Generating Realisations of Stationary Gaussian Random Fields by Circulant Embedding.

Catherine E. Powell, School of Mathematics, University of Manchester, UK.

Random fields are families of random variables, indexed by a d-dimensional parameter $\mathbf{x}$ with $d > 1$. They are important in many applications and are used, for example, to model properties of biological tissue, velocity fields in turbulent flows and permeability coefficients of rocks. Mark 24 of the NAG Fortran library includes new routines for generating realisations of stationary Gaussian random fields using the method of *circulant embedding*. This short note illustrates the main ideas behind circulant embedding and how to use the routines `g05zr` and `g05zs` in the NAG toolbox for MATLAB. The routines `g05zm`, `g05zn` and `g05zp` can also be used to generate realisations of stationary Gaussian *stochastic processes* (the $d = 1$ case).[1]

Random Fields

For a two-dimensional domain $D \subset \mathbb{R}^2$, a (real-valued) *random field* $\{Z(\mathbf{x}) : \mathbf{x} \in D\}$, also written $Z(\mathbf{x}, \omega)$, is a set of real-valued random variables on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. That is, for each $\mathbf{x} \in D$, $Z(\mathbf{x}) : \Omega \to \mathbb{R}$ is a random variable. The random field is *second-order* if $Z(\mathbf{x})$ has finite variance for each $\mathbf{x} \in D$ and for such fields we can define the *mean* function $\mu(\mathbf{x}) = \mathbb{E}\left[Z(\mathbf{x})\right]$ and the *covariance* function

$$C(\mathbf{x}, \mathbf{y}) = \mathrm{Cov}(Z(\mathbf{x}), Z(\mathbf{y})) = \mathbb{E}\left[\Big(Z(\mathbf{x}) - \mu(\mathbf{x})\Big)\Big(Z(\mathbf{y}) - \mu(\mathbf{y})\Big)\right], \qquad \mathbf{x}, \mathbf{y} \in D.$$

Important cases are *stationary* random fields, where $\mu(\mathbf{x})$ is constant and the covariance depends on $\mathbf{x} - \mathbf{y}$ and *isotropic* random fields, where the covariance depends on $\|\mathbf{x} - \mathbf{y}\|_2$.

For a fixed $\omega \in \Omega$, the associated *realisation* of a random field is the deterministic function $f\colon D \to \mathbb{R}$ defined by $f(\mathbf{x}) := Z(\mathbf{x}, \omega)$ for $\mathbf{x} \in D$. Thus, a realisation represents one possibility for the quantity Z as a function of $\mathbf{x}$. Examples are given in Figures 2 and 3, later. If $Z(\mathbf{x}, \omega)$ represents an input for a mathematical model such as a system of PDEs, then we often need to generate multiple realisations of $Z(\mathbf{x}, \omega)$ so that statistical analysis of the solution can be performed (e.g., using Monte Carlo methods).

Gaussian Random Fields

A *Gaussian* random field is a second-order field such that the vector of random variables

$$\mathbf{Z} = [Z(\mathbf{x}_1), Z(\mathbf{x}_2), \ldots, Z(\mathbf{x}_N)]^T$$

follows the multivariate Gaussian distribution for any $\mathbf{x}_1, \ldots, \mathbf{x}_N \in D$. That is, $\mathbf{Z} \sim N(\boldsymbol{\mu}, C)$ where the mean vector $\boldsymbol{\mu}$ and the covariance matrix C have entries $\mu_i = \mu(\mathbf{x}_i)$ and

$$c_{ij} = C(\mathbf{x}_i, \mathbf{x}_j), \qquad i, j = 1, \ldots, N.$$

The matrix C is, by definition, symmetric and nonnegative definite. By choosing N sample points $\mathbf{x}_i$ on D, we can generate discrete realisations of Gaussian random fields $Z(\mathbf{x}, \omega)$ by drawing samples of $\mathbf{Z}$. We focus on the case $\mu(\mathbf{x}) = 0$ so that $\boldsymbol{\mu} = \mathbf{0}$. A useful observation is that a *pair* of independent samples from $N(\mathbf{0}, C)$ can be drawn simultaneously by taking the real and complex parts of $\mathbf{Y} \sim CN(\mathbf{0}, 2C)$ (a sample from the complex Gaussian distribution).

[1]These routines were developed as part of the NAG-sponsored PhD project of Phillip Taylor at the University of Manchester, supervised by Tony Shardlow (University of Bath) and Catherine Powell.

Toeplitz and Circulant Matrices

An $N \times N$ matrix C is *Toeplitz* if the entries along each diagonal are the same. A *circulant* matrix is a Toeplitz matrix for which each column is a circular shift of the elements in the preceding column (so that the last entry becomes the first entry). Consider,

$$\begin{pmatrix} 1 & 2 & 3 \\ 2 & 1 & 2 \\ 3 & 2 & 1 \end{pmatrix}, \qquad \begin{pmatrix} 1 & 2 & 2 \\ 2 & 1 & 2 \\ 2 & 2 & 1 \end{pmatrix}.$$

The matrix on the left is symmetric and Toeplitz while the matrix on the right is symmetric and circulant. Symmetric Toeplitz matrices can always be extended to give symmetric circulant matrices by padding them with extra rows and columns. For example,

$$\begin{pmatrix} 1 & 2 & 3 \\ 2 & 1 & 2 \\ 3 & 2 & 1 \end{pmatrix} \qquad \longrightarrow \qquad \left(\begin{array}{ccc|c} 1 & 2 & 3 & 2 \\ 2 & 1 & 2 & 3 \\ 3 & 2 & 1 & 2 \\ \hline 2 & 3 & 2 & 1 \end{array}\right).$$

Notice that we only need to store the first column of a symmetric Toeplitz or circulant matrix to generate the whole matrix. Similarly, symmetric block Toeplitz matrices with Toeplitz blocks (BTTB matrices) can always be extended to form symmetric block circulant matrices with circulant blocks (BCCB matrices). These can also be stored in a handy reduced format.

Circulant and BCCB matrices can be factorised using discrete Fourier transforms. In particular, any BCCB matrix B has the decomposition $B = W\Lambda W^*$ where W is the two-dimensional Fourier matrix and Λ is the diagonal matrix of eigenvalues. If B is a valid covariance matrix (i.e., symmetric and nonnegative definite), then samples can be drawn from $CN(\mathbf{0}, 2B)$ easily. Since none of the eigenvalues are negative, $\Lambda^{1/2}$ is well defined. We can then compute

$$\mathbf{Y} = W\Lambda^{1/2}\boldsymbol{\xi}, \qquad \text{where } \boldsymbol{\xi} \sim CN(\mathbf{0}, 2I).$$

It is easy to show that $\mathbf{Y} \sim CN(\mathbf{0}, 2B)$ and hence $\mathbf{Z}_1 = \mathrm{Re}(\mathbf{Y})$ and $\mathbf{Z}_2 = \mathrm{Im}(\mathbf{Y})$ are independent $N(\mathbf{0}, B)$ samples. Multiplications with W can be performed by applying discrete Fourier transforms and Λ can be obtained by applying an inverse discrete Fourier transform to the reduced version of B. These observations are the basis of the *circulant embedding* algorithm.

Gaussian Random Fields with Stationary Covariance

In general, the $N \times N$ covariance matrix C associated with a Gaussian random field on a two-dimensional domain D is *not* a BCCB matrix. However, if the sample points $\mathbf{x}_i$ are uniformly spaced and the covariance function is stationary, then C is always a BTTB matrix.

Suppose $D = [x_{min}, x_{max}] \times [y_{min}, y_{max}]$ and divide D into $n_1 \times n_2$ rectangular elements. Choose the $N = n_1 n_2$ sample points to be the mid-points and order these lexicographically. If the covariance function is stationary then $C(\mathbf{x}_i, \mathbf{x}_j) = \gamma(\mathbf{x}_i - \mathbf{x}_j)$ where $\gamma(\mathbf{x})$ is a one-parameter function. Since the spacings in the x and y directions are constant, C is a BTTB matrix with $n_2 \times n_2$ Toeplitz blocks of size $n_1 \times n_1$. Consider the exponential covariance function

$$C(\mathbf{x}, \mathbf{y}) = \sigma^2 \exp\left(-\frac{\|\mathbf{x} - \mathbf{y}\|_2}{\ell}\right) = \gamma(\mathbf{x} - \mathbf{y}),$$

where $\gamma(\mathbf{x}) = \sigma^2 \exp(-\|\mathbf{x}\|_2/\ell)$ and choose the variance and correlation length, respectively, to be $\sigma^2 = 1$ and $\ell = 1$. Figure 1 illustrates the BTTB covariance matrices associated with two distinct uniform grids on $D = [0,1] \times [0,1]$. The Toeplitz structure is clearly visible.

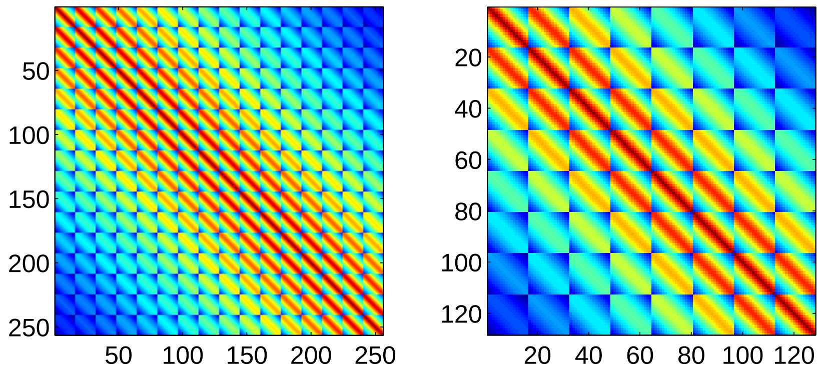

Figure 1: MATLAB `imagesc` plots of the $N \times N$ covariance matrices C associated with $\gamma(\mathbf{x}) = \exp(-\|\mathbf{x}\|_2)$ and a uniform grid of $N = n_1 \times n_2$ points on $D = [0,1] \times [0,1]$ with $n_1 = n_2 = 16$ (left) and $n_1 = 16, n_2 = 8$ (right). The matrix on the left is BTTB with 16×16 Toeplitz blocks of size 16×16. The matrix on the right is BTTB with 8×8 Toeplitz blocks of size 16×16.

Circulant Embedding

Circulant embedding can be used to generate realisations of stationary Gaussian random fields provided the sample points are uniformly spaced. Since the associated $N \times N$ covariance matrix C is BTTB and any BTTB matrix can be embedded inside a larger $M \times M$ BCCB matrix B, then samples of $\mathbf{Z} \sim N(\mathbf{0}, C)$ can be obtained from samples of $\widetilde{\mathbf{Z}} \sim N(\mathbf{0}, B)$ (by simply discarding some entries of the vector). Independent samples of $\widetilde{\mathbf{Z}}$ can be obtained from the real and imaginary parts of $\mathbf{Y} \sim CN(\mathbf{0}, 2B)$ using discrete Fourier transforms, as explained above. The only difficulty is that for a fixed M, the extended BCCB matrix B is not necessarily a valid covariance matrix, so $\Lambda^{1/2}$ may not be well defined. B needs to be nonnegative definite. A check on the eigenvalues reveals whether this is the case. If there are negative eigenvalues, then a *larger* BCCB matrix B must be found that is nonnegative definite.

For some stationary covariance functions $\gamma(\mathbf{x})$, it is not possible to construct a nonnegative definite BCCB embedding matrix B of a feasible size. In that case, we must be content with samples from an *approximate* distribution $N(\mathbf{0}, \widehat{C})$ where $\widehat{C}$ is 'close' to the true covariance matrix C. Given a BCCB extension matrix B, we have

$$B = W\Lambda_+ W^* + W\Lambda_- W^* = B_+ + B_-$$

where $\Lambda = \Lambda_+ + \Lambda_-$ is decomposed into two diagonal matrices. Λ_+ contains the positive eigenvalues of B and has a zero on the diagonal where negative values occur in Λ. Similarly, Λ_- contains the negative eigenvalues. One possibility is to approximate B by B_+. The samples obtained using the BCCB matrix B_+ come from an approximate distribution $N(\mathbf{0}, \widehat{C})$ where the covariance error $\|C - \widehat{C}\|_2$ depends on the size of the neglected negative eigenvalues.

NAG routines g05zr and g05zs

For a library of standard stationary covariance functions $\gamma(\mathbf{x})$ (including the exponential, Gaussian, Whittle-Matérn and Bessel covariance functions), the NAG routine `g05zr` constructs a BCCB extension B of the BTTB covariance matrix C associated with a user-defined grid of

$N = n_1 n_2$ points on $D = [x_{min}, x_{max}] \times [y_{min}, y_{max}]$. It is also possible to work with other stationary covariance functions by using the routine `g05zq` to define a non-standard $\gamma(\mathbf{x})$. The maximum size allowed for the BCCB extension matrix is $M \times M$, where $M = m_1 m_2$ and m_1, m_2 are supplied by the user. The square roots of the eigenvalues of B are returned by `g05zr`. The NAG routine `g05zs` can then be used to produce samples from $N(\mathbf{0}, C)$ if all the eigenvalues are nonnegative or from an approximate distribution $N(\mathbf{0}, \widehat{C})$, otherwise.

Example: Exponential Covariance

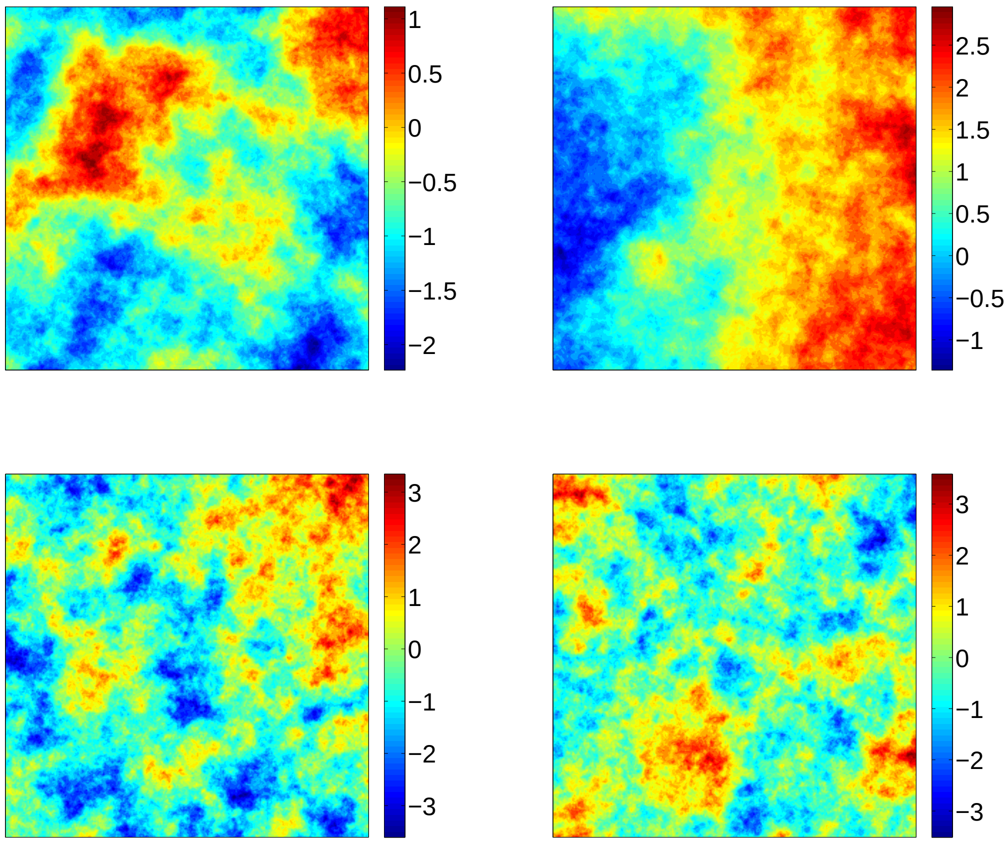

Figure 2: Realisations of a mean zero Gaussian random field on $D = [0,1] \times [0,1]$ with the exponential covariance function $\gamma(\mathbf{x}) = \exp(-\|\mathbf{x}\|_2/\ell)$ with correlation length $\ell = 1$ (top) and $\ell = 1/10$ (bottom).

Let $D = [0,1] \times [0,1]$ and select a grid of $N = 256 \times 256$ uniformly spaced sample points. To generate realisations of a mean zero Gaussian random field with the stationary covariance function $C(\mathbf{x}, \mathbf{y}) = \gamma(\mathbf{x} - \mathbf{y})$ where $\gamma(\mathbf{x})$ is the exponential function

$$\gamma(\mathbf{x}) = \sigma^2 \exp\left(-\frac{\|\mathbf{x}\|_2}{\ell}\right) = \sigma^2 \exp\left(-\sqrt{\left(\frac{x_1}{\ell}\right)^2 + \left(\frac{x_2}{\ell}\right)^2}\right),$$

we must first call `g05zr`. In the NAG toolbox for MATLAB, this can be done by calling the function nag_rand_field_2d_predef_setup. To define the sample points and choose the exponential covariance with $\sigma^2 = 1$ and $\ell = 1$, the required inputs are:

```
>> ns=[int64(256),int64(256)]; xmin=0; xmax=1; ymin=0; ymax=1;
>> var=1; icov2=int64(4); params=[1,1]; norm=int64(2);
```

To define the maximum size of the circulant embedding matrix, we can set, for example,

```
>> maxm=[int64(4096),int64(4096)];
```

This produces a BCCB matrix B which is no larger than $4096^2 \times 4096^2$. Finally, to control how approximate embedding is done, we need to specify icorr. To replace B by B_+, we set

```
>> icorr=int64(2);
```

With all the inputs defined, we can now call g05zs via the MATLAB command

```
>> [lam, xx, yy, m, approx, rho, icount, eig, ifail] = ...
        nag_rand_field_2d_predef_setup(ns, xmin, xmax, ymin, ymax, maxm, ...
                      var, icov2, params, 'icorr', icorr,'norm_p', int64(2));
```

The eigenvalues of B are stored in lam, the grid points in xx and yy and the output approx indicates whether approximation is necessary. In this case, approx=0 and no approximation is needed. Note however that the smallest BCCB embedding matrix that is nonnegative definite is of size $4096^2 \times 4096^2$ so if smaller dimensions are set for maxm, then approximate embedding *is* performed. By specifying the number of realisations s to generate, and the current state of the random number generator, g05zs can be called as follows.

```
>> [state, z, ifail] = nag_rand_field_2d_generate(ns, s, m, lam, rho, state);
```

The resulting realisations of the mean zero Gaussian random field are shown in Figure 2.

Example: Gaussian Covariance

Once again, let $D = [0,1] \times [0,1]$ and select $N = 256 \times 256$ uniformly spaced sample points. Consider $C(\mathbf{x}, \mathbf{y}) = \gamma(\mathbf{x} - \mathbf{y})$ where $\gamma(\mathbf{x})$ is the so-called Gaussian function

$$\gamma(\mathbf{x}) = \sigma^2 \exp\left(-\frac{\|\mathbf{x}\|_2^2}{\ell^2}\right) = \sigma^2 \exp\left(-\left(\frac{x_1}{\ell}\right)^2 - \left(\frac{x_2}{\ell}\right)^2\right).$$

To define this covariance function with $\sigma^2 = 1$ and $\ell^2 = 1/10$, the inputs for g05zr are:

```
>> var=1; icov2=int64(5); params=[1/sqrt(10),1/sqrt(10)]; norm=int64(2);
```

This time, if we set maxm=[int64(4096),int64(4096)], then approximation is required. The BCCB extension matrix B of dimension $4096^2 \times 4096^2$ has negative eigenvalues. However, the smallest negative eigenvalue is $\mathcal{O}(10^{-11})$. By specifying icorr=int64(2), approximate samples are generated by replacing B with B_+. Since all the negative eigenvalues of B are actually very small, the covariance error is acceptable. The resulting approximate realisations of the mean zero Gaussian random field are shown in Figure 3.

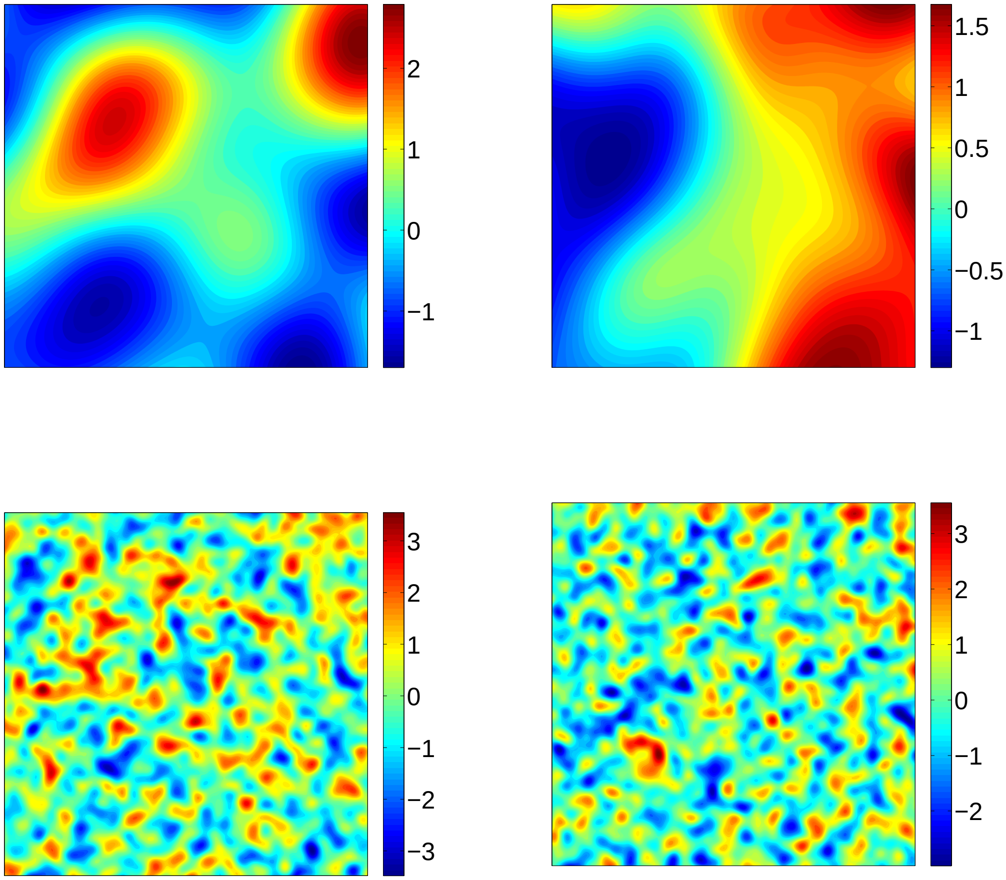

Figure 3: Approximate realisations of a mean zero Gaussian random field on $D = [0,1] \times [0,1]$ with the Gaussian covariance function $\gamma(\mathbf{x}) = \exp(-\|\mathbf{x}\|_2^2/\ell^2)$ with correlation length $\ell^2 = 1/10$ (top) and $\ell^2 = 1/1000$ (bottom).

References

[1] Dietrich C.R. and Newsam G.N., Fast and exact simulation of stationary Gaussian processes through circulant embedding of the covariance matrix. *SIAM J. Sci. Comput.,* 18, pp.1088-1107 (1997).

[2] Lord, G.J., Powell, C.E., Shardlow, T., *An Introduction to Computational Stochastic PDEs*. Cambridge Texts in Applied Mathematics, Cambridge University Press. To appear, April 2014.

[3] Taylor, P., Simulating Gaussian Random Fields and Solving Stochastic Differential Equations using Bounded Wiener Increments, *PhD Thesis*, University of Manchester, 2013.

[4] Wood A.T.A and Chan G., Simulation of stationary Gaussian processes in $[0,1]^d$. *J. Comput. Graph. Statist.,* 3(4), pp.409–432 (1994).