eqs 18–21). As our calculation of predicted parameter covariance (eq. 12) depends on these terms, the expected behaviour of the model estimates ultimately hinges on the sampling window $w(\mathbf{x})$. As we wrote in Section 2.3, the debiasing effect of the sampling pattern blurring the spectral density can be

understood through the spectral window $|w(\mathbf{k})|^2$ and its interplay with $\bar{\mathcal{S}}_{\boldsymbol{\theta}}(\mathbf{k})$ and $\partial\bar{\mathcal{S}}_{\boldsymbol{\theta}}(\mathbf{k})/\partial\theta$ and their asymptotics already discussed in Section 8.

Gaining an intuitive appreciation for how sampling geometry influences estimator biases and uncertainties is relevant for a variety of signal processing tasks, for example, spectral estimation methods for irregularly sampled processes (e.g. T.P. Bronez 1988; L.D. Abreu *et al.* 2016), function approximation on bounded domains (e.g. F.J. Simons & D.V. Wang 2011; R. Matthysen & D. Huybrechs 2018) and basis representations of covariance spaces (e.g. C.E. Buell 1975, 1979; M.B. Richman 1986). The imprints of the sampling domain need to be accounted for lest they manifest and be misinterpreted as signal (e.g. C. Lehr & T.L. Hohenbrink 2025). Other disciplines that offer valuable perspectives for viewing the structure imparted by spatial sampling patterns include topology (L. Ibáñez *et al.* 1996) and graph signal processing (S. Chen *et al.* 2016). In the analogue setting of (Fourier) optics, a given aperture $w(\mathbf{x})$ can be shaped like a lens, and then the spectral window $|w(\mathbf{k})|^2$ is a far-field diffraction pattern.

Moreover, if we could design a sampling geometry for a spatial field from scratch, or if we had the opportunity to modify a given sampling scenario by the addition of more sampling points, where should they be allotted? In field campaigns or remote experiments, pragmatics and logistics and a prudent use of resources may be additional considerations (E.L. Evans *et al.* 2022; M. Zhang *et al.* 2025) but even in the least restricted of experimental settings, where samples are cheap and the region-of-interest is broad, appealing to sampling designs that prioritize information gain will reduce redundancy in samples collected and retained (M. Zhang & Y. Li 2024), alleviating the downstream burdens of large data in terms of storage and modelling.

For parametric covariance modelling, the first concern of data collection should be adequately sampling over an extent that is broad enough to capture several repetitions of the correlation length yet dense enough to prevent its aliasing (Fig. 18, *left*). Ideally, the field or remote-sensing scientist will be aware of the spatial bounds at which the assumption of local stationarity holds during their sample acquisition. Subsequent data collection campaigns can consider whether employing a growing-, fixed- or mixed-domain sampling plan would be most useful (e.g. in terms of parameter resolution) and whether these approaches are viable (e.g. whether the region-of-interest is geographically confined). From our studies of growing- and fixed-domain asymptotics for full and partial, regular and irregular grids (Figs 19–22), we observed that increasing the number of samples generally reduces parameter bias and (co)variance. An exception to this behaviour was observed for estimates of the variance σ^2 and range ρ in the infill experiments (Figs 21 and 22), whose performance either remained unchanged or negligibly degraded, worsening for experiments that included random missingness, and/or large correlation lengths ρ_0. While more samples will improve the behaviour of spatial covariance model estimates, we have shown that the rate at which they do is determined by the sampling pattern. For the regularly sampled grid, the sampling pattern can be modified to tune the recovery of our estimator by densifying the discretization or by physically expanding its size. For the partial grid, the geometry of sample distribution across a domain is an additional factor that must be considered in the design of spatial sampling experiments.

Optimal spatial sampling (K. Smith 1918; V.V. Fedorov 1971) can be designed to determine sample size and locations by which to minimize parameter uncertainties or maximize information gain. Many optimality criteria (N. Cressie 1993) can be implemented, and for K samples in the plane, one could consider as many as $\sum_{n=1}^{K} {}_nC_K$ combinatorial sampling patterns. To quell this mass of possibilities, the intuition and understanding built up so far can provide general guidance, and our theory and software for estimation, simulation and (data-less) uncertainty quantification for (ir)regularly sampled fields, fully equips us for a variety of constructions. Optimizing a sampling pattern for parameter uncertainty reduction can be geometrically understood as designing an aperture mask that concentrates and centres signal in the spectral domain, and reduces destructive interference by apodization. Such a mask would reduce irregularities, in particular those from random stippling, in the sampling pattern, which act like optical aberrations, producing clutter within the spectral window that diffuses signal.

Figs 23–26 illustrate the appearance of several sampling geometries as spatial windows and their corresponding spectral windows. We configure the spatial and spectral window pairs from left to right for three scenarios of increasing sample size following the three mechanisms we studied in Section 8, that is, for a growing domain, fixed domain and reduced missingness. We show two styles of dispersed sampling as uniform random missingness (Fig. 23) and checkerboard (Fig. 24) patterns. For the random missingness pattern, increasing the number of samples by any mechanism results in the amplitude of the spectral window focusing near the zero wavenumber. This is particularly impactful in the sparse sampling pattern shown in the bottom left for a 2 per cent observed 16×16 grid whose spectral window has transitioned to a lack of structure that will efface any modelled spectrum entirely. For the checkerboard pattern, the orientation of the longest sampled diagonal, which might be purely a matter of arbitrary availability, generates orthogonal diagonals in the spectral window, which will impart an indelible signature on the recovered spectrum. We revisit sampling contiguous patches, including within the interior (Fig. 25) and the exterior (Fig. 26) of an irregular boundary. In general, we have observed parameter (co)variance to be reduced for sampling patterns that follow contiguous patches rather than dispersive stippling (e.g. Fig. 12). Signal is concentrated over a broader range of low wavenumbers compared to the dispersed sampling examples, with clutter at higher wavenumbers reduced in magnitude for larger sample sizes. For both end-member (dispersed and contiguous) irregular sampling scenarios, the imprint of the sampling window $w(\mathbf{x})$ induces a diffraction pattern, plainly visible in the spectral window $w(\mathbf{k})$, and mathematically carried by the very terms that control the estimator and shape its uncertainty, that is, the blurred likelihood and its various derivatives.

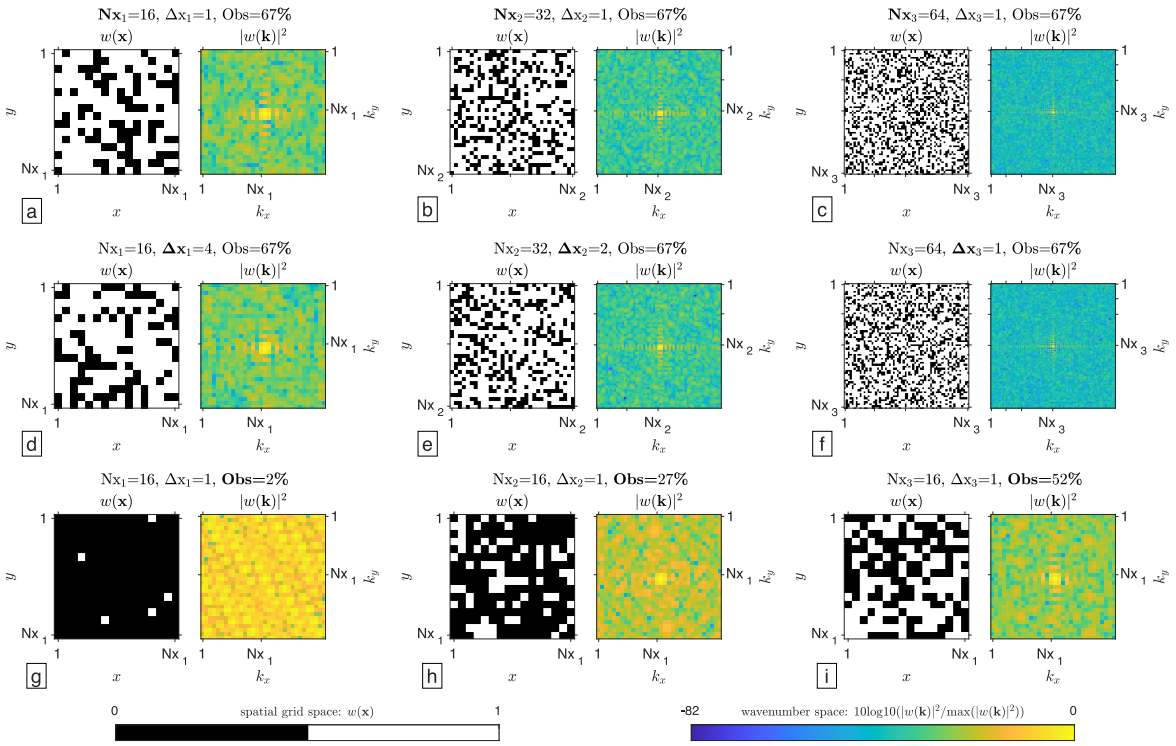


Figure 23. Spatial $w(\mathbf{x})$ (*left-most subfigure of each column pair*) and spectral $|w(\mathbf{k})|^2$ (*right-most of each column pair*) windows computed with a factor of two zero-padding in both directions for incompletely sampled grids. Sampling patterns organized as panels for missingness distributed as a uniformly random pattern. The three rows in each panel show allocation strategies for increasing sample size such that the physical grid increases in a growing domain (a–c), spacing decreases in a fixed domain (d–f) and missingness is reduced in an irregularly filled fixed domain (g–i). Grid parameters for the number of samples N_{x_i}, spacing Δx_i and the per cent observed 'Obs' of the regular grid are noted in the subfigure titles. The subfigures share a common colourmap.

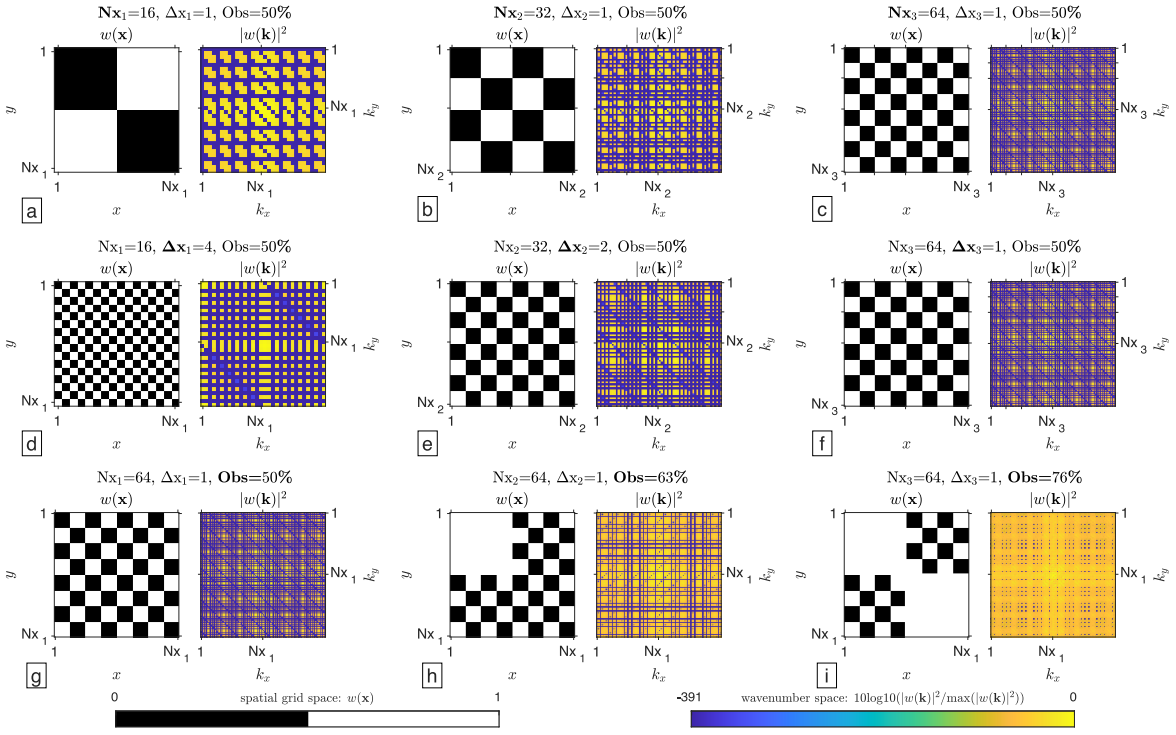


Figure 24. Spatial $w(\mathbf{x})$ (*left-most subfigure of each column pair*) and spectral $|w(\mathbf{k})|^2$ (*right-most of each column pair*) windows computed with a factor of two zero-padding in both directions for incompletely sampled grids following a checkerboard pattern. Organization of rows within the panels is as in Fig. 23 with grid parameters listed in the subfigure titles.

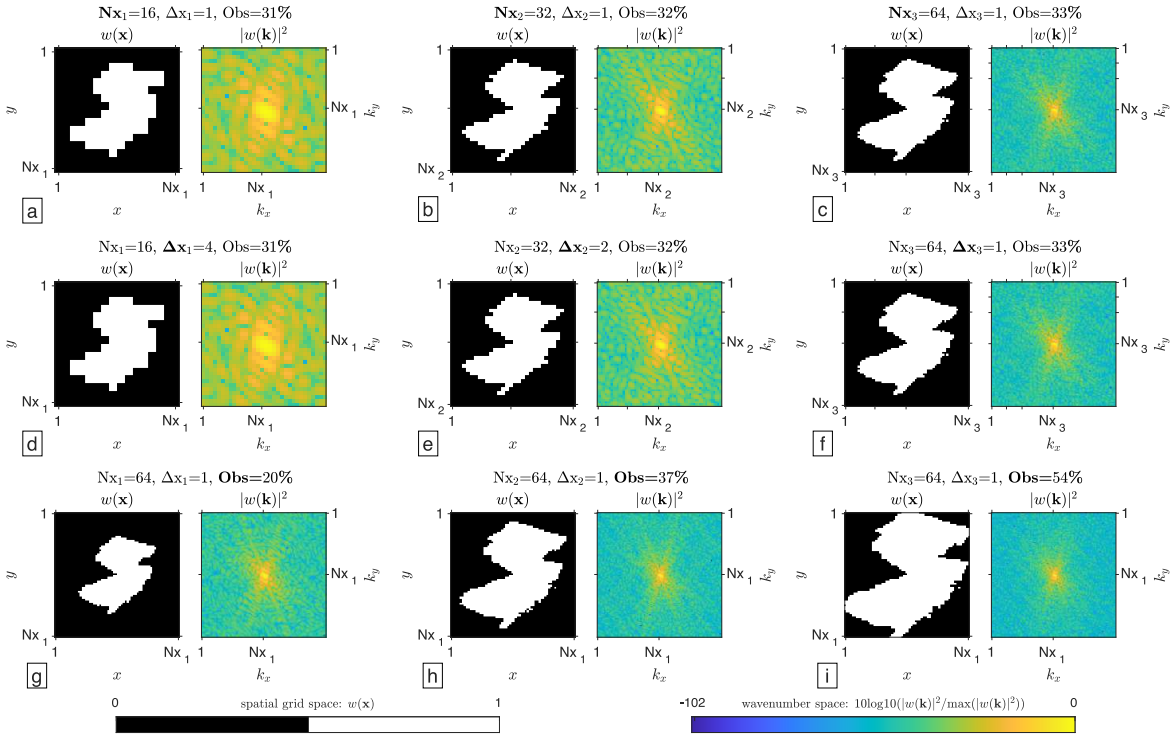


Figure 25. Spatial $w(\mathbf{x})$ (*left-most subfigure of each column pair*) and spectral $|w(\mathbf{k})|^2$ (*right-most of each column pair*) windows computed with a factor of two zero-padding in both directions for incompletely sampled grids constrained to the interior of the boundary of contiguous geographic patches. Organization of rows within the panels is as in Fig. 23 with grid parameters listed in the subfigure titles.

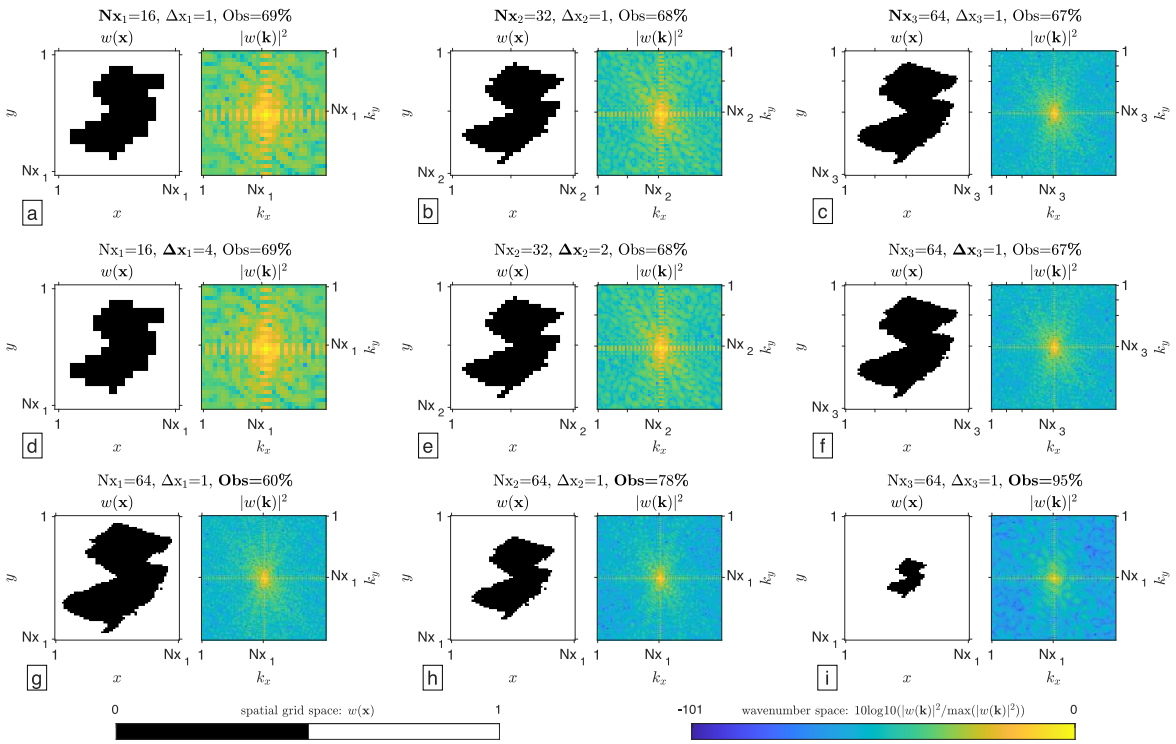


Figure 26. Spatial $w(\mathbf{x})$ (*left-most subfigure of each column pair*) and spectral $|w(\mathbf{k})|^2$ (*right-most of each column pair*) windows computed with a factor of two zero-padding in both directions for incompletely sampled grids constrained to the exterior of the boundary of contiguous geographic patches. Organization of rows within the panels is as in Fig. 23 with grid parameters listed in the subfigure titles.

10 CONCLUSIONS

Irregularly sampled spatial data are pervasive in the geophysical sciences. Treating such observations as random fields presents the opportunity to characterize them as single realizations of a statistical process. We have demonstrated how to properly account for sampling patterns, edges and omissions, by incorporating a (weighted) window of observation into a spectral-domain maximum-likelihood analysis. Given incomplete and irregularly spaced observations we estimate the covariance structure of a locally non-varying-mean field as a stationary, isotropic, parametric Matérn form. Our debiased Whittle maximum-likelihood estimation relies on the exact blurring of the modelled spectral density as the expected periodogram to balance the finite-field effects that affect the (windowed) periodogram of the data. While we show that we indeed can exclude wave-vector correlations for efficient parameter estimation, we also did demonstrate that these correlations can be completely calculated and correctly incorporated through the periodogram covariance for an uncertainty assessment that fully matches empirically observed behaviours. In realistic simulation scenarios and by analysis of real geophysical data, we substantiate that Matérn parameter estimates, their aleatoric uncertainty and epistemic model goodness-of-fit are fully and reliably quantifiable. Correctly fixing the smoothness (ν) to special familiar models that only parametrize variance (σ^2) and range (ρ) yields estimate ensembles and parameter covariances comparable to when all three parameters of the general Matérn class are inverted for. Crucially, incorrect *a priori* assumptions impart considerable bias. Little is gained beyond functional simplicity in selecting a special case over the general Matérn form: We should fully estimate it in real-data scenarios when we are without knowledge of the underlying process. We provide the numerical tools for implementation, estimation and simulation, and the analytical expressions for connecting the general Matérn with its special cases in both spatial and spectral representations to make this suggestion practical. Spatially merged processes are separable when their boundaries, however irregular, are known *a priori* and accounted for, but their analysis on an encompassing rectangular grid, violating the modelling assumption of stationarity, leads to strong parameter bias. Theory and simulation reveal that sample size and geometry play a considerable role in parameter (co)variance, which we understand through the spectral window. Simulations under growing-domain and infill asymptotic regimes for full or partial observations, with sampling irregularities in the way of boundaries, structure, or missingness, reveal the rates at which estimator performance improve with sample size, setting the stage for experimental design and (optimal) sample allocation. Analysis and synthesis of planetary-scale geophysical data, tapestries of time and terrain, as random fields, is efficient under our framework, and our ability to critically assess it will enable relaxing our current null hypotheses that such data are Gaussian, stationary, isotropic, random, Matérn, fields.

ACKNOWLEDGMENTS

OLW gratefully acknowledges financial support from the Schmidt DataX Fund, grant number 22-008, at Princeton University, made possible through a major gift from the Schmidt Futures Foundation, NSF grants EAR-2341811 and EAR-2422649, and Gordon and Betty Moore Foundation grant 13614 to FJS. We thank the Associate Editor, Carl Tape, and the reviewers, Andrew Valentine and John Goff, for very constructive reviews that helped improve the manuscript.

DATA AVAILABILITY

All the code used to conduct the calculations and produce the figures in this paper made in MATLAB by the authors is openly available as Release 2.0.0 from https://github.com/csdms-contrib/slepian_juliet, doi: `10.5281/zenodo.4085253`. For an alternative code base in Python, see also A.P. Guillaumin *et al.* (2026). Also compare with the *R* package by C.J. Paciorek (2007). The seafloor bathymetry model `GEBCO_2024` and its Type Identifier Grid are available from https://www.gebco.net/data-products/gridded-bathymetry-data/gebco2024-grid; both data products are described by GEBCO Bathymetric Compilation Group (2024). The lunar digital elevation model `SLDEM2015` is available from https://pgda.gsfc.nasa.gov/products/54 and is described by M.K. Barker *et al.* (2016). The crater catalogue of S.J. Robbins (2018) is available from https://astrogeology.usgs.gov/search/map/moon_crater_database_v1_robbins. The remote sensing data are available from https://earthexplorer.usgs.gov for the Collection 2, Level-2 data products of the path 014, row 032 tile captured on 2025 July 20 and are further described by Earth Resources Observation and Science (EROS) Center (2019, 2020).

REFERENCES

Abramowitz, M. & Stegun, I.A., 1965. Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables, Dover, New York.

Abreu, L.D., Gröchenig, K. & Romero, J.L., 2016. On accumulated spectrograms, *Trans. Am. Math. Soc.,* **368**(5), 3629–3649.

Adler, R.J., 1981. *The Geometry of Random Fields,* Wiley, New York.

Barker, M.K., Mazarico, E., Neumann, G.A., Zuber, M., Haruyama, J. & Smith, D.E., 2016. A new lunar digital elevation model from the Lunar Orbiter Laser Altimeter and SELENE Terrain Camera, *Icarus,* **237,** 346–355.

Beyaztas, U., Shang, H.L. & Yaseen, Z.M., 2021. A functional autoregressive model based on exogenous hydrometeorological variables for river flow prediction, *J. Hydrol.,* **598,** 126 380.

Bronez, T.P., 1988. Spectral estimation of irregularly sampled multidimensional processes by generalized prolate spheroidal sequences, *IEEE T. Acoust. Speech Signal Process.,* **36**(12), 1862–1873.

Buell, C.E., 1975. The topography of empirical orthogonal functions, *Preprints Fourth Conf. on Prob. and Stats. in Atmos. Sci.,* Vol. **188,** Amer. Meteor. Soc.

Buell, C.E., 1979. On the physical interpretation of empirical orthogonal functions. *Preprints Sixth Conf. on Prob. and Stats. in Amos. Sci.,* Vol. **112,** Amer. Meteor. Soc.
Bui-Thanh, T., Ghattas, O., Martin, J. & Stadler, G., 2014. A computational framework for infinite-dimensional Bayesian inverse problems, Part I: the linearized case, with application to global seismic inversion, *SIAM J. Sci. Comput.,* **35**(6), A2494–A2523.
Cárdenas-Avendaño, A., Lupsasca, A. & Zhu, H., 2023. Adaptive analytical ray tracing of black hole photon rings, *Phys. Rev. D,* **107,** 043 030.
Chan, G. & Wood, A.T.A., 1999. Simulation of stationary Gaussian vector fields, *Stat. Comput.,* **9**(4), 265–268.
Chen, S., Varma, R., Singh, A. & Kovačević, J., 2016. Signal recovery on graphs: fundamental limits of sampling strategies, *IEEE Trans. Signal Inf. Process. Netw.,* **2**(4), 539–554.
Christakos, G., 1992. *Random Field Models in Earth Sciences,* 2nd edn., Academic Press, San Diego, Calif.
Chu, T., 2023. Mixed domain asymptotics for geostatistical processes, *Stat. Sin.,* **33**(1), 551–571.
Cramér, H., 1942. On harmonic analysis in certain fuctional spaces, *Arkiv Mat. Astr. Fysik,* **28B,** 1–7.
Cramér, H., 1946. *Mathematical Methods of Statistics,* Princeton Univ. Press.
Cressie, N., 1993. *Statistics for Spatial Data,* John Wiley, London, UK.
Dahlhaus, R. & Künsch, H., 1987. Edge effects and efficient parameter estimation for stationary random fields, *Biometrika,* **74**(4), 877–882.
de Karman, T. & Howarth, L., 1938. On the statistical theory of isotropic turbulence, *Proc. R. Soc. London, Ser. A,* **164**(917), 192–215.
Deb, S., Pourahmadi, M. & Wu, W.B., 2017. An asymptotic theory for spectral analysis of random fields, *Electron. J. Stat.,* **11,** 4297–4322.
Dietrich, C.R. & Newsam, G.N., 1997. Fast and exact simulation of stationary Gaussian processes through circulant embedding of the covariance matrix, *SIAM J. Sci. Comput.,* **18**(4), 1088–1107.
Earth Resources Observation and Science (EROS) Center, 2019. *Landsat 8 (L8) Data Users Handbook, Tech. Rep.,* Department of the Interior, U.S. Geological Survey.
Earth Resources Observation and Science (EROS) Center, 2020. Landsat 8-9 operational land imager / thermal infrared sensor level-2, collection 2 [dataset], Tech. Rep., U.S. Geological Survey.
Erdélyi, A., 1953. *Higher Transcendental Functions, Vol. II,* McGraw-Hill, New York.
Evans, E.L., Minson, S.E. & Chadwell, C.D., 2022. Imaging the next Cascadia earthquake: optimal design for a seafloor GNSS-A network, *Geophys. J. Int.,* **228,** 944–957.
Fedorov, V.V., 1971. The design of experiments in the multiresponse case, *Theory Probab. Appl.,* **16**(2), 323–332.
Frankel, A. & Clayton, R.W., 1986. Finite difference simulations of seismic scattering: implications for the propagation of short-period seismic waves in the crust and models of crustal heterogeneity, *J. geophys. Res.,* **91**(B6), 6465–6489.
GEBCO Bathymetric Compilation Group, 2024. The GEBCO_2024 grid—A continuous terrain model of the global oceans and land, Tech. Rep., NERC EDS British Oceanographic Data Centre, National Oceanography Centre, NOC.
Goff, J.A. & Jordan, T.H., 1988. Stochastic modeling of seafloor morphology: inversion of Sea Beam data for second-order statistics, *J. geophys. Res.,* **93**(B11), 13 589–13 608.
Gradshteyn, I.S. & Ryzhik, I.M., 2000. *Tables of Integrals, Series, and Products,* 6th edn., Academic Press, San Diego, Calif.
Guillaumin, A.P., Sykulski, A.M., Olhede, S.C., Early, J.J. & Lilly, J.M., 2017. Analysis of non-stationary modulated time series with applications to oceanographic surface flow measurements, *J. Time Ser. Anal.,* **38**(5), 668–710.
Guillaumin, A.P., Sykulski, A.M., Olhede, S.C. & Simons, F.J., 2022. The debiased spatial Whittle likelihood, *J. R. Stat. Soc., Ser. B,* **84**(4), 1526–1557.
Guillaumin, A.P., Goodwin, T., Walbert, O.L., Sykulski, A.M., Olhede, S.C. & Simons, F.J., 2026. DSWL package: a Python implementation of the Debiased Spatial Whittle Likelihood, *J. Open Source Softw.,* **11**(119), 8323.
Guttorp, P. & Gneiting, T., 2006. Studies in the history of probability and statistics XLIX. On the Matérn correlation family, *Biometrika,* **93**(4), 989–995.
Handcock, M.S. & Wallis, J.R., 1994. An approach to statistical spatial-temporal modeling of meteorological fields, *J. Am. Stat. Assoc.,* **89**(426), 368–378.
Häselich, M., Arends, M., Wojke, N., Neuhaus, F. & Paulus, D., 2013. Probabilistic terrain classification in unstructured environments, *Robot. Auton. Syst.,* **61**(10), 1051–1059.
Isserlis, L., 1918. On a formula for the product-moment coefficient of any order of a normal frequency distribution in any number of variables, *Biometrika,* **12**(1–2), 134–139.
Kroese, D.P. & Botev, Z.I., 2015 Spatial process simulation, edited by *Stochastic Geometry, Spatial Statistics and Random Fields, Chap. 12 ,* pp. 369–404, ed., Schmidt, , V., Springer, Heidelberg, Germany.
Lehr, C. & Hohenbrink, T.L., 2025. Technical note: an illustrative introduction to the domain dependence of spatial principal component patterns, *Hydrol. Earth Syst. Sci.,* **29,** 6735–6760.
Matérn, B., 1960. *Spatial Variation: Stochastic Models and Their Application to Some Problems in Forest Surveys and Other Sampling Investigations,* Vol. **49,** Statens Skogsforskningsintitut, Stockholm, Sweden.
Matthysen, R. & Huybrechs, D., 2018. Function approximation on arbitrary domains using Fourier extension frames, *SIAM J. Numer. Anal.,* **56**(3), 1360–1385.
Monsalve-Bravo, G.M. *et al.*, 2022. Analysis of sloppiness in model simulations: unveiling parameter uncertainty when mathematical models are fitted to data, *Sci. Adv.,* **8,** eabm5952.
Muir, J.B. & Ross, Z.E., 2023. A deep Gaussian process model for seismicity background rates, *Geophys. J. Int.,* **234,** 427–438.
Ibáñez, L.I., Hamitouche, C. & Roux, C., 1996. Determination of discrete sampling grids with optimal topological and spectral properties, Discrete Geometry for Computer Imagery. DGCI 1996., Vol. 1176 of Lecture Notes in Computer Science, pp. 177–192 10.1007/3–540–62005–2_15, Springer, Berlin, Germany.
North, G.R., Wang, J. & Genton, M.G., 2011. Correlation models for temperature fields, *J. Clim.,* **24,** 5850–5862.
Olhede, S.C., McCoy, E.J. & Stephens, D.A., 2004. Large-sample properties of the periodogram estimator of seasonally persistent processes, *Biometrika,* **91**(3), 613–628.
Paciorek, C.J., 2007. Bayesian smoothing with Gaussian processes using Fourier basis functions in the spectralGP package, *J. Stat. Softw.,* **19**(2), nihpa22751.
Pardo-Igúzquiza, E. & Dowd, P., 2022. The roughness of martian topography: A metre-scale fractal analysis of six selected areas, *Icarus,* **384**(115109), 1–14.
Percival, D.B. & Walden, A.T., 1993. *Spectral Analysis for Physical Applications, Multitaper and Conventional Univariate Techniques,* Cambridge Univ. Press, New York.

Przybilla, J., Wegler, U. & Korn, M., 2009. Estimation of crustal scattering parameters with elastic radiative transfer theory, *Geophys. J. Int.,* **178**(2), 1105–1111.
Rao, C.R., 1945. Information and the accuracy attainable in the estimation of statistical parameters, *Bull. Calcutta Math. Soc.,* **37,** 81–89.
Richman, M.B., 1986. Rotation of principal components, *J. Climatol.,* **6,** 293–335.
Robbins, S.J., 2018. A new global database of lunar impact craters > 1–2 km: 1. Crater locations and sizes, comparisons with published databases, and global analysis, *J. geophys. Res.,* **124**(4), 871–892.
Sandwell, D.T. *et al.*, 2022. Improved bathymetric prediction using geological information: SYNBATH, *Earth Space Sci.,* **9**(2), e2021EA002069.
Schwaiger, T., Jault, D., Gillet, N., Schaeffer, N. & Mandea, M., 2023. Local estimation of quasi-geostrophic flows in Earth's core, *Geophys. J. Int.,* **234,** 494–511.
Simons, F.J. & Olhede, S.C., 2013. Maximum-likelihood estimation of lithospheric flexural rigidity, initial-loading fraction and load correlation, under isotropy, *Geophys. J. Int.,* **193**(3), 1300–1342.
Simons, F.J. & Wang, D.V., 2011. Spatiospectral concentration in the Cartesian plane, *Int. J. Geomath.,* **2**(1), 1–36.
Simons, F.J., Zuber, M.T. & Korenaga, J., 2000. Isostatic response of the Australian lithosphere: estimation of effective elastic thickness and anisotropy using multitaper spectral analysis, *J. geophys. Res.,* **105**(B8), 19 163–19 184.
Simons, F.J., Walbert, O.L., Guillaumin, A.P., Eggers, G.L., Lewis, K.W. & Olhede, S.C., 2026. Maximum-likelihood estimation of the Matérn covariance structure of isotropic spatial random fields on finite, sampled grids, *Geophys. J. Int.,* **245**(2), 1–24.
Slobbe, D.C., Simons, F.J. & Klees, R., 2012. The spherical Slepian basis as a means to obtain spectral consistency between mean sea level and the geoid, *J. Geod.,* **86**(8), 609–628.
Smith, K., 1918. On the standard deviations of adjusted and interpolated values of an observed polynomial function and its constants and the guidance they give towards a proper choice of the distribution of observations, *Biometrika,* **12**(1 and 2), 1–85.
Smith, W.H.F. & Sandwell, D.T., 1997. Global sea floor topography from satellite altimetry and ship depth soundings, *Science,* **277**(5334), 1956–1962.
Stein, M.L., 1995. Fixed-domain asymptotics for spatial periodograms, *J. Am. Stat. Assoc.,* **90**(432), 1277–1288.
Stein, M.L., 1999. *Interpolation of Spatial Data: Some Theory for Kriging, Springer Series in Statistics,* Springer-Verlag, Berlin.
Sykulski, A.M., Olhede, S.C., Guillaumin, A.P., Lilly, J.M. & Early, J.J., 2019. The debiased Whittle likelihood, *Biometrika,* **106,** 251–266.
Tarantola, A. & Nercessian, A., 1984. Three-dimensional inversion without blocks, *Geophys. J. R. Astron. Soc.,* **76**(2), 299–306.
Thrun, S., Montemerlo, M. & Aron, A., 2006. Probabilistic terrain analysis for high-speed desert driving, Robot.: Sci. Syst.,16–19.
Tunnicliffe Wilson, G., Haywood, J. & Petherick, L., 2022. Modeling cycles and interdependence in irregularly sampled geophysical time series, *Environmetrics,* **33**(2), e2708.
Valentine, A.P. & Davies, D.R., 2020. Global models from sparse data: a robust estimate of Earth's residual topography spectrum, *Geochem. Geophys. Geosys.,* **21**(8), e2020GC009240.
Vuong, Q.H., 1989. Likelihood ratio tests for model selection and non-nested hypotheses, *Econometrica,* **57**(2), 307–333.
Walden, A.T., McCoy, E. & Percival, D.B., 1994. The variance of multitaper spectrum estimates for real Gaussian processes, *IEEE Trans. Signal Process.,* **42,** 479–482.
Whittle, P., 1954. On stationary processes in the plane, *Biometrika,* **41**(3–4), 434–449.
Zhang, H., 2004. Inconsistent estimation and asymptotically equal interpolations in model-based geostatistics, *J. Am. Stat. Assoc.,* **99**(465), 250–261.
Zhang, H. & Zimmerman, D.L., 2005. Towards reconciling two asymptotic frameworks in spatial statistics, *Biometrika,* **92**(4), 921–936.
Zhang, M. & Li, Y., 2024. Ergodic sampling: acquisition design to maximize information from limited samples, *Geophys. Prospect.,* **72**(2), 435–467.
Zhang, M., Battig, E. & Li, Y., 2025. First field trial of ergodic sampling for airborne magnetic survey, *Leading Edge,* **44**(11), 831–837.
Zhang, Y., Wang, D., Li, H. & Gao, J., 2023. S-wave velocity prediction using physical model-driven Gaussian process regression: a case study of tight sandstone reservoir, *Geophysics,* **88**(2), D85–D93.

APPENDIX A

Tables

Table A1. Analytic forms of the isotropic Matérn covariance and 2-D spectral density for special values of ν. The d-D spectral density expressions are provided in Table A2. All simplifications are trivial, with $\lim_{\nu\to\infty}$ requiring the application of L'Hospital's rule in triplicate.

Smoothness ν	Normalized spatial covariance $\mathcal{C}_{\boldsymbol{\theta}}(r)/\sigma^2$	Normalized spectral density $\mathcal{S}^2_{\boldsymbol{\theta}}(k)/\sigma^2$	Street name
1/3	$\frac{4^{1/3}}{\Gamma(\frac{1}{3})}\left(\frac{2\sqrt{3}}{3\pi\rho}r\right)^{1/3}K_{1/3}\left(\frac{2\sqrt{3}}{3\pi\rho}r\right)$	$\frac{4^{1/3}\pi\rho^2}{(4+3\pi^2\rho^2k^2)^{4/3}}$	von Kármán
1	$\left(\frac{2}{\pi\rho}r\right)K_1\left(\frac{2}{\pi\rho}r\right)$	$\frac{4\pi\rho^2}{(4+\pi^2\rho^2k^2)^2}$	Whittle
1/2	$\exp\left(-\frac{\sqrt{2}}{\pi\rho}r\right)$	$\frac{2\pi\rho^2}{(4+2\pi^2\rho^2k^2)^{3/2}}$	Exponential
3/2	$\exp\left(-\frac{\sqrt{6}}{\pi\rho}r\right)\left(1+\frac{\sqrt{6}}{\pi\rho}r\right)$	$\frac{9\sqrt{6}\pi\rho^2}{(6+\pi^2\rho^2k^2)^{5/2}}$	Autoregressive (second order)
5/2	$\exp\left(-\frac{\sqrt{10}}{\pi\rho}r\right)\left(1+\frac{\sqrt{10}}{\pi\rho}r+\frac{10}{3\pi^2\rho^2}r^2\right)$	$\frac{250\sqrt{10}\pi\rho^2}{(10+\pi^2\rho^2k^2)^{7/2}}$	Autoregressive (third order)
∞	$\exp\left(-\frac{1}{\pi^2\rho^2}r^2\right)$	$\left(\frac{\pi\rho^2}{4}\right)\exp\left(-\frac{\pi^2\rho^2k^2}{4}\right)$	Squared exponential

Table A2. Simplified analytic forms of the isotropic Matérn covariance and d-D spectral density for special values of ν for comparison with the 2-D spectral density Table A1, with the addition of the general half-integer smoothness case, $\nu = (n+1)/2$.

Smoothness ν	Normalized spatial covariance $C_{\boldsymbol{\theta}}(r)/\sigma^2$	Normalized spectral density $S^d_{\boldsymbol{\theta}}(k)/\sigma^2$
1/3	$\frac{2^{2/3}}{\Gamma(\frac{1}{3})}\left(\frac{2\sqrt{3}}{3\pi\rho}r\right)^{1/3} K_{1/3}\left(\frac{2\sqrt{3}}{3\pi\rho}r\right)$	$\frac{\Gamma(1/3+d/2)}{\Gamma(1/3)}\frac{4^{1/3}(3\pi\rho^2)^{d/2}}{(4+3\pi^2\rho^2k^2)^{1/3+d/2}}$
1	$\left(\frac{2}{\pi\rho}r\right)K_1\left(\frac{2}{\pi\rho}r\right)$	$\Gamma(1+d/2)\frac{4(\pi\rho^2)^{d/2}}{(4+\pi^2\rho^2k^2)^{1+d/2}}$
1/2	$\exp\left(-\frac{\sqrt{2}}{\pi\rho}r\right)$	$\frac{\Gamma(1/2+d/2)}{\sqrt{\pi}}\frac{\sqrt{2}(\pi\rho^2)^{d/2}}{(2+\pi^2\rho^2k^2)^{(1+d)/2}}$
3/2	$\exp\left(-\frac{\sqrt{6}}{\pi\rho}r\right)\left(1+\frac{\sqrt{6}}{\pi\rho}r\right)$	$\frac{2\,\Gamma(3/2+d/2)}{\sqrt{\pi}}\frac{6\sqrt{6}(\pi\rho^2)^{d/2}}{(6+\pi^2\rho^2k^2)^{(3+d)/2}}$
5/2	$\exp\left(-\frac{\sqrt{10}}{\pi\rho}r\right)\left(1+\frac{\sqrt{10}}{\pi\rho}r+\frac{10}{3\pi^2\rho^2}r^2\right)$	$\frac{4\,\Gamma(5/2+d/2)}{3\sqrt{\pi}}\frac{100\sqrt{10}(\pi\rho^2)^{d/2}}{(10+\pi^2\rho^2k^2)^{(5+d)/2}}$
$(n+1)/2$	$\exp\left(-\frac{2\sqrt{n+1/2}}{\pi\rho}r\right)\frac{n!}{(2n)!}\sum_{k=0}^{n}\frac{(n+k)!}{k!(n-k)!}\left(\frac{4\sqrt{n+1/2}}{\pi\rho}r\right)^{n-k}$	$\frac{\Gamma(n+(1+d)/2)}{\Gamma(n+1/2)}\left(\frac{(4(n+1/2))^{n+1/2}(\pi\rho^2)^{d/2}}{4(n+1/2)+\pi^2\rho^2k^2}\right)^{n+(1+d)/2}$
∞	$\exp\left(-\frac{1}{\pi^2\rho^2}r^2\right)$	$\left(\frac{\pi\rho^2}{4}\right)^{d/2}\exp\left(-\frac{\pi^2\rho^2k^2}{4}\right)$

Table A3. Results from the simulation experiments presented in Fig. 12 for 100 realizations of irregular sampling patterns. For each sampling scenario, the experiments are ordered by the inversion for the three Matérn parameters (3p) and the two parameters for correctly (2c) and incorrectly (2i) fixed ν. The columns list the sample means for each of the estimated Matérn parameters ($\overline{\theta}_i$, $i = 1, 2, 3$, that is, $\overline{\sigma^2}$, $\overline{\nu}$ and $\overline{\rho}$), the parameter standard deviations (analytically calculated from eq. 12 at the mean estimate, followed by the empirical values for the ensemble), both listed in per cent of the parameter itself and the parameter correlation in per cent calculated from the normalized (analytical | empirical) parameter covariance. All experiment details are as in Fig. 12.

Taper	$\boldsymbol{\theta}$	Per cent obs	$\overline{\sigma^2}$ units^2	±Per cent	$\overline{\nu}$	±Per cent	$\overline{\rho}\times 10$ km	±Per cent	$\{\sigma^2,\nu\}$	$\{\sigma^2,\rho\}$	$\{\nu,\rho\}$
Full	3p	100	999.79	19.7 |19.8	1.001	1.3 |1.5	1.986	11.4 |12.1	−50.5 |−58.7	97.6 |97.7	−67.2 |−74.3
Full	2c	100	999.95	19.8 |18.7	*1*	–|–	1.986	10.1 |9.6	–|–	–|–	–|–
Full	2i	100	1028.08	32.6 |11.3	*0.5*	–|–	10.404	165.3 |11.4	–|–	–|–	–|–
Bound	3p	33.6	994.57	17.1 |16.1	1.006	4.1 |3.3	1.986	10.1 |9.3	10.2 |2.3	94.2 |95.2	−18.3 |−19.5
Bound	2c	33.6	988.62	17.1 |15.1	*1*	–|–	1.980	9.7 |8.6	–|–	96.4 |95.9	–|–
Bound	2i	33.6	695.62	10.0 |10.4	*0.5*	–|–	5.806	10.6 |38.6	–|–	98.4 |93.4	–|–
Struct	3p	33.6	996.77	9.2 |8.9	1.005	13.6 |13.0	1.991	10.5 |10.0	−0.92 |−8.4	65.9 |61.8	−62.5 |−69.5
Struct	2c	33.6	991.95	10.1 |9.5	*1*	–|–	1.997	8.7 |8.1	–|–	83.2 |81.8	–|–
Struct	2i	33.6	833.08	8.6 |9.2	*0.5*	–|–	3.246	11.8 |16.5	–|–	89.9 |81.4	–|–
Rand	3p	33.6	987.82	6.8 |6.7	1.025	11.6 |11.1	1.974	7.4 |6.3	−18.6 |−9.5	70.0 |73.7	−68.6 |−53.4
Rand	2c	33.6	1006.43	7.0 |7.0	*1*	–|–	1.997	5.5 |5.8	– |–	82.5 |80.4	–|–
Rand	2i	33.6	971.84	6.3 |6.8	*0.5*	–|–	2.410	8.9 |7.8	– |–	85.3 |83.6	–|–
Anti-Bnd	3p	66.4	999.24	17.1 |14.9	1.003	3.6 |3.5	1.991	10.2 |8.9	35.5 |33.5	92.1 |90.6	2.7 |1.1
Anti-Bnd	2c	66.4	989.04	15.7 |14.4	*1*	– |–	1.980	10.7 |9.8	–|–	95.7 |95.1	–|–
Anti-Bnd	2i	66.4	759.05	12.9 |9.3	*0.5*	– |–	6.892	13.8 |49.5	– |–	99.6 |93.0	–|–
Anti-Str	3p	66.4	978.19	8.9 |8.9	1.020	8.7 |9.3	1.957	8.2 |7.5	0.4|14.0	72.3 |70.9	−59.8 |−52.2
Anti-Str	2c	66.4	1007.72	9.0 |9.7	*1*	–|–	2.009	6.3 |6.5	–|–	87.3 |83.9	–|–
Anti-Str	2i	66.4	790.13	8.4 |6.7	*0.5*	–|–	3.733	9.9 |15.1	– |–	95.0 |83.6	–|–
Anti-Rnd	3p	66.4	995.30	7.0 |7.2	1.002	6.7 |5.8	1.997	6.5 |6.2	−21.1 |−17.5	78.7 |81.2	−68.4 |−60.6
Anti-Rnd	2c	66.4	995.18	6.9 |7.2	*1*	– |–	1.991	4.7 |5.1	–|–	90.9 |92.1	–|–
Anti-Rnd	2i	66.4	905.53	6.2 |6.6	*0.5*	– |–	2.711	7.8 |8.1	– |–	92.7 |90.0	–|–

Table A4. Results from the simulation experiments presented in Fig. 13 for 10 scenarios over 200 realizations of (spatially merged) Matérn processes (*I*, *II*). For each process the experiments are ordered as labelled (a–j) in Fig. 13. As in Table A3, the columns list the sample means for each Matérn parameter $\overline{\theta_i}$, the relative standard deviation (analytically calculated from eq. 12 at the mean estimate, and empirically obtained from the ensemble) and the parameter correlations from the (analytical | empirical) parameter covariance. All experiment details are as in Fig. 13.

Exp	Obs	$w(\mathbf{x})$	$\overline{\sigma^2}$ units^2	±per cent	$\overline{\nu}$	±per cent	$\overline{\rho}$ km	±per cent	$\{\sigma^2, \nu\}$	$\{\sigma^2, \rho\}$	$\{\nu, \rho\}$
Ia	full	unit	1.14	8.1 |7.3	0.65	2.7 |2.6	1.788	8.0 |7.7	−28.0 |−32.7	90.5 |89.5	−63.5|−68.2
Ib	mix	unit	1.47	19.0 |10.8	0.65	2.6 |2.8	2.436	18.9 |9.6	−41.0 |−16.2	95.9 |76.3	−63.7|−43.0
Ic	mask	unit	0.42	2.4 |2.8	0.62	2.7 |3.7	1.601	6.4 |7.4	−24.7 |−24.7	87.6 |81.6	−65.0 |−67.4
Id	anti	unit	0.73	4.9 |6.8	0.64	2.6 |3.2	1.747	7.7 |10.9	−26.6 |−32.9	89.7 |91.1	−63.6 |−65.0
Ie	sm.mask	unit	0.35	2.1 |2.5	0.66	2.8 |4.2	1.609	6.3 |7.7	−25.9 |−22.9	88.0 |83.8	−65.4 |−67.0
If	sm.anti	unit	0.68	4.8 |7.2	0.66	2.7 |3.4	1.785	8.0 |12.1	−28.5 |−33.4	90.6 |92.1	−63.7|−64.1
Ig	mask	mask	1.13	10.2 |9.2	0.65	4.3 |4.2	1.773	11.2 |10.4	−22.0 |−24.8	87.6 |87.0	−59.8 |−61.7
Ih	anti	anti	1.14	10.2 |9.2	0.65	3.7 |3.4	1.791	11.1 |10.0	−23.9 |−15.0	89.5 |87.9	−60.0|−55.1
Ii	sm.mask	sm.mask	1.13	10.5 |10.0	0.65	4.2 |4.6	1.771	11.2 |11.3	−27.0 |−34.1	88.7 |88.9	−65.2 |−69.0
Ij	sm.anti	sm.anti	1.14	10.9 |10.0	0.65	3.8 |3.5	1.792	11.6 |10.7	−32.4 |−25.0	90.5 |88.6	−66.9 |−63.5
IIa	full	unit	2.84	26.9 |26.3	1.76	3.6 |3.5	2.942	11.3 |11.1	−40.5 |−36.9	89.5 |84.4	−72.0|−74.6
IIb	mix	unit	2.60	25.1 |23.9	0.70	1.2 |1.5	3.890	31.1 |29.0	−46.2 |−9.3	97.0 |91.6	−65.8|−36.1
IIc	mask	unit	0.51	3.4 |2.5	0.74	1.1 |2.8	2.665	14.2 |5.8	−50.1 |−57.3	96.9 |64.2	−69.0 |−45.9
IId	anti	unit	2.38	27.4 |25.1	0.87	1.9 |3.7	4.676	37.2 |30.8	−36.4 |15.2	94.2 |80.1	−65.2 |−24.4
IIe	sm.mask	unit	0.32	2.2 |1.5	1.06	1.6 |4.9	2.314	9.6 |4.9	−58.4 |−63.3	96.5 |66.5	−76.8 |−62.4
IIf	sm.anti	unit	2.55	28.4 |27.4	1.20	2.6 |5.2	4.073	24.5 |22.5	−35.5 |9.2	90.9 |76.0	−70.4|−38.3
IIg	mask	mask	2.78	22.9 |20.8	1.81	18.5 |17.4	2.933	13.4 |12.4	−16.0 |−4.8	74.4 |67.9	−72.1 |−67.0
IIh	anti	anti	2.76	18.4 |19.2	1.80	23.3 |22.7	2.977	16.4 |15.2	−18.0 |10.8	53.3 |41.9	−84.4|−69.5
IIi	sm.mask	sm.mask	2.82	30.2 |31.7	1.80	13.9 |12.1	2.918	14.5 |14.5	−4.1 |1.8	85.1 |85.8	−51.4 |−42.7
IIj	sm.anti	sm.anti	2.75	20.2 |21.7	1.83	21.8 |20.6	2.927	15.8 |14.7	0.9 |16.7	51.1 |50.2	−76.3 |−61.8

© The Author(s) 2026. Published by Oxford University Press on behalf of The Royal Astronomical Society. All rights reserved. For commercial re-use, please contact reprints@oup.com for reprints and translation rights for reprints. All other permissions can be obtained through our RightsLink service via the Permissions link on the article page on our site-for further information please contact journals.permissions@oup.com