Smoothing windows for the synthesis of Gaussian stationary random fields using circulant matrix embedding [*†‡]

Hannes Helgason

KTH Royal Institute of Technology, Stockholm

Vladas Pipiras

University of North Carolina, Chapel Hill

Patrice Abry

École Normale Supérieure, Lyon

June 8, 2012

Abstract

When generating Gaussian stationary random fields, a standard method based on circulant matrix embedding usually fails because some of the associated eigenvalues are negative. The eigenvalues can be shown to be nonnegative in the limit of increasing sample size. Computationally feasible large sample sizes, however, rarely lead to nonnegative eigenvalues. Another solution is to extend suitably the covariance

[*]Hannes Helgason would like to thank Department of Statistics and Operations Research, UNC, for its hospitality during his visit there when this work was undertaken.

[†]*AMS Subject classification.* Primary: 62M10, 62M40. Secondary: 60G10, 60G15, 60G60

[‡]*Keywords and phrases*: numerical synthesis, Gaussian fields, circulant matrix embedding, stationarity.

function of interest so that the eigenvalues of the embedded circulant matrix become nonnegative in theory. Though such extensions have been found for a number of examples of stationary fields, the method depends on nontrivial constructions in specific cases.

In this work, the embedded circulant matrix is smoothed at the boundary by using a cutoff window or overlapping windows over a transition region. The windows are not specific to particular examples of stationary fields. The resulting method modifies the standard circulant embedding, and is easy to use. It is shown that this straightforward approach works for many examples of interest, with the overlapping windows performing consistently better. The method even outperforms in the cases where extending covariance leads to nonnegative eigenvalues in theory, in the sense that the transition region is considerably smaller. The MATLAB code implementing the method is publicly available at `www.hermir.org`.

1 Introduction

We are interested in simulation of Gaussian stationary random fields by using circulant matrix embedding. The focus is on two-dimensional fields throughout, though our ideas extend easily to a general dimension (and possible to vector random fields). Circulant matrix embedding is one of the most popular and superior methods to generate Gaussian stationary univariate and multivariate series, and random fields. In one dimension, the method seems to go back to Davies and Harte (1987). It has been studied more extensively in Dembo et al. (1989), Dietrich and Newsam (1993, 1997), Gneiting (2000), Stein (2001), Craigmile (2003), and Percival (2006). Extension to multivariate (vector) series is considered in Chan and Wood (1999), Helgason et al. (2011), and that to random fields in Dietrich and Newsam (1993), Wood and Chan (1994, 1997), Stein (2002, 2012), Gneiting et al. (2006). Generation of Gaussian stationary fields, in particular, is of interest in a range of scientific disciplines, most notably in Geostatistics (e.g., Chilès and Delfiner (1999)).

We consider the generation of Gaussian stationary random *fields* because available approaches based on circulant matrix embedding are not completely satisfactory. In the *standard* circulant matrix embedding, as considered in Wood and Chan (1994, 1997), the covariance matrix of interest is embedded into a circulant matrix. The eigenvalues of the circulant matrix can be computed efficiently using FFT. When all of them are nonnegative, a Gaussian field can be generated easily having the circulant matrix for the covariance. Its subfield is then the desired random field. It is known, however, that the eigenvalues are rarely all nonnegative in practice. The eigenvalues can be shown to be nonnegative in the limit of increasing sample size. For computationally feasible large sample sizes, however, some of the eigenvalues remain negative, and the standard method continues to fail.

The *cutoff* circulant embedding was suggested in Stein (2002), Gneiting et al. (2006), as a way to overcome the problem of negative eigenvalues. It consists of extending the covariance function of interest so that its circulant embedding matrix has *always* nonnegative eigenvalues in theory. Such extensions were found for several classes of stationary random fields, for example, the powered exponential. The difficulty with this approach is that constructing such extensions depends on particular classes of stationary random fields at hand, and is generally nontrivial.

In this work, we propose a new approach, called *smoothing* circulant embedding. One reason the negative eigenvalues arise with the standard embedding is that the circulant embedding matrix is not "smooth" at the boundary (of periodization). The basic idea is to increase the circulant embedding matrix over a transition region by smoothing it at the boundary. We define and investigate two types of smoothing based on: *nonoverlapping windows* and *overlapping windows*. Smoothing circulant embedding with nonoverlapping windows is similar to cutoff embedding in the sense that it can be viewed as extending the original covariance function to zero and then performing the circulant embedding. The difference is that the extension is no longer specific to the covariance function in question. On the other hand, the embedding with overlapping windows suitably smooths the circulant

embedding matrix directly.

We show numerically on a number of classes of stationary random fields that the proposed method works very well. The surprising finding is that when using overlapping windows, the method actually greatly outperforms the cutoff embedding for which the eigenvalues are nonnegative in theory. For nonoverlapping windows, we generally find our method performing similarly to the cutoff embedding. These findings suggest that the exact form of the extension in the cutoff embedding may not be that important, and that what may really matter is to have smoothness at the boundary.

We do not provide a theoretical justification for our method (though some discussion is given in Section 5). Nevertheless, we find the method appealing for a number of reasons: it works very well (especially with overlapping windows); it outperforms available alternatives (even cutoff embedding); it does not depend on the stationary covariance of interest; it is easy to implement (by modifying the standard circulant embedding). The MATLAB code implementing our method is publicly available at `www.hermir.org`.

The structure of the paper is as follows. Some preliminaries, and available standard and cutoff circulant embeddings are discussed in Section 2. The proposed smoothing embedding can be found in Section 3. The performance of the proposed method is examined on several stationary random fields in Section 4. The method is further discussed in Section 5.

2 Available circulant embeddings

Consider a random field $X = \{X_n, n \in \mathbb{Z}^2\} = \{X_{(n_1,n_2)}, n_1, n_2 \in \mathbb{Z}\}$, where random variables $X_{(n_1,n_2)}$ are real-valued. Suppose the random field X is Gaussian, that is, any of its finite-dimensional distributions $(X_{n_1}, \ldots, X_{n_\ell})$, $n_1, \ldots, n_\ell \in \mathbb{Z}^2$, is multivariate Gaussian. Suppose also that the random field X is stationary. Taking the mean of X zero without loss of generality, this means that the autocovariance

$$r(h) = r(h_1, h_2) = EX_0X_h = EX_nX_{n+h} \tag{2.1}$$

does not depend on $n \in \mathbb{Z}^2$. Note, in particular, that stationarity implies

$$r(h) = r(-h), \qquad h \in \mathbb{Z}^2. \tag{2.2}$$

We are interested in synthesis of the Gaussian stationary field X. More precisely, suppose the goal is to generate the field X on the *square* grid

$$G(N) = \{n \in \mathbb{Z}^d : 0 \le n_1, n_2 \le N-1\}, \tag{2.3}$$

that is, to generate

$$X_n = X_{(n_1,n_2)}, \qquad n \in G(N). \tag{2.4}$$

(Though we consider a square grid, the methods described below extend easily to rectangular grids.) In this section, we recall two synthesis methods based on circulant matrix embedding: the standard embedding (Section 2.1) and the cutoff embedding (Section 2.2).

2.1 Standard embedding

Set

$$M = 2N - 1, \tag{2.5}$$

where N is the (side) size of the grid $G(N)$ in (2.3). M will stand for the (side) size of a larger, embedding grid $G(M)$ where an embedding field will be generated. Set also

$$\widetilde{r}(n) = r(\xi(n)), \qquad n \in G(M), \tag{2.6}$$

where $\xi(n) = (\xi_1(n_1), \xi_2(n_2))$ is defined by

$$\xi_1(u) = \xi_2(u) = \begin{cases} u, & \text{if } 0 \le u \le N-1, \\ u - M, & \text{if } N-1 < u < M-1. \end{cases} \tag{2.7}$$

Extend $\widetilde{r}$ periodically by period M in both dimensions. Note that $\widetilde{r}$ can also be defined as the function which is M-periodic in both dimensions, and satisfies

$$\widetilde{r}(n) = r(\xi(n)), \quad \text{with } \xi(n) = n, \quad n \in \widetilde{G}(N), \tag{2.8}$$

where

$$\widetilde{G}(N) = \{n \in \mathbb{Z}^2 : -(N-1) \le n_1, n_2 \le N-1\}. \tag{2.9}$$

One can describe the *standard circulant embedding* (*SCE*, for short) by using either matrices or operators. We adopt a more convenient operator viewpoint. Consider a circular convolution operator $\widetilde{\Sigma}$ acting on vectors $v(n)$, $n \in G(M)$, and defined by

$$\widetilde{\Sigma}v(n) = \sum_{m \in G(M)} \widetilde{r}(m-n)v(m), \qquad n \in G(M). \tag{2.10}$$

The two-dimensional DFT basis $\{e^{-\imath 2\pi k\cdot(n/M)},\ n \in G(M)\}$, $k \in G(M)$, where $n/M = (n_1/M, n_2/M)$ and $\cdot$ is the usual inner product, diagonalizes the operator $\widetilde{\Sigma}$. Let

$$\lambda_k = \sum_{n \in G(M)} \widetilde{r}(n)e^{-\imath 2\pi k\cdot(n/M)} = \sum_{n \in \widetilde{G}(M)} \widetilde{r}(n)e^{-\imath 2\pi k\cdot(n/M)}, \qquad k \in G(M), \tag{2.11}$$

be the eigenvalues of the operator $\widetilde{\Sigma}$, computed using the 2-dimensional FFT.

Supposing that

$$\lambda_k \ge 0, \qquad k \in G(M), \tag{2.12}$$

the SCE method uses FFT to compute variables

$$\widetilde{X}_n = M^{-1} \sum_{k \in G(M)} \lambda_k^{1/2}(Z_k^0 + \imath Z_k^1)e^{-\imath 2\pi n\cdot(k/M)}, \qquad n \in G(M), \tag{2.13}$$

where Z_k^0, Z_k^1, $k \in G(M)$, are independent standard Gaussian variables. One can show that $\{\Re(\widetilde{X}_n), n \in G(M)\}$ and $\{\Im(\widetilde{X}_n), n \in G(M)\}$ are independent random fields with the

covariance structure

$$E\Re(\widetilde{X}_n)\Re(\widetilde{X}_{n+h}) = E\Im(\widetilde{X}_n)\Im(\widetilde{X}_{n+h}) = \widetilde{r}(h), \quad n, n+h \in G(M). \tag{2.14}$$

Since $\widetilde{r}(n) = r(n)$, for $n \in \widetilde{G}(N)$, taking $X_n = \Re(\widetilde{X}_n)$ or $X_n = \Im(\widetilde{X}_n)$, $n \in G(N)$, then ensures the desired covariance structure.

It is often the case that the condition (2.12) does not hold (for a number of $k \in G(M)$) and hence the SCE method fails. One common way to try to make it work is to increase N to some $\widetilde{N}$. It is convenient to think of $G(\widetilde{N}) \setminus G(N)$ as the *transition region*. Increasing N to $\widetilde{N}$ in SCE can thus be thought as using the covariance function $r(n)$ for extension over the transition region. Since increasing N to any computationally feasible $\widetilde{N}$ often does not ensure (2.12), it is natural to consider other ways to extend $r(n)$ over the transition region $G(\widetilde{N}) \setminus G(N)$. The rest of the methods discussed in this work (Sections 2.2 and 3 below) propose different extension schemes.

2.2 Cutoff embedding

The *cutoff circulant embedding* method (*CCE*, for short) considers two extension schemes over the transition region. Let $\|n\|_2 = (n_1^2 + n_2^2)^{1/2}$. In one CCE extension (Gneiting et al. (2006)), following (2.8), the extension $\widetilde{r}$ is defined as

$$\widetilde{r}(n) = \begin{cases} r(n), & \text{if } 0 \le \|n\|_2 \le \sqrt{2}N, \\ b_1(a_1 - \|n\|_2^{1/2}), & \text{if } \sqrt{2}N \le \|n\|_2 \le \widetilde{N} - 1. \end{cases} \tag{2.15}$$

This extension is used with isotropic covariances $r(n) = \psi(\|n\|_2)$. The constants a_1, b_1, and $\widetilde{N}$ are chosen as

$$a_1 = (\sqrt{2}N)^{1/2} - (\sqrt{2}N)^{-1/2}\frac{1}{2}\frac{\psi(\sqrt{2}N)}{\psi'(\sqrt{2}N)}, \qquad b_1 = -2(\sqrt{2}N)^{1/2}\psi'(\sqrt{2}N), \qquad \widetilde{N} = [a_1^2], \tag{2.16}$$

where $[x]$ denotes the integer part of x. The choices (2.16) ensure that the function $f(t)$ with $f(t) = \psi(t)$, $0 < t \leq \sqrt{2}N$, and $f(t) = b_1(a_1 - t^{1/2})$, $t > \sqrt{2}N$, is once continuously differentiable at $t = \sqrt{2}N$, and satisfies $f(a_1^2) = 0$. The choice of $\sqrt{2}N$ in (2.15) is to have $\widetilde{r}(n) = r(n)$ for $n \in \widetilde{G}(N)$.

In another CCE extension, the embedding $\widetilde{r}$ is defined as

$$\widetilde{r}(n) = \begin{cases} r(n), & \text{if } 0 \leq \|n\|_2 \leq \sqrt{2}N, \\ b_2(a_2 - \|n\|_2)^2, & \text{if } \sqrt{2}N \leq \|n\|_2 \leq \widetilde{N} - 1, \end{cases} \tag{2.17}$$

where, assuming isotropic covariances $r(n) = \psi(\|n\|_2)$,

$$a_2 = \sqrt{2}N - 2\frac{\psi(\sqrt{2}N)}{\psi'(\sqrt{2}N)}, \qquad b_2 = \frac{(\psi'(\sqrt{2}N))^2}{4\psi(\sqrt{2}N)}, \qquad \widetilde{N} = [a_2]. \tag{2.18}$$

After extending r into $\widetilde{r}$ as in (2.15) or (2.17), the rest of CCE is the same as the algorithm in (2.10)–(2.13) with $M = 2\widetilde{N}-1$. As shown in Gneiting et al. (2006), Theorems 1 and 2, the extensions (2.15) and (2.17) ensure that the condition (2.12) holds for a number of classes of covariance models.

3 Proposed smoothing embedding

We propose here two methods to extend the covariance function over the transition region (see the end of Section 2.1). The basic idea is that SCE method fails partly because of the "discontinuities" along the extension boundaries, that is, at $n_1 = N - 1$ or $n_2 = N - 1$ (see, for example, Figure 4 below). To "smooth" the covariance at the boundary, we shall use smoothing windows and call the resulting method the *smoothing window circulant embedding* (*SWCE*, for short).

We first define what we mean by smoothing windows. Let $L > 0$ and consider a rectangle

$$A(L) = [-L, L] \times [-L, L] \subset \mathbb{R}^2. \tag{3.1}$$

Consider the function $\varphi(x)$, $x \in \mathbb{R}^2$, defined by

$$\varphi(x) = \begin{cases} Ce^{-\frac{1}{1-\|x\|_2^2}}, & \text{if } \|x\|_2 < 1, \\ 0, & \text{if } \|x\|_2 \geq 1, \end{cases} \tag{3.2}$$

where $C > 0$ is such that $\int_{\mathbb{R}^2} \varphi(x)dx = 1$. The function φ is compactly supported, nonnegative, symmetric, and infinitely differentiable. For $K > 0$, define

$$\varphi_K(x) = K^{-2}\varphi(K^{-1}x). \tag{3.3}$$

Our smoothing windows are defined as

$$w(x) = \left(1_{A(L)} * \varphi_K\right)(x) = \int_{\mathbb{R}^2} 1_{A(L)}(y)\varphi_K(x-y)dy = \int_{A(L)} \varphi_K(x-y)dy, \tag{3.4}$$

where 1_A is an indicator function of a set A. The smoothing window $w(x)$ is a continuous and smooth function satisfying

$$w(x) = \begin{cases} 1, & \text{if } x \in A(L-K), \\ \rho(x), & \text{if } x \in A(L+K) \setminus A(L-K), \\ 0, & \text{if } x \in \mathbb{R}^2 \setminus A(L+K), \end{cases} \tag{3.5}$$

where $\rho(x)$ decays smoothly from 1 to 0 when moving from the boundary of $A(L-K)$ towards that of $A(L+K)$.

3.1 Non-overlapping windows

Non-overlapping SWCE works as the SCE method in Section 2.1 with the following modifications. The transition region is taken as at the end of Section 2.1 determined by $\widetilde{N} \geq N$, and the embedding size is set to $M = 2\widetilde{N} - 1$ instead of (2.5). Instead of the embedding (2.8), we set

$$\widetilde{r}(n) = r(n)w(n), \qquad n \in \widetilde{G}(\widetilde{N}), \tag{3.6}$$

where w is the smoothing window in (3.4)–(3.5) with

$$L - K = N, \qquad L + K = \widetilde{N}, \tag{3.7}$$

that is, $L = (N + \widetilde{N})/2$ and $K = (\widetilde{N} - N)/2$. Since $w(n) = 1$ for $n \in \widetilde{G}(N)$, we have $\widetilde{r}(n) = r(n)$ for $n \in \widetilde{G}(N)$. The rest of the algorithm is the same as SCE in Section 2.1. The basic idea behind (3.6) is that $w(n)$ smooths $\widetilde{r}(n)$ at the extension boundary by forcing it to zero.

Note that (3.6) is the embedding on $\widetilde{G}(\widetilde{N})$, which is then periodized in both directions. It can equivalently be defined on $G(M)$ and then periodized by setting

$$\begin{aligned}\widetilde{r}(n) =& r(n_1, n_2)w_1(n_1, n_2) + r(n_1, n_2 - M)w_2(n_1, n_2) \\ &+ r(n_1 - M, n_2 - M)w_3(n_1, n_2) + r(n_1 - M, n_2)w_4(n_1, n_2), \quad n \in G(M),\end{aligned} \tag{3.8}$$

where

$$\begin{aligned} w_1(n_1, n_2) &= w(n_1, n_2), & w_2(n_1, n_2) &= w(n_1, n_2 - M), \\ w_3(n_1, n_2) &= w(n_1 - M, n_2 - M), & w_4(n_1, n_2) &= w(n_1 - M, n_2). \end{aligned} \tag{3.9}$$

To see the equivalence, note that by (3.7) only one of the $w_i(n)$'s can be different from zero; see also the left plot on Figure 1 for an illustration of the sections $w_1(n_1, 0)$ and

$w_2(n_1, 0)$ for $0 \le n_1 \le M-1$. The representation (3.8)–(3.9) is also that used in computer implementation of the algorithm.

3.2 Overlapping windows

In the overlapping SWCE, we will smooth the embedding $\widetilde{r}(n)$ at the boundary without forcing it to be zero as with nonoverlapping SWCE. The basic idea is to use the embedding given by (3.8)–(3.9) but choosing a larger value of $L+K$, so that the windows $w_i(n)$'s now start to overlap. More specifically, consider the embedding (3.8)–(3.9) with

$$L - K = N \quad \text{and} \quad L + K = \widetilde{N} + (\widetilde{N} - N) = 2\widetilde{N} - N, \tag{3.10}$$

that is, $L = \widetilde{N}$ and $K = \widetilde{N} - N$; the right plot on Figure 1 illustrates the sections $w_1(n_1, 0)$ and $w_2(n_1, 0)$ for $0 \le n_1 \le M-1$. The fact that $L + K > \widetilde{N}$ in (3.10) ensures the overlap of $w_i(n)$'s. On the other hand, note by (3.7) that $w(n) = w(n_1, n_2) = 0$ when $n_1 \ge L + K = 2\widetilde{N} - N = M - N - 1$ or $n_2 \ge M - N - 1$. This implies that $w_2(n) = w_3(n) = w_4(n) = 0$ for $n \in G(N)$. Since $w_1(n) = 1$ for $n \in G(N)$, it follows that $\widetilde{r}(n) = r(n)$ for $n \in G(N)$. One can argue similarly that more generally, under (3.10), $\widetilde{r}(n) = r(n)$ for $n \in \widetilde{G}(N)$. The rest of the algorithm is the same as SCE in Section 2.1 using the embedding (3.8)–(3.9) with (3.10).

Remark: Under (3.10), we do not believe that (3.8)–(3.9) can be equivalently expressed in a simple form on $\widetilde{G}(\widetilde{N})$ as (3.6) under (3.7).

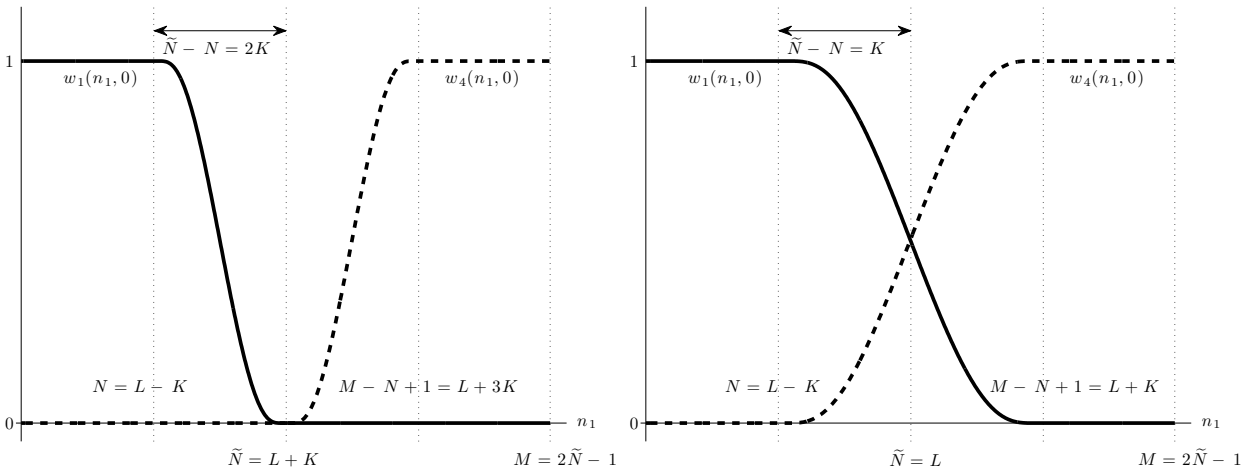

Figure 1: Diagrams illustrating windows $w_1(n)$ (solid line) and $w_4(n)$ (broken line) from Sections 3.1-3.2. Left plot corresponds to nonoverlapping SWCE; right plot corresponds to overlapping SWCE. Both plots show the section corresponding to $n = (n_1, n_2)$ with $n_2 = 0$.

4 Examples and simulation studies

4.1 Smoothing embeddings versus cutoff embedding

We compare here the performance of the proposed SWCE procedures from Section 3 with CCE procedures from Section 2.2. In the comparisons, we consider the following question: For an embedding grid $G(M)$ of fixed size $M = 2\widetilde{N} - 1$ and a given circulant embedding procedure, what is the largest grid $G(N)$ where one can generate the field with a targeted covariance $r(n)$, $n \in \widetilde{G}(N)$? Denote the size of the largest grid by N^* and define the *efficiency* of the embedding as

$$\gamma = \gamma(M) = \frac{2N^* - 1}{M}. \tag{4.1}$$

The efficiency γ satisfies $0 < \gamma \leq 1$ and the closer γ is to 1, the larger the grid $G(N)$ one can synthesize the field on; $\gamma = 1$ corresponds to SCE, hence where the transition region $G(\widetilde{N}) \setminus G(N)$ is empty.

For the CCE embeddings in (2.15)–(2.18), the parameter choices $\widetilde{N} = [a_1^2]$ and $\widetilde{N} = [a_2]$ will determine the grid size N for a fixed embedding size $M = 2\widetilde{N} - 1$. The purpose of this

choice is to ensure that the extension $\widetilde{r}(n)$ reaches zero as $\|n\|_2$ approaches $\widetilde{N}$. We do not impose this behavior in the numerical comparisons, and allow N and the parameters a_j, b_j, $j = 1, 2$, to vary independent of the value of $\widetilde{N}$. More specifically, we use the first two equations in (2.16) and (2.18) for determining the parameters a_j, b_j, $j = 1, 2$, as functions of N (this makes the cutoff embedding smooth when $\|n\|_2$ crosses $\sqrt{2}N$). Then, for a fixed value of $\widetilde{N}$, we find the largest N such that the cutoff embeddings with these parameter choices succeed. (Note that the efficiency of these methods is always bounded from above by $\sqrt{2}/2 \approx 0.71$, where the bound is achieved when $\sqrt{2}N \approx \widetilde{N} - 1$.) Allowing N, a_j, and b_j, to vary independently of $\widetilde{N}$ can only increase efficiency. This is also helpful when the theorems established in Gneiting et al. (2006) do not guarantee that CCE satisfies condition (2.12), such as with several of the examples considered below.

For the comparison we consider isotropic covariance functions of the powered exponential form

$$r(n) = e^{-(\theta\|n\|_2)^\alpha}, \qquad 0 < \alpha \leq 2, \quad \theta > 0, \tag{4.2}$$

and isotropic Cauchy covariance functions of the form

$$r(n) = (1 + (\theta\|n\|_{\ell_2})^\alpha)^{-\beta/\alpha}, \qquad 0 < \alpha \leq 2, \quad \beta > 0, \quad \theta > 0. \tag{4.3}$$

The size of the embedding grid $G(M)$ is kept fixed at $M = 2049$.

Figure 2 compares embedding efficiency γ for powered exponential covariance (4.2) as a function of $M\theta$ (this choice, rather than θ, is just to have a nicer scale for the values on the x-axis). The comparisons were done for $\alpha = 0.5, 1.0, 1.5$ and $\theta = k \cdot 5 \cdot 10^{-4}$, $k = 1, \ldots, 14$. Note that for covariance of the form (4.2), the conditions of Theorem 1 in Gneiting et al. (2006) (concerning embedding (2.15)) hold if and only if $\alpha \leq 0.5$; the conditions of Theorem 2 (concerning embedding (2.17)) hold if and only if $\alpha \leq 1$. These theorems give only sufficient conditions so that as considered here, the embeddings might work outside these ranges.

The points on the curves where the SWCE methods give $\gamma = 1$ is where SCE succeded (i.e. the transition region is empty). In all other cases, SWCE based on the overlapping windows was dramatically better than all of the other embeddings. The difference is greatest for $\alpha = 0.5$, where it has efficiency close to 90% while the other methods have efficiency ranging from 5% to 40%; note that the embedding based on (2.15) fails for both $\alpha = 1.0$ and $\alpha = 1.5$. The proposed embedding using nonoverlapping windows performs well when compared to the parametric cutoff embeddings; for $\alpha = 1.0$, it performs significantly better.

Figure 3 compares embedding efficiency γ for Cauchy covariance (4.3) as a function of β (again, the points where the SWCE methods give $\gamma = 1$ is where SCE succeded). In all the comparisons shown here we took $\theta = 0.001$. The three plots in the figure correspond to the parameter values $\alpha = 0.5, 1.0, 1.5$ respectively; the parameter β is varied. Similar conclusions can be drawn as in the case of powered exponential covariance: SWCE with overlapping windows is superior to all of the other embeddings.

If the goal was to generate fields in quarter-circle domains $\{n : \|n\|_{\ell_2} \leq N - 1, n_1 \geq 0, n_2 \geq 0\}$ instead of square grids $G(N)$, then the efficiency of the CCE procedures would improve by a multiplicative factor $\sqrt{2}$; the efficiency of the proposed smoothing procedures would stay the same. Note that even in this case, smoothing using overlapping windows would perform dramatically better than the other methods.

4.2 Anisotropic covariance

The proposed methods from Section 3 do not need any specific considerations for dealing with anisotropic covariances. Here we give one example demonstrating that the overlapping SWCE has good efficiency in case of an anisotropic covariance of the powered exponential form

$$r(n) = e^{-(\theta\|n\|_W)^\alpha}, \quad 0 < \alpha \leq 2, \quad \theta > 0, \tag{4.4}$$

where $\|t\|_W := \sqrt{t'Wt}$ for column vectors t and W is a nonnegative definite 2×2 matrix. In the example presented here we take

$$W = \begin{pmatrix} 1 & 1 \\ 1 & 2 \end{pmatrix}, \qquad \theta = 0.01, \qquad \alpha = 1; \tag{4.5}$$

a plot of this covariance is shown in Figure 4. For embedding size $M = 1025$, the largest grid size is $N^* = 459$ and the corresponding efficiency is $\gamma \approx 0.89$. Figure 4 shows SCE with $M = 1025$ and SWCE with overlapping windows for these embedding sizes; notice the sharp discontinuities of the SCE along the vertical line through $n_1 = (M+1)/2$ and the horizontal line through $n_2 = (M+1)/2$ – compare this with the SWCE which has smooth transition in the vicinity of these lines. The figure also shows a realization of the Gaussian field with the targeted covariance on the grid $G(N^*)$ where one can observe correlation in the direction of 45° (this is in accordance with the parameter W in (4.5)). Note that one will would need to take the embedding size at least as large as $M = 3339$ for SCE to succeed for this choice of covariance (SCE breaks down for values smaller than $M = 3339$).

5 Discussion of smoothing embeddings

Here we give some insight into why smoothing embedding can be helpful in cases where SCE fails. To simplify the mathematical notation, we consider the one-dimensional case but the arguments given here can be extended to higher dimensions. We also provide a quantitative explanation of how overlapping windows improve upon nonoverlapping windows.

Consider the case of a Gaussian stationary series X_n, $n \in \mathbb{Z}$, with covariance $r(n)$, $n \in \mathbb{Z}$. The goal is to synthesize X_n for $n = 0, \ldots, N-1$. A one-dimensional version of the nonoverlapping SWCE starts by smoothing the covariance $r(n)$ with a window $w(n)$ satisfying

(i) $w(n)$ is symmetric: $w(-n) = w(n)$ for all n,

(ii) $w(n) = 1$ for $|n| \leq N-1$,

(iii) $w(n)$ decays "smoothly" to zero as $|n|$ approaches $\widetilde{N} > N$, and is zero for $|n| \geq \widetilde{N}$.

Write $M = 2\widetilde{N}$ and let $\widetilde{r}$ be the M-periodic series with

$$\widetilde{r}(n) = r(n)w(n), \qquad -M/2 < n \leq M/2. \tag{5.1}$$

If $\widetilde{r}(n)$ is a positive definite series, hence a valid covariance for circular stationary Gaussian series, then one can use it to synthesize the series $\{X_n, 0 \leq n \leq N-1\}$ efficiently using the FFT algorithm. The series $\widetilde{r}$ being a valid covariance is equivalent to its DFT being nonnegative:

$$DFT(\widetilde{r})(k) = \sum_{n=0}^{M-1} \widetilde{r}(n) e^{-\imath 2\pi \frac{nk}{M}} \geq 0, \qquad 0 \leq k < M. \tag{5.2}$$

The problem of finding conditions for $DFT(\widetilde{r})(k)$ to be nonnegative can be studied through a (2π-periodic) convolution:

$$\begin{aligned} DFT(\widetilde{r})(k) &= \sum_{n=-\widetilde{N}+1}^{\widetilde{N}} \widetilde{r}(n) e^{-\imath 2\pi \frac{nk}{M}} = \sum_{n=-\infty}^{\infty} r(n)w(n)e^{-\imath \omega_k n} \\ &= \frac{1}{2\pi} \int_{-\pi}^{\pi} \widehat{r}(\theta)\widehat{w}(\omega_k - \theta)d\theta =: \widehat{r} * \widehat{w}(\omega_k), \qquad \omega_k = 2\pi k/M, \end{aligned} \tag{5.3}$$

where the discrete-time Fourier transform (DTFT) for series $s(n)$, $n \in \mathbb{Z}$, is denoted by $\widehat{s}(\omega) = \sum_{n=-\infty}^{\infty} s(n)e^{-\imath\omega n}$, $\omega \in (-\pi, \pi]$; in the above we used $r(n)w(n) = 0$ for $|n| \geq \widetilde{N}$ and $\widehat{rw} = \widehat{r} * \widehat{w}$.

For the overlapping SWCE, we take $\widetilde{r}$ as the M-periodic series satisfying (5.1) or equivalently

$$\widetilde{r}(n) = r(n)w(n) + r(n-M)w(n-M), \qquad 0 \leq n \leq M-1, \tag{5.4}$$

with $M = 2\widetilde{N}$. The DFT of $\widetilde{r}$ in (5.4) is of the same form as for the nonoverlapping SWCE (5.1). The difference between nonoverlapping and overlapping SWCEs lies in different choices of window w. In the overlapping SWCE, one can take $w(n)$ and $w(n-M)$ to be

overlapping in the transition region without affecting $\widetilde{r}(n)$ for $|n| \leq N-1$ (i.e., $\widetilde{r}(n) = r(n)$ holds for $|n| \leq N-1$). We will argue next in which sense this allows for a more efficient use of the available transition region (supposing M and N are fixed).

Assume that the window w is constructed in a similar way as the two-dimensional smoothing window in Section 3, that is, by taking

$$w(n) = 1_{A(L)} * \varphi_K(n), \qquad n \in \mathbb{Z}, \qquad A(L) = [-L, L], \tag{5.5}$$

where $\varphi_K(n) = K^{-1}\varphi(n/K)$ with $\varphi(x) = C\exp(-1/(1-|x|^2)$ if $|x| < 1$ and 0 if $|x| \geq 1$; $C > 0$ is such that $\int_{\mathbb{R}} \varphi(x)dx = 1$. (We take discrete rather than continuous convolution in (5.5) for simplicity.) Using a standard property of the Fourier transform for convolutions, we have

$$\widehat{w}(\omega) = \widehat{1_{A(L)}}(\omega)\widehat{\varphi_K}(\omega) = D_L(\omega)\widehat{\varphi_K}(\omega), \tag{5.6}$$

where D_L is the Dirichlet kernel

$$D_L(\omega) := \sum_{n=-L}^{L} e^{-\imath\omega n} = \begin{cases} \frac{\sin(\omega(L+1/2))}{\sin(\omega/2)} & \text{if } \omega \neq 2m\pi \text{ (m an integer)}, \\ 2L+1 & \text{if } \omega = 2m\pi \text{ (m an integer)}. \end{cases} \tag{5.7}$$

We also have

$$\begin{aligned} \widehat{\varphi_K}(\omega) &= \sum_{n=-K}^{K} K^{-1}\varphi(n/K)e^{-\imath\omega n} = \sum_{n=-K}^{K} K^{-1}\varphi(n/K)e^{-\imath K\omega(n/K)} \\ &\approx \mathcal{F}(\varphi)(K\omega), \end{aligned} \tag{5.8}$$

where $\mathcal{F}(f)(\omega) = \int_{\mathbb{R}} f(x)e^{-\imath\omega x}dx$, $\omega \in \mathbb{R}$, denotes the continuous-time Fourier transform of an integrable function $f(x)$, $x \in \mathbb{R}$, and the approximation holds for K large enough. With some abuse of notation, we write $\mathcal{F}(\varphi)(K\omega) = \widehat{\varphi}(K\omega)$ and

$$\widehat{w}(\omega) = D_L(\omega)\widehat{\varphi}(K\omega). \tag{5.9}$$

Consider now the relation $DFT(\widetilde{r})(k) = \widehat{r} * \widehat{w}(\omega)$ in (5.3) in the light of (5.9). In the case of SCE, we simply have $\widehat{w}(\omega) = D_L(\omega)$ with $L = N - 1$. The slow decay of the Dirichlet kernel, and its negative side-lobes near $\omega = 0$, are the main causes which can make $\widehat{r} * \widehat{w}(\omega)$ become negative. One way to overcome this is to take the embedding size M larger, hence increasing L. This will make D_L more concentrated around $\omega = 0$, in hope that $\widehat{r} * D_L(\omega)$ will get closer to $\widehat{r}(\omega) > 0$. But how is SWCE helpful?

Since $\varphi(x)$, $x \in \mathbb{R}$, is infinitely differentiable on $\mathbb{R}$, its Fourier transform $\mathcal{F}(\varphi)(\omega)$ decays to zero faster than any polynomial as ω moves away from zero. The larger K, the faster $\widehat{\varphi}(K\omega)$ decays, suppressing the negative side-lobes of the Dirichlet kernel in $\widehat{w}(\omega) = D_L(\omega)\widehat{\varphi}(K\omega)$, and improving the decay rate. In other words, this makes $\widehat{w}(\omega)$ a better approximate identity (i.e., a 2π-periodic Dirac delta function) than only having the Dirichlet kernel $D_L(\omega)$ in SCE, hence making $\widehat{r} * \widehat{w}(\omega)$ closer to $\widehat{r}(\omega) > 0$.

The argument shows that it is desirable to take L and K as large as possible. Assume the requested domain size N and the embedding size $M = 2\widetilde{N}$ to be fixed. For nonoverlapping SWCE, we have

$$L = \frac{N + \widetilde{N}}{2}, \qquad K = \frac{\widetilde{N} - N}{2}.$$

and for overlapping SWCE, we have

$$L = \widetilde{N}, \qquad K = \widetilde{N} - N.$$

The fact that both parameters L and K can be taken larger for the overlapping SWCE will improve the decay of (5.9). This is the reason why overlapping SWCE performs better than the nonoverlapping SWCE.

References

Chan, G. & Wood, A. (1997), 'An algorithm for simulating stationary Gaussian random fields', *Applied Statistics, Algorithm Section* **46**, 171–181.

Chan, G. & Wood, A. (1999), 'Simulation of stationary Gaussian vector fields', *Statistics and Computing* **9**(4), 265–268.

Chilés, J. & Delfiner, P. (1999), *Geostatistics: Modeling Spatial Uncertainty*, Wiley, New York.

Craigmile, P. F. (2003), 'Simulating a class of stationary Gaussian processes using the Davies-Harte algorithm, with application to long memory processes', *Journal of Time Series Analysis* **24**(5), 505–511.

Davies, R. & Harte, D. (1987), 'Tests for Hurst effect', *Biometrika* **74**(4), 95–101.

Dembo, A., Mallows, C. & Shepp, L. (1989), 'Embedding nonnegative definite Toeplitz matrices in nonnegative definite circulant matrices, with applications to covariance estimation', *IEEE Transactions on Information Theory* **35**, 1206–1212.

Dietrich, C. & Newsam, G. (1993), 'A fast and exact simulation for multidimensional Gaussian stochastic simulations', *Water Resources Research* **29**(8), 2861–2869.

Dietrich, C. & Newsam, G. (1997), 'Fast and exact simulation of stationary Gaussian processes through circulant embedding of the covariance matrix', *SIAM Journal on Scientific Computing* **18**(4), 1088–1107.

Gneiting, T. (2000), 'Power-law correlations, related models for long-range dependence and their simulation', *Journal of Applied Probability* **37**(4), 1104–1109.

Gneiting, T., Ševčíková, H., Percival, D., Schlather, M. & Jiang, Y. (2006), 'Fast and exact simulation of large Gaussian lattice systems in R^2: Exploring the limits', *Journal of Computational and Graphical Statistics* **15**(3), 483–501.

Helgason, H., Pipiras, V. & Abry, P. (2011), 'Fast and exact synthesis of stationary multivariate Gaussian time series using circulant embedding', *Signal Processing* **91**(5), 1123–1133.

Percival, D. (2006), 'Exact simulation of complex-valued Gaussian stationary processes via circulant embedding', *Signal Processing* **86**(7), 1470–1476.

Stein, M. (2001), 'Local stationarity and simulation of self-affine intrinsic random functions', *IEEE Transactions on Information Theory* **47**(4), 1385–1390.

Stein, M. (2002), 'Fast and exact simulation of fractional Brownian surfaces', *Journal of Computational and Graphical Statistics* **11**(3), 587–599.

Stein, M. (2012), 'Simulation of Gaussian random fields with one derivative', *Journal of Computational and Graphical Statistics* **21**(1), 155–173.

Wood, A. & Chan, G. (1994), 'Simulation of stationary Gaussian processes in $[0, 1]^d$', *Journal of Computational and Graphical Statistics* **3**(4), 409–432.

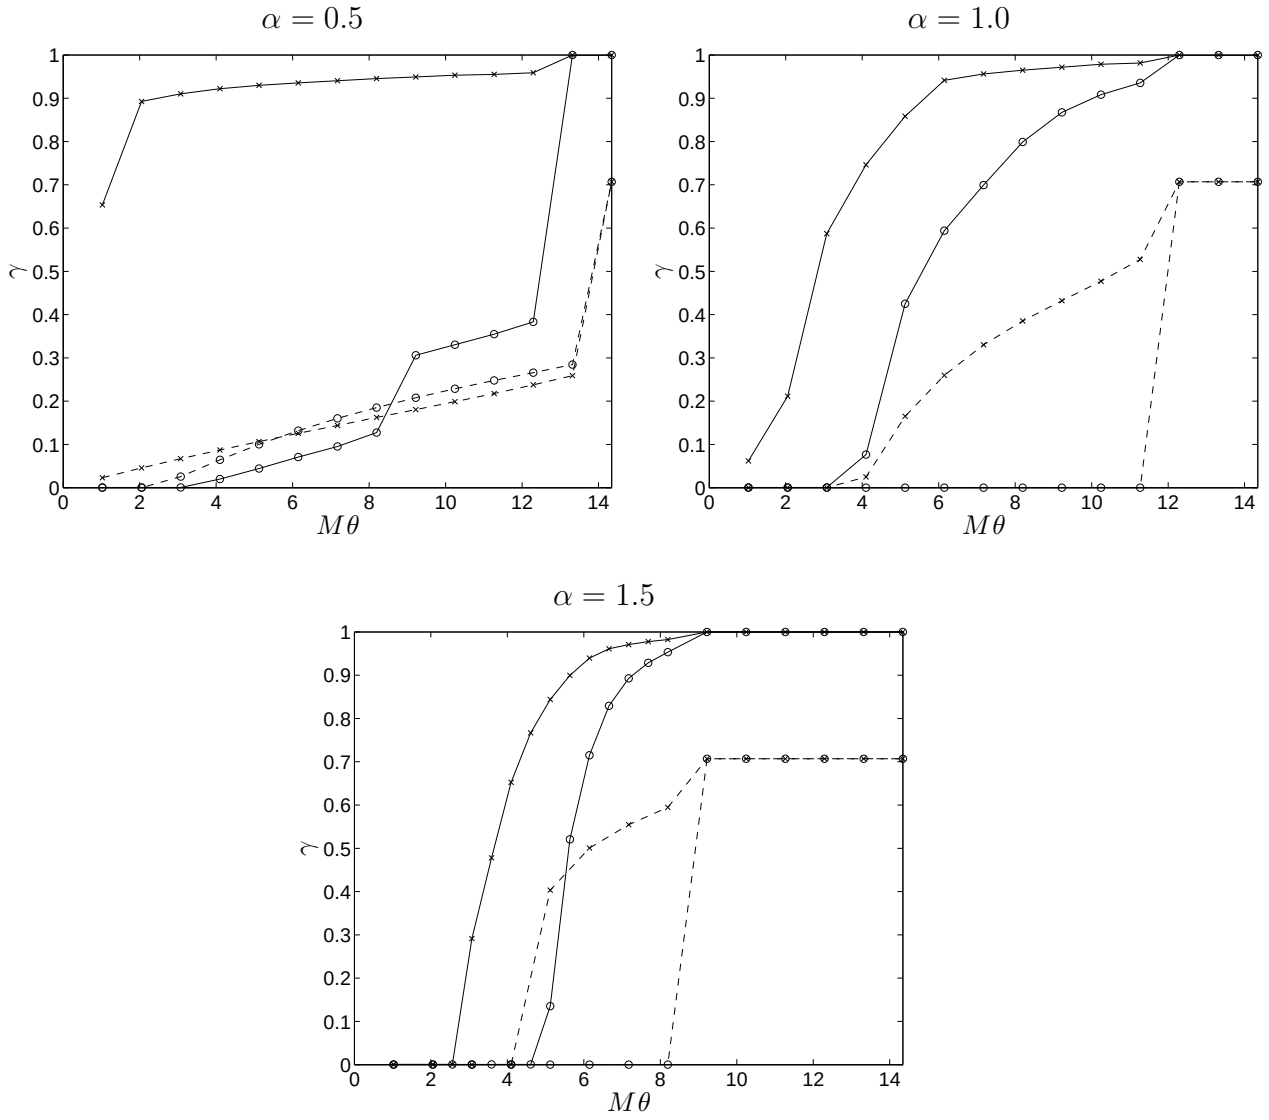

Figure 2: Comparison of embedding efficiency (4.1) for powered exponential covariance (4.2) as a function of $M\theta$ with $M = 2049$. Top left plot ($\alpha = 0.5$); top right plot ($\alpha = 1.0$); bottom plot ($\alpha = 1.5$). SWCE with overlapping windows: solid line with labels $\times$; SWCE with nonoverlapping windows: solid line with labels $\bigcirc$; CCE of type (2.17): broken line with labels $\times$; CCE of type (2.15): broken line with labels $\bigcirc$.

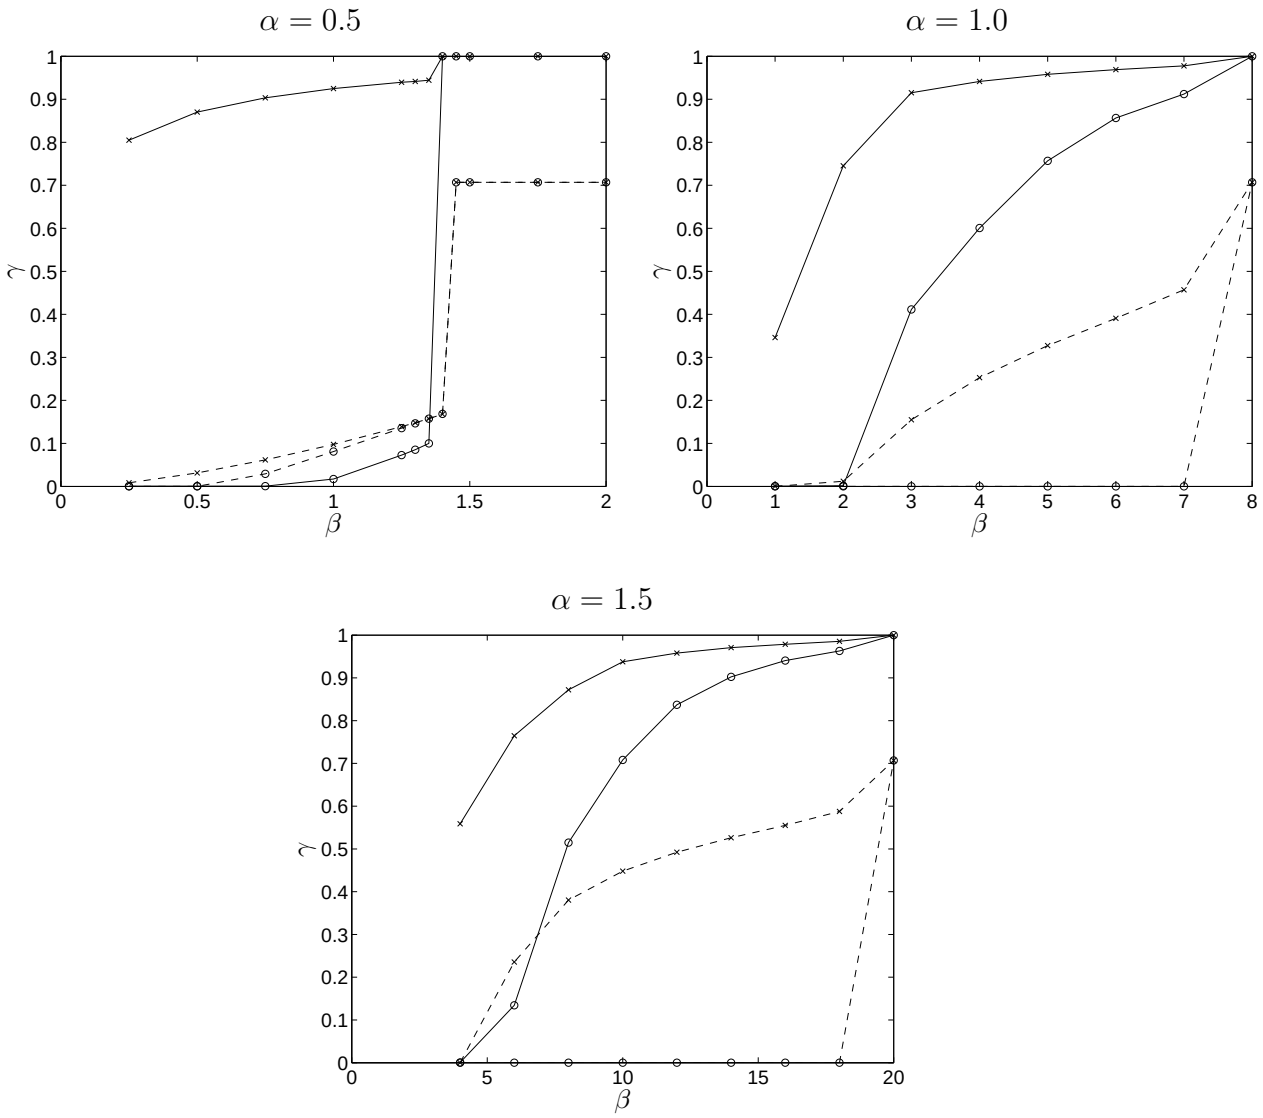

Figure 3: Comparison of embedding efficiency (4.1) for Cauchy covariance (4.3) as a function of β with $M = 2049$ and $\theta = 0.001$. Top left plot ($\alpha = 0.5$); top right plot ($\alpha = 1.0$); bottom plot ($\alpha = 1.5$). SWCE with overlapping windows: solid line with labels $\times$; SWCE with nonoverlapping windows: solid line with labels $\bigcirc$; CCE of type (2.17): broken line with labels $\times$; CCE of type (2.15): broken line with labels $\bigcirc$.

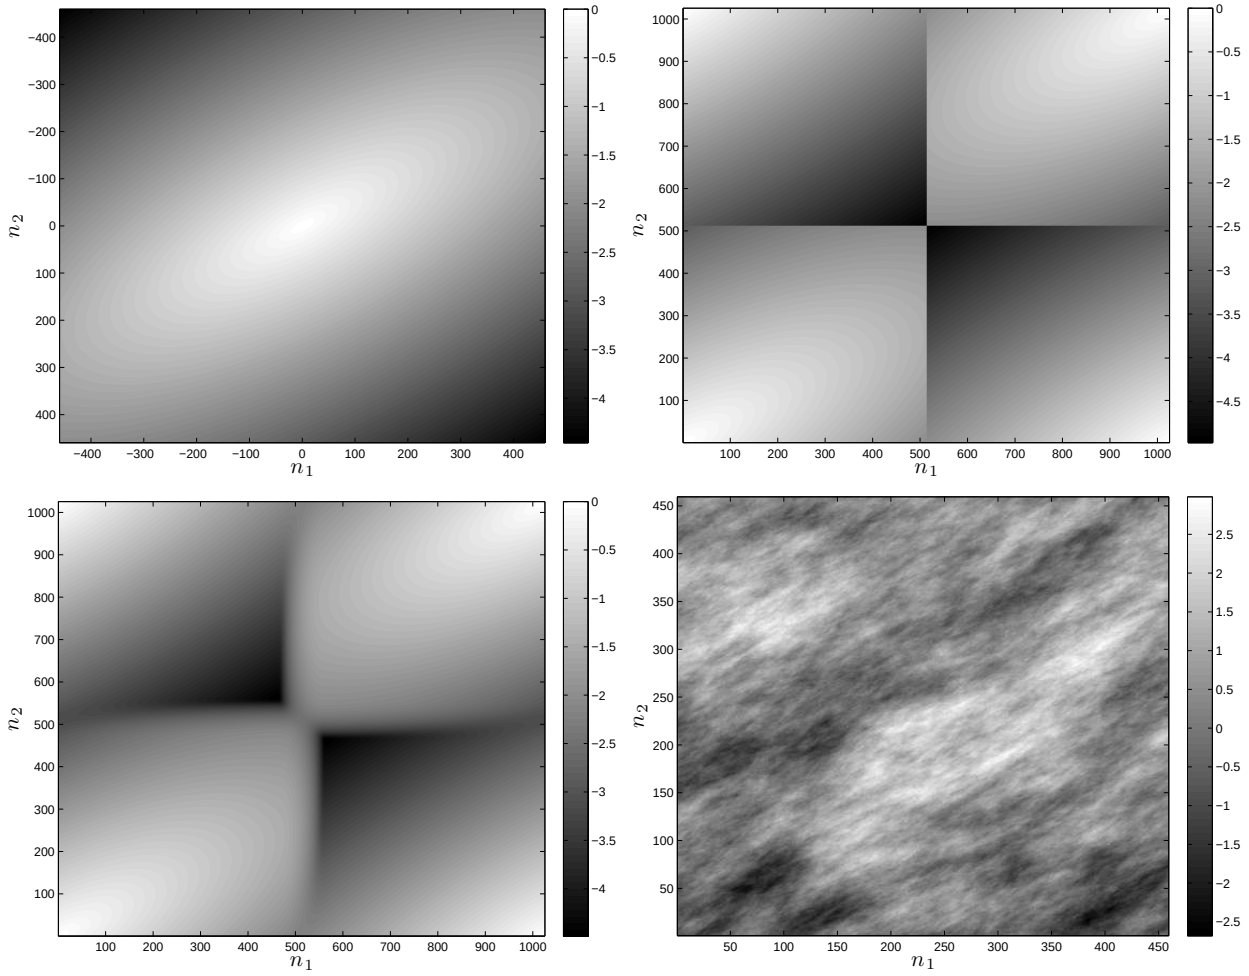

Figure 4: Plots for the anisotropic covariance example in Section 4.2. Top left: Covariance $r(n)$ in (4.4) with parameters (4.5); top right: SCE of $r(n)$ for $M = 1025$; bottom left: SWCE with overlapping windows for $M = 1025$ and $N = 459$; bottom right: realization of the Gaussian field with covariance $r(n)$ on $G(N)$ with $N = 459$. (The values of the covariance and embeddings are on log-10 scale).